Choices and Their Consequences: Using Generative Models of Cognition in Understanding Voluntary Decision Making

PhD Thesis

Wojciech Zajkowski

10.09.2021

Contents

Abstract **4**

1 Introduction **6**

1.1 Basic Terminology and Concepts 6

1.2 Research Questions 11

1.3 Methodological Approach 16

2 Literature Review of Voluntary Choices **18**

2.1 What is volition? 21

2.2 Why to study voluntary choices? 29

2.3 When to utilize the free choice paradigm? 32

2.4 How to model voluntary choices? 39

3 Methodological Review **42**

3.1 Integrative approach 43

3.2 Role of generative cognitive modelling 47

3.3 Hierarchical Modelling 50

3.4 Practical example 53

3.5 Summary 69

4 Choices between equal options: impact of varying levels of choice freedom and discriminability (fMRI study) 70

4.1 Introduction 71

4.2 Aims and Hypotheses 73

4.3 Methods 75

4.4 Results 85

4.5 Discussion 94

5 Breaking Deadlocks: Reward Probability and Spontaneous Preference Shape Voluntary Decisions and Electrophysiological Signals in Humans 101

5.1 Note on Contributions 104

5.2 Introduction 104

5.3 Materials and Methods 111

5.4 Results 129

5.5 Discussion 151

5.6 Conclusion 161

6 Within and Across-Domain Effects of Choice Induced Bias (Behavioural Experiments) 164

6.1 Introduction 164

6.2 General Methods 170

6.3 Experiment 1 177

6.4 Experiment 2 183

6.5 Experiment 3 198

6.6 Experiment 4 201

6.7 Experiment 5 204

6.8 Discussion 208

7 Summary 217

7.1 General Implications and Limitations 218

7.2 Future Directions 226

7.3 Conclusion 230

References 232

Appendix 295

Chapter 4 295

Chapter 6 296

Abstract

In this work, I integrate three research projects focused on exploring the neurocognitive basis of voluntary decision making in humans. The thesis consists of two main parts.

First part (Chapters 1 to 3) focuses on reviewing the theoretical context of my later work. Chapter 1 introduces the main concepts of interest and scientific questions that motivate my work. Chapter 2 reviews the scientific contributions, such as previous findings, theoretical conceptualizations, and experimental frameworks on which my empirical work is built upon. Chapter 3 discusses the methodological philosophy my work follows. Each theoretical chapter starts with a brief introduction, giving the reader a glimpse of the topics to come, followed by sections discussing these points in more depth.

Second part (Chapters 4 to 7) describes the empirical research. Chapter 4 reports the findings of an fMRI experiment on the brain correlates of voluntary choices between equal options in a perceptual discrimination task (later referred to as *Project 1*). Chapter 5 describes an EEG study on the mechanisms of breaking decisional deadlocks, when faced with equally valuable options in a probabilistic choice task (later referred to as *Project*

2). Chapter 6 reviews a set of 5 behavioural studies on how a confirmatory bias induced by voluntary decisions affects future choices and evaluations (later referred to as *Project 3*). Chapter 7 summarizes the three projects and discusses their contributions and limitations.

Chapter 1

Introduction

In this work, I explore the behavioural, cognitive, and neural mechanisms of voluntary decision making in humans. This compact description hides beneath a very concrete set of study designs, problems, and methodological assumptions with a scope that barely scratches the vastness of human cognition and decision making. For this reason, in the initial chapter I take upon itself a rather laborious task of building the rationale of my work from bottom-up. This includes deriving meaningful definitions of voluntary choice, clarifying the scope of my investigation, and providing a bird-eye overview of the methodology. While defining terms such as voluntary choice might seem like a futile exercise in rhetoric, I believe that it is essential for developing a proper understanding, and therefore, rigorous scientific investigation.

1.1 Basic Terminology and Concepts

Decision Making

Decision making refers to the voluntary processes of choosing among a set of alternatives, leading to the act of commitment (Payne, Bettman, and Johnson 1993). While different models for the stages of the decisional process have been proposed (Busemeyer and Johnson 2004; Rangel, Camerer, and Montague 2008; Sutton and Barto 1998), most of them agree that common cognitive steps can be extracted, such as forming a representation of the options available for choice, evaluating those options, and setting a rule that defines when enough information is gathered to form commitment (*i.e.,* choice termination). My thesis focuses on difficult voluntary choices and their influence on future choices and judgments. Before delving into the specific research questions, I will first sketch the basic ideas and concepts behind these terms. A more in-depth literature review will then be provided in Chapter 2.

Voluntary Choices

Voluntary decision making, seemingly a pleonasm, is nonetheless often used in the literature to distinguish choices that are driven internally (by our internal values or preferences) from externally driven ones. In contrast, involuntary choices (also referred to as forced, specified, or instructed; Fleming et al. 2009; Phillips et al. 2018; Richardson, Pfister, and Fournier 2020) refer to the situation when one is forced to make a predefined choice. The conceptual waters in this case are a little murky – one might argue that an involuntary choice is not a decision at all. In practice, forced choices almost always

do require a choice to be made. One might be coerced or forced to do something he disagrees, or dislikes yet still have the freedom to refuse and face the consequences. Conversely, in a laboratory task, when asked to make a forced choice, a participant has the will to refuse, risking being excluded from the experiment. In between these two extremes, there also exists a finer spectrum, where the degree of freedom is defined by the degree towards which the choice is driven by internal (subjective representations, memories, goals or motivations) or external factors (properties of the stimulus or the circumstances; Fleming et al. 2009).

For example, perceptual choices, by their more objective nature (*i.e.,* the existence of a factually correct answer), are more externally driven compared to preferential decisions. Nuance also exists within the domains: a choice between two very distinct and well defined options is arguably more externally driven (*i.e.,* driven by those well-defined external characteristics), compared to when the definitions are blurred and options are more arbitrary. Let us imagine choosing between an apple and an orange. In the first scenario, one is presented with concrete instances of the two fruits, which they can touch and examine, before picking their favourite. In the second case, one is presented only with abstract semantic information in the form of words 'apple' and 'orange'. In the first case, one's choice will likely be driven by external characteristics of the fruit instances, such as colour, texture, smell, or ripeness. In the second case, one's choice will be more dependent on her internal, abstract representations of those fruits. This concept becomes more subtle when considering the perceptual domain, as here choices should ideally be made exclusively based on external sensory signals. This idealized

case is however far from reality, where perpetual decisions can be influenced by a wide range of biases and internal expectations (Berlemont and Nadal 2019; Glickman, Moran, and Usher 2020; Linares, Aguilar-Lleyda, and Lopez-Moliner 2019; Summerfield and de Lange 2014; Talluri et al. 2018). A common way of manipulating a perceptual choice so that it is driven by internal processes is assigning identical values to different perceptual stimuli (Lau et al. 2004; Rowe et al. 2008; Thimm et al. 2012), or increasing the discrimination difficulty or option value to a point where the options are almost or exactly identical (Bode, Bogler, and Haynes 2013).

Voluntary Action

While the thesis focuses predominantly on voluntary choices, due to the intertwined nature of choices and actions, it is impossible to completely separate the two.

Voluntary action is closely associated with voluntary choice, as the former requires the latter. In this sense, an action is a potential output of a choice process. Since choice does not require an physical action one can make a voluntary choice *not* to act (Rae et al. 2014), or commit to choice not associated with any immediate action, and voluntary action cannot be performed without a choice to do so, choice can be seen as encompassing a broader category of circumstances.

Another crucial line of distinction lies between the implementation of choices and actions. Choices involve processes described in the above sections, such as information sampling, evidence weighting or setting a choice threshold (Forstmann, Ratcliff, and Wagenmakers 2016) actions are a motor act of

choice execution (Haggard 2008). Consequently, studies of voluntary choices and actions focus on different aspects of neural processing. However, often such delineation is difficult to make, as some areas and processes can be involved in both. An excellent example of this idea is the intermediary role of the pre-supplementary motor area (preSMA) contributing to both choice formation and motor execution (Haggard 2008; Nachev, Kennard, and Husain 2008, for a detailed discussion see Chapters 2 and 4).

Compared to choices, voluntary actions are much often discussed in the cognitive literature (Haggard, Clark, and Kalogeras 2002; Haggard 2008, 2019), perhaps due to their more tangible nature (one can directly observe an action, while choice often needs to be inferred from said action). As such, in the voluntary action literature, *actions* can be often understood as a broader process, which includes the decision leading to the action (Haggard 2008).

Difficulty

Throughout the thesis, difficulty will be understood as the closeness in value between choice options. This closeness can refer either to perceptual discriminability, or the subjective value of choosing one option over the other. An extreme of this is a choice between 2 or more options with the same value or outcome – a paradigm adapted in many studies of voluntary choice (Barlas, Hockley, and Obhi 2017; Beck, Di Costa, and Haggard 2017; Krieghoff et al. 2009; Zhang, Hughes, and Rowe 2012). Two of the three projects described here use precisely this paradigm, asking the question: *what happens, both in terms of behaviour and neural activity, when one is faced with seemingly impossible choice between multiple equivalent options?*

This definition is constrained to a specific type of 2-alternative forced choice (2AFC) problems, where choice options are assumed to be compared on a single dimension, such as a perpetual attribute *e.g.* size (Rafael Polania et al. 2014), motion direction (Mulder et al. 2012), phase coherence (Ratcliff, Philiastides, and Sajda 2009) brightness (Teodorescu, Moran, and Usher 2016); or subjective value (Rafael Polania et al. 2014; Fontanesi et al. 2019).

In multi-attribute problems, difficulty might reflect aspects like emotional valence (Luce, Payne, and Bettman 1999), number (Dhar 1994) and similarity of choice-relevant attributes (Cheng and González-Vallejo 2018). Difficulty in multi-alternative problems, where participants choose between 3 or more options, becomes even harder to define, as apart from factors such as closeness in value and comparison across different attributes, one also has to consider also the number of alternatives (more options require more processing time, Proctor and Schneider 2018) and choice heuristics (Gigerenzer and Gaissmaier 2011). Hence, while the current definition of difficulty is appropriate when considering a limited 2AFC case using experimentally-controlled stimuli, it is important to remember that in more complex real life situations it is only one of many factors that can influence how hard or easy a choice is.

1.2 Research Questions

The research projects contained in this thesis follow a flow of starting from the most broad questions relating to describing basic correlates of free choice and their neural signatures (Project 1), to delving deeper into finding the mechanisms of solving decisional deadlocks (Project 2), and finally understanding

how choices can influence value, leading to a self-reinforcing mechanism of valuing more the things that we choose (Project 3).

In the broadest terms, my thesis focuses on 3 main research questions:

1. Describe correlates of equal options made (an fMRI study; Chapter 4)
2. Understand how decisional deadlocks are broken (an EEG study; Chapter 5)
3. Describe how choices influence subsequent evaluations (a set of 5 behavioural studies; Chapter 6)

Project 1

The first project focuses on perceptual choices between equally salient stimuli, manipulating the number of options available to choose, from 1 to 3. Option availability manipulation is usually associated with prolonged response times, a phenomenon widely recognized in psychological science (Proctor and Schneider 2018). Prolonged responses are associated with the abundance of options to choose from, even when all of them lead to a similar outcome. This phenomenon is tangentially related to the *paradox of choice* effect found in value-based literature, where more choice alternatives can lead to more indecision and less choice satisfaction among consumers (Dar-Nimrod et al. 2009). Previous studies using this variation of a perceptual free-choice paradigm have shown that contrasting free choices with forced ones slows down reaction times and activates a set of brain regions associated with internally-driven action selection, most notably the pre-supplementary motor cortex (preSMA) and dorsolateral prefrontal cortex (dlPFC; Rae et al. 2014; Si, Rowe, and Zhang 2020). Apart from replicating these known effects, this

study aims at answering 3 novel questions:

- How does stimulus saliency influence choices between equal options
- Whether introducing an intermediate level of choice freedom would produce a graded BOLD response
- Are more 'objectively free' choices also subjectively perceived as more 'free'

Saliency here refers to the strength, or ease of discriminability of the stimulus (Teodorescu, Moran, and Usher 2016). Objective choice freedom was manipulated by controlling the number of available response options. Subjective freedom refers to explicit feeling of freedom ratings, spread across the experimental trials following a choice (Filevich et al. 2013). Despite a broad literature on the topic (for reviews see: Haggard 2008, 2019; Si, Rowe, and Zhang 2020), these questions still remain not fully explored. Describing these basic features of voluntary choice is a necessary building block in understanding and modelling voluntary choice as a coherent and integrated process.

Project 2

The second project takes this initial idea a bit further and poses the question of how such decisional deadlocks are broken. Here, we aim to describe the effects of equal choices when manipulating the probability of obtaining a reward, and finding the cognitive mechanisms used for breaking such decisional deadlocks, together with their neural underpinnings. The study focuses on a probabilistic reward task, where participants make 2AFC choices between abstract symbols representing reward probabilities, while their EEG

signal is recorded. The main focus of the study lies in choices made between symbols representing identical probabilities (equal choice condition).

This project is the first to provide evidence of magnitude effects in probability-based reward task, as well as provides a candidate neuro-cognitive mechanism for resolving a famous choice paradox, referred to as the *Buridan's ass* problem (Inwagen 1989). In this hypothetical paradox, a donkey faced with two identical haystacks starves to death, as a result of inability to choose the better option. In practice, humans need to be equipped with mechanisms able to break out of such decisional deadlocks (Pais et al. 2013). The project establishes a potential mechanism, but also uncovers an interesting phenomenon, that people exhibit spontaneous preference for certain options, even when the option outcomes are identical (preference effect; Lebovich et al. 2019). The study has been published (Zajkowski et al. 2020).

Project 3

Project 3 deals with the question of how one's choices can shape the emerging phenomenon of solving difficult dilemmas by spreading alternative values.

The third project expands on this finding by asking how people develop such spontaneous preferences. The premise for this study is that internal values are not stationary, but rather object to fluctuations, that can be biased by one's choices. The study drifts away from the equal-option setup and compares how choices in preference and perceptual domains shape internal estimations. In contrast to previous research on this topic focused mostly on preferential choice (Hornsby and Love 2020; Sharot et al. 2012; Vinckier et al. 2019), this project compares and contrasts both preference and perception, asking whether a preferential choice can induce a change in perception, and

vice versa.

To my knowledge, this is the first study to directly look at across-domain effects (i.e. how choices in one domain can influence an unrelated domain) and find such effects. The results build upon a vast literature of literature on choice-induces bias (Izuma et al. 2010; Sharot, Velasquez, and Dolan 2010; Talluri et al. 2018; Voigt, Murawski, and Bode 2017; Voigt et al. 2019). The modelling approach used provides a link between cognitive dissonance (Festinger 1957) and active inference based on one's choices (Bem 1967; Friston et al. 2016), showing that both mechanisms - a conflict-driven adjustment (Brehm 1956) as well as a choice-driven permanent value-update (*e.g.* Chammat et al. 2017) contribute to the spread of alternative values following a choice.

The three projects together can be seen as a progression, from describing free choice (Project 1), to understanding how free-choice conflicts can be resolved (Project 2), to delving deeper into the consequences of voluntary choices and describing how they can affect future decisions and evaluations (Project 3).

While being similar in the respect of the main topic of study (voluntary choices and the neuro-cognitive mechanisms supporting them) as well as the methodological approach used (see next section), the projects vary in terms of experimental methodology, types of tasks and some of the definitions. The main differences in paradigms are introduced in the next chapter and summarized in the concluding Chapter 7 (Table 7.1), together with the discussion about the limitations related to comparing different task modalities (perceptual, value-based, and preference-based).

1.3 Methodological Approach

Experimental design and data analysis methods are often underappreciated in psychological science. While flashy results and speculative implications are exposed at the forefront, their methodological foundations are often hidden from view (Francis 2012), even though they are the cornerstone of any inferential process. In reality, methodological oversights were one of the main reasons for a replicability crisis taking place across domains of psychological science over the last decade (Lilienfeld 2017; Romero 2019; Tackett et al. 2019). Strong arguments have been made for moving away from traditional cookie-cutter linear methods of analysis and towards custom models prioritizing the data-generating process (Haynes 2019; Rooij and Baggio 2021; Yarkoni 2019).

This line of thinking follows from the assumption that any behaviour is a manifestation of an underlying cognitive process, which is constrained by one's biology. Therefore, a meaningful understanding of a psychological phenomenon requires an integrative approach of traversing between observed behavior, latent cognitive processes generating it, and the neural mechanisms via which the cognitive processes are implemented. In accordance with this philosophy, where possible, my work follows three methodological tenets:

1) Integrative Approach

- between levels of analysis
- between modalities (combining behavioural and neural data)

2) Generative Modelling

3) Hierarchical Bayesian Estimation

Levels of analysis refer to the seminal work by Marr (1982), and consist of *computational* (*what* is computed; highest-level, qualitative description of a process), *algorithmic* (*how* is the process computed; describes the intermediate, abstract algorithm employed to realise the high-level goal), and *implementational* (*how* is the algorithm implemented, given neural constraints) levels.

Generative modelling, at it's heart, is a way of improving understanding of the data-generating process, and therefore scientific inference (not to be confused with statistical inference; see: Szollosi et al. 2020).

Hierarchical Bayesian estimation aids in statistical inference by solving common estimation problems, such as providing full parameter distributions and partial pooling (Gelman 2006; for details see Chapter 3).

These tenets are further elaborated upon in *Chapter 3*, which discusses the methodological approach in depth.

Chapter 2

Literature Review of Voluntary Choices

A prominent model specifies voluntary action in terms of three crucial components: *whether*, *what* and *when* (the *WWW* model; Brass and Haggard 2008), which can be also extended to voluntary choices. *Whether* asks the most general question: a decision whether an action should be performed or withheld. This concept is related to impulse control (Aron, Robbins, and Poldrack 2004) *i.e.,* the ability to inhibit prepotent responses when faced with a situation that requires it (*e.g.* yawning at an important meeting or showing visible disgust when tasting an exotic dish in a fancy restaurant). Above impulse control, the *whether* component also encompasses more deliberate decisions related more broadly to self-control (Brass and Haggard 2007), such as whether to go the gym, or write a thesis on a weekend. A

popular framework for studying the *whether* component includes the variants of the *stop and go* task, as well as deliberate stopping paradigms (Brass and Haggard 2007, 2008).

The *what* component asks which action should be chosen among a set of alternatives. It aims to explain the core elements of the voluntary decision-making process, such as information selection, decision rule and commitment. It is often studied using force-choice designs, where a varying number of alternatives is presented to the participant. Due to its straightforward conceptualization, this type of choice has been studied extensively both under the guise of the study of volition (Brass and Haggard 2008), as well as perceptual (Heekeren, Marrett, and Ungerleider 2008), or preference-based (Philiastides and Ratcliff 2013) decision making.

The *when* component asks the question when exactly the chosen action should be executed. Such choice requires a temporal specification of action performance, when the *whether* and *what* choices have already been made. Most famously studied by Benjamin Libet (Libet et al. 1983), who described a brain potential (so called *readiness potential*) which reliably preceded conscious action. By linking the readiness potential to free will, Libet triggered a great debate about moral responsibility in light of this new facet of determinism (Harris 2012; Lavazza 2016; Roskies 2006; Wegner 2004). The argument being "*How can one be responsible for his actions, if he does not consciously decide to act?*". Libet's controversial claim has been since mostly rejected, both on its philosophical (Nachev and Hacker 2014) as well as methodological merits (Schurger, Mylopoulos, and Rosenthal 2016), yet the study of the *when* component is still a thriving area of volitional research

(Haggard 2019).

Together, the model provides a concise way to partition the science of voluntary choice and action into three qualitatively different problems which can be studied within their respective domains. As a tribute to this theoretical conceptualization, current chapter is organized analogously to WWW model, with sections representing a variation on the WWW theme:

- *What is volition?* section will aim at explaining the theoretical considerations, nuances and pitfalls of what is and what is not considered voluntary.
- *Why to study voluntary choices?* section will aim to convince the reader as to why studying voluntary action is important, and what can it bring to the table in terms of scientific advancement. An attentive reader will notice that *why* is not included in the original WWW model. The change from *whether* to *why* illustrates my optimism and conviction in the necessity in continuing this avenue of scientific pursuit.
- *When to utilize the free choice paradigm?* section will discuss the subtleties in using two popular experimental designs used to study voluntary choices.

An additional follow-up section, not accommodated by the WWW metaphor due it's alliterative deficiency, *How to model voluntary choice?* will discuss and examine popular theoretical frameworks used for explaining the cognitive machinery involved in producing voluntary choices.

2.1 What is volition?

In order to provide an appropriate context for further discussion, I will embed the topic of voluntary choices in a general discussion of human volition: how is it understood and studied through the multidisciplinary lens of cognitive neuroscience, drawing its influences from the fields of philosophy, psychology and biology. Throughout this section I will focus on a critical axis of distinction between *voluntary action* and the *subjective experience* of intentionality. Along the way, I will contextualize this general framework by providing specific examples of studies from the realm of decision-making.

Volition is difficult to define. This is due to the fact that, similarly to other concepts studied by both psychologists and philosophers, such as consciousness or free will, a subjective phenomenological experience is at the core of its meaning. In the terms of folk psychology (Lewis 1972), to act voluntarily means to feel responsible for one's actions, or to feel control over their consequences (Haggard 2008; Nahmias et al. 2005).

In their critique of contemporary research on free will, Nachev and Hacker (2014) stress the difference between *acting*, that is, responding voluntarily to the circumstances, and *doing*, which includes involuntary reactions. They categorize the first as a motivated property exclusive to sentient beings, while the second as a more general term referring to both voluntary acts as well as non-autonomous processes such as photosynthesis or falling asleep. Acting however is just part of the story, since voluntary behaviour can be expressed not only through action, but also through inaction: an autonomous agent can also voluntarily inhibit an inappropriate response (Aron, Robbins, and

Poldrack 2004; Chikazoe et al. 2009).

This understanding of volition is tightly coupled with the notion of intentionality (Marken 2016), as well as the sense of agency (Haggard 2017). While similar in scope, intentionality refers to the more objective process of acting upon the environment (C. K. Turner 2017), whereas sense of agency refers to the subjective experience associated with action (Hayden and Haggard 2017). Crucially, an intended action is not always perceived as such, while an unintended action can be associated with a strong sense of agency (Sato and Yasuda 2005). The latter has been shown in famous perceptual illusion where people perceive movements of a rubber hand as their own (Kalckert and Ehrsson 2014), as well as in cognitive dissonance literature, where people falsely attribute choices they never made as their own (Henkel and Mather 2007). Finally, volition is also strongly linked to the concept of free will and personal responsibility related to it (Haggard 2008), a topic more thoroughly discussed in the *Why study voluntary behavior* section.

Given this wide range of phenomena and competing constructs bundled together under the guise of volition, one might wonder how to study such a complex concept. Taking all of these considerations into account makes it particularly challenging to come up with a concise definition and operationalization that would satisfy the constraints of an experimental investigation. Psychological research has developed two main ways to operationalize and study volition: one related to the more objectively defined intention, and the other related to the phenomenological experience of acting freely.

2.1.1 Volition as exogenously driven stimulus-independent action

This conceptualization aims at providing an objective and scientifically measurable definition. It proposes that voluntariness of action can be defined on a continuum, ranging from completely involuntary, such as simple reflexes that can be driven by a neurological reflexes (*e.g.* the knee-jerk reaction), to stimulus-independent, free from immediacy action (Fried et al. 2017; Gold and Shadlen 2001). Main strength of this approach is that it allows to experimentally manipulate the degree of voluntariness by manipulating the number of available alternatives, hence making actions more or less stimulus-dependent.

Early work showed that self-generated choices were associated with the activity of the pre-supplementary motor association cortex (preSMA; Lau et al. 2004; Krieghoff et al. 2011), pointing to it being a critical region responsible for transitioning from intention to action (Haggard 2008). The studies consistently indicated preSMA being more active in the voluntary, compared to instructed movement condition. Later work suggested that preSMA might be specifically involved in accumulating evidence in favour of a given option (Rowe, Hughes, and Nimmo-Smith 2010; Tomassini et al. 2019; Zhang, Hughes, and Rowe 2012). Some studies also suggested that the organization of the dorsomedial frontal region consisting of the preSMA, SMA and the dorso-anterior part of the cingulate cortex follows a functional, rosto-caudal gradient, where more anterior regions process more abstract information (*whether* and *what*) while the more posterior ones are engaged in solving more concrete problems (*when*; Krieghoff et al. 2009). In

contrast, lateral parts of the frontal cortex, involving dorsolateral cortices and the inferior frontal gyrus, are associated with switching and inhibiting the previous response, respectively (Rowe, Hughes, and Nimmo-Smith 2010; Zhang, Hughes, and Rowe 2012). Meta-analytic studies further indicated a set of regions associated with voluntary choice, including the preSMA, dorsolateral prefrontal cortex (DLPFC), inferior parietal lobule (IPL) and left anterior insular cortex (AIC; Si, Rowe, and Zhang 2020).

While this approach proved fruitful in discovering a functional network underlying intentional choice, it also comes with a set of weaknesses related to its ecological validity. Firstly, it is not certain whether more stimulus-independent actions are always experienced as more voluntary, therefore severing the link between this operationalization, and an intuitive understanding of volition. Secondly, as pointed out by Haggard (2008), instructing participants to perform a free choice can be perceived as paradoxical.

2.1.2 Volition as the feeling of freedom

The subjective feeling of voluntariness is related to the feeling of freedom (Filevich et al. 2013) and a sense of agency (Gallagher 2000). While similar in scope and often used interchangeably, it is important to note that sense of agency implies both a feeling of intentionality, and control over action consequences (Chambon, Sidarus, and Haggard 2014). Cognitive science has developed explicit and implicit measures to study this phenomenon. Explicit measures constitute self-reports about the feeling of control or freedom (Linser and Goschke 2007; Metcalfe and Greene 2007; Sato and Yasuda 2005). Implicit measures aim at quantifying this feeling without overtly

querying the participants about it. Intentional binding (Haggard, Clark, and Kalogeras 2002) is arguably the most popular of such methods. The technique measures the time difference between action generation and outcome, where the estimated closeness in time of the two events approximates the sense of agency (Moore and Obhi 2012).

While referring to the same underlying concept, some suggest that explicit measures are related to high-level metacognitive processes, *i.e.,* inferred agency, while implicit ones to a low-level, non-conceptual and mostly experiential feeling, related to simple behaviour, such as switching a light switch or grabbing a fork (Synofzik, Vosgerau, and Newen 2008). Studies show that explicit and implicit measures covary as a function of number of available alternatives (Barlas, Hockley, and Obhi 2017; Barlas and Kopp 2018; Filevich et al. 2013; but see also Dewey and Knoblich 2014 for the opposite claim), but also deviate in certain conditions: self-reports being more sensitive to action-outcome congruency (Ebert and Wegner 2010) and outcome pleasantness (Barlas, Hockley, and Obhi 2017), while intentional binding to choice repetition (Moore et al. 2012). The discrepancies between the two are often attributed to two different components driving both measures: a prospective component (experienced choice fluency) influencing intentional binding, and a retrospective component (rationalized; metacognitive inference) associated with explicit assessments (Haggard 2017).

The neural correlates of sense of freedom are currently a contested issue. Recent meta-analysis study including 20 fMRI experiments revealed no significant clusters positively associated with agency (Zito, Wiest, and Aybek 2020). A meta-analysis of 6 transcranial Direct Current Stimulation (tDCS)

experiments indicated stronger intentional binding is associated with activity in DLPFC (Khalighinejad, Di Costa, and Haggard 2016). Results on negative predictors of agency (*i.e.,* lack of agency) have been more fruitful, showing consistent activity of temporo-parietal junction (TPJ) across different experiments and paradigms (Sperduti et al. 2011; Zito, Wiest, and Aybek 2020). One meta-analysis study of 15 experiments also reported increased activity within preSMA, precunuous and dorsomedial prefrontal cortex (Sperduti et al. 2011), suggesting potential overlap with the representation of intentionality.

The subjective definition of volition also has its fair share of limitations. An important question relates to construct validity: how accurate do implicit and explicit measures really measure the subjective experience? While this issue might be impossible to tackle empirically due to the subjective nature of the definition, the heterogeneity of operationalizations and measures makes the studies very difficult to compare. The fact that a branch of the literature is devoted to comparing the different measures (Barlas, Hockley, and Obhi 2017; Barlas and Kopp 2018; Dewey and Knoblich 2014; Moore and Obhi 2012; Sidarus, Vuorre, and Haggard 2017) illustrates this problem. Additionally, study designs vary in whether choice autonomy is a dichotomous or a graded variable, making the across-study comparisons problematic even for studies utilizing the same measure of experienced freedom.

2.1.3 Combining the two meanings

Both perspectives on volition: the study of exogenously-driven intention, and the study of subjective experience bring complementary insights into our understanding of voluntary action. Frith (2013) offers an interesting

perspective by describing these two approaches as a third-person vs first-person view of voluntary behavior. Hence, an outstanding question arises: how are these two perspectives related?

While there has been a considerable theoretical interest in the relation between the two definitions (Frith 2013; Nahmias et al. 2004), only a handful of studies have attempted to address it experimentally. Haggard et al. (2004) used hypnotic suggestion to make a voluntary choice seem involuntary. This masked voluntariness condition had no effect on subjective experience, but significantly affected the perceived time between action initiation and outcome – an implicit marker of sense of agency (Haggard, Clark, and Kalogeras 2002). A planned contrast also revealed that voluntary movements were also perceived as more voluntary, as compared to hypnotically induced ones. The authors claimed this as evidence for a qualitative difference in how voluntary and involuntary (when the intention is not experienced as endogenous) actions are experienced. Further investigations indicated that choice fluency (or smoothness; *i.e.,* involving less conflict) predicts agency judgments (Chambon and Haggard 2012; Wenke, Fleming, and Haggard 2010). In a similar vein, in a set of experiments, Barlas and colleagues (Barlas, Hockley, and Obhi 2017; Barlas and Kopp 2018; Barlas and Obhi 2013) showed that increasing the number of available alternatives also influences both implicit and explicit (feeling of control) measures of subjective freedom. On the level of the brain, transcranial stimulation studies indicated that stimulating preSMA results in an urge to move, suggesting the area is involved not only in intention generation, but also subjective experience (Desmurget et al. 2009; Fried et al. 1991). This consistent picture was put in

doubt by Filevich et al. (2013), who showed that while number of available alternatives significantly predicts subjective feeling of freedom, it's neural signature was in stark opposition - brain regions involved in coding intention (SMA, DLPFC, inferior parietal lobule (IPL) and left premotor cortex (PMC) were inversely related to the subjective experience of freedom. Only activity of the precunuous was consistent across both measures.

2.1.4 Other Considerations

While studies on voluntary choice have been focused on tasks requiring participants to choose between equivalent options (equal choice paradigm; see *When to utilize the free choice paradigm?* section), one might argue that this approach is rather limited. Studying only the extreme end of the spectrum (choices being equal) might not necessarily generalize to more nuanced scenarios, where actions are driven by a mixture of exogenous and endogenous drives. This experimental design also brings about a philosophical conundrum: how free is choosing between alternatives that lead to an identical outcome? Using a coarse analogy, one might argue that choosing between a left and right door, where both of the doors lead you off a cliff, is not really a free choice. Control over choice consequences is an important component of agency (Haggard 2017).

Another restricting factor in many studies on voluntary choice is the assumption of action independence, *i.e.,* that subsequent choices do not influence one other in any temporally meaningful way. Research in perceptual (Kayser and Kayser 2018) and value-based (Senftleben et al. 2019) decision making however strongly suggests that our decisions have causal effects on future

choices, as well as on underlying values. Hence, a more comprehensive understanding of voluntary actions requires understanding how values and preferences fluctuate in time - a topic explored to a greater depth in the third project of this thesis.

2.2 Why to study voluntary choices?

Understanding voluntary behavior can bring progress in a range of fields, from theoretical insight into one of the most profound philosophical issues, to concrete advancements in treating psychological and medical disorders. Below I list motivating factors for studying volition and it's consequences in the domains of philosophy, law and psychiatry.

2.2.1 Philosophy

The question of whether people have free will has been a subject of human curiosity for centuries, absorbing the minds of ancient philosophers and theologists alike (Frede 2011). Overall, people have a strong sense of being able to act freely (Hallett 2007) and that the world is non-deterministic (Sarkissian et al. 2010). Many have argued that free will is compatible with our current scientific knowledge (Dennett 1984; Frankfurt 1988; Klemm 2010). This view is challenged by other thinkers, who claim this feeling to be nothing more than an illusion (Harris 2012; Wegner 2004) or an unnecessary epiphenomenon, devoid of causal influence (Hallett 2016). While psychology and neuroscience are unlikely to provide a definite answer to this philosophical dispute, they can enrich our understanding of the cognitive

and neural processes giving rise to intention and a sense of freedom, as well as explain causal structures driving actions (Nachev and Hacker 2014). For example, the famous Libet's clock experiment (Libet et al. 1983) was able to shift the tone of the discussion on free will in both scientific literature and popular media for years to come (Schlegel et al. 2015).

2.2.2 Law

The problem of free will is strongly tied to how people perceive and interpret moral responsibility. Our jurisdictional system is based on the assumption that people are responsible for their voluntary actions, and the verdicts are based on judgments of agency (Haggard 2017). The perpetrator can be only held accountable for a crime if he consciously decides to perform it, while being aware of the probable outcome (Summers 1969). Law acknowledges diminished responsibility in cases where agency could be considered incomplete, such as neurological or psychiatric impairments, or obeying orders. Assessing agency in such cases can extremely difficult. A newly emergent field of neurolaw aims at aiding this process (Meynen 2016). The literature is focused on three main topics: whether (and if so, to what extent) advances in neuroscience should influence the word of law (Greene and Cohen 2004); how brain data can aid in the assessment of people (e.g. does a psychiatric or neurological disorder diminish personal responsibility; Greely 2011); and interventions focused on developing potential treatments that reduce the risk of reoffending. Assessment issue is the most relevant from the point of view of cognitive neuroscience of voluntary choice. Current research suggests that limiting the number of available action alternatives (Barlas, Hockley, and

Obhi 2017; Barlas and Obhi 2013; Filevich et al. 2013) as well as coercion (Caspar et al. 2016) can significantly constrain experienced agency. A deeper understanding of neurocognitive mechanisms of intention and agency can lead to the development of diagnostic tools sensitive in assessing the degree of freedom.

2.2.3 Neurology and psychiatry

Several psychiatric and neurological disorders are connected to dysfunctions related to voluntary behaviour. Deficits in intentionality have been linked to a variety of neuropsychiatric symptoms, such as phantom limbs, anarchic hands or utilization behaviour (*i.e.,* inappropriate use of objects; Blakemore, Wolpert, and Frith 2002). Depression has been hypothesized to be related to a reduced sense of control over own actions, suggesting a diminished (but more realistic) sense of agency (Alloy and Abramson 1979). Symptoms of schizophrenia such as delusions of control or misattributions of thoughts are related to impaired ability to distinguish self and other-generated actions (Daprati et al. 1997), indicating a lack of awareness of aspects of motor control (Frith, Blakemore, and Wolpert 2000). A reduced sense of agency due to misattribution of own movements has also been observed in patients with Parkinson's disease (Saito et al. 2017). The patients exhibited diminished agency based on both explicit and implicit measures. Other studies have also indicated that Parkinson's patients' experience of intention is significantly delayed, the effect being mediated by dopamine and reduced preSMA activity (Di Costa et al. 2020; Jacobs et al. 2009; Tabu et al. 2015). Similar deficits have been observed with patients with Tourette syndrome, a movement

disorder characterized by involuntary ticks (Moretto et al. 2011).

Further work on the cognitive and neural mechanisms of voluntary behaviour can lead to the discovery of more effective treatments for symptoms of many disorders. While not many direct attempts have yet been made, recent advances in computational neuropsychiatry (Adams, Huys, and Roiser 2016) bring promise of developing effective interventions in the near future.

2.3 When to utilize the free choice paradigm?

The name *free choice paradigm*, quite confusingly, has been assigned to two entirely different experimental designs. One originates from voluntary choice literature and refers to a design where participants choose one action among a number of equally valuable alternatives (Si, Rowe, and Zhang 2020). The second originates from the study of preferences and confirmation bias. Introduced in 1956 by Brehm, the task utilizes a rating-choice-rating design to measure how choices can influence evaluations. To avoid further confusion, I will henceforth refer to the former as equal-choice paradigm (ECP) and the latter as free choice paradigm (FCP; see Figure 2.1). Since my thesis utilizes both: 2 variants of the ECP (projects 1 and 2) and a new variation of FCP (project 3), I will next provide a brief introduction into both experimental traditions.

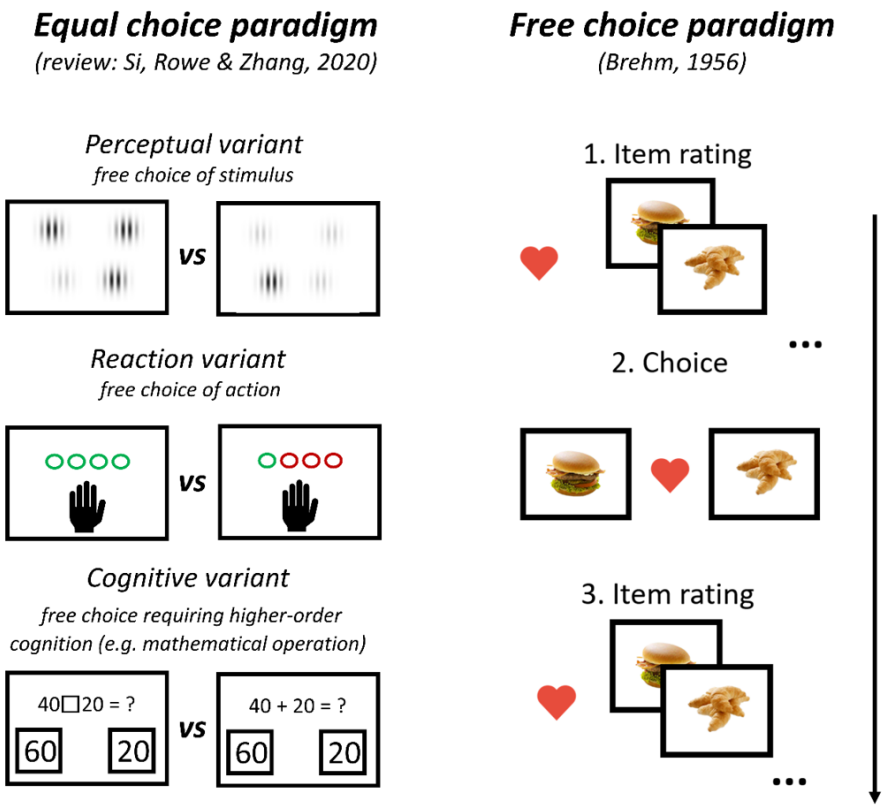

Figure 2.1: Equal (left) vs Free Choice (right) paradigm. In the ECP column, left screens exemplify free choices (more than 1 possible response), while right - instructed responses (only 1 viable alternative). In the perceptual variant example, high-contrast Gabor patches represent alternatives available for choice.

2.3.1 Equal Choice Paradigm

The equal choice paradigm (ECP) was initially developed in the early nineties (Deiber et al. 1991; Frith et al. 1991; Jahanshahi et al. 1995) as a means of studying the neural mechanisms of voluntary choice. In these studies, researchers compared self-generated finger tapping responses to pre-specified sequences of taps, in an attempt to extract the difference in brain activity between free, internally generated action and instructed sequence repetition. These studies identified a rather broad volitional functional network involving the pre-motor, supplementary motor, dorsomedial and dorsolateral frontal cortices. These early attempts however were later shown to be confounded (Lau et al. 2004). The unconstrained nature of choice in the *free* condition made it also more demanding of attentional resources. Future iterations of the design aimed at controlling the attentional confound by either introducing more demanding specified action condition (Lau et al. 2004), reducing the complexity of the free choice condition by constraining the free movement space to only few alternatives (Rowe, Hughes, and Nimmo-Smith 2010; Rowe et al. 2005; Zhang, Hughes, and Rowe 2012) or both (Rowe et al. 2008).

More variants of the task have been developed, including perceptual discrimination (*e.g.* Thimm et al. 2012), voluntary *go/no go* (Karch et al. 2009) and choices including higher cognitive functions, like choosing meaningful mathematical operators (Wisniewski, Goschke, and Haynes 2016) or attention redirection without movement (Taylor, Rushworth, and Nobre 2008). The studies began to show a growing consensus on a key involvement of preSMA cortex in generating free actions (Haggard 2008, see also *Volition as exogenously driven stimulus-independent action* section). Recent metanalysis

involving 39 studies across all variants of the task revealed an intention network consisting of preSMA, bilateral DLPFC, bilateral IPL and right AIC (Si, Rowe, and Zhang 2020).

One often underutilized feature of ECP is manipulating the extent of choice freedom. The task design allows for having many options available to choose on screen (usually limited to 4 in fMRI studies, so that responses can be mapped to 4 fingers (e.g. van Eimeren et al. 2006; Zhang, Hughes, and Rowe 2012), yet most studies use an 'all-or-nothing' design, where only one level of decisional freedom is contrasted against a specified action (for exceptions see: Forstmann et al. 2006; Lau et al. 2004; van Eimeren et al. 2006). Project 1 aims at tackling this issue, combining it with testing the effect of perceptual discriminability and the relation between these systematically manipulated factors and subjective experience of freedom. Study 2 further modifies the ECP, adapting it to the realm of probabilistic value-based choices in order to differentiate the effect of value on breaking decisional deadlocks.

2.3.2 Free Choice Paradigm

The free choice paradigm (FCP) was first introduced by Jack Brehm (1956) in one of the first recorded experiments testing Festinger's cognitive dissonance theory (1957). The theory proposed that holding opposing cognitions produces a negatively valenced feeling of dissonance, which motivates one to reduce it by either changing behavior or the underlying cognitions (McGrath 2017). For example, during a pandemic, a person may feel an internal dissonance from not wearing a mask while knowing it minimizes the risk of spreading the virus. In order to minimize the experienced dissonance, he may

then change his behavior (start wearing a mask), or one of the underlying cognitions (*e.g. 'masks don't work'*, *'the mask can suffocate me'*, or even *'the virus is a hoax'*).

One of the theoretical predictions of the theory was that choices between options of similar value should induce a cognitive conflict, which would then lead to a dissonance reduction *via* changing one's cognitions about the choice options. In particular, the chosen item should be viewed more positively, while the rejected one more negatively, as compared to before the choice was made - a phenomenon referred to as *spreading of alternatives* (E. Harmon-Jones and Mills 2019). Brehm's task consisted of 3 parts: first participants rated a set of kitchen appliances one by one on a 8-point scale. They were then presented with a single choice between two of the items. The choice was consequential, *i.e.,* the participant could take the chosen item home. Afterwards, participants rated all items again. In line with theoretical predictions, results revealed the chosen item was rated significantly higher in the second rating, while the rejected one significantly lower.

Similarly to the ECP, early findings using FCP have also been subjected to a serious confound. Chen and Risen (2010) used simulations to show that the spreading effect can arise even without an underlying change in preference. To understand this, we must realise that choices made in the task are not independent and cannot be experimentally manipulated. Since initial ratings are inherently noisy, by observing one's choice, we gain more information about his underlying values. This means choices can also reveal preexisting preferences, something that early research did not take into account. Let us look at a concrete example: in phase 1 participant rates an apple as more

preferred than an orange; in phase 2, when faced with a choice between the two, he chooses the orange; in phase 3 he rates orange as more preferred. A classic explanation of this effect would be than the choice influenced item values. However, it could also mean that the participant preferred the orange from the start and the estimation noise in phase 1 was why he rated it lower. What happened in phase 3 was a regression to the mean effect, where both items were rated closer their true underlying values (Chen and Risen 2010; Izuma and Murayama 2013).

These results shook the field of choice-induce preference change, with focus put on confirming whether the effect exists and designing improved experimental paradigms able to control for the confound, including blind choice (*e.g.* Egan, Bloom, and Santos 2010; Sharot, Velasquez, and Dolan 2010) and a control condition where second rating comes before the choice (*e.g.* Koster, Duzel, and Dolan 2015; Voigt, Murawski, and Bode 2017).

Today, a large body of literature have confirmed the existence of the effect in both preference-based (Vinckier et al. 2019) and perceptual (Luu and Stocker 2018) task, and scientific efforts have shifted back to explaining the effect. Since the introduction of FCP, a plethora of explanations for the effect have been proposed, most of which can be grouped within 2 broad categories. First stems from the cognitive dissonance theory (Festinger 1957), suggesting a motivated, conflict and consistency-driven adjustment (Gawronski and Brannon 2019; E. Harmon-Jones and Harmon-Jones 2019; Kenworthy et al. 2011). The second originates from autoperception theory (Bem 1967), suggesting the effect is driven by learning about ones preferences by sampling previous choices ("*since I remember choosing A, I must like it*"; Ariely and

Norton 2007; Kaaronen and Dale 2018; Kruglanski et al. 2018). Evidence for one or the other mechanism driving the effect have been mixed (Chammat et al. 2017; Izuma et al. 2010; Lieberman et al. 2001; Luettgau et al. 2020; Sharot, De Martino, and Dolan 2009; van Veen et al. 2009; Voigt et al. 2019).

An outstanding issue that still awaits investigation is whether the effects observed in preference-based studies are comparable with the ones from perceptual research, both in terms of qualitative (cognitive processes driving the effect) and quantitative (effect size) features. Additionally, such comparison allows for testing *across* domains, *i.e.,* whether a perceptual choice can influence preference, and *vice versa*. This question is vital from both theoretical and empirical perspectives. From a theoretical point of view, the existence of an across-domain bias would cast a serious doubt on many models proposing its adaptiveness (Lee and Daunizeau 2020; Harmon-Jones, Harmon-Jones, and Levy 2015; Kaaronen and Dale 2018; Kruglanski et al. 2018; Peters 2020). One would be hard-pressed to imagine a scenario where a context-irrelevant choice bias would be beneficial. From an applied perspective, existence of across-domain bias could be a powerful explanatory tool in studying many real-life phenomena that can be perceived as irrational, such as voting preferences inconsistent with self-interest (Fishbein and Coombs 1974), or irrational economic behaviors (Becker 1962).

In Project 3, I aim at directly comparing perceptual and preference-induced choice bias as well as their interactions, in order to address questions above. Additionally, I propose a cognitive model based on Reinforcement Learning theory (Sutton and Barto 1998) to account for how the bias arises. The model can account for 2 types of processes: a domain independent, consistency-

driven bias, and a domain-specific value update mechanism, representing dissonance and autoperception learning theories respectively.

2.4 How to model voluntary choices?

One of the core themes of this thesis is linking overt behaviour with measures of neural activity and generative models of cognition. In this work, I utilize two computational frameworks, often used in modelling decisions in free-choice and equal-choice tasks: evidence accumulation (Evans and Wagenmakers 2019) and reinforcement learning (Sutton and Barto 1998).

2.4.1 Evidence Accumulation

Evidence accumulation is utilized by the family off sequential sampling models to describe how choices are made. Such models assume that decision-relevant evidence is sampled sequentially in time and accumulated until it crosses a predefined decision threshold. The state of evidence in any given moment in time is represented by evidence accumulators. This dynamic feature sets sequential sampling models apart from other frequently encountered models in psychology such as process dissociation (Jacoby 1991) or signal detection (Green and Swets 1966), as it allows to incorporate both accuracy and decision time into a single model. Particular sequential sampling models differ in terms of accumulator dependency, shape of the accumulation (linear and non-linear variants) function or the dynamics of decision bound (is the threshold constant or does it shrink in time, representing an urgency

signal; for reviews see Forstmann, Ratcliff, and Wagenmakers 2016; Gold and Shadlen 2007; Ratcliff and Smith 2004; Usher et al. 2013).

In Projects 1 and 2 I use the Linear Ballistic Model (Brown and Heathcote 2008), which assumes that each choice option is represented by one linear and independent accumulator. Ballistic (*i.e.,* deterministic) accumulation is a simplifying assumption allowing for an analytic solution of the likelihood equation (Brown and Heathcote 2008). Psychologically meaningful parameters of the model include the speed and variability of evidence accumulation for each option, decision threshold (metric of decision caution) and non-decision time (combined time of decision unrelated processes such as sensory processing and motor execution). Similar models have been previously applied to voluntary choices in ECP paradigm (Rowe, Hughes, and Nimmo-Smith 2010; Tomassini et al. 2019; Zhang, Hughes, and Rowe 2012).

2.4.2 Reinforcement Learning

Reinforcement Learning provides a mathematical description of how internal values are updated in the face of environmental feedback. The model typically assumes learning can be explained by a Markov Decision Process in which an agent interacts with the environment in order to obtain rewards. Each action results in feedback (positive or negative reward), which provides a learning signal for the agent (Sutton and Barto 1998).

This feedback can be utilized to different degrees, dependent on how sensitive to it the agent is. Imagine a child touching a hot stove. If a child is a *perfect learner*, it will never touch a hot stove again, as it will associate the experience

with a negative outcome and stay away from all hot stoves for the rest of its life. On the other extreme, if the child is completely insensitive to feedback, then no amount of negative experiences will make a difference, as no learning will occur and the child will keep touching hot stoves never expecting the negative outcome. This intuition is captured by the Rescorla-Wagner model (Rescorla 1972), which estimates the agent's learning rate as a fraction of the feedback signal (between 0 and 1), such that

$$v_{t+1} = v_t + \alpha R \tag{2.1}$$

where v represents an action value, R represents reward, α represents the learning rate and t represents the time point.

While Reinforcement Learning has been most often applied to test how external feedback influences the internal value (L. Zhang et al. 2020), I use a modified approach, where one's choices serve as an internal feedback mechanism that can alter subsequent choices and evaluations. This process can account of choice-induced value updating in FCP tasks, proposed by many theoretical models (Bem 1967; Ariely and Norton 2007; Kaaronen and Dale 2018; Kruglanski et al. 2018).

Chapter 3

Methodological Review

This chapter provides a review of the integrative modelling approach applied in analysing the projects contained in this thesis. I first provide a brief summary of the methodological philosophy: integrative approach to modelling, the role of generative cognitive modelling, hierarchical and Bayesian analyses. I then provide a hypothetical toy example, where a synthetic dataset is simulated and analysed utilizing the Bayesian workflow guidelines (Schad, Betancourt, and Vasishth 2020). This example, similarly to the bulk of the work in chapters 4-6, is performed using *Stan* programming language (Carpenter et al. 2017) in conjunction with R (R Core Team 2018). The code for reproducing all the steps is available at https://osf.io/a3vxw.

3.1 Integrative approach

In 1982, David Marr proposed that complex systems can be grouped into three hierarchical levels of organization (Marr, 1982). The most abstract *computational* level describes the goal of the computation, *i.e., what* is the problem that the system is trying to solve - the *concept*. For many cognitive scientists this is the most crucial insight one might obtain. Brass's and Haggard's framework for volitional choice (2008) is an example of computational-level theory which describes, in the broadest terms, the goals of voluntary action, and how they can be realized. Computational-level hypotheses might also relate to more specific problems, such as how people perform categorization (Rouder and Ratcliff 2004) or how they reach a decision threshold (Evans et al. 2017).

A computational-level theory is only verifiable if its operationalisation can generate clear predictions. These can be specified at the *algorhitmic* level. The *algorhitmic* level represents the processes used to solve the computational problem - the *mechanism*. These can be described in a precise mathematical form using set of equations that represent how the computation is carried out. The algorithmic level is the most potent battlefield for hypotheses to clash. A centuries-long, yet still not concluded debate between exemplar and prototype categorization is one of the most prominent examples, where researchers specify progressively more precise models to pit the hypotheses against each other (for review see: Murphy 2016). In the realm of decision making, one of recent debates revolves around the mechanism leading to decision formation - whether it is driven by accumulating a conclusive amount of incoming evidence (Smith and Ratcliff 2004), a time-dependent urgency signal (Cisek,

Puskas, and El-Murr 2009), or a combination of the two (Miletic 2016). One might also test specific algorithms under a single computational concept. An example of this can be found in evidence accumulation literature, where assumption of accumulation is uncontested, but specific implementations of the model are compared (Donkin et al. 2011; Kloosterman et al. 2019; Urai et al. 2019; Verdonck and Tuerlinckx 2013).

Finally, the *implementation* level describes how the algorithm can be realized in a physical system. In cognitive science, the brain is the system of interest and much research is devoted to finding the neural implementation of both high level concepts (*e.g.* Gold and Shadlen 2007) and specific algorithms (*e.g.* B. Forstmann et al. 2010). Together, the levels provide a concise and informative representation of a complex system - from conceptual high-level understanding, through mechanistic latent processes, to mapping of thereof to a physical space.

Research focusing on a single isolated level, while can be successful, is also associated with a range of limitations. The most glaring example of this is the implementation level - it's nearly impossible to study the brain mechanism of a process that is not defined on a higher conceptual level. Although some attempt to derive new concepts solely from brain data (Buzsaki 2019), these ambitious attempts are few and in between (Love 2015).

Similarly, omitting the algorithmic level and mapping concepts from computational level directly to brain activity can be problematic. A historical example of this practice is mapping a high-level psychological concept (emotions), to the limbic system; an idea dating to the 1930's (Rajmohan and Mohandas 2007). Today, approaches such as this have fallen out of favour,

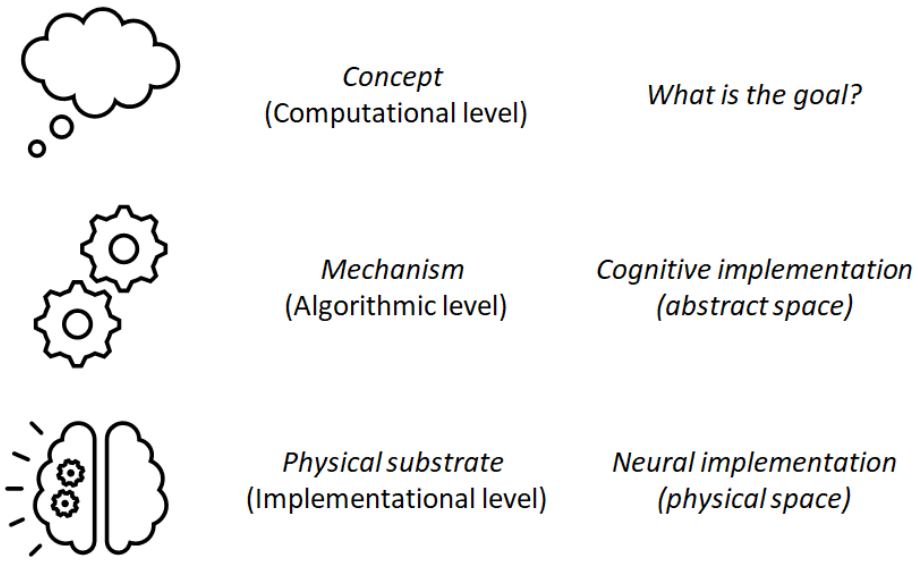

Figure 3.1: Marr's levels of analysis in relation to cognitive science. The nomenclature is a byproduct of Marr's area of expertise - computer science, and hence might be somewhat misleading from the perspective of a cognitive scientist. Firstly, in cognitive parlance, computation is often synonymous with low-level calculation or algorithm, so understanding it as a high-level concept is not obvious. Secondly, both *algorithmic* and *implementation* levels are, in fact, implementations of the computational level in abstract and physical space, respectively.

and it is easy to see why when looking from Marr's levels perspective. First and foremost, this approach provides no information about the causal mechanism driving the association. This means it has also little value in terms of practical prediction or intervention. Misconstruing purely correctional findings as causal (*"limbic system is associated with emotional experience, therefore it is their cause"*) might lead to disastrous consequences, such as performing lobotomy as a universal cure for mental illness (Kelly 1973). For these reasons, it has been referred to by some as *modern phrenology* (Love and Gureckis 2007). Additionally, similar brain-behaviour investigations run the risk of being misinterpreted in terms of the directionality of the inference. For example, amygdala activity is associated with fear (Davis 1992). Direct inference dictates conditioning amygdala response on fear induction: *if* fear *then* amygdala. However, a common fallacy is to perform a *reverse inference* (Poldrack 2006) and condition fear on the amygdala activation: *if* amygdala, *then* fear. This fallacious line of reasoning leads to false claims both in the scientific literature as well as the public eye (for review see Poldrack 2011). In reality, amygdala activation can be associated with many different mental states and emotions (Janak and Tye 2015).

Working only on computational or algorithmic levels is associated with its own hurdles. Different theories and models can make similar, or even identical predictions (Myung and Pitt 2018). In such cases, constraining the conceptual models using neural data can enable valid model comparison and favour one hypothesis over other alternatives. Recent example of this in the decision-making literature is how the previously mentioned *urgency gating* hypothesis (Cisek, Puskas, and El-Murr 2009), having a relatively

weak support relative to its direct competitor - evidence accumulation (Evans et al. 2017), received substantial empirical support when neural data was taken into account (Miletic 2016; Yau, Hinault, et al. 2020; Yau, Dadar, et al. 2020). These issues can be successfully sidestepped by filling the *missing link*: a generative cognitive model, representing the algorithmic level.

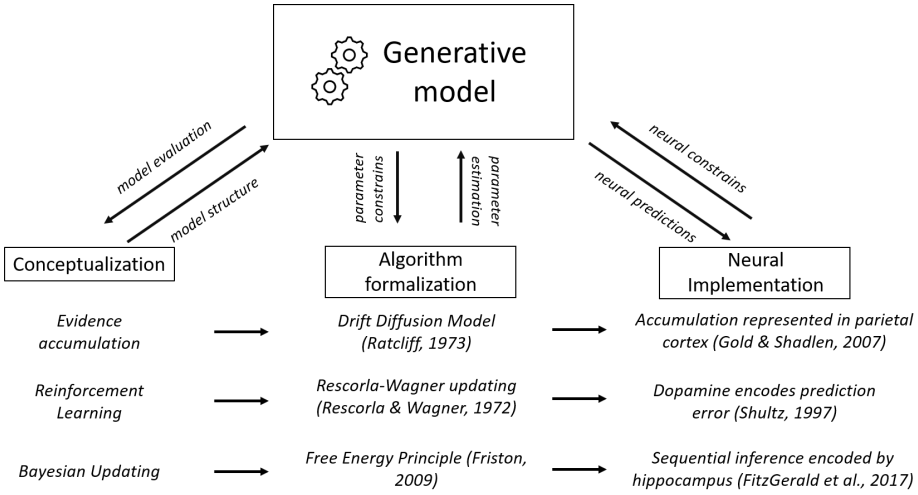

Figure 3.2: Central role of the generative model in integrative analysis framework. Here, the generative model is a formalization of the algorithmic level and can inform all levels of analysis. Conversely, all levels can constrain the model. Below are three examples of successful theories (evidence accumulation, reinforcement learning and Bayesian updating), their algorithmical specifications, and neural correlates of specific parameters.

3.2 Role of generative cognitive modelling

Generative modelling is defined as a statistical approach where modelling is intentionally aligned with theory (Haynes 2019). In other words, it puts emphasis on finding and fitting a model which can explain how the observed

data was generated. This is contrasted with more traditional modelling approach which is focused on finding the best possible fit, without claims about how true the model is with regards to the data generating process.

