Prediction Errors of Decision Demands Influence Cost-Benefit Computations in Reasoning

by

Chad C. Williams

MSc, University of Victoria, 2018
BSc (Honours), University of Victoria, 2016

A Dissertation Submitted in Partial Fulfillment of the
Requirements for the Degree of

DOCTOR OF PHILOSOPHY

in the Division of Medical Sciences

Supervisory Committee

Prediction Errors of Decision Demands Influence Cost-Benefit Computations in Reasoning

By

Chad C. Williams

MSc, University of Victoria, 2018
BSc (Honours), University of Victoria, 2016

Supervisory Committee

Dr. Olave Krigolson, **Supervisor**
Division of Medical Sciences

Dr. Bruce Wright, **Committee Member**
Division of Medical Sciences

Dr. Michael Inzlicht, **Outside Member**
Department of Psychology, University of Toronto

Abstract

For each decision we make, we must first determine the degree of effort that we are going to exert, and this can range from no effort to full effort. To select a reasoning strategy (e.g., withholding or exerting effort), it has been proposed that we must first integrate internal and external factors to compute the degree of effort necessary and solve the problem at hand. In this dissertation, I sought to determine the mechanisms underlying selecting such reasoning strategies by leveraging electroencephalographic imaging techniques. My investigations began by exploring neural correlates of effortful contemplation and evolved to test assumptions of prediction errors as it became apparent that they were an influential factor. I then tied this mechanism to the strategy selection phase of reasoning and cost-benefit computations. From these findings, I proposed that prediction errors of decision demands function to lessen or remove the burden of cost-benefit computations. Specifically, repeated encounters of the same or similar decisions provide an opportunity to develop expectations of the prospective costs and benefits of those judgments and these expectations facilitate the reasoning process. I consider two possible explanations as to how prediction errors may influence reasoning: first, our expectations provide our cost-benefit computations with a starting point to be adjusted if necessary, and second, our expectations act as a gating mechanism for cost-benefit computations. Although more research is needed to test these hypotheses, I hope my work provides grounds for advancing this field of study.

Table of Contents

List of Tables

List of Figures

Territorial Acknowledgement

I acknowledge and respect the ləkʷəŋən peoples on whose traditional territory the university stands and the Songhees, Esquimalt and W̱SÁNEĆ peoples whose historical relationships with the land continue to this day.

Acknowledgements

I could have never done this without my friends and family. Thank you to my families for always supporting and believing in me. Thank you to Cameron Hassall for teaching me how to be a great researcher, to Thomas Ferguson for always being encouraging, and to Mathew Hammerstrom for the endless enthusiasm. Thank you to Rob Trska for always making life positive, to Drew Commandeur for always pushing me to learn, and to the Wizards of the Coast for always sparking my creativity. Thank you to the countless lab members, colleagues, and fellow students that have supported me along the way. Thank you to Bruce Wright for always helping me see how my research may impact the world, and to Mickey Inzlicht for the insightful conversations. Indeed, thank you to Olav Krigolson who not only taught me how to be a neuroscientist, but also an educator and mentor. Finally, thank you to the person who has taught me to always do what I want – thank you to Nic Cage.

Dedications

I whole-heartedly dedicate this dissertation to my wife, Ashley, who is the most inspiring person I know. Not only has she single-handedly kept my mental health afloat, but she has always taught me not to take myself, or life, too seriously.

Chapter 1: General Introduction

Outline

This is a manuscript-based dissertation, meaning that the research described (i.e., Chapters 2-5) are meant to be stand-alone articles for publication. The general introduction (i.e., Chapter 1) has been constructed by pulling from each of these manuscripts, and thus is highly redundant with them, albeit with more detail. Moreover, the general introduction was written with the inclusion of the four manuscripts and as such serves as both an introduction to, and a discussion of, the research described here. The general discussion, therefore, serves as a general summary of findings and a platform for speculation as to how these findings may progress considering future research.

It should be noted that, due to the journal of interest for each article, the manuscripts do often use different terminology to refer to effortless and effortful reasoning. Sometimes, I draw from the dual-process literature that depicts effortless reasoning as System 1 and effortful reasoning as System 2. Other times, I draw from the reasoning literature that refers to intuitive and analytical reasoning as effortless and effortful reasoning, respectively. These differences of terminology will be found across the manuscripts; however, the general introduction and discussion conform these terminologies to use effortless and effortful reasoning.

"Thirty white horses on a red hill,

First they champ,

Then they stamp,

Then they stand still"

(Tolkien, 1937)

What are they?

Well, did you solve the riddle? Better yet – did you *try* to solve the riddle, or did you rather decide to move on? Either choice is fine – sure, it might be fun, and you might feel accomplished by solving the riddle, but it might also just be better for you to save your time and energy and get to the paper that you are here to read. The choice to exert effort or not is surprisingly complicated and that is the focus of the work presented here. Oh, and by the way – the answer is teeth.

When presented with a decision, such as a riddle, we first determine how we will approach the problem, and then we engage with it to reach a solution (Alexander & Brown, 2011; J. W. Brown & Alexander, 2017; Collins & Shenhav, 2022; Frömer & Shenhav, 2021; Kool et al., 2017; Kool & Botvinick, 2018; Shenhav et al., 2013, 2017, 2021; Vassena, Holroyd, et al., 2017; Vassena et al., 2020; Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Hassall, et al., 2022). In terms of the riddle, you would have first determined whether you would bother trying to solve it at all and then either attempt to solve it or skip ahead. In this introduction we will discuss these two concepts – the strategy selection phase and the engagement phase of a decision – in reverse order. With regard to the engagement phase, we will discuss effort and its limitations; in terms of the strategy selection phase, we will focus on hypotheses that ground strategy selection as a cost-benefit computation. We will then investigate current electroencephalographic (EEG) findings in reasoning and end by providing a brief summary of the four manuscripts presented here.

The Engagement Phase: Exerting and Withholding Effort

Effort

Exerting or withholding effort in a decision is a process that we are all very familiar with, and whether you spent the effort to solve the riddle or not has very different consequences. Our understanding of effort and reasoning largely stems from a body of research focused on the resolution of word problems. For example, Frederick's (Frederick, 2005) classic widget problem: "*It takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets?"* Such problems have two common answers: one incorrect if effort is withheld (i.e., 100 minutes) and the other correct if effort is exerted (i.e., 5 minutes). It is commonly thought that exerting effort requires more time and mental resources, but with the advantage of being more accurate (DeNeys, 2017; DeNeys & Pennycook, 2019; Evans & Stanovich, 2013; Kahneman, 2011; Pennycook, 2017). In contrast, withholding effort is quick and requires little-to-no mental resources, yet is prone to error (DeNeys, 2017; DeNeys & Pennycook, 2019; Evans & Stanovich, 2013; Kahneman, 2011; Pennycook, 2017). Despite its propensity for mistakes, the majority of our daily decisions are done so effortlessly (Kahneman, 2011). We can often rely on effortless reasoning because it only fails when encountering atypical or novel, rather than highly familiarized, decisions (Cavanagh & Frank, 2014; Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Hassall, et al., 2022). When encountered with mundane and well-practiced decisions, such as when to go to the bathroom, effortless decisions do just fine (Williams, Hassall, et al., 2022).

Just as there are times where effortless reasoning is optimal, there are times where exerting effort is necessary (Williams, Hassall, et al., 2022) – for example, when deciding whether to move to a new city for a job. Indeed, it is unlikely that we are going to make such a life-changing decision on a whim without much thought. Instead, we will contrast the pros and the cons of this move to eventually reach a solution. That is not to say, however, that effortful reasoning is without its faults. Indeed, errors have also been demonstrated with effortful contemplation (Norman, 2000; Norman et al., 2017). One explanation for errors in effortful reasoning is premature closure of reasoning – with the implication that decision making is an enduring process, if motivation to exert effort deteriorates during reasoning, a decision must be made with incomplete consideration (Inzlicht et al., 2014; Inzlicht & Schmeichel, 2012; Müller & Apps, 2018). In other words, errors with effortful reasoning are in part due to an abandonment of the reasoning process. The limitations of each decision making strategy then indicates that rather than always relying on either an effortless or effortful strategy, we must become proficient in identifying when to use which.