This means that the main difference lies in the intention of the modeler and his belief whether his model can capture the data generating process. A simple linear regression is usually used as a descriptive or predictive tool, to show which factors are meaningfully associated with the outcome, or to attain the best possible future prediction. It could however be also considered a generative model, in the specific case when the modeler believes that the cognitive process he is describing is aking to a linear square-error minimization (McElreath 2016). A generative way of thinking about a simple regression process in described in the Toy Example in section 3.4.1. Similarly, many sophisticated models that fit behavior well, such as signal detection theory in describing metacognitive confidence (Fleming and Dolan 2014), neuronal cosine tuning functions (Todorov 2002) or representational similarity matrices (Kriegeskorte, Mur, and Bandettini 2008) make no claims about the data generating process, but simply do very well in succinctly describing the data.

Cognitive processes are *latent*, *i.e.,* they cannot be directly observed. Generative models of cognition require a set of assumptions about the data generating processes, and use behavioural and neural measures as proxies to infer these latent mechanisms. Throughout the years, the insufficiency of using purely data-driven approaches has been more widely recognized and theoretically-driven generative models become more prevalent in psychology (Farmer, Brown, and Tanenhaus 2013; Haynes 2019; Navarro 2020; Strube 2000) and neuroscience (Ahn, Haines, and Zhang 2017; Betzel and Bassett

2017; Frässle et al. 2018; Friston and Price 2001).

From the perspective of an integrative approach, generative modelling plays a crucial role in the analysis framework as it provides a formal causal structure that can be tested directly (Love 2015).Generative modelling spans across and influences all levels of Marr's framework (Figure 3.2). Model theory and general assumptions are specified on the computational level. Based on the conceptualization, specific algorithms can be derived, simulated, and tested against real data. A well designed model should generate accurate predictions about behaviour and neural activity. For example, the Drift Diffusion Model (Ratcliff 1978) predicts that higher decision threshold should be associated with increased accuracy and reaction times. Neural extentions of the model further show that increased threshold is determined by the functional coupling between the subthalamic nucleus oscillations and the medial prefrontal cortex activity (Herz et al. 2016). Data can also constrain the model. An obvious example is constraining the parameter values by fitting. Sometimes however, the model might not be able to reproduce important data patterns, in which case the model needs to be adjusted, or refuted altogether. For example, independent evidence accumulation cannot account for Hick's Law principle, stating that increasing the number of alternatives leads to slower, not faster responses (Proctor and Schneider 2018). To accommodate this, independent accumulation can be discarded in favour of approaches suggesting dependent accumulation or cross-inhibition (Ratcliff, Voskuilen, and Teodorescu 2018; Usher and McClelland 2001), modified to accommodate the effect (van Ravenzwaaij et al. 2020), or limited to inference within specific situations where the number of alternatives is

fixed. Alternatively, a better fitting or simpler model can be build, rendering the previously favoured alternative obsolete. Recent work on joint modelling (de Hollander, Forstmann, and Brown 2016; Brandon M. Turner et al. 2016; B. Turner et al. 2016; Brandon M. Turner, Wang, and Merkle 2017) provides a rich framework on how to combine neural and behavioural data in ways that are appropriate for the experimental goal, by either constraining the models with neural data, inferring brain function by conditioning on the model, or modelling both jointly.

3.3 Hierarchical Modelling

Hierarchical models (in frequentist literature also often referred to as mixed models; Bates et al. 2014) aim to accurately model the variance structure in the data by clustering together correlated observations, where measurement units can be organized into groups (Hofmann 1997; Stryhn and Christensen 2014). For example, observations such as individuals from a certain country or grades from a certain school are drawn from the same pool, hence belong to a single cluster. This assumption makes sense, since we can expect that people from different countries might behave differently, and different schools might have different grading criteria. In a typical cognitive experiment, a hierarchical structure involves group-level measures and individual-level measures, enabling to distinguish *between* and *within-person* effects. Thanks to this, in a *repeated measures* design where participants engage in multiple trials of a single condition, the model can account for every single observation.

The grouping (or lack of thereof) of particular observations is referred to

as *pooling*. A model with complete pooling assumes all observations belong to a single group, while a model with no pooling assumes all data-points are independent. Hierarchical models utilize partial pooling, which improves predictive validity of the models (Anders, Oravecz, and Alario 2018; Barr et al. 2013; Baayen, Davidson, and Bates 2008). This can be contrasted with a traditional ANOVA approach, where due to the assumption of independence, each individual can only be represented by a single value (usually the mean). Conversely, a fixed-effects (non-hierarchical) linear regression would use all datapoints but assume they all belong to a single group. This approach should be avoided when dealing with hierarchical data, as it breaks the assumption of independence.

3.3.1 Bayesian approaches

It is important to distinguish between Bayesian *analysis* and Bayesian *theory of cognition*. The former is a set of mathematical tools constructed to analyse data in accordance with Bayes Theorem. As such, it can be applied to any statistical test or model, irrespective of the modelling approach or the level of analysis. The latter represent the belief that human cognition approximately follows the tenets of Bayesian probabilistic inference (Colombo and Hartmann 2017). My work utilizes Bayesian approach to data analyses, however remains agnostic with regards to Bayesian Cognition.

Bayesian analysis utilizes Bayes Theorem, an axiom probability theory which indicates how to update one's beliefs in the face of new evidence. The theorem derives posterior belief from the factorization of one's prior belief and the likelihood of the observations, divided by the normalizing constant

(accounting for all possible events). In data analysis, it is most often used in estimating the posterior probability of a hypothesis (operationalised as a set of parameters or a model), given the likelihood of the data (probability of the data given the hypothesis) and the priors. The priors represent an inherently subjective belief about the state of the world. In contrast, frequentist analysis rejects subjectivity, focusing entirely on analysing frequencies of events (for review see Dienes 2011).

Bayesian approach to data analysis presents several advantages compared to its frequentist counterpart. First of all, Bayesian analysis gives a straightforward answer for the question researchers are interested in: what is the probability of a hypothesis being true, given the observed data. Frequentist approaches can only provide a reverse inference about the probability of the data given a hypothesis, *i.e.,* the likelihood. This likelihood is conditioned on the null effect being true. Consequently, a p-value represents the probability of an outcome being this extreme, given that the null hypothesis is true. Many argue this approach is both counterintuitive and impractical, since it gives an answer to a question we are not interested in (Kruschke and Liddell 2018; Wagenmakers et al. 2018). This is why the interpretation of frequentist p-values is often misconstrued, even among researchers and educators (Gigerenzer 2004).

Secondly, Bayesian analysis naturally incorporates uncertainty into parameter estimation, resulting in posterior parameter distributions, which variance accounts for estimation uncertainty. Frequentist methods can only provide parameter point estimates, and require additional assumptions to derive indirect measures such as confidence intervals: similar to p-values, they

need to be conditioned on the assumptions of repeated sampling from a null distribution.

Other, practical advantages of Bayesian analysis include: the ability to directly compare evidence in favour of the null (and not only against it; Wagenmakers et al. 2018; Dienes and Mclatchie 2018); not being tied to a fixed sampling plan (Bayesian researcher can stop the sampling process at any arbitrary moment, which will not affect the inference process; Hackenberger 2019; Wagenmakers et al. 2018) and a wide range of applications where Bayesian estimation provides better parameter estimation, especially when using hierarchical models (Ferreira et al. 2020; Hong et al. 2013; Piray et al. 2019; Rouder and Lu 2005).

While the benefits of Bayesian approach are getting more widely recognised and its applications have risen exponentially in recent years (Hackenberger 2019), it is important to recognize that Bayesian model estimation can be very computationally expensive and time consuming, hence difficult to implement in some data-rich areas (Green et al. 2015).

3.4 Practical example

Throughout this thesis, I use *Stan* programming language (Carpenter et al. 2017) for model building and estimation, together with *R* language (R Core Team 2018) for all other analyses, such as data preprocessing, model-free analyses, plotting and simulating model predictions. This section focuses on *Stan*, providing a high-level overview of modelling using a toy example, which is a simplified problem thematically relevant to Project 3. All code to

this section can be found at https://osf.io/a3vxw.

Stan is a probabilistic programming language designed for specifying and estimating statistical models. Although not exclusively, it is most commonly used for Bayesian inference (Steyvers, Lee, and Wagenmakers 2009). The choice of *Stan* in my work is guided by a few considerations which in conjunction make it, in my belief, one of the best currently available alternatives.

Inherently Bayesian

In section 3.3.1 I have described why Bayesian estimation is often considered more desirable for generative cognitive modelling, compared to traditional maximum likelihood approach. Stan is designed and optimized around the idea of Bayesian Estimation. The language is expressed in terms of probabilistic sampling statements and has wide range of built-in functions for estimating priors, likelihoods, and posterior predictions (Carpenter et al. 2017).

Efficient Sampling

Sampling is what in my view sets *Stan* apart from some of its Bayesian counterparts. *Stan* utilizes *Hamiltonian Monte Carlo (HMC)* Sampling with *No-U-Turn* sampler (Betancourt 2017). In terms of sampling efficiency, HMC performs better than Metropolis-Hastings Mote Carlo or Gibbs Sampling used in other software predating Stan, such as WinBUGS (Lunn et al. 2000) or JAGS (Plummer, 2003; for a comparison see: Annis, Miller, and Palmeri, 2017). Specifically, it is efficient in sampling from models with a larger

number of intercorrelated parameters – a common case in many cognitive models (Annis, Miller, and Palmeri, 2017). Other popular alternatives include *Dynamic Models of Choice (DMC) R* package (Heathcote et al. 2019) specializing in Bayesian Estimation of cognitive models which utilizes a population MCMC algorithm called *differential evolution Monte Carlo* (Turner et al. 2013; although from personal experience, not as nearly as efficient as HMC), and *PyMC* (Salvatier, Wiecki, and Fonnesbeck 2016), a general purpose Python library designed for Bayesian Modelling.

Accessibility

Stan has an extensive manual, and provides interactive interface packages for *R*, *Python* or the command line. In addition, it has been widely adopted in recent cognitive literature, providing many excellent empirical papers (Romeu et al. 2020; L. Zhang and Gläscher 2020) as well as tutorials (Ahn, Haines, and Zhang 2017; Annis, Miller, and Palmeri 2017), which mitigates the initial hurdle of learning a new programming language.

In the demonstration, I will *1)* present a toy problem, *2)* propose a data generative process and simulate a synthetic dataset *3)* build a generative model capable of recovering the underlying process *4)* run prior predictive simulations *5)* fit the model to the simulated data *6)* generate posterior predictions and compare them with the synthetic data.

3.4.1 Toy Problem

Choices are known to influence item evaluation. In a famous study, Brehm (1956) asked participants to rate a list of household appliances in terms of subjective value. Following the evaluation, participants were presented with a choice between 2 of the items. Finally, participants rated all the items again. Brehm found that on the second rating the chosen item was rated higher, while the rejected one lower, as compared to the initial rating. Although there has been a debate about the legitimacy of this effect in recent literature (Chen and Risen 2010), the current consensus seems to support it (see Chapters 2 and 6).

Following this line of reasoning, we want to test whether there is a linear effect between the amount of times a given item was chosen and the difference in its estimated value before and after the choices were made. The hypothesis to test is that the more times a given item was chosen, the greater the increase in it's subjective value. For that, we use a design where participants first rate a set of 50 items on a continuous scale (0-100; referred to as *rating 1*) in terms of their desirability, then make repeated 2-alternative forced choices between the items (all possible item pairs give a total of 50 * 49 / 2 = 1225 trials), and finally rate each item again (*rating 2*). Our hypothesis states that there is a linear relationship between how often an item was chosen and the increase in its value measured by the difference between the second and first rating. For this hypothetical example, we are using a rather unrealistically large number of trials. This is done for the purpose of having enough datapoints (50) to reliably simulate the assumed correlation structure between the variables. For the sake of brevity, we will perform the simulation and analysis only on

a single hypothetical participant and then discuss the extensions necessary for performing a hierarchical group-level analysis.

3.4.1.1 Data generative process

Ratings. We can now translate our hypothesis about a linear relationship to a statement about how the data is generated. Our variable of interest is the difference between rating 2 and rating 1. Let us assume that the true difference in ratings is centered at 0 (i.e., there is no systematic effect of value mean between ratings) with a standard deviation of 15:

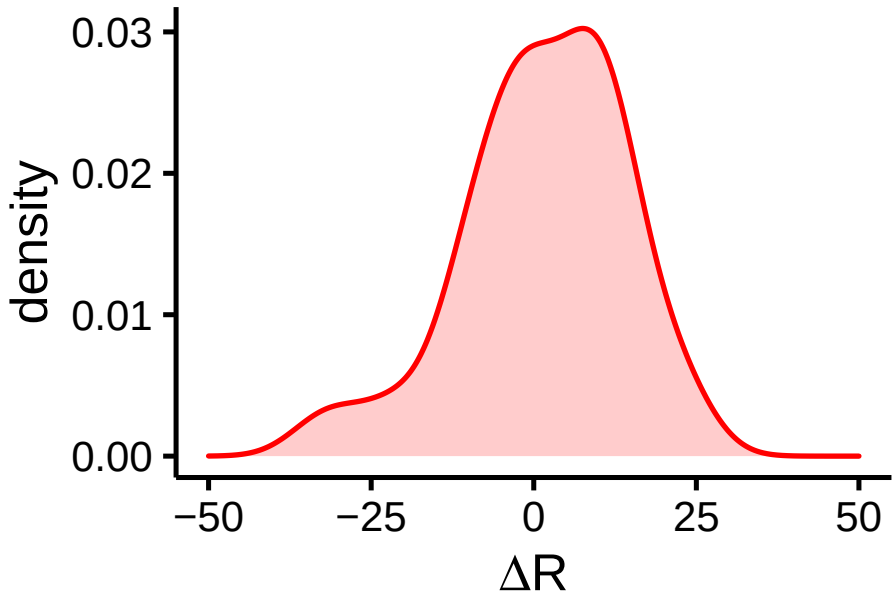

Figure 3.3: Density of the simulated difference in ratings

$$\Delta R_k \sim N(0, 15) \tag{3.1}$$

where ΔR_k is the difference in rating for the *k-th* item. Positive values indicate higher rating values in rating 2. Since our rating scale is finite (0-100), we need to make sure that values do not exceed the possible bounds. To do this the distribution needs to be truncated so that no values are larger than 100 or smaller than -100. The simulated distribution is plotted in Figure 3.3. The heavy left tail observed in the data, even though it was sampled from a symmetrical Gaussian distribution, reflects the variability associated with a limited number of samples.

Choices. The 2AFC task consists of 1225 choices, where each item is presented 49 times, so that 49 is the maximum number an item can be chosen. We can simulate a random distribution of binary choices by modelling them as drawn from a binomial distribution:

$$\Delta Chosen_k \sim Binom(N = 50, prob = 0.5) \tag{3.2}$$

where $Chosen_k$ is the number of times *k-th* item was chosen, N is the total number of items and *prob* is the probability of choice. Additionally, we assume that the choices are correlated with the ΔR_k variable (for detailed implementation see attached code). This synthetic dataset represents our beliefs in the data-generating process, *i.e.,* that choosing an item linearly increases it's subjective value.

Notice some simplifying assumptions we have made when specifying the

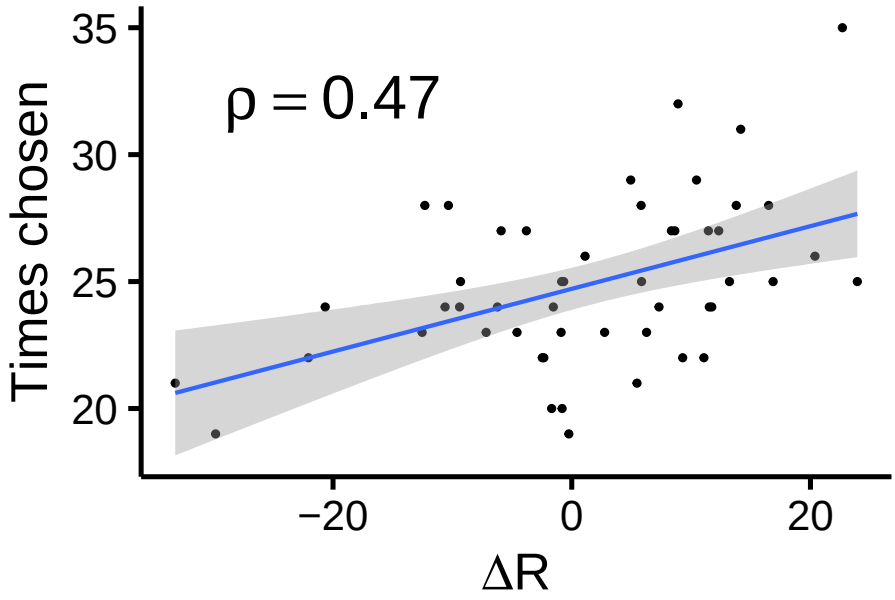

Figure 3.4: Correlation between the difference in ratings and choices (total number of times an item was chosen).

model:

- we do not model individual ratings, but both of them together (*i.e.,* only the difference)
- we aggregate over all choices made in favour of a given item, simulating only the sum of how many times it was chosen, and not individual choices
- we assume choices are independent of item value

While the first two simplifications should not affect our inference process too strongly, the third one is more problematic, as we are running the risk of simulating unrealistic data. While this is not a major issue when performing an illustrative example, an applied simulation study performed for the purpose of designing an experiment would require a more involved approach modelling each individual choice and accounting for the value-choice correlation structure (an example of such approach was implemented in Project 3).

3.4.1.2 Building a Generative Model

The next step is to build a model and check whether the model fit can recover true data-generating parameters. In Stan, we should explicitly specify a *reasonable* set of priors, and a likelihood function.

The likelihood expresses our belief about the data-generating process. We believe that the difference in value between two ratings (outcome variable y)

is a linear function of choice frequencies (predictor variable x). This can be expressed by a sampling statement:

$$y \sim N(\alpha + \beta * x, \sigma) \tag{3.3}$$

where α , β and σ represent the intercept, slope and error scaling, respectively (for exact derivation see McElreath 2016, chap. 4).

In addition, Bayesian models must have a prior. Priors are not influenced by the data and therefore are inherently subjective. Priors should reflect our initial beliefs while providing reasonable constraints. They should not be too narrow, as we do not want them to influence the outcome too strongly. They should also not be too broad, as this will lead to inefficient sampling and, especially in more complex cases, may lead to model unidentifiability (Smid et al. 2020). We will use generic informative priors (Gelman, Lee, and Guo 2015) which are both realistic and sceptical (centered at null effect), but also broad enough for the data to have a decisive influence:

$$\alpha \sim N(0, 1) \tag{3.4}$$

$$\beta \sim N(0, 2.5) \tag{3.5}$$

$$\sigma \sim cauchy(0, 2) \tag{3.6}$$

The intercept is centered at 0, since we will be using a scaled predictor - we will subtract the mean from the number of choices and divide it by its standard deviation (sd), making sure it is centered at 0 and has a sd of 1.

This has three main advantages when using linear models. Firstly, it reduces the collinearity between predictors in multiple regression models. Secondly, it makes the model intercept interpretable, as it becomes the mean of the outcome variable. Thirdly, it allows for using "safe" standard priors for any predictors, independent of their true mean and variability.

Error is drawn from a cauchy distribution, which has a half-gaussian shape with parameters describing its mode and scale. This is a popular choice for an error prior, since variability cannot be negative.

3.4.1.3 Prior Predictive Simulations

Before fitting the model to data, it is important to first examine whether the priors are reasonable. Prior predictive checks allow to generate predictions from the model even before supplying the outcome variable. This technique is very useful especially in more complex models, as it can identify misspecified priors. If the generated data are unrealistic or inconsistent with the domain of expertise, it is likely the priors need to be adjusted (Schad, Betancourt, and Vasishth 2020).

Our priors are skeptical (centered at 0), so we should expect to simulate a null effect. The observations should also be realistic. This means that, ideally, they should be indistinguishable from real data. Practically, one needs to define a set of metrics assessing the similarity. In our case, we would want our simulated ΔR_k values to have a decently large spread around 0, so that the prior can accommodate a moderately large effect, if one is found in the data.

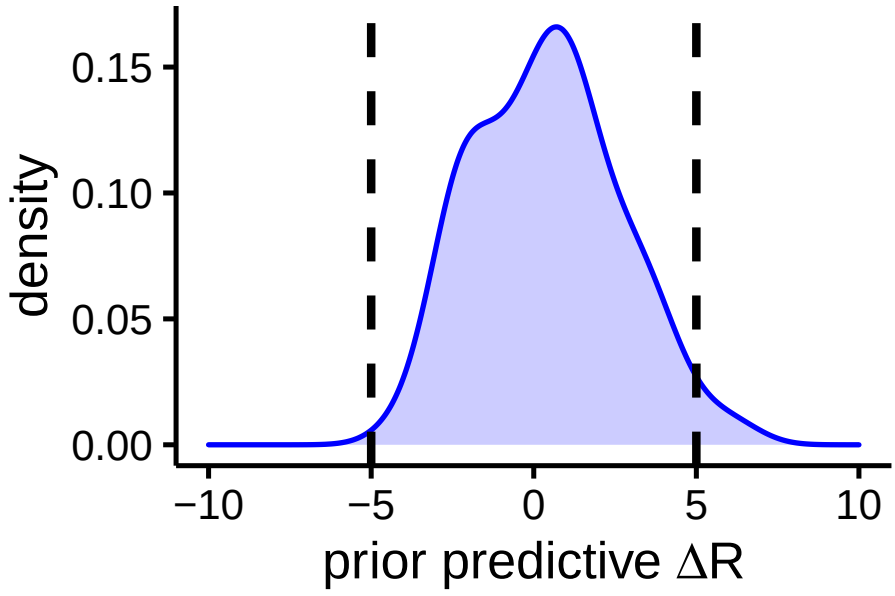

Figure 3.5: Prior predictive density for the regression model.

By plotting the model prediction, we can see that the prior is skeptical of effects larger than a 5 point difference between the two ratings, assigning them only ~2% probability. While this might be a perfectly reasonable prior, not knowing the actual effect sizes, we should err on the size of caution, and allow the prior account for more extreme effect sizes with more probability. To achieve this, we can modify the prior on σ, increasing the estimation error:

$$\sigma \sim cauchy(0, 6) \tag{3.7}$$

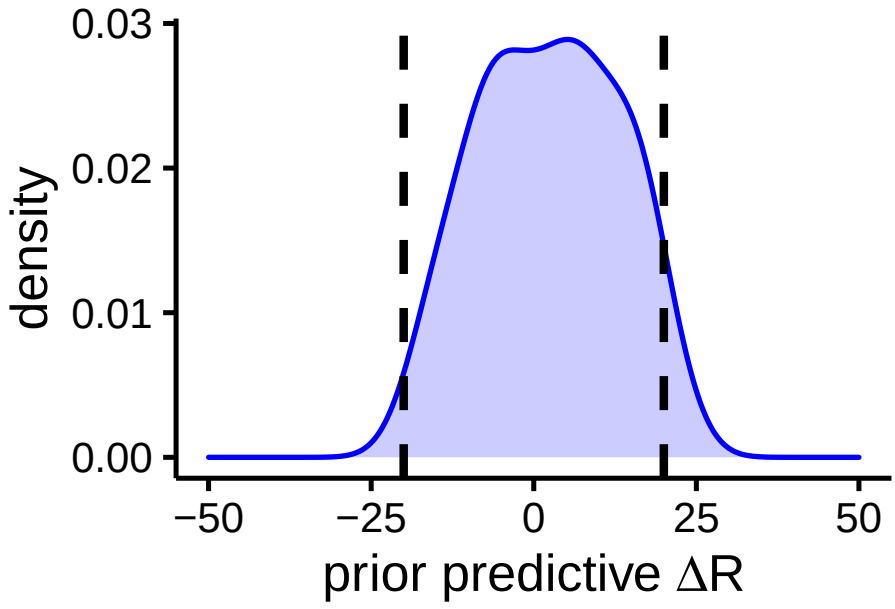

Figure 3.6: Prior predictive density for the second model with increased error term distribution.

Table 3.1: Posterior estimates of fitted parameters.

	mean	se	2.5%	50%	97.5%	Rhat
alpha	0.421	0.014	-1.239	0.423	2.130	1
beta	4.135	0.025	1.297	4.170	6.785	1
sigma	11.285	0.021	9.250	11.193	13.856	1

Our new predictive prior is much more agnostic, allowing for both large effects with high probability (~2% of probability for values larger or smaller than 20; Figure 3.6). It seems more appropriate for a novel paradigm, when we and do not know what effect size to expect. This illustrated the importance of running a prior predictive checks, especially when using default priors to a new problem.

3.4.1.4 Fitting the Model and Parameter Inference

We can now fit the our model to the recorded responses using 4 Markov Chain Monte Carlo (MCMC) chains, 4000 samples with a 2000 burn-in each. The fit provides posterior distributions for parameters of interest: α, β and σ (Table 3.1).

The Gelman-Rubin $\hat{R}$ statistic indicates how well the independent chains have mixed (Gelman and Rubin 1992). Intuitively, they can be interpreted as a ratio of between to within-chain variance. Values below 1.1 are considered as indicative of good convergence (Annis, Miller, and Palmeri 2017). To infer whether the model predicts a significant linear relation between the number of choices and difference in ratings, we examine the β parameter. Median $\beta = 1.297$, with a 95% Bayesian Credible Interval of [0.025, 4.170]

and probability of directionality (pd) = 0.996. Pd is the fraction of posterior distribution > 0 and can be interpreted equivalently to a frequentist p-value (Makowski et al. 2019).

3.4.1.5 Generating posterior predictions

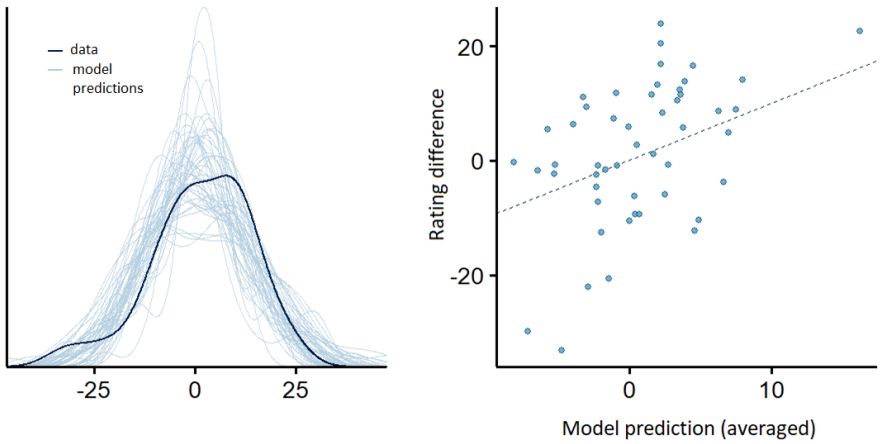

Figure 3.7: Model posterior predictions. Left panel shows the density distribution of data (darker color) compared to 50 randomly sampled posterior samples. Right panel indicates the correlation between the data (y axis) and posterior model predictions, averaged across 2000 samples.

Finally, we should look whether the estimated parameters can accurately predict the data. In our simple case, we can look how well sampled predictions approximate actual data in terms of the distribution overlap and correlation (Figure 3.7; see also: L. Zhang et al. 2020) Left panel shows that the model can reproduce the shape of the real data very well. Averaged model predictions ideally should reduce to the line of best fit, which we can see on the right panel.

3.4.2 Advantages of a Bayesian Approach

In such simple example, one could question the practicality of a Bayesian approach is dubious. Is it really worth the effort to go through all these steps, when a similar frequentist analysis would provide similar results and conclusions?

Indeed, more apparent differences could be observed if the problem was hierarchical (*e.g.* by adding group-level parameters) - in such cases Bayesian methods tend to shine the most, often outperforming frequentist methods with regards to the accuracy of parameter estimation (Ferreira et al. 2020; Hong et al. 2013; Piray et al. 2019; Rouder and Lu 2005). Even in this simple case however, two arguments can be made in favor of using the Bayesian approach. A theoretical argument relates to the philosophy of statistical analysis. The Bayesian approach offers a straightforward answer to the question we are interested, *i.e.* does the number of choices influence the difference in value. In contrast, in order to perform a frequentist analysis one has to *a priori* assume that the null hypothesis is in fact true, and test the probability of the data under that assumption. While the answer is similar, it is also clear that the question we can ask using the frequentist method is not the one we are really interested in (Kruschke and Liddell 2018; Wagenmakers et al. 2018). A practical argument is that by estimating parameter distributions instead of single points, Bayesian estimation naturally naturally provides uncertainties around parameter estimates, which can only be indirectly inferred when using the frequentist approach.

3.4.3 Next Steps

This toy example showcases the basic procedure of building a cognitive model. Such simulations can often inform future experimental design. Model simulations can reveal if, for example, the number of trials, participants or conditions is not sufficient for distinguishing between different mechanisms, of having enough power for finding the effect of interest. In real applications, next step would involve making the model hierarchical, *i.e.,* adding group-level parameters and running the model for a whole experimental cohort.

While building a model, one should also be very mindful of its realism. The toy example has a crucial inaccuracy - the simplification that choices are independent of value, which is not true for real data. People's choices tend to be rational, which in this case means that higher rated options should have a greater probability of being chosen. Modelling every choice would require using a more involved approach, where we would track all values of individual items. The choices then could be generated by a softmax function, taking as input items' values on a trial-by-trial level (Ahn, Haines, and Zhang 2017). Such approach was implemented in Project 3.

Finally, for many applications it is often advantageous to consider more than one generative process, and use model comparison techniques (e.g. Vehtari, Gelman, and Gabry 2017) to estimate which can account better for the observed data. For example, a good test in this case would be an *artifact effect* model, assuming that the change in rating is simply driven by a regression to the mean (Chen and Risen 2010; Izuma and Murayama 2013).

3.5 Summary

This chapter provided an overview of the methodological approach used throughout my thesis. All projects use a hierarchical approach to data fitting. All also utilize generative modelling, albeit to a different extent, since the modelling goals are different. In Projects 1 and 2, a well-established cognitive model (LBA) is used to describe how choices arise, and linking the latent process with neural data on a subject (Project 1) or trial (Project 2) level. In Project 1, the model is estimated on behavioural data, which parameters are then correlated with neural activity, with the goal of inferring a neural implementation of the latent process. In Project 2, the model is constrained by incorporating the neural data (trial-level EEG signal) as a predictor of model parameters. Different model variants are compared to obtain the most accurate description of the underlying computation. Modelling in Project 3 focuses on building an entirely new model, capable of explaining the experimental effects of interest.

Chapter 4

Choices between equal options: impact of varying levels of choice freedom and discriminability (fMRI study)

4.1 Introduction

Human volition is often described as the feeling of freedom over own actions, as well as the control over their implementation. Acting out of your own volition requires us to make a series of decisions regarding *whether*, *when*, and *how* to act (Brass and Haggard 2008). For example, consciously choosing to move a leg requires an explicit intention which specifies the direction and strength of the motion in accordance with our goal. Such intention needs to be followed by correct execution, mediated by activating appropriate motor action program.

Volition refers to both the act itself (Haggard 2019), as well as the subjective experience of the act (Haggard 2017), which is a metacognitive judgment of agency and freedom. Voluntary choice is usually studied using paradigms that present participant with freedom in one of free domains: whether to perform an action, when to perform an action and what to choose (Brass and Haggard 2008). Here, we focus on the third paradigm - that is the study of volition under the constraints of a decision-making paradigm. In this domain of study, the focus lies in eliminating the confound of external stimuli properties on choice. In order to do so, a typical task involves choosing between a number of identical options, compared to a control condition, when only one option is available to choose (*e.g.* Rowe, Hughes, and Nimmo-Smith 2010; Zhang, Hughes, and Rowe 2012).

Early attempts at describing the neural underpinnings of free choice trace back to famous experiment done by Libet and colleagues (1983), indicating that the EEG readiness potential over central electrodes associated with

an upcoming movement precedes the conscious experience of action. More recently, attempts have been made at describing the neural network underlying internally driven choices. Pre-supplementary motor area (preSMA) is believed to be a key hub of this network - it integrates inputs from basal ganglia and prefrontal cortex and relays them to motor cortex, which executes the action (Nachev, Kennard, and Husain 2008). PreSMA has been shown to play a critical role in controlling motor action in situations of response conflict (Nachev, Kennard, and Husain 2008), switching between actions (Nachev et al. 2005) and inhibiting automatic responses (Sumner et al. 2007). Due to its strategic neuroanatomy and task-related activity, preSMA has been proposed to serve the role of binding intention and action, which is achieved by controlling motor output via motor cortex inhibition (Haggard 2008). Modelling approach suggests that its activity in equal-choice tasks can be well approximated by an accumulation-to-bound process (Rowe, Hughes, and Nimmo-Smith 2010; Tomassini et al. 2019; Zhang, Hughes, and Rowe 2012). In this model, evidence in favour of each choice is accumulated to the point when one of accumulators reaches a decisional threshold, after which the choice is executed (Forstmann, Ratcliff, and Wagenmakers 2016). We aim to build upon these results by testing the effects varying levels of freedom (1, 2 or 3 available alternatives), perceptual discriminability (high vs low), and the link between these experimental manipulations and subjectively experienced freedom.

4.2 Aims and Hypotheses

The study aims at describing the three factors associated with voluntary choice together with their neural correlates. These include:

1) Testing 3 levels of choice freedom
2) Examining the influence of stimulus discriminability on the voluntary choice process
3) Testing how objectively defined freedom of choice translates to subjectively experienced feeling of freedom

Since our goals include describining both the behavioural effects, as well as their neural correlates, we utilize the functional MRI technology, which allows us to test the neural underpinning of these effects, reflected in the blood-oxygen-level-dependent (BOLD) signal. FMRI technology is a noninvasive brain imaging technique, known for it's high spatial resolution, compared to other noninvasive techniques, such as EEG or MEG (Aine 1995). This makes fMRI well suited for studies such as this, which focus on finding specific functional areas involved in the process of interest.

1) Experiments that aim at explaining the brain mechanisms of free choice using a multiple-option decision paradigm usually dichotomize the freedom by presenting participants with either an instructed or a free choice (notable exceptions being Forstmann et al. 2006; Lau et al. 2004; van Eimeren et al. 2006). By using only one level of free choice however we cannot differentiate between a set of possible mechanisms driving the brain activity. Possible explanations for increase in preSMA activity with free choice include:

threshold mechanism (higher activity when choice is free, regardless of the degree of freedom), a linear increase (more freedom leads to more activity), or a different nonlinear function. In line with studies suggesting preSMA involved in evidence accumulation (Rowe, Hughes, and Nimmo-Smith 2010; Tomassini et al. 2019; Zhang, Hughes, and Rowe 2012) we predict that preSMA will be sensitive to the number of available alternatives. We also expect to observe longer reaction times (RTs) with increasing number of options, an effect consistent with Hick's Law (Proctor and Schneider 2018).

2) Stimulus discriminability affects behavior – more salient stimuli are associated with faster response times (Teodorescu, Moran & Usher, 2016). Our design allows to test whether discriminability also affects brain processing of voluntary choice. We expect lower discriminability (LD) condition to be associated with longer RTs and activity in the network related to increased difficulty and conflict, including dorsolateral prefrontal cortex (DLPFC), the anterior cingulate cortex, insula and the inferior parietal lobule (Cole and Schneider 2007; Ho, Brown, and Serences 2009; Keuken et al. 2014).

3) We strive to see how freedom of choice, defined as the number of options available to choose, relates to the subjective experience of freedom. Rating subjective feeling of freedom and agency is a metacognitive process of assessing one's own control over the action (Sidarus and Haggard 2016).

Previous studies have identified a negative predictor of agency in the temporo-parietal junction (TPJ; Sperduti et al. 2011; Zito, Wiest, and Aybek 2020), but no consensus was established regarding brain areas positively associated with agency (Zito, Wiest, and Aybek 2020). By asking participants to assess their freedom after making a choice we aim to establish a) how does objectively

defined freedom influence one's subjective perception b) what are the neural correlates supporting this judgment. Behaviourally, we expect to observe effects of number of available alternatives (from 1 to 3) and discriminability on freedom rating. In particular, we predict HD condition will be associated with higher perceived freedom, due to lower difficulty (Sidarus and Haggard 2016).

4.3 Methods

4.3.1 Participants

Participants were recruited from Cardiff University participant panel. Our sampling plan included obtaining data from 30 participants. In total, we recruited 38 participants, since 4 did not show up for the fMRI sessions and another 4 were excluded due to poor quality of the fMRI data due to excessive movement or falling asleep in the scanner. The final sample consisted of 30 participants (24 females, 24 right-handed, $M_{age} = 20.89$, $SD_{age} = 2.51$). All participants had normal or corrected-to-normal vision, and none reported a history of neurological or psychiatric illness. Written consent was obtained from all participants. The study was approved by the Cardiff University School of Psychology Research Ethics Committee.

4.3.2 Behavioural Task

Motivation

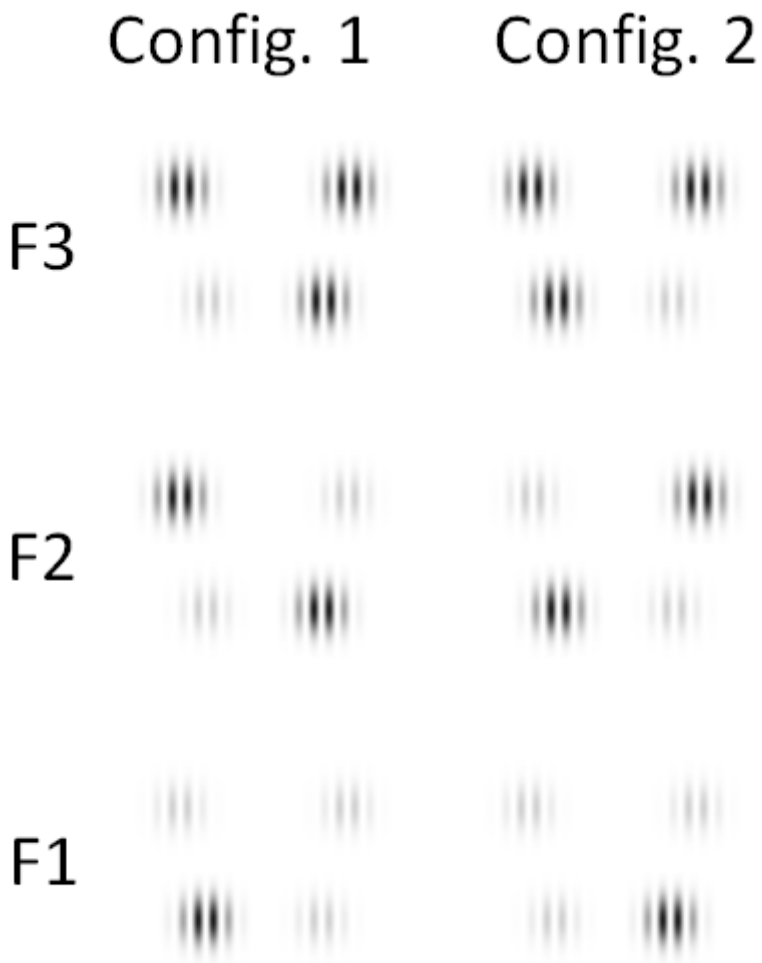

Figure 4.1: Possible configurations of available Gabor patches per condition. High contrast patches were available to select. The contrast in the figure represents contrast used in the HD condition (100% vs 20%). The configurations were balanced so that overall, each patch was equally likely to be available.

The behavioural task was designed to test how two main factors of choice freedom (3 levels) and discriminability (2 levels) affect reaction times and accuracies in an equal choice paradigm. Our goal was to set the discriminability bar relatively low in order to obtain high levels of accuracy. The idea behind this approach was to avoid the task being perceived by the participants as a perceptual discrimination task, where getting the correct answer is the main goal, but rather a free choice task, where a brief glance can make it obvious how many alternatives are available, and one can focus on the processes of choosing between equal options. This assumption made manipulating discriminability challenging, as we aimed at both high accuracies and a measurable difference in difficulty, as observed by longer reaction times in the lower disciminability condition. In order to achieve this balance, we first performed a series of behavioural pilots, manipulating the disciminability levels. Apart from Gabor patch contrast levels, which is how the disciminability was manipulated in the final version of the task, we also tested patch orientation (i.e. choice-available patches being signified by the degree to which the patch was tilted). Manipulating contrast provided better results (desired high accuracies in all conditions similar to the ones observed in the main study; see *Results* section), hence this design was used.

Final Task In the final version of the experimental task participants chose between 4 Gabor patches displayed on a gray background. After every few trials participants rated their subjective feeling of freedom of the last choice on a discrete 4-point scale. A rating trial occurred after 3 to 5 choices, drawn pseudo-randomly from a uniform distribution so that the ratio of choice-to-rating ratio was exactly 4:1. Option availability (whether

a given patch is available to choose) was signaled by the patch contrast. Participants were instructed to choose any high-contrast patch and avoid choosing low-contrast patches. Unavailable patches had 20% contrast, while available patches had either 100% (high discriminability condition; HD) or 60% contrast (low discriminability condition; LD). Available patches in any given trial had always the same contrast. Additionally, we controlled for available configurations, so that only 2 unique configurations were possible for each cell design (see Figure 4.1).

4.3.3 Design

We have set a 3 (Alternative Availability) by 2 (Discriminability) factorial design. Alternative Availability factor controls the number of options that are available for the participant to choose in a given trial, from 1 to 3 (further denoted as F1, F2 and F3). Discriminability controls for the ease of differentiating between available and non-available options. Each condition consisted of 60 trials in a given experimental session (240 choice trials per session plus 60 rating trials) Each session was divided into 6 40-trial blocks, in between which participants could take a short break.

4.3.4 Procedure

A white fixation cross appeared at the start of each choice trial for 400 ms, followed by the 4 patches appearing on the screen for 1000 ms. Participants had 2000 ms (including the initial 1000 ms during which the stimuli were present on the screen) to make their choice, which was signaled by the cross

turning green once the choice has been recorded. A 800-1200 ms period (randomly drawn from a uniform distribution) with a white fixation cross in the centre followed each choice, followed by a rating trial (every 3 to 5 trials for 2000 ms) and a jitter period drawn from a uniform distribution between 2000 and 5000 ms.

4.3.5 Stimuli

The batches were 2.25 visual degree in size, placed 5 visual degrees from the centre of the screen. Visual angles of the batches (starting from left to right) were set to -63 degrees, -136.5 degrees, 136.5 degrees and 63 degrees (see Figure 4.1). The horizontally asymmetric patch placement was utilized so that the patch location would easily correspond to 4 fingers (from left to right: left middle finger, left index finger, right index finger and right middle finger) which participants used to respond. The responses were provided using a *NATA* button box.

The freedom scale was displayed as a white line with 4 points, going from 1 to 4, where 1 represented the lowest and 4 the highest level of perceived freedom. Similarly to the choice task, participants used the 4 fingers as corresponding the 4 scale values, mapped from left to right.

The rating trials did not have any text prompt. Participants were only instructed prior to the task, that the scale relates to 'How free they feel their previous choice was' with 1 being the least, and 4 being the most free. The instructions emphasized that the rating relates to a subjective feeling one might experience.

4.3.6 Experimental paradigm

The experiment consisted of 3 sessions held on different days for each participant. On the first session participants were trained in the task and performed 240 trials of the main task. Second session consisted of performing 240 trials of the main task (27 minutes), a structural T1 image acquisition (15 minutes) and two visual field mapping tasks (20 minutes), all done inside the MRI scanner. Third session consisted of 240 trials of the main task (27 minutes), a resting state scan (10 minutes), an DWI scan (18 minutes) inside the scanner and two short post-experimental tasks in the behavioural lab, to control for choice bias (20 minutes). All sessions were performed at The Cardiff University's Brain Imaging Centre (*CUBRIC*). Sessions 2 and 3 were carried out in the 7T Siemens Magnetom MRI scanner. Visual stimuli were presented using *PsychoPy* (Peirce 2007).

4.3.7 fMRI data acquisition

Measurements were performed on a whole-body 7 Tesla research MR-system (Magnetom, Siemens Healthcare GmbH, Erlangen, Germany) with 32-channel head receive/volume transmit (Nova Medical, Wilmington MA). BOLD sensitive T2* weighted EPI images in sequential descending order in a rapid event-related design (TR=1800 ms, TE=26 ms, FA=70°, 256 × 256 × 192 isotropic 1 mm voxels, 60×2 mm slices with slice separation 2 mm). 900 volumes were acquired each session and the first six of which were discarded to allow for steady-state magnetization.

4.3.8 fMRI data preprocessing

MRI data was processed using SPM12 (www.fil.ion.ucl.ac.uk/spm). fMRI data were converted from DICOM to NIFTII format, unwarped, spatially realigned to the first image, and corrected for acquisition delay by sinc interpolation with references to the middle slice. The mean fMRI volume and T1-weighted structural image were coregistered using mutual information, and the T1-weighted image was segmented and normalized to the Montreal Neurological Institute (MNI) T1 template by linear and non-linear deformations. The normalization parameters were applied to all spatio-temporally realigned functional images obtaining normalized volumes with a voxel size of 1×1×1 mm. Normalized fMRI data were smoothed with an isotropic Gaussian kernel with full-width half-maximum of 4 mm.

4.3.9 fMRI data analysis

4.3.9.1 Whole Brain Analysis.

First-level and second-level whole brain general linear models (GLM) were build using SPM12. We labeled choice trials according to alternative availability (3 levels) and discriminability (2 levels). A single parametrically modulated regressor was added for freedom rating trials. Additionally, we included motion parameters as nuisance regressors. Each session was modelled separately. A high-pass filter (128-s cutoff) was applied to remove low-frequency drifts. Contrasts were applied to test effects of interest: *free – forced* (choices with more than one alternative vs instructed), *F3–F2* (3 available alternatives – 2) and *HD–LD* (comparing whole brain effects of

Table 4.1: Peak MNI coordinates of ROI associated with free choice in perceptual tasks from Si, Rowe and Zhang, 2020.

ROI	x	y	z
preSMA	0	18	48
Left DLPFC	-44	32	30
Right DLPFC	44	34	30
Left IPL	-44	-50	50
Right IPL	54	-38	48
Left AIC	-34	14	2

discriminability).

To accommodate for random effects across participants, first level contrasts were entered into a second level analysis and compared using one-sample t-tests, allowing for population-level inference. We assessed group-level significance by applying one-sample t-tests against 0 to the first-level contrast images. We report clusters significant at $p<0.05$, FWE-corrected for multiple comparisons, with a cluster-defining threshold of $p<0.001$, uncorrected.

4.3.9.2 Regions of Interest.

Region of Interest (ROI) masks for preSMA, DLPFC, inferior parietal lobule (IPL) and left insula (AIC) were created by defining 10 mm spheres centered at peak coordinate activations based on a recent meta-analysis study of equal-choice paradigm (Si, Rowe, and Zhang 2020; Table 4.1).

4.3.10 Behavioural Modelling

We report the results from Bayesian Mixed-Effects models using standard priors from the *brms* R package (Burkner 2017). Random effects structure includes intercepts and slopes for all regressors of interest (see exact model specification below). We report posterior distribution medians (med), 95% posterior credible intervals (CI), probability of directionality (pd; % of posterior density larger or smaller than 0). Analogous to frequentist analysis, we define significance as pd > 0.95. Reaction times and ratings were modelled using linear regression, while accuracies using logistic regression with Gaussian linking function.

4.3.10.1 Model specification.

All syntax represented using *lme* syntax format (Bates et al. 2014), where stars denote interaction terms and random effect terms are contain within the parentheses.

1. Reaction times model: RT ~ option availability * discriminability + (1+configuration|participant) + (1|session)
2. Accuracy model: accuracy ~ alternative availability * discriminability + (1+configuration|participant) + (1|session), family = Bernoulli(link='logit').
3. Freedom Rating model ratings ~ alternative availability * discriminability + (1+configuration|participant) + (1|session),

Note on random effects. Two sourced of random effects were modelled:

experimental session number and participant. In order to account for different preferences regarding either spatial or finger-responding biases across participants we added configuration slopes to the participant random effects.

4.3.11 Cognitive Modelling

Similarly to precious work (Rowe, Hughes, and Nimmo-Smith 2010; Tomassini et al. 2019; Zhang, Hughes, and Rowe 2012) we assume the volitional choice process can be accounted by a linear integration to threshold. In particular, we use the Linear Ballistic Accumulator (LBA) model (Brown and Heathcote 2008). LBA assumes evidence for each choice option is sampled independently in time and the choice is made when the integrated evidence for one of the accumulators reaches a predefined decision threshold. We assume each available choice option is associated with an accumulator, which can vary per condition and per participant. Our model excludes the instructed (F1) trials and artificially fixes the inactive accumulators (*i.e.,* representing unavailable options) to 0, due to very low error rate, making the model parameters exceptionally difficult to estimate in such cases (Wiecki, Sofer, and Frank 2013).

We allow the accumulation rate and non-decision time parameters to vary between conditions, keeping other parameters fixed. We assume that the threshold cannot vary between conditions, following the assumption that it has to be set before observing the stimuli, disallowing for strategical adjustments during the duration of the trial (Ratcliff and Smith 2004). Due to the reviewer comments asking to test this possibility, additional analysis comparing a model with varying thresholds per conditions is added in the

Appendix.

We use hierarchical Bayesian fitting procedure with Markov Chain Monte Carlo (MCMC) parameter estimation routine to estimate the posterior distributions of the model parameters. Our main aim in fitting the model was to extract subject-level estimates of accumulation rate and correlate them with brain activity within the preSMA. Due to its partial pooling property, Hierarchical estimation aids in attaining more generalizable subject-level predictions (McElreath 2016).

We use the Stan language (Carpenter et al. 2017) for hierarchical implementation of the LBA model. We generate four independent chains of 4,000 samples from the joint posterior distribution of the model parameters using Hamiltonian Monte Carlo (HMC), an efficient method suitable for exploring high-dimensional joint probability distributions (Betancourt 2017). The initial 2,000 samples were discarded as burn-in. To assess the convergence of the Markov chains, we calculate Gelman-Rubin convergence diagnostic $\widehat{R}$ of each model (Gelman and Rubin 1992) and use $\widehat{R} < 1.1$ as a stringent criterion of convergence (Annis, Miller, and Palmeri 2017). In order to extract a single accumulation rate per subject in a given condition type, we average individual accumulation rates across active accumulators.

4.4 Results

4.4.1 Behavioural Results Analysis

behavioural analysis includes all 3 sessions (initial behavioural and two MRI-based). Session effects are modelled as random effects (see *Model Specification*).

4.4.1.1 Reaction Times.

Effect of alternative availability on RT indicated responses in the F3 condition $\text{med}_{\texttt{F3 RT}} = 740$ ms, 95% CI = [692 ms, 790 ms] to be significantly longer than in F2 $\text{med}_{\texttt{F2 RT}} = 701$ ms, 95% CI = [652 ms, 750 ms], pd > 0.999, which in turn were significantly longer than F1 $\text{med}_{\texttt{F1 RT}} = 659$ ms, 95% CI = [623 ms, 712 ms], pd > 0.999. RTs in HD condition $\text{med}_{\texttt{HD RT}} = 659$ ms, 95% CI = [617 ms, 705 ms] were significantly longer than in LD $\text{med}_{\texttt{LD RT}} =$ 617 ms, 95% CI = [563 ms, 661 ms], pd > 0.999. No interactions reached significance.

4.4.1.2 Accuracies.

Effect of alternative availability on accuracy indicated accuracy in the F3 condition $\text{med}_{\texttt{F3 ACC}} = 96.8\%$, 95% CI = [94.1%, 98.1%] to be significantly higher than in F2 $\text{med}_{\texttt{F2 ACC}} = 95.8\%$, 95% CI = [92.6%, 97.4%], pd $= 0.980$ and F1 $\text{med}_{\texttt{F1 ACC}} = 94.8\%$, 95% CI = [91.3%, 96.9%], pd > 0.999. Difference between F2 and F1 did not reach significance $\text{pd}_{\texttt{F2>F1}} = 0.942$. Accuracy in HD condition $\text{med}_{\texttt{HD ACC}} = 96.6\%$, 95% CI = [94.1%, 97.9%] was significantly higher than in LD $\text{med}_{\texttt{LD ACC}} = 94.8\%$, 95% CI = [91.4%, 96.9%], pd > 0.999. No interactions reached significance.

4.4.1.3 Freedom Ratings.

Effect of alternative availability on freedom ratings indicated perceived freedom in the F3 condition $\text{med}_{\texttt{F3 RATING}} = 2.81$, 95% CI = [2.58, 3.03] to be significantly higher than in F2 $\text{med}_{\texttt{F2 RATING}} = 2.60$, 95% CI = [2.38, 2.82] pd > 0.999, which in turn was higher than in F1 $\text{med}_{\texttt{F1 RATING}} = 2.45$, 95% CI = [2.22, 2.69] pd > 0.999. Freedom ratings did not differ significantly between levels of discriminability pd = 0.76. No interactions reached significance.

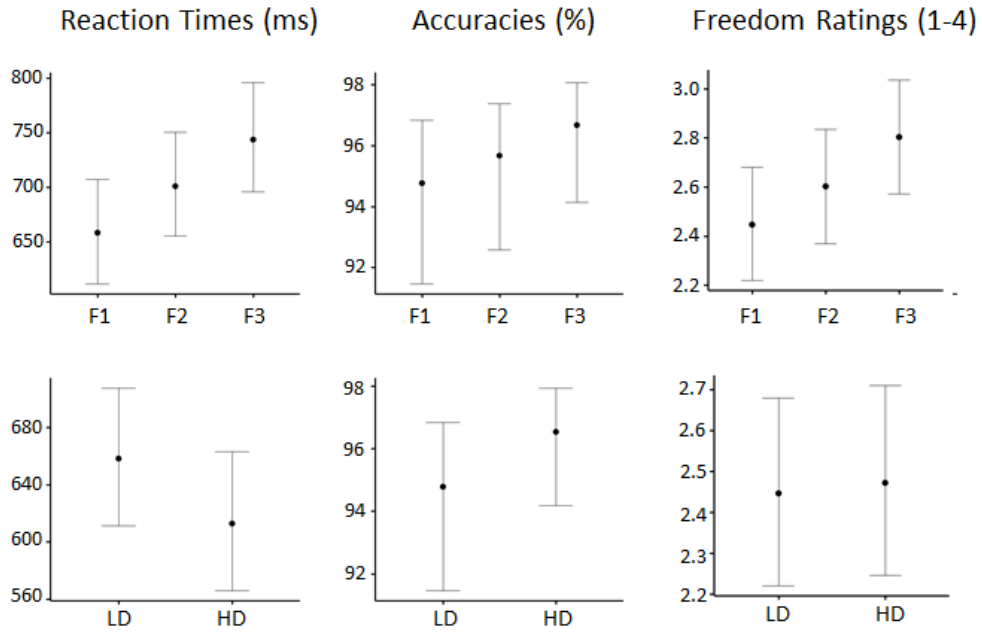

Figure 4.2: Behavioural Results. Plots show group-level parameter estimations centered at the median. Panels represent main effects of choice availability (upper row) and discriminability (lower row) on reaction times (left column), accuracies (middle column) and freedom ratings (right column). Error bars represent 95% credible intervals.

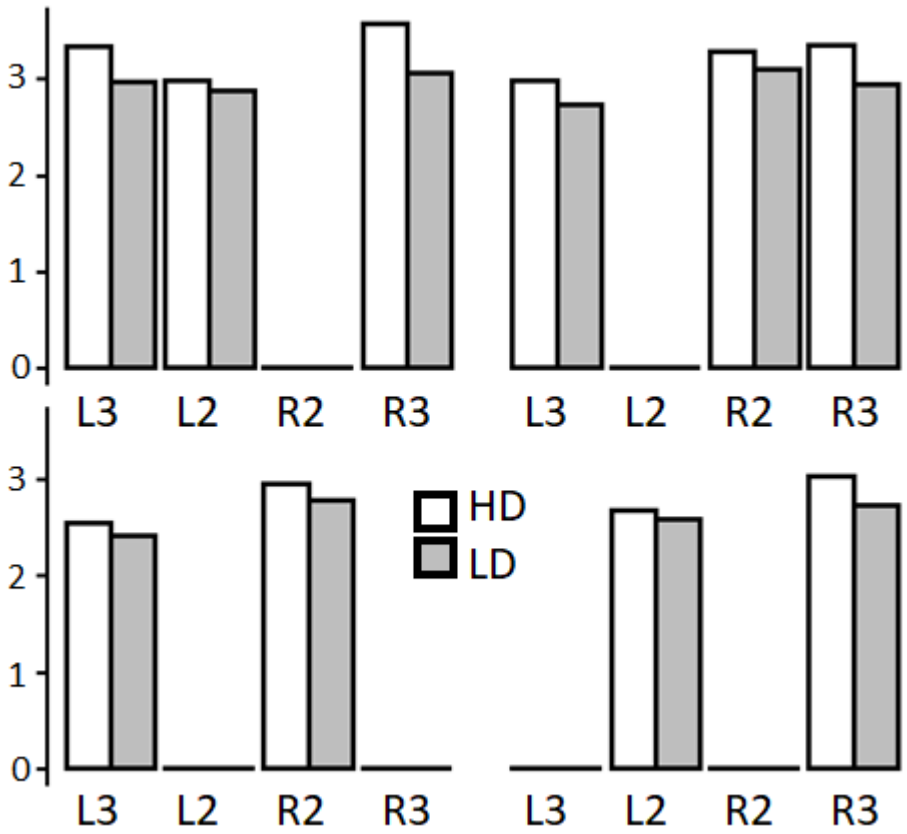

Figure 4.3: Modelling Results. Accumulation rate across conditions. Accumulator labels correspond to fingers used for response: L2 = left index, L3 = left middle, R2 = right index, R3 = right middle. Accumulators set to 0 represent unavailable alternatives (not included in the model).

4.4.2 LBA Generative Modelling

We fitted a hierarchical LBA model to the responses and reaction times with accumulation rates varying by condition type. Each accumulator represented an available alternative in a given condition type. Since the goal of the modelling exercise is to correlate model parameters with BOLD response (specifically, individual accumulation rates with activity in the preSMA region of interest), initial behavioural-only session was excluded from the fitting procedure.

The model provided a good fit to the data ($\widehat{R} < 1.05$ for all parameter values). Figure 4.3 presents group-level accumulation rate values across conditions. Comparison of group-level posterior distributions of accumulation rates across discriminability levels revealed faster accumulation in the HD condition (pd > 0.999). Non-decision times did not differ significantly pd = 0.685 for LD $>$ HD.

4.4.3 Whole-brain analysis

4.4.3.1 Free vs Instructed.

The contrast *free* $>$ *instructed* revealed increased BOLD signal in the SMA/preSMA t(1,29) = 6.46, left pre-central/post-central gyrus t(1,29) = 7.80, right pre-central gyrus t(1,29) = 6.16 and left superior/inferior parietal lobule t(1,29) = 5.88 (Table 4.2). Opposite contrast (*instructed* $>$ *free*) revealed no significant clusters. An interaction contrast revealed no significant clusters, suggesting that the difference between free and instructed

Table 4.2: Clusters of increased activity for planned contrasts (MNI coordinates). Cluster sizes in voxels. Significant at $p < 0.05$, FWE-corrected for multiple comparisons, with a cluster-defining threshold of $p < 0.001$, uncorrected.