These examples highlight the two extremes of the effort spectrum: cases where we would not fathom spending time and effort to reach a solution, such as going to the bathroom, and other cases where we would be foolish to decide without a thought, such as moving for a job (Alexander & Brown, 2011; Botvinick & Cohen, 2014; Egner, 2017; Nigg, 2017; Shenhav et al., 2013). But there are a lot of decisions in-between these two extremes that could go either way depending on who is making the decision and how they were feeling in the moment of making that decision (Alexander & Brown, 2011; J. W. Brown & Alexander, 2017; Collins & Shenhav, 2022; Frömer & Shenhav, 2021; Kool et al., 2017; Kool & Botvinick, 2018; Shenhav et al., 2013, 2017, 2021; Vassena, Holroyd, et al., 2017; Vassena et al., 2020; Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Hassall, et al., 2022). For example, what's for dinner? Answering with "anything is fine" versus "what are our options?" have very different demands. Whereas the former relinquishes the need to effortfully decide what to eat (such as those that skipped the riddle), it also comes with the risk of getting something that you do not want. In contrast, the latter requires the effortful consideration of a lot of factors to reach a solution (such as those that engaged with the riddle) yet ensures that you get the meal you are craving. Why, then, don't we always just put forth effort to make these types of decisions? Afterall, wouldn't we always want to satisfy our cravings? It turns out that we are limited in how many effortful decisions we can make in a given amount of time and so we must choose carefully what decisions we will effortfully engage with (Inzlicht & Friese, 2019; Inzlicht & Schmeichel, 2012; Shenhav et al., 2017).

The Limited Capacity of Mental Energy

The depletion of mental energy during effortful reasoning is similar to the depletion of physical energy during a workout (Baumeister et al., 1998; Baumeister & Heatherton, 1996; Inzlicht & Friese, 2019; Inzlicht & Schmeichel, 2012; Muraven et al., 1998; Shenhav et al., 2017). In both cases, the exertion of energy results in a lessened likelihood of again exerting energy. In the physical domain, we know that one major contributor to physical depletion is glucose and that we can quickly replenish our physical capacity by consuming sugar (Baumeister et al., 1998; Baumeister & Heatherton, 1996; Muraven et al., 1998). In the mental domain, it is not this simple. There is a rich history of research that considered this same substance, blood glucose, to be proportional to our mental capacity (Gailliot et al., 2007), yet this claim has since been criticized and is no longer valid (Dang, 2016; Finley et al., 2019; Kurzban, 2010; Kurzban et al., 2013; Vadillo et al., 2016). That is not to say that the concept of a metabolic constraint in itself has been refuted, but simply that glucose is not the substance underlying our mental capacity. In fact, others have proposed astrocytic glycogen to be proportional (Christie & Schrater, 2015), and amyloid-Beta neurotoxin build up to be inversely proportional (Holroyd, 2015), to mental capacity. None-the-less, these latter substances are still under early investigation and so we do not yet know the extent to which they influence our capacity to engage in effortful reasoning (Shenhav et al., 2017).

Some have rather proposed our mental capacity to be limited by our motivation (Inzlicht et al., 2014; Inzlicht & Berkman, 2015; Inzlicht & Friese, 2019; Inzlicht & Schmeichel, 2012; Kool & Botvinick, 2014). In other words, it is not an inability to continuously engage in effortful reasoning, but rather a lack of motivation to do so. Even another account has suggested that there rather exists a structural capacity – that is, effort is guided by a single mechanism that has computational limitations (Cowan, 2012; Luck & Vogel, 1997; Ma et al., 2014; Oberauer et al., 2016). Although this hypothesis draws on the well-studied limitations of working memory, it proposes a single mechanism as the controller for all of human effort and has been criticized to be biologically implausible (Cowan, 2012; Luck & Vogel, 1997; Ma et al., 2014; Oberauer et al., 2016). Indeed, others have reframed this concept to be system-wide in that exerting effort places strain on information-processing pathways that persists and interferes with our ability to re-engage in another effortful decision (Feng et al., 2014; Musslick, Jang, et al., 2018; Musslick & Cohen, 2019, 2021). Together, it seems that engaging in an effortful decision results in a reduced chance of engaging in another effortful decision, but it’s still unclear why this happens biologically.

The Strategy Selection Phase: The Expected Costs and Benefits of a Decision

Models of Strategy Selection

To select a reasoning strategy (e.g., withholding or exerting effort), it has been proposed that we must first integrate internal (e.g., motivation) and external (e.g., problem difficulty) factors to compute the degree of effort necessary and solve the problem at hand (Evans, 2011; Kahneman & Frederick, 2002). Some have advocated for a dual-process decision making framework wherein effortless and effortful reasoning rely on distinct cognitive systems (Kahneman, 2011; Kahneman & Frederick, 2002). This model posits that effortless processes are engaged as the main operator of the brain, but may be overridden by effortful processes, if deemed necessary. Albeit a popular theory, it lacks any formulized description of what occurs for effortful systems to interrupt effortless systems. In other words, it does not describe how a decision making strategy is selected.

One potential solution to this gap has been proposed in that within decision making multiple potential responses rise and are compared as to whether they are similarly activated (e.g., in terms of value or likelihood) (Bago & Neys, 2018, 2019; DeNeys & Pennycook, 2019; Pennycook, 2017; Pennycook et al., 2015, 2018). If one response is significantly more activated than another, effortless decision making is adopted, but if two or more responses are equally activated, effortful decision making is recruited to resolve the conflict. Although promising, another prominent model describes strategy selection to be much more complex and to rely on cost-benefit computations (Braver et al., 2014; Grahek et al., 2020; Kool et al., 2010, 2013, 2017; Kool & Botvinick, 2014, 2018; Musslick, Cohen, et al., 2018; Musslick et al., 2015, 2019; Shenhav et al., 2013, 2017).

The Expected Value of Control (EVC) model frames strategy selection in reasoning to rely on the expected costs and benefits of a decision (Collins & Shenhav, 2022; Frömer & Shenhav, 2021; Kool et al., 2017; Shenhav et al., 2013, 2017, 2021). It poses that we compare our *expected* benefits and our *expected* costs of engaging effortfully and if the benefits outweigh the costs, then we exert effort, but if the costs outweigh the benefits, then we withhold effort. This claim turns out to be complicated because benefits and costs are not objective truths of the world, but subjective estimations that we compute (Collins & Shenhav, 2022; Frömer & Shenhav, 2021; Kool et al., 2017; Shenhav et al., 2013, 2017, 2021). What are the expected costs of the riddle? Well, this certainly depends on one's ability to solve riddles. Riddles are easier for some than others and so the costs vary from person to person – but also can vary from moment to moment. Solving a riddle when you are sick and tired is more difficult than when you are well and rested. Further, these expectations are encompassed in our personal experiences.

Indeed, experience can decrease the degree of effort needed to complete the same decisions (Ericsson et al., 1993). We have all heard that *practice makes perfect* and that is exactly what we are talking about here. Let us think about learning a new skill, such as driving. The first time you sat behind the wheel of a car, you likely had more to think about than your brain could handle. You were confronted with a load of gauges, the pedals, and of course the steering wheel. Moreover, you do not know what to pay attention to, so you try and pay attention to all of it. Every decision you make, for example pulling onto the street from your driveway, is a stressful process requiring cognitive effort. However, driving gets easier with practice and now as an expert driver you can hold a conversation with your passenger while listening to the radio and drinking a coffee all as you commute along a busy highway.

Expected Costs and Benefits

Sometimes benefits can seem objective and easy to predict – for example, your paycheque. There is often no ambiguity in what your paycheque will be, but that is not the only benefit for doing a good job is it? Alongside a paycheque, we could get praise and respect, but also a promotion. Similarly, doing a good job avoids getting fired – another benefit. But again, all of these benefits are subjective and vary across people and across time (Frömer & Shenhav, 2021; Kool & Botvinick, 2018; Shenhav et al., 2013, 2021; Westbrook et al., 2013; Yee et al., 2021). Money is more important for some than others and at some times than other times, and so the expected benefit of that paycheque varies accordingly. Likewise, some may be aiming for that promotion, while others could not care less if they get fired. And those are only a few examples as to how our expectations of benefits can be pulled this way and that.

In terms of benefits, the EVC model borrowed from the reinforcement learning literature (Apps et al., 2015; Braver et al., 2014; Kool et al., 2017; Kool & Botvinick, 2014, 2018; Müller & Apps, 2018). Specifically, internal representations of environment-action contingencies are developed through exposures, or trial-and-error. Reinforcement learning relies on explicit external feedback and operates to learn the outcomes of actions to maximize future rewards (Holroyd & Krigolson, 2007; Krigolson, 2018; Proudfit, 2015; Sambrook & Goslin, 2015; Williams et al., 2017; Williams, Hassall, et al., 2019). Over many exposures of a given context, we can learn the actions that will best result in future reward depending on our goals (Holroyd & Krigolson, 2007; Krigolson, 2018; Luft, 2014; Walsh & Anderson, 2012; Williams, Hassall, et al., 2019). The values learned then serve as a proxy of the benefits that are considered by the EVC. Other research has also posited intrinsic motivation as a source of benefit – i.e., the process of overcoming a challenge is, in itself, rewarding (Inzlicht et al., 2018).

It should be no surprise that our expectations of costs are also complicated. Some have begun to disentangle the factors that increase the difficulty of a decision, and thus the cost of engaging with it (Grahek et al., 2020; Musslick, Cohen, et al., 2018; Musslick et al., 2015, 2019). Conflict is one such source of decision difficulty. Conflict occurs when two or more options are similarly likely (Botvinick & Cohen, 2014; Egner, 2011; Egner, 2017; Nigg, 2017). To make a decision we consider different response options each with a value (Krajbich et al., 2010; Krajbich & Rangel, 2011; Tajima et al., 2016) and we typically choose the most valued of options. However, if the values are similar it is difficult to choose between them and response conflict arises (Botvinick & Cohen, 2014; Egner, 2011; Egner, 2017; Nigg, 2017). In other words, easy decision making can rely on effortless reasoning systems (De Neys, 2017; Evans & Stanovich, 2013; Kahneman, 2011; Pennycook, 2017), but this strategy only functions as long as there is a clear response (i.e., no conflict). If conflict is present, we rather need to engage effortfully with the problem to be solved (Evans & Stanovich, 2013; Kahneman, 2011). Within this cost-benefit framework, conflict operates as a cost of decision making – enforcing the need for increased incentives in fear of relying on faulty reasoning (Shenhav et al., 2014).

We have so far considered costs in terms of the mental resources drained when making an effortful decision, and although a major factor, this is far from the whole picture (Frömer & Shenhav, 2021; Kool & Botvinick, 2018; Shenhav et al., 2013, 2017, 2021; Westbrook et al., 2013; Yee et al., 2021). Another source of cost is the time it takes to engage effortfully. Have you ever waited hours to see a doctor only to be checked over for a few minutes and then sent away? Emergency room doctors need to make an incredible number of decisions all day long. When seeing you as a new patient, they must decide whether to exert effort and consider all possible aspects of your health before resulting in a diagnosis, or whether they withhold effort and go with their gut-hunch. As patients, we are quick to say that we would prefer they take their time and make sure our cough is not a symptom of an underlying, potentially fatal, disease. But even if doctors had an unlimited bank of mental energy, they still do not have the time. By spending extra time on their patients that portray typical symptoms of a common flu, they are taking time away from seeing other patients that may have life-threatening symptoms. These opportunity costs become an important factor in deciding whether we should exert effort (Frömer & Shenhav, 2021; Kool & Botvinick, 2018; Kurzban et al., 2013; Lin et al., 2022; Shenhav et al., 2013, 2017, 2021; Westbrook et al., 2013; Yee et al., 2021). Similarly, if you had set your entire day aside to read this paper, then maybe you found the time cost of engaging the riddle to be negligible, but if you instead have a stack of papers to read, assignments to grade, and grants to write, then maybe you simply could not eat the time cost of that riddle.

The Interaction of Costs and Benefits

Altogether, in the strategy selection phase of reasoning, we weigh the expected costs and benefits of exerting effort, both of which are shrouded in complexity. Yet, the story still does not end there – costs and benefits also interact with each other. When deciding whether to exert effort in a decision, we must also consider the efficacy of exerting that effort. First, performance efficacy refers to whether exerting effort will have any influence on attaining more benefits (Frömer et al., 2021; Frömer & Shenhav, 2021; Schevernels et al., 2014; Shenhav et al., 2013, 2021). When performance efficacy exists, the more effort exerted, the more benefits we may expect. For example, when deciding whether to exert effort in doing the riddle, you also determine *how much* effort to exert. The more effort you exert, the higher the chances that you will solve the riddle and feel that sense of accomplishment. More specifically, the amount of effort exerted increases our expectations of benefits. If there is no performance efficacy, then the likelihood of achieving the benefits stays the same, no matter how much effort we exert. In these cases, we are not motivated to exert effort. Second, control efficacy regards whether exerting effort will have any influence on our performance (Frömer et al., 2021; Frömer & Shenhav, 2021; Schevernels et al., 2014; Shenhav et al., 2013, 2021). When control efficacy exists, exerting more effort leads to increases in performance, and thus a higher likelihood of attaining the benefits. In other words, our beliefs about whether we are even able to solve the decision at all matters. If you have never been able to solve a riddle in your life, you might believe that no amount of effort will be enough to solve our riddle and achieve the sense of accomplishment. In this case, there is no benefit of exerting effort as there is no possibility of actually attaining it.

Proactive Expectations of Decision Demands

We have one final consideration – our proactive expectations of decision demands. Here, decision demands correspond to a decision's expected costs and benefits, and it seems we build expectations for future decision demands (Alexander & Brown, 2011; J. W. Brown & Alexander, 2017; Vassena et al., 2020; Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Hassall, et al., 2022). These expectations are the *long-run average* of decision demands and are proactive, meaning that they are determined before encountering a decision. When presented with a decision, if the demands match our expectations, we continue with our predisposed reasoning strategy (Frömer & Shenhav, 2021; Shenhav et al., 2021; Williams, Hassall, et al., 2022). However, if the decision demands deviate from our expectations then a prediction error occurs. Prediction errors occur when our expectations are violated and signify that we must re-assess our predisposed strategy (Alexander & Brown, 2011; J. W. Brown & Alexander, 2017; Holroyd & Krigolson, 2007; Krigolson, 2018; Proudfit, 2015; Vassena et al., 2020; Williams et al., 2017; Williams, Ferguson, Hassall, Abimbola, et al., 2021; Williams, Hassall, et al., 2022). In reasoning, prediction errors could signify a violation of our expectations of decision demands and occur both when a decision is more or less demanding than expected (Williams, Hassall, et al., 2022). In other words, in consistent environments, we may settle into a strategy, whether that is to employ effort or to withhold effort. We continue to employ this strategy until demands do not meet our expectations, and then we re-assess to determine whether we should change our strategy.

If there is a flu outbreak, doctors in the emergency room learn to expect the symptoms of their next patient before ever seeing any medical charts – in this case, they expect the classic symptoms of a cough, runny nose, and fever. Correspondingly, they expect that the costs of exerting effort will outweigh the benefits and that they will not need to exert effort. If the doctor's next patient indeed presents with flu-like symptoms, the doctor continues to effortlessly prescribe flu fighting treatments. If the doctor's next patient rather presents with a severed finger, then the doctor must assess whether to change their strategy from effortless to effortful with a cost-benefit computation.

The proactive expectation of decision demands is in line with the Predicted Response Outcome (PRO) model (Alexander & Brown, 2011, 2014; Vassena, Deraeve, et al., 2017; Vassena et al., 2020), which poses a monitoring process that tracks the likelihood of events. Researchers often pit the EVC and PRO models against each other in pursuit of a single account of effortful decision making (Alexander & Brown, 2011; Brown & Alexander, 2017; Kool et al., 2017; Kool & Botvinick, 2018; Shenhav et al., 2013, 2014; Vassena et al., 2020; Vassena, Deraeve, et al., 2017). The intent of this approach is to determine which of these frameworks independently explain neural patterns of effortful decision making. Some have concluded neural data to fit one model over the other (Alexander & Vassena, 2020; Vassena et al., 2020);

however, others have indicated concurrent evidence for both (Shenhav et al., 2020), indicating a need to further consider both mechanisms in parallel. Indeed, we have recently proposed that both the PRO and EVC hypotheses may work in tandem during reasoning (Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Hammerstrom, et al., 2022; Williams, Hassall, et al., 2022). Specifically, the PRO model acts to proactively account for decision demands while the EVC model acts to reactively determine whether to exert effort in a decision. In this framework, the PRO model mediates cost-benefit computations.