Area	x	y	z	Cluster size
Free-Instructed				
SMA/preSMA	-4	-8	58	110
Left pre-central/post-central gyrus	-44	-32	54	1717
Right pre-central gyrus	24	-8	48	248
Left Superior/Inferior Parietal Lobule	-18	-66	62	279
HD-LD				
Left pre-central gyrus	-28	-10	62	236
Left post-central gyrus	-38	-30	38	416
Right inferior occipital cortex	42	-66	0	118
Right Fusiform Gyrus	32	-70	-12	384
Right cerebellar cortex	26	-48	-22	112
F3-F2				
Dorsal Posterior Cingulate Cortex	0	30	38	168
Right/Central Cerebellar Cortex	18	-40	-16	74

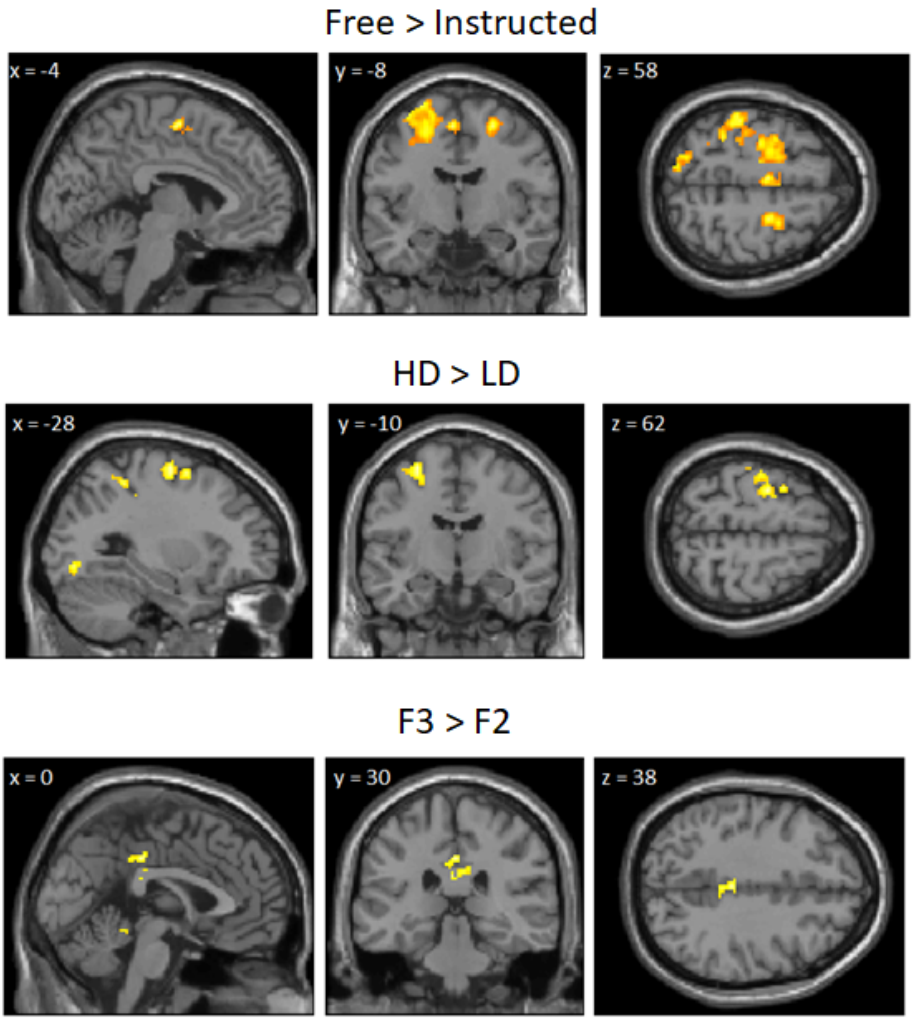

Figure 4.4: Whole-brain analysis results for contrast of interest. Each contrast represented in 3 slices (sagittal, coronal and horizontal), centered at peak coordinates of one of the significant clusters: SMA/preSMA for Free > Instructed; left pre-central gyrus for HD > LD, and posterior cingulate cortex for F3 > F2.

was not modulated by discriminability.

4.4.3.2 HD vs LD.

The contrast *HD > LD* revealed increased BOLD signal in the left pre-central t(1,29) = 5.34, and post-central t(1,29) = 5.10, gyri, right inferior occipital cortex t(1,29) = 4.57, right fusiform gyrus t(1,29) = 6.03, and right cerebellar cortex t(1,29) = 4.84 (Table 4.2). The opposite contrast, which was hypothesized to associated with areas involved with processing difficulty or conflict, did not show any significant activations. An interaction contrast revealed no significant clusters, suggesting that the difference between HD and LD was not modulated by alternative availability.

4.4.3.3 F3 vs F2.

The contrast *F3 > F2* revealed increased BOLD signal in the dorsal posterior cingulate cortex t(1,29) = 4.94, and the right/central cerebellar cortex t(1,29) = 5.67 (Table 4.2). Opposite contrast (*F2 > F3*) revealed no significant clusters. An interaction contrast revealed no significant clusters, suggesting that the difference between F3 and F2 was not modulated by discriminability.

4.4.3.4 Rating Freedom.

Due to a methodological error (see Discussion) this contrast has been excluded from the analysis.

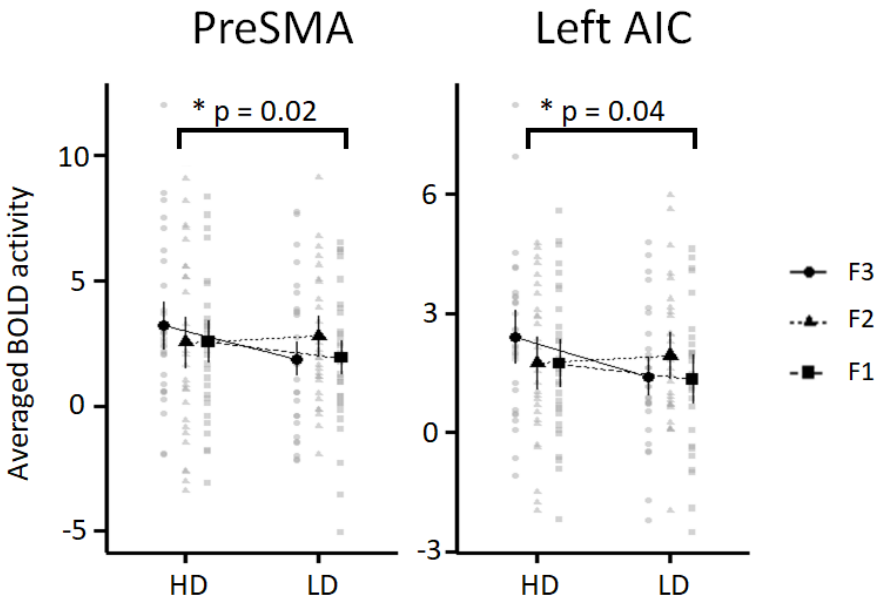

Figure 4.5: ROI ANOVA results from preSMA and left AIC. Only the main effect of discriminability was found to be significant. Solid black shapes represent group means per condition; transparent ones represent individual means. Vertical bars represent standard errors.

4.4.4 ROI Analysis

We run a within-person ANOVA (Availability x Discriminability) for each of 6 predefined clusters of interest. We found a significant effect of discriminability (HD>LD) in preSMA $F(1,29) = 6.16$, $p = 0.02$, $ges = 0.012$, and left insula (AIC) $F(1,29) = 4.72$, $p = 0.04$, $ges = 0.08$ (Figure 4.5). Discriminability also approached significance in the left Inferior Parietal Lobule (IPL), $F(1,29) = 3.68$, $p = 0.06$, $ges = 0.007$. No other effects or interactions reached significance. We also tested whether activity within the preSMA ROI can be predicted by individual accumulation rates, estimated using the LBA model. In order to obtain a single accumulation rate per condition which could be then regressed onto preSMA activity, we averaged the accumulation rates per participant. The regression analysis using availability, discriminability and individual accumulation rate estimates on preSMA ROI actvity revealed no effect of accumulation speed $\beta = -0.65$, $p = 0.72$.

4.5 Discussion

This study aimed to describe the effects of voluntary choice between equally salient perceptual stimuli, controlling for the number of alternatives available to choose, and assessing the subjective feeling of freedom. We replicate some of the findings from the literature, while our novel predictions were not supported by the data. This section discusses our findings in detail.

Our behavioural findings replicate well known effects. Both accuracies and reaction times were modulated by the number of available alternatives (Bar-

las, Hockley, and Obhi 2017) and discriminability (Teodorescu, Moran, and Usher 2016). Very high accuracies across conditions indicate the perpetual task was easy in both discriminability conditions. Subjective freedom ratings were predicted by the number of available alternatives, indicating that manipulating the number of options, even though they were represented by identical stimuli and led to identical outcomes, did significantly affect the experienced freedom of choice. This finding supports the link between objective and subjective definitions of choice freedom (Filevich et al. 2013; Frith 2013; Nahmias et al. 2004).

Subjective freedom was not modulated by discriminability, which is in opposition to studies showing that lower difficulty (Sidarus and Haggard 2016) or higher fluency (Chambon, Sidarus, and Haggard 2014) leads to a greater feeling of agency. The absence of this effect might be associated with the ceiling effect due to low overall task difficulty. Alternatively, it could relate to differences between how people perceive freedom and agency. While related, these concepts might convey slightly different meanings to participants. A potential explanation might be that agency puts an emphasis the experiential, low-level feeling of making the choice, while freedom assessment relates more into the retrospective, metacognitive inference about previous choice (Synofzik, Vosgerau, and Newen 2008). Further work is necessary to distinguish how humans differentiate between the two terms.

The *free vs instructed* contrast reveals some of typical activations known from the literature, including the SMA/preSMA and parietal lobule (Si, Rowe, and Zhang 2020). We however found no support for hypothesized involvement of preSMA in evidence accumulation, as evidenced by the ROI analyses, nor

sensitivity to the difference between 2 and 3 available alternatives.

The *F3 vs F2* contrast revealed a significant cluster in the medial posterior cingulate cortex (PCC). This finding is inconsistent with previous studies comparing more than 2 levels of choice freedom (Forstmann et al. 2006; Lau et al. 2004; van Eimeren et al. 2006). We believe this effect is associated with the fact that in contrast to the mentioned experiments our design introduced an element of perceptual difficulty.

Activation in the PCC can be explained by its sensitivity to the load of response selection (Badgaiyan and Posner 1998). This interpretation is congruent with another study comparing the effect the number of available alternatives (Woo and Lee 2007). In their paradigm, a cue informed participants whether the trial will contain 1, 2, or 4 possible responses. In contrast to the equal-choice paradigm, Woo and Lee's task had a one-to-one stimulus-response mapping, meaning each trial was associated with only one correct response. Congruently, neurons in PCC have been associated with selectivity for spatial choice in rhesus monkeys (Li et al. 2019).

An alternative account stems from theories hypothesizing PCC role in internal, goal-directed cognition (Leech and Sharp 2014; Pearson et al. 2011; Spreng et al. 2010). Based on this account, the PCC could be related to intentional action planning when more alternatives are available to choose. This interpretation however seems less likely due to lack of experimental support from previously mentioned studies. Our result suggests PCC might be an important, previously overlooked component of action selection between multiple alternatives. To distinguish its function, future work could additionally control for the type of task being performed (perceptual vs

reactional intention).

Contrary to our hypotheses, we did not find any clusters related to the more difficult LD condition. We expected areas involved in cognitive control and perceptual difficulty such as anterior cingulate cortex, DLPFC, preSMA, insula or the SPL (Cole and Schneider 2007; Ho, Brown, and Serences 2009; Keuken et al. 2014) to be more active during the LD condition. In contrast, we found clusters in the precentral and postcentral gyri, as well occipital and cerebellar cortices being more responsive to the HD condition. One explanation of this effect is that the difference in difficulty was too small for an increase in cognitive control to occur. This is partially supported by the behavioural results, showing only slightly reduced accuracy, from 97% to 95%, across discriminability levels. This ceiling effect could contribute to the LD condition not being perceived as challenging, but rather the HD condition as being more salient, linking this the difference in contrast to a magnitude effect (Pais et al. 2013), a ecologically motivated phenomenon where sensitivity to absolute stimulus values leads to faster responses. According to this view, the activations reflect a network involved in processing perceptual magnitude (Skagerlund, Karlsson, and Traaff 2016). Interestingly, the ROI analysis indicated a significant effect of discriminability (HD>LD) in 2 out of 6 areas related to the free choice network (preSMA and AIC), with IPL also trending towards significance. This finding supports the link between choice fluency and the low-level experience of freedom (Chambon, Sidarus, and Haggard 2014; Sidarus and Haggard 2016), suggesting that the encoding of free choice and fluency shares overlapping representations, from which agency is derived. Importantly however, as noted in the behavioural section, this difference was

not observed in the freedom ratings. A regression analysis using availability, discriminability and individual accumulation rate estimates averaged across trials for all accumulators revealed no effect of accumulation speed on preSMA activity, showing no evidence for the region being involved in the evidence accumulation process.

Our study design was associated with a series of limitations. Limiting the number of possible configurations (Figure 4.1) might reduce the level of conflict and cause participants to respond more habitually, as opposed to freely. Secondly, the first full behavioural session could lead to overtraining, and, as a result, an attenuated cortical response, an effect found in perceptual learning literature (Yotsumoto, Watanabe, and Sasaki 2008). Additionally, the optimization of the EPI scanning sequence might also be of crucial importance. The current study was the first functional experiment on the 7T Siemiens Cubric machine, hence no prior efficient solutions were established. One specific issue might relate to the lack of *z-shimming*, a procedure designed to prevent signal dropout due to variations in magnetic susceptibility (Weiskopf et al. 2006). This issue is amplified in high magnetic fields, and is especially prevalent in the anterior parts of the cortex, which was in fact the case in the current dataset. These limitations might have contributed to why we did not find some of the expected activations patterns, especially relating to the more anterior parts of the brain such as DLPFC, ACC and anterior part of the preSMA.

While the differences observed in subjective freedom between conditions were significant, one might argue that the effect sizes were rather small. This might be due to the difficulty in both individual definitions of freedom,

as well as the method of prompting it in the experiment. While the first factor is a welcomed component of between-subject variability, the second is problematic, as different phrasing of the same or very similar concept can potentially result in different responses (*e.g.* Henley et al. 2016). Similar to Filevich et al (2013), we adopted a rather high-level question, asking directly about the level of perceived subjective freedom. Most common alternate phrasing used in the literature refers to the feeling of control (Barlas, Hockley, and Obhi 2017; Chambon, Sidarus, and Haggard 2014; Barlas and Kopp 2018; Linser and Goschke 2007). Since the range of the scale (1 to 4) was the same as the number of targets in each trial, we cannot exclude the possibility that some of the participants misconstrued it as a test question and responded by reporting the perceived number of available targets in the previous trial. This interpretation is however unlikely, since the initial instructions were clear on the scale being related to a subjective feeling, and this fact was additionally stressed to the participants verbally during the task tutorial before session 1. Overall, the small effect sizes observed suggest that more research is needed to discern the influence of option availability and discriminability on subjective feeling of freedom, possibly testing different phrasings and prompts.

While, consistently with the definition stated at the start of the thesis, we associate discriminability with difficulty, it is important to note that both high and low discriminability conditions in this study were in fact very easy by design (see *Methods* section). This is because in most of such studies, the more difficult conditions are associated with a much greater gap in accuracy. For instance, the difficult condition in previously mentioned studies oscillated around 86% (Cole and Schneider 2007) and 66% (Ho,

Brown, and Serences 2009), compared to 95% in the low discriminability condition here. This makes comparing our discriminability effects with studies manipulating difficulty potentially problematic. From this point of view, our discriminability manipulation is more similar to controlling for stimulus saliency (Teodorescu, Moran & Usher, 2016) than difficulty.

A design flaw disabled us to dissociate the modulatory effect of subjective freedom from motor activity. Since the freedom ratings were mapped on the response pad to go from 1 to 4 from the left-most to the right-most button, the activations correspond to both freedom as well as to motor output, confounding the resulting contrasts. Due to this confound we do not report the findings resulting from this comparison. A solution of this design flow would involve counterbalancing the ratings between subjects, such that for half the ratings would be mapped from left to right, while for the others from right to left.

Based on these experiences, we recommend future researchers to increase choice conflict by including a larger set of possible response configurations, as well as including more task-naïve subjects, to counteract the potential cortical habituation effects associated with overtraining. In case of using a new experimental scanning sequences, we strongly recommend gathering a larger sample of pilot data before committing to the final design. While no set in stone rule can be applied here, an reasonable approach would include a sample size large enough to be used for power calculation in the main experiment (Soares et al. 2016). In cases of a novel design or sequence, a pilot sample size of minimum 12 participants is recommended (Desmond and Glover 2002).

Chapter 5

Breaking Deadlocks: Reward Probability and Spontaneous Preference Shape Voluntary Decisions and Electrophysiological Signals in Humans

Choosing between equally valued options is a common conundrum, for which classical decision theories predicted a prolonged response time (RT). This

contrasts with the notion that an optimal decision maker in a stable environment should make fast and random choices, as the outcomes are indifferent. Here, we characterize the neurocognitive processes underlying such voluntary decisions by integrating cognitive modelling of behavioural responses and EEG recordings in a probabilistic reward task. Human participants performed binary choices between pairs of unambiguous cues associated with identical reward probabilities at different levels. Higher reward probability accelerated RT, and participants chose one cue faster and more frequent over the other at each probability level. The behavioural effects on RT persisted in simple reactions to single cues. By using hierarchical Bayesian parameter estimation for an accumulator model, we showed that the probability and preference effects were independently associated with changes in the speed of evidence accumulation, but not with visual encoding or motor execution latencies. Time-resolved MVPA of EEG evoked responses identified significant representations of reward certainty and preference as early as 120 ms after stimulus onset, with spatial relevance patterns maximal in middle central and parietal electrodes. Furthermore, EEG-informed computational modelling showed that the rate of change between N100 and P300 event-related potentials modulated accumulation rates on a trial-by-trial basis. Our findings suggest that reward probability and spontaneous preference collectively shape voluntary decisions between equal options, providing a mechanism to prevent indecision or random behaviour.

5.1 Note on Contributions

The study described in this chapter is a collaborative effort of different members of our lab. Initial study design and data collection was performed by Jacopo Barone and Jiaxiang Zhang, before the start of my PhD. Due to this, part of the methods section involving experimental design are archival in nature. The sections describing these archival elements have been written by me. behavioural and modelling sections analyzing the results are a product of my own work and ideas, and therefore can be considered as my unique contributions. Additionally, parts of the analysis involving EEG preprocessing, ERP analysis and MVPA analysis have been performed and written by Dominik Krzeminski.

Both the archival elements and analyses performed not by me are included for completeness, and clearly marked at the start of a given section. It is important to stress that due to an integrative nature of many of the analyses (combining EEG, MVPA and modelling results) including all those sections is necessary for understanding the the full extend of the novel insights.

5.2 Introduction

Cognitive flexibility enables decision strategies to adapt to environmental and motivational needs (Schiebener and Brand 2015). One characteristic of this ability is that harder decisions often take longer. Evidence from neurophysiology (Gold and Shadlen 2001), neuroimaging (Heekeren, Marrett, and Ungerleider 2008) and modelling (Ratcliff and Smith 2004) suggest an

evidence accumulation process for decision-making: information is accumulated over time, and a decision is made when the accumulated evidence reached a threshold (Gold and Shadlen 2007). This process can accommodate paradigms consisting of noisy stimuli (perceptual choices), as well as a rich variety of tasks with unambiguous stimuli (value-based: Pisauro et al. 2017; or memory-based choices: Ratcliff 1978). For perceptual choices, evidence is derived from the sensory properties of the stimuli; for value or preference-based choices, it originates from internal value evaluation and comparison (Krajbich et al. 2012); while for memory-dependent choices, from sampling memory traces (Ratcliff 1978; Shadlen and Shohamy 2016). According to this framework, decision difficulty, and in turn response time (RT), is proportional to the relative difference in the evidence supporting each option, consistent with results from perceptual (Ditterich, Mazurek, and Shadlen 2003), value-based (R Polania et al. 2014; Oud et al. 2016) and memory based decisions (Ratcliff and McKoon 2008).

Making difficult choices requires more evidence, and hence longer deliberation can be an advantageous decision strategy. Yet, scaling deliberation with difficulty is beneficial only to a certain point. What happens if decision difficulty reaches a tipping point with values of options being indistinguishable? In the hypothetical paradox of *Buridan's ass* (Inwagen 1989), a donkey which cannot choose between two identical haystacks would, as a result of its indecision, starve to death. This view is consistent with the classical drift-diffusion model (DDM; Ratcliff and McKoon 2008), which encodes the relative difference of evidence in favour of two options as a single accumulation process between two absorbing boundaries. Such a model would predict

a deadlock or indecision between two equal alternatives, because there is zero difference in the mean evidence supporting each choice (*e.g.* two identical haystacks), and the decision process is dominated by noise accumulated over time, resulting in prolonged RT (Teodorescu, Moran, and Usher 2016; but see Ratcliff, Voskuilen, and Teodorescu 2018 for a recent model modification that addresses this theoretical limitation).

On the other hand, economic analysis suggests that choices between equal alternatives should be made as fast as possible. The benefit of *rushing to decisions* comes from being able to relocate our cognitive resources elsewhere (Rustichini 2009). If evidence cannot bring us closer to a better choice, deliberative thinking becomes an expensive and unnecessary luxury. This effect can be modelled using stochastic decision models with multiple accumulators, each encoding the accumulated evidence in favour of one choice, such as the Linear Ballistic Accumulator model (Brown and Heathcote 2008) and the Leaky Competing Accumulator model (Usher and McClelland 2001; Bogacz et al. 2007). For those models, multiple accumulators compete against each other on the basis of multiple sources of evidence inputs, which by default eliminates the scenario of indecision between equal alternatives.

In reality, individuals can make timely choices between equally valued options. For example, in preference-based decisions, it took under 2 seconds for one to choose between two snack food stimuli that had similar valuations (Voigt et al. 2019). In both humans and non-human primates, higher reward magnitude facilitates RT in perceptual and value-based decisions between equal choices (Pirrone, Azab, Hayden, Stafford, and Marshall 2018a). Intuitively, Buridan's donkey would be motivated to make faster decisions if the haystacks are

fresh, compared to when they are stale. This magnitude effect is in line with ecological incentives: high rewards may imply a resource-rich environment, for which one needs to exploit as early as possible; low rewards may imply a resource-poor environment in which it is worth waiting for a better option (Pirrone, Azab, Hayden, Stafford, and Marshall 2018a). Furthermore, if choices are based purely on expected rewards, one may choose any of the equal-valued options with the same frequency, leading to random behavior. Nevertheless, previous studies (Zhang and Rowe 2015; Phillips et al. 2018) showed that in a sequence of voluntary action decisions, humans deviated from a random pattern of choice and exhibited low choice entropy across trials. A similar conclusion has been reached in consumer decisions, where brand loyalties are driven by seemingly irrational preferences (Wheeler 1974). These findings suggest a possible preference bias between equal options, which renders some options more likely to be chosen than others.

We focus on three issues that have been unresolved in previous research on choices between equal alternatives. First, we aim to explore the effect of reward probability on RT. We expect that, similar to magnitude (Teodorescu, Moran, and Usher 2016; Pirrone, Azab, Hayden, Stafford, and Marshall 2018a), higher reward probability accelerates RTs. This prediction is not trivial, since probability and magnitude can have different effects on behavior. For example, (Young et al. 2014) showed that magnitude discounting follows a power law, while probability is discounted hyperbolically. Unlike magnitude, probability has a upper bound at 100%, which acts in a qualitatively distinct way on behavior (Tversky and Kahneman 1989). We expect this increase in speed to be non-linear, with choices between two certain (100% probability)

options being disproportionately faster compared to choices between two uncertain ones.

Second, in the evidence accumulation framework, both the rate of the accumulation and the non-decision time can influence a model's prediction of reaction time, the former encoding the strength of evidence and the latter reflecting the latencies of visual encoding and motor execution. During perceptual learning, the accumulation rate increases along with behavioural improvements (Jia et al. 2018), while the non-decision time remains unchanged in the late stage of training (Zhang and Rowe 2014). Furthermore, the accumulation rate is associated with the individual differences in working memory (Schmiedek et al. 2007) and attention (Nunez, Vandekerckhove, and Srinivasan 2017), while the non-decision time is faster in individuals with higher diffusion MRI derived neurite density in the corticospinal tract, the primary motor output pathway (Karahan et al. 2019). Recent research showed that both parameters can be influenced by reward magnitude (Wagner et al. 2020), and the current study will examine further whether reward probability and preference influence the two model parameters.

Third, we aim to describe the macroscopic pattern of brain activities associated with differences in behaviour: it is temporal evolution and relation to model-derived parameters. Functional imaging studies have localized the mesocorticolimbic dopaminergic network to be involved in both reward certainty and preference processing (Tobler et al. 2007; Abler et al. 2009), but little is known about how these relate to global activations across the scalp. Pinpointing when EEG activity diverges between different levels of reward probability and assessing whether these differences are transient or

sustained can further inform our computational model, giving deeper insight into the cognitive underpinnings of the decision process.

Here, we address these questions by combining advanced computational modelling and EEG in a probabilistic reward task. Participants memorized six unambiguous cues associated with three levels of reward probability, a certain reward level (*i.e.*, 100%) and two levels of uncertain reward probabilities (80% and 20%). Participants made two-alternative forced choices between cues with equal reward probability (Figure 5.1). The inclusion of the 100% reward probability condition allowed us to investigate whether cues with definitive rewards are processed in a different manner than the uncertain cues (Esber and Haselgrove 2011). Additional task conditions involved binary decisions between cues with different reward probability (unequal trials) and unitary responses to single cues (single-option trials). This design enabled us to focus on the neurocognitive processes underlying choices between equal options, while participants maintained a clear understanding of cue values for rational decisions between unequal options.

We first examine how reward probability influences behavior and whether a preference bias between equal options is present. We then fit an accumulator model of decision-making (Brown and Heathcote 2008) to the behavioural performance across reward probability levels. Posterior group parameters from hierarchical Bayesian model fitting procedure were used to infer whether the behavioural effects were driven by evidence accumulation or non-decisional components of the process. EEG data were analyzed with time-resolved multivariate pattern classification for decoding spatiotemporal representations of reward probability and preference. To establish a link between the decision

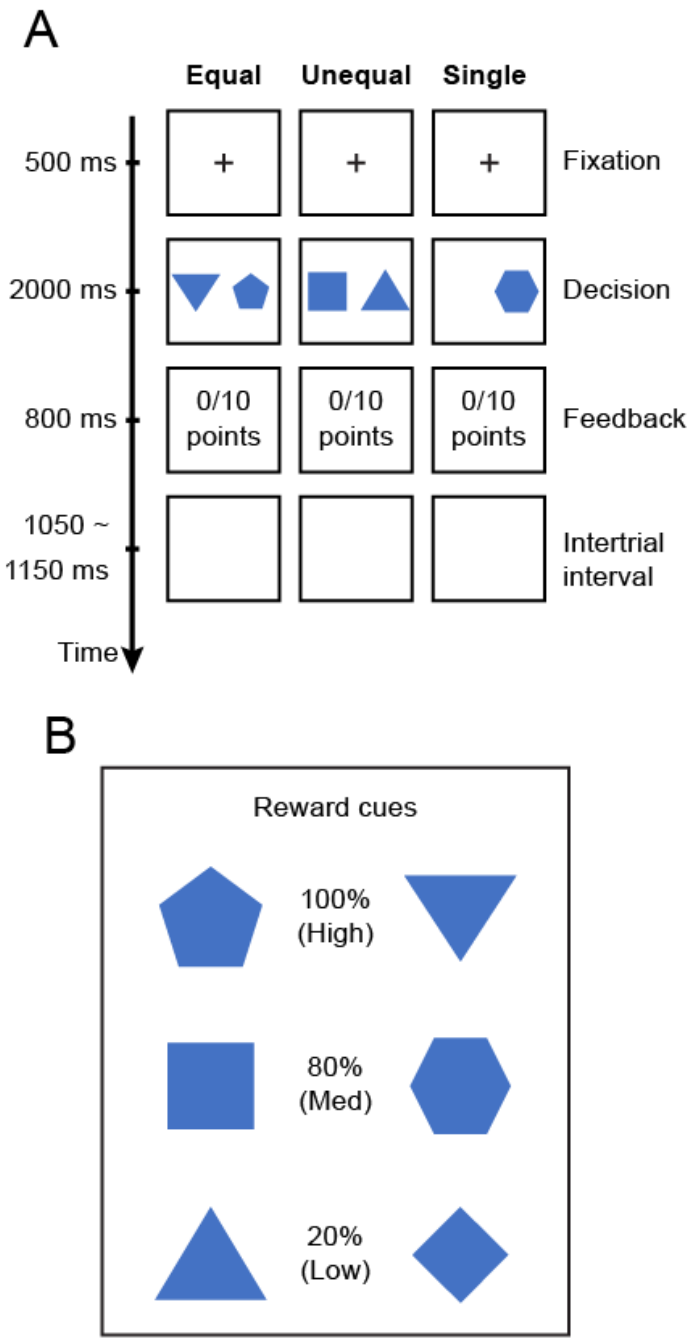

Figure 5.1: Experimental paradigm of the probabilistic reward task. A. Participants were instructed to decide between two reward cues (equal and unequal trials) or respond to a single cue (single-option trials). B. A total of six reward cues were randomly assigned to three levels of reward probability.

process and its EEG signatures, we integrated behavioural and EEG data into a joint hierarchical Bayesian model and tested the hypothesis that electrophysiological activity reflects trial-by-trial changes in the speed of evidence accumulation for decisions (D. M. Twomey et al. 2015).

We demonstrate that reward probability and spontaneous preference independently shape RTs and choices when deciding between equal alternatives. These behavioural effects affect the decision process and evoke a distinct electrophysiological pattern. Together, our findings contribute to the understanding of how decision deadlocks between two equally probable rewards can be overcome.

5.3 Materials and Methods

5.3.1 Participants

Twenty-three healthy participants were recruited from Cardiff University School of Psychology participant panel (20 females; age range 19-32, mean age 22.7 years; 22 right-handed). All participants had normal or corrected-to-normal vision, and none reported a history of neurological or psychiatric illness. Written consent was obtained from all participants. The study was approved by the Cardiff University School of Psychology Research Ethics Committee.

5.3.2 Apparatus

The experiment was conducted in a dedicated EEG testing room. A computer was used to control visual stimulus delivery and record behavioural responses. Visual stimuli were presented on a 24-inch LED monitor (ASUS VG248) with a resolution of 1920 by 1080 pixels and a refresh rate of 60 Hz, located approximately 100 cm in front of participants. Participants' responses were collected from a response box (NATA technologies). The experiment was written in Matlab (Mathworks; RRID: SCR_001622) and used the Psychophysics Toolbox Version 3 extensions (Kleiner, Brainard, and Pelli 2007).

5.3.3 Experimental design

All participants performed a decision-making task with probabilistic rewards during EEG recording (Fig. 5.1A). Before the task, participants memorized 6 unambiguous cues represented by different symbols and their associated probabilities of receiving a reward (Figure 5.1B; see *Procedure*). All the cues had the same color (RGB = 246, 242, 92) on a black background (100% contrast). Each cue was mapped onto one of the three reward probability levels: *high* (a reward probability of 100%, *i.e.*, always rewarded), *medium* (a reward probability of 80%) and *low* (a reward probability of 20%), and hence there were two different cues associated with each reward probability.

Participants were instructed to maximize the total accumulated reward in the decision-making task. The task contained three types of trials: equal, unequal and single-option. On an equal trial, two different cues with the

same reward probability appeared on the left and right sides of a central fixation point (*e.g.* 100% vs. 100%, 80% vs. 80% or 20% vs. 20%). On an unequal trial, two cues with different reward probability levels appeared on both sides of the central fixation point (*e.g.* 100% vs. 20%, 100% vs. 80% or 80% vs. 20%). On a single-option trial, one of the six cues appeared on either the left or right side of the fixation point. In equal and unequal trials, participants chose the left or right cue via button presses with the right-hand index and middle fingers. In single-option trials, participants responded to which side the single cue was presented (*i.e.*, left or right). In all trials, the reward was operationalized as 10 virtual "game points" that did not have any tangible value. The probability of receiving the reward in a trial was either 100%, 80% or 20%, which was determined by the chosen cue. It is worth noting that, in equal trials, participants' decisions did not actually affect the probability of receiving the reward because both options had equal reward probability. In single-option trials, if participants chose the wrong side with no cue presented (0.1% across all single-option trials), no reward was given. Feedback for rewarded (a "10 points" text message on the screen) or not rewarded (blank screen) choices was given after each trial. The total game points awarded were presented at the bottom of the screen throughout the experiment.

5.3.4 Procedure

Each experimental session comprised 640 trials, which were divided into 4 blocks of 160 trials. Participants took short breaks between blocks and after every 40 trials within a block. The mapping between the six reward cues

and three levels of reward probability was randomized across participants. During breaks, the cues-reward mappings were explicitly presented on the screen (Figure 5.1B), and the participants could take as much time as they needed to memorize them. After the first two blocks, all the cues were re-mapped to different reward probabilities. For example, for the pair of two cues that were associated with 100% reward probability in the first and second blocks, one of the two cues would be associated with 80% reward probability in the third and fourth blocks, and the other associated with 20% reward probability. Participants were encouraged to memorize the altered cue-probability associations prior to the third block. This remapping procedure reduced the potential bias associated with specific cues. No explicit memory tests were performed.

Each block contained 64 equal trials (32 for 100% vs. 100%, 16 for 80% vs. 80% and 16 for 20% vs. 20%); 64 unequal trials (32 for 80% vs. 20%, 16 for 100% vs. 80% and 16 for 100% vs. 20%) and 32 single-option trials (16 for 100%, 8 for 80% and 8 for 20%) at a randomized order. This design ensured the same number of trials with and without cues with the highest reward probability (100%). Note however that individual cues did not differ much in terms of frequency of occurrence: each 100% cue appeared 56 times, compared to 48 for each non-certain cue. This makes it unlikely that observed differences can be explained by occurrence frequency alone. Because two cues were bound to every probability level, different cue positions and combinations can result in the same reward probability pair (*e.g.* there are 4 possible combinations for 80% vs. 20% unequal trials). These combinations were counterbalanced across trials.

Each trial began with the presentation of a fixation point at the center of the screen for 500 ms. After the fixation period, in the equal and unequal trials, two reward cues appeared on the left and right sides of the screen with a horizontal distance of 4.34° from the fixation point. Both cues were vertically centered. In single-option trials, only one reward cue appeared on one side of the screen, and the side of cue appearance was randomized and counterbalanced across trials. Cues were presented for a maximum of 2000 ms, during which participants were instructed to make a left or right button press. The cues disappeared as soon as a response was made, or the maximum duration was reached. The reaction time (RT) on each trial was measured from the cue onset to button press. Reward feedback was given 200 ms. after the reward cue offset and lasted 800 ms, followed by a random intertrial interval uniformly distributed between 1050 and 1150 ms. As in our previous study (Zhang and Rowe 2014), if the participant failed to respond within 2000 ms or responded within 100 ms, no reward was given and a warning message "Too slow" or "Too fast" was presented for 1500 ms.

5.3.5 Behavioural Analysis

We excluded trials with RT faster than 200 ms (fast guesses). For each participant, trials with RTs longer than 2.5 standard deviations from the mean RT were also excluded from subsequent analysis. The discarded trials accounted for 1.5% of all trials.

We first analyzed the proportion of choices in equal trials to establish the existence of a preference bias. In the equal condition, by definition, there was no "correct" or "incorrect" response, since the cues had the same reward

probability. For each pair of cues with the same reward probability, we defined the preferred cue as the one chosen more frequently than the other (non-preferred) in equal trials. The categorization of preferred and non-preferred cues was estimated separately between the first two and the last two blocks, because of the cue-probability remapping after the first two blocks. At each level of reward probability, a preference bias was then quantified as the proportion of trials where the preferred cue was chosen. The preference bias had a lower bound of 50%, at which both cues were chosen with equal frequency.

In the unequal condition, we defined decision accuracy as the proportion of choosing the cue with higher reward probability, separately for each combination of reward probabilities (100% vs. 80%, 100% vs. 20% and 80% vs. 20%). Two-tailed one-sample t-tests compared the decision accuracy in the unequal condition against a chance level of 50%, which would indicate irrational decisions (*i.e.*, both high and low reward cues were chosen in 50% of trials).

To determine how reward probability, preferences and other experimental factors influence RT, we analyzed single-trial RT data with linear mixed-effects models (LMMs) using the *lme4* package (Bates et al. 2015) in *R* (RRID: SCR_001905). The LMM is a hierarchical regression method that distinguishes between fixed and random effects (Gueorguieva and Krystal 2004). LMMs take into account all single-trial data without averaging across trials and offer better control of type 1 and type 2 errors than ANOVA (Baayen, Davidson, and Bates 2008). Therefore, statistical inferences from LMMs are robust to experimental designs with unbalanced trials across

conditions (Bagiella, Sloan, and Heitjan 2000), which is an important feature suitable for the current study.

We designed two LMMs with different dependent variables and factors (Table 5.1). Model 1 analyzed the RTs from equal and single-option trials, including choice type (equal or single-option), reward probability (high, medium or low), cue remapping (before and after), preference (whether the chosen cue was preferred) and right-side bias (whether the chosen cue was on the right side of the screen) as factors. Right-side bias was included to control for spatial bias relating to preference for stimuli presented on the right or left side of the screen. For the unequal condition, because each trial had two cues with different levels of reward probability that cannot be directly compared with equal or single-option trials, the RTs were analyzed separately in Model 2. Here we used similar predictors with exception of probability, which was captured by two additional factors: the sum and the absolute difference of the two reward probabilities, as they both have been shown to affect choice behavior (Thaler 1991; Ballard et al. 2017; Teodorescu, Moran, and Usher 2016).

In all the LMMs, fixed effects structures included hypothesis-driven, design-relevant factors and their interactions, and individual participants were included as the source of random variance (*random effect*). We used a standard data-driven approach to identify the random effects structure justified by the experimental design, which resulted in good generalization performance (Barr et al. 2013). This approach starts with the maximal random effects structure (*i.e.* including all random slopes, intercepts and interactions) and systematically simplified it until the LMM reaches convergence. Table 5.1

Table 5.1: The linear mixed-effects models of RT. Model 1 analyzed single-trial RT in equal and single-option trials. Model 2 analyzed single-trial RT in unequal trials. In both models, preference was a predictor indicating whether the preferred cue was selected in each trial. Cue-remapping was a predictor indicating whether each trial was before or after cure-probability remapping in the second half of each session. Right-bias indicated whether the cue on the right size of the screen was chosen in each trial, modeling a possible response bias.

	Model 1	Model 2
Dependent Variables	RT	RT
Main Effects	reward probability preference cue-remapping choice (equal of single-option) right-bias	sum of reward probability difference of reward probability preference cue-remapping right-bias
Interaction Terms	probability * choice probability * preference probability * cue-remapping choice * preference choice * cue-remapping preference * cue-remapping probability * choice * preference probability * choice * cue-remapping probability * cue-remapping * preference	sum of reward probability * preference difference of reward probability * preference sum of reward probability * cue-remapping difference of reward probability * cue-remapping preference * cue-remapping
Random Effects (correlated slopes) and intercepts)	reward probability preference cue-remapping choice right-bias	sum of reward probability difference of reward probability preference cue-remapping right-bias

lists the simplified random effects structures. The correlation structures of each fitted LMM was assessed to avoid overfitting (Matuschek et al. 2017).

5.3.6 A Cognitive Model of Voluntary Decision-Making

We further analyzed the behavioural data using the Linear Ballistic Accumulator (LBA) model (Brown and Heathcote 2008). LBA model is a simplified implementation of a large family of sequential sampling models of decision-making (Ratcliff and Smith 2004; Bogacz et al. 2006; Gold and Shadlen 2007; Zhang 2012) which assumes an independent accumulation process for each choice option. Our model-based analysis has three stages. First, we fit a family of LBA models with various model complexity to the behavioural data of individual participants in equal trials. By identifying the best-fitting model, we infer how reward probability and preference modulated subcomponents of the evidence accumulation process during decision-making. Next, we simulate the best fitted LBA model and examine whether model simulations are consistent with the experimental data in single-option and unequal conditions. This is a stringent test of model generalizability because the experimental data in single-option and unequal trials are unseen by the model fitting procedure. Finally, we link the cognitive processes identified by the LBA model to brain activity by incorporating a trial-by-trial measure of EEG activity regressors into the best-fitted model (Cavanagh et al. 2011; Nunez, Vandekerckhove, and Srinivasan 2017; Nunez et al. 2019).

The LBA model assumes that the decision of when and which option to choose is governed by a "horse race" competition between two accumulators $i \in \{1, 2\}$ that accumulate evidence over time supporting the two choice

options. One accumulator is in favor of the preferred cue and the other of the non-preferred cue. The activations of the accumulators represent the accumulated evidence. At the beginning of each trial, the initial activations of the two accumulators are independently drawn from a uniform distribution between 0 and A. The activation of each accumulator then increases linearly over time, and the speed of accumulation (*i.e.* accumulation rate) varies as a Gaussian random variable with mean v_i and standard deviation S_i across trials. The accumulation process terminates when the activation of any accumulator reaches a response threshold B ($B > A$) and the choice corresponding to the winning accumulator is selected. The model prediction of RT (measured in seconds) is the sum of the duration of the accumulation process and a constant non-decision time T_{er}, with the latter accounts for the latency associated with other processes including stimulus encoding and action execution (Brown and Heathcote 2008; Nunez et al. 2019; Karahan et al. 2019).

5.3.7 Model Parameter Estimation and Model Selection

LBA model has five key parameters: mean v and standard deviation S of the accumulation rate across trials, decision threshold B, starting point variability A and non-decision time T_{er}. To accommodate the empirical data, one or more model parameters need to vary between conditions. We evaluated a total of 21 variants of the LBA model with different parameter constraints (Figure 5.3). First, the accumulation process may differ between the preferred and non-preferred options, leading v or S to vary between accumulators (preferred, non-preferred). Second, reward probability could

modulate the accumulation process or visuomotor latencies unrelated to decisions, leading to v, S or T_{er} to vary between three levels of reward probability. Third, the decision threshold B and starting point A were fixed between conditions, because the trial order was randomized, and we do not expect the participants to systematically vary their decision threshold before knowing the cues to be presented (Ratcliff and Smith 2004). Fourth, decision threshold B and starting point A were fixed across preference levels, since participants could not predict which cue would appear on which side of the screen. During model-fitting, the decision threshold was fixed at 3 as the scaling parameter (Brown and Heathcote 2008), and all the other parameters allowed to vary between participants. Theoretically, the scaling parameter can be set to an arbitrary value, which does not influence the parameter inference, as long as the priors of other parameters remain realistic. Finally, because the participants showed behavioural differences between reward probability levels and between preferred/non-preferred choices, we only estimated models that could bapture these data features, where least one parameter varied between reward probability levels (v, S or T_{er}) and at least one parameter varied between accumulators (v or S).

We use a hierarchical Bayesian model estimation procedure to fit each LBA model variant to individual participant's choices (the proportion of preferred and non-preferred choices) and RT distributions in equal trials. The hierarchical model assumes that model parameters at the individual-participant level are random samples drawn from group-level parameter distributions. Given the observed data, Bayesian model estimation uses Markov chain Monte Carlo (MCMC) methods to simultaneously estimate posterior parameter

distributions at both the group level and the individual-participant level. The hierarchical Bayesian approach has been shown to be more robust in recovering model parameters than conventional maximum likelihood estimation (Jahfari, Ridderinkhof, and Scholte 2013; Zhang et al. 2016).

For group-level parameters (v, S, A and T_{er}), similar to previous studies (Annis, Miller, and Palmeri 2017), we used weakly informed priors for their means $E(.)$ and standard deviations $std(.)$:

$$\begin{aligned} E(v) &\sim N(2.5, 1), & std(v) &\sim \gamma(1, 1), \\ E(S) &\sim N(1, 0.75), & std(S) &\sim \gamma(1, 1), \\ E(A) &\sim N(2.5, 1), & std(A) &\sim \gamma(1, 1), \\ E(T_{er}) &\sim N(0.5, 0.2), & std(A) &\sim \gamma(1, 1), \end{aligned} \tag{5.1}$$

where N represents a positive normal distribution (truncated at 0) with parameterized mean and standard deviation, and γ represents a gamma distribution with parameterized mean and standard deviation.

We used the *hBayesDM* package (Ahn, Haines, and Zhang 2017) in *R* for the hierarchical implementation of the LBA model. For each of the 21 model variants, we generated four independent chains of 7,500 samples from the joint posterior distribution of the model parameters using Hamiltonian Monte Carlo (HMC) sampling in Stan (Carpenter et al. 2017). HMC is an efficient method suitable for exploring high-dimensional joint probability distributions (Betancourt 2017). The initial 2,500 samples were discarded as burn-in. To assess the convergence of the Markov chains, we calculated Gelman-Rubin convergence diagnostic $\hat{R}$ of each model (Gelman and Rubin 1992) and used

$\hat{R} < 1.1$ as a stringent criterion of convergence (Annis, Miller, and Palmeri 2017). We compared the fitted LBA model variants using Bayesian leave-one-out information criterion (LOOIC). LOOIC evaluates the model fit while considering model complexity, with lower values of LOOIC indicating better out-of-sample model prediction performance (Vehtari, Gelman, and Gabry 2017).

5.3.8 EEG Data Acquisition and Processing[1]

EEG data were collected using a 32-channel Biosemi ActiveTwo device (BioSemi, Amsterdam). Due to technical issues, EEG data collection was not successful in two participants, and therefore all EEG data analyses were performed on the remaining 21 participants. EEG electrodes were positioned at standard scalp locations from the International 10-20 system. Vertical and horizontal eye movements were recorded using bipolar electrooculogram (EOG) electrodes above and below the left eye as well as from the outer canthi. Additional electrodes were placed on the mastoid processes. EEG recordings (range DC-419 Hz; sampling rate 2048 Hz) were referenced to linked electrodes located midway between POz and PO3/PO4 respectively and re-referenced off-line to linked mastoids. Additional electrodes were placed on the mastoid processes. EEG (range DC-419 Hz; sampling rate 2048 Hz) was collected with respect to an active electrode (CMS; common mode sense) and a passive electrode (DRL; driven right leg), which were located midway between POz and PO3/PO4 respectively, to form a ground-like feedback loop.

[1]data collected by Jacopo Barone & Jiaxiang Zhang, processing done by Dominik Krzeminski, included here for completeness

EEG data were pre-processed using EEGLab toolbox 13.4.4b (Delorme and Makeig, 2004; RRID: SCR_007292) in Matlab. The raw EEG data were high-pass filtered at 0.1 Hz, low-pass filtered at 100 Hz using Butterworth filters and downsampled to 250 Hz. An additional 50 Hz notch filter was used to remove mains interference. We applied Independent Component Analysis (ICA) to decompose continuous EEG data into 32 spatial components, using runica function from the *EEGLab* toolbox. Independent components reflecting eye movement artifacts were identified by the linear correlation coefficients between the time courses of independent components and vertical and horizontal EOG recordings. Additional noise components were identified by visual inspection of the components' activities and scalp topographies. Artefactual components were discarded, and the remaining components were projected back to the data space.

After artifact rejection using ICA, the EEG data were low-pass filtered at 40 Hz and epoched from -400 ms to 1000 ms, time-locked to the onset of the stimulus (*i.e.* reward cues) in each trial. Every epoch was baseline corrected by subtracting the mean signal from -100 ms to 0 ms relative to the onset of reward cues.

5.3.9 Multivariate Pattern Analysis[2]

We use time-resolved Multi-Voxel Pattern Analysis (MVPA) on pre-processed, stimulus-locked EEG data to assess reward-specific and preference-specific information throughout the time course of a trial. In contrast to univariate ERP analysis, MVPA combines information represented across multiple

[2]analysis performed by Dominik Krzeminski, included for completeness

electrodes, which has been shown to be sensitive in decoding information representation from multi-channel human electrophysiological data (Cichy, Pantazis, and Oliva 2014; Dima et al. 2018).

We conduct three MVPA analysis to identify the latency and spatial distribution of the EEG multivariate information. The first to decode reward probability levels in equal choices (*e.g.* equal trials with two 100% reward cues versus equal trials with two 80% cues). The second to decode preferred versus non-preferred choices in equal trials. The third to decode between equal and single-option choices with the same reward probability (*e.g.* equal trials with two 100% cues versus single-option trials with a 100% cue).

Each analysis is formed as one or multiple binary classification problems, and the data feature for classification included EEG recordings from all 32 electrodes. In each analysis, at each sampled time point (-400 ms to 1000 ms) and for each participant, we train linear support vector machines (SVM) (Garrett et al. 2003) using the 32-channel EEG data and calculate the mean classification accuracy following a stratified ten-fold cross-validation procedure. In all MVPA, we include the EEG data from 400 ms before cue onset as a sanity check, because one would not expect significant classification before the onset of reward cues.

In each cross-validation, 90% of the data issued as a training set, and the remaining 10% as a test set. In some analysis (*e.g.* equal trials with 100% cues versus equal trials with 80% cues), the number of samples belonging to the two classes is unbalanced in the training set. We use a data-driven over-sampling approach to generate synthetic instances for the minor class until the two classes had balanced samples (Zhang and Wang 2011). The

synthetic instances are generated from Gaussian distributions with the same mean and variance as in the original minority class data. Training set data were standardized with z-score normalization to have a standard normal distribution for each feature. The normalization parameters estimated from the training set was then applied separately to the test set to avoid overfitting. To reduce data dimensionality, we perform principal component analysis to the training set data and selected the number of components that explained over 99% of the variance in the training set. The test set data are projected to the same space with reduced dimensions by applying the eigenvectors of the chosen principal components. We then train SVM to distinguish between the two classes (*i.e.* conditions) and evaluate the classification accuracy using the test set data. The procedure is repeated ten times with different training and test sets, and the classification accuracies are averaged from the ten-fold cross-validation. We use the SVM implementation in MATLAB Machine Learning and Statistics Toolbox. The trade-off between errors of the SVM on training data and margin maximization is set to 1.

To estimate the significance of the classification performance, we use two-tailed one-sample t-test to compare classification accuracies across participants against the 50% chance level. To account for the number of statistical tests at multiple time points, we use cluster-based permutation (Maris and Oostenveld 2007) to control the family-wise error rate at the cluster level from 2000 permutations.

5.3.10 Estimation of Single-Trial ERP Components[3]

We estimate two ERP components from single-trial EEG data in equal trials: N100 and P300, which are subsequently used to inform cognitive modelling. The visual N100 is related to visual processing (Mangun and Hillyard 1991) and the P300 is related to evidence accumulation during decision making (Kelly and O'Connell 2013; D. M. Twomey et al. 2015).

To improve the signal-to-noise ratio of single-trial ERP estimates, we use a procedure similar to previous studies (Kayser and Tenke 2003; Parra et al. 2005; Nunez et al. 2019). For each participant, we first performed singular value decomposition (SVD) to the grand averaged ERP data across all trials from the same experimental condition. SVD decomposes the trial-averaged ERP data $A_{k\times p}$ (where k is a number of channels and p is a number of time points) into independent principal components. Each component consists of a time series of that component and a weighing function of all channels, defining the spatial distribution (or spatial filter) of that component. Because the ERP waveform is the most dominant feature of the trial-averaged ERP data, the time course of the first principal component (*i.e.* the one that explains the most variance) represents a cleaned trial-average ERP waveform (Nunez et al. 2019), and its weight vector provides an optimal spatial filter to detect the ERP waveforms across EEG channels. We then applied the spatial filter from the first principal component as a channel weighting function to single-trial EEG data to improve the signal-to-noise ratio.

The single-trial EEG data filtered with the SVD-based weighting function is then used to identify the peak-latency and peak-amplitude of the N100 and

[3]analysed together with Dominik Krzeminski, who performed the component extraction

P300 components. For N100, we search for the peak negative amplitude in a window centered at the group-level N100 latency (112 ms) and started at 60 ms. The lower bound of the search window was determined by the evidence that the visual onset latency is ~60 ms in V1 (Schmolesky et al. 1998). For P300, we search for a peak positive amplitude in a window centered at the group-level P300 latency (324 ms). For both N100 and P300, the search window has a length of 104 ms, similar to a previous study (Nunez et al. 2019).

5.3.11 EEG-Informed Cognitive Modelling

Recent studies showed that the variability of the P300 component closely relates to the rate of evidence accumulation during decision making (D. M. Twomey et al. 2015). We therefore extend the best fitting LBA model with EEG-informed, single-trial regressors, which estimates the effect of trial-by-trial variability in EEG activity on the mean accumulation rate (Hawkins et al. 2015; Nunez, Vandekerckhove, and Srinivasan 2017).

The main regressor of interest is the slope of change between the N100 and P300 components, which is defined as the ratio of the P300-N100 peak-amplitude difference and the P300-N100 peak-latency difference in each equal trial. We also test four additional regressors from individual ERP components: P300 amplitude, P300 latency, N100 amplitude and N100 latency. All the EEG regressors are obtained from the estimations of single-trial ERP components in equal choice trials. To obtain a meaningful intercept, the regressors are mean-centered and rescaled to have a unit standard deviation.

Each EEG regressor is tested in a linear regression model, using the same Bayesian hierarchical model estimation procedure as in the behavioural modelling analyses. For each regression model, we assume that the mean accumulation rates of both accumulators $v_1(t)$ and $v_2(t)$ (*i.e.* the one in favor of the preferred option and the other one in favor of the non-preferred option) are influenced by the EEG regressor of interest on a trial-by-trial basis:

$$v_1(t) = \widetilde{v_1} + \beta \times EEG(t), \quad v_2(t) = \widetilde{v_2} + \beta \times EEG(t), \tag{5.2}$$

where $t = 1, 2, 3, \ldots$ represents the equal choice trials, and $\widetilde{v_1}$ and $\widetilde{v_2}$ are the intercepts. The regression coefficient β represents the effect of EEG regressor on the mean accumulation rates.

The rationale of estimating an EEG regressor to the mean drift rate is twofold. First, this approach allows quantifying the trial-by-trial change over the intercept (*i.e.* the mean drift rate), independent of its trial-by-trial variability (parameter S). Second, one would not expect the sensor level EEG signal has sufficient spatial resolution to distinguish between the two accumulators encoding two options. Therefore, we estimated a single EEG regressor across both accumulators.

5.4 Results

We examined the effects of reward probability and spontaneous preference on behavior and EEG activity during voluntary decisions. In a probabilistic reward task (Figure 5.1B), participants chose between two options with

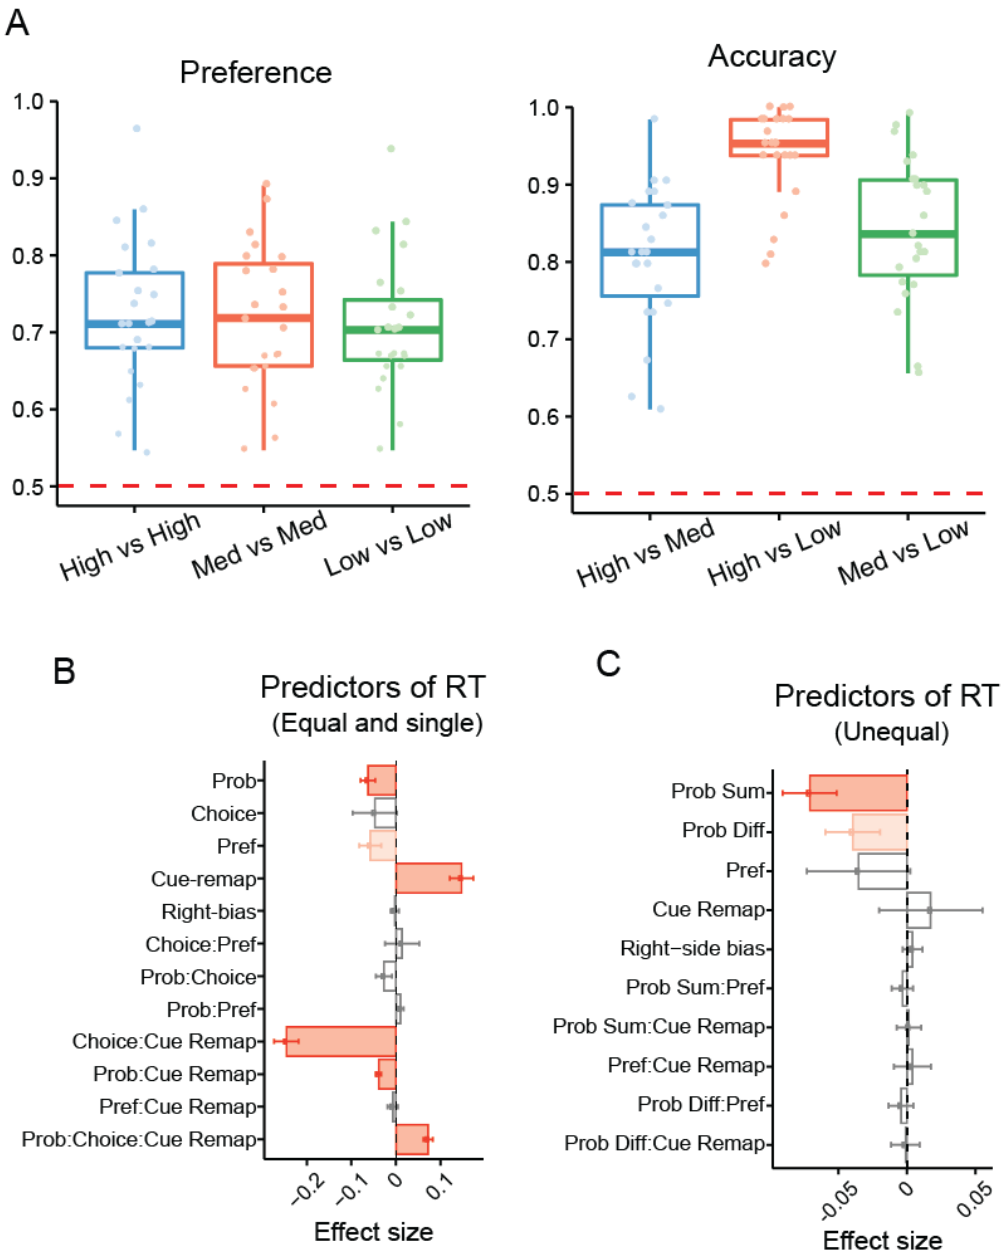

Figure 5.2: behavioural results. A. Preference bias across reward probability levels in equal trials (left) and decision accuracy across reward probability levels in unequal trials (right). B. Linear mixed-effects model results for Model 1 in Table 5.1. Dark red bars represent significant effects with $p < 0.001$. Light red bars represent significant effects with $p < 0.05$. Grey bars represent non-significant factors and interactions. Error bars represent standard errors across participants. C. Linear mixed-effects model results for Model 2 in Table 5.1.