Although we speculate that our expectations of decision demands function to relieve the burden of cost-benefit computations (Bustamante et al., 2021; Grahek et al., 2020; Lieder et al., 2018; Musslick, Cohen, et al., 2018; Musslick et al., 2015; Musslick & Cohen, 2019), we caution that we do not know exactly how this works. Perhaps it is that cost-benefit computations draw from our expectations to better fine tune our reasoning strategies. Expectations then provide our cost-benefit computations with a starting point to be adjusted if necessary (Musslick, Cohen, et al., 2018; Musslick et al., 2015). Another possibility could be that our expectations act as a gating mechanism for cost-benefit computations (Frömer & Shenhav, 2021; Shenhav et al., 2021). Specifically, if our expectations of decision demands are met when encountering a decision, no prediction error occurs and there is no need for a cost-benefit computation as we will exert our predisposed amount of effort. Yet if our expectations of decision demands are violated, a prediction error occurs, and a cost-benefit assessment is computed to determine the appropriate strategy for the decision at hand. In this hypothesis, cost-benefit computations are only conducted when decision demands do not meet our expectations. If everything is going as we predicted, then we simply adopt the same strategy as always.

Electroencephalography in Reasoning

EEG imaging may provide unique insights into the mechanisms underlying reasoning. EEG techniques that utilize Event-Related-Potentials (ERP) in reasoning began by examining the neural underpinnings of human reasoning during the deliberation of word problems (Bago et al., 2018; Banks, 2017). Although this field of research found neural responses that distinguished gross differences in reasoning, their analyses did not investigate the strategy selection mechanisms of reasoning. As such, the neural systems that underlie strategy selection, while hypothesized, are still unclear.

ERP techniques, however, have not always proven to be the most robust approach of disentangling cognitive mechanisms. Research has rather pointed towards the power of EEG oscillatory analyses to dissociate underlying mechanisms of cognition (Bernat et al., 2008, 2011, 2015; Foti et al., 2015; Nelson et al., 2011; Williams, Ferguson, Hassall, Abimbola, et al., 2021). One such oscillation is frontal theta activity, which has seen rapid development in recent years. According to some, frontal theta activity reflects signals of cognitive control – novelty, conflict, punishment, and errors – and has been linked to control demands (Cavanagh & Shackman, 2015; Eisma et al., 2021). Others have additionally demonstrated it to be a signal of salience or surprise (Lin et al., 2020).

Indeed, theta activity has also been linked to effortful reasoning (Cavanagh, Figueroa, et al., 2012; Cavanagh & Frank, 2014; Cavanagh & Shackman, 2015; Eisma et al., 2021; Umemoto et al., 2019; Williams, Kappen, et al., 2019; Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, VanOorschot, & Krigolson, 2021; Williams, Hammerstrom, et al., 2022; Williams, Hassall, et al., 2022). For example, in 2019 (Williams, Kappen, et al., 2019) we provided evidence that increased frontal theta power corresponded to effortful rather than effortless reasoning. In 2021 (Williams, VanOorschot, & Krigolson, 2021), we linked theta power to a person's tendency to respond in line with effortful processing when completing word problems. We further demonstrated that, during reasoning, increased theta power was an indicator of expectancy violations, or in other words prediction errors of decision demands (Williams, Ferguson, Hassall, Wright, et al., 2021). Strengthening these findings, we found that theta signals indicating prediction errors of decision demands diminished as one stabilized their expectations through learning (Williams, Hassall, et al., 2022). Finally, in 2022 we found that theta activity reflected prediction errors of decision demands tied to cost-benefit computations during the strategy selection phase of reasoning (Williams, Hammerstrom, et al., 2022).

Overview of Experiments

The goal of this dissertation was to shed light on the mechanisms involved in selecting a reasoning strategy – for example, the degree of effort to exert in a decision. It is a continuation of my master's work, which found frontal theta activity to reflect effortful consideration (Williams, Kappen, et al., 2019). At the beginning of this dissertation, the goal was to identify neural signatures of the EVC model, which would support this theory in the brain and potentially allow for the neural tracking of human reasoning in the real-world. Specifically, I was in search of neural signatures that reflected expected costs, expected benefits, and the cost-benefit computation. In light of past research on frontal theta activity, I found it to be a good candidate of focus and predicted that it reflected our conception of expected costs and/or the cost-benefit computation (with the intent of later searching for the neural indicator of expected benefits).

As such, my research primarily focuses on this neural oscillation – and thus it is present in all manuscripts presented here, albeit alongside other oscillations in some of the research. However, theta activity is a robust signal that seems to manifest across many tasks. Moreover, it reflects different functionality across contexts, implying that it is either a complicated signal reflecting many functions, or that different theta activity across contexts do not represent the same underlying construct. Although theta activity was the focus of this dissertation, the goal was not to further investigate the underlying function of this neural oscillation. Rather, my approach was to use theta activity as a tool to investigate assumptions underlying computational models of reasoning – specifically, the EVC and PRO models. In other words, I sought to determine whether theta activity would reflect costs, benefits, or cost-benefit computations regarding the EVC model, and/or reflect prediction errors in line with the PRO model.

Experiment 1: My first experiment was designed with my original hypothesis in mind – that frontal theta activity reflected expected costs and/or cost-benefit computations. As such, it draws from the reasoning literature by having participants complete base-rate word problems. These problems have two responses – one thought to reflect effortless reasoning and the other effortful reasoning. As such, we tracked theta activity in light of a person's propensity to reason effortfully and found evidence that theta activity reflected effortful reasoning.

Experiment 2: My second experiment was my first to investigate the underlying factors of the EVC model. As I believed that theta activity reflected a signal for the need to engage effortfully in reasoning, I began to manipulate the costs of a decision. I had participants diagnose a series of patients based upon a single diagnostic metric. Whereas low values within this metric would indicate one diagnosis, high values in this metric would indicate another diagnosis. The difficulty of the decision, and thus the expected cost of it, was manipulated as the proximity the presented metric was from the border that separated the two diagnoses. Values near this border would reflect diagnostic conflict and so as the values approached the border, they became more difficult, increasing their cost to engage effortfully. In contrast to my account that theta activity would reflect this conflict and thus a signal of expected costs and/or cost-benefit computations, theta activity was related to deviations from the long-run averaged metric for each diagnosis. I concluded that theta activity then may reflect surprise signals, which my later manuscripts reframed as prediction errors.

Experiment 3: My third experiment was then designed to provide additional support for what was found in experiment two – that theta activity reflected prediction errors. From my work in the reinforcement learning literature, I knew that prediction errors diminish with learning, as our expectations converge towards real truths of the worlds. As such, I used past data that had participants learn to classify different complex shapes into their respective families. I reasoned if the prediction error account were true that theta activity would decrease across learning, as expectations stabilized. Indeed, this was what I found, supporting the proposition that theta activity reflected prediction errors. Moreover, halfway through this experiment, participants were presented with new families to categorize. This then meant that participants were seeing some familiar shapes, which have been learned and thus should not elicit prediction errors, and some novel shapes, which have yet to be learned and thus should elicit prediction errors. Theta activity was dissociated between the familiar and novel shapes, further supporting that it reflected prediction errors during reasoning.

Experiment 4: At this point, my research signified the existence of prediction errors during reasoning. However, experiments 1-3 tracked participant's neural activity while making a decision, rather than while determining the strategy they would employ to make that decision. This approach confounds the strategy selection phase of reasoning with the engagement phase of reasoning. As we did not dissociate the two, we were unable to assertively conclude that these prediction error signals in fact related to selecting strategies through cost-benefit computations.

As such, my final experiment explicitly dissociated the strategy selection from the engagement phase of reasoning by having participants first decide whether to withhold or exert effort in an upcoming judgment, each for different monetary rewards. This paradigm then allowed us to investigate signals during the strategy selection phase of reasoning prior to engaging with a judgment. In this paradigm, we manipulated the benefits of reasoning so that participants encountered decisions with equal costs and benefits as well as decisions that were undervalued and overvalued – these latter two both leading to prediction errors. Theta activity indeed reflected prediction errors in the strategy selection phase of reasoning, corroborating its potential involvement with cost-benefit computations.

Chapter 2: Experiment 1

A Window into the Rational Mind:

The Neural Underpinnings of Human Reasoning

Chad C. Williams*, Folkert Van Oorschot, and Olave E. Krigolson

The Theoretical and Applied Neuroscience Laboratory, University of Victoria, Victoria, British Columbia, V8W 2Y2, Canada

Under Revisions in Thinking and Reasoning, 2022

Williams, C., Van Oorschot, F., & Krigolson, O. (2022). A Window into the Rational Mind: The Neural Underpinnings of Human Reasoning. Thinking and Reasoning. https://doi.org/10.31234/osf.io/gc6u9

Abstract

Humans reason intuitively by relying on gut hunches or rationally through analytical contemplation. The majority of research on human reasoning has relied on behavioural data and thus the neural underpinnings of this process remain unclear. To address this, we had participants perform a classic reasoning task while electroencephalographic (EEG) data was recorded. Within our reasoning task, participants completed a series of base-rate word problems wherein their decisions were either biased by a provided stereotype or based on statistical probability. Post experiment, we defined participant rationality as the percentage of responses that were made based on likelihood. We then examined frontal theta neural oscillations and found that increased power in this frequency range was associated with increased rationality. Our findings suggest that theta oscillations are sensitive to rationality and further that rational reasoning involves a diverse brain network relative to intuitive reasoning.

Introduction

Reasoning is an art that often fails us. We all, on occasion, have acted instinctively and ended up kicking ourselves in hindsight for making faulty judgments. To save ourselves from poor decision making and potential embarrassment, should we abandon intuitive thinking in favour of rationality? Of course not, as we have all also analytically overthought a decision, and still made the wrong choice. A balance between intuitive and rational reasoning is necessary for optimal decision making.

Intuitive reasoning is fast and effortless but can lead to mistakes (Evans & Stanovich, 2013; Kahneman, 2011). In contrast, rational reasoning is slower and more accurate but takes considerable effort (Evans & Stanovich, 2013; Kahneman, 2011). We should not discredit intuitive reasoning, however, simply because it is potentially less accurate. Indeed, we rely on intuitive reasoning for the majority of our decisions as it most often leads to an acceptable outcome (Kahneman, 2011; Norman et al., 2017). Yet, intuitive reasoning can fail when an unexpected or difficult decision requires more thought and consideration (DeNeys & Pennycook, 2019; Norman et al., 2017; Pennycook, 2017). In these instances, we typically turn to rational reasoning to arrive at a suitable decision. As such, there is an intimate balance between intuitive and rational reasoning where we operate to conserve energy until we must exert effort when faced with a difficult decision (Inzlicht et al., 2014; Inzlicht & Schmeichel, 2012).

Our understanding of intuitive and rational decision making mostly stems from a large body of research focused on the resolution of word problems. For example, Frederick's (Frederick, 2005) classic widget problem: "*It takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets?*" Such problems have two common answers: one intuitive and incorrect (i.e., 100 minutes) and the other rational and correct (i.e., 5 minutes). Another example of reasoning problems are base-rate problems (Kahneman & Tversky, 1973; Pennycook et al., 2014): "*There are clowns and accountants. Person A is funny. There are 30 clowns and 70 accountants. Is person A a clown or an accountant?"* Here, both stereotype (Person A is funny) and base-rate (there are more accountants than clowns) information are presented – whereas an intuitive response adheres to the stereotype information, a rational response follows the base-rate information. Unlike Frederick's (Frederick, 2005) widget problem, base-rate problems do not have right or wrong answers – the person being a clown or an accountant are both possible – but instead these problems measure people's tendency to reason intuitively or rationally depending on the piece of information they base their decision on (Pennycook et al., 2014). When the presented stereotype and base-rate information are incongruent, strictly intuitive reasoners will respond using the stereotype information and strictly rational reasoners will respond using the base-rate information. Reasoners can also fall in between these two extremes – sometimes making decisions using the stereotype information and other times using the base-rate information. Thus, people who make the majority of their decisions using base-rate information are proposed to reason more rationally than those that rely more on stereotype information.

Neuroimaging has begun to provide insight into the neural systems that underlie intuitive and rational reasoning. For example, Bago and colleagues (Bago et al., 2018) used electroencephalography (EEG) to examine the neural underpinnings of human reasoning during the deliberation of base-rate problems. In this task, the stereotype and base-rate information may be congruent or incongruent, depending on whether they indicate the same response or not, respectively. Although the authors found neural responses that distinguished congruent from incongruent trials, their analyses did not specifically investigate differences between intuitive and rational reasoning (see also (Banks, 2017)). Paralleling these findings, Vartanian and colleagues (Vartanian et al., 2018) determined via functional magnetic resonance imaging (fRMI) that activity within the anterior and posterior cingulate cortices differed between congruent and incongruent base-rate problems (see also (DeNeys et al., 2008)). However, they too did not examine differences in neural activity between intuitive and rational responses. As such, the neural systems that underlie intuitive and rational reasoning, while hypothesized, are still not clear.

To address this gap, in the current study we used electroencephalography to probe the neural correlates of human rationality. We accomplished this by having participants complete base-rate problems (e.g., (Bago et al., 2018)) so that we could examine whether neural oscillations are sensitive to intuitive and rational reasoning (Williams, Kappen, et al., 2019). Furthermore, neural oscillations are theorized to correspond to the degree of communication across the brain – fast oscillations reflect localized brain activity and slow oscillations reflect the coordination of structures across a diverse brain network (Dai et al., 2017). As rational thinking involves coordination across multiple brain regions, we focused our analyses on neural oscillations in the theta band (4 – 8 Hz) (Williams, Kappen, et al., 2019) and found a strong relationship between these oscillations and rationality.

Methods

Participants

Thirty-three undergraduate students participated in this study. One participant was removed due to excessive noise in their EEG data and two others were removed as outliers determined via an evaluation of Cook's distances during analyses (see below). Thus, here we report 30 participants (M_{age} = 20.53 years old [19.61 years old, 21.46 years old]; 15 female, 15 male). Cognitive control and reward learning literature indicate theta effect sizes to be large and thus require sample sizes of 10 (Cohen & Donner, 2013) and 23 (Williams, Ferguson, Hassall, Abimbola, et al., 2021), respectively, to achieve 95% power – confirming our sample size which exceeds these calculations. All participants were recruited through an online sign-up system wherein they earned course credit within psychology courses, were native English speakers, and had normal or corrected-to-normal vision. All participants provided informed consent and this study was approved by the Human Research Ethics Board at the University of Victoria (Ethics Protocol Number: 16-428).

Apparatus and Procedure

Participants sat in a dark, sound-dampened room in front of a 22” LCD computer monitor and responded to the task using a ResponsePixx controller (VPixx, Vision Science Solutions, Quebec, Canada). The task (written in MATLAB version 8.6, Mathworks, Natick, MA, using the Psychophysics Toolbox Extension version 3.0.8 (Brainard, 1997)) consisted of a series of base-rate questions taken from work by Pennycook and colleagues (Pennycook et al., 2014). Base-rate questions introduce two populations (e.g., clowns and accountants), provides stereotype information (e.g., Person A is funny), provides base-rate information (e.g., there are 5 clowns and 995 accountants), and ask the participant to determine which population the indicated target (Person A) is from. With base-rate problems, the stereotype and base-rate information may be congruent or incongruent. For example, a congruent problem would indicate that Person A is funny and that there are 995 clowns and 5 accountants, while an incongruent problem would indicate that Person A is funny but that there are 5 clowns and 995 accountants. Within the congruent condition, there is an easy and obvious response; however, within the incongruent condition a participant may formulate their response dependent on the stereotype or the base-rate information.

On each trial, population information, stereotype information, base-rate information, and response options were presented sequentially with each lasting 1900 to 2100 ms. Problem information was presented sequentially with new information being appended to previously presented information. Presented information was spaced so that the base-rate information was presented in the center of the monitor (where participants were instructed to keep their gaze). Participants were provided up to 3000 ms to indicate (left, right on response controller) which population (e.g., accountant or clown) the target person was from. If they did not respond within this time limit, they were told they were 'TOO SLOW' via feedback presented on the computer display; note that too slow trials were removed from all analyses. Base-rate information was manipulated across extreme, moderate, and non-informative conditions: within the extreme condition, base-rates were 994/6, 995/5, 996/4, 6/994, 5/995, 4/996; within the moderate condition, base-rates were 69/31, 70/30, 71/29, 31/69, 30/70, 29/71; within the non-informative condition, base-rates were 50/50. Conditions were pseudo-randomly selected each trial across 11 blocks of 18 problems with the restriction that all conditions were presented 66 times, half of which were congruent and the other half incongruent. Prior to the task, participants were presented two practice problems. The current article considers a subset of data of the larger project just described. As such, the current analyses only consider the extreme and moderate base-rate conditions.