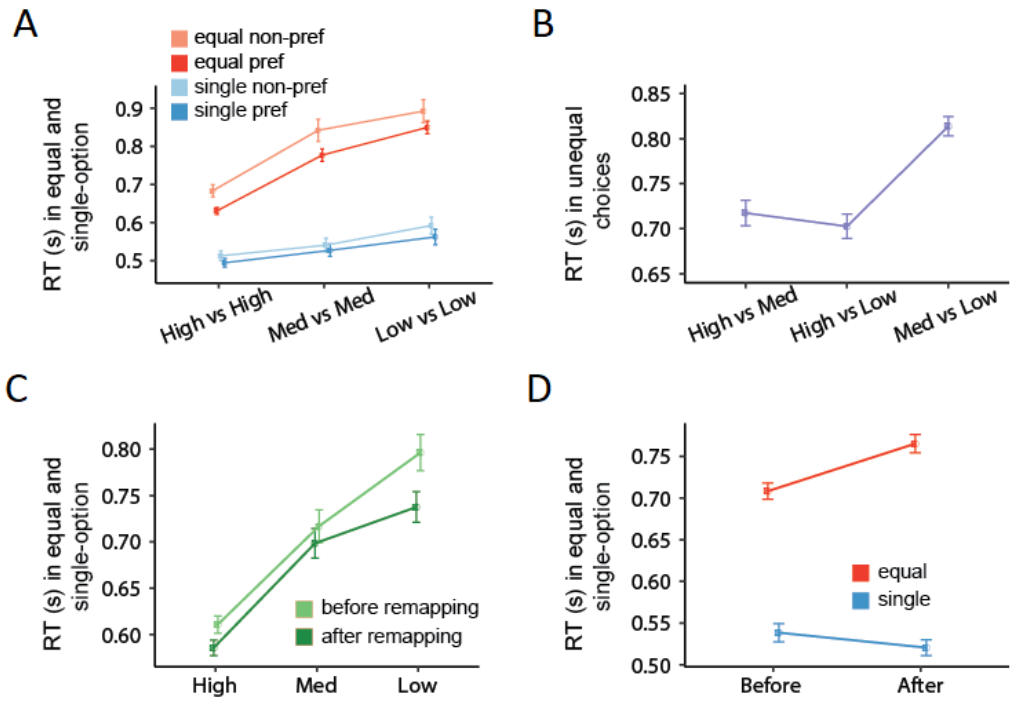

Figure 5.3: Main effects and interactions. Significant effects and interactions in RT from Model 1 (Table 5.1) were presented separately for reward probability and preference in equal and single-option trials (A); before and after cue-remapping at different reward probability levels (C); before and after cue-remapping in equal and single-option trials (D). Significant main effects in Model 2 were presented in panel B. Error bars represent standard errors across participants.

the same reward probability (equal trials) at high (100%), medium (80%) or low (20%) levels. In two control conditions, participants made binary choices between options with different levels of reward probability (unequal trials) or responded to the location of a single reward cue (single-option trials). Below, we first report behavioural results. We then fit linear ballistic accumulator (LBA) models to the choices and RT distributions of equal trials and infer about the underlying cognitive processes based on best-fitting model parameters. Next, we perform univariate and multivariate analyses of EEG data to identify spatio-temporal representations of reward probability and preference information as well as their time courses. We then extend the best-fitted LBA model with single-trial measures of EEG activity to test whether trial-to-trial variations in EEG data relates to the rate of evidence accumulation across trials.

5.4.1 Behavioural Results

Choices

For each pair of cues with the same reward probability, we defined the preferred cue as the one chosen more frequently than the other (non-preferred) in the equal choice trials (see *Behavioural Analysis* in *Methods* section). We found a strong preference bias (>50%) for choosing one reward cue over the other at each level of reward probability (Figure 5.2A; high: 95% CI [0.682, 0.765]; medium: 95% CI [0.679, 0.759]; low: 95% CI [0.669, 0.745]). A repeated-measures ANOVA showed no significant difference in preference between reward probability levels ($F(2, 44) = 0.2, p = 0.81$). Therefore,

although the two options were associated with the same level of reward probability, participants did not make their choices randomly. We further used a linear mixed-effects model (LMM) to evaluate the preference bias as a function of cue remapping (*before* vs *after*) and trial order in each testing block. The preference bias was smaller after cue remapping (Supplementary Figure 5.11, $\beta = -0.181$, 95% CI [-0.01, -0.348], $p = 0.03$), but was not influenced by trial order ($\beta = 0.037$, 95% CI [-0.170, 0.243], $p = 0.73$). These results imply that, for a given set of cue-probability associations, the extent of preference bias did not significantly vary over time. Because the cue-probability mapping was randomized across participants and re-mapped within each session, the observed preference bias is unlikely to be explained by a group-level preference towards any specific cue, but rather a spontaneous preference at the individual level. Additionally, to check if preference from first half of the experiment affected preference after remapping, we calculated the proportion of any cue being preferred in both sessions on a subject level. We found that preference was consistent only in 51.5% of cases, rendering no support for preference transfer after remapping (one-sided binomial-test $p = 0.5$, 95% $CI = [0.361, 1]$).

In unequal choice trials, as expected, the cues with higher reward probability were chosen more often, as evidenced by the above-chance decision accuracies in all conditions (Figure 5.2B; high vs. medium: $t(22) = 16.08$, 95% CI [0.774, 1], $p < 0.001$; high vs. low: $t(22) = 23.31$, 95% CI [0.862, 1], $p < 0.001$; medium vs. low: $t(22) = 20.97$, 95% CI [0.834, 1], $p < 0.001$; one-sample t-test against the 0.5 chance level). A repeated-measures ANOVA showed significant differences in decision accuracy between reward probability levels

($F(2, 44) = 28.17, p < 0.001$). Post-hoc pairwise comparison with Tukey's correction indicated that accuracy in the high vs. low probability condition (93.8%) was significantly higher than in the high vs. medium (84.3%) ($t(44) = 5.267, p < 0.001$) and the medium vs. low (80.7%) ($t(44) = 7.265, p < 0.001$). Similar to the analysis of preference, we used a LMM to evaluate decision accuracy in unequal trials as a function of cue remapping and trial order and found no significant associations (Supplementary Figure 5.11, cue remapping: $\beta = 0.022$, 95% CI [-0.235, 0.1901], $p = 0.84$; trial order: $\beta = 0.078$, 95% CI [0.013, 0.169], $p = 0.1$). These results suggested that participants memorized the cue-probability associations for rational choice behavior and maintained the decision accuracy throughout the experiment.

Response Times

We used a LMM to quantify the influence of experimental factors on RTs in equal and single-option choices (Figure 5.2B, Model 1 in Table 5.1). The fixed effects included reward probability, choice type (equal vs. single-option), preference (choosing the preferred vs. the non-preferred option), cue remapping and their meaningful interactions (Figure 5.3). Participants were faster when choosing the preferred than the non-preferred option ($\beta = -0.063$, 95% CI [-0.027, -0.991], $p < 0.05$) and RTs decreased as the reward probability increased ($\beta = -0.101$, 95% CI [-0.067, -0.135], $p < 0.001$). The RT in equal choice trials were longer than that in single-option trials ($\beta = -0.292$, 95% CI [-0.201, -0.384], $p < 0.001$). The effect of reward probability on RT was stronger in equal compared to single-option choices, supported by a significant interaction between the two main effects ($\beta = 0.045$, 95% CI [0.025, 0.066],

$p < 0.001$).

Participants had slower responses after memorizing a new set of cue-probability associations, indicated by a significant main effect in RT before and after cue remapping ($\beta = 0.149$, 95% CI [0.096, 0.201], $p < 0.001$). The significant interaction between cue remapping and reward probability suggested that the increase in RT was more pronounced in trials with lower reward probability ($\beta = -0.039$, 95% CI [-0.051, -0.026], $p < 0.001$). The interaction between cue remapping and choice type ($\beta = -0.247$, 95% CI [-0.192, -0.302], $p < 0.001$) indicated that this pattern was mainly associated with equal trials. Because evaluating reward probability of a cue was likely associated with additional cognitive load after cue remapping, the observed RT difference before and after cue remapping implies that participants evaluated both cues throughout the experimental session.

In a second LMM, we analyzed RTs in unequal trials (Model 2 in Table 5.1), including the sum and difference of the reward probability of two cues in each trial as fixed effects. The sum of two reward probabilities in unequal trials was negatively associated with RT ($\beta = -0.071$, 95% CI [-0.032, -0.110], $p < 0.001$), consistent with previous studies that the total reward magnitude influences decision-making (Pirrone, Azab, Hayden, Stafford, and Marshall 2018a; Teodorescu, Moran, and Usher 2016). Additionally, the difference of two reward probabilities was also a significant predictor at a more lenient threshold ($\beta = -0.028$, 95% CI [-0.001, -0.055], $p < 0.05$). No other effects or interactions reached significance, further solidifying that the cue-probability associations were well remembered in both halves of the experiment.

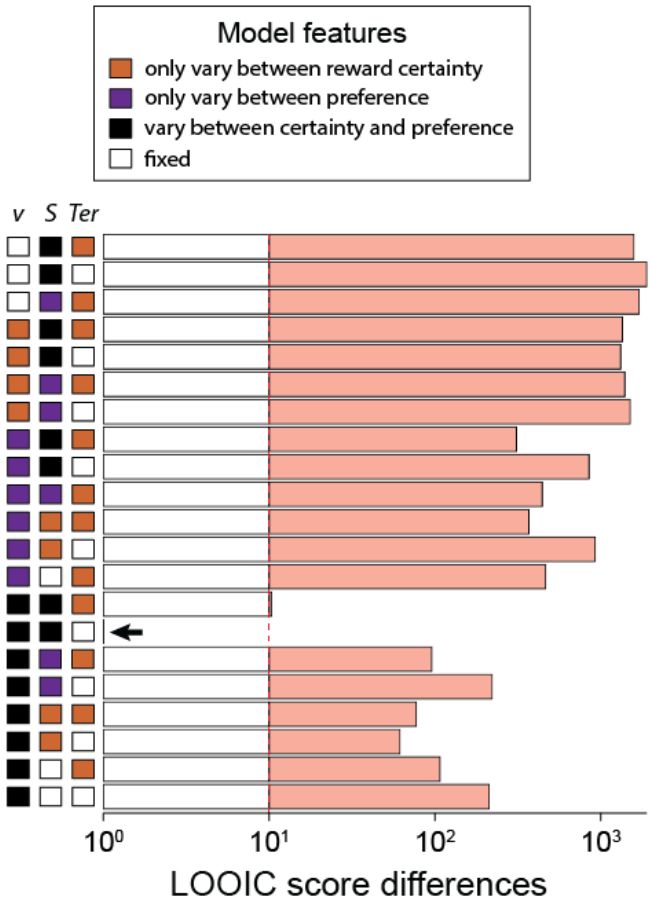

Figure 5.4: Model comparison. LOOIC scores of 21 LBA model variants. The LOOIC score differences between all models and the best model are plotted against corresponding model structures, which were illustrated on the left of the figure. The model structure specified how the mean accumulation rate v, the standard deviation S of the accumulation rate and the non-decision time T_{er} could vary between conditions. A black filled square indicated that the corresponding parameter could vary between reward probability levels and preferred/non-preferred options. An orange or purple filled square indicated that the corresponding parameter could only vary between reward probability levels or preferred/non-preferred options, respectively. Unfilled (white) squares indicated that the parameter remained fixed between conditions. Bar color indicates whether the difference in LOOIC scores is considered substantial (over 10): white part of the bar corresponds to score up to 10, orange to the amount exceeding 10. The best model was shown with a LOOIC score difference of zero (indicated by the red arrow).

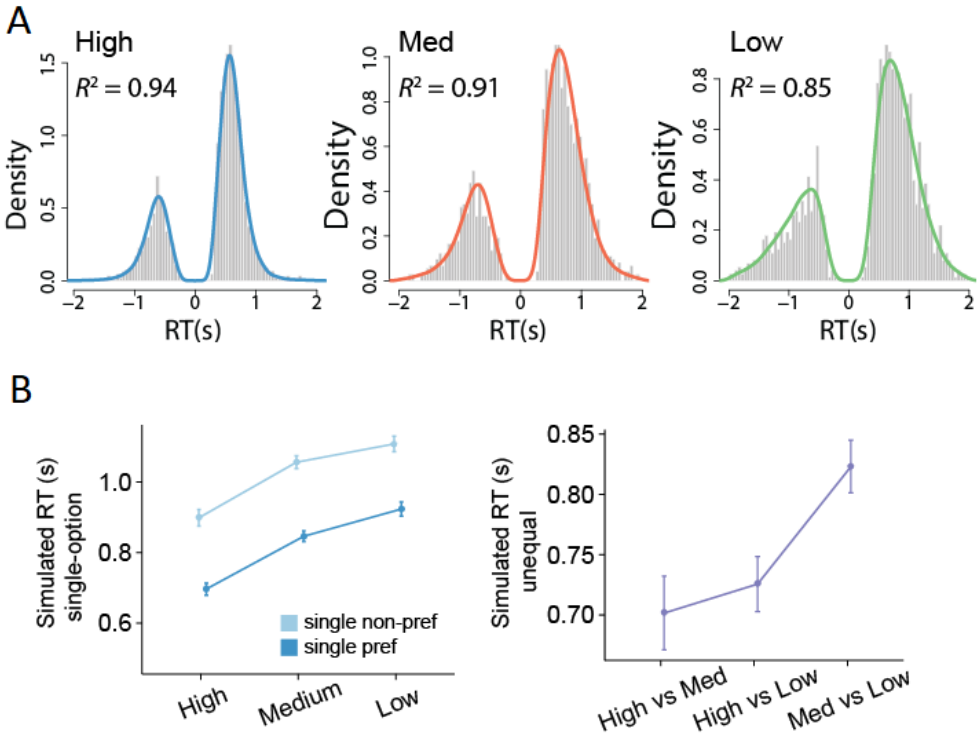

Figure 5.5: Model fit and simulations. A. Simulations of RTs in equal choices, generated from the posterior distribution of the best fitted model for high (left), medium (middle) and low (right) reward probability levels. Histograms represent experimental data and density distributions represent model simulation from 100 iterations. Negative values represent RTs for non-preferred choices. B. Simulation of RTs in single-option (left) and unequal (right) choices from 100 iterations. Error bars represent standard errors across participants.

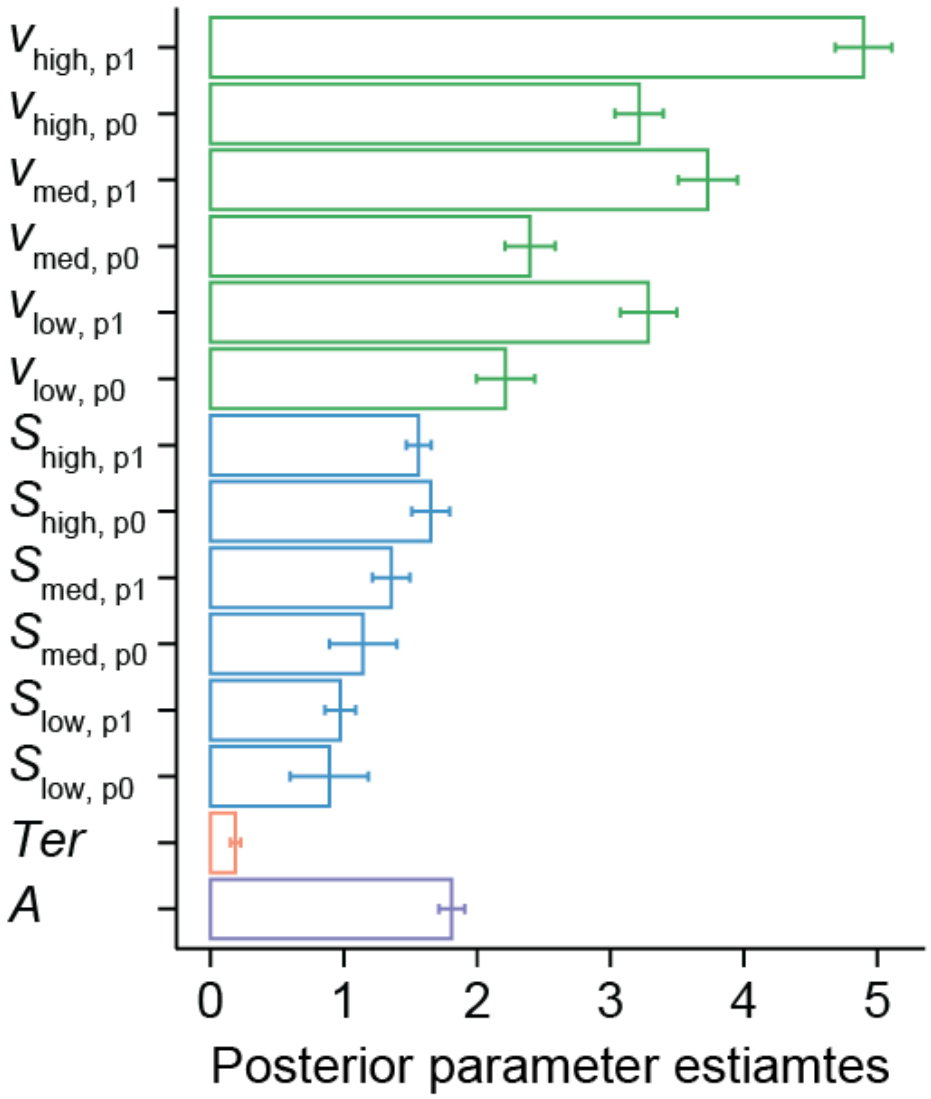

Figure 5.6: Posterior model parameters and inferences. Group-level LBA model parameters of the best fitting model: means of accumulation rates (v, green), standard deviations of accumulation rates (S, blue), non-decision time (T_{er}; orange) and starting point (A, purple). Error bars represent standard deviations of posterior distributions of parameter values. The means and standard deviations of accumulation rates were shown separately for each reward probability level (high, medium and low) and accumulator (p1, preferred option; p0, non-preferred option).

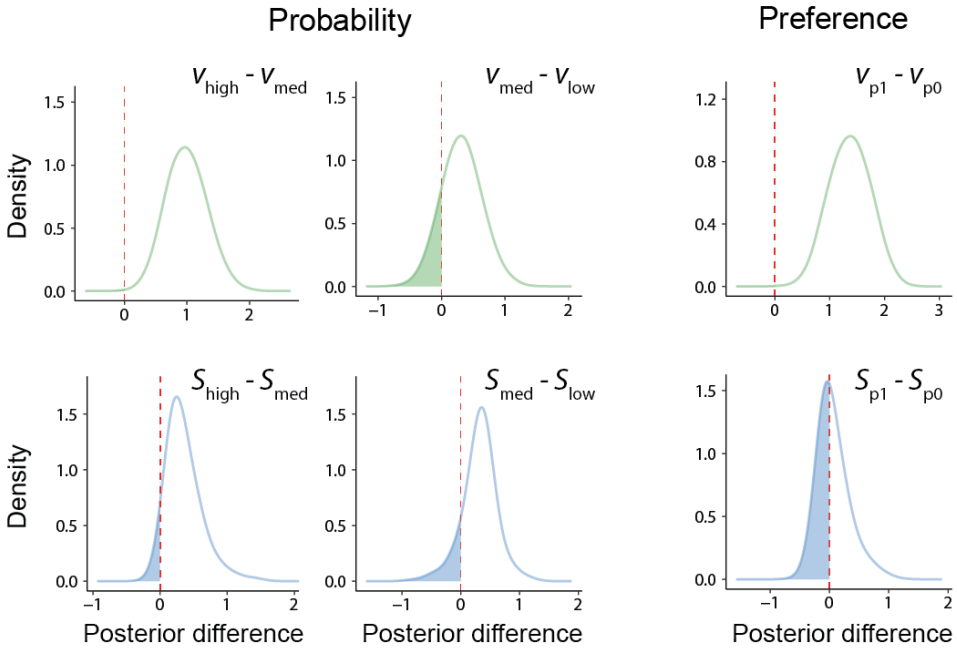

Figure 5.7: Parameter comparison. Differences of posterior parameter estimates across probability levels (left and middle columns) and preference levels (right column). The proportion of posterior difference distributions above zero suggested higher parameter values for higher probability level or more preferred options.

5.4.2 Cognitive Modelling of Behavioural Data

To identify the cognitive processes that led to the observed behavioural differences, we compared 21 variants of the LBA model. The model variants differed systematically in their constraints on whether the rate of evidence accumulation and non-decision time could change between reward probability levels or preferred/non-preferred options. For each model variant, we used hierarchical Bayesian modelling with Markov chain Monte Carlo (MCMC) parameter estimation routine to estimate the posterior distributions of the model parameters, given the observed choice and RT distribution from individual participants (see *Model parameter estimation and model selection*). To identify the model with the best fit, we calculated the Bayesian leave-one-out information criterion (LOOIC) score for each model (Vehtari, Gelman, and Gabry 2017).

MCMC chains representing posterior parameter estimates in all the 21 model variants reached high levels of convergence (Gelman-Rubin convergence diagnostic $\hat{R} \leq 1.02$ for all parameters in all models). The LOOIC scores suggested that the models with the mean accumulation rate varying between reward probability levels and between preference levels fitted the data better than others model variants. The best-fitting model (*i.e.* the one with the lowest LOOIC score; Figure 5.4) had fixed group-level non-decision time with the standard deviation of the accumulation rate varying between reward probability levels and preferred/non-preferred options. To evaluate the model fit to the empirical data in equal trials, we calculate the posterior prediction of the best fitting model by averaging 100 iterations of model simulation using posterior parameter estimates. Averaging across multiple iterations

reduces potential biases when sampling from posterior parameter estimates. Each of the 100 iterations generates simulated behavioural responses (RTs and choices) of individual participants, with the same number of trials per condition as in the actual experiment. There was a good agreement between the observed data and the model simulations across reward probability levels and choice preferences (Figure 5.5A).

We use Bayesian inference to analyze the posterior distributions of group-level model parameters (Bayarri and Berger 2004). To evaluate if a parameter varies substantially between any two conditions, we calculate the proportion of posterior samples in which the parameter value for one condition was greater than the other. To test if a parameter differs from a threshold value, we calculate the proportion of the posteriors greater or smaller than the threshold. To avoid confusion, we use p to refer to classical frequentist p-values, and $P_{p|D}$ to refer to Bayesian inference results based on the proportion of posteriors supporting the testing hypothesis, given the observed data.

For the best fitting model (Figure 5.4), we compared the posterior estimates of the group-level parameters between conditions (Figures 5.6 and 5.7). We found strong evidence for choices with high reward probability to have higher mean (v) and standard deviation (S) of the accumulation rate than choices with medium ($v_{high} > v_{medium} : P_{p|D} = 0.999$; $S_{high} > S_{medium} : P_{p|D} = 0.954$) or low medium ($v_{high} > v_{low} : P_{p|D} = 1$; $S_{high} > S_{low} : P_{p|D} > 0.999$) reward probability. The mean and standard deviation of accumulation rates between choices with medium and low reward probabilities were inconclusive ($v_{medium} > v_{low} : P_{p|D} = 0.839$; $S_{medium} > S_{low} : P_{p|D} = 0.877$). Furthermore, there was also strong evidence for a higher mean

accumulation rate for the preferred than the non-preferred options ($P_{p|D} = 0.999$), and no evidence for a difference in the standard deviation of the accumulation rate ($P_{p|D} = 0.532$). These results supported the claim that preferred and certain (100%) cues were recalled and processed faster than non-preferred cues. Certain cues were also associated with more variable accumulation rate. Model comparisons further suggested that the latencies of early visual encoding and motor execution were not influenced by reward probability nor preference as the models with varying non-decision time parameter did not fit the data as well.

Next, we evaluated whether the best fitting model could reproduce qualitative RT patterns in the single-option and unequal choices, which were unseen by the parameter estimation procedure. This allows us to evaluate whether the model that fits to the equal choice data can also characterize behavioural patterns in other conditions. For unequal choices, two accumulators representing two cues with different reward probability levels compete to reach the decision threshold, with their parameters set to the posterior estimates from the fitted LBA model. For single-option choices, a single accumulator is set to reach to the decision threshold. Similar to the simulation of equal choices, we average the predicted behavioural responses of unequal and single-option choices for each participant from 100 iterations of simulation. Each iteration contains the same number of trials as in the experiment.

For unequal trials, the simulated RT showed similar patterns to the observed data, in which choosing between medium and low probability cues led to the longest RT (Figure 5.5B). For single-option choices, similar to the observed data, higher reward probability and preferred cues were associated with faster

RT in simulation. However, simulated RT in single-option choices was longer than the experimental data, suggesting that simple reactions to a single cue may engage distinct cognitive processes beyond the current model.

5.4.3 EEG Results

We focused our EEG analysis on equal trials (with additional control analysis on EEG data from single-option trials), because both reward probability and preference bias played major roles in shaping the behavioural performance of that condition.

Event-Related Potentials

We examine univariate differences in evoked responses between conditions in single EEG electrodes. For each participant, trial-averaged ERPs are calculated from epochs of equal or single-option choices, with epochs time-locked to reward cue onset. For both equal and single-option conditions, we test for differences in ERPs between three levels of reward probability using a one-way repeated-measures ANOVA. Furthermore, we test for differences in ERPs between preferred and non-preferred choices in equal trials using a paired t-test. We perform statistical tests on all electrodes and all time points. Cluster-based permutation tests (2000 iterations with maximum statistics) are used to correct for multiple comparisons across electrodes and time points (Maris and Oostenveld 2007).

Different reward probability levels produced similar grand-average ERP waveforms during equal (Figure 5.8A) and single-option (Figure 5.8B) choices,

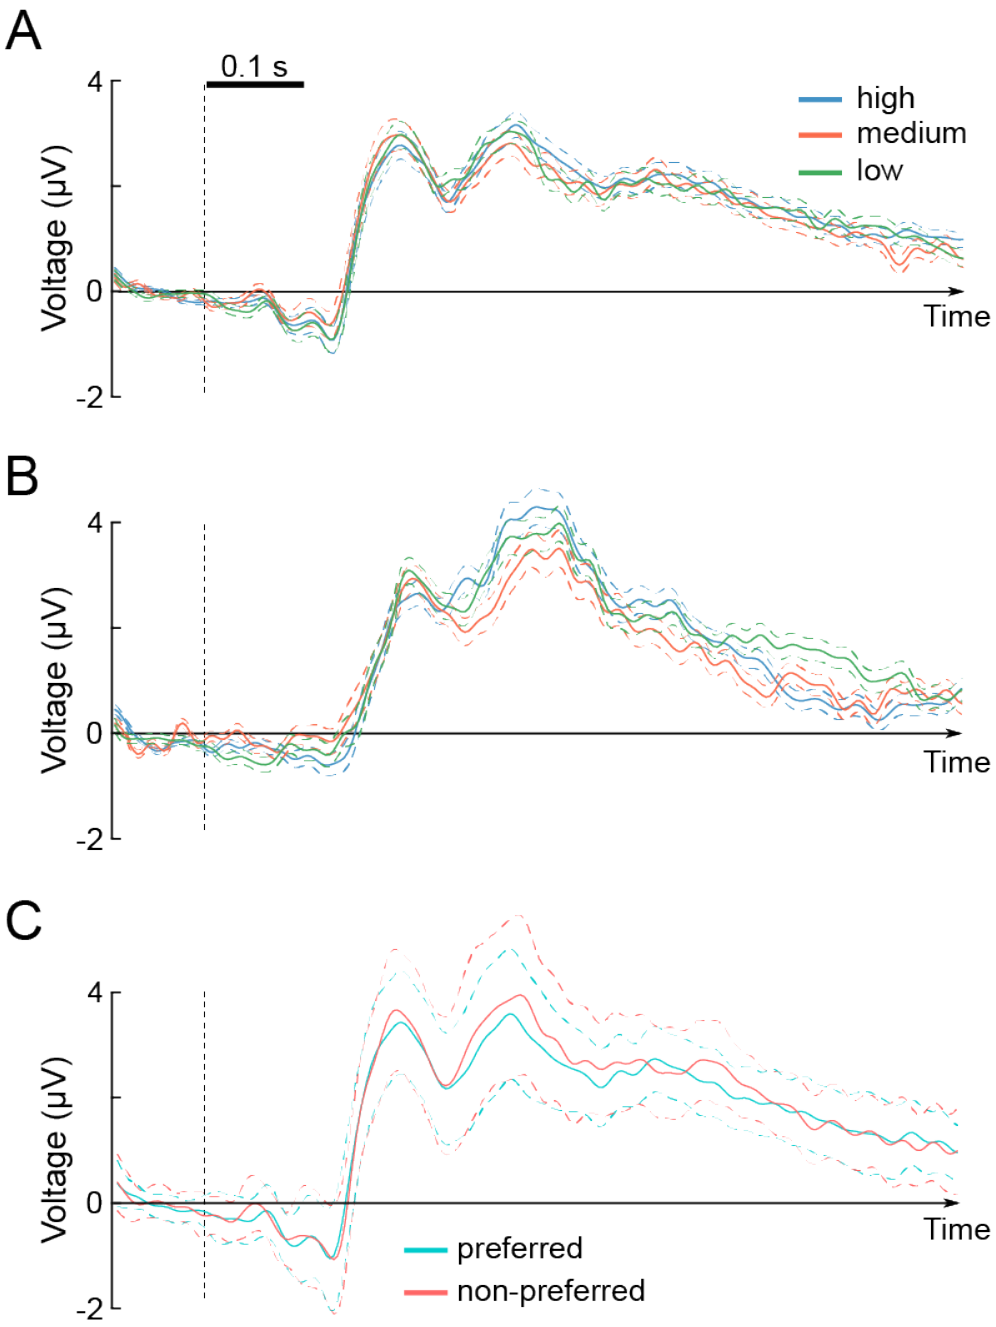

Figure 5.8: Grand-average stimulus-locked ERPs across all EEG electrodes. A. ERPs from high, medium and low reward probability in equal trials. B. ERPs from high, medium and low reward probability in single-option trials. C. ERPs from equal trials in which the preferred or non-preferred cue was chosen. In all panels, the dashed lines represent standard errors across participants.

with a negative peak in the 100 – 150 ms time window (the N100 component) and a positive peak in the 300 – 400 ms time window (the P300 component).

When assessing the effect of reward probability on ERPs, we found no univariate differences survived the correction for multiple comparisons in equal ($p = 0.552$ at all time points, cluster-level permutation test across electrodes and time points) or single-option trials ($p = 0.175$, cluster-level permutation test). For equal trials, we found no significant difference in ERPs between preferred and non-preferred choices (Figure 5.8C, $p = 0.208$, cluster-level permutation test). Therefore, in the current study, univariate ERPs were not sensitive to reward probability or preferred/non-preferred choices.

Multivariate Patterns in Equal Choices

To decode multivariate information representing reward probability in equal choice trials, we applied the linear SVM on multivariate EEG patterns across all electrodes (see Multivariate pattern analysis). Binary classification between high and medium reward probability was significantly above chance ($p = 0.01$, cluster permutation correction, non-parametric Wilcoxon test) from 144 ms after cue onset (Figure 5.9A). Similarly, the information between high and low reward probability was decodable above chance from 192 ms after cue onset ($p < 0.05$, cluster permutation correction). We found no significant classification accuracy between medium and low reward probability ($p > 0.16$ in all time points, uncorrected). Therefore, choices associated with certain (100%) rewards were distinguishable from those with uncertain reward probabilities.

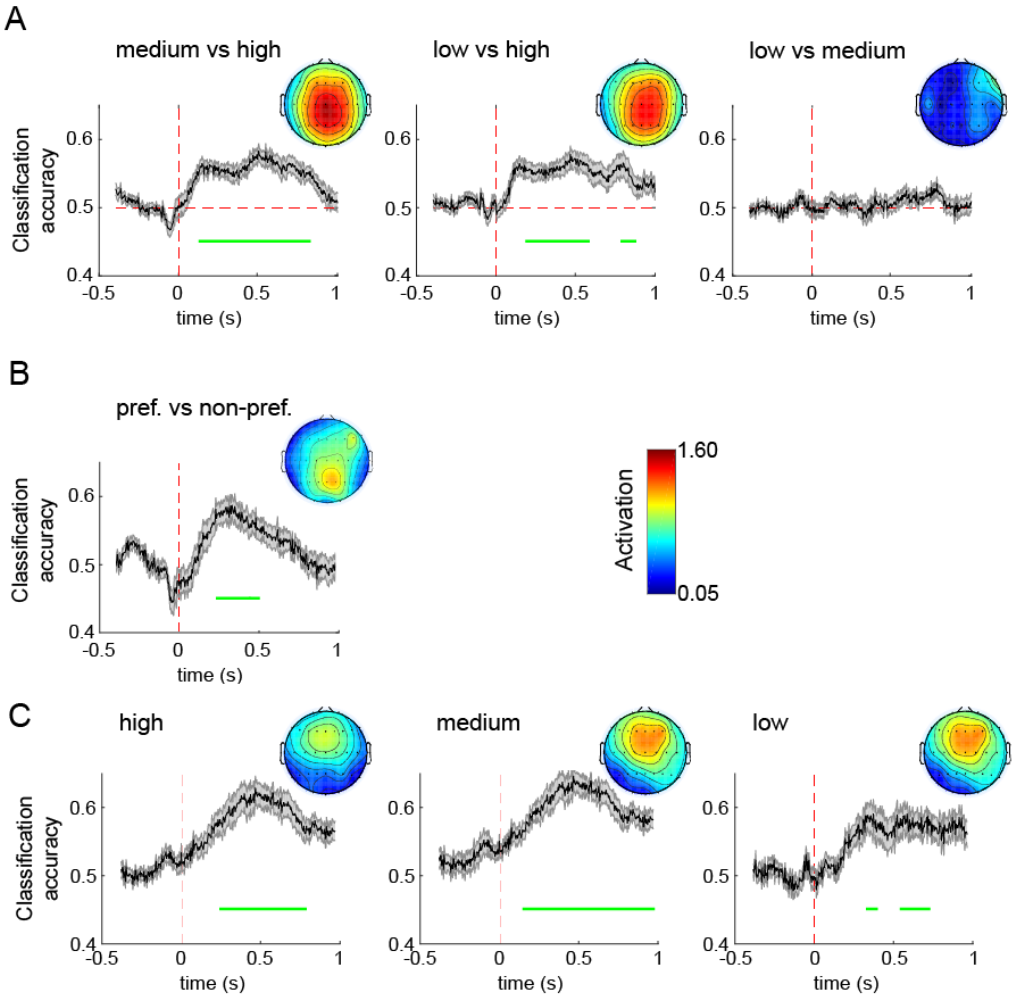

Figure 5.9: MVPA results. A. Classification accuracies across time-points between equal choices with different levels of reward probability. B. Classification accuracies across time-points between equal trials with preferred and non-preferred choices. C. Classification accuracies across time-points between equal and single-option choices with the same level of reward probability. In all panels, the black lines denote classification accuracies from a stratified 10-fold cross-validation and the gray areas denote standard errors. Significant decoding time windows (green horizontal bars) were determined from cluster-level permutation tests ($p < 0.05$, corrected). Topographic maps represent activation patterns from classification weights, which indicate the contribution of different EEG channels to overall classification accuracies.

We applied a similar classification procedure to decode the information between equal trials in which the participants chose their preferred or non-preferred choices across reward probability levels. The information about preferred versus non-preferred choices was decodable from 316 ms to 472 ms after cue onset ($p = 0.009$, cluster permutation correction).

To evaluate the relative importance of each feature (*i.e.* EEG electrode) to the classification performance, we calculated the weight vector of SVMs. For each classification problem, we retrained the SVM at each time point with all the data included in the training set and obtained the SVM weight vector. The weight vectors were then transformed into interpretable spatial patterns by multiplying the data covariance matrix (Haufe et al. 2014). The group spatial patterns were calculated by averaging across participants and from all time points which had significant classification accuracy. Relevance spatial patterns based on SVM's weight vector showed that mid-line central and posterior electrodes contained the most information for significant classification (Figure 5.9).

5.4.4 EEG-informed Cognitive Modelling

P300 component is a strong candidate for a marker of evidence accumulation. It's amplitude has been associated with attention (Datta et al. 2007), working memory (Kok 2001) and it's amplitude with task difficulty (Kok 2001). Prominent models propose it reflects build-to-threshold of the decision variable (D. M. Twomey et al. 2015; Kelly and O'Connell 2013) or marks the conclusion of internal decision-making process (Nieuwenhuis, Aston-Jones, and Cohen 2005). Considering that the latency of early visual processing

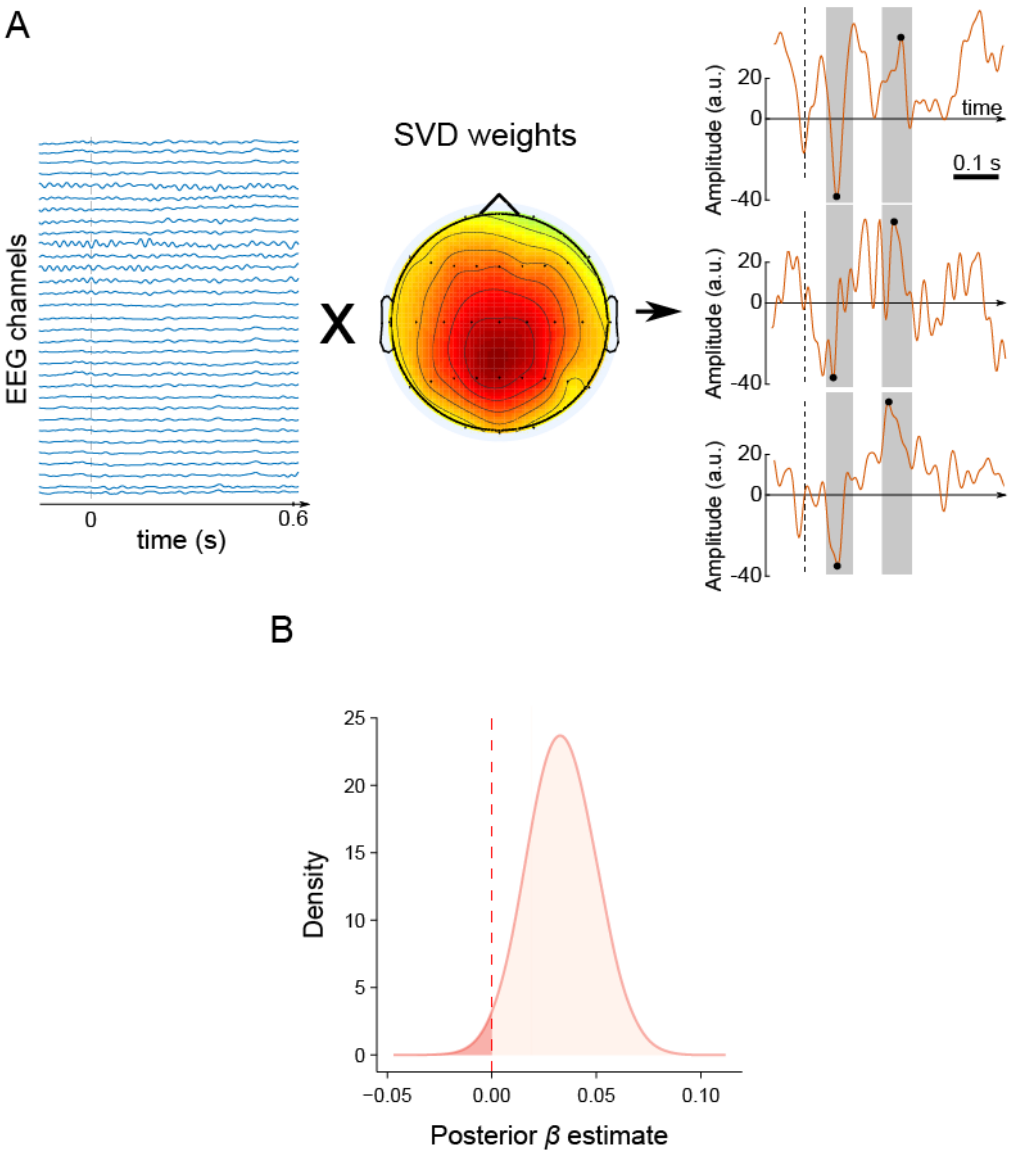

Figure 5.10: EEG-informed modelling. A. The schematic diagram of extracting single-trial ERP components. 32-channel EEG signals from a single trial were multiplied by the weights of the first SVD component, calculated from the grand-averaged ERP. Next, the N100 and P300 components in that trial were identified by searching for the peak amplitude in a time of 60-164 ms for the N100 component, and 272-376 ms for the P300 component, respectively. ERP marks in three representative trials were illustrated in the right column of the panel. The ratio between N100-P300 peak amplitude difference and N100-P300 peak latency difference was calculated as a single-trial regressor for modelling. B. Posterior estimates of the coefficient between the EEG-informed single-trial regressor (the rising slope of N100-P300 components) and changes in the accumulation rate.

is a part of non-decision time (Nunez et al. 2019), we further hypothesized that the evidence accumulation process initiates at N100 peak latency. This assumption is based on combined inferences from the MVPA analysis, showing very early components differentiating between conditions, and LBA modelling, indicating no differences in non-decision time, which suggests that the differences are related to evidence accumulation. Such early start of accumulation could be a product of the simple and unambiguous stimulus set consisting of basic shapes with no perceptual noise. This led to a theoretical prediction that the slope of the rise in EEG activity between N100 and P300 peak amplitudes reflected the accumulation rate on a trial-by-trial basis. To validate this prediction, we estimated the N100 and P300 components from single trials of equal choices (Figure 5.10A), using an SVD-based spatial filter to improve the signal-to-noise ratio of single-trial ERPs (see Estimation of single-trial ERP components). This single-trial EEG estimate was then added as a linear regressor (Equation 1) of the mean accumulation rate to the LBA model variant with the best fit to behavioural data (*i.e.* model 15 in Figure 5.4).

We used the same MCMC procedure to fit the extended LBA model with the EEG-informed regressor to the equal trial data. The extended LBA model showed good convergence ($\hat{R} \leq 1.02$ for all parameters) and provided a better fit, with a lower LOOIC score 2687 than the model without the EEG-informed regressor (LOOIC score 2796), suggesting that the rising slope of N100-P300 indeed affected the decision process. The posterior estimate of the regression coefficient β provided strong evidence for a positive single-trial effect (Figure 5.10B, $P_{p|D} = 0.983$), indicating that a bigger N100-300 slope

is associated with faster accumulation rate.

5.4.5 Additional Analyses: Alternative EEG Regressors and Representations of Choice Types

Is it possible that a simpler EEG-based regressor based on a single ERP component could provide a better model fit than the N100-P300 slope? To test this possibility, we fitted four additional extended LBA models with different single-trial EEG regressors applied to the mean accumulation rate: N100 peak latency, N100 peak amplitude, P300 peak latency and P300 peak latency. All the alternative regression models showed inferior fits (LOOIC scores larger than 2700) than the N100-P300 slope model. We therefore conclude the effects of single-trial EEG activity on the accumulation rate were related to both ERP components.

We did not observe above-chance classification between equal trials with the two levels of uncertain reward probability (Figure 5.9A). One may concern whether the lack of significant classification was due to the small number of trials in those conditions. To rule out this possibility, we conducted binary classifications to discriminate equal and single-option trials. The information about trial types (equal vs. single-option) was decodable at every level of reward probability (Figure 5.9C, $p < 0.05$, cluster corrected), including the one with the least number of trials (*i.e.* the low reward probability). This result was expected, given the large difference in stimulus presentation and behavioural performance between the two types of choices. SVM-based relevance patterns highlighted the middle central and frontal electrodes to

contain most of the information of trial types. These results suggested that the difference in classification accuracies between certain and uncertain reward conditions could not be readily caused by differences in the number of trials.

5.5 Discussion

We provide novel evidence that reward probability and spontaneous preference influence choices between equally probable alternatives and their electrophysiological signatures. We observed two patterns that were consistently distinct at behavioural, cognitive and neural levels: a *certainty effect*, distinguishing choices between cues with 100% reward probability and cues with uncertain reward probabilities (80% or 20%), and a *preference effect*, differentiating between equally valued options. At the behavioural level, reward certainty (*i.e.* 100% reward vs. non 100% rewards) resulted in disproportionately faster reaction times, while preference biased both choice frequency and RT, resulting in more frequent and faster responses for preferred cues. Using hierarchical Bayesian implementation of a cognitive model, we showed that reward certainty and preference bias were associated with changes in the accumulation rate, a model-derived parameter to account for the speed of evidence accumulation during decision-making. At the electrophysiological level, the information of certainty and preference could be reliably decoded from multivariate ERP patterns early during decisions, but not from univariate EEG activities. The accumulation rate was further affected by the slope of the rise in ERPs between the N100 and P300 components on a trial by

trial basis. Together, the current study provides insight into neurocognitive mechanisms driving choices in a deadlock situation, where there is no clear advantage in choosing one option over the other.

The certainty effect implies a monotonic but nonlinear relationship between reward probability and RT in equal choices: the difference between certain (100%) and uncertain (80% and 20%) reward was greater than that between the two uncertain conditions. This points to a special status of the 100% reward certainty distinct from lower reward probabilities, as the latter always carries a non-zero risk of no reward. The salient representation of the 100% reward certainty is further highlighted by the lack of significant EEG pattern classification between the two uncertain reward probabilities (80% vs. 20%, Figure 5.9A). Here, the certainty effect in rapid voluntary decisions resembles risk-averse behavior in economic decisions (Tversky and Kahneman 1989), which overweights outcomes with 100% certainty relative to probable ones.

One has to be cautious while interpreting the differences in MVPA pattern between conditions. While it is tempting to assume that 100% reward certainty is driving the observed EEG differences, it is also possible that uncertain cues drove the observed choice biases. This is due to the fact that MVPA cannot indicate in which condition signal was stronger, but rather whether differences in response pattern exist (Stelzer, Chen, and Turner 2013). Given no differences in grand-averaged ERP signal, it is impossible to tell which condition drove the differences in brain activation. While our current paradigm does not provide a clear answer to this question, a dedicated experimental design systematically manipulating different levels of uncertainty, and a analysis using spatially integrated MVPA output with

EEG signals passed through the weights could resolve this ambiguity.

Interestingly, reward probability affected RTs across all trial types. It persisted from equal choices to simple reactions to cue locations in single-option trials (Figure 5.3A). In unequal choices, there was also a negative association between RT and the sum of reward probability of the two choices (Figure 5.3B). Therefore, even though the reward was not contingent upon RT in the current study, we observed a general tendency of accelerating ones' responses in the presence of more certain reward. These results are akin to the effect of reward magnitude, which also demonstrates a facilitating effect on RT (Schurman and Belcher 1974; Chen and Kwak 2017). In non-human primates, the phasic activation of dopamine neurons in the ventral midbrain has similar response profiles to changes in reward probability and magnitude (Fiorillo, Tobler, and Schultz 2003), suggesting a common mesolimbic dopaminergic pathway underlying different facets of reward processing that affect decision-making.

Bayesian model comparison identified specific effects of reward probability on accumulation rates, highlighting two possible cognitive origins of the certainty effect. First, in equal choices, cues with 100% certain reward resulted in larger mean accumulation rates than those with uncertain reward probabilities (Figure 5.7). Accumulation rate has been linked to the allocation of attention on the task (Schmiedek et al. 2007). Because reward plays a key role in setting both voluntary (top-down) and stimulus-driven (bottom-up) attentional priority (Libera and Chelazzi 2006; Raymond and O'Brien 2009; Krebs, Boehler, and Woldorff 2010; Won and Leber 2016), high reward probability may boost the attentional resources allocated to sensory processing

for more rapid decisions. Second, reward probability affected the variability of accumulation rates across trials. Higher accumulation rate variability has been associated with better-memorized items (Starns and Ratcliff 2014; Osth, Dennis, and Heathcote 2017; Tillman et al. 2017). It is possible that stimuli associated with 100% certain reward were memorized more strongly (Miendlarzewska, Bavelier, and Schwartz 2016), a hypothesis to be confirmed in future studies.

Furthermore, MVPA of stimulus-locked ERPs showed multivariate EEG patterns distinguishing between cues with 100% certain reward and other uncertain reward probabilities as early as ~150 ms after stimulus onset (Fig 5.9A; see also Thomas, Vanni-Mercier, and Dreher (2013)), and model comparisons found no evidence to support for non-decision time to vary between reward probability levels (Figure 5.4). Considering the average RT of 600 $\sim$ 900 ms in equal choices, our results did not support the latency of post-decision motor preparation, which constitutes a part of the non-decision time (Karahan et al. 2019), to be the source of the certainty effect. This result is consistent with the view that motor action implementation is independent of the stimulus value (Marshall, Bogacz, and Gilchrist 2012). Instead, the certainty effect possibly originates from evidence accumulation during the decision process, as supported by the changes in the accumulation rate.

When choosing between equally valued options, classical evidence accumulation theories predict a deadlock scenario with a prolonged decision process (Bogacz et al. 2006). This was not supported by recent experimental findings in value-based decisions (Pirrone, Azab, Hayden, Stafford, and Marshall 2018a; Teodorescu, Moran, and Usher 2016), including the current study, in

which equal choices took no longer than unequal ones. Our behavioural, modelling and EEG analyses indicated a preference bias which could effectively serve as a cognitive mechanism to break the decision deadlock. Compared with non-preferred options, preferred decisions facilitated RTs, were associated with larger accumulation rates and evoked distinct EEG multivariate patterns. Here, we did not aim to provide a mechanistic interpretation of preference (*i.e.* why or how the preference bias originated). Instead, our results demonstrated a consistent presence of preference bias before and after cue-probability re-mapping, independently across reward probabilities (Figure 5.2A) and maintained in single-option trials (Figure 5.3A), which we considered as a novel finding in the literature of voluntary choice.

What can induce a preference bias? Because the cue-probability association was initially randomized and later changed within each session, and no differences in shape preference were found, this bias was not due to stimulus salience but established spontaneously (Voigt et al. 2019). Multiple factors may contribute to the establishment of preferred options. Preference might arise as a function of early choices and outcome frequencies (Izuma et al. 2010; Bakkour et al. 2018), which shape future beliefs or alter the memory trace of certain cue-probability bindings. This interpretation is consistent with an irrational bias, which favors previously rewarded stimuli, even when controlling for their value (Scholl et al. 2015). Alternatively, some cue-value associations might be remembered more reliably due to a deliberate cognitive strategy of memory resource allocation.

Our results provide little evidence to either support or refute these hypotheses. However, memory strength alone cannot explain the full set of results in the

current study. First, it is worth noting that the stimulus-reward mapping was presented a total of 16 times throughout each session (at the beginning of each block and after every 40 trials), and participants took as much time as they needed before the next set of trials.

Secondly, purely memory-based interpretation cannot support both equal and unequal condition results taken together (Figure 5.2). Response times in equal trials would suggest that memorization is a function of value, since the 100% probability condition was associated with fastest RTs, while the 20% probability with the slowest. Looking at the pattern of accuracies in the the unequal condition paints a different picture however: here the accuracy is a function of value difference. A purely memory-based interpretation of this pattern would suggest that medium value symbols are remembered the least, leading to a higher error rate in trials involving them.

Similarly, the linear mixed-effect models found significant effects of preference on RT only in equal and single-option trials, but not in unequal trials (Figure 5.2). If we were to believe memorization of items to be different between two cues of the same reward probability, we would expect this to be reflected also in the unequal condition, which was not the case. This pattern of results might suggest differences related to the strategy of choice. An potential interpretation for lower accuracy rate in the high vs medium condition could be a satisficing strategy, where 80% (medium) rewards are considered *good enough* and sometimes chosen over the 100%. Such strategy would make sense from resource allocation perspective, given that items are examined sequentially - then simply observing only a single cue, choosing the 80% symbol without paying attention to the alternative can be rational.

Such strategy adjustments could have effects on the interpretation of our modelling and EEG findings, such as testing more sophisticated models which can account for dynamical threshold adjustment during the accumulation process.

Future studies could validate these hypotheses by employing more frequent cue-probability remapping throughout experiments and controlling for memory effects. Furthermore, all trials in the current studies were randomized and participants did not have prior knowledge of upcoming stimuli. One future extension would be to evaluate whether presenting prior information of reward probability in an upcoming trial would modulate boundary separation in voluntary decisions, similar to the effect of prior bias on perceptual decisions (Mulder et al. 2012).

The current study considered a simplified form of decision, in which the amount of reward was fixed (10 game points). In traditional value-based decisions assumed by the prospect theory, a decision-maker needs to integrate the value and probability of gain or loss to obtain an expected utility for each option (Tversky and Kahneman 1992). Together, our results here and previous studies (Wagner et al. 2020) provide converging evidence that both reward value and probability can influence RT in equal choices. This raises the intriguing possibility of our results to be generalized to choices with the same expected utility but the different combinatory of value and probability. Interestingly, the multiattribute extension of the LBA model (Trueblood, Brown, and Heathcote 2014) has been fitted to RTs from such tasks (Cohen, Kang, and Leise 2017), suggesting that our modelling and EEG approaches could also be extended to explore more complex decision problems.

Our study highlights the advantages of EEG-informed cognitive modelling to inform behavioural data. Hierarchical Bayesian parameter estimation of the LBA model provides a robust fit to an individual's behavioural performance with less experimental data needed than other model-fitting methods (Vandekerckhove, Tuerlinckx, and Lee 2011; Wiecki, Sofer, and Frank 2013; Zhang et al. 2016). By integrating single-trial EEG regressors with the cognitive model, we identified the accumulation rate to be affected by the rate of EEG activity changes between visual N100 and P300 components. This result contributes to a growing literature of EEG markers of evidence accumulation processes, including ERP components (D. M. Twomey et al. 2015; Loughnane et al. 2016; Nunez, Vandekerckhove, and Srinivasan 2017), readiness potential (Lui et al. 2018) and oscillatory power (Vugt et al. 2012). It further consolidates the validity of evidence accumulation as a common computational mechanism leading to voluntary choices of rewarding stimuli (Summerfield and Tsetsos 2012; Afacan-Seref et al. 2018; Maoz et al. 2019), beyond its common applications to perceptually difficult and temporally extended paradigms.

The EEG-informed modelling builds upon the known functional link between the P300 component and evidence accumulation for decisions (Polich, Ellerson, and Cohen 1996; Verleger, Jaśkowski, and Wascher 2005; D. M. Twomey et al. 2015). A new extension in the current study was to consider the accumulation process begins at the peak latency of the visual N100 component. The slope regressor provided superior fit in comparison to other alternatives, such as latencies and peaks of the individual components, providing an argument for the accumulation slope being affected, instead of the starting point,

termination, or total evidence. Theoretically, the delayed initiation of the decision process accounts for information transmission time of $60 \sim 80$ ms from the retina (Schmolesky et al. 1998). Single-unit recording concur with this pre-decision delay, as neurons in putative evidence accumulation regions exhibit a transient dip and recovery activity independent of decisions approximately 90 ms after stimulus onset (Roitman and Shadlen 2002). Practically, our EEG data has a clear N100 component, and time-resolved MVPA identified significant pattern differentiating between task conditions at a similar latency. The relatively early start of the accumulation process in our experiment might be explained by the easily discriminable nature of the cues, consisting of basic shapes with no perceptual noise. Longer visual processing stage has been reported in an experiment involving more complex processing of visual information (Nunez et al. 2019). An alternative possibility is that the differences caught very early reflect non-decision time, consistent with literature suggesting later accumulation start (Nunez, Vandekerckhove, and Srinivasan 2017; D. M. Twomey et al. 2015). While this interpretation stands against our modelling results, it is not impossible for our model not be sensitive enough to catch small differences in early visual processing (Goldfarb et al. 2014).. Further research could dissect the non-decision time (White et al. 2014; Tomassini et al. 2019) and compare latencies of visual encoding across decision tasks with stimuli at different levels of complexity. Decoupling early visual processing from motor execution (Lui et al. 2018) could bring new insights into this issue.

Several issues require further consideration. First, our cognitive modeling was not meant to reproduce all the rich behavioural features in the data. To

include sufficient observations for model-fitting, we combined the data before and after cue-probability remapping. As a result, our model did not account for behavioural changes related to cue remapping. Future studies could employ a multi-session design to investigate how learning new cue-probability associations influence model parameters (Zhang and Rowe 2014).

Second, we focused on the certainty and preference effects by fitting the LBA model only to the data from equal choices. Although simulations indicated that the fitted model provided similar behavioural patterns as in the empirical data in unequal and single-option choices, it was not fitted directly to the experiment data in those two choice conditions. A more parsimonious model for all three types of choices would require additional assumptions, which is beyond the scope of the current study. For example, to incorporate the large RT discrepancy between equal and single-option choices, one could assume that the urgency signal (Boehm et al. 2016; Thura and Cisek 2017) plays a more dominant role in accelerating RT when no apparent comparisons are needed in single-option choices.

Third, our model selection procedure does not encompass all conceivable model types that might account for this data. Independent accumulation is consistent with findings on brain mechanisms of probability-based choices in humans (Kolling, Wittmann, and Rushworth 2014; Scholl et al. 2015), as well as choice behavior in rats (Ojeda, Murphy, and Kacelnik 2018). Alternative explanations of certainty and preference effects can be provided by urgency gating (Thura et al. 2012), collapsing threshold (Ratcliff and Frank 2012) or cross-inhibition (Pais et al. 2013; Usher and McClelland 2001). Depending on the parametrization, interpretations based on these models could slightly

vary. It is also important to note that there may be no straightforward way to disentangle the interpretations provided by these different models (Miletić and Maanen 2019). These potential differences however, although important, would not challenge the main conclusions of this paper.

Finally, our model-based analysis is unavoidably constrained by the choice of model, and one needs to be cautious when extending findings to different models. There is an ongoing debate on how accurately different models can mimic each other when estimating the non-decision time (Goldfarb et al. 2014; Lerche and Voss 2018). The DDM, for example, tends to predict longer non-decision times than LBA (Dutilh et al. 2019), as well as might be more susceptible to urgency manipulations (Evans 2020). Although an extended DDM has been shown to account for magnitude effects (Ratcliff, Voskuilen, and Teodorescu 2018), the drift rate of a DDM represents the relative signal difference between the two options. As a result, without fitting a new DDM to each condition, the DDM cannot directly describe all conditions in our current study (*i.e.* the unequal and single-option trials). LBA, on the other hand, assumes an independent accumulator for each option, offering a parsimonious account to our task and the capacity to produce the qualitative features of responses in all conditions, as demonstrated in our model simulations.

5.6 Conclusion

When choosing between equally probable reward outcomes, probability and preference selectively modulate the decision processes and their electrophysiological signatures, providing a mechanism for breaking a decision deadlock.

These findings extend and substantiate the computational framework of evidence accumulation for voluntary decisions. Our results further highlight the intricate nature of human behavior, as susceptible to external factors as well as endogenous heuristics.