Data Acquisition and Processing

Behavioural data were collected and recorded using the ResponsePixx response box and MATLAB, respectively. EEG data were collected using a 32-electrode system with a standard

10-20 layout (ActiCAP, Brainproducts GmbH, Munich, Germany) and recorded using Brain Vision Recorder (Version 2.1., Brainproducts GmbH, Munich Germany). EEG data were sampled at 500 Hz, amplified, and filtered using an antialiasing low-pass filter of 245 Hz (ActiCHamp, Revision 2, Brainproducts GmbH, Germany). To ensure data quality, electrode impedances were on average kept below 10 kΩ, and to ensure temporal accuracy of event markers we used a DATAPixx stimulus timing unit (VPixx, Vision Science Solutions, Quebec, Canada).

Processing of EEG data was conducted using custom MATLAB scripts which can be found on github.com/neuro-tools (MATLAB-EEG-fileIO toolbox, MATLAB-EEG-preProcessing toolbox, and MATLAB-EEG-timeFrequencyAnalysis toolbox) and functions from EEGLAB (Delorme & Makeig, 2004). First, excessively noisy or faulty electrodes were removed, data were re-referenced to an averaged mastoid electrode, and data were filtered using a Butterworth filter (passband: 0.1 to 30 Hz, 4th order) and a notch filter (60 Hz, 2nd order). Next, eye blinks were removed using infomax independent component analysis (ICA) where ICA components were first determined and identified as eye blinks via visual inspection and then EEG data were reconstructed without the eye blink components. Then, removed electrodes in step one were interpolated using spherical splines, data was segmented from -700 to 1500 ms (note that the window of interest was -200 to 1000 ms and that we here added 500 ms on either side of this range to accommodate for time-frequency edge artifacts) surrounding markers of interest (the onset of the base-rate information), and baseline corrected using a -200 to 0 ms time window. Trials that violated artifact rejection criteria of 150 µV max-min were removed (total percentage of trials removed: 16% [12%, 20%]). After processing, we transformed EEG data into time-frequency data using a Gaussian-windowed complex Sine wave with a Morlet parameter of

6 cycles, a baseline correction of -200 to 0 ms, and normalization into 30 frequencies in linear steps from 1 to 30 Hz. Here, we only considered theta frequencies (4 – 8 Hz).

Data Analysis

The experimental paradigm afforded us with two within-subject factors: congruency (congruent, incongruent) between the stereotype and base-rate information and base-rate extremity (extreme, moderate, non-informative). As we were here investigating behavioural and neural responses when presented two sources of information (stereotype, base-rate), the non-informative base-rate conditions were not considered. Thus, we were left with congruent extreme, incongruent extreme, congruent moderate, and incongruent moderate conditions.

First, we determined participants' behavioural tendencies to respond intuitively (i.e., focusing on the stereotype information in incongruent trials) or rationally (i.e., focusing on the base-rate information in incongruent trials). We determined the percentage of rational responses as the percentage of times participants responded in line with the base-rate information on incongruent trials. Thus, 0% means that participants always responded in accord with stereotype information, 100% means that participants always responded using the base-rate information, and 50% means that participants responded with equal investment in both. We also extracted reaction time data for each participant as the average reaction time of incongruent extreme and moderate conditions (the same conditions used to determine the percentage of rational responses). To determine whether there was a relationship between rationality and reaction time, we conducted a Pearson R correlation (using the *corrcoef* function in MATLAB, alpha = .05). Assumptions of normality and homoskedasticity were met for these analyses.

Here, we focused our research question to determine if EEG activity was in general related to participant's tendency to respond rationally. EEG data extraction involved two steps:

1) determining clusters of interest, and 2) extracting oscillatory activity on a participant level. To determine clusters of interest, we first extracted participant level time-frequency activity by averaging each condition separately and then baseline correcting the incongruent conditions with the corresponding congruent conditions (e.g., incongruent moderate minus congruent moderate). These data were then averaged across conditions and participants and time-frequency clusters were determined as any data greater than 0.10 dB. We restricted clusters to frontal theta activity (4-8 Hz at electrode FCz) given our *a priori* hypotheses. Theta power was then extracted and averaged (across frequency and time) for each participant using the identified cluster. To determine whether there was a relationship between rationality and reaction time/theta power, we performed Pearson R correlation analyses (using the *corrcoef* function in MATLAB, alpha = .05). Assumptions of normality and homoskedasticity were met for these analyses.

Data and Code Availability

Data and code for this study can be found at https://osf.io/nyavt/.

Results

No Relationship Between Rationality and Reaction Times

To determine whether rationality was reflected in behavioural observations, we compared the percentage of times that participants responded rationally (*M* = 62% [54%, 70%], range = 8% to 98%) with reaction times. We found no relationship between rationality and reaction times, *r*(28) = 0.07, *p* = .7042, see Figure 1.

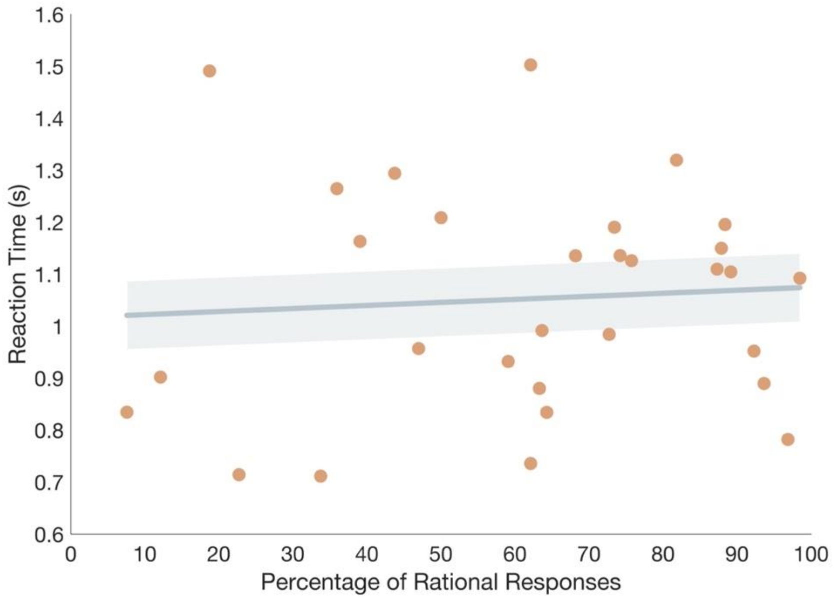

Figure 1: Experiment 1 (Figure 1)

No relationship between rationality and reaction time. Participant rationality was determined as the number of responses a participant made in line with base-rate rather than stereotype information in the incongruent trials. Blue line corresponds to a linear regression with a 95% confidence interval.

Strong Relationship Between Rationality and Theta Activity

Next, we determined whether participant rationality was related to frontal theta activity. Indeed, we found a theta cluster of interest at electrode FCz with the frequency range of 4 to 8 Hz and the time range of ~100 to ~600 ms post-base-rate information onset (see Figure 2 and Figure 3). There was a strong positive relationship between this frontal theta activity and participants' rationality, *r*(28) = 0.64, *p* = .0001 (see Figure 4).