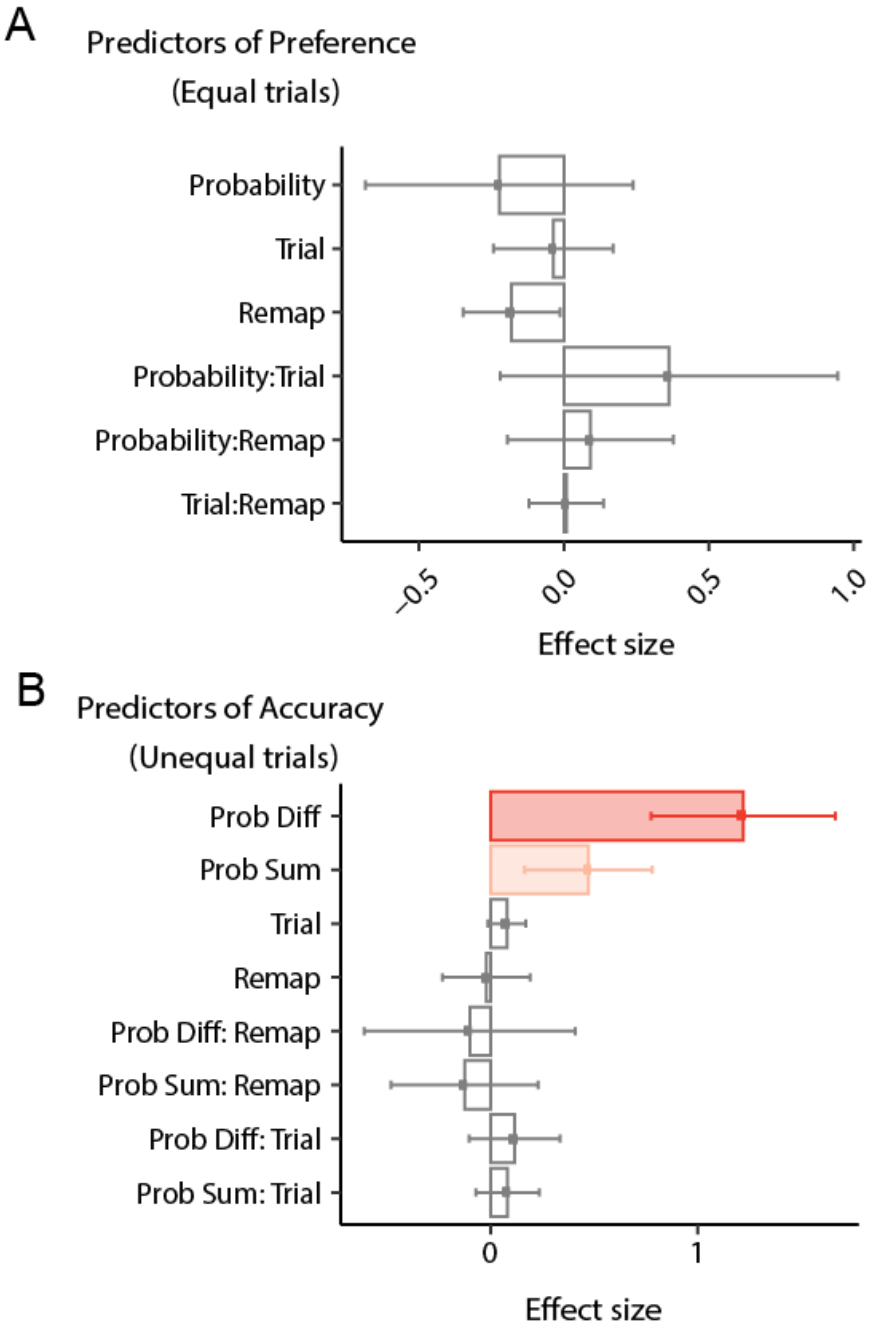

Figure 5.11: Effects of trial order and cue remapping on (A) preference in equal trials and (B) accuracy in unequal trials. Preference and accuracy were coded as binary dependent variables (whether the choice trial was preferred / associated with a higher reward probability). Rows represent tested fixed effects in each linear mixed-effects model. Remapping was coded as a binary variable (whether the trial occurred pre or post remapping). Trial order was defined with respect to each block (from 1 to 64). Both models contain varying individual intercepts as random effects. Dark red bars represent significant effects with $p < 0.001$. Light red bars represent significant effects with $p < 0.05$. Grey bars represent non-significant predictors. Error bars represent 95% confidence intervals.

Chapter 6

Within and Across-Domain Effects of Choice Induced Bias (Behavioural Experiments)

6.1 Introduction

According to economic theory, choices are a passive reflection of underlying preferences (Rangel, Camerer, and Montague 2008). This view has been challenged by psychological research, indicating that choices can also have causal power over preferences (Vinckier et al. 2019; Koster, Duzel, and Dolan

2015; Sharot, Velasquez, and Dolan 2010; Voigt, Murawski, and Bode 2017) leading to a positive feedback loop: the more an option is chosen, the greater its value, the more likely it is to be chosen in the future.

The effects of choice-induced bias (CIB) have been demonstrated in different domains (subjective preference: Vinckier et al. 2019; Harmon-Jones et al. 2008; Koster, Duzel, and Dolan 2015; Voigt et al. 2019; perception: Akaishi et al. 2014; Bode et al. 2012; Bronfman et al. 2015), and across timeframes, from immediate, trial-level effects (Glickman, Moran, and Usher 2020; Talluri et al. 2018), to long-lasting changes in preference (Sharot et al. 2012). CIB can affect both cognitive representations (Mather, Shafir, and Johnson 2000, 2003; Salti et al. 2014) and neural activity (Chammat et al. 2017; Izuma et al. 2010; Luettgau et al. 2020; Sharot, De Martino, and Dolan 2009). Combined, these findings suggest that the phenomenon is robust. It remains unclear however whether the effect found in perceptual studies is driven by similar mechanisms to that described in the preference literature.

The comparison of the CIB between different domains leads to another untested issue: is it possible for the bias to transfer from one domain to another? For example, a hypothetical voter is motivated to vote for candidate A over candidate B due to his charisma. Firstly, the act of choice (vote) can further widen the gap in perceived charisma in favour of the chosen candidate. This is a classic case of CIB, where the affected beliefs are causally associated with the domain of choice (*"I find candidate A more charismatic, therefore I vote for him, which leads me to find him even more charismatic"*). Additionally, the vote could potentially affect how one views the candidates in unrelated domains, such as proposed policies (*"I find candidate A more*

charismatic, therefore I vote for him, which leads me to view his policies more favourably"). We refer to these as a *within-domain*, and *across-domain* effects, respectively.

CIB has been commonly associated with two alternative mechanisms: a motivated conflict resolution *via* dissonance reduction (Aronson, Blanton, and Cooper 1995; Brehm 1956; E. Harmon-Jones and Mills 2019; Kenworthy et al. 2011), or inferring value based on one's own choices (Akaishi et al. 2014; Cockburn, Collins, and Frank 2014; Chammat et al. 2017; Miyagi, Miyatani, and Nakao 2017). The first mechanism stems from cognitive dissonance theory (Festinger 1957), arguing that people are intrinsically motivated to keep an internal consistency between their choices and judgments, even when facing contradictory evidence. The second proposal originates from Bem's autoperception theory (Bem 1967) and argues for an epistemic interpretation (Kruglanski et al. 2018) where conflict or an affective experience are unnecessary, and the bias reflects choice-driven expectancy update (Friston et al. 2016) or inference based on the explicit memory of previous choices (*"I remember choosing it, so I must like it"*).

Here, we compare the behavioural effects and cognitive mechanisms of CIB *between* and *across* preference and perceptual domains in a set of 5 experiments. We use a task where choices between 2 items in either domain are immediately followed by a comparative judgment of the same two items on a continuous scale, reflecting the difference in their value estimations. The design allows us to test both within-domain (when choice and judgment domains are consistent across the trial) and across-domain effects (when choice and judgment are inconsistent across the trial; *e.g.* size choice followed

by preference judgment).

Firstly, we compare the qualitative (driving factors) and quantitative (effect sizes) differences between within-domain effects in preference and perception. An important qualitative distinction is that preference is more prone to time-dependent fluctuations than perception (one's appetite for pizza might be different in the morning compared to the afternoon, but the judgment of it's size will likely remain similar), hence one could expect internal preference values to exhibit more malleability towards choice. Preference CIB has also been shown to have a long-term effect (*e.g.* Sharot et al. 2012), while perceptual CIB has been studied mostly in terms of immediate trial-level effects (Akaishi et al. 2014; Bronfman et al. 2015; Glickman, Moran, and Usher 2020). Hence it remains an open question whether CIB can be influenced by a longer history of choices.

Other important predictors of CIB involve choice difficulty and magnitude (Voigt, Murawski, and Bode 2017). In accordance with earlier chapters and similarly to how the concept is used in related literature (Liberman and Förster 2006), difficulty here will be conceptualized as the difference in value between options, while magnitude as the sum of option values (see: *Chapter 1*). In relation to studies of CIB in the cognitive dissonance framework, difficulty is often equivalent to conflict, which drives and scales with bias (Coppin et al. 2014; Izuma et al. 2010; Koster, Duzel, and Dolan 2015). Magnitude on the other hand is similar to the concept of choice importance, as higher-valued options are often perceived as more important (Pais et al. 2013; Pirrone, Stafford, and Marshall 2014a), which can change the subsequent choice strategy (Pirrone, Azab, Hayden, Stafford, and Marshall

2018a; Teodorescu, Moran, and Usher 2016). Similarly, my previous research has indicated that probability magnitude can drastically reduce decision times (see: *Chapter 5*). These findings suggest that choice magnitude might also influence choice-induced bias.

Secondly, we test whether a significant effect of an *across-domain* bias can be determined, and if so, describe its characteristics. Sequential choice effects in the perceptual literature indicate a limited capacity of the human mind in the assessment of subsequent independent events in a single domain. Whether such effects can transfer across independent domains remains to be tested. Existence of an *across-domain* effect would be challenging for hypotheses postulating an adaptive role of the bias (Lee and Daunizeau 2020; Harmon-Jones, Harmon-Jones, and Levy 2015; Kaaronen and Dale 2018; Kruglanski et al. 2018; Peters 2020). It could also provide a framework for understanding a range of real life phenomena where certain choices or actions can lead to developing seemingly irrational attitudes, such as voting preferences inconsistent with self-interest (Fishbein and Coombs 1974), or irrational economic behaviors (Becker 1962).

Thirdly, we model the generative process, distinguishing between two mechanisms responsible for the bias: a consistency-driven domain-general effect which affects immediate evaluation but has no effect on the underlying values, and a domain-sensitive value-update mechanism. First of the proposed mechanisms, the consistency-driven effect, is similar to the time-dependency effects found in many perceptual studies, where a choice modulates the gain of upcoming sensory information (Bronfman et al. 2015; Talluri et al. 2018; Urai et al. 2019), and can be related to the cognitive dissonance theory. In

this view, the choice conflict and importance elevate experienced discomfort (Harmon-Jones 1985; E. Harmon-Jones and Harmon-Jones 2019), which can be reduced by adjusting one's expectations in favour of the chosen option or against the rejected one. The value-update mechanism, on the other hand, is a rational update of one's beliefs based on remembered choices, similar to a reinforcement learning mechanism (Chambon et al. 2020; Cockburn, Collins, and Frank 2014) with implicit rewards. The rationale being that one can learn their own preferences in a similar way they can learn the structure of the surrounding environment.

Two crucial factors that differentiate the two explanations are context-sensitivity and effect longevity. If the bias is driven by a need for choice consistency irrespective of the context, one might expect it to be domain-general, *i.e.* independent of the type of choice, as well as short-lived, as consistency exhibited in the past no longer holds relevance when future choices are made. In contrast, value-sensitive updates should be sensitive to the domain of choice (person's preference-based choices should ideally affect only his preference values, but not perceptual estimations, and *vice versa*) while its temporal effects are constrained by one's memory capacity for previous choices.

Finally, we use different variations of the task design to test some of the limits and alternative explanations of the CIB effect. Specifically, we test whether voluntary choice is necessary or can it be evoked by an instructed choice (Egan, Bloom, and Santos 2010; Sharot, Velasquez, and Dolan 2010) (Experiment 3), compare the strength of the effect independently for the chosen (overestimation) and rejected (underestimation) items (Experiment

Table 6.1: Demographic information across experiments. *N* column contains the total number of participants, followed by the number of males within that sample (m). *Exc* column sums up all participants that were excluded, based on preregistered exclusion criteria. Columns 3 and 4 represent the mean and standard deviation of participants per experiment, respectively.

Exp	N	Exc	M_{age}	SD_{age}
1	23 (5m)	0	21.65	2.21
2	250 (155m)	14	26.35	8.20
3	50 (36m)	1	24.96	5.94
4	50 (34m)	1	24.67	7.04
5	50 (30m)	8	25.75	9.30

4), and testing to what extent these effects can be accounted by a positive reference (Hunt et al. 2016), *i.e.,* interpreting the chosen item as the default option (Experiment 5).

6.2 General Methods

This section contains methods shared across all of the experiments, including the stimuli, general task structure, operationalization of variables of interest and behavioural modeling methodology. Exact specifications regarding each experiment can be found in the experiment-specific method sections.

6.2.1 Stimuli

We used food items from *Food-Pics* database (Blechert et al. 2019).

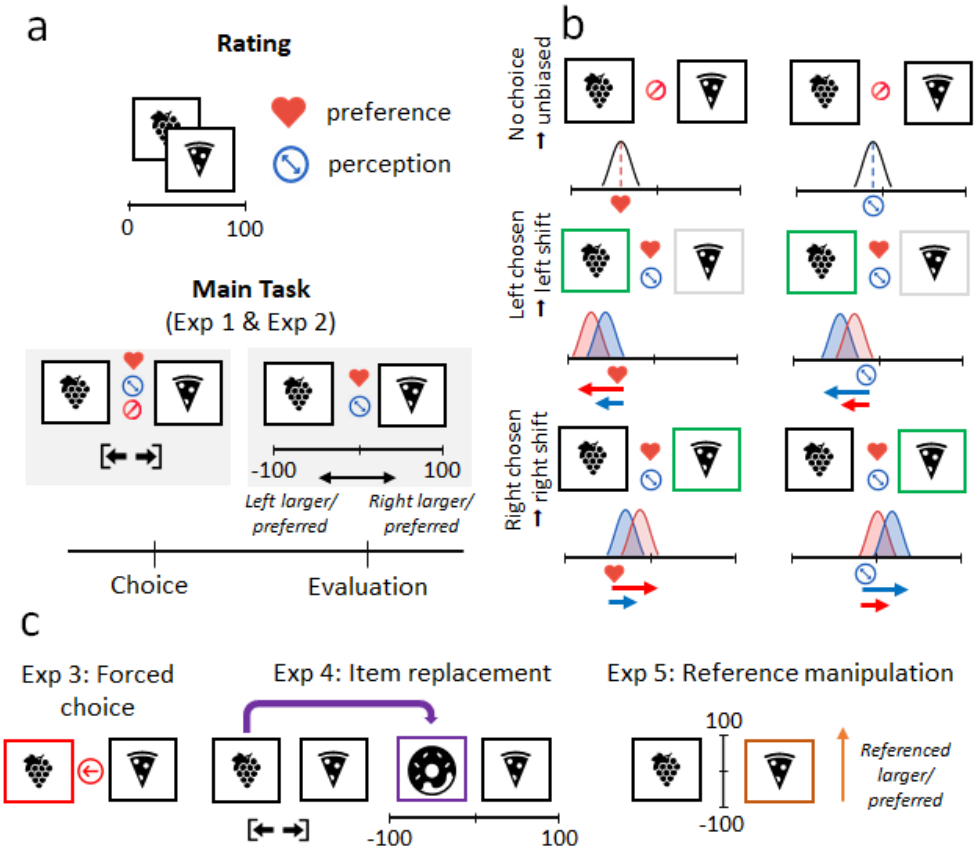

Figure 6.1: Task design and predicted effects. Panel a) represents the task structure, consisting of item ratings (repeated twice before the main task and once after) and the main task. Each trial of the task consisted of a choice followed by a judgment. Panel b) illustrates expected effects on of choice on the judgment phase. Columns represent judgment types; rows: choice types. Panel c) represents design adjustments in Experiments 3-5. Preference (heart) and perception (blue circle with arrow) symbols represent the types of choice and judgments. Red crossed circle represents the no-choice condition present in experiments 1 and 2. Red arrow within a circle (panel c) represents the forced choice condition.

Table 6.2: Design differences across experiments. Choice types refer to 2-alternative force choice. Judgments refer to estimation of difference between item values on a continuous scale from -100 to 100. *Pref*, *Size* and *No* refer to preference-based, size-based or no-choice conditions, respectively. *Dom* refers to the domain of the judgment. *Stay* judgment factor refers to which item out of a unique pair was replaced (and which stayed) during the judgment. *Ref* refers to which item was considered the referenced item out of each unique pair. Unique pair column specifies how often (and in which conditions) was a unique pair of two items repeated.

Exp	Choice Types	Judgment Types	Trials	Unique Pair Reps
1	Pref Size No	Dom. (Pref, Size)	906	6 times (1 x cond)
2	Pref Size No	Dom. (Pref, Size)	198	3 times (1 x choice type)
3	Pref Size Forced	Dom. (Pref, Size)	264	4 times (1 x choice type)
4	Pref Size No	Dom. (Pref, Size) x Stay (Item1, Item2)	264	4 times (1 x judg. type)
5	Pref Size No	Dom. (Pref, Size) x Ref. (Item1, Item2)	264	4 times (1 x judg. type)

6.2.2 Procedure

The experiment consisted of three main stages (Figure 6.1). In the first and third stages, participants rated their preference and the size of each stimulus. The second stage was the main task, where on each trial participants made a choice between two items, followed by a comparative judgment of the items on a continuous scale. Choice and judgment conditions varied across experiments (see Table 6.2).

Stage 1: Initial rating. Participants performed initial preference and size-based ratings of food items on a continuous scale from 1 to 100. For preference-based ratings, participants rated item desirability from strong dislike (1) to strong liking (100). In the perceptual-based rating, participants rated the size of each food picture with respect to the white background (*i.e.* the percentage of non-white pixels). Each participant completed two iterations of both types of ratings. Rating type order was counterbalanced across participants and item order was randomized within each rating.

On each rating trial, a food stimulus was presented in the center of the screen. A task cue (red heart for preference or a blue circle with outward arrows for size; see Figure 6.1) was presented on top of the stimulus to indicate the type of rating. A white rating scale was presented below the food stimulus with markers placed on both ends.

Stage 2: Main Task. Each trial of the main task consisted of a choice between 2 items followed by a judgment of difference between item values on a continuous scale from -100 to 100 (Figure 6.1).

Experiments consisted of different types of choices (preference, size, no-choice

or forced choice) and two types of judgments (preference or size). No-choice condition required participants to withhold a response while observing the items on screen. Trials in which choice and judgment are from the size domain (both referring to size or both refrering to preference) are referred to as *congruent*, while trials in which choice and judgment are from different domains (size-based choice followed by preference-based judgment or *vice versa*) are referred to as *incongruent*. No-choice trials are referred to as *neutral*.

Stage 3: Final rating. After completing Stage 2, participants performed the preference and perceptual-based ratings for the third time for all items.

6.2.3 Operarationalization of Variables of Interest

Right-side judgment bias (RB) was defined as a difference between the judgment J and the difference between right and left item values in choice domain d on trial t of the main phase of each experiment:

$$RB_t = J_t - (\bar{I}_{d,right} - \bar{I}_{d,left}) \tag{6.1}$$

where $\bar{I}_{d,right}$ and $\bar{I}_{d,left}$ are the means of initial ratings for the right and left item respectively. Similarly, choice-induced bias (CB) was calculated by conditioning the sign on the choice:

$$CB_t = \begin{cases} J_t - (\bar{I}_{d,right} - \bar{I}_{d,left}) & \text{if } choice_t = right \\ -(J_t - (\bar{I}_{d,right} - \bar{I}_{d,left})) & \text{if } choice_t = left \end{cases} \tag{6.2}$$

Choice-induced bias was calculated in respect to the initial ratings to the initial scale and also refers to the judgment domain (perceptual bias can occur during size judgments, while preference bias can occur during preference judgments. Within-domain bias can occur in trials when both choice and judgment belong to the same domain, while across-domain bias, when preference choice was followed by a size judgment, or *vice versa*.

Choice or judgment *difficulty* was defined as the inverse of the absolute value difference in values between items, such that:

$$Diff_{t,d} = 100 - |(\bar{I}_{d,right} - \bar{I}_{d,left})| \tag{6.3}$$

where $Diff_{t,d}$ represents the difficulty on trial t in domain d. The closer the item values are, the higher the difficulty. Difficulty values are naturally confined to values between 0 and 100, since the largest possible difference between initial ratings is equal to 100, and the lowest to 0.

Choice or judgment *magnitude* is defined as the sum of item values, such that:

$$Mag_{t,d} = \bar{I}_{d,right} + \bar{I}_{d,left} \tag{6.4}$$

where $Mag_{t,d}$ represents the magnitude on trial t in domain d. Magnitude values are constrained to values between 0 (when both items were rated at 0) and 200 (when both items were rated at 100).

Choice consistency was defined as the percentage of times the item, which was rated as more valuable during the rating phase on a given scale, had

been chosen.

Judgment consistency was assessed by binarizing the judgments (negative values indicate judgment towards the left item; positive - towards the right one) and then, similarly to choice consistency, calculated the percentage of judgments consistent with the initial ratings.

6.2.4 Behavioural Modelling

For all experiments, we report the results from Bayesian Mixed-Effects models using standard priors from the *brms* R package (Burkner 2017). Random effects structure includes intercepts and slopes for all regressors of interest (exact model specification can be found in *Appendix*). We report posterior distribution medians (*med*), 95% posterior credible intervals (CI), probability of directionality (pd; % of posterior density larger or smaller than 0), and, where the result is found to be significant, % of posterior in the Region of Practical Equivalence (ROPE) (Makowski et al. 2019). Analogous to frequentist analysis, we define significance as pd > 0.95. ROPE can be interpreted as an arbitrary null region where effect size is small enough not to be meaningful (Kruschke and Liddell 2018). We set the ROPE range to [-2, 2] ($\pm 1\%$ of the scale), as finer differences would be difficult for participants to distinguish. Frequentist analyses can be found in the *Appendix*.

6.3 Experiment 1

In our first experiment we tested the hypotheses regarding within-domain and across-domain CIB and their potential predictors (difficulty, magnitude) in a lab-based setting.

Participants performed a 2-step main task, in which they made binary choices, followed by judgments of difference between the item values on a continuous scale (Figure 6.1). Choices and judgments could belong to either perceptual (size) or preference domains. Additionally, on some trials participants only examined the items without making a choice (no-choice condition).

This design allowed us to test how choices influence subsequent judgment within-domain (when preference choice was followed by a preference judgment, or size-based choice was followed by a size-based judgment), as well as cross-domain (where preference choice was followed by size-based judgment, or *vice versa*).

6.3.1 Methods

6.3.1.1 Participants

Experiment 1 involved 23 participants recruited from Cardiff University School of Psychology participant panel for lab-based tasks. Consent was obtained from all participants. The study was approved by the Cardiff University School of Psychology Research Ethics Committee. Participants were compensated with course credits.

6.3.1.2 Stimuli

We selected 18 food pictures from the *Food-Pics* database (Blechert et al. 2019). The food items were chosen to be diverse and equally represent different categories. These categories contained: sweet snacks, savory dishes, and healthy foods. Each item was presented on a squared white background (350×350 pixels). Each food category contained two relatively small (<35% of non-white pixels), two medium (36-45% of non-white pixels) and two relatively large items (>46% of non-white pixels).

During the main task, two items were presented on the opposite sides of the screen with a symbol indicating *response type* (choice or judgment) placed centrally above the items; *domain* (size, preference, no-choice or forced choice) centrally in the gap between the items and a scale below (judgments only). The judgment scales were horizontal and had the width spanning the width of the 2 items. Only the two ends and the middle of the scales had markers. No labels were present.

6.3.1.3 Procedure

Participants completed two behavioural sessions conducted on different days, each session taking between 75 and 90 minutes. The experiment was written and conducted in *PsychoPy* v3.1.2 (Peirce 2007).

Stage 1: Initial rating. Participants performed initial preference and size-based ratings of 18 food items on a continuous scale from 1 to 100. Each participant completed a total of 72 rating trials, with two iterations of both types of ratings. Rating type order was counterbalanced across participants

and item order was randomized within each rating. Participants used *z* and *m* keys to move the value indicator to the left or right. There was no time limit for response.

Stage 2: Main Task. Each trial of the main task consisted of a choice between 2 items followed by a judgment of difference between item values on a continuous scale from -100 to 100. The scale was horizontal, and -100 was maximally favouring the left item, and 100 was maximally in favour of the right item.

The experiment consisted of three types of choices (preference, size, or no-choice) and two types of judgments (preference or size). No-choice condition required participants to withhold a response while observing the items on screen.

The task was divided into 9 51-trial blocks per session. Participants took short self-paced breaks between blocks. Participants had a time limit of 2250 ms to choose and 6000 ms to make a judgment. Similar to the rating phase, participants used either *z* and *m* keys to choose and move left and right along the scale. A choice was indicated with a green border surrounding the chosen item for the length of the choice phase. The choice screen remained visible for the full length (2250 ms), irrespective of reaction time speed, followed immediately by a judgment screen. The trial was completed after confirming the judgment by pressing the space bar. If no response was provided in time, a prompt saying *"Too slow"* was presented for 500 ms, after which the next trial was presented.

Stage 3: Final rating. After completing Stage 2, the participants performed the preference and perceptual-based ratings for the third time for all 18 items

once.

6.3.2 Results

6.3.2.1 Consistency and Performance Quality

We assume that the initial ratings are a consistent and unbiased measure of individual value estimation. To verify this, we tested *a)* the consistency between a set of objective and task-derived measures and the initial ratings, and *b)* how consistent task performance was in relation to the initial estimations.

We correlated judgments with objective item sizes using Spearman's rank correlation coefficient (r_S) per session. Size estimations were reflective of actual item size rank order M = 0.85, SD = 0.14, t(44) = 38.25 p < 0.001. High accuracy in item size ranking indicates good understanding of the task requirements.

Pre-task and post-task ratings were significantly correlated with one another M_{pref} = 0.81, SD_{pref} = 0.14 t(44) = 45.49 p < 0.001, M_{size} = 0.78, SD_{size} = 0.19 t(44) = 28.74 p < 0.001.

Choice consistency (calculated *per* session) was significantly higher than chance level (50%) in both domains M_{pref} = 77.7%, SD_{pref} = 6.9% t(44) = 44.52 p < 0.001, M_{size} = 79.7% SD_{size} = 6.0% t(44) = 52.91 p < 0.001,

Judgment consistency was also significantly above chance M_{pref} = 72.2%, SD_{pref} = 6.0% t(44) = 56.11 p < 0.001, M_{size} = 73.1%, SD_{size} = 6.4% t(44) = 59.74 p < 0.001.

Overall, high consistency measures indicate that participants understood

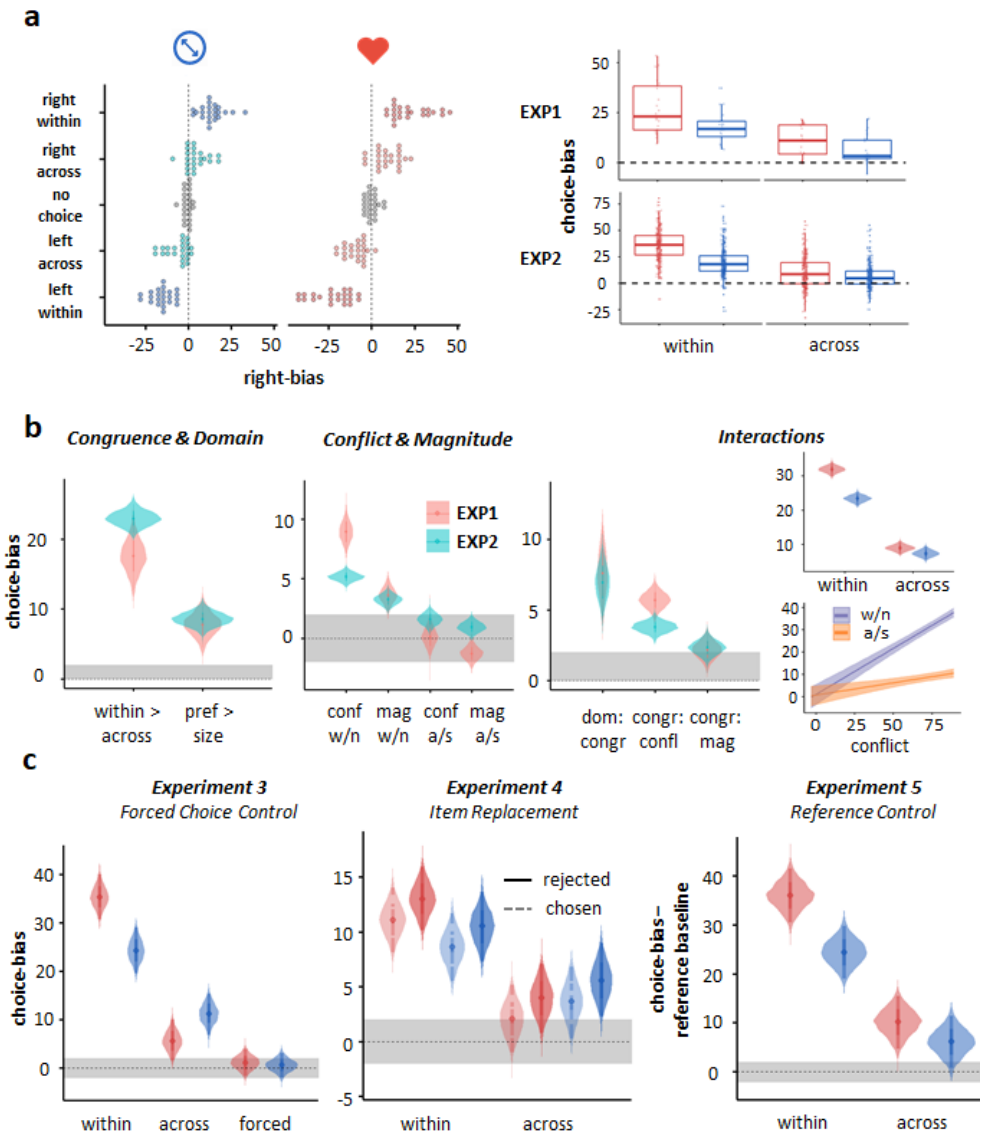

Figure 6.2: . Behavioural Results. a) Visualization of the main findings: right-bias as a function of choice and trial type in Experiment 1 (left), choice induced bias across trial types (error bars represent standard errors) in Experiments 1 (upper right) and 2 (lower right). Red and blue represent preference and size domains respectively. b) Posterior group parameter estimations for main predictors across Experiments 1 & 2. Gray area represents Region of Practical Equivalence (ROPE). c) Posterior group parameter estimations for crucial predictors in Experiments 3-5. Gray area represents ROPE, red and blue represent preference and size domains respectively.

the task and their initial value estimations were relatively stable across the length of the experiment.

6.3.2.2 Across and Within-Domain CIB

We found overwhelming evidence for the existence of choice-induced bias (CIB) $med = 26.63$, 95% CI [22.01, 31.72], pd > 0.999, 0% in ROPE. The main effect was larger for within-domain compared to across-domain conditions $med = 17.63$, 95% CI [12.47, 21.98], pd > 0.999, 0% in ROPE, and for preference judgments $med = 7.78$, 95% CI [4.42, 10.82], pd > 0.999, 0% in ROPE.

Bias was also significantly driven by congruent judgment difficulty (absolute difference in estimated value between items in the domain of choice) $med = 8.94$, 95% CI [6.96, 10.69] pd > 0.999, 0% in ROPE, and the congruent judgment magnitude (sum of the item values) $med = 3.59$, 95% CI [2.13, 5.01], pd > 0.999, 1.57% in ROPE.

We did not find sufficient evidence for the influence of the difficulty or magnitude of the incongruent domain (that is, domain inconsistent with the current choice) $med_{\text{Difficulty-across}} = 0.10$, 95% CI [-1.37, 1.71], pd = 0.56, $med_{\text{Magnitude-across}} = -1.30$, 95% CI [-2.42, 0.22], pd = 91.9. The effect of domain was only present in the within-domain conditions $med_{\texttt{congruency x domain}} = 7.78$, 95% CI [4.42, 10.82], pd > 0.999, 0% in ROPE.

Congruent difficulty had a stronger effect in the within-domain conditions $med_{\texttt{conguency x difficulty}} = 5.69$, 95% CI [4.56, 6.82], pd > 0.999 indicating choice conflict had a significant effect on subsequent judgment. Similarly,

congruent magnitude influenced judgment more strongly in the within-domain trials $med_{\texttt{conguency x magnitude}} = 2.16$, 95% CI [0.93, 3.32], pd > 0.999, 39.73% in ROPE.

6.4 Experiment 2

The goals of Experiment 2 were to replicate our initial findings on a larger and more diverse online sample and test individual differences related to effect sizes. The task used was a shortened version of the one used in Experiment 1, adapted for an online setting. To make sure our initial findings are not an artifact associated with a specific item set, we selected a new, non-overlapping set of food items. The large sample size also allowed for reliable generative modelling of the effects.

Additionally, we aimed to test a set of hypotheses regarding individual differences:

Confirmation Bias

Since choice in an unrelated domain should provide no judgment information, cross-domain CIB can be thought of as a pure measure of confirmation bias. We hypothesized that cross-domain CIB would be associated with individual differences in confirmation bias, measured with the Confirmation Bias Inventory (Rassin 2008) as well as Extroversion scale (Matz, Hofstedt, and Wood 2008), since the construct of extroversion was previously shown to be positively related to confirmation bias.

Consistency

We hypothesized that consistency, as measured by consistency measures in the task, could be associated with individual differences in Preference for Consistency (Cialdini 2014).

Extreme-Value Aversion

We hypothesized that aversion to strong commitment, one of the core features of individual differences in Action Control (Kuhl and Beckmann 1994) could be reflected in our task by a more conservative usage of the judgment scale during the main task - namely preferring values close to the middle of the scale, not favouring strongly any of the alternatives (extreme-value aversion).

The study was preregistered (https://osf.io/v5pdb).

6.4.1 Methods

6.4.1.1 Participants

Experiment involved 250 subjects (M_{age}=26.31 years SD_{age}=8.29; 93 female) from a participant recruitment portal *Prolific* (https://prolific.co). Consent was obtained from all participants. The study was approved by the Cardiff University School of Psychology Research Ethics Committee. Participants received cash payments for their participation based on an hourly rate (7£ per hour).

Power analyses, exclusion criteria, experiment procedures and further analysis plans were preregistered (see: *Online Resources*). A sample size of N=250 provides 80% power to detect a moderate effect of correlations between behavioural performance and trait measures (Pearson's R=0.2, $\alpha = 0.01$;

Gignac and Szodorai 2016).

6.4.1.2 Stimuli

A new set 24 food pictures from the *Food-Pics* database (Blechert et al. 2019) was selected (for images see: *Appendix*). The food items were chosen to be diverse and equally represent different categories. These categories contained: sweet snacks, savory dishes, fruits and vegetables, making up a total of 4 categories. Each item was presented on a squared white background (350×350 pixels). The size of each stimulus was defined as the proportion of the area taken by the actual food picture (*i.e.* non-white pixels). Each food category contained two relatively small (<35% of non-white pixels), two medium (36-45% of non-white pixels) and two relatively large items (>46% of non-white pixels).

During the choice and judgment phase, two items were presented on the opposite sides of the screen with a symbol indicating *response type* (choice or judgment) placed centrally above the items; *domain* (size, preference, no-choice or forced choice) centrally in the gap between the items and a scale below (judgments only). The judgment scales were horizontal and had the width spanning the width of the 2 items Only two ends and the middle of the scales had markers. No labels were present.

6.4.1.3 Procedure

The Experiment followed a similar 3-step procedure as Experiment 1 (Figure 6.1) with addition of a final part, in which participants filled in a set of

personality questionnaires. Experiment took the participants between 45-75 minutes to complete. The experiment was written and conducted in *jspsych* v6.0.5 (Leeuw 2014).

Stage 1: Initial rating. Participants performed initial preference and size-based ratings of 24 food items on a continuous scale from 1 to 100. For preference-based ratings, participants rated item desirability from strong dislike (1) to strong liking (100). In the perceptual-based rating, participants rated the size of each food picture with respect to the white background, in terms of the proportion of non-white pixels. Each participant completed a total of 96 rating trials, with two iterations of both types of ratings. Rating type order was counterbalanced across participants and item order was randomized within each rating. Unlike Experiment 1, only the top 12 items on the preference scale were used in the main task.

On each rating trial, a food stimulus was presented in the center of the screen. A task cue (red heart for preference or a blue circle with outward arrows for size; see Figure 6.1) was presented on top of the stimulus to indicate the type of rating. A white rating scale was presented below the food stimulus with markers placed on both ends. A small circle on the rating scale indicates the current rating, with its default position in the middle of the scale. Participants used a mouse to drag or click on the scale. There was no time limit for response.

Stage 2: Main Task. The main structure of the task was identical to Experiment 1, with differences related to the number of trials and interface, described below. The Task was divided into three 66-trial blocks. Participants took short self-paced breaks between blocks. Participants had a time

limit of 4000 ms to choose and or 5500 ms to make a judgment. Similar to the rating phase, participants used mouse clicks to choose and move left and right along the scale. A choice was indicated with a green border surrounding the chosen item for the length of the choice phase. The choice screen disappeared 500 ms after making a choice. The trial was completed after confirming the judgment with clicking the *"confirm"* button (Experiments 2-5). If no response was provided in time, a prompt saying *"Too slow"* was presented for 500 ms, after which the next trial was presented.

Stage 3: Final rating. After completing Stage 2, participants performed the preference and perceptual-based ratings for the third time, involving all 24 items.

Questionnaires. After the final rating stage, participants filled five questionnaires measuring Preference for Consistency (Cialdini 2014), Extraversion (BFI-2-S; Soto and John 2017), Susceptibility to Confirmation Bias (CI; Rassin 2008), Positive and Negative Affect (I-PANAS-SF; Thompson 2007) and Action Control (ACS-90; Kuhl and Beckmann 1994). Questionnaire order was randomized across participants.

6.4.1.4 Computational Modelling

Our model aims at explaining peoples' choices and judgments by estimating the internal values associated with item preference and size from which the data is generated, and how they change in time. The model assumes that values fluctuate as a function of choice, so that choosing or rejecting an item updates its estimation by a fraction of its value (proportional update). An item value at any point is not allowed to exceed the boundaries of the scale

(0-100). The starting preference and size values of each item are derived from the mean of the two initial ratings on the appropriate scales. We model 4 distinct phases on each trial: *choice, value update, judgment,* and *value attenuation*.

Choice. The choice between 2 items is modelled using the *softmax* function, which converts item values into choice probabilities:

$$p(A) = \frac{e^{\tau} V_{d,t}(A)}{e^{\tau} V_{d,t}(A) + e^{\tau} V_{d,t}(B)} \tag{6.5}$$

where $V_{d,t}(A)$ and $V_{d,t}(B)$ represent value estimates for items A and B in a given domain d (preference or size) on trial t, and τ represents inverse temperature parameter reflecting how deterministic are the choices, given the estimations of subjective values (Ahn, Haines, and Zhang 2017). This parameter is often interpreted as the noisiness of choice.

Update. After each choice, both chosen I_{ch} and rejected I_{rej} item values are updated:

$$V_{d,t+1}(I_{ch}) = V_{d,t}(I_{ch}) + \alpha V_{d,t}(I_{ch}) \tag{6.6}$$

$$V_{d,t+1}(I_{rej}) = V_{d,t}(I_{rej}) + \beta V_{d,t}(I_{rej}) \tag{6.7}$$

where α and β represent proportional updates for the chosen and rejected items respectively. Both update parameters can take values between -1 and 1. α and β can vary dependent on choice domain (size, preference) and congruency (whether choice-congruent or choice-incongruent domain is updated). Together, this gives a maximum 8 update parameters.

Judgment. Judgment is assumed to follow a normal distribution, centered at the difference of the true subjective value between the right and left items:

$$J_t \sim N(V_{r,d,t+1} - V_{l,d,t+1} + CB, \sigma_i) \tag{6.8}$$

where $V_{r,d,t+1}$ and $V_{l,d,t+1}$ represent the current (updated) values of the right and left items in domain d on trial t, CB represents a general consistency bias term, and σ_i represents individual-level variability. The valence of CB is dependent on choice, so that it takes positive values if right item was chosen, negative values if left was chosen, and 0 when no choice was made. The parameter represents a consistency-driven adjustment that does not affect the underlying value estimations.

Attenuation. Finally, item values are attenuated towards their initial estimates:

$$V_{att,d,t+1}(A) = V_{d,t+1}(A) - k(V_{d,t+1}(A) - V_{d,init}(A)) \tag{6.9}$$

$$V_{att,d,t+1}(B) = V_{d,t+1}(B) - k(V_{d,t+1}(B) - V_{d,init}(B)) \tag{6.10}$$

where V_{att} represents the attenuated value of the item, and k is the discounting parameter and can take values between 0 and 1. k represents the strength of the decay effect, where the current item value *V_{d,t+1}* is discounted proportionally to it's distance to the initial value *V_{d,init}*. One can imagine this mechanism akin to a rubber band - the greater the deviation of the current updated item value from the initial one, the stronger the pull back towards the initial value becomes.

This accounts for the possibility that choice-driven value updates might be

temporary and decay when earlier choices get forgotten. A k of 1 indicates perfect discounting, and the model is reduced to a fixed-value model where values are never updated. A k of 0 indicates no discounting (permanent update). Values of k in-between suggest a decay effect, where choices made longer in the past have less influence on current value.

6.4.1.5 Model fitting procedure

We fitted the model using Bayesian Hierarchical Modelling, estimating group level and subject-level parameters. Following previous work in this domain (Ahn, Haines, and Zhang 2017; L. Zhang and Gläscher 2020) we used uniform priors on a realistic constrained range:

$$\alpha \sim U(0, 0.5) \tag{6.11}$$

$$\beta \sim U(0, 0.5) \tag{6.12}$$

$$\tau \sim U(0, 10) \tag{6.13}$$

$$k \sim U(0, 1) \tag{6.14}$$

We used *Stan* programming language (Carpenter et al. 2017) for the hierarchical implementation of our model. For each model, we generated 2 independent chains of 2000 samples from the joint posterior distributions of the model parameters, using Hamiltonian Monte Carlo (HMC) sampling (Carpenter et al. 2017). The initial 1000 samples were discarded as burn-in. To assess model convergence of the chains we calculated the Gelman-Rubin convergence diagnostic $\widehat{R}$ for each model parameter. Similar to previous work

(Annis, Miller, and Palmeri 2017; Zajkowski et al. 2020) we used $\widehat{R} < 1.1$ as a criterion for good convergence.

6.4.1.6 Model comparison

We compare 4 theoretically-driven models fitted to data from Experiment 2. *Null* model fixes all updating parameters to 0 and k to 1, resulting in no value updating or consistency bias. *Consistency-Only* model assumes only a general consistency bias and no value-updating. Congruent Update model assumes only choice-congruent values are updated. *Update-Only* model assumes all choices induce a value update, but no consistency bias. Full Update model allows all update parameters to vary, together with a general consistency bias. Direct model fit was measured using Leave-One-Out (LOO) information criterion (Vehtari, Gelman, and Gabry 2017). LOO evaluates the model fit while controlling for model complexity, with lower value indicating better out-of-sample prediction. We use LOO scores to determine model stacking weights, which give the predictive probability of each model being best (Yao et al. 2017). When comparing 2 models, a probability of .9 or larger in favour of one of them is considered decisive (L. Zhang and Gläscher 2020); otherwise a simpler model is preferred.

6.4.1.7 Model evaluation

To evaluate the winning model in terms of prediction accuracy we *1)* generate posterior predictive distribution and correlate model predictions with participant choices across subjects, and *2)* regress the updated item values

onto final rating values and compare whether the model-derived values can predict final ratings better than the initial ratings.

6.4.2 Results

6.4.2.1 Consistency and Performance Quality

Similarly to Experiment 1, performance and consistency were significantly above chance (calculated per participant), indicating that participants understood the task and their initial value estimations were relatively stable across the length of the experiment.

Size estimations were reflective of actual item size rank order M = 0.66, SD = 0.19 t(249) = 43.31, p<0.001 (one-way against chance level of 0).

Pre-task and post-task ratings were significantly correlated with one another $M_{pref} = 0.85$, $SD_{pref} = 0.13$ t(249) = 106.58 $p < 0.001$, $M_{size} = 0.73$, $SD_{size} = 0.19$ t(249) = 57.74 $p < 0.001$ (one-way against chance level of 0).

Choice consistencies were significantly higher than chance level $M_{pref} = 71.3\%$, $SD_{pref} = 11.0\%$ t(249) = 103.05 $p < 0.001$, $M_{size} = 76.7\%$, $SD_{size} = 10.0\%$ t(249) = 121.79 $p < 0.001$ (one-way against chance level of 0).

Judgment consistency was also significantly above chance $M_{pref} = 75.4\%$, $SD_{pref} = 11.0\%$ t(149) = 157.47 $p < 0.001$, $M_{size} = 74.4\%$, $SD_{size} = 11.0\%$ t(149) = 149.71 $p < 0.001$ (one-way against chance level of 0).

6.4.2.2 Across and Within-Domain CIB

Experiment 2 replicated all main behavioural findings regarding choice-induced bias from Experiment 1 (Figure 6.2). Overall, the CIB effect was strong med = 31.90, 95% CI [30.03, 33.80], pd > 0.999, 0% in ROPE. The main effect was larger for within-domain compared to across-domain conditions med = 22.95, 95% CI [20.89, 25.07], pd > 0.999, 0% in ROPE, and for preference judgments med = 8.54, 95% CI [6.77, 10.55], pd > 0.999, 0% in ROPE. Bias was also significantly driven by congruent judgment difficulty (absolute difference in estimated value between items) med = 5.16, 95% CI [4.52, 5.87] pd > 0.999, 0% in ROPE. and the congruent judgment magnitude (sum of the item values) med = 3.24, 95% CI [0.18, 1.77], pd > 0.999, 0.20% in ROPE.

We did not find sufficient evidence for the influence of the difficulty or magnitude of the incongruent domain $med_{\text{Difficulty-across}}$ = 1.58, 95% CI [0.60, 2.46], pd = 80.0, 100% in ROPE $med_{\text{Magnitude-across}}$ = 0.93, 95% CI [0.18, 1.77], pd = 98.5, 99.38% in ROPE.

The effect of domain was only present in the within-domain conditions $med_{\texttt{congruency x domain}}$ = 8.54, 95% CI [6.77, 10.55], pd > 0.999, 0% in ROPE. Congruent difficulty had a stronger effect in the within-domain conditions $med_{\texttt{conguency x difficulty}}$ = 6.94, 95% CI [4.49, 9.22], pd > 0.999, 0% in ROPE indicating choice conflict had a significant effect on subsequent judgment. Similarly, congruent magnitude influenced judgment more strongly in the within-domain trials $med_{\texttt{conguency x magnitude}}$ = 2.32, 95% CI [1.38, 3.29], pd > 0.999, 26.07% in ROPE.

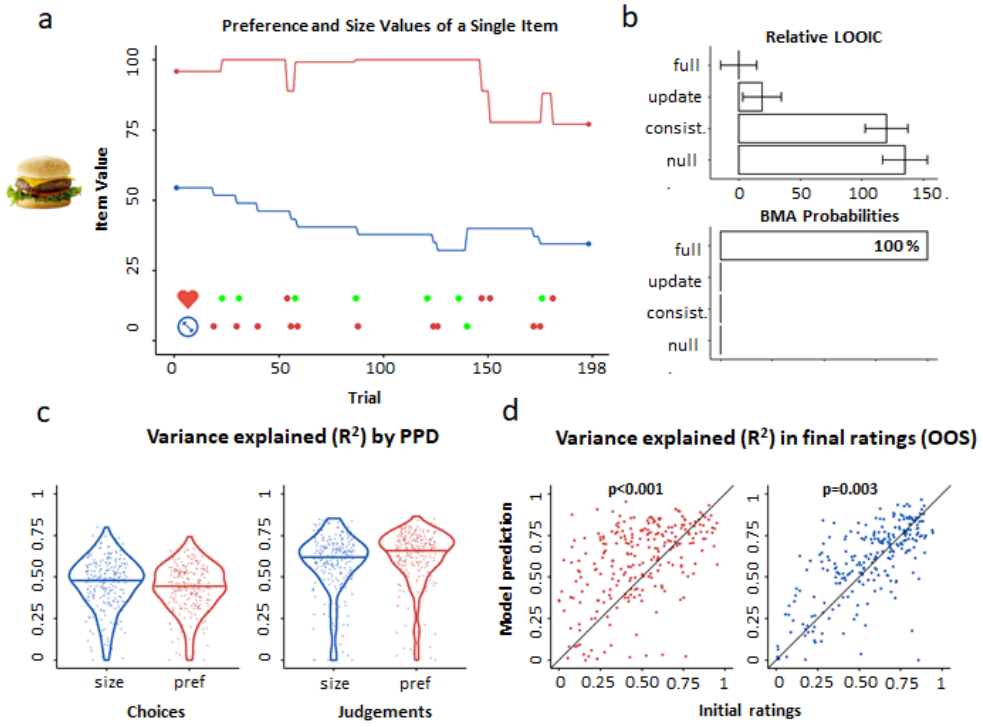

Figure 6.3: Generative model fitting. a) An instance of dynamic value updating for a single item. The internal value fluctuates as a function of choices. Green points show trials where item was chosen; red when it was rejected. b) Model comparison. Upper panel shows relative LOOIC scores, error bars represent SE. Lower panel shows the probabilities of each model providing the best fit based on Bayesian Model Averaging of LOOIC scores (Yao et al., 2018). c) Posterior Predictive model validation for choices (left) and judgments (right) in terms of R2. Each dot represents a single participant. d) Out-of-sample (OOS) model-derived prediction of final ratings (Y-axis) plotted against prediction derived from initial ratings for preference (left) and size (right). P-values indicate that model-based prediction was significantly better.

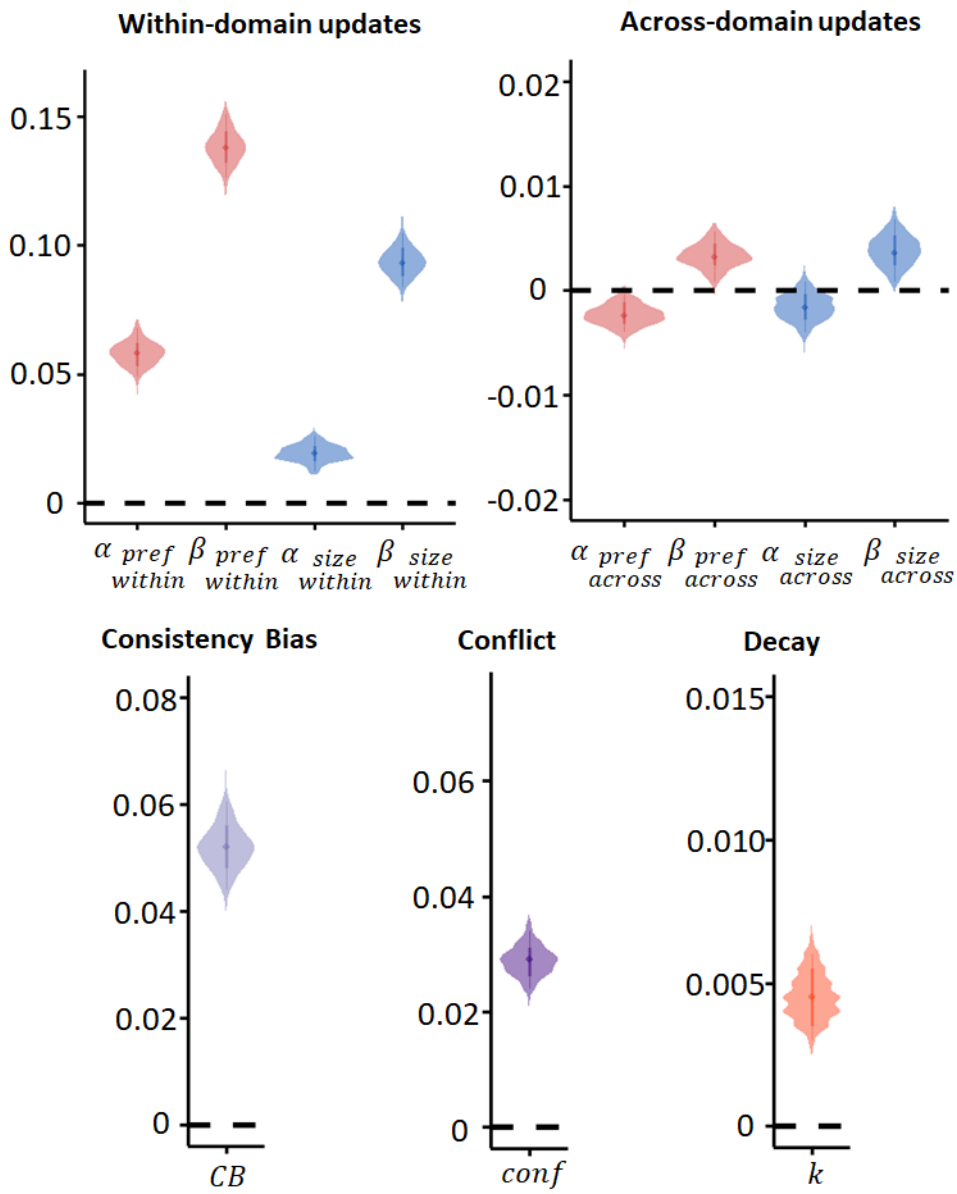

Figure 6.4: Parameters of the best fitting model. α and β represent the positive and negative updates, respecitvely.

6.4.2.3 Generative Modelling

To understand the generative processes giving rise to the choice-induced bias, we built a hierarchical Bayesian model introducing two potential sources of CIB: a domain-general consistency bias and a domain-specific value update (see: *Methods*), and fitted it to the data from Experiment 2. We tested 4 models with varying assumptions: *Null* (stable values; no sources of CIB present in data generating process), *Consistency-Only*, *Update-Only* and *Full* (both sources of CIB present). Model comparison using Leave-One-Out Information Criterion (Vehtari, Gelman, and Gabry 2017) indicated that the full model fitted the data best LOO = 269.9, SE = 29.4, $P_{best} > 0.999$ based on Bayesian model averaging (Yao et al. 2017; Figure 6.3b). Strong correlations (per participant) between model-derived final value predictions and post-task ratings M = 0.750, SD = 0.226 for preference and M = 0.739 SD = 0.228 for size indicated the model can predict the final ratings well. Compared to the initial ratings, model-derived values predicted the final ratings more accurately: $M_{DeltaR^2} = 0.11$, t(243) = 7.80, $p < 0.001$ for preference and $M_{DeltaR^2} = 0.03$, t(243) = 2.878, p = 0.002 for size (Figure 6.3c).

Consistency Bias. The winning model displayed a consistency effect *med*= 0.065 95% CI [0.055, 0.076], pd > 0.999, which was modulated by decision conflict *med*= 0.029 95% CI [0.023 ,0.033], pd > 0.999.

Value Update. Within-domain update parameter values were all greater than 0 with pd > 0.999: positive preference-within update *med* = 0.058, 95% CI [0.049, 0.067]; negative preference-within update *med* = 0.138, 95% CI

[0.126, 0.152]; positive size-within update *med* = 0.019, 95% CI [0.013, 0.026]; negative size-within update *med* = 0.093, 95% CI [0.084, 0.105];

Across-domain parameter updates were all smaller than 0.01 (1% of item value), suggesting their nature is inconsequential for the observed bias: positive preference-across update *med* = -0.004, 95% CI [-0.010, 0.002], pd = 0.09; negative preference-across update *med* = 0.001, 95% CI [0.001, 0.015], pd = 0.99; positive size-across update *med* = -0.006, 95% CI [-0.011, -0.001], pd = 0.99; negative size-across update *med* = 0.006, 95% CI [0.000, 0.011], pd = 0.97;

Value Discounting. Parameter k values were close to 0 *med* = 0.009, 95% CI [0.006, 0.012] indicating only marginal decay.

6.4.2.4 Individual Differences

In Experiment 2 we also tested a set of hypotheses correlating individual differences with behavioural measures. Choice-induced bias in the *across-domain* conditions was not significantly correlated with scores on Confirmation Bias Inventory (Rassin 2008) r = 0.09, p = 0.14; or Extraversion (Matz, Hofstedt, and Wood 2008) r = 0.11, p = 0.09. Decision-Judgment consistency was not significantly correlated with Preference for Consistency (Cialdini 2014) r = -0.02, p = 0.75; and middle-scale aversion (quantified as the mean absolute value of the judgments) was not correlated with Action Control (Kuhl and Beckmann 1994).

6.4.2.5 Confound analysis

An alternative interpretation of CIB is that it can arise spontaneously due to a regression to the mean, effectively revealing true preferences and beliefs, instead of shaping them (Chen and Risen 2010; Izuma and Murayama 2013). To account for this effect, we performed an extensive set of simulations, varying parameters related to the ratings, choices and judgments (*Appendix*), which revealed that given a reasonable set of assumptions a large bias is extremely unlikely to arise due to this confound. In all but the most extreme parameter settings within-domain CIB was not larger than 1 point (0.5% of the judgment scale), while the across-domain effect was centered at 0 (see: *Simulation* section in the *Appendix*).

6.5 Experiment 3

The main goal of Experiment 3 was testing whether the CIB effects found in previous iterations of the study can be induced by a exogenously dictated (forced) choice. Our prediction is that if the effect exists, it would be weaker and would not be as strongly influenced by previous choice history, compared to when the choice is driven endogenously. In order to test this, the *no-choice* condition from Experiments 1 and 2 was replaced by a *forced choice* condition, where the choice item was dictated randomly by the computer. Voluntariness of choice have been postulated to be necessary for a CIB to occur in preference-based tasks (Egan, Bloom, and Santos 2010; Sharot, Velasquez, and Dolan 2010), it has however not been systematically compared across domains. The experiment was preregistered (https://osf.io/tkxg5).

6.5.1 Methods

6.5.1.1 Participants

Experiment involved 50 subjects (M_{age}=25.95 years SD_{age}=5.95; 14 female) from a participant recruitment portal *Prolific* (https://prolific.co). Consent was obtained from all participants. The study was approved by the Cardiff University School of Psychology Research Ethics Committee Participants received cash payments for their participation based on an hourly rate (7£ per hour).

6.5.1.2 Stimuli

The same set of 24 stimuli from the *Food-Pics* database (Blechert et al. 2019) as in Experiment 2 were used (for images see: *Appendix*).

6.5.1.3 Procedure

The Experiment followed a similar 3-step procedure as previous experiments (Figure 6.1) with modified Main Task, where the *no-choice* condition was replaced with *forced choice*. Experiment took the participants between 45-75 minutes to complete. The experiment was written and conducted in *jspsych* v6.0.5 (Leeuw 2014).

Stage 1: Initial rating. Initial stage was identical to Experiment 2.

Stage 2: Main Task. The main structure of the task was similar to Experiment 2, with differences related condition specifications and trial number, described below (see Table 6.2). The Task was divided into four 66-trial blocks.

The *forced-choice* condition required participants to choose one randomly predetermined item in the choice phase. A red arrow enclosed within a circle replaced the domain symbol, and the direction of the arrow indicated the item to be chosen (Figure 6.1). The judgment part of the trial remained identical to previous experiments. The number of forced and free choice conditions was balanced, each containing 132 trials. Similarily, the number od force-left and force-right choices was balanced equally (66 trials each).