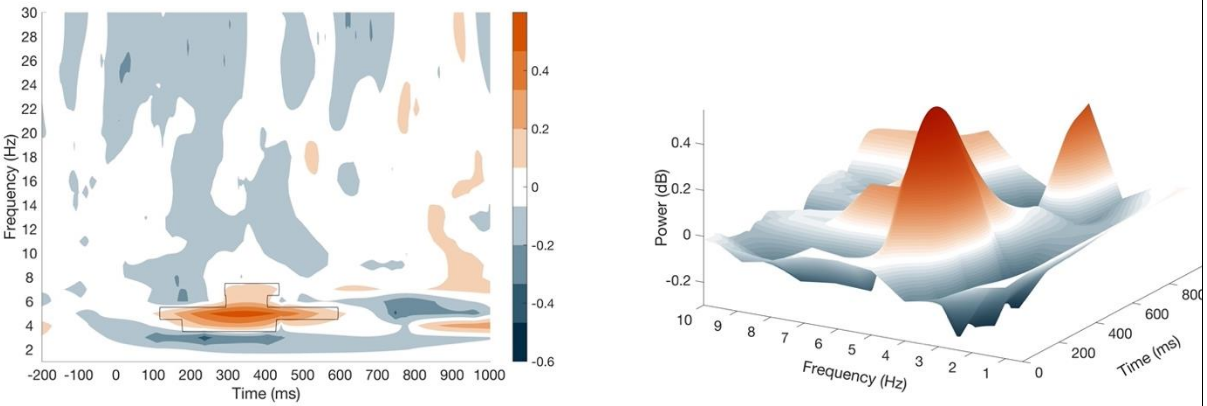

Figure 2: Experiment 1 (Figure 2)

Two and three-dimensional time-frequency wavelet plots demonstrating a frontal theta cluster of interest at electrode FCz. Activity are in decibel units. The cluster border is designated as any activity greater that 0.10 dB within the theta range (4 – 8 Hz). To better illustrate theta effects, the data within this figure was scaled using an asymmetric Hamming window which peaked at 5 Hz. An interactive supplemental figure for this data can be found at https://osf.io/nyavt/.

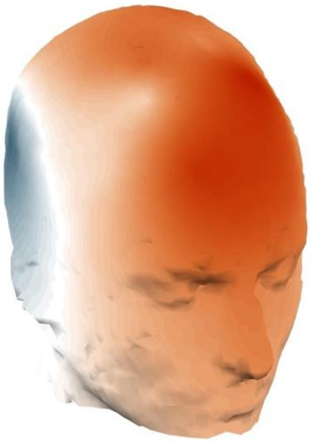

Figure 3: Experiment 1 (Figure 3)

Three-dimensional topographic map of theta activity illustrating activity to be frontal-central. An interactive supplemental figure for this data can be found at https://osf.io/nyavt/.

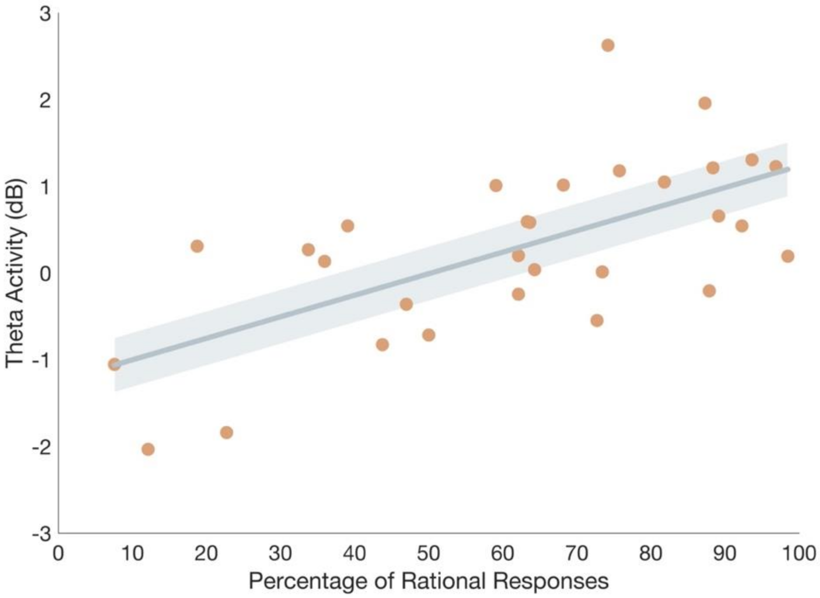

Figure 4: Experiment 1 (Figure 4)

A strong positive relationship between rationality and frontal theta activity. Participant rationality was determined as the number of responses a participant made in line with base-rate rather than stereotype information in the incongruent trials. Theta activity was extracted as the cluster of activity at electrode FCz as illustrated in Figure 2. Blue line corresponds to a linear regression with a 95% confidence interval. An interactive supplemental figure for this data can be found at https://osf.io/nyavt/.

Discussion

Our principal finding was that frontal theta oscillations were predictive of rationality. Specifically, increased frontal theta power was associated with increased rationality and vice versa. The relationship between theta oscillations and rationality that we observed here reinforces the proposition that high-level cognitive mechanisms including cognitive control and working memory play a critical role in reasoning (Cavanagh & Frank, 2014; DeLoof et al., 2019; Hsieh & Ranganath, 2014; Williams, Kappen, et al., 2019). Furthermore, as slow-wave brain oscillations such as those in the theta band are posited to reflect the communication of disparate brain regions, the relationship between theta oscillations and rationality also implies that rational reasoning involves increased coordination across the whole-brain in comparison to intuitive reasoning (Dai et al., 2017).

Our findings are in line with recent work (Williams, Kappen, et al., 2019) replicating and extending Kahneman and colleagues' (Kahneman et al., 1968) seminal findings on reasoning. Kahneman et al. presented participants with four numbers and had them either report those same four numbers (add-zero condition) or report the four numbers after adding one to each of them (add-one condition). Their pupillometry findings determined the add-one condition to require more effortful and rational thinking than the add-zero condition. Williams and colleagues (Williams, Kappen, et al., 2019) replicated these findings and further determined frontal theta activity to be larger for the add-one condition than the add-zero condition, signifying the involvement of cognitive control and working memory in reasoning. Moreover, Cavanagh and colleagues (Cavanagh & Shackman, 2015) investigated the neural underpinnings of conflict detection, a pre-cursor to cognitive control, in a Simon task. This task had participants making rightward and leftward responses contingent on the presented shape. These shapes may cause conflict when they are presented on the opposite side of the screen than the response. Indeed, frontal theta activity was enhanced in the conflict condition relative to the non-conflict condition, implying the need for cognitive control. Congruent with our findings here, these aforementioned studies concluded rational thinking to involve mechanisms of cognitive control and working memory as well as a coordination across the brain.

Our findings add to a growing body of EEG and fMRI literature (Bago et al., 2018; Banks, 2017; DeNeys et al., 2008; Vartanian et al., 2018) investigating the neural basis of reasoning. Our results are also in line with studies examining cognitive control and executive functioning that have probed the underlying mechanisms of reasoning by employing niche tasks to assess specific theorized mechanisms (Cohen, 2017). For example, the Stroop task has people reading the name of a colour that is printed in either the same or a different colour as the word; findings have shown reaction times to be longer when the name-colour pairings were incongruent (Monahan, 2001). This task is theorized to investigate focused attention, a core mechanism of reasoning (Monahan, 2001). Although the field of executive functioning is invaluable, the distilled nature of their corresponding tasks and findings can only anecdotally describe real-world reasoning (Cohen, 2017). As such, it is important to expand this field of research by investigating natural reasoning.

In summary, we demonstrated here that frontal theta oscillations were predictive of rationality. As such, theta activity may be a neural measure of rationality and could potentially be used in future research to further decipher the mental operations of reasoning. Moreover, as slow wave oscillations such as theta imply long-range communication across the brain, our results suggest that rational reasoning requires a diverse whole-brain network in comparison to intuitive reasoning.

Chapter 3: Experiment 2

Dissociated Neural Signals of Conflict and Surprise in Effortful Decision Making: Theta Activity Reflects Surprise while Alpha and Beta Activity Reflect Conflict.

Chad C. Williams[a]*, Thomas D. Ferguson[a], Cameron D. Hassall[a], Bruce Wright[b], and Olave E. Krigolson[a]

[a]Centre for Biomedical Research, University of Victoria, Victoria, British Columbia, V8W 2Y2, Canada

[b]Division of Medical Science, University of Victoria, Victoria, British Columbia, V8W 2Y2, Canada

Published in Neuropsychologia, 2021

Williams, C. C., Ferguson, T. D., Hassall, C. D., Wright, B., & Krigolson, O. E. (2021). Dissociated neural signals of conflict and surprise in effortful decision Making: Theta activity reflects surprise while alpha and beta activity reflect conflict. Neuropsychologia, 155, 107793.