Stage 3: Final rating. Stage was identical to Experiment 2.

6.5.2 Results

6.5.2.1 Consistency and Performance Quality

Similarly to previous experiments, performance and consistency (calculated per participant) were significantly above chance, indicating that participants understood the task and their initial value estimations were relatively stable across the length of the experiment.

Size estimations were reflective of actual item size rank order $M = 0.69$, $SD = 0.16$ $t(49) = 25.14$, $p < 0.001$.

Pre-task and post-task ratings were significantly correlated with one another $M_{pref} = 0.89$, $SD_{pref} = 0.10$ $t(49) = 62.19$ $p < 0.001$, $M_{size} = 0.67$, $SD_{size} = 0.16$ $t(49) = 21.87$ $p < 0.001$.

Choice consistencies were significantly higher than chance level

$M_{pref} = 73.0\%$, $SD_{pref} = 8.7\%$ $t(49) = 59.25$ $p < 0.001$ $M_{size} = 76.4\%$, $SD_{size} = 9.3\%$ $t(49) = 57.78$ $p < 0.001$.

Judgment consistency was also significantly above chance $M_{pref} = 76.1\%$, $SD_{pref} = 9.7\%$ $t(49) = 99.65$ $p < 0.001$ $M_{size} = 78.4\%$, $SD_{size} = 9.7\%$ $t(49) = 104.09$ $p < 0.001$.

6.5.2.2 The Need for Voluntary Choices

To test the effect of voluntary choice on CIB, Experiment 3 introduced forced-choice control condition (Figure 6.1c) in which participants were instructed to choose a specified item. In contrast to voluntary decision conditions which replicated findings from our previous experiments (Figure 6.2c), the bias following forced choices was not significantly different from 0: $med = 1.08$ 95% CI [-1.29, 3.43], $pd = 0.77$ suggesting that a voluntary choice is necessary for the bias to occur.

6.6 Experiment 4

The goal of Experiment 4 was to test to what extent is the CIB effect driven by increasing the value of the chosen alternative in comparison to decreasing the value of the rejected one. In order to test this, we modified the experimental procedure so that one of the items (chosen or rejected) was randomly replaced within trial, after the choice, but before the judgment. This allowed us to compare to what extent is the CIB effect driven by overvaluing the chosen item compared to undervaluing the rejected one. Previous research in this area was done only in preference literature, and provided inconsistent results, from undervaluation effects being stronger (Brehm 1956; Izuma et al. 2010), both being roughly equal (Harmon-Jones et al. 2008; Jarcho, Berkman,

and Lieberman 2011; Luo and Yu 2016) to overvaluation being stronger (Sharot, Velasquez, and Dolan 2010). The experiment was preregistered (https://osf.io/qhmg2).

6.6.1 Methods

6.6.1.1 Participants

Experiment involved 50 subjects (M_{age}=25.95 years SD_{age}=7.05; 15 female) from a participant recruitment portal *Prolific* (https://prolific.co). Consent was obtained from all the participants. The study was approved by the Cardiff University School of Psychology Research Ethics Committee Participants received cash payments for their participation based on an hourly rate (7£ per hour).

6.6.1.2 Stimuli

The same set of 24 stimuli from the *Food-Pics* database (Blechert et al. 2019) as in Experiments 2, 3 and 4 were used (for images see: *Appendix*).

6.6.1.3 Procedure

The Experiment followed a similar 3-step procedure as previous experiments (Figure 6.1) modified Main Task, where the one of the items was randomly replaced during the judgment phase of the trial. Experiment took the participants between 45-75 minutes to complete. The experiment was written and conducted in *jspsych* v6.0.5 (Leeuw 2014).

Stage 1: Initial rating. Initial stage was identical to Experiments 2 and 3.

Stage 2: Main Task. The main structure of the task was similar to Experiment 2, with differences related the judgment phase and trial number, described below (see Table 6.2). The Task was divided into four 66-trial blocks. One of the items was replaced by a random item right after choice. The number of chosen and rejected items that were replaced was counterbalanced (66 trials each).

Stage 3: Final rating. Stage was identical to Experiments 2 and 3.

6.6.2 Results

6.6.2.1 Consistency and Performance Quality

Similarly to previous experiments, performance and consistency (calculated per participant) were significantly above chance,indicating that participants understood the task and their initial value estimations were relatively stable across the length of the experiment.

Size estimations were reflective of actual item size rank order $M = 0.64$, $SD = 0.22$ $t(49) = 19.49$, $p < 0.001$.

Pre-task and post-task ratings were significantly correlated with one another $M_{pref} = 0.85$, $SD_{pref} = 0.20$ $t(49) = 34.11$, $p < 0.001$, $M_{size} = 0.73$, $SD_{size} = 0.16$ $t(49) = 40.07$, $p < 0.001$.

Choice consistencies were significantly higher than chance level $M_{pref} = 71.5\%$, $SD_{pref} = 9.3\%$ $t(49) = 54.30$, $p < 0.001$, $M_{size} = 75.1\%$, $SD_{size} = 12.4\%$ $t(49) = 42.71$, $p < 0.001$.

Judgment consistency was also significantly above chance $M_{pref} = 72.7\%$, $SD_{pref} = 9.5\%$ $t(49) = 70.94$, $p < 0.001$, $M_{size} = 76.7\%$, $SD_{size} = 12.2\%$ $t(49) = 78.94$, $p < 0.001$.

6.6.2.2 Overestimation vs Underestimation

In Experiment 4, we tested to what extent is the bias driven by overvaluation the chosen option, as compared to undervaluating the rejected one. For this purpose, one of items (either the chosen or rejected) was replaced after each choice (Figure 6.1c). We found a modest yet significant effect of item replacement, $med_{\text{rejected-chosen}} = 1.89$, CI = [0.09, 3.75], pd = 0.98, 53.85% in ROPE (Figure 6.2c), indicating that the undervaluation of the rejected item being stronger than overvaluation of the chosen one.

6.7 Experiment 5

In Experiment 5 we manipulate which item should be considered the 'default' (so called *reference* item) to which the other is compared during the judgment phase. This manipulation accounts for one possible source of the bias - a bias towards the status quo (Kahneman, Knetsch, and Thaler 1991). If the chosen item is considered the reference, CIB could be explained by information sampling biased towards positive evidence (Hunt et al. 2016). In this view, participants use their choice as an implicit reference, which drives the bias. By explicitly manipulating the reference item during the judgment phase we can dissociate between the effects of choice and reference.

6.7.1 Methods

6.7.1.1 Participants

Experiment involved 50 subjects (M_{age}=27.96 years SD_{age}=9.29; 18 female) from a participant recruitment portal *Prolific* (https://prolific.co). Consent was obtained from all participants. The study was approved by the Cardiff University School of Psychology Research Ethics Committee Participants received cash payments for their participation based on an hourly rate (7£ per hour).

6.7.1.2 Stimuli

The same set of 24 stimuli from the *Food-Pics* database (Blechert et al. 2019) as in Experiments 2, 3 and 4 were used (for images see: *Appendix*).

In contrast to previous Experiments, the judgment scale was vertical and placed in the middle of the screen between the two items (Figure 6.1). This was done in order to decouple the items (displayed on the right and left) and the scale (going up and down) and enabling manipulating which item is associated with which end of the scale.

6.7.1.3 Procedure

The Experiment followed a similar 3-step procedure as previous experiments (Figure 6.1) with modified judgment phase of Main Task, where one of the items was designated as a *reference item*, to which the other item was compared. Experiment took the participants between 45-75 minutes to

complete. The experiment was written and conducted in *jspsych* v6.0.5 (Leeuw 2014).

Stage 1: Initial rating. Initial stage was identical to Experiments 2, 3 and 4.

Stage 2: Main Task. The main structure of the task was similar to Experiments 2 ,3 and 4, with differences related the judgment phase and trial number, described below (see Table 6.2). The Task was divided into four 66-trial blocks. The reference manipulation was done such that after each choice one of the items was randomly assigned as the *reference item*. This was signified by an orange framing around the item. The judgment task was reframed to evaluate whether the reference item is larger (perceptual domain) /more valuable (preference domain), compared to the non-reference item, so that positive evidence always favored the reference item. The number of times the left and right items were assigned as reference was balanced (66 trials each).

Stage 3: Final rating. Stage was identical to Experiments 2, 3 and 4.

6.7.1.4 Consistency and Performance Quality

Similarly to previous experiments, performance and consistency (calculated per participant) were significantly above chance,indicating that participants understood the task and their initial value estimations were relatively stable across the length of the experiment.

Size estimations were reflective of actual item size rank order $M = 0.63$, $SD = 0.19$ $t(49) = 22.86$, $p < 0.001$.

Pre-task and post-task ratings were significantly correlated with one another

$M_{pref} = 0.85$, $SD_{pref} = 0.14$ $t(49) = 34.11$, $p < 0.001$ $M_{size} = 0.73$, $SD_{size} = 0.14$ $t(49) = 40.71$, $p < 0.001$.

Choice consistencies were significantly higher than chance level $M_{pref} = 73.0\%$, $SD_{pref} = 11.6\%$ $t(49) = 44.52$, $p < 0.001$. $M_{size} = 73.0\%$, $SD_{size} = 9.8\%$ $t(49) = 52.91$, $p < 0.001$.

Judgment consistency was also significantly above chance $M_{pref} = 72.6\%$, $SD_{pref} = 12.2\%$ $t(49) = 56.11$, $p < 0.001$, $M_{size} = 72.8\%$, $SD_{size} = 11.3\%$ $t(49) = 59.74$, $p < 0.001$.

6.7.1.5 Positive reference effects

To isolate the positive reference effect, we looked at reference bias in no-choice trials, finding a strong effect of positive reference on judgment *med* = 11.81, 95% CI [6.88, 17.02], pd > 0.999, 0.02% in ROPE.

Then, we analyzed whether the choice-induced effect prevailed in choice trials. No-choice condition served as a baseline for reference bias, to which all other trials types were compared to. This enabled to estimate the effect of choice independent of the frame of reference *med* = 36.00, 95% CI [30.34, 41.29], pd > 0.999, 0% in ROPE. The differences in choice-induced bias between reference-chosen and reference-rejected conditions after removing reference baseline were non-significant *med*= 4.51, 95% CI [-1.70, 11.12], pd = 0.91.

6.8 Discussion

People's judgments are biased by their choices. In a suite of 5 experiments, we show that this phenomenon is not only ubiquitous within perceptual and preference domains, but also that choices of one type can affect evaluation in an unrelated domain. Our findings bridge the literature on CIB in perception and preference, which up to this point has been studied separately.

The similarities include both being driven by choice difficulty (how close in value the items are) and magnitude (the sum of item values). Conflict-driven bias has been proposed by Festinger (1957) and shown to influence dissonance resolution after preference-based choice (Brehm 2007; Izuma et al. 2010; van Veen et al. 2009). Magnitude effects (or *value sensitivity*; Pirrone, Stafford, and Marshall 2014b) are known to influence choices (Pirrone, Azab, Hayden, Stafford, and Marshall 2018b; Teodorescu, Moran, and Usher 2016; Zajkowski et al. 2020) as well as confidence (Lebreton et al. 2018), suggesting that absolute value of choice options inflates decision noise, but reduces uncertainty. In the context of CIB, effects of magnitude have been postulated as potentially influencing the post-choice alternative spread (Festinger 1957). This claim has previously met with mixed evidence (Greenwald 1969). One explanation of the effect, as implemented in our model, is that magnitude can be incorporated naturally into the value-update mechanism but assuming proportional updates.

Our main findings demonstrate the existence of an across-domain CIB, whose strength was consistent across all experiments. The effect was similar across judgment domains, suggesting that the across-domain effect is qualitatively

distinct from domain-dependent within-domain effects. Across-domain effects are especially interesting, as they illustrate a a clear deviation from optimality: there is no apparent reason as for how such bias could be beneficial for the agent. While within-domain effects can be explained by value inference conditioned on ones choices (Bem 1967; Kruglanski et al. 2018), similar interpretation is irrelevant in the context of across-domain effects, suggesting a prioritization of consistency, even at the cost of being biased. One potential mechanism for a similar effect have been also recently proposed by Horsby and Love (2020), who suggest that all relevant dimensions of a choice option are represented in a shared, continuous space, where values are updated though coherency optimization. The model however does not tackle representations of inherently objective dimensions, such as assessment of perceptual attributes.

For our knowledge, this study is the first to directly experimentally demonstrate across domain choice-induced effects. Previous literature has demonstrated malleability of preference for choice attributes in the face of a decision, promoting choice-consistent values (Holyoak and Simon 1999; Simon, Krawczyk, and Holyoak 2004; Simon et al. 2008). Those studies found choice-coherent preference shifts in a variety of domains, such inductive reasoning, social reasoning, and resolving ambiguous situations (Simon et al. 2008).

The authors explain the effect by constraint satisfaction mechanism, which allows the system to impose a coherent interpretation on an initially ambiguous set of inputs (Holyoak and Simon 1999). Such adjustments would be considered as within-domain, as the attributes used in those studies were relevant to the subsequent choice. Crucially, the constrain-satisfaction hypothesis

could be expanded to cross-domain effects found here.

Our generative modelling supports the conclusion that CIB is driven by at least two separate mechanisms: a conflict-driven, domain-general Consistency Bias (CB) and a domain-specific Value Update (VU).

CB is quite common in perceptual literature, being associated with a range of mechanisms, such as history effects (Palminteri et al. 2017; Urai et al. 2019) or decisional inertia (Akaishi et al. 2014). In relation to preference-based choice, it can be compared to a conflict-driven dissonance reduction mechanism postulated by the dissonance reduction theory (Festinger 1957; Brehm 2007; Harmon-Jones, Harmon-Jones, and Levy 2015). In our specification, this mechanism is dependent only on the difficulty of the choice and does not affect underlying item evalution.

VU mechanism is based on the proposal that our value inference process is conditioned on the memory of previous choices (Bem 1967). The idea is that our previous choices serve as an indication of our preferences or the objective state of the world. Such choice driven inference has been implicated in many theoretical models (Friston et al. 2016; Kruglanski et al. 2018). We assume this mechanism can be described as a proportional update. Additionally, we model a memory decay process accounting for forgetting, so that choices made earlier in time can have a weaker effect on the current choice and judgment. In contrast to CB, this mechanism serves to improve one's estimation, hence it should be both domain-specific and long-lasting. In addition, unless there is an inherent correlation between dimensions, an across-domain update would by definition be irrational.

In line with this theoretical distinction, we find that within-domain effects

can be best explained as a conjunction of these two mechanisms, while across-domain effects as driven only by CB. Analysis of model parameters suggests that negative updates are larger than positive ones (consistent with Experiment 4, as well as some findings from preference-based literature: Brehm 1956; Izuma et al. 2010), and that the value-updates exhibit only a very marginal decay throughout the length of the task.

While VU has a clear adaptive interpretation, it is less clear why one should exhibit CB, or any type of across-domain bias. The constrain satisfaction hypothesis (Simon, Krawczyk, and Holyoak 2004) could provide a common mechanistic explanation for both types of adjustments, however it does not clearly explain why such mechanism would be beneficial. One prominent theory suggests that it might serve to facilitate action implementation (Harmon-Jones, Harmon-Jones, and Levy 2015). Strikingly, a recent paper suggests that CB-driven evidence accumulation can sometimes enhance performance from purely computational perspective in both perceptual and higher-order inference tasks (Glickman, Moran, and Usher 2020). On a broader level, CIB is often considered as an instance of confirmation bias (Nickerson 1998; Talluri et al. 2018), which has been proposed to to facilitate group cooperation and stability (Peters 2020).

A different line of argumentation is that human brains are not well suited for solving task consisting of independent samples, which provide a constraint on our rational processing capacities (Griffiths, Lieder, and Goodman 2015; Yoo, Hayden, and Pearson 2021). Such an inductive bias might be a product of evolutionary adaptation to temporally-dependent environments, such as foraging patches (Stephens and Anderson 2001).

Further research is needed to distinguish between these competing hypotheses and shed light on the mystery of this seemingly irrational bias.

In a set of additional experiments, we test a set of alternative explanations, as well as constraints of the CIB effect. In Experiment 3 we find no effect of involuntary choices on CIB, finding which is consistent with the preference literature (Sharot, Velasquez, and Dolan 2010). This shows that a simple motor action is not enough for the bias to occur, but rather, a free (voluntary) choice is necessary, a finding consistent with an action-driven theory of CIB (Harmon-Jones, Harmon-Jones, and Levy 2015). Additionally, this proves that CIB is not a result of a simple motor bias towards the just performed action due to motor costs associated with switching (Orban de Xivry and Lefèvre 2016).

In Experiment 4, after each binary choice, we replace one of the two items to test the effects of chosen and rejected item separately. Previous literature on post-choice dissonance produced a set of inconsistent results with regards to this issue, from undervaluation effects being stronger, (Brehm 1956, Izuma et al. 2010), both being roughly equal (Harmon-Jones et al. 2008; Jarcho, Berkman, and Lieberman 2011; Luo and Yu 2016) to overvaluation being stronger (Sharot, Velasquez, and Dolan 2010). Our results indicate that both processes may contribute to a similar extent, with a small but significant difference in favour of rejected item devaluation being stronger than chosen item overvaluation. Additionally, the design allowed us to test a potential attentional confound. In particular, visual attention biases evidence accumulation in favor of the fixated item (Tavares, Perona, and Rangel 2017). If attention is biased towards the chosen item which then drives the bias, then

we should see a significant reduction of CIB when the chosen item is replaced with an alternative, compared to when it remains on the screen, which was not the case.

In Experiment 5, we test another potential interpretation of CIB that is unrelated to conflict-driven processing or value updating - the positivity effect (Hunt et al. 2016). Assuming participants treat the chosen option as the default, a bias could arise due to uneven weighting of positive (pro-choice) vs negative evidence. By controlling for which item is the positive reference, we entangle these effects, showing that positivity bias and CB act independently, both contributing strongly to the observed bias. Since the reference-independent CB was similar in size to the one observed in all previous experiments, we conclude that the positivity effect cannot account for the effect of choice.

Current task was designed to be a robust frame for comparing CIB between domains and choice-judgment congruencies. This comes at a cost, limiting the ability to test more specific mechanisms. One weakness of the current design is that it does not allow to directly compare the two types of inferred mechanisms. Inference based on generative modelling provides a very powerful way of compensating for this deficit. Since however the generative model space in theory can be infinite, it can provide only indirect, relative evidence (since there can always exist a model with different parametrization that fits the data better). A more stringent test of this would involve a modified, dedicated design, which is a potential future direction we believe should be pursued.

Another direction for future studies involves a deeper exploration of how

choices influence value uncertainties. Current model assumes that the value-update mechanism affects the central tendencies of value estimation, but it is also possible that choices reduce the uncertainty around the estimates. Comparing the current model with a Bayesian value update model (*e.g.* Meyniel, Sigman, and Mainen 2015), where both means and uncertainties fluctuate as a function of choices made can lead to a more accurate preditions and a better understanding of the cognitive process. One testable prediction is that VU mechanism should influence the uncertainty parameter, while CIB should not.

Yet another unexplored direction revolves around a deeper understanding on the temporal dynamics of CIB. A promising approach of studying this would be to substitute the simple *softmax* choice rule (Ahn, Haines, and Zhang 2017) with a more cognitively feasible process, such as sequential sampling (Ratcliff and Smith 2004). Such approach would allow for testing a prediction that updated values should be reflected in accumulation rates. Recent modelling advancements (Kvam and Turner 2021; Ratcliff and McKoon 2020; Smith et al. 2020) bring promise for modelling continuous judgments in a similar fashion, and allowing to test whether the hypothesis that VU should affect the accumulation rate, while CB should only influence the accumulation starting point.

Future research should also focus on identifying a neural signature of both CB and VU mechanisms. Previous research shows preference value-updating involves dorsal striatal (Izuma et al. 2010) and hippocampal activity during reevaluation (Chammat et al. 2017), suggesting choices affect both immediate value as well as modify memory represtantions; while choice conflict is tracked

by anterior cingulate cortex and dorsolateral prefrontal cortex (Izuma et al. 2010). This distinction corresponds well with the idea of two separate mechanisms driving CIB. A careful study design and methodology could provide evidence of these two mechanisms driving CIB also on a neural level. One promising way of approaching this problem is using multivariate voxel-pattern analysis (MVPA) for distinguishing shared and non-overlapping representations (e.g. Vermeylen et al. 2020) of *within-domain* and *across-domain* adjustments.

6.8.1 Conclusion

We show that people can be influenced by their choices not only when the choices are relevant to the evaluation (*within-domain*), but also when they are not (*across-domains*). CIB is specific to voluntary choices, asymmetric, such that rejected items are undervalued more strongly than chosen items are overvalued, and independent of reference. The existence of *across-domain* CIB and the mechanisms driving it can help us understand seemingly irrational, yet socially relevant behaviours such as changing political opinions based on voting or the emergence of polarized in-group worldviews, driven by one's choices and actions. While previous studies attempted to explain CIB with a single process, our generative modelling suggests that CIB is driven by two separate mechanisms: a conflict driven consistency bias contributing to any evaluation immediately following a choice, and a value-update, affecting only choice-congruent judgments. Both mechanisms have different temporal dynamics and functional interpretations. While value-updating is only beneficial when choice information is relevant to the evaluation, a

general consistency bias can serve to reduce uncertainty, and facilitate action implementation.

Chapter 7

Summary

My work reveals some of the influencing factors (Project 1) and mechanisms (Project 2) of voluntary choices, as well as their consequences (Project 3). While focusing on the same broad concept and using a common analytical framework (for details see *Chapter 3*), the Projects differ in how they approach the issue, exploring different facets of voluntary choices.

All three Projects focus on the study of the *what* component of volitional choices, *i.e.,* describing the cognitive processes (Projects 1, 2 and 3), behavioural signatures (1, 2 and 3), neural correlates (1 and 2), goals (2), and consequences (1 and 3). In addition, Projects 1 and 2 also include the *when* component by analysing reaction times. Most Projects utilize only the 1st person definition of voluntariness (*i.e.,* objectively defined; as opposed to subjectively experienced; Frith 2013) with the exception of Project 1, which also examines the link between them. Tasks consist of perceptual

(1) and preference (2) variants of the equal-choice paradigm, as well as a modified version the free-choice paradigm (3). Projects 1 and 2 put focus on how responses are generated, assuming a linear accumulation to threshold mechanism. Project 3 puts emphasis on how choices influence future choices and evaluations, which is modelled using a reinforcement learning mechanism with a modified version of Rescorla-Wagner update rule, assuming internal feedback.

7.1 General Implications and Limitations

The Projects described in the thesis provide new insights into the study of voluntary choice, linking some of the behavioural, cognitive and neural processes during and after choice. The current section will summarise the implications and limitations of the three Projects. In contrast to Chapters focusing on individual Projects, here I will discuss the implications that link the Projects and discuss the common themes of the thesis, dividing them by inference type, distinguishing between behavioural findings, inferred cognitive mechanisms, and neural underpinnings.

7.1.1 Behavioural Findings

Projects 1 & 2 show how voluntary choice behavior is modulated by a variety of factors, such as option availability and discriminability (Project 1), or reward probability (Project 2), as well as how these factors affect the subjective feeling of freedom (Project 1). Advanced behavioural modelling in Project 2 also revealed a spontaneous preference among equally valuable

options, which biased the choice frequency and facilitated response times. Additionally, Project 1 also assessed the subjective experience of freedom, which was weakly modulated by both number of alternatives and stimuli discriminability. These findings provide elements of a behavioural fingerprint of voluntary decision making in a forced-choice setting.

Across all Projects, manipulations of difficulty (defined as the difference in value between the options) and magnitude (defined as the total value of options available to choose) played an important role during the choice process, affecting reaction times (Projects 1 & 2) as well as influencing post-choice evaluation by driving the choice-induced bias (Project 3). Magnitude, operationalized as the absolute value of the two reward probabilities available for choice, was of central importance in Project 2, showing a strong facilitating effect on reaction times. A similar operationalization was used in Project 3, where magnitude and difficulty related to either estimations of size or subjective value. The results show that both difficulty (or conflict) and magnitude positively modulated the choice-induced bias effect, consistent with ideas suggested in the cognitive dissonance literature (E. Harmon-Jones and Harmon-Jones 2019). The situation with Project 1 is a bit more complicated, as the effects of discriminability can be interpreted either in terms of difficulty (difference in contrast value between the available and unavailable alternatives) or magnitude (the absolute value of contrast of the available alternatives). Assuming the first case, the manipulation is related to the difficulty in detection between available and unavailable options, making it similar to classical paradigms, manipulating difficulty by controlling perceptual noise levels (*e.g.* Mulder, Van Maanen, and Forstmann 2014;

Liberman and Förster 2006). Assuming the second, the facilitatory effects of high discriminability are akin to the magnitude effect found in Project 2, or in similar perceptual studies (Teodorescu, Moran, and Usher 2016; Ratcliff, Voskuilen, and Teodorescu 2018). While the only way to disentangle the two would be to manipulate them independenty (*e.g.* Teodorescu, Moran, and Usher 2016; Ratcliff, Voskuilen, and Teodorescu 2018), it is reasonable to assume that the observed effect is a product of both phenomena. Project 3 provides further evidence of the post-choice processing being influenced by both, with choice-induced bias going up when the conflict is greater and the magnitude is larger. This was true for both domains (perception and preference), as well as for both *within* and *across-domain* effects.

While the effects of difficulty and magnitude across all three Projects are consistent and provide a coherent narrative of both factors significantly influencing both choice and underlying subjective value, it is important to remember that this understanding is constrained by the nature of the tasks and the definitions used. Especially in the case of difficulty, it is rather unclear how well these operationalizations generalize to real world problems, where difficulty might be influenced by factors other than the difference on a unified subjective value scale. The assumption of such conceptualisation of difficulty is that value can be understood as a unified dimension (Rafael Polania et al. 2014; Fontanesi et al. 2019).

This highlights a significant difference between the perceptual and preference-based paradigms. While such conceptualisation is rather natural for perceptual tasks, where items are usually assessed only on a single dimension, such as size (Rafael Polania et al. 2014), motion direction (Mulder et al.

2012) or phase coherence (Ratcliff, Philiastides, and Sajda 2009), it is not necessarily true for preference, where value might be a reflection of different item attributes (Luce, Payne, and Bettman 1999; Dhar 1994; Cheng and González-Vallejo 2018), and it is not obvious how these attributes can be integrated into a single value. For instance, in case of choosing a meal, the attributes might consist of taste, price, and the aesthetics of the food item. These attributes might be assigned different levels of importance, and a choice might be reached by the means of heuristics, instead of subjective value calculation (Gigerenzer and Gaissmaier 2011). Similarly, heuristics such as *context*, *similarity* of *compromise* effects might play a significant role when choosing between multiple alternatives, irrespective of choice domain (Spektor, Bhatia, and Gluth 2021). In fact, real-life decisions, such as buying groceries or choosing a career, are usually quite more complex than highly specific experimental paradigms which put rigid constrains on both the number, and contents of available options, as well as the time to respond. This arguably makes such choices a subject to a larger number of influences, affecting the difficulty and choice strategies.

Overall, while the closeness in subjective value is a clear indicator of choice difficulty in a highly constrained experimental paradigms used in this thesis, one has to be careful not to neglect the impact of individual attributes and other factors influencing real-life decisions, when generalising the findings to more ecologically valid contexts.

7.1.2 Inferred Cognitive Mechanisms

Projects 1 and 2 indicate that the decision process between equal options can be reliably modelled using the evidence accumulation framework. In both Projects, a linear accumulator model provided good fits for the observed response behavior. Additionally, in Project 2, the estimated trial-level evidence accumulation rate was predictive of middle prefrontal EEG signal. This is consistent with an emerging consensus that evidence accumulation is a reliable method used by the brain to make choices not only in the perceptual (Gold and Shadlen 2007), but also value-based (Pisauro et al. 2017; Rafael Polania et al. 2014) and memory-based (Supekar et al. 2021) domains. While the linear ballistic model provided a good fit to both paradigms, it is important to note that it was not formally compared against other potential mechanisms (*e.g.* urgency gating: Yau, Hinault, et al. 2020) which could provide a better fit, so one has to be careful in inferring about the true mechanism driving choice. The specifications of the accumulation process also remain debatable. For example, both Projects assumed that the boundary parameter cannot vary between conditions, unless the condition type is known *a priori* (Ratcliff and Smith 2004; Forstmann et al. 2010). This assumption however can be challenged (Simen 2012). It is possible to imagine that thresholds could be set in an early decoding stage of the pre-decisional processing. To accommodate for this possibility, I perform an additional analysis involving modelling Project 1 data using LBA with threshold parameter varying *per* condition (see *Appendix*). Other realistic possibilities involve combining evidence accumulation with a time-dependent urgency signal (Miletic 2016) which can interact with early visual processing. For

example, an alternative interpretation of the disciminability effect in Project 1, or magnitude effect observed in Project 2 could relate to the urgency (and not evidence or threshold) being amplified by stimulus saliency.

In Project 3, I propose that voluntary choices influence our future value estimations in both the short term, driven by a need for consistency (Festinger 1957; Brehm 2007; Harmon-Jones, Harmon-Jones, and Levy 2015), as well as in the long term, by updating value estimations by one's choices. Such a value-updating mechanism is similar to the active inference hypothesis (Bem 1967; Kruglanski et al. 2018) driven by choices, using a simple Rescorla-Wagner updating rule (Rescorla 1972), found in a range of other paradigms studying human learning with explicit rewards (Ahn, Haines, and Zhang 2017; Zajkowski, Kossut, and Wilson 2017; L. Zhang et al. 2020).

From a generative perspective, it is interesting to see whether, consistently with Projects 1 and 2, modelling choices using accumulation evidence framework could provide a better fit than a softmax choice rule (Ahn, Haines, and Zhang 2017). Similar models which combine the two frameworks have started to gain attention in recent research (Fontanesi et al. 2019; Miletić, Boag, and Forstmann 2020). Such an approach would allow for richer inferences regarding the specific elements of the choice process (accumulation rates, thresholds, starting points) being affected by consistency-driven bias and value-update mechanisms, found to drive choice-induced bias. Similarly, modelling judgments using evidence accumulation models adjusted for a continuous response space (Kvam and Turner 2021; Ratcliff and McKoon 2020; Smith et al. 2020) could differentiate between different processes guiding the judgment estimation. A particularly interesting hypothesis derived from the

two mechanisms driving choice-induced bias is that short-term conflict-driven effects would only bias the judgment accumulation starting point, while long-term value updates would be reflected in accumulation rates, consistent with studies showing trial-level perturbations of accumulation speed by subjective value (Krajbich et al. 2012).

7.1.3 Neural Underpinnings

Projects 1 & 2 deal with neural underpinnings of voluntary choice. In Project 1, I show that voluntary choices (as opposed to instructed ones, where only one alternative is available) are associated with the activity of the SMA/preSMA cortex and the parietal lobule, as measured by the BOLD signal in an fMRI experiment. This finding is consistent with a recent meta-analysis of studies using the ECP (Si, Rowe, and Zhang 2020). Similarly located medial central cortical region differentiated between reward probability conditions in Project 2. Both of these findings suggest a crucial role of the medial prefrontal cortex (mPFC) playing a important role in executive control of voluntary decisions (Shenhav, Botvinick, and Cohen 2013). The exact mechanism however is uncertain.

In Project 2 we find evidence for the slope of single-trial EEG signal between the N100 and P300 components being associated with the speed of evidence accumulation, as modelled by a linear accumulator model, consistent with literature suggesting the medial P300 EEG component being a marker of accumulation (S. Kelly and O'Connell 2013; D. Twomey et al. 2015). We however did not find support for a similar accumulation mechanism to be associated with preSMA activation on a subject-level in Project 1. This

might have been caused by several reasons. Firstly, preSMA might in fact be guided by a different principle, such as processing conflict (Iannaccone et al. 2015) or adjusting decisional threshold (Forstmann et al. 2010). According to the first view, preSMA is involved in resolving conflict between competing action plans by triggering response switching (Nachev et al. 2005). This is evidenced by a negative N200 amplitude EEG component, which is a marker of conflict monitoring, being associated with preSMA activation (Iannaccone et al. 2015). According to the second view, preSMA activity regulates decision threshold by modulating striatal excitability *via* a direct white matter pathway (Forstmann et al. 2010). Assuming that the 2 and 3-alternative conditions require a similar level of cognitive control (which would be consistent with very similar levels of accuracy and reaction times across the 2 and 3 available option conditions), such explanations would be compatible with no differences in BOLD activations observed between these two conditions. Alternatively, the preSMA/SMA region might in fact be driven by evidence accumulation (arguments for this view in the literature are currently mixed, see: Tomassini et al. 2019; Tosun et al. 2017; Berkay et al. 2018), but the methods used in Project 1 were not sensitive enough to detect the effect.

Reconciling the seemingly inconsistent results of the preSMA cluster not being related to evidence accumulation in Project 1, but medial PFC region being associated with accumulation speed in Project 2 is possible, given an anatomical distinction between the two regions, with the more frontal and lateral preSMA cluster being involved in conflict-driven boundary setting, while the slightly more posterior and central mPFC being involved in accumu-

lation (Pisauro et al. 2017), supported by the region's functional connection with the ventral prefrontal cortex, associated with a global subjective value signal (Hare et al. 2010).

7.2 Future Directions

The diversity of the Projects described in this thesis lends itself to a long vector of possible future research directions, many of which have been proposed in the Chapters describing individual Projects. From a broader perspective, I believe that a long-run goal of this line of research should be to provide a cohesive model explaining the behavioural, cognitive and neural mechanisms of voluntary decision making.

To achieve this, it is necessary to fill in the gaps between *whether*, *what* and *when* components of the choice process (Brass and Haggard 2008), as well as between the computational, algorithmic and implementational levels of explanation (Marr 1982). The following sections are organized with regards to these distinctions, describing the possible follow-up research directions with regards to the *whether* and *when* component of voluntary choice, and the computational, algorithmic and implementational levels of analysis.[1]

7.2.1 *Whether* Component of Voluntary Choice

Studies described here focus mostly on the *what* and *when* elements of choice (Table 7.1), omitting the decision *whether* an action should be performed.

[1]The *what* component has been purposefully omitted here, since describing future directions would be reiterating discussion contents from Chapters 4-6.

This underexplored component could be studied using modified version of the paradigms studied here. For example, Project 1 design could include a *no-go* condition, where none of the options were available to choose. In Project 3, a *no-go* condition could include making a choice but withholding the response (for a similar design see: Chambon et al. 2020). Theoretical accounts predict that refraining from choice should reduce choice-induced bias (Harmon-Jones, Harmon-Jones & Levy, 2015) as well as decrease the choice calibration (Kvam, Pleskac & Busemeyer, 2015).

7.2.2 *When* Component of Voluntary Choice

The *when* component remains unexplored in Project 3. This is due to two fundamental limitations of the design. Firstly, judgment scales in all Experiments were made on a straight horizontal (or vertical: Experiment 5) scale, and since the mouse position was centered, it takes longer to move the mouse in the direction of more extreme values, located further from the center, necessarily making the reaction time estimation biased. Secondly, in order to provide a more familiar and consistent user experience (participants are more comfortable to respond by dragging the mouse for both choices and judgments, than switching between key-value mappings and mouse dragging interchangeably in every trial) the online Experiments used mouse-guided responses in both choices and judgments, rendering reaction times more prone to noise (Plant, Hammond, and Whitehouse 2003). Potential resolution to the first issue is designing a semicircular judgment scale with the cursor placed in the origin, so that at the start of the judgment all points on the scale are equidistant and reaction times are not biased by the distance needed for the

mouse to traverse to any point on the scale. The resolution of the second issue would require using a different means for providing responses, which would be comfortable for both binary choices and continuous judgments. One possibility would be training participants in using a joystick.

7.2.3 Computational Level of Analysis

In this work, I used two computational (as in referring to the computational level of analysis) frameworks. In Projects 1 and 2, I model the choice process using an evidence accumulation framework, assuming that perceptual, preferential and memory-based evidence is sampled and integrated until a predefined threshold is reached. In Project 3, I model the temporal evolution of choice-induced bias similar to a reinforcement learning process with one's choices playing the role of implicit rewards. While these approaches proved to fit the data well, it is possible that testing alternative frameworks could provide a more accurate description of the data. Future studies could focus on comparing different possible mechanisms driving choice (such as urgency gating, Thura et al. 2012). The design of Project 3 makes it also a good candidate for linking both mechanisms by modelling choice using evidence accumulation, and value-update using a learning model (Miletić, Boag, and Forstmann 2020).

7.2.4 Algoritmic Level of Analysis

Future research regarding all of the Projects contained here would benefit from more in-depth work on modelling the algorithmic level of cognitive processing.

In Projects 1 and 2 I utilize linear ballistic accumulation, a model assuming that each option is represented by a independent accumulator in a race towards a common threshold (Brown and Heathcote 2008). Many alternative representations of the accumulation process are also used in the literature, such as the drift-diffusion model, which assumes that the noisy difference in value is accumulated (Ratcliff 1978), Leaky Accumulator Model, which assumes that the accumulators decay can mutually inhibit each other (Usher and McClelland 2001), as well as other variants of thereof (*e.g.* non-static threshold model, Evans 2020). While some tests have been performed using classical perceptual designs (Hawkins et al. 2015), a systematic comparison among these models in an ECP task could result in providing a more accurate description of the observed data.

In Project 3, using the Rescorla-Wagner rule to update choice value is only one of the possibilities. An interesting alternative involves modelling value representations as distributions, instead of single points. Such approach would allow for modelling the choice-induced update process similar to Bayesian Learning, where value uncertainties play a critical role (*e.g.* Daw et al. 2006; Zajkowski, Kossut, and Wilson 2017).

7.2.5 Implementational Level of Analysis

Project 3 is the only one devoid from any implementational elements, as the set of experiments was behavioural in nature. Future work should explore how the proposed mechanisms are realised in the brain. Candidate brain regions involved in biasing judgments towards choice *via* a consistency bias include anterior cingulate cortex and dorsolateral prefrontal cortex (Izuma

et al. 2010). Areas potentially involved in long-term value-updating include dorsal striatal (Izuma et al. 2010) and the hippocampus (Chammat et al. 2017).

7.2.6 Further Linking the Levels of Analysis

All of the Projects provide a link between at least two levels of analysis. Additional links could be acquired by using more diverse imaging methodology. Combining EEG with fMRI could help in attaining high temporal (EEG) and spatial (fMRI) resolution (B. Turner et al. 2016). Such approach could help in reconciling the dilemma regarding the roles of preSMA *vs* mPFC cortices in accumulation discussed in the previous section. Adding eye-tracking to either fMRI or EEG could allow for precise modelling of attention and a better estimate of accumulation rates (Krajbich et al. 2012), which would allow mapping it in the brain with higher precision. As described in Chapter 3, neural data can serve not only as a means of describing how a process is implemented in the brain, but also to actively constrain the algorithmic model, aiding in model selection (Brandon M. Turner et al. 2016). Future work should consider these possibilities as means of obtaining better models on both the algorithmic and implementational levels.

7.3 Conclusion

In this thesis I explore the behavioural, cognitive, and neural mechanisms of voluntary decision making. The diverse set of results from the three Projects, including voluntary choices between: abstract perceptual stimuli

(Project 1), symbols associated with reward probabilities (Project 2), and concrete food items, based on their perceptual and preferential qualities (Project 3), provide an insight into the decision making process and how it affects future choices and evaluations. My results indicate that voluntary choices are influenced by such factors as the number of available alternatives, choice difficulty, magnitude of the value alternatives, as well as an evolving subjective preference, driven by previous choices.

Cognitive modelling using a linear ballistic accumulation to threshold (LBA, Brown and Heathcote 2008) indicates that the activity in the medial prefrontal cortex, area in general associated with voluntary choices (Project 1), is predictive of evidence accumulation speed on a trial-to-trial level (Project 2). Finally, in Project 3 I show how voluntary choices influence future choices and value estimations, not only within the domains of perception and preference, but also across domains, *i.e.* such that preference choice can influence perceptual evaluation, and *vice versa*. This effect is driven by both, a short-term conflict driven consistency bias, and a long-term value update, similar to learning from one's previous choices. The results described above provide further insights into the complex nature of voluntary choice.

References

Abler, Birgit, Bärbel Herrnberger, Georg Grön, and Manfred Spitzer. 2009. "From Uncertainty to Reward: BOLD Characteristics Differentiate Signaling Pathways." *BMC Neuroscience* 10 (1): 154.

Adams, Rick A., Quentin J. M. Huys, and Jonathan P. Roiser. 2016. "Computational Psychiatry: Towards a Mathematically Informed Understanding of Mental Illness." *J Neurol Neurosurg Psychiatry* 87 (1): 53–63. https://doi.org/10.1136/jnnp-2015-310737.

Afacan-Seref, Kivilcim, Natalie A Steinemann, Annabelle Blangero, and Simon P Kelly. 2018. "Dynamic Interplay of Value and Sensory Information in High-Speed Decision Making." *Current Biology* 28 (5): 795–802.

Ahn, Woo-Young, Nathaniel Haines, and Lei Zhang. 2017. "Revealing Neurocomputational Mechanisms of Reinforcement Learning and Decision-Making with the hBayesDM Package." *Computational Psychiatry (Cambridge, Mass.)* 1 (October): 24–57. https://doi.org/10.1162/CPSY_a_00002.

Aine, C. J. 1995. "A Conceptual Overview and Critique of Functional Neuroimaging Techniques in Humans: I. MRI/FMRI and PET." *Crit Rev*

Neurobiol 9 (2-3): 229–309.

Akaishi, Rei, Kazumasa Umeda, Asako Nagase, and Katsuyuki Sakai. 2014. "Article Autonomous Mechanism of Internal Choice Estimate Underlies Decision Inertia." *NEURON* 81 (1): 195–206. https://doi.org/10.1016/j.neuron.2013.10.018.

Alloy, L. B., and L. Y. Abramson. 1979. "Judgment of Contingency in Depressed and Nondepressed Students: Sadder but Wiser?" *J Exp Psychol Gen* 108 (4): 441–85. https://doi.org/10.1037//0096-3445.108.4.441.

Anders, R., Z. Oravecz, and F.-X. Alario. 2018. "Improved Information Pooling for Hierarchical Cognitive Models Through Multiple and Covaried Regression." *Behav Res* 50 (3): 989–1010. https://doi.org/10.3758/s13428-017-0921-7.

Annis, Jeffrey, Brent J. Miller, and Thomas J. Palmeri. 2017. "Bayesian Inference with Stan: A Tutorial on Adding Custom Distributions." *Behavior Research Methods* 49 (3): 863–86. https://doi.org/10.3758/s13428-016-0746-9.

Ariely, Dan, and Michael I Norton. 2007. "How Actions Create – Not Just Reveal – Preferences," no. December: 13–16. https://doi.org/10.1016/j.tics.2007.10.008.

Aron, Adam R., Trevor W. Robbins, and Russell A. Poldrack. 2004. "Inhibition and the Right Inferior Frontal Cortex." *Trends Cogn Sci* 8 (4): 170–77. https://doi.org/10.1016/j.tics.2004.02.010.

Aronson, J., H. Blanton, and J. Cooper. 1995. "From Dissonance to Disidentification: Selectivity in the Self-Affirmation Process." https://doi.org/10.

1037/0022-3514.68.6.986.

Baayen, R. H., D. J. Davidson, and D. M. Bates. 2008. "Mixed-Effects Modeling with Crossed Random Effects for Subjects and Items." *Journal of Memory and Language* 59 (4): 390–412. https://doi.org/10.1016/j.jml.2007.12.005.

Badgaiyan, R. D., and M. I. Posner. 1998. "Mapping the Cingulate Cortex in Response Selection and Monitoring." *Neuroimage* 7 (3): 255–60. https://doi.org/10.1006/nimg.1998.0326.

Bagiella, Emilia, Richard P Sloan, and Daniel F Heitjan. 2000. "Mixed-Effects Models in Psychophysiology." *Psychophysiology* 37 (1): 13–20.

Bakkour, Akram, Ariel Zylberberg, Michael N Shadlen, and Daphna Shohamy. 2018. "Value-Based Decisions Involve Sequential Sampling from Memory." *BioRxiv*, 269290.

Ballard, Ian C, Bokyung Kim, Anthony Liatsis, Gökhan Aydogan, Jonathan D Cohen, and Samuel M McClure. 2017. "More Is Meaningful: The Magnitude Effect in Intertemporal Choice Depends on Self-Control." *Psychological Science* 28 (10): 1443–54.

Barlas, Zeynep, William E Hockley, and Sukhvinder S Obhi. 2017. "Acta Psychologica the E Ff Ects of Freedom of Choice in Action Selection on Perceived Mental E Ff Ort and the Sense of Agency" 180 (August): 122–29. https://doi.org/10.1016/j.actpsy.2017.09.004.

Barlas, Zeynep, and Stefan Kopp. 2018. "Action Choice and Outcome Congruency Independently Affect Intentional Binding and Feeling of Control Judgments." *Front. Hum. Neurosci.* 12. https://doi.org/10.3389/fnhum.

2018.00137.

Barlas, Zeynep, and Sukhvinder Obhi. 2013. "Freedom, Choice, and the Sense of Agency." *Front. Hum. Neurosci.* 7. https://doi.org/10.3389/fnhum.2013.00514.

Barr, Dale J, Roger Levy, Christoph Scheepers, and Harry J Tily. 2013. "Random Effects Structure for Confirmatory Hypothesis Testing: Keep It Maximal." *Journal of Memory and Language* 68 (3): 10.1016/j.jml.2012.11.001. https://doi.org/10.1016/j.jml.2012.11.001.

Bates, Douglas, Reinhold Kliegl, Shravan Vasishth, and Harald Baayen. 2015. "Parsimonious Mixed Models." *arXiv Preprint arXiv:1506.04967*.

Bates, Douglas, Martin Machler, Ben Bolker, and Steve Walker. 2014. "Fitting Linear Mixed-Effects Models Using Lme4." 2014. https://doi.org/10.18637/jss.v067.i01.

Bayarri, M Jésus, and James O Berger. 2004. "The Interplay of Bayesian and Frequentist Analysis." *Statistical Science*, 58–80.

Beck, Brianna, Steven Di Costa, and Patrick Haggard. 2017. "Having Control over the External World Increases the Implicit Sense of Agency." *Cognition* 162 (May): 54–60. https://doi.org/10.1016/j.cognition.2017.02.002.

Becker, Gary S. 1962. "Irrational Behavior and Economic Theory." *Journal of Political Economy* 70 (1): 1–13. http://www.jstor.org/stable/1827018.

Bem, Daryl J. 1967. "Self-Perception: An Alternative Interpretation of Cognitive Dissonance Phenomena." *Psychological Review* 74 (3): 183–200. https://doi.org/10.1037/h0024835.

Berkay, Dilara, Hale Yapıcı Eser, Alexander T. Sack, Yusuf Özgür Çakmak, and Fuat Balcı. 2018. "The Modulatory Role of Pre-SMA in Speed-Accuracy Tradeoff: A Bi-Directional TMS Study." *Neuropsychologia* 109 (January): 255–61. https://doi.org/10.1016/j.neuropsychologia.2017.12.031.

Berlemont, Kevin, and Jean-Pierre Nadal. 2019. "Perceptual Decision-Making: Biases in Post-Error Reaction Times Explained by Attractor Network Dynamics." *J. Neurosci.* 39 (5): 833–53. https://doi.org/10.1523/JNEUROSCI.1015-18.2018.

Betancourt, Michael. 2017. "A Conceptual Introduction to Hamiltonian Monte Carlo." 2017. http://arxiv.org/abs/1701.02434.

Betzel, Richard F., and Danielle S. Bassett. 2017. "Generative Models for Network Neuroscience: Prospects and Promise." *Journal of the Royal Society Interface* 14 (136): 20170623. https://doi.org/10.1098/rsif.2017.0623.

Blakemore, Sarah Jayne, Daniel M. Wolpert, and Christopher D. Frith. 2002. "Abnormalities in the Awareness of Action." *Trends Cogn Sci* 6 (6): 237–42. https://doi.org/10.1016/s1364-6613(02)01907-1.

Blechert, Jens, Anja Lender, Sarah Polk, Niko A Busch, and Kathrin Ohla. 2019. "Food-Pics_Extended—an Image Database for Experimental Research on Eating and Appetite: Additional Images, Normative Ratings and an Updated Review." *Frontiers in Psychology* 10: 307. https://doi.org/10.3389/fpsyg.2019.00307.

Bode, Stefan, Carsten Bogler, and John-Dylan Haynes. 2013. "Similar Neural Mechanisms for Perceptual Guesses and Free Decisions." *Neuroimage* 65 (January): 456–65. https://doi.org/10.1016/j.neuroimage.2012.09.064.

Bode, Stefan, David K. Sewell, Simon Lilburn, Jason D. Forte, Philip L. Smith, and Jutta Stahl. 2012. "Predicting Perceptual Decision Biases from Early Brain Activity." *J Neurosci* 32 (36): 12488–98. https://doi.org/10.1523/JNEUROSCI.1708-12.2012.

Boehm, Udo, Guy E Hawkins, Scott Brown, Hedderik van Rijn, and Eric-Jan Wagenmakers. 2016. "Of Monkeys and Men: Impatience in Perceptual Decision-Making." *Psychonomic Bulletin & Review* 23 (3): 738–49.

Bogacz, Rafal, Eric Brown, Jeff Moehlis, Philip Holmes, and Jonathan D Cohen. 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.

Bogacz, Rafal, Marius Usher, Jiaxiang Zhang, and James L McClelland. 2007. "Extending a Biologically Inspired Model of Choice: Multi-Alternatives, Nonlinearity and Value-Based Multidimensional Choice." *Philosophical Transactions of the Royal Society B: Biological Sciences* 362 (1485): 1655–70.

Brass, Marcel, and Patrick Haggard. 2008. "The What, When, Whether Model of Intentional Action." *Neuroscientist* 14 (4): 319–25. https://doi.org/10.1177/1073858408317417.

———. 2007. "To Do or Not to Do: The Neural Signature of Self-Control." *J. Neurosci.* 27 (34): 9141–5. https://doi.org/10.1523/JNEUROSCI.0924-07.2007.

Brehm, Jack W. 2007. "A Brief History of Dissonance Theory" 1: 381–91.

Brehm, J. W. 1956. "Postdecision Changes in the Desirability of Alternatives." *J Abnorm Psychol* 52 (3): 384–89. https://doi.org/10.1037/h0041006.

Bronfman, Zohar, Noam Brezis, Rani Moran, Konstantinos Tsetsos, Tobias Donner, and Marius Usher. 2015. "Decisions Reduce Sensitivity to Subsequent Information." *Proceedings of the Royal Society B: Biological Sciences* 282 (1810): 20150228. https://doi.org/10.1098/rspb.2015.0228.

Brown, Scott D, and Andrew Heathcote. 2008. "The Simplest Complete Model of Choice Response Time: Linear Ballistic Accumulation." *Cognitive Psychology* 57 (3): 153–78.

Burkner, Paul-Christian. 2017. "Brms: An R Package for Bayesian Multilevel Models Using Stan." *Journal of Statistical Software* 80 (1, 1): 1–28. https://doi.org/10.18637/jss.v080.i01.

Busemeyer, Jerome R., and Joseph G. Johnson. 2004. "Computational Models of Decision Making." In *Blackwell Handbook of Judgment and Decision Making*, 133–54. Malden: Blackwell Publishing. https://doi.org/10.1002/9780470752937.ch7.

Buzsaki, Gyorgy. 2019. *The Brain from Inside Out.* Oxford University Press.

Carpenter, Bob, Andrew Gelman, Matthew D. Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2017. "*Stan* : A Probabilistic Programming Language." *Journal of Statistical Software* 76 (1). https://doi.org/10.18637/jss.v076.i01.

Caspar, Emilie A., Julia F. Christensen, Axel Cleeremans, and Patrick Haggard. 2016. "Coercion Changes the Sense of Agency in the Human Brain." *Curr Biol* 26 (5): 585–92. https://doi.org/10.1016/j.cub.2015.12.067.

Cavanagh, James F, Thomas V Wiecki, Michael X Cohen, Christina M Figueroa, Johan Samanta, Scott J Sherman, and Michael J Frank. 2011.

"Subthalamic Nucleus Stimulation Reverses Mediofrontal Influence over Decision Threshold." *Nature Neuroscience* 14 (11): 1462.

Chambon, Valerian, and Patrick Haggard. 2012. "Sense of Control Depends on Fluency of Action Selection, Not Motor Performance." *Cognition* 125 (3): 441–51. https://doi.org/10.1016/j.cognition.2012.07.011.

Chambon, Valerian, Nura Sidarus, and Patrick Haggard. 2014. "From Action Intentions to Action Effects: How Does the Sense of Agency Come About?" *Front Hum Neurosci* 8 (May). https://doi.org/10.3389/fnhum.2014.00320.

Chambon, Valérian, Héloïse Théro, Marie Vidal, Henri Vandendriessche, Patrick Haggard, and Stefano Palminteri. 2020. "Information About Action Outcomes Differentially Affects Learning from Self-Determined Versus Imposed Choices." *Nat Hum Behav* 4 (10): 1067–79. https://doi.org/10.1038/s41562-020-0919-5.

Chammat, Mariam, Imen El Karoui, Sébastien Allali, Joshua Hagège, Katia Lehongre, Dominique Hasboun, Michel Baulac, and Stéphane Epelbaum. 2017. "Cognitive Dissonance Resolution Depends on Episodic Memory." *Nature Publishing Group*, no. September 2016: 1–10. https://doi.org/10.1038/srep41320.

Chen, M Keith, and Jane L Risen. 2010. "How Choice Affects and Reflects Preferences : Revisiting the Free-Choice Paradigm" 99 (4): 573–94. https://doi.org/10.1037/a0020217.

Chen, Xing-jie, and Youngbin Kwak. 2017. "What Makes You Go Faster?: The Effect of Reward on Speeded Action Under Risk." *Frontiers in Psychology* 8: 1057.

Cheng, Jiuqing, and Claudia González-Vallejo. 2018. "Unpacking Decision Difficulty: Testing Action Dynamics in Intertemporal, Gamble, and Consumer Choices." *Acta Psychologica* 190 (October): 199–216. https://doi.org/10.1016/j.actpsy.2018.08.002.

Chikazoe, Junichi, Koji Jimura, Satoshi Hirose, Ken-ichiro Yamashita, Yasushi Miyashita, and Seiki Konishi. 2009. "Preparation to Inhibit a Response Complements Response Inhibition During Performance of a Stop-Signal Task." *J. Neurosci.* 29 (50): 15870–7. https://doi.org/10.1523/JNEUROSCI.3645-09.2009.

Cialdini, Robert B. 2014. "Preference for Consistency : The Development of a Valid Measure and the Discovery of Surprising Behavioral Implications," no. July. https://doi.org/10.1037/0022-3514.69.2.318.

Cichy, Radoslaw Martin, Dimitrios Pantazis, and Aude Oliva. 2014. "Resolving Human Object Recognition in Space and Time." *Nature Neuroscience* 17 (3): 455.

Cisek, Paul, Genevi'eve Aude Puskas, and Stephany El-Murr. 2009. "Decisions in Changing Conditions: The Urgency-Gating Model." *Journal of Neuroscience* 29 (37): 11560–71. https://doi.org/10.1523/JNEUROSCI.1844-09.2009.

Cockburn, Jeffrey, Anne G E Collins, and Michael J Frank. 2014. "A Reinforcement Learning Mechanism Responsible for the Valuation of Free Choice." *Neuron* 83 (3): 551–57. https://doi.org/10.1016/j.neuron.2014.06.035.

Cohen, Andrew L, Namyi Kang, and Tanya L Leise. 2017. "Multi-Attribute,

Multi-Alternative Models of Choice: Choice, Reaction Time, and Process Tracing." *Cognitive Psychology* 98: 45–72.

Cole, Michael W., and Walter Schneider. 2007. "The Cognitive Control Network: Integrated Cortical Regions with Dissociable Functions." *NeuroImage* 37 (1): 343–60. https://doi.org/10.1016/j.neuroimage.2007.03.071.

Colombo, Matteo, and Stephan Hartmann. 2017. "Bayesian Cognitive Science, Unification, and Explanation." *The British Journal for the Philosophy of Science* 68 (2): 451–84. https://doi.org/10.1093/bjps/axv036.

Coppin, Géraldine, Sylvain Delplanque, Charlène Bernard, Christelle Porcherot, Isabelle Cayeux, and David Sander. 2014. "The Quarterly Journal of Experimental Choice Both Affects and Re Fl Ects Preferences," no. October: 37–41. https://doi.org/10.1080/17470218.2013.863953.

Daprati, E., N. Franck, N. Georgieff, J. Proust, E. Pacherie, J. Dalery, and M. Jeannerod. 1997. "Looking for the Agent: An Investigation into Consciousness of Action and Self-Consciousness in Schizophrenic Patients." *Cognition* 65 (1): 71–86. https://doi.org/10.1016/s0010-0277(97)00039-5.

Dar-Nimrod, Ilan, Catherine D. Rawn, Darrin R. Lehman, and Barry Schwartz. 2009. "The Maximization Paradox: The Costs of Seeking Alternatives." *Personality and Individual Differences* 46 (5): 631–35. https://doi.org/10.1016/j.paid.2009.01.007.

Datta, Avijit, Rhodri Cusack, Kari Hawkins, Joost Heutink, Chris Rorden, Ian H Robertson, and Tom Manly. 2007. "The P300 as a Marker of Waning Attention and Error Propensity." *Computational Intelligence and Neuroscience* 2007.

Davis, M. 1992. "The Role of the Amygdala in Fear and Anxiety." *Annual Review of Neuroscience* 15 (1): 353–75. https://doi.org/10.1146/annurev.ne.15.030192.002033.

Daw, N D, J P O'Doherty, P Dayan, R J Dolan, and Ben Seymour. 2006. "Cortical Substrates for Exploratory Decisions in Humans." *Nature* 441 (7095): 876–9. https://doi.org/10.1038/nature04766.

Deiber, M. -P., R. E. Passingham, J. G. Colebatch, K. J. Friston, P. D. Nixon, and R. S. J. Frackowiak. 1991. "Cortical Areas and the Selection of Movement: A Study with Positron Emission Tomography." *Exp Brain Res* 84 (2): 393–402. https://doi.org/10.1007/BF00231461.

Dennett, Daniel. 1984. "Elbow Room | the MIT Press." The MIT Press. 1984. https://mitpress.mit.edu/books/elbow-room.

Desmond, John E., and Gary H. Glover. 2002. "Estimating Sample Size in Functional MRI (fMRI) Neuroimaging Studies: Statistical Power Analyses." *J Neurosci Methods* 118 (2): 115–28. https://doi.org/10.1016/s0165-0270(02)00121-8.

Desmurget, Michel, Karen T. Reilly, Nathalie Richard, Alexandru Szathmari, Carmine Mottolese, and Angela Sirigu. 2009. "Movement Intention After Parietal Cortex Stimulation in Humans." *Science* 324 (5928): 811–13. https://doi.org/10.1126/science.1169896.

Dewey, John A., and Günther Knoblich. 2014. "Do Implicit and Explicit Measures of the Sense of Agency Measure the Same Thing?" *PLOS ONE* 9 (10): e110118. https://doi.org/10.1371/journal.pone.0110118.

Dhar, Ravi. 1994. "Decision Difficulty and Uncertain Preferences: Im-

plications for Consumer Choice." *ACR North American Advances* NA-21. https://www.acrwebsite.org/volumes/7555/volumes/v21/NA-21/full.

Di Costa, Steven, Ewgenia Barow, Ute Hidding, Tina Mainka, Monika Pötter-Nerger, Carsten Buhmann, Christian K. E. Moll, Patrick Haggard, and Christos Ganos. 2020. "Dopamine Boosts Intention and Action Awareness in Parkinson's Disease." *Exp Brain Res* 238 (9): 1989–95. https://doi.org/10.1007/s00221-020-05847-2.

Dienes, Zoltan. 2011. "Bayesian Versus Orthodox Statistics: Which Side Are You on?" *Perspect Psychol Sci* 6 (3): 274–90. https://doi.org/10.1177/1745691611406920.

Dienes, Zoltan, and Neil Mclatchie. 2018. "Four Reasons to Prefer Bayesian Analyses over Significance Testing." *Psychon Bull Rev* 25 (1): 207–18. https://doi.org/10.3758/s13423-017-1266-z.

Dima, Diana C, Gavin Perry, Eirini Messaritaki, Jiaxiang Zhang, and Krish D Singh. 2018. "Spatiotemporal Dynamics in Human Visual Cortex Rapidly Encode the Emotional Content of Faces." *Human Brain Mapping* 39 (10): 3993–4006.

Ditterich, Jochen, Mark E Mazurek, and Michael N Shadlen. 2003. "Microstimulation of Visual Cortex Affects the Speed of Perceptual Decisions." *Nature Neuroscience* 6 (8): 891–98.

Donkin, Chris, Scott Brown, Andrew Heathcote, and Eric Jan Wagenmakers. 2011. "Diffusion Versus Linear Ballistic Accumulation: Different Models but the Same Conclusions About Psychological Processes?" *Psychonomic Bulletin and Review* 18 (1): 61–69. https://doi.org/10.3758/s13423-010-0022-4.

Dutilh, Gilles, Jeffrey Annis, Scott D. Brown, Peter Cassey, Nathan J. Evans, Raoul P. P. P. Grasman, Guy E. Hawkins, et al. 2019. "The Quality of Response Time Data Inference: A Blinded, Collaborative Assessment of the Validity of Cognitive Models." *Psychonomic Bulletin & Review* 26 (4): 1051–69. https://doi.org/10.3758/s13423-017-1417-2.

Ebert, Jeffrey P., and Daniel M. Wegner. 2010. "Time Warp: Authorship Shapes the Perceived Timing of Actions and Events." *Conscious Cogn* 19 (1): 481–89. https://doi.org/10.1016/j.concog.2009.10.002.

Egan, Louisa C, Paul Bloom, and Laurie R Santos. 2010. "Journal of Experimental Social Psychology Choice-Induced Preferences in the Absence of Choice : Evidence from a Blind Two Choice Paradigm with Young Children and Capuchin Monkeys" 46: 204–7. https://doi.org/10.1016/j.jesp.2009.08.014.

Eimeren, Thilo van, Thomas Wolbers, Alexander Münchau, Christian Büchel, Cornelius Weiller, and Hartwig Roman Siebner. 2006. "Implementation of Visuospatial Cues in Response Selection." *Neuroimage* 29 (1): 286–94. https://doi.org/10.1016/j.neuroimage.2005.07.014.

Esber, Guillem R, and Mark Haselgrove. 2011. "Reconciling the Influence of Predictiveness and Uncertainty on Stimulus Salience: A Model of Attention in Associative Learning." *Proceedings of the Royal Society B: Biological Sciences* 278 (1718): 2553–61.

Evans, Nathan. 2020. "Think Fast! The Implications of Emphasizing Urgency in Decision-Making." *Psyarchiv*. https://doi.org/10.31234/osf.io/pfrb4.

Evans, Nathan J., Guy E. Hawkins, Udo Boehm, Eric-Jan Wagenmakers, and

Scott D. Brown. 2017. "The Computations That Support Simple Decision-Making: A Comparison Between the Diffusion and Urgency-Gating Models." *Scientific Reports* 7 (1): 16433. https://doi.org/10.1038/s41598-017-16694-7.

Evans, Nathan J., and Eric-Jan Wagenmakers. 2019. "Evidence Accumulation Models: Current Limitations and Future Directions." https://doi.org/10.31234/osf.io/74df9.

Farmer, Thomas A., Meredith Brown, and Michael K. Tanenhaus. 2013. "Prediction, Explanation, and the Role of Generative Models in Language Processing." *Behav Brain Sci* 36 (3): 211–12. https://doi.org/10.1017/S0140525X12002312.

Ferreira, David, Mael Barthoulot, Julien Pottecher, Klaus D. Torp, Pierre Diemunsch, and Nicolas Meyer. 2020. "Theory and Practical Use of Bayesian Methods in Interpreting Clinical Trial Data: A Narrative Review." *British Journal of Anaesthesia* 125 (2): 201–7. https://doi.org/10.1016/j.bja.2020.04.092.

Festinger, Leon. 1957. "A Theory of Cognitive Dissonance." *A Theory of Cognitive Dissonance.*, xi, 291–xi, 291.

Filevich, Elisa, Patricia Vanneste, Marcel Brass, Wim Fias, Patrick Haggard, and Simone Kühn. 2013. "Brain Correlates of Subjective Freedom of Choice." *Conscious Cogn* 22 (4): 1271–84. https://doi.org/10.1016/j.concog.2013.08.011.

Fiorillo, Christopher D, Philippe N Tobler, and Wolfram Schultz. 2003. "Discrete Coding of Reward Probability and Uncertainty by Dopamine Neurons." *Science* 299 (5614): 1898–1902.

Fishbein, Martin, and Fred S. Coombs. 1974. "Basis for Decision: An Attitudinal Analysis of Voting Behavior1." *Journal of Applied Social Psychology* 4 (2): 95–124. https://doi.org/10.1111/j.1559-1816.1974.tb00662.x.

Fleming, Stephen M., and Raymond J. Dolan. 2014. "The Neural Basis of Metacognitive Ability." *The Cognitive Neuroscience of Metacognition* 9783642451: 245–65. https://doi.org/10.1007/978-3-642-45190-4_11.

Fleming, Stephen M., Rogier B. Mars, Thomas E. Gladwin, and Patrick Haggard. 2009. "When the Brain Changes Its Mind: Flexibility of Action Selection in Instructed and Free Choices." *Cereb Cortex* 19 (10): 2352–60. https://doi.org/10.1093/cercor/bhn252.

Fontanesi, Laura, Sebastian Gluth, Mikhail S. Spektor, and Jörg Rieskamp. 2019. "A Reinforcement Learning Diffusion Decision Model for Value-Based Decisions." *Psychon Bull Rev* 26 (4): 1099–1121. https://doi.org/10.3758/s13423-018-1554-2.

Forstmann, B, Alfred Anwander, Andreas Schäfer, Jane Neumann, Scott Brown, Eric-Jan Wagenmakers, Rafal Bogacz, and Robert Turner. 2010. "Cortico-Striatal Connections Predict Control over Speed and Accuracy in Perceptual Decision Making." *Proceedings of the National Academy of Sciences* 107 (36): 15916 LP–15920. https://doi.org/10.1073/pnas.1004932107.

Forstmann, Birte U, Alfred Anwander, Andreas Schäfer, Jane Neumann, Scott Brown, Eric-Jan Wagenmakers, Rafal Bogacz, and Robert Turner. 2010. "Cortico-Striatal Connections Predict Control over Speed and Accuracy in Perceptual Decision Making." *Proceedings of the National Academy of Sciences* 107 (36): 15916–20.

Forstmann, Birte U., Marcel Brass, Iring Koch, and D. Yves von Cramon. 2006. "Voluntary Selection of Task Sets Revealed by Functional Magnetic Resonance Imaging." *Journal of Cognitive Neuroscience* 18 (3): 388–98. https://doi.org/10.1162/jocn.2006.18.3.388.

Forstmann, B U, R Ratcliff, and E-J Wagenmakers. 2016. "Sequential Sampling Models in Cognitive Neuroscience: Advantages, Applications, and Extensions." *Annual Review of Psychology* 67: 641–66. https://doi.org/10.1146/annurev-psych-122414-033645.

Francis, Gregory. 2012. "Publication Bias and the Failure of Replication in Experimental Psychology." *Psychon Bull Rev* 19 (6): 975–91. https://doi.org/10.3758/s13423-012-0322-y.

Frankfurt, Harry G. 1988. "Freedom of the Will and the Concept of a Person." In *What Is a Person?*, edited by Michael F. Goodman, 127–44. Contemporary Issues in Biomedicine, Ethics, and Society. Totowa, NJ: Humana Press. https://doi.org/10.1007/978-1-4612-3950-5_6.

Frässle, Stefan, Yu Yao, Dario Schöbi, Eduardo A. Aponte, Jakob Heinzle, and Klaas E. Stephan. 2018. "Generative Models for Clinical Applications in Computational Psychiatry." *Wiley Interdiscip Rev Cogn Sci* 9 (3): e1460. https://doi.org/10.1002/wcs.1460.

Frede, Michael. 2011. *A Free Will: Origins of the Notion in Ancient Thought*. 1st ed. University of California Press. http://www.jstor.org/stable/10.1525/j.ctt1ppdd9.

Fried, I., A. Katz, G. McCarthy, K. J. Sass, P. Williamson, S. S. Spencer, and D. D. Spencer. 1991. "Functional Organization of Human Supplementary

Motor Cortex Studied by Electrical Stimulation." *J Neurosci* 11 (11): 3656–66.

Fried, Itzhak, Patrick Haggard, Biyu J. He, and Aaron Schurger. 2017. "Volition and Action in the Human Brain: Processes, Pathologies, and Reasons." *J Neurosci* 37 (45): 10842–7. https://doi.org/10.1523/JNEUROSCI.2584-17.2017.

Friston, Karl, Thomas FitzGerald, Francesco Rigoli, Philipp Schwartenbeck, John O'Doherty, and Giovanni Pezzulo. 2016. "Active Inference and Learning." *Neuroscience & Biobehavioral Reviews* 68 (September): 862–79. https://doi.org/10.1016/j.neubiorev.2016.06.022.

Friston, K. J., and C. J. Price. 2001. "Generative Models, Brain Function and Neuroimaging." *Scand J Psychol* 42 (3): 167–77. https://doi.org/10.1111/1467-9450.00228.

Frith, C. D., S. Blakemore, and D. M. Wolpert. 2000. "Explaining the Symptoms of Schizophrenia: Abnormalities in the Awareness of Action." *Brain Res Brain Res Rev* 31 (2-3): 357–63. https://doi.org/10.1016/s0165-0173(99)00052-1.

Frith, C. D., K. Friston, P. F. Liddle, and R. S. Frackowiak. 1991. "Willed Action and the Prefrontal Cortex in Man: A Study with PET." *Proc Biol Sci* 244 (1311): 241–46. https://doi.org/10.1098/rspb.1991.0077.

Frith, Chris. 2013. "The Psychology of Volition." *Exp Brain Res* 229 (3): 289–99. https://doi.org/10.1007/s00221-013-3407-6.

Gallagher, null. 2000. "Philosophical Conceptions of the Self: Implications for Cognitive Science." *Trends Cogn Sci* 4 (1): 14–21. https://doi.org/10.

1016/s1364-6613(99)01417-5.

Garrett, Deon, David A Peterson, Charles W Anderson, and Michael H Thaut. 2003. "Comparison of Linear, Nonlinear, and Feature Selection Methods for Eeg Signal Classification." *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 11 (2): 141–44.

Gawronski, Bertram, and Skylar M Brannon. 2019. "What Is Cognitive Consistency, and Why Does It Matter?" 91–116.

Gelman, Andrew. 2006. "Multilevel (Hierarchical) Modeling: What It Can and Cannot Do." *Technometrics* 48 (3): 432–35. https://doi.org/10.1198/004017005000000661.

Gelman, Andrew, Daniel Lee, and Jiqiang Guo. 2015. "Stan: A Probabilistic Programming Language for Bayesian Inference and Optimization." *Journal of Educational and Behavioral Statistics* 40 (5): 530–43. https://doi.org/10.3102/1076998615606113.

Gelman, Andrew, and Donald B Rubin. 1992. "Inference from Iterative Simulation Using Multiple Sequences." *Statist. Sci.* 7 (4): 457–72. https://doi.org/10.1214/ss/1177011136.

Gigerenzer, Gerd. 2004. "Mindless Statistics." *The Journal of Socio-Economics*, Statistical Significance, 33 (5): 587–606. https://doi.org/10.1016/j.socec.2004.09.033.

Gigerenzer, Gerd, and Wolfgang Gaissmaier. 2011. "Heuristic Decision Making." *Annu. Rev. Psychol.* 62 (1): 451–82. https://doi.org/10.1146/annurev-psych-120709-145346.

Gignac, Gilles, and Eva Szodorai. 2016. "Effect Size Guidelines for Individual Differences Researchers." *Personality and Individual Differences* 102 (November): 74–78. https://doi.org/10.1016/j.paid.2016.06.069.

Glickman, Moshe, Rani Moran, and Marius Usher. 2020. "Evidence Integration and Decision-Confidence Are Modulated by Stimulus Consistency." *bioRxiv*, October, 2020.10.12.335943. https://doi.org/10.1101/2020.10.12.335943.

Gold, Joshua I, and Michael N Shadlen. 2001. "Neural Computations That Underlie Decisions About Sensory Stimuli." *Trends in Cognitive Sciences* 5 (1): 10–16.

———. 2007. "The Neural Basis of Decision Making." *Annual Review of Neuroscience* 30.

Goldfarb, Stephanie, Naomi E. Leonard, Patrick Simen, Carlos H. Caicedo-Núñez, and Philip Holmes. 2014. "A Comparative Study of Drift Diffusion and Linear Ballistic Accumulator Models in a Reward Maximization Perceptual Choice Task." *Frontiers in Neuroscience* 8 (August): 148–48. https://doi.org/10.3389/fnins.2014.00148.

Greely, Henry T. 2011. *Mind Reading, Neuroscience, and the Law. A Primer on Criminal Law and Neuroscience.* Oxford University Press. https://oxford.universitypressscholarship.com/view/10.1093/acprof:oso/9780199859177.001.0001/acprof-9780199859177-chapter-5.

Green, David M., and John A. Swets. 1966. *Signal Detection Theory and Psychophysics.* Signal Detection Theory and Psychophysics. Oxford, England: John Wiley.

Green, Peter J., Krzysztof Łatuszyński, Marcelo Pereyra, and Christian P. Robert. 2015. "Bayesian Computation: A Summary of the Current State, and Samples Backwards and Forwards." *Stat Comput* 25 (4): 835–62. https://doi.org/10.1007/s11222-015-9574-5.

Greene, Joshua, and Jonathan Cohen. 2004. "For the Law, Neuroscience Changes Nothing and Everything." *Philos Trans R Soc Lond B Biol Sci* 359 (1451): 1775–85. https://doi.org/10.1098/rstb.2004.1546.

Greenwald, Herbert J. 1969. "Dissonance and Relative Versus Absolute Attractiveness of Decision Alternatives." *Journal of Personality and Social Psychology* 11 (4): 328–33. https://doi.org/10.1037/h0027369.

Griffiths, Thomas L., Falk Lieder, and Noah D. Goodman. 2015. "Rational Use of Cognitive Resources: Levels of Analysis Between the Computational and the Algorithmic." *Top Cogn Sci* 7 (2): 217–29. https://doi.org/10.1111/tops.12142.

Gueorguieva, Ralitza, and John H Krystal. 2004. "Move over Anova: Progress in Analyzing Repeated-Measures Data Andits Reflection in Papers Published in the Archives of General Psychiatry." *Archives of General Psychiatry* 61 (3): 310–17.

Hackenberger, Branimir K. 2019. "Bayes or Not Bayes, Is This the Question?" *Croat Med J* 60 (1): 50–52. https://doi.org/10.3325/cmj.2019.60.50.

Haggard, Patrick. 2008. "Human Volition: Towards a Neuroscience of Will." *Nature Reviews Neuroscience* 9 (12, 12): 934–46. https://doi.org/10.1038/nrn2497.

———. 2017. "Sense of Agency in the Human Brain." *Nat Rev Neurosci* 18

(4): 196–207. https://doi.org/10.1038/nrn.2017.14.

———. 2019. "The Neurocognitive Bases of Human Volition." *Annu Rev Psychol* 70 (January): 9–28. https://doi.org/10.1146/annurev-psych-010418-103348.

Haggard, Patrick, Peter Cartledge, Meilyr Dafydd, and David A. Oakley. 2004. "Anomalous Control: When 'Free-Will' Is Not Conscious." *Conscious Cogn* 13 (3): 646–54. https://doi.org/10.1016/j.concog.2004.06.001.

Haggard, Patrick, Sam Clark, and Jeri Kalogeras. 2002. "Voluntary Action and Conscious Awareness." *Nature Neuroscience* 5 (4, 4): 382–85. https://doi.org/10.1038/nn827.

Hallett, Mark. 2007. "Volitional Control of Movement: The Physiology of Free Will." *Clin Neurophysiol* 118 (6): 1179–92. https://doi.org/10.1016/j.clinph.2007.03.019.

———. 2016. "Physiology of Free Will." *Ann Neurol* 80 (1): 5–12. https://doi.org/10.1002/ana.24657.

Hare, Todd a, Colin F Camerer, Daniel T Knoepfle, and Antonio Rangel. 2010. "Value Computations in Ventral Medial Prefrontal Cortex During Charitable Decision Making Incorporate Input from Regions Involved in Social Cognition." *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience* 30 (2): 583–90. https://doi.org/10.1523/JNEUROSCI.4089-09.2010.

Harmon-Jones, Eddie. 1985. "Cognitive Dissonance and Experienced Negative Affect : Evidence That Dissonance Increases Experienced Negative Affect Even in the Absence of Aversive Consequences," 1490–1501.

Harmon-Jones, Eddie, and Cindy Harmon-Jones. 2019. "Understanding the Motivation Underlying Dissonance Effects," 63–89.

Harmon-Jones, Eddie, Cindy Harmon-Jones, Meghan Fearn, Jonathan D Sigelman, and Peter Johnson. 2008. "Left Frontal Cortical Activation and Spreading of Alternatives: Tests of the Action-Based Model of Dissonance." *Journal of Personality and Social Psychology* 94 (1): 1–15. https://doi.org/10.1037/0022-3514.94.1.1.

Harmon-Jones, Eddie, Cindy Harmon-Jones, and Nicholas Levy. 2015. "An Action-Based Model of Cognitive-Dissonance Processes." *Current Directions in Psychological Science* 24 (3): 184–89. https://doi.org/10.1177/0963721414566449.

Harmon-Jones, Eddie, and Judson Mills. 2019. "An Introduction to Cognitive Dissonance Theory and an Overview of Current Perspectives on the Theory," 3–24.

Harris, Sam. 2012. *Free Will*. New York: Free Press.

Haufe, Stefan, Frank Meinecke, Kai Görgen, Sven Dähne, John-Dylan Haynes, Benjamin Blankertz, and Felix Bießmann. 2014. "On the Interpretation of Weight Vectors of Linear Models in Multivariate Neuroimaging." *Neuroimage* 87: 96–110.

Hawkins, Guy E, Eric-Jan Wagenmakers, Roger Ratcliff, and Scott D Brown. 2015. "Discriminating Evidence Accumulation from Urgency Signals in Speeded Decision Making." American Physiological Society Bethesda, MD.

Hayden, Benjamin Y., and Patrick Haggard. 2017. "Neuroscience: Decision, Insight and Intention." *Current Biology* 27 (15): R750–R753. https://doi.

org/10.1016/j.cub.2017.06.065.

Haynes, John-Dylan. 2019. "Thinking Generatively: Why Do We Use Atheoretical Statistical Models to Test Substantive Psychological Theories? | Computational Psychology." September 5, 2019. http://haines-lab.com/post/thinking-generatively-why-do-we-use-atheoretical-statistical-models-to-test-substantive-psychological-theories/.

Heathcote, Andrew, Yi-Shin Lin, Angus Reynolds, Luke Strickland, Matthew Gretton, and Dora Matzke. 2019. "Dynamic Models of Choice." *Behav Res* 51 (2): 961–85. https://doi.org/10.3758/s13428-018-1067-y.

Heekeren, Hauke R, Sean Marrett, and Leslie G Ungerleider. 2008. "The Neural Systems That Mediate Human Perceptual Decision Making." *Nature Reviews Neuroscience* 9 (6): 467–79. https://doi.org/10.1038/nrn2374.

Henkel, Linda A., and Mara Mather. 2007. "Memory Attributions for Choices: How Beliefs Shape Our Memories." *Journal of Memory and Language* 57 (2): 163–76. https://doi.org/10.1016/j.jml.2006.08.012.

Henley, Amy J., Jason M. Hirst, Florence D. DiGennaro Reed, Amel Becirevic, and Derek D. Reed. 2016. "Function-Altering Effects of Rule Phrasing in the Modulation of Instructional Control." *Anal Verbal Behav* 33 (1): 24–40. https://doi.org/10.1007/s40616-016-0063-5.

Herz, Damian M., Baltazar A. Zavala, Rafal Bogacz, and Peter Brown. 2016. "Neural Correlates of Decision Thresholds in the Human Subthalamic Nucleus." *Curr Biol* 26 (7): 916–20. https://doi.org/10.1016/j.cub.2016.01.051.

Ho, Tiffany C, Scott Brown, and John T Serences. 2009. "Domain General Mechanisms of Perceptual Decision Making in Human Cortex." *The Journal*

of Neuroscience : The Official Journal of the Society for Neuroscience 29 (27): 8675–87. https://doi.org/10.1523/JNEUROSCI.5984-08.2009.

Hofmann, David A. 1997. "An Overview of the Logic and Rationale of Hierarchical Linear Models." *Journal of Management* 23 (6): 723–44. https://doi.org/10.1177/014920639702300602.

Hollander, Gilles de, Birte U. Forstmann, and Scott D. Brown. 2016. "Different Ways of Linking Behavioral and Neural Data via Computational Cognitive Models." *Biol Psychiatry Cogn Neurosci Neuroimaging* 1 (2): 101–9. https://doi.org/10.1016/j.bpsc.2015.11.004.

Holyoak, Keith J., and Dan Simon. 1999. "Bidirectional Reasoning in Decision Making by Constraint Satisfaction." *Journal of Experimental Psychology: General* 128: 3–31.

Hong, Hwanhee, Bradley P. Carlin, Tatyana A. Shamliyan, Jean F. Wyman, Rema Ramakrishnan, François Sainfort, and Robert L. Kane. 2013. "Comparing Bayesian and Frequentist Approaches for Multiple Outcome Mixed Treatment Comparisons." *Med Decis Making* 33 (5): 702–14. https://doi.org/10.1177/0272989X13481110.

Hornsby, Adam N., and Bradley C. Love. 2020. "How Decisions and the Desire for Coherency Shape Subjective Preferences over Time." *Cognition* 200 (July): 104244. https://doi.org/10.1016/j.cognition.2020.104244.

Hunt, Laurence T, Robb B Rutledge, W M Nishantha Malalasekera, Steven W Kennerley, and Raymond J Dolan. 2016. "Approach-Induced Biases in Human Information Sampling." *PLOS Biology* 14 (11): e2000638. https://doi.org/10.1371/journal.pbio.2000638.

Iannaccone, Reto, Tobias U. Hauser, Philipp Staempfli, Susanne Walitza, Daniel Brandeis, and Silvia Brem. 2015. "Conflict Monitoring and Error Processing: New Insights from Simultaneous EEG–fMRI." *NeuroImage* 105 (January): 395–407. https://doi.org/10.1016/j.neuroimage.2014.10.028.

Inwagen, Peter van. 1989. "When Is the Will Free?" *Philosophical Perspectives* 3: 399–422. http://www.jstor.org/stable/2214275.

Izuma, Keise, Madoka Matsumoto, Kou Murayama, Kazuyuki Samejima, Norihiro Sadato, and Kenji Matsumoto. 2010. "Neural Correlates of Cognitive Dissonance and Choice-Induced Preference Change" 107 (51): 22014–9. https://doi.org/10.1073/pnas.1011879108.

Izuma, Keise, and Kou Murayama. 2013. "Choice-Induced Preference Change in the Free-Choice Paradigm : A Critical Methodological Review" 4 (February): 1–12. https://doi.org/10.3389/fpsyg.2013.00041.

Jacobs, Jesse V., Jau-Shin Lou, Jeff A. Kraakevik, and Fay B. Horak. 2009. "The Supplementary Motor Area Contributes to the Timing of the Anticipatory Postural Adjustment During Step Initiation in Participants with and Without Parkinson's Disease." *Neuroscience* 164 (2): 877–85. https://doi.org/10.1016/j.neuroscience.2009.08.002.

Jacoby, Larry L. 1991. "A Process Dissociation Framework: Separating Automatic from Intentional Uses of Memory." *Journal of Memory and Language* 30 (5): 513–41. https://doi.org/10.1016/0749-596X(91)90025-F.

Jahanshahi, M., I. H. Jenkins, R. G. Brown, C. D. Marsden, R. E. Passingham, and D. J. Brooks. 1995. "Self-Initiated Versus Externally Triggered Movements. I. An Investigation Using Measurement of Regional

Cerebral Blood Flow with PET and Movement-Related Potentials in Normal and Parkinson's Disease Subjects." *Brain* 118 (Pt 4) (August): 913–33. https://doi.org/10.1093/brain/118.4.913.

Jahfari, Sara, K Richard Ridderinkhof, and H Steven Scholte. 2013. "Spatial Frequency Information Modulates Response Inhibition and Decision-Making Processes." *PLoS One* 8 (10).

Janak, Patricia H., and Kay M. Tye. 2015. "From Circuits to Behaviour in the Amygdala." *Nature* 517 (7534): 284–92. https://doi.org/10.1038/nature14188.

Jarcho, Johanna M., Elliot T. Berkman, and Matthew D. Lieberman. 2011. "The Neural Basis of Rationalization: Cognitive Dissonance Reduction During Decision-Making." *Social Cognitive and Affective Neuroscience* 6 (4): 460–67. https://doi.org/10.1093/scan/nsq054.

Jia, Ke, Xin Xue, Jong-Hwan Lee, Fang Fang, Jiaxiang Zhang, and Sheng Li. 2018. "Visual Perceptual Learning Modulates Decision Network in the Human Brain: The Evidence from Psychophysics, Modeling, and Functional Magnetic Resonance Imaging." *Journal of Vision* 18 (12): 9–9.

Kaaronen, Roope Oskari, and Rick Dale. 2018. "A Theory of Predictive Dissonance : Predictive Processing Presents a New Take on Cognitive Dissonance" 9 (November): 1–15. https://doi.org/10.3389/fpsyg.2018.02218.

Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. 1991. "Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias." *Journal of Economic Perspectives* 5 (1): 193–206. https://doi.org/10.1257/jep.5.1.193.

Kalckert, Andreas, and H. Henrik Ehrsson. 2014. "The Moving Rubber Hand

Illusion Revisited: Comparing Movements and Visuotactile Stimulation to Induce Illusory Ownership." *Consciousness and Cognition* 26 (May): 117–32. https://doi.org/10.1016/j.concog.2014.02.003.

Karahan, Esin, Alison G Costigan, Kim S Graham, Andrew D Lawrence, and Jiaxiang Zhang. 2019. "Cognitive and White-Matter Compartment Models Reveal Selective Relations Between Corticospinal Tract Microstructure and Simple Reaction Time." *Journal of Neuroscience* 39 (30): 5910–21.

Karch, Susanne, Christoph Mulert, Tobias Thalmeier, Juergen Lutz, Gregor Leicht, Thomas Meindl, Hans-Juergen Moeller, Lorenz Jaeger, and Oliver Pogarell. 2009. "The Free Choice Whether or Not to Respond After Stimulus Presentation." *Hum Brain Mapp* 30 (9): 2971–85. https://doi.org/10.1002/hbm.20722.

Kayser, Jürgen, and Craig E Tenke. 2003. "Optimizing Pca Methodology for Erp Component Identification and Measurement: Theoretical Rationale and Empirical Evaluation." *Clinical Neurophysiology* 114 (12): 2307–25.

Kayser, Stephanie J., and Christoph Kayser. 2018. "Trial by Trial Dependencies in Multisensory Perception and Their Correlates in Dynamic Brain Activity." *Scientific Reports* 8 (1, 1): 3742. https://doi.org/10.1038/s41598-018-22137-8.

Kelly, Desmond. 1973. "Psychosurgery and the Limbic System." *Postgraduate Medical Journal* 49 (578): 825–33.

Kelly, Simon P, and Redmond G O'Connell. 2013. "Internal and External Influences on the Rate of Sensory Evidence Accumulation in the Human Brain." *Journal of Neuroscience* 33 (50): 19434–41.

Kelly, S., and R. G. O'Connell. 2013. "Internal and External Influences on the Rate of Sensory Evidence Accumulation in the Human Brain." *Journal of Neuroscience* 33 (50): 19434–41. https://doi.org/10.1523/jneurosci.3355-13.2013.

Kenworthy, Jared B., Norman Miller, Barry E. Collins, Stephen J. Read, and Mitchell Earleywine. 2011. "A Trans-Paradigm Theoretical Synthesis of Cognitive Dissonance Theory: Illuminating the Nature of Discomfort." *European Review of Social Psychology* 22 (1): 36–113. https://doi.org/10.1080/10463283.2011.580155.

Keuken, Max C., Christa Müller-Axt, Robert Langner, Simon B. Eickhoff, Birte U. Forstmann, and Jane Neumann. 2014. "Brain Networks of Perceptual Decision-Making: An fMRI ALE Meta-Analysis." *Front. Hum. Neurosci.* 8. https://doi.org/10.3389/fnhum.2014.00445.

Khalighinejad, Nima, Steven Di Costa, and Patrick Haggard. 2016. "Endogenous Action Selection Processes in Dorsolateral Prefrontal Cortex Contribute to Sense of Agency: A Meta-Analysis of tDCS Studies of 'Intentional Binding'." *Brain Stimul* 9 (3): 372–79. https://doi.org/10.1016/j.brs.2016.01.005.

Kleiner, Mario, David Brainard, and Denis Pelli. 2007. "What's New in Psychtoolbox-3?"

Klemm, W. R. 2010. "Free Will Debates: Simple Experiments Are Not so Simple." *Adv Cogn Psychol* 6 (August): 47–65. https://doi.org/10.2478/v10053-008-0076-2.

Kloosterman, Niels A, Jan Willem de Gee, Markus Werkle-Bergner, Ulman Lindenberger, Douglas D Garrett, and Johannes Jacobus Fahrenfort. 2019.

"Humans Strategically Shift Decision Bias by Flexibly Adjusting Sensory Evidence Accumulation." Edited by Michael J Frank. *eLife* 8: e37321. https://doi.org/10.7554/eLife.37321.

Kok, Albert. 2001. "On the Utility of P3 Amplitude as a Measure of Processing Capacity." *Psychophysiology* 38 (3): 557–77.

Kolling, Nils, Marco Wittmann, and Matthew FS Rushworth. 2014. "Multiple Neural Mechanisms of Decision Making and Their Competition Under Changing Risk Pressure." *Neuron* 81 (5): 1190–1202.

Koster, Raphael, Emrah Duzel, and Raymond J Dolan. 2015. "Action and Valence Modulate Choice and Choice-Induced Preference Change," 1–10. https://doi.org/10.1371/journal.pone.0119682.

Krajbich, Ian, Dingchao Lu, Colin Camerer, and Antonio Rangel. 2012. "The Attentional Drift-Diffusion Model Extends to Simple Purchasing Decisions." *Frontiers in Psychology* 3: 193.

Krebs, Ruth M, Carsten N Boehler, and Marty G Woldorff. 2010. "The Influence of Reward Associations on Conflict Processing in the Stroop Task." *Cognition* 117 (3): 341–47.

Kriegeskorte, Nikolaus, Marieke Mur, and Peter Bandettini. 2008. "Representational Similarity Analysis - Connecting the Branches of Systems Neuroscience." *Frontiers in Systems Neuroscience* 2: 4. https://doi.org/10.3389/neuro.06.004.2008.

Krieghoff, Veronika, Marcel Brass, Wolfgang Prinz, and Florian Waszak. 2009. "Dissociating What and When of Intentional Actions." *Front Hum Neurosci* 3: 3. https://doi.org/10.3389/neuro.09.003.2009.

Krieghoff, Veronika, Florian Waszak, Wolfgang Prinz, and Marcel Brass. 2011. "Neural and Behavioral Correlates of Intentional Actions." *Neuropsychologia* 49 (5): 767–76. https://doi.org/10.1016/j.neuropsychologia.2011.01.025.

Kruglanski, Arie W., Katarzyna Jasko, Maxim Milyavsky, Marina Chernikova, David Webber, Antonio Pierro, and Daniela di Santo. 2018. "Cognitive Consistency Theory in Social Psychology: A Paradigm Reconsidered." *Psychological Inquiry* 29 (2): 45–59. https://doi.org/10.1080/1047840X.2018.1480619.

Kruschke, John K., and Torrin M. Liddell. 2018. "Bayesian Data Analysis for Newcomers." *Psychon Bull Rev* 25 (1): 155–77. https://doi.org/10.3758/s13423-017-1272-1.

Kuhl, J., and J. Beckmann. 1994. "Volition and Personality: Action Versus State Orientation." undefined. 1994. /paper/Volition-and-Personality%3A-Action-Versus-State-Kuhl-Beckmann/2224fe6778830d5e2bb9a02fabaf7ca0daef99cd.

Kvam, Peter D., and Brandon M. Turner. 2021. "Reconciling Similarity Across Models of Continuous Selections." *Psychological Review* 128 (4): 766–86. https://doi.org/10.1037/rev0000296.

Lau, H. C, R. D Rogers, N Ramnani, and R. E Passingham. 2004. "Willed Action and Attention to the Selection of Action." *NeuroImage* 21 (4): 1407–15. https://doi.org/10.1016/j.neuroimage.2003.10.034.

Lavazza, Andrea. 2016. "Free Will and Neuroscience: From Explaining Freedom Away to New Ways of Operationalizing and Measuring It." *Front. Hum. Neurosci.* 10. https://doi.org/10.3389/fnhum.2016.00262.

Lebovich, Lior, Ran Darshan, Yoni Lavi, David Hansel, and Yonatan Loewenstein. 2019. "Idiosyncratic Choice Bias Naturally Emerges from Intrinsic

Stochasticity in Neuronal Dynamics." *Nature Human Behaviour* 3 (11, 11): 1190–1202. https://doi.org/10.1038/s41562-019-0682-7.

Lebreton, Maël, Shari Langdon, Matthijs J. Slieker, Jip S. Nooitgedacht, Anna E. Goudriaan, Damiaan Denys, Ruth J. van Holst, and Judy Luigjes. 2018. "Two Sides of the Same Coin: Monetary Incentives Concurrently Improve and Bias Confidence Judgments." *Sci Adv* 4 (5): eaaq0668. https://doi.org/10.1126/sciadv.aaq0668.

Lee, Douglas, and Jean Daunizeau. 2020. "Choosing What We Like Vs Liking What We Choose: How Choice-Induced Preference Change Might Actually Be Instrumental to Decision-Making." *PLoS One* 15 (5). https://doi.org/10.1371/journal.pone.0231081.

Leech, Robert, and David J. Sharp. 2014. "The Role of the Posterior Cingulate Cortex in Cognition and Disease." *Brain* 137 (1): 12–32. https://doi.org/10.1093/brain/awt162.

Leeuw, Joshua. 2014. "jsPsych: A JavaScript Library for Creating Behavioral Experiments in a Web Browser." *Behavior Research Methods* 47 (March). https://doi.org/10.3758/s13428-014-0458-y.

Lerche, Veronika, and Andreas Voss. 2018. "Speed–Accuracy Manipulations and Diffusion Modeling: Lack of Discriminant Validity of the Manipulation or of the Parameter Estimates?" *Behavior Research Methods* 50 (6): 2568–85. https://doi.org/10.3758/s13428-018-1034-7.

Lewis, David. 1972. "Psychophysical and Theoretical Identifications." *Australasian Journal of Philosophy* 50 (3): 249–58. https://doi.org/10.1080/00048407212341301.

Li, Yin S., Matthew R. Nassar, Joseph W. Kable, and Joshua I Gold. 2019. "Individual Neurons in the Cingulate Cortex Encode Action Monitoring, Not Selection, During Adaptive Decision-Making." *J. Neurosci.* 39 (34): 6668–83. https://doi.org/10.1523/JNEUROSCI.0159-19.2019.

Libera, Chiara Della, and Leonardo Chelazzi. 2006. "Visual Selective Attention and the Effects of Monetary Rewards." *Psychological Science* 17 (3): 222–27.

Liberman, Nira, and Jens Förster. 2006. "Inferences from Decision Difficulty." *Journal of Experimental Social Psychology* 42 (3): 290–301. https://doi.org/10.1016/j.jesp.2005.04.007.

Libet, B., C. A. Gleason, E. W. Wright, and D. K. Pearl. 1983. "Time of Conscious Intention to Act in Relation to Onset of Cerebral Activity (Readiness-Potential). The Unconscious Initiation of a Freely Voluntary Act." *Brain* 106 (Pt 3) (September): 623–42. https://doi.org/10.1093/brain/106.3.623.

Lieberman, Matthew D, Kevin N Ochsner, Daniel T Gilbert, and Daniel L Schacter. 2001. "Research Article DO AMNESICS EXHIBIT COGNITIVE DISSONANCE REDUCTION ? The Role of Explicit Memory and Attention in Attitude Change" 12 (2): 135–40.

Lilienfeld, Scott O. 2017. "Psychology's Replication Crisis and the Grant Culture: Righting the Ship." *Perspect Psychol Sci* 12 (4): 660–64. https://doi.org/10.1177/1745691616687745.

Linares, Daniel, David Aguilar-Lleyda, and Joan Lopez-Moliner. 2019. "Decoupling Sensory from Decisional Choice Biases in Perceptual Decision Mak-

ing." Edited by Joshua I Gold, Maik C Stüttgen, and Daniel Merfeld. *eLife* 8 (March): e43994. https://doi.org/10.7554/eLife.43994.

Linser, Katrin, and Thomas Goschke. 2007. "Unconscious Modulation of the Conscious Experience of Voluntary Control." *Cognition* 104 (3): 459–75. https://doi.org/10.1016/j.cognition.2006.07.009.

Loughnane, Gerard M, Daniel P Newman, Mark A Bellgrove, Edmund C Lalor, Simon P Kelly, and Redmond G O'Connell. 2016. "Target Selection Signals Influence Perceptual Decisions by Modulating the Onset and Rate of Evidence Accumulation." *Current Biology* 26 (4): 496–502.

Love, Bradley C. 2015. "The Algorithmic Level Is the Bridge Between Computation and Brain." *Topics in Cognitive Science* 7 (2): 230–42. https://doi.org/10.1111/tops.12131.

Love, Bradley C., and Todd M. Gureckis. 2007. "Models in Search of a Brain." *Cognitive, Affective, & Behavioral Neuroscience* 7 (2): 90–108. https://doi.org/10.3758/CABN.7.2.90.

Luce, Mary Frances, John W. Payne, and James R. Bettman. 1999. "Emotional Trade-Off Difficulty and Choice." *Journal of Marketing Research* 36 (2): 143–59. https://doi.org/10.1177/002224379903600201.

Luettgau, Lennart, Claus Tempelmann, Luca Franziska Kaiser, and Gerhard Jocham. 2020. "Decisions Bias Future Choices by Modifying Hippocampal Associative Memories." *Nature Communications* 11 (1, 1): 3318. https://doi.org/10.1038/s41467-020-17192-7.

Lui, Kitty K, Michael D Nunez, Jessica M Cassidy, Joachim Vandekerckhove, Steven C Cramer, and Ramesh Srinivasan. 2018. "Timing of Readiness

Potentials Reflect a Decision-Making Process in the Human Brain." *bioRxiv*, 338806.

Lunn, David J., Andrew Thomas, Nicky Best, and David Spiegelhalter. 2000. "WinBUGS - A Bayesian Modelling Framework: Concepts, Structure, and Extensibility." *Statistics and Computing* 10 (4): 325–37. https://doi.org/10.1023/A:1008929526011.

Luo, Jiayi, and Rongjun Yu. 2016. "The Spreading of Alternatives : Is It the Perceived Choice or Actual Choice That Changes Our Preference ?" https://doi.org/10.1002/bdm.1967.

Luu, Long, and Alan A Stocker. 2018. "Post-Decision Biases Reveal a Self-Consistency Principle in Perceptual Inference," no. 2010: 1–24.

Makowski, Dominique, Mattan S. Ben-Shachar, S. H. Annabel Chen, and Daniel Lüdecke. 2019. "Indices of Effect Existence and Significance in the Bayesian Framework." *Front Psychol* 10: 2767. https://doi.org/10.3389/fpsyg.2019.02767.

Mangun, George R, and Steven A Hillyard. 1991. "Modulations of Sensory-Evoked Brain Potentials Indicate Changes in Perceptual Processing During Visual-Spatial Priming." *Journal of Experimental Psychology: Human Perception and Performance* 17 (4): 1057.

Maoz, Uri, Gideon Yaffe, Christof Koch, and Liad Mudrik. 2019. "Neural Precursors of Decisions That Matter—an Erp Study of Deliberate and Arbitrary Choice." *eLife* 8.

Maris, Eric, and Robert Oostenveld. 2007. "Nonparametric Statistical Testing of Eeg-and Meg-Data." *Journal of Neuroscience Methods* 164 (1):

177–90.

Marken, Richard. 2016. "Intentional and Accidental Behavior: A Control Theory Analysis:" *Psychological Reports*, August. https://doi.org/10.2466/pr0.1982.50.2.647.

Marr, David. 1982. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information.* San Francisco: W.H. Freeman.

Marshall, James AR, Rafal Bogacz, and Iain D Gilchrist. 2012. "Consistent Implementation of Decisions in the Brain." *PloS One* 7 (9).

Mather, Mara, Eldar Shafir, and Marcia K. Johnson. 2003. "Remembering Chosen and Assigned Options." *Memory & Cognition* 31 (3): 422–33. https://doi.org/10.3758/BF03194400.

Mather, M., E. Shafir, and M. K. Johnson. 2000. "Misremembrance of Options Past: Source Monitoring and Choice." *Psychol Sci* 11 (2): 132–38. https://doi.org/10.1111/1467-9280.00228.

Matuschek, Hannes, Reinhold Kliegl, Shravan Vasishth, Harald Baayen, and Douglas Bates. 2017. "Balancing Type I Error and Power in Linear Mixed Models." *Journal of Memory and Language* 94: 305–15.

Matz, David C, Petra M Hofstedt, and Wendy Wood. 2008. "Extraversion as a Moderator of the Cognitive Dissonance Associated with Disagreement" 45: 401–5. https://doi.org/10.1016/j.paid.2008.05.014.

McElreath, Richard. 2016. *Statistical Rethinking: A Bayesian Course with Examples in R and Stan.* 1st edition. Boca Raton: Chapman and Hall/CRC.

McGrath, April. 2017. "Dealing with Dissonance: A Review of Cognitive Dissonance Reduction." *Social and Personality Psychology Compass* 11 (12): e12362. https://doi.org/10.1111/spc3.12362.

Metcalfe, Janet, and Matthew Jason Greene. 2007. "Metacognition of Agency." *J Exp Psychol Gen* 136 (2): 184–99. https://doi.org/10.1037/0096-3445.136.2.184.

Meynen, Gerben. 2016. "Neurolaw: Recognizing Opportunities and Challenges for Psychiatry." *J Psychiatry Neurosci* 41 (1): 3–5. https://doi.org/10.1503/jpn.150317.

Meyniel, Florent, Mariano Sigman, and Zachary F. Mainen. 2015. "Confidence as Bayesian Probability: From Neural Origins to Behavior." *Neuron* 88 (1): 78–92. https://doi.org/10.1016/j.neuron.2015.09.039.

Miendlarzewska, Ewa A, Daphne Bavelier, and Sophie Schwartz. 2016. "Influence of Reward Motivation on Human Declarative Memory." *Neuroscience & Biobehavioral Reviews* 61: 156–76.

Miletic, Steven. 2016. "Neural Evidence for a Role of Urgency in the Speed-Accuracy Trade-Off in Perceptual Decision-Making." *Journal of Neuroscience* 36 (22): 5909–10. https://doi.org/10.1523/JNEUROSCI.0894-16.2016.

Miletić, Steven, Russell J. Boag, and Birte U. Forstmann. 2020. "Mutual Benefits: Combining Reinforcement Learning with Sequential Sampling Models." *Neuropsychologia* 136 (January): 107261. https://doi.org/10.1016/j.neuropsychologia.2019.107261.

Miletić, Steven, and Leendert van Maanen. 2019. "Caution in Decision-Making Under Time Pressure Is Mediated by Timing Ability." *Cognitive*

Psychology 110: 16–29. https://doi.org/10.1016/j.cogpsych.2019.01.002.

Miyagi, Madoka, Makoto Miyatani, and Takashi Nakao. 2017. "Relation Between Choice-Induced Preference Change and Depression," 1–14.

Moore, James W, D. Middleton, P. Haggard, and P. C. Fletcher. 2012. "Exploring Implicit and Explicit Aspects of Sense of Agency." *Consciousness and Cognition* 21 (4): 1748–53. https://doi.org/10.1016/j.concog.2012.10.005.

Moore, James W, and Sukhvinder S. Obhi. 2012. "Intentional Binding and the Sense of Agency: A Review." *Conscious Cogn* 21 (1): 546–61. https://doi.org/10.1016/j.concog.2011.12.002.

Moretto, Giovanna, Petra Schwingenschuh, Petra Katschnig, Kailash P. Bhatia, and Patrick Haggard. 2011. "Delayed Experience of Volition in Gilles de La Tourette Syndrome." *J Neurol Neurosurg Psychiatry* 82 (12): 1324–7. https://doi.org/10.1136/jnnp.2010.221143.

Mulder, Martijn J, Eric-Jan Wagenmakers, Roger Ratcliff, Wouter Boekel, and Birte U Forstmann. 2012. "Bias in the Brain: A Diffusion Model Analysis of Prior Probability and Potential Payoff." *Journal of Neuroscience* 32 (7): 2335–43.

Mulder, MJ, L Van Maanen, and BU Forstmann. 2014. "Perceptual Decision Neurosciences–a Model-Based Review." *Neuroscience* 277: 872–84.

Murphy, Gregory L. 2016. "Is There an Exemplar Theory of Concepts?" *Psychonomic Bulletin & Review* 23 (4): 1035–42. https://doi.org/10.3758/s13423-015-0834-3.

Myung, Jay I., and Mark A. Pitt. 2018. "Model Comparison in Psychology." In *Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience*,

1–34. American Cancer Society. https://doi.org/10.1002/9781119170174.epcn503.

Nachev, Parashkev, and Peter Hacker. 2014. "The Neural Antecedents to Voluntary Action: A Conceptual Analysis." *Cognitive Neuroscience* 5 (3-4): 193–208. https://doi.org/10.1080/17588928.2014.934215.

Nachev, Parashkev, Christopher Kennard, and Masud Husain. 2008. "Functional Role of the Supplementary and Pre-Supplementary Motor Areas." *Nature Reviews. Neuroscience* 9 (11): 856–69. https://doi.org/10.1038/nrn2478.

Nachev, Parashkev, Geraint Rees, Andrew Parton, Christopher Kennard, and Masud Husain. 2005. "Volition and Conflict in Human Medial Frontal Cortex." *Curr Biol* 15 (2): 122–28. https://doi.org/10.1016/j.cub.2005.01.006.

Nahmias, Eddy, Stephen Morris, Thomas Nadelhoffer, and Jason Turner. 2004. "The Phenomenology of Free Will." *Journal of Consciousness Studies* 11 (7-8): 162–79.

———. 2005. "Surveying Freedom: Folk Intuitions About Free Will and Moral Responsibility." *Philosophical Psychology* 18 (5): 561–84. https://doi.org/10.1080/09515080500264180.

Navarro, Danielle. 2020. "If Mathematical Psychology Did Not Exist We Might Need to Invent It: A Comment on Theory Building in Psychology." https://doi.org/10.31234/osf.io/ygbjp.

Nickerson, Raymond S. 1998. "Confirmation Bias : A Ubiquitous Phenomenon in Many Guises" 2 (2): 175–220.

Nieuwenhuis, Sander, Gary Aston-Jones, and Jonathan D Cohen. 2005.

"Decision Making, the P3, and the Locus Coeruleus–Norepinephrine System." *Psychological Bulletin* 131 (4): 510.

Nunez, Michael D, Aishwarya Gosai, Joachim Vandekerckhove, and Ramesh Srinivasan. 2019. "The Latency of a Visual Evoked Potential Tracks the Onset of Decision Making." *Neuroimage* 197: 93–108.

Nunez, Michael D, Joachim Vandekerckhove, and Ramesh Srinivasan. 2017. "How Attention Influences Perceptual Decision Making: Single-Trial Eeg Correlates of Drift-Diffusion Model Parameters." *Journal of Mathematical Psychology* 76: 117–30.

Ojeda, Andrés, Robin A Murphy, and Alex Kacelnik. 2018. "Paradoxical Choice in Rats: Subjective Valuation and Mechanism of Choice." *Behavioural Processes* 152: 73–80.

Orban de Xivry, Jean-Jacques, and Philippe Lefèvre. 2016. "A Switching Cost for Motor Planning." *J Neurophysiol* 116 (6): 2857–68. https://doi.org/10.1152/jn.00319.2016.

Osth, Adam F, Simon Dennis, and Andrew Heathcote. 2017. "Likelihood Ratio Sequential Sampling Models of Recognition Memory." *Cognitive Psychology* 92: 101–26.

Oud, Bastiaan, Ian Krajbich, Kevin Miller, Jin Hyun Cheong, Matthew Botvinick, and Ernst Fehr. 2016. "Irrational Time Allocation in Decision-Making." *Proceedings of the Royal Society B: Biological Sciences* 283 (1822): 20151439.

Pais, Darren, Patrick M Hogan, Thomas Schlegel, Nigel R Franks, Naomi E Leonard, and James A R Marshall. 2013. "A Mechanism for Value-Sensitive

Decision-Making." *PloS One* 8 (9): e73216–e73216. https://doi.org/10.1371/journal.pone.0073216.

Palminteri, Stefano, Germain Lefebvre, Emma J. Kilford, and Sarah-Jayne Blakemore. 2017. "Confirmation Bias in Human Reinforcement Learning: Evidence from Counterfactual Feedback Processing." *PLOS Computational Biology* 13 (8): e1005684. https://doi.org/10.1371/journal.pcbi.1005684.

Parra, Lucas C, Clay D Spence, Adam D Gerson, and Paul Sajda. 2005. "Recipes for the Linear Analysis of Eeg." *Neuroimage* 28 (2): 326–41.

Payne, John W., James R. Bettman, and Eric J. Johnson. 1993. *The Adaptive Decision Maker*. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9781139173933.

Pearson, John M., Sarah R. Heilbronner, David L. Barack, Benjamin Y. Hayden, and Michael L. Platt. 2011. "Posterior Cingulate Cortex: Adapting Behavior to a Changing World." *Trends Cogn Sci* 15 (4): 143–51. https://doi.org/10.1016/j.tics.2011.02.002.

Peirce, Jonathan W. 2007. "PsychoPy-Psychophysics Software in Python." *Journal of Neuroscience Methods* 162 (1-2): 8–13. https://doi.org/10.1016/j.jneumeth.2006.11.017.

Peters, Uwe. 2020. "What Is the Function of Confirmation Bias?" *Erkenn*, April. https://doi.org/10.1007/s10670-020-00252-1.

Philiastides, Marios G., and Roger Ratcliff. 2013. "Influence of Branding on Preference-Based Decision Making." *Psychol Sci* 24 (7): 1208–15. https://doi.org/10.1177/0956797612470701.

Phillips, HN, TE Cope, LE Hughes, Jiaxiang Zhang, and JB Rowe. 2018. "Monitoring the Past and Choosing the Future: The Prefrontal Cortical Influences on Voluntary Action." *Scientific Reports* 8 (1): 1–13.

Piray, Payam, Amir Dezfouli, Tom Heskes, Michael J. Frank, and Nathaniel D. Daw. 2019. "Hierarchical Bayesian Inference for Concurrent Model Fitting and Comparison for Group Studies." *PLOS Computational Biology* 15 (6): e1007043. https://doi.org/10.1371/journal.pcbi.1007043.

Pirrone, Angelo, Habiba Azab, Benjamin Y Hayden, Tom Stafford, and James AR Marshall. 2018a. "Evidence for the Speed–Value Trade-Off: Human and Monkey Decision Making Is Magnitude Sensitive." *Decision* 5 (2): 129.

Pirrone, Angelo, Habiba Azab, Benjamin Y Hayden, Tom Stafford, and James A R Marshall. 2018b. "Evidence for the Speed–Value Trade-Off: Human and Monkey Decision Making Is Magnitude Sensitive." *Decision* 5 (2): 129–42. https://doi.org/10.1037/dec0000075.

Pirrone, Angelo, Tom Stafford, and James AR Marshall. 2014a. "When Natural Selection Should Optimize Speed-Accuracy Trade-Offs." *Frontiers in Neuroscience* 8: 73.

Pirrone, Angelo, Tom Stafford, and James A. R. Marshall. 2014b. "When Natural Selection Should Optimize Speed-Accuracy Trade-Offs." *Frontiers in Neuroscience* 8 (April): 1–5. https://doi.org/10.3389/fnins.2014.00073.

Pisauro, M Andrea, Elsa Fouragnan, Chris Retzler, and Marios G Philiastides. 2017. "Neural Correlates of Evidence Accumulation During Value-Based Decisions Revealed via Simultaneous Eeg-fMRI." *Nature Communications* 8 (1): 1–9.

Plant, Richard R., Nick Hammond, and Tom Whitehouse. 2003. "How Choice of Mouse May Affect Response Timing in Psychological Studies." *Behav Res Methods Instrum Comput* 35 (2): 276–84. https://doi.org/10.3758/bf03202553.

Polania, Rafael, Ian Krajbich, Marcus Grueschow, and Christian C Ruff. 2014. "Neural Oscillations and Synchronization Differentially Support Evidence Accumulation in Perceptual and Value-Based Decision Making." *Neuron* 82 (3): 709–20.

Polania, R, Ian Krajbich, Marcus Grueschow, and Christian C. Ruff. 2014. "Neural Oscillations and Synchronization Differentially Support Evidence Accumulation in Perceptual and Value-Based Decision Making." *Neuron* 82 (3): 709–20. https://doi.org/10.1016/j.neuron.2014.03.014.

Poldrack, Russell A. 2006. "Can Cognitive Processes Be Inferred from Neuroimaging Data?" *Trends in Cognitive Sciences* 10 (2): 59–63.

———. 2011. "Inferring Mental States from Neuroimaging Data: From Reverse Inference to Large-Scale Decoding." *Neuron* 72 (5): 692–97. https://doi.org/10.1016/j.neuron.2011.11.001.

Polich, John, Patricia Crane Ellerson, and Jill Cohen. 1996. "P300, Stimulus Intensity, Modality, and Probability." *International Journal of Psychophysiology* 23 (1-2): 55–62.

Proctor, Robert W., and Darryl W. Schneider. 2018. "Hick's Law for Choice Reaction Time: A Review." *Q J Exp Psychol (Hove)* 71 (6): 1281–99. https://doi.org/10.1080/17470218.2017.1322622.

Rae, Charlotte L., Laura E. Hughes, Chelan Weaver, Michael C. Anderson,

and James B. Rowe. 2014. "Selection and Stopping in Voluntary Action: A Meta-Analysis and Combined fMRI Study." *Neuroimage* 86 (February): 381–91. https://doi.org/10.1016/j.neuroimage.2013.10.012.

Rajmohan, V., and E. Mohandas. 2007. "The Limbic System." *Indian Journal of Psychiatry* 49 (2): 132–39. https://doi.org/10.4103/0019-5545.33264.

Rangel, Antonio, Colin Camerer, and P. Read Montague. 2008. "A Framework for Studying the Neurobiology of Value-Based Decision Making." *Nature Reviews Neuroscience* 9 (7, 7): 545–56. https://doi.org/10.1038/nrn2357.

Rassin, Eric. 2008. "Individual Differences in the Susceptibility to Confirmation Bias."

Ratcliff, Roger. 1978. "A Theory of Memory Retrieval." *Psychological Review* 85 (2): 59–108. https://doi.org/10.1037/0033-295X.85.2.59.

Ratcliff, Roger, and Michael J Frank. 2012. "Reinforcement-Based Decision Making in Corticostriatal Circuits: Mutual Constraints by Neurocomputational and Diffusion Models." *Neural Computation* 24 (5): 1186–1229.

Ratcliff, Roger, and Gail McKoon. 2008. "The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks." *Neural Computation* 20 (4): 873–922.

———. 2020. "Decision Making in Numeracy Tasks with Spatially Continuous Scales." *Cognitive Psychology* 116. https://doi.org/10.1016/j.cogpsych.2019.101259.

Ratcliff, Roger, and Philip L Smith. 2004. "A Comparison of Sequential

Sampling Models for Two-Choice Reaction Time." *Psychological Review* 111 (2): 333–67. https://doi.org/10.1037/0033-295X.111.2.333.

Ratcliff, Roger, Chelsea Voskuilen, and Andrei R Teodorescu. 2018. "Modeling 2-Alternative Forced-Choice Tasks: Accounting for Both Magnitude and Difference Effects." *Cogn Psychol* 103 (June): 1–22. https://doi.org/10.1016/j.cogpsych.2018.02.002.

Ratcliff, R., M. G. Philiastides, and P. Sajda. 2009. "Quality of Evidence for Perceptual Decision Making Is Indexed by Trial-to-Trial Variability of the EEG." *Proceedings of the National Academy of Sciences* 106 (16): 6539–44. https://doi.org/10.1073/pnas.0812589106.

Ravenzwaaij, Don van, Scott D. Brown, A. a. J. Marley, and Andrew Heathcote. 2020. "Accumulating Advantages: A New Conceptualization of Rapid Multiple Choice." *Psychol Rev* 127 (2): 186–215. https://doi.org/10.1037/rev0000166.

Raymond, Jane E, and Jennifer L O'Brien. 2009. "Selective Visual Attention and Motivation: The Consequences of Value Learning in an Attentional Blink Task." *Psychological Science* 20 (8): 981–88.

R Core Team. 2018. *R: A Language and Environment for Statistical Computing*. Manual. Vienna, Austria: R Foundation for Statistical Computing. https://www.r-project.org/.

Rescorla, R. 1972. "A Theory of Pavlovian Conditioning : Variations in the Effectiveness of Reinforcement and Nonreinforcement." undefined. 1972. /paper/A-theory-of-Pavlovian-conditioning-%3A-Variations-in-Rescorla/afaf65883ff75cc19926f61f181a687927789ad1.

Richardson, Benjamin, Roland Pfister, and Lisa R. Fournier. 2020. "Free-Choice and Forced-Choice Actions: Shared Representations and Conservation of Cognitive Effort." *Atten Percept Psychophys* 82 (5): 2516–30. https://doi.org/10.3758/s13414-020-01986-4.

Roitman, Jamie D, and Michael N Shadlen. 2002. "Response of Neurons in the Lateral Intraparietal Area During a Combined Visual Discrimination Reaction Time Task." *Journal of Neuroscience* 22 (21): 9475–89.

Romero, Felipe. 2019. "Philosophy of Science and the Replicability Crisis." *Philosophy Compass* 14 (11): e12633. https://doi.org/10.1111/phc3.12633.

Romeu, Ricardo J., Nathaniel Haines, Woo-Young Ahn, Jerome R. Busemeyer, and Jasmin Vassileva. 2020. "A Computational Model of the Cambridge Gambling Task with Applications to Substance Use Disorders." *Drug and Alcohol Dependence* 206 (January): 107711. https://doi.org/10.1016/j.drugalcdep.2019.107711.

Rooij, Iris van, and Giosue Baggio. 2021. "Theory Before the Test: How to Build High-Verisimilitude Explanatory Theories in Psychological Science." *Perspectives on Psychological Science* 0 (0): 1745691620970604. https://doi.org/10.1177/1745691620970604.

Roskies, Adina. 2006. "Neuroscientific Challenges to Free Will and Responsibility." *Trends Cogn Sci* 10 (9): 419–23. https://doi.org/10.1016/j.tics.2006.07.011.

Rouder, Jeffrey N., and Jun Lu. 2005. "An Introduction to Bayesian Hierarchical Models with an Application in the Theory of Signal Detection." *Psychonomic Bulletin & Review* 12 (4): 573–604. https://doi.org/10.3758/

BF03196750.

Rouder, Jeffrey N., and Roger Ratcliff. 2004. "Comparing Categorization Models." *Journal of Experimental Psychology. General* 133 (1): 63. https://doi.org/10.1037/0096-3445.133.1.63.

Rowe, James B., Klaas E. Stephan, Karl Friston, Richard S. J. Frackowiak, and Richard E. Passingham. 2005. "The Prefrontal Cortex Shows Context-Specific Changes in Effective Connectivity to Motor or Visual Cortex During the Selection of Action or Colour." *Cereb Cortex* 15 (1): 85–95. https://doi.org/10.1093/cercor/bhh111.

Rowe, J. B., L. Hughes, and I. Nimmo-Smith. 2010. "Action Selection: A Race Model for Selected and Non-Selected Actions Distinguishes the Contribution of Premotor and Prefrontal Areas." *Neuroimage* 51 (2): 888–96. https://doi.org/10.1016/j.neuroimage.2010.02.045.

Rowe, J., L. Hughes, D. Eckstein, and A. M. Owen. 2008. "Rule-Selection and Action-Selection Have a Shared Neuroanatomical Basis in the Human Prefrontal and Parietal Cortex." *Cereb Cortex* 18 (10): 2275–85. https://doi.org/10.1093/cercor/bhm249.

Rustichini, Aldo. 2009. "Neuroeconomics:: Formal Models of Decision Making and Cognitive Neuroscience." In *Neuroeconomics*, 33–46. Elsevier.

Saito, Naho, Keisuke Takahata, Hodaka Yamakado, Nobukatsu Sawamoto, Satoshi Saito, Ryosuke Takahashi, Toshiya Murai, and Hidehiko Takahashi. 2017. "Altered Awareness of Action in Parkinson's Disease: Evaluations by Explicit and Implicit Measures." *Scientific Reports* 7 (1, 1): 8019. https://doi.org/10.1038/s41598-017-08482-0.

Salti, Moti, Imen El Karoui, Mathurin Maillet, and Lionel Naccache. 2014. "Cognitive Dissonance Resolution Is Related to Episodic Memory" 9 (9): 1–8. https://doi.org/10.1371/journal.pone.0108579.

Salvatier, John, Thomas V. Wiecki, and Christopher Fonnesbeck. 2016. "Probabilistic Programming in Python Using PyMC3." *PeerJ Computer Science* 2 (April): e55. https://doi.org/10.7717/peerj-cs.55.

Sarkissian, Hagop, Amita Chatterjee, Felipe De Brigard, Joshua Knobe, Shaun Nichols, and Smita Sirker. 2010. "Is Belief in Free Will a Cultural Universal?" *Mind & Language* 25 (3): 346–58. https://doi.org/10.1111/j.1468-0017.2010.01393.x.

Sato, Atsushi, and Asako Yasuda. 2005. "Illusion of Sense of Self-Agency: Discrepancy Between the Predicted and Actual Sensory Consequences of Actions Modulates the Sense of Self-Agency, but Not the Sense of Self-Ownership." *Cognition* 94 (3): 241–55. https://doi.org/10.1016/j.cognition.2004.04.003.

Schad, Daniel J., Michael Betancourt, and Shravan Vasishth. 2020. "Toward a Principled Bayesian Workflow in Cognitive Science." *Psychological Methods*, no. 2013. https://doi.org/10.1037/met0000275.

Schiebener, Johannes, and Matthias Brand. 2015. "Decision Making Under Objective Risk Conditions–a Review of Cognitive and Emotional Correlates, Strategies, Feedback Processing, and External Influences." *Neuropsychology Review* 25 (2): 171–98.

Schlegel, Alexander, Prescott Alexander, Walter Sinnott-Armstrong, Adina Roskies, Peter Ulric Tse, and Thalia Wheatley. 2015. "Hypnotizing Libet:

Readiness Potentials with Non-Conscious Volition." *Conscious Cogn* 33 (May): 196–203. https://doi.org/10.1016/j.concog.2015.01.002.

Schmiedek, Florian, Klaus Oberauer, Oliver Wilhelm, Heinz-Martin Süß, and Werner W Wittmann. 2007. "Individual Differences in Components of Reaction Time Distributions and Their Relations to Working Memory and Intelligence." *Journal of Experimental Psychology: General* 136 (3): 414.

Schmolesky, Matthew T, Youngchang Wang, Doug P Hanes, Kirk G Thompson, Stefan Leutgeb, Jeffrey D Schall, and Audie G Leventhal. 1998. "Signal Timing Across the Macaque Visual System." *Journal of Neurophysiology* 79 (6): 3272–8.

Scholl, Jacqueline, Nils Kolling, Natalie Nelissen, Marco K Wittmann, Catherine J Harmer, and Matthew FS Rushworth. 2015. "The Good, the Bad, and the Irrelevant: Neural Mechanisms of Learning Real and Hypothetical Rewards and Effort." *Journal of Neuroscience* 35 (32): 11233–51.

Schurger, Aaron, Myrto Mylopoulos, and David Rosenthal. 2016. "Neural Antecedents of Spontaneous Voluntary Movement: A New Perspective." *Trends in Cognitive Sciences* 20 (2): 77–79. https://doi.org/10.1016/j.tics.2015.11.003.

Schurman, DL, and JP Belcher. 1974. "Noncontingent Reward Magnitude Effects on Reaction Time: A Replication and Extension." *Bulletin of the Psychonomic Society* 3 (2): 104–6.

Senftleben, Ulrike, Martin Schoemann, Diana Schwenke, Sarah Richter, Maja Dshemuchadse, and Stefan Scherbaum. 2019. "Choice Perseveration in Value-Based Decision Making: The Impact of Inter-Trial Interval and Mood."

Acta Psychol (Amst) 198 (July): 102876. https://doi.org/10.1016/j.actpsy.2019.102876.

Shadlen, Michael N, and Daphna Shohamy. 2016. "Decision Making and Sequential Sampling from Memory." *Neuron* 90 (5): 927–39.

Sharot, Tali, Benedetto De Martino, and Raymond J. Dolan. 2009. "How Choice Reveals and Shapes Expected Hedonic Outcome." *J Neurosci* 29 (12): 3760–5. https://doi.org/10.1523/JNEUROSCI.4972-08.2009.

Sharot, Tali, Stephen M. Fleming, Xiaoyu Yu, Raphael Koster, and Raymond J. Dolan. 2012. "Is Choice-Induced Preference Change Long Lasting?" *Psychol Sci* 23 (10): 1123–9. https://doi.org/10.1177/0956797612438733.

Sharot, Tali, Cristina M Velasquez, and Raymond J Dolan. 2010. "Do Decisions Shape Preference? Evidence from Blind Choice." *Psychological Science* 21 (9): 1231–5. https://doi.org/10.1177/0956797610379235.

Shenhav, Amitai, MatthewM Botvinick, and JonathanD Cohen. 2013. "The Expected Value of Control: An Integrative Theory of Anterior Cingulate Cortex Function." *Neuron* 79 (2): 217–40. https://doi.org/10.1016/j.neuron.2013.07.007.

Si, Ruoguang, James B. Rowe, and Jiaxiang Zhang. 2020. "Functional Localization and Categorization of Intentional Decisions in Humans: A Meta-Analysis of Brain Imaging Studies." *bioRxiv*, November, 2020.11.27.401208. https://doi.org/10.1101/2020.11.27.401208.

Sidarus, Nura, and Patrick Haggard. 2016. "Difficult Action Decisions Reduce the Sense of Agency: A Study Using the Eriksen Flanker Task." *Acta Psychol (Amst)* 166 (May): 1–11. https://doi.org/10.1016/j.actpsy.2016.03.003.

Sidarus, Nura, Matti Vuorre, and Patrick Haggard. 2017. "Integrating Prospective and Retrospective Cues to the Sense of Agency: A Multi-Study Investigation." *Neurosci Conscious* 2017 (1): nix012. https://doi.org/10.1093/nc/nix012.

Simen, Patrick. 2012. "Evidence Accumulator or Decision Threshold – Which Cortical Mechanism Are We Observing?" *Frontiers in Psychology* 3: 183. https://doi.org/10.3389/fpsyg.2012.00183.

Simon, Dan, Daniel C. Krawczyk, Airom Bleicher, and Keith J. Holyoak. 2008. "The Transience of Constructed Preferences." *Journal of Behavioral Decision Making* 21 (1): 1–14. https://doi.org/10.1002/bdm.575.

Simon, Dan, Daniel C Krawczyk, and Keith J Holyoak. 2004. "Construction of Prefrences by Constraint Satisfaction." *Psychological Science* 15 (5): 331–6. https://doi.org/10.1111/j.0956-7976.2004.00678.x.

Skagerlund, Kenny, Thomas Karlsson, and Ulf Traaff. 2016. "Magnitude Processing in the Brain: An fMRI Study of Time, Space, and Numerosity as a Shared Cortical System." *Front Hum Neurosci* 10: 500. https://doi.org/10.3389/fnhum.2016.00500.

Smid, Sanne C., Daniel McNeish, Milica Miočevi'c, and Rens van de Schoot. 2020. "Bayesian Versus Frequentist Estimation for Structural Equation Models in Small Sample Contexts: A Systematic Review." *Structural Equation Modeling: A Multidisciplinary Journal* 27 (1): 131–61. https://doi.org/10.1080/10705511.2019.1577140.

Smith, Philip L, and Roger Ratcliff. 2004. "Psychology and Neurobiology of Simple Decisions." *Trends in Neurosciences* 27 (3): 161–8. https://doi.org/

10.1016/j.tins.2004.01.006.

Smith, Philip L., Saam Saber, Elaine A. Corbett, and Simon D. Lilburn. 2020. "Modeling Continuous Outcome Color Decisions with the Circular Diffusion Model: Metric and Categorical Properties." *Psychological Review* 127 (4): 562–90. https://doi.org/10.1037/rev0000185.

Soares, José M., Ricardo Magalhães, Pedro S. Moreira, Alexandre Sousa, Edward Ganz, Adriana Sampaio, Victor Alves, Paulo Marques, and Nuno Sousa. 2016. "A Hitchhiker's Guide to Functional Magnetic Resonance Imaging." *Front Neurosci* 10 (November). https://doi.org/10.3389/fnins.2016.00515.

Soto, Christopher J, and Oliver P John. 2017. "Short and Extra-Short Forms of the Big Five Inventory – 2 : The BFI-2-S and." *Journal of Research in Personality* 68: 69–81. https://doi.org/10.1016/j.jrp.2017.02.004.

Spektor, Mikhail S., Sudeep Bhatia, and Sebastian Gluth. 2021. "The Elusiveness of Context Effects in Decision Making." *Trends in Cognitive Sciences*, September. https://doi.org/10.1016/j.tics.2021.07.011.

Sperduti, Marco, Pauline Delaveau, Philippe Fossati, and Jaqueline Nadel. 2011. "Different Brain Structures Related to Self- and External-Agency Attribution: A Brief Review and Meta-Analysis." *Brain Struct Funct* 216 (2): 151–57. https://doi.org/10.1007/s00429-010-0298-1.

Spreng, R. Nathan, W. Dale Stevens, Jon P. Chamberlain, Adrian W. Gilmore, and Daniel L. Schacter. 2010. "Default Network Activity, Coupled with the Frontoparietal Control Network, Supports Goal-Directed Cognition." *Neuroimage* 53 (1): 303–17. https://doi.org/10.1016/j.neuroimage.2010.06.

016.

Starns, Jeffrey J, and Roger Ratcliff. 2014. "Validating the Unequal-Variance Assumption in Recognition Memory Using Response Time Distributions Instead of Roc Functions: A Diffusion Model Analysis." *Journal of Memory and Language* 70: 36–52.

Stelzer, Johannes, Yi Chen, and Robert Turner. 2013. "Statistical Inference and Multiple Testing Correction in Classification-Based Multi-Voxel Pattern Analysis (MVPA): Random Permutations and Cluster Size Control." *NeuroImage* 65: 69–82. https://doi.org/10.1016/j.neuroimage.2012.09.063.

Stephens, David W., and Dack Anderson. 2001. "The Adaptive Value of Preference for Immediacy: When Shortsighted Rules Have Farsighted Consequences." *Behavioral Ecology* 12 (3): 330–39. https://doi.org/10.1093/beheco/12.3.330.

Steyvers, Mark, Michael D. Lee, and Eric Jan Wagenmakers. 2009. "A Bayesian Analysis of Human Decision-Making on Bandit Problems." *Journal of Mathematical Psychology* 53 (3): 168–79. https://doi.org/10.1016/j.jmp.2008.11.002.

Strube, Gerhard. 2000. "Generative Theories in Cognitive Psychology." *Theory & Psychology* 10 (1): 117–25. https://doi.org/10.1177/0959354300010001606.

Stryhn, Henrik, and Jette Christensen. 2014. "The Analysis–Hierarchical Models: Past, Present and Future." *Prev Vet Med* 113 (3): 304–12. https://doi.org/10.1016/j.prevetmed.2013.10.001.

Summerfield, Christopher, and Floris P. de Lange. 2014. "Expectation

in Perceptual Decision Making: Neural and Computational Mechanisms." *Nature Reviews Neuroscience* 15 (11, 11): 745–56. https://doi.org/10.1038/nrn3838.

Summerfield, Christopher, and Konstantinos Tsetsos. 2012. "Building Bridges Between Perceptual and Economic Decision-Making: Neural and Computational Mechanisms." *Frontiers in Neuroscience* 6: 70.

Summers, Robert S. 1969. "Review of Punishment and Responsibility." *The University of Toronto Law Journal* 19 (4): 642–53. https://doi.org/10.2307/824999.

Sumner, Petroc, Parashkev Nachev, Peter Morris, Andrew M. Peters, Stephen R. Jackson, Christopher Kennard, and Masud Husain. 2007. "Human Medial Frontal Cortex Mediates Unconscious Inhibition of Voluntary Action." *Neuron* 54 (5): 697–711. https://doi.org/10.1016/j.neuron.2007.05.016.

Supekar, Kaustubh, Hyesang Chang, Percy K. Mistry, Teresa Iuculano, and Vinod Menon. 2021. "Neurocognitive Modeling of Latent Memory Processes Reveals Reorganization of Hippocampal-Cortical Circuits Underlying Learning and Efficient Strategies." *Communications Biology* 4 (1): 405. https://doi.org/10.1038/s42003-021-01872-1.

Sutton, Richard S., and Francis Barto. 1998. *Reinforcement Learning: An Introduction*. Second edition. Cambridge, Mass: MIT Press.

Synofzik, Matthis, Gottfried Vosgerau, and Albert Newen. 2008. "Beyond the Comparator Model: A Multifactorial Two-Step Account of Agency." *Conscious Cogn* 17 (1): 219–39. https://doi.org/10.1016/j.concog.2007.03.010.

Szollosi, Aba, David Kellen, Danielle J. Navarro, Richard Shiffrin, Iris van Rooij, Trisha Van Zandt, and Chris Donkin. 2020. "Is Preregistration Worthwhile?" *Trends Cogn Sci* 24 (2): 94–95. https://doi.org/10.1016/j.tics.2019.11.009.

Tabu, Hayato, Toshihiko Aso, Masao Matsuhashi, Yoshino Ueki, Ryosuke Takahashi, Hidenao Fukuyama, Hiroshi Shibasaki, and Tatsuya Mima. 2015. "Parkinson's Disease Patients Showed Delayed Awareness of Motor Intention." *Neurosci Res* 95 (June): 74–77. https://doi.org/10.1016/j.neures.2015.01.012.

Tackett, Jennifer L., Cassandra M. Brandes, Kevin M. King, and Kristian E. Markon. 2019. "Psychology's Replication Crisis and Clinical Psychological Science." *Annu. Rev. Clin. Psychol.* 15 (1): 579–604. https://doi.org/10.1146/annurev-clinpsy-050718-095710.

Talluri, Bharath Chandra, Anne E. Urai, Konstantinos Tsetsos, Marius Usher, and Tobias H. Donner. 2018. "Confirmation Bias Through Selective Overweighting of Choice-Consistent Evidence." *Current Biology* 28 (19): 3128–3135.e8. https://doi.org/10.1016/j.cub.2018.07.052.

Tavares, Gabriela, Pietro Perona, and Antonio Rangel. 2017. "The Attentional Drift Diffusion Model of Simple Perceptual Decision-Making." *Front. Neurosci.* 11. https://doi.org/10.3389/fnins.2017.00468.

Taylor, Paul C. J., Matthew F. S. Rushworth, and Anna C. Nobre. 2008. "Choosing Where to Attend and the Medial Frontal Cortex: An FMRI Study." *J Neurophysiol* 100 (3): 1397–1406. https://doi.org/10.1152/jn.90241.2008.

Teodorescu, Andrei R, Rani Moran, and Marius Usher. 2016. "Absolutely Relative or Relatively Absolute: Violations of Value Invariance in Human

Decision Making." *Psychonomic Bulletin and Review* 23 (1): 22–38. https://doi.org/10.3758/s13423-015-0858-8.

Thaler, Richard H. 1991. "'Some Empirical Evidence on Dynamic Inconsistency." *Quasi Rational Economics* 1: 127–36.

Thimm, Markus, Ralph Weidner, Gereon R. Fink, and Walter Sturm. 2012. "Neural Mechanisms Underlying Freedom to Choose an Object." *Human Brain Mapping* 33 (11): 2686–93. https://doi.org/10.1002/hbm.21393.

Thomas, Julie, Giovanna Vanni-Mercier, and Jean-Claude Dreher. 2013. "Neural Dynamics of Reward Probability Coding: A Magnetoencephalographic Study in Humans." *Frontiers in Neuroscience* 7: 214.

Thompson, Edmund R. 2007. "DEVELOPMENT AND VALIDATION OF AN INTERNATIONALLY RELIABLE SHORT-FORM OF THE POSITIVE AND NEGATIVE AFFECT SCHEDULE (PANAS)" 38 (2): 227–42. https://doi.org/10.1177/0022022106297301.

Thura, David, Julie Beauregard-Racine, Charles-William Fradet, and Paul Cisek. 2012. "Decision Making by Urgency Gating: Theory and Experimental Support." *Journal of Neurophysiology* 108 (11): 2912–30.

Thura, David, and Paul Cisek. 2017. "The Basal Ganglia Do Not Select Reach Targets but Control the Urgency of Commitment." *Neuron* 95 (5): 1160–70.

Tillman, Gabriel, Adam F Osth, Don van Ravenzwaaij, and Andrew Heathcote. 2017. "A Diffusion Decision Model Analysis of Evidence Variability in the Lexical Decision Task." *Psychonomic Bulletin & Review* 24 (6): 1949–56.

Tobler, Philippe N, John P O'Doherty, Raymond J Dolan, and Wolfram Schultz. 2007. "Reward Value Coding Distinct from Risk Attitude-Related Uncertainty Coding in Human Reward Systems." *Journal of Neurophysiology* 97 (2): 1621–32.

Todorov, Emanuel. 2002. "Cosine Tuning Minimizes Motor Errors." *Neural Comput* 14 (6): 1233–60. https://doi.org/10.1162/089976602753712918.

Tomassini, Alessandro, Darren Price, Jiaxiang Zhang, and James B Rowe. 2019. "On the Evolution of Neural Decisions from Uncertain Visual Input to Uncertain Actions." *bioRxiv*, January, 803049. https://doi.org/10.1101/803049.

Tosun, Tuğçe, Dilara Berkay, Alexander T. Sack, Yusuf Ö Çakmak, and Fuat Balcı. 2017. "Inhibition of Pre-Supplementary Motor Area by Continuous Theta Burst Stimulation Leads to More Cautious Decision-Making and More Efficient Sensory Evidence Integration." *J Cogn Neurosci* 29 (8): 1433–44. https://doi.org/10.1162/jocn_a_01134.

Trueblood, Jennifer S, Scott D Brown, and Andrew Heathcote. 2014. "The Multiattribute Linear Ballistic Accumulator Model of Context Effects in Multialternative Choice." *Psychological Review* 121 (2): 179.

Turner, Brandon M., Birte U. Forstmann, Bradley C. Love, Thomas J. Palmeri, and Leendert Van Maanen. 2016. "Approaches to Analysis in Model-Based Cognitive Neuroscience." *Journal of Mathematical Psychology*, 1–15. https://doi.org/10.1016/j.jmp.2016.01.001.

Turner, Brandon M, Per B Sederberg, Scott D Brown, and Mark Steyvers. 2013. "A Method for Efficiently Sampling from Distributions with Correlated

Dimensions." *Psychological Methods* 18 (3): 368–84. https://doi.org/10.1037/a0032222.

Turner, Brandon M., Ting Wang, and Edgar C. Merkle. 2017. "Factor Analysis Linking Functions for Simultaneously Modeling Neural and Behavioral Data." *Neuroimage* 153 (June): 28–48. https://doi.org/10.1016/j.neuroimage.2017.03.044.

Turner, B, Christian A. Rodriguez, Tony M. Norcia, Samuel M. McClure, and Mark Steyvers. 2016. "Why More Is Better: Simultaneous Modeling of EEG, fMRI, and Behavioral Data." *Neuroimage* 128 (March): 96–115. https://doi.org/10.1016/j.neuroimage.2015.12.030.

Turner, Charles K. 2017. "A Principle of Intentionality." *Front. Psychol.* 8. https://doi.org/10.3389/fpsyg.2017.00137.

Tversky, Amos, and Daniel Kahneman. 1989. "Rational Choice and the Framing of Decisions." In *Multiple Criteria Decision Making and Risk Analysis Using Microcomputers*, 81–126. Springer.

———. 1992. "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of Risk and Uncertainty* 5 (4): 297–323.

Twomey, Deirdre M, Peter R Murphy, Simon P Kelly, and Redmond G O'Connell. 2015. "The Classic P300 Encodes a Build-to-Threshold Decision Variable." *European Journal of Neuroscience* 42 (1): 1636–43.

Twomey, Deirdre, Peter R. Murphy, Simon P. Kelly, and Redmond G. O'Connell. 2015. "The Classic P300 Encodes a Build-to-Threshold Decision Variable." *European Journal of Neuroscience* 42 (1): 1636–43. https://doi.org/10.1111/ejn.12936.

Urai, Anne E, Jan Willem de Gee, Konstantinos Tsetsos, and Tobias H Donner. 2019. "Choice History Biases Subsequent Evidence Accumulation." Edited by Timothy Verstynen, Barbara G Shinn-Cunningham, and Timothy Verstynen. *eLife* 8: e46331. https://doi.org/10.7554/eLife.46331.

Usher, Marius, and James L McClelland. 2001. "The Time Course of Perceptual Choice: The Leaky, Competing Accumulator Model." *Psychological Review* 108 (3): 550–92. https://doi.org/10.1037/0033-295X.108.3.550.

Usher, Marius, Konstantinos Tsetsos, David Albert Lagnado, and Erica Yu. 2013. "Dynamics of Decision-Making: From Evidence Accumulation to Preference and Belief." *Front. Psychol.* 4. https://doi.org/10.3389/fpsyg.2013.00758.

Vandekerckhove, Joachim, Francis Tuerlinckx, and Michael D Lee. 2011. "Hierarchical Diffusion Models for Two-Choice Response Times." *Psychological Methods* 16 (1): 44.

Veen, Vincent van, Marie K MK Krug, Jonathan W JW Schooler, Cameron S CS Carter, and Vincent Van Veen. 2009. "Neural Activity Predicts Attitude Change in Cognitive Dissonance." *Nature Neuroscience* 12 (11): 1469–74. https://doi.org/10.1038/nn.2413.

Vehtari, Aki, Andrew Gelman, and Jonah Gabry. 2017. "Practical Bayesian Model Evaluation Using Leave-One-Out Cross-Validation and WAIC." *Statistics and Computing* 27 (5): 1413–32. https://doi.org/10.1007/s11222-016-9696-4.

Verdonck, Stijn, and Francis Tuerlinckx. 2013. "Factoring Out Non-Decision Time in Choice RT Data : Theory and Implications" 123 (2): 1–13.

Verleger, Rolf, Piotr Jaśkowski, and Edmund Wascher. 2005. "Evidence for an Integrative Role of P3b in Linking Reaction to Perception." *Journal of Psychophysiology* 19 (3): 165–81.

Vermeylen, Luc, David Wisniewski, Carlos González-García, Vincent Hoofs, Wim Notebaert, and Senne Braem. 2020. "Shared Neural Representations of Cognitive Conflict and Negative Affect in the Medial Frontal Cortex." *J. Neurosci.* 40 (45): 8715–25. https://doi.org/10.1523/JNEUROSCI.1744-20.2020.

Vinckier, Fabien, Lionel Rigoux, Irma T. Kurniawan, Chen Hu, Sacha Bourgeois-Gironde, Jean Daunizeau, and Mathias Pessiglione. 2019. "Sour Grapes and Sweet Victories: How Actions Shape Preferences." *PLoS Comput Biol* 15 (1). https://doi.org/10.1371/journal.pcbi.1006499.

Voigt, Katharina, Carsten Murawski, and Stefan Bode. 2017. "Endogenous Formation of Preferences : Choices Systematically Change Willingness-to-Pay for Goods" 43 (12): 1872–82.

Voigt, Katharina, Carsten Murawski, Sebastian Speer, and Stefan Bode. 2019. "Hard Decisions Shape the Neural Coding of Preferences." *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience* 39 (4): 718–26. https://doi.org/10.1523/JNEUROSCI.1681-18.2018.

Vugt, Marieke Karlijn van, Patrick Simen, Leigh E Nystrom, Philip Holmes, and Jonathan D Cohen. 2012. "EEG Oscillations Reveal Neural Correlates of Evidence Accumulation." *Frontiers in Neuroscience* 6: 106.

Wagenmakers, Eric-Jan, Maarten Marsman, Tahira Jamil, Alexander Ly, Josine Verhagen, Jonathon Love, Ravi Selker, et al. 2018. "Bayesian Inference

for Psychology. Part I: Theoretical Advantages and Practical Ramifications." *Psychon Bull Rev* 25 (1): 35–57. https://doi.org/10.3758/s13423-017-1343-3.

Wagner, Ben Jonathan, Mareike Clos, Tobias Sommer, and Jan Peters. 2020. "Dopaminergic Modulation of Human Inter-Temporal Choice: A Diffusion Model Analysis Using the D2-Receptor-Antagonist Haloperidol." *bioRxiv*.

Wegner, Daniel M. 2004. "Précis of the Illusion of Conscious Will." *Behavioral and Brain Sciences* 27 (5): 649–59. https://doi.org/10.1017/S0140525X04000159.

Weiskopf, Nikolaus, Chloe Hutton, Oliver Josephs, and Ralf Deichmann. 2006. "Optimal EPI Parameters for Reduction of Susceptibility-Induced BOLD Sensitivity Losses: A Whole-Brain Analysis at 3 T and 1.5 T." *Neuroimage* 33 (2): 493–504. https://doi.org/10.1016/j.neuroimage.2006.07.029.

Wenke, Dorit, Stephen M. Fleming, and Patrick Haggard. 2010. "Subliminal Priming of Actions Influences Sense of Control over Effects of Action." *Cognition* 115 (1): 26–38. https://doi.org/10.1016/j.cognition.2009.10.016.

Wheeler, David R. 1974. "Brand Loyalties: Qualitative, Quantitative, or Both?" *Journal of the Academy of Marketing Science* 2 (4): 651–58.

White, Corey N, Eliza Congdon, Jeanette A Mumford, Katherine H Karlsgodt, Fred W Sabb, Nelson B Freimer, Edythe D London, Tyrone D Cannon, Robert M Bilder, and Russell A Poldrack. 2014. "Decomposing Decision Components in the Stop-Signal Task: A Model-Based Approach to Individual Differences in Inhibitory Control." *Journal of Cognitive Neuroscience* 26 (8): 1601–14.

Wiecki, Thomas V, Imri Sofer, and Michael J Frank. 2013. "HDDM: Hierarchical Bayesian Estimation of the Drift-Diffusion Model in Python."

Frontiers in Neuroinformatics 7 (August): 14. https://doi.org/10.3389/fninf.2013.00014.

Wisniewski, David, Thomas Goschke, and John-Dylan Haynes. 2016. "Similar Coding of Freely Chosen and Externally Cued Intentions in a Fronto-Parietal Network." *Neuroimage* 134 (July): 450–58. https://doi.org/10.1016/j.neuroimage.2016.04.044.

Won, Bo-Yeong, and Andrew B Leber. 2016. "How Do Magnitude and Frequency of Monetary Reward Guide Visual Search?" *Attention, Perception, & Psychophysics* 78 (5): 1221–31.

Woo, Sung-Ho, and Kyoung-Min Lee. 2007. "Effect of the Number of Response Alternatives on Brain Activity in Response Selection." *Hum Brain Mapp* 28 (10): 950–58. https://doi.org/10.1002/hbm.20317.

Yao, Yuling, Aki Vehtari, Daniel Simpson, and Andrew Gelman. 2017. "Using Stacking to Average Bayesian Predictive Distributions." September 15, 2017. https://doi.org/10.1214/17-BA1091.

Yarkoni, Tal. 2019. "The Generalizability Crisis." PsyArXiv. https://doi.org/10.31234/osf.io/jqw35.

Yau, Y., M. Dadar, M. Taylor, Y. Zeighami, L. K. Fellows, P. Cisek, and A. Dagher. 2020. "Neural Correlates of Evidence and Urgency During Human Perceptual Decision-Making in Dynamically Changing Conditions." *Cerebral Cortex (New York, N.Y.: 1991)* 30 (10): 5471–83. https://doi.org/10.1093/cercor/bhaa129.

Yau, Y., T. Hinault, M. Taylor, P. Cisek, L. K. Fellows, and A. Dagher. 2020. "Evidence and Urgency Related EEG Signals During Dynamic Decision-

Making in Humans." *bioRxiv*, October, 2020.10.02.323683. https://doi.org/10.1101/2020.10.02.323683.

Yoo, Seng Bum Michael, Benjamin Yost Hayden, and John M. Pearson. 2021. "Continuous Decisions." *Philosophical Transactions of the Royal Society B: Biological Sciences* 376 (1819): 20190664. https://doi.org/10.1098/rstb.2019.0664.

Yotsumoto, Yuko, Takeo Watanabe, and Yuka Sasaki. 2008. "Different Dynamics of Performance and Brain Activation in the Time Course of Perceptual Learning." *Neuron* 57 (6): 827–33. https://doi.org/10.1016/j.neuron.2008.02.034.

Young, Michael E, Tara L Webb, Jillian M Rung, and Anthony W McCoy. 2014. "Outcome Probability Versus Magnitude: When Waiting Benefits One at the Cost of the Other." *PloS One* 9 (6).

Zajkowski, Wojciech K., Malgorzata Kossut, and Robert C. Wilson. 2017. "A Causal Role for Right Frontopolar Cortex in Directed, but Not Random, Exploration." *eLife* 6: 1–18. https://doi.org/10.7554/eLife.27430.

Zajkowski, Wojciech, Dominik Krzemiński, Jacopo Barone, Lisa H. Evans, and Jiaxiang Zhang. 2020. "Breaking Deadlocks: Reward Probability and Spontaneous Preference Shape Voluntary Decisions and Electrophysiological Signals in Humans." *Comput Brain Behav*, November. https://doi.org/10.1007/s42113-020-00096-6.

Zhang, Huaxiang, and Zhichao Wang. 2011. "A Normal Distribution-Based over-Sampling Approach to Imbalanced Data Classification." In *International Conference on Advanced Data Mining and Applications*, 83–96. Springer.

Zhang, Jiaxiang. 2012. "The Effects of Evidence Bounds on Decision-Making: Theoretical and Empirical Developments." *Frontiers in Psychology* 3: 263.

Zhang, Jiaxiang, Laura Hughes, and James Rowe. 2012. "Selection and Inhibition Mechanisms for Human Voluntary Action Decisions." *NeuroImage* 63 (July): 392–402. https://doi.org/10.1016/j.neuroimage.2012.06.058.

Zhang, Jiaxiang, Timothy Rittman, Cristina Nombela, Alessandro Fois, Ian Coyle-Gilchrist, Roger A Barker, Laura E Hughes, and James B Rowe. 2016. "Different Decision Deficits Impair Response Inhibition in Progressive Supranuclear Palsy and Parkinson's Disease." *Brain* 139 (1): 161–73.

Zhang, Jiaxiang, and James B Rowe. 2014. "Dissociable Mechanisms of Speed-Accuracy Tradeoff During Visual Perceptual Learning Are Revealed by a Hierarchical Drift-Diffusion Model." *Frontiers in Neuroscience* 8: 69.

———. 2015. "The Neural Signature of Information Regularity in Temporally Extended Event Sequences." *Neuroimage* 107: 266–76.

Zhang, Lei, and Jan Gläscher. 2020. "A Brain Network Supporting Social Influences in Human Decision-Making." *Science Advances* 6 (34): eabb4159. https://doi.org/10.1126/sciadv.abb4159.

Zhang, Lei, Lukas Lengersdorff, Nace Mikus, Jan Gläscher, and Claus Lamm. 2020. "Using Reinforcement Learning Models in Social Neuroscience: Frameworks, Pitfalls and Suggestions of Best Practices." *Social Cognitive and Affective Neuroscience* 15 (6): 695–707. https://doi.org/10.1093/scan/nsaa089.

Zito, Giuseppe A., Roland Wiest, and Selma Aybek. 2020. "Neural Correlates of Sense of Agency in Motor Control: A Neuroimaging Meta-Analysis." *PLoS One* 15 (6): e0234321. https://doi.org/10.1371/journal.pone.0234321.

Appendix

Chapter 4

LBA model including threshold varying per condition

To account for the suggested possibility of strategic threshold adjustment during the trials, I fitted two additional models which included the possibility of threshold parameter varying dependent on condition type. First one included varying accumulation rates v and thresholds B per condition (*VB model*) and second one additionally included varying non-decision time *Ter* per condition (*VBTer model*). I then compared these models with the original model, which varied v and *Ter* per conditon (*VTer model*) using the LOO information Criterion score (Vehtari, Gelman, and Gabry 2017), where lower scores indicate better fit. The comparison revealed that the original *VTer* model provided a better fit of LOO = 12421, compared to both *VB* (LOO = 15113) and *VBTer* (LOO = 12678) models. This indicates that varying the threshold parameter *per* condition resulted in worse fits, so the models were discarded from further analysis.

Chapter 6

Exact Behavioural Model Specifications

All syntax represented using lme syntax format (Bates et al. 2014), where stars denote interaction terms and random effect terms are contain within the parentheses.

Experiments 1-3. Choice-Induced Bias ~ congruence * domain + congruence * congruent difficulty + congruence * incongruent difficulty + congruence * congruent magnitude + congruence * incongruent magnitude + (1 + congruence + domain + congruent difficulty + incongruent difficulty + congruent_magnitude + incongruent difficulty | participant)

Experiment 4. Choice-Induced Bias ~ congruence * domain + congruence * congruent difficulty + congruence * incongruent difficulty + congruence * congruent magnitude + congruence * incongruent magnitude + replaced item (1 + congruence + domain + congruent difficulty + incongruent difficulty + congruent_magnitude + incongruent difficulty | participant)

Experiment 5. a) Reference Bias ~ domain (1+domain | participant) (using only no-choice trials) b) Choice-Induced Bias ~ domain * congruence * reference type + (1|participant) *(data aggregated per participant; non-hierarchical structure was used, since our design does not allow us to estimate single-trial positivity scores without confounding the choice-bias estimate)*

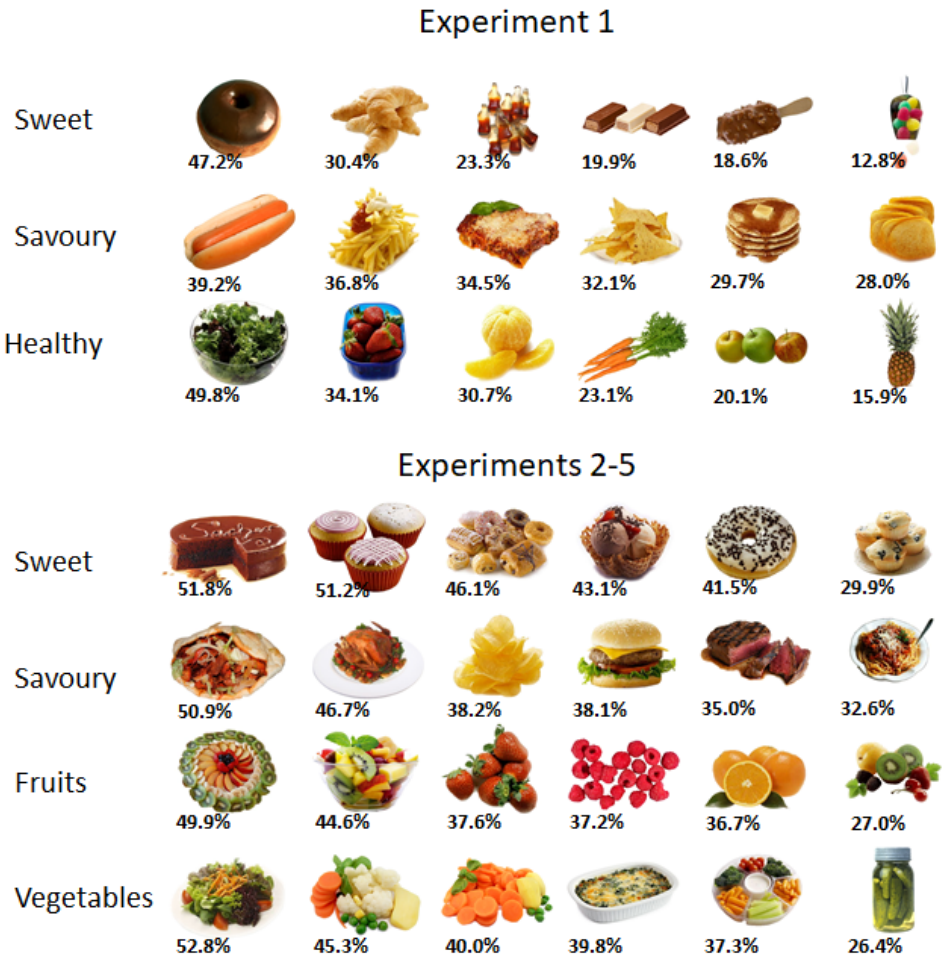

Figure 7.1: Stimuli used across experiments. Below each item the true percentage of non-white pixels (item size). All items were taken from Food-Pics online database (Blechert et al., 2019).

Preference vs Size Ratings

Preference ratings displayed a moderate significant correlation across sessions in all but Experiment 1: M = 0.062, SD = 0.265 t(22) = 1.120, p = 0.27 (Experiment 1); M = 0.120, SD = 0.228 t(243) = 8.178, p < 0.001 (Experiment 2); M = 0.130, SD = 0.193 t(49) = 4.765, p < 0.001 (Experiment 3); M = 0.180, SD = 0.230 t(48) = 5.544, p < 0.001 (Experiment 4); M = 0.204, SD = 0.266 t(49) = 5.367, p < 0.001 (Experiment 5);

This correlation increased in the post-task rating, compared to pre ,with the increase being significant in 2 of the Experiments: M = 0.015, SD = 0.165 t(22) = -0.434, p = 0.67 (Experiment 1); M = 0.038, SD = 0.215 t(243) = -2.788, p = 0.005 (Experiment 2); M = 0.086, SD = 0.193 t(49) = -3.163, p = 0.002 (Experiment 3); M = 0.038, SD = 0.217 t(48) = -1.251, p = 0.22 (Experiment 4); M = 0.018, SD = 0.233 t(49) = -0.553, p = 0.58 (Experiment 5);

Frequentist effect sizes of choice-induced bias across experiments and conditions

To measure the bias effects independently across domain and congruence, we aggregate the data per participant and test the bias against 0. In experiment 1, we found strong evidence for both within and *across-domain* bias in both domains: $M_{\text{pref-within}}$ = 27.3, CI = [21.55, 33.09], t(22)=9.82, p < 0.001, d = 2.05; $M_{\text{size-within}}$ = 17.4, CI = [14.11, 20.72], t(22)=10.92, p < 0.001, d = 2.28; $M_{\text{pref-between}}$ = 10.8, CI = [7.45, 14.09], t(22)=6.73 , p < 0.001, d = 1.40; $M_{\text{size-between}}$ = 6.33, CI = [3.14, 9.52], t(22)=4.12, p < 0.001, d = 0.86.

We replicated this effect in online experiment 2, performed on a much larger sample (250), using a different set of stimuli and a shorter paradigm (12 items; 198 trials; see *Methods*): $M_{\text{pref-within}} = 35.8$, CI = [34.02, 37.63], t(250)=39.11, $p < 0.001$, d = 2.49; $M_{\text{size-within}} = 19.6$, CI = [17.92, 21.36], t(250)=22.52, $p < 0.001$, d = 1.41; $M_{\text{pref-between}} = 10.3$, CI = [8.52, 12.17], t(250)=11.16 $p < 0.001$, d = 0.69; $M_{\text{size-between}} = 6.1$, CI = [4.73, 7.57], t(250)=8.55, $p < 0.001$, d = 0.53.

In Experiment 3, the effects of forced-choice bias was not significant: $M_{\text{pref-forced}} = 0.9$, CI = [-1.93, 3.88], t(49)=0.67, p = 0.5, $M_{\text{size-forced}} = 0.6$, CI = [-1.31, 2.64], t(49)=0.68, p = 0.5.

In Experiment 4, the effect rejected-item undervaluation was greater than chosen-item overvaluation: F(1,49) = 4.95, p = 0.03, d = 0.320.

In Experiment 5, referencing increased the judgments of the referenced item in the no-choice condition: $M_{\text{reference-pref}} = 11.9$, CI = [6.7, 17.02], t(49)=4.62, $p < 0.001$, d = 0.649; $M_{\text{reference-size}} = 7.3$, CI = [2.96, 11.7], t(49)=3.37, $p < 0.001$, d = 0.470. The baselined choice-bias was significantly greater than 0 for all conditions: $M_{\text{pref-within}} = 38.3$, CI = [34.17, 42.45], t(49)=13.73, $p < 0.001$, d = 1.84; $M_{\text{size-within}} = 27.6$, CI = [21.96, 33.26], t(49)=9.82, $p < 0.001$, d = 1.36; $M_{\text{pref-between}} = 6.64$, CI = [1.35, 11.93], t(49)=2.52, $p < 0.01$, d = 0.432; $M_{\text{size-between}} = 8.18$, CI = [3.01, 13.36], t(49)=3.18, p = 0.001, d = 0.393. After removing positivity baseline, the differences in choice-induced bias between reference-chosen and reference-rejected conditions were non-significant F(1,49) = 0.64, p =0.43.

Spontaneous Bias confound simulation

An alternative interpretation of choice-induced bias effects assumes that choices simply reveal preexisting stable preferences (Chen and Risen 2010). This is due to the fact that initial item ratings are a noisy measurement and give imperfect information about the underlying values. Future choices might therefore not as much induce a change of underlying values, but rather reveal them. A simple example would be when item A's 'true' value is 50, while item B's true value is 45, but due to rating variability, A is rated at 45 and B at 50. Based on ratings, a choice of A followed by a higher preference for A would indicate choice-induced value change, while in reality it might be an artifact reflecting regression to the mean effect. To test to what extent this spontaneous bias artifact can affect our design, we performed a series of simulations with a conservative assumption that true values are indeed stationary and choices can only reveal, but not influence them. We then measured the spontaneous bias generated from this model, manipulating key data generating parameters: the within-item rating variability, choice inverse temperature, and judgment standard deviation. All of these relate to the precision with which items are assessed, and hence should be crucial when estimating the confound.

True Value Estimation/Rating generation/Simulation

We take a data-driven approach, where true values and ratings are derived from ratings in experiment 2. We first aggregate the rating data per participant and empirically derive the group-level distributions for 3 key parameters:

true value mean *VM*, true value variability between items *VVB*, and rating variability within items *RVW*, within each domain. We assume that these parameters can be considered a reliable approximation of the underlying true values. We further assume that individual level parameters are distributed normally, centered at the group-level mean and scaled with group group-level variability, truncated at the ends of the ratings scale:

$$VM_{s,d} \sim N(\mu VM_d, \sigma VM_d) \tag{7.1}$$

$$VVB_{s,d} \sim N(\mu VVB_d, \sigma VB_d) \tag{7.2}$$

$$RVW_{s,d} \sim N(\mu RVW_d, \sigma RVW_d) \tag{7.3}$$

where μ and σ represent group-level mean and variability respectively estimated per domain d, while individual parameters are estimated per participant s.

We additionally derive group-level parameter controlling for the correlation between the two rating types. This is because spontaneous bias in *across-domain* trials can only occur if the true preference and size values are positively correlated (otherwise, a preference choice could not provide any information about size, or vice versa). Similarly to above, we assume that individual-level correlation is drawn from the group-level distribution:

$$Vcor_s \sim N(\mu Vcor, \sigma Vcor_d) \tag{7.4}$$

We therefore assume that the true values for preference and size are drawn from a bivariate normal distribution centered at the respective means, with variance-covariance matrix given by the *VVB* and *Vcor* parameters:

$$\begin{pmatrix} V_{pref,s} \\ V_{size,s} \end{pmatrix} = N\left[\begin{pmatrix} VVM_{pref,s} \\ VVM_{size,s} \end{pmatrix}\begin{pmatrix} VVB_{pref,s} & Vcor_s \\ Vcor_s & VVB_{size,s} \end{pmatrix}\right] \tag{7.5}$$

Given the true values, we assume that the ratings for each item i are distributed normally, centered at the true mean $V_{i,d}$ and scaled by the within-item variability $VVW_{i,d}$:

$$Ri, s, d, n \sim N(\mu VMi, d, \sigma RVWi, d) \tag{7.6}$$

where $R_{i,s,d,n}$ represents the rating for item i and participant s in domain d and rating number n (as in the experiment, we simulate two ratings, which are then averaged).

Choices and Judgments Generation

Similar to our model, we assume choices can be approximated by a softmax decision rule with parameter τ controlling choice variability, while judgments are distributed normally, centered at their true value with a scaling parameter ξ (see Methods for equations).

We estimate the two variability parameters by fitting a version of the Null model with hierarchical estimation of τ and ξ, and extracting their group-level

Table 7.1: Data-driven group-level parameters used for generating synthetic data.

Value-related	Size-related	Value and Size-related
$\mu VM_{pref} = 79.48$	$\mu VM_{size} = 63.50$	$\mu Vcor = 0.07$
$\sigma VM_{pref} = 7.94$	$\sigma VM_{size} = 10.34$	$\sigma Vcor = 0.28$
$\mu VVB_{pref} = 10.48$	$\mu VVB_{size} = 16.19$	$\mu\tau = 8.75$
$\sigma VVB_{pref} = 3.93$	$\sigma VVB_{size} = 5.74$	$\sigma\tau = 3.16$
$\mu RVW_{pref} = 6.90$	$\mu RVW_{size} = 9.04$	$\mu\xi = 39.08$
$\sigma RVW_{pref} = 3.54$	$\sigma RVW_{size} = 3.34$	$\sigma\xi = 11.22$

means and standard deviations parameter fit:

$$\tau_i \sim N(\mu\tau, \sigma\tau) \tag{7.7}$$

$$\xi_i \sim N(\mu\xi, \sigma\xi) \tag{7.8}$$

All group-level parameter values are summarized in table 7.1. Data-driven group-level parameters used for generating synthetic data.

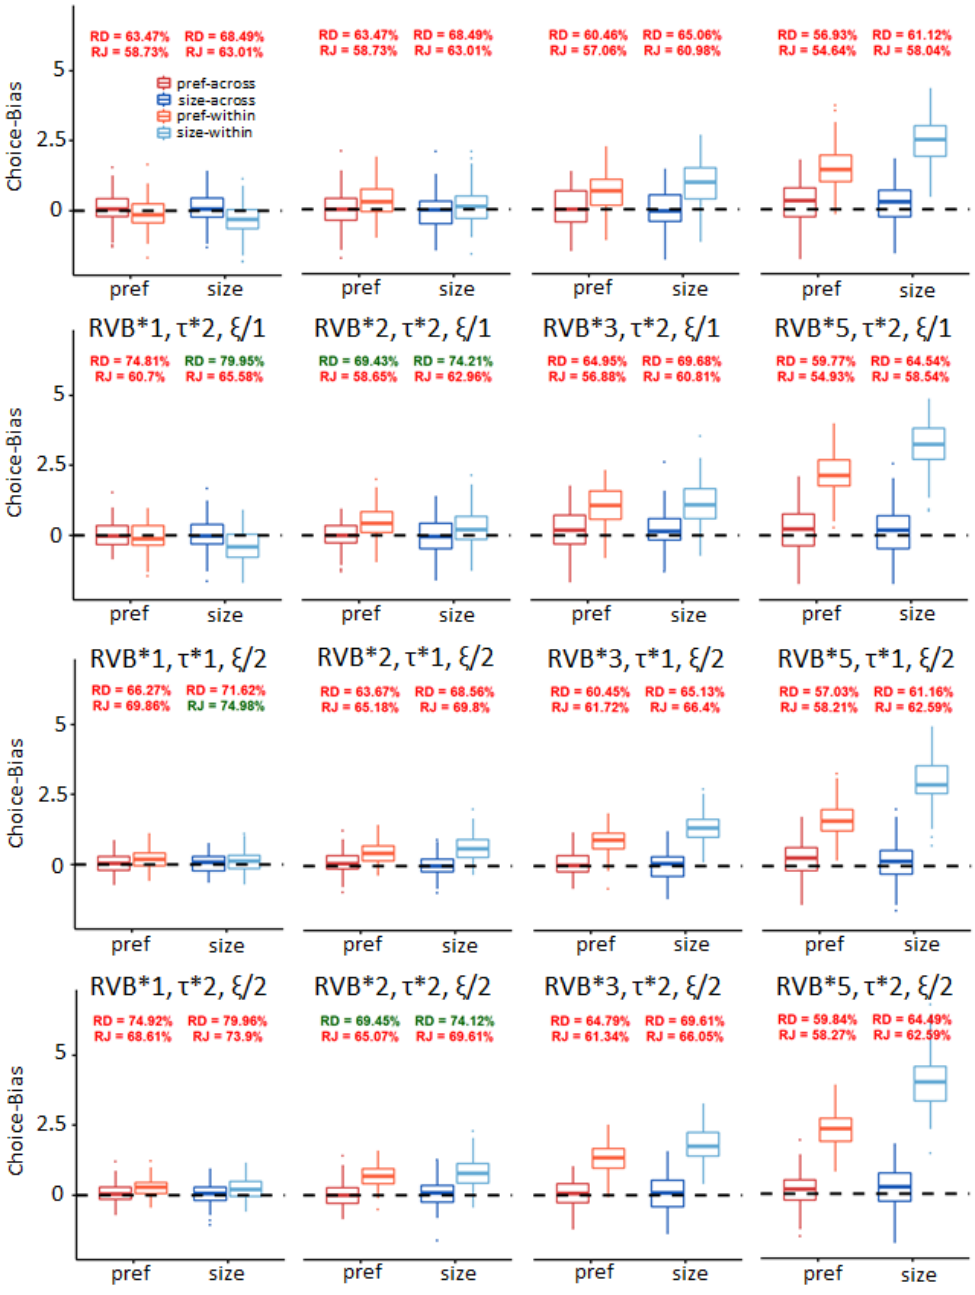
RD = 63.47% RD = 68.49%
RJ = 58.73% RJ = 63.01%
pref-across
size-across
pref-within
size-within
RD = 63.47% RD = 68.49%
RJ = 58.73% RJ = 63.01%
RD = 60.46% RD = 65.06%
RJ = 57.06% RJ = 60.98%
RD = 56.93% RD = 61.12%
RJ = 54.64% RJ = 58.04%
RVB*1, τ*2, ξ/1
RD = 74.81% RD = 79.95%
RJ = 60.7% RJ = 65.58%
RVB*2, τ*2, ξ/1
RD = 69.43% RD = 74.21%
RJ = 58.65% RJ = 62.96%
RVB*3, τ*2, ξ/1
RD = 64.95% RD = 69.68%
RJ = 56.88% RJ = 60.81%
RVB*5, τ*2, ξ/1
RD = 59.77% RD = 64.54%
RJ = 54.93% RJ = 58.54%
RVB*1, τ*1, ξ/2
RD = 66.27% RD = 71.62%
RJ = 69.86% RJ = 74.98%
RVB*2, τ*1, ξ/2
RD = 63.67% RD = 68.56%
RJ = 65.18% RJ = 69.8%
RVB*3, τ*1, ξ/2
RD = 60.45% RD = 65.13%
RJ = 61.72% RJ = 66.4%
RVB*5, τ*1, ξ/2
RD = 57.03% RD = 61.16%
RJ = 58.21% RJ = 62.59%
RVB*1, τ*2, ξ/2
RD = 74.92% RD = 79.96%
RJ = 68.61% RJ = 73.9%
RVB*2, τ*2, ξ/2
RD = 69.45% RD = 74.12%
RJ = 65.07% RJ = 69.61%
RVB*3, τ*2, ξ/2
RD = 64.79% RD = 69.61%
RJ = 61.34% RJ = 66.05%
RVB*5, τ*2, ξ/2
RD = 59.84% RD = 64.49%
RJ = 58.27% RJ = 62.59%
Choice-Bias
5
2.5
0
pref
size

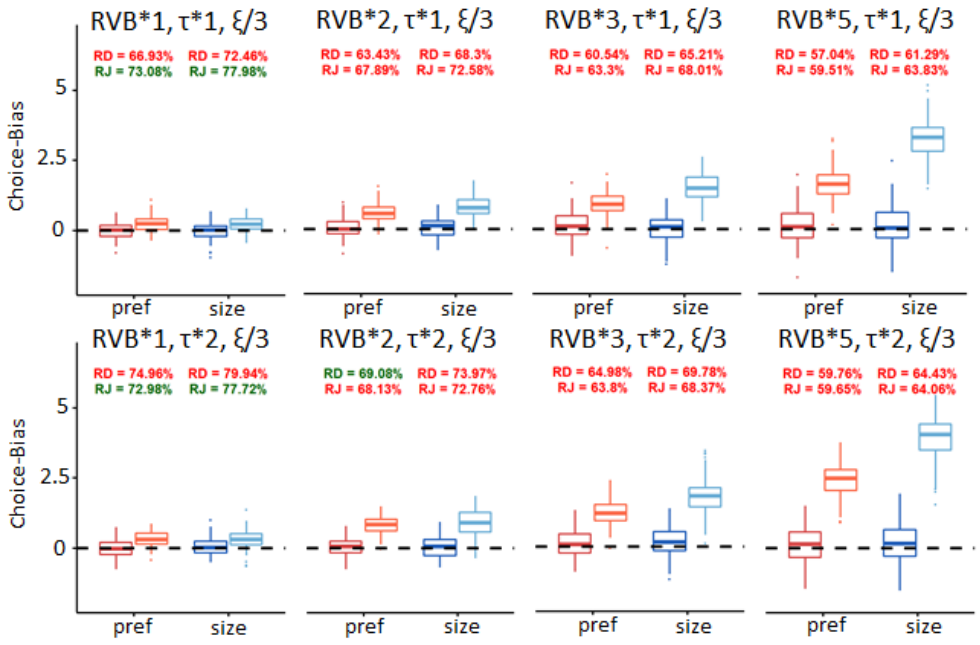

Figure 7.2: Spontaneous bias simulation. Each plot presents one simulation condition (100 simulations), each data-point is the mean bias across a single simulation per condition (250 synthetic participants per simulation). Parameters RVB, τ, and ξ represent empirically estimated rating variability between-items, decision inverse temperature, and judgment standard deviation. RD and RJ are rating-decision and rating-judgment consistency measures averaged across simulation condition (left column: size condition; right column: preference). Green indicates that the 95% confidence interval estimated from the simulation intersected with the 95% interval calculated for the real data (simulation produced plausible values of a given metric).

Simulation

We test 24 conditions based on varying levels of within-item rating variability (4 levels), decision inverse temperate (3 levels) and decision variability (2 levels).

The RVW parameter is crucial, as it controls the precision with which the ratings reflect true values. Intuitively, the less precise the ratings, the less we can infer about the underlying distributions and the larger the confounding bias should be (for a more thorough argument see: Chen and Risen 2010; Izuma and Murayama 2013). Since estimating scale parameters from only 2 observations is rather noisy, we err on the side of caution and additionally test it at 3 more conservative levels: twice (RVBx2), thrice (RVBx3), and five times (RVBx5) as large as their empirical estimations. This conservative approach allows us to account for the worst case scenario: how large can the spontaneous bias be given extremely unreliable ratings.

We also test more conservative levels of τ and ξ. Unlike in the case of real data, here we have direct access to the true item values and therefore do not need to make the simplifying assumption that ratings are their perfect representation. It is therefore likely that knowing the true values would lead to smaller error and variability terms. To account for this, we test three levels of τ: empirically estimated (τx1), twice and thrice as large (τx3), as well as 2 levels of ξ: empirically estimated (ξx1), and twice as small (ξx1). The parameters are scaled in opposite directions, since greater τ represents higher accuracy, while the opposite is true for ξ.

We perform 100 simulations per tested condition. For each simulation, we

generate ratings and data for 250 synthetic participants. To test whether simulation condition produces a realistic output, we calculate the rating-choice and rating-judgment consistency measures (see *Results* section for definitions) and test whether the 95% confidence intervals estimated from simulation in each conditions overlap with the 95% confidence interval estimated from the real data.

Our analysis shows that the design is robust against the confounding effect of spontaneous bias. Given realistic assumptions that can reproduce key data patterns, the bias is negligible at most (< 1 point; Figure 7.2). Consistency analysis indicates that high levels of RVB lead to unrealistically low consistency levels, even when choice and judgment precision is very high. Even in highly unrealistic scenarios, the bias for within-trials did not exceed 5 points.