Abstract

What makes a decision difficult? Two key factors are conflict and surprise: conflict emerges with multiple competing responses and surprise occurs with unexpected events. Conflict and surprise, however, are often thought of as parsimonious accounts of decision making rather than an integrated narrative. We sought to determine whether conflict and/or surprise concurrently or independently elicit effortful decision making. Participants made a series of diagnostic decisions from physiological readings while electroencephalographic (EEG) data were recorded. To induce conflict and surprise, we manipulated task difficulty by varying the distance between a presented physiological reading and the category border that separated the two diagnoses. Whereas frontal theta oscillations reflected surprise – when presented readings were far from the expected mean, parietal alpha and beta oscillations indicated conflict – when readings were near the category border. Our findings provide neural evidence that both conflict and surprise engage cognitive control to employ effort in decision making.

Introduction

We make countless decisions every day - some of which are easy while others are difficult. But what makes a decision difficult? It turns out that this is a complicated question entrenched in debate. Consider a clinician. Clinicians assess patient symptoms to develop diagnostic hypotheses. A sore throat may indicate a cold, a flu, or the measles. Measles is uncommon and can often be disregarded, but cold and flu are both likely diagnoses and a clinician may be conflicted when deciding between them. Likewise, clinicians are predisposed with expectations of what symptoms their patients may exhibit even prior to meeting them. Just by sheer statistics, clinicians may, for example, expect the average patient to arrive with a cough and a runny nose. When patients exhibit uncommon symptoms, however, the decision-making process becomes more complicated.

We then return to the pressing question: what makes a decision difficult? Decision difficulty refers to judgments that require additional cognitive effort and two key factors contributing to the effort of a decision are conflict and surprise. On one hand, conflict occurs when two or more options are similarly likely (Botvinick & Cohen, 2014; Egner, 2011, 2017; Nigg, 2017) – for example, when the clinician was first deliberating between a cold and a flu. To make a decision we consider different response options (e.g., diagnoses) each with a likelihood (or value) (Krajbich et al., 2010; Krajbich & Rangel, 2011; Tajima et al., 2016) and typically choose the most likely option over less likely options. However, if the likelihood for response options are similar it is difficult to choose between them and response conflict arises (Botvinick & Cohen, 2014; Egner, 2011; Nigg, 2017). On the other hand, we become surprised with unexpected events (Alexander & Brown, 2011; J. W. Brown, 2013; J. W. Brown & Alexander, 2017; Vassena, Deraeve, et al., 2017; Vassena et al., 2020) – such as when a patient demonstrates rare symptoms. In other words, easy decision making can rely on intuitive heuristics (DeNeys, 2017; Evans & Stanovich, 2013; Kahneman, 2011; Pennycook, 2017) but this strategy only functions as long as there is a clear response (i.e., no conflict) and everything subscribes to what is expected (i.e., nothing surprising). If conflict or surprise are present, intuitive decision systems are superseded by analytical decision systems that engage cognitive control (Evans & Stanovich, 2013; Kahneman, 2011).

Neural decision systems determine whether it is necessary to exert effortful top-down control (e.g., to resolve conflict or surprise) and, if so, direct the brain in doing so (Egner, 2017). One computational framework of decision making posits that determining whether to exert control involves a cost-benefit assessment (Alexander & Brown, 2011; J. W. Brown & Alexander, 2017; Kool et al., 2017; Kool & Botvinick, 2018; Shenhav et al., 2013, 2014; Vassena, Holroyd, et al., 2017; Vassena et al., 2020). Whereas cost refers to the expenditure of resources needed to make a difficult decision, benefit is the progress that would be achieved by exerting top-down influence. This model describes that if the benefit outweighs the cost, cognitive control and top-down bias is exerted, but if the cost outweighs the benefit, decisions are made without additional effort (Alexander & Brown, 2011; J. W. Brown & Alexander, 2017; Kool et al., 2017; Kool & Botvinick, 2018; Shenhav et al., 2013, 2014; Vassena, Holroyd, et al., 2017; Vassena et al., 2020).

Within this cost-benefit framework, conflict operates as a cost of decision making – enforcing the need for increased incentives in fear of relying on faulty intuitions (Shenhav et al., 2014). Surprise signals, on the other hand, affect decision making in a nuanced way. As described by another computational framework of decision making (Alexander & Brown, 2011), humans have a pre-determined degree of effort which we expect to employ during decision making. If these predictions are mistaken, surprise signals index a discrepancy between our expected effort and the actual needed effort and we update our expectations for future decision making.

In sum, computational accounts of decision making posit that conflict signals may operate to resolve decision making in the moment (Shenhav et al., 2014). and surprise signals may function to reduce future difficulty (Alexander & Brown, 2011). With that said, rather than considering conflict and surprise in parallel, researchers have pitted them against each other in pursuit of a single account of effortful decision making (Alexander & Brown, 2011; J. W. Brown & Alexander, 2017; Kool et al., 2017; Kool & Botvinick, 2018; Shenhav et al., 2013, 2014; Vassena, Holroyd, et al., 2017; Vassena et al., 2020). For example, during an investigation using fMRI, Vassena and colleagues (Alexander & Vassena, 2020; Vassena et al., 2020) discerned patterns of conflict and surprise through computational modelling wherein conflict was modelled using the Expected Value of Control (EVC) framework (Shenhav et al., 2013) and surprise was modelled as the Predicted Response Outcome (PRO) framework (Alexander & Brown, 2011). Their intent was to determine which of these frameworks independently explained neural patterns of effortful decision making. They concluded their neural data to fit the PRO model rather than the EVC model and thus effortful decision making to reflect surprising events rather than conflicting responses. Others have, however, re-interpreted their findings to have reflected a combination of conflict and surprise (Shenhav et al., 2020), indicating a need to further consider both mechanisms in parallel.

In the current study, we examined whether conflict and surprise signals independently or concurrently influence control demands and the employment of effort in decision making. In other words, we investigated whether one of conflict or surprise guides effortful decision making alone or whether the two operate in parallel. Participants were to diagnose virtual patients with one of two diseases based on a physiological reading. Each disease was characterized by a unique range of the reading and task difficulty was manipulated as the distance between the presented reading and the category border that separated the two diseases (see Figure 5). Thus, difficulty was highest near this border and decreased as readings diverged from it.

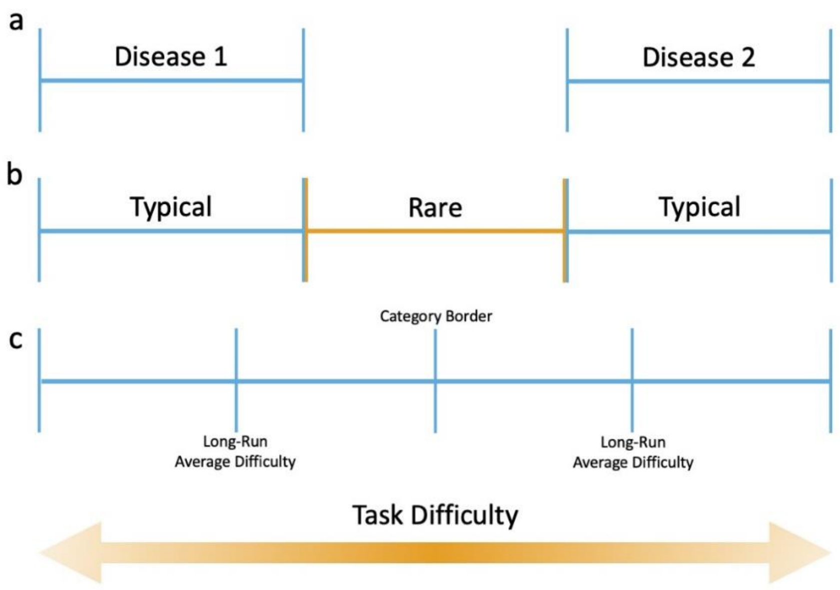

Figure 5: Experiment 2 (Figure 1)

Depictions of disease ranges in the learning phase (a) and the decision making phase (b) and how they correspond to task difficulty (c). Category border signifies where the range of one disease ended and the other began. The double-sided arrow is a model where the brightness corresponds to the degree of difficulty.

Our analyses first focused on determining whether our task varied control demands, which would prompt the adjustment of decision making strategies via cognitive control (Jiang et al., 2015). The demand for control increases when there exists conflict between responses

(Shenhav et al., 2013) – in our task the two diagnoses conflicted when the readings were near the category border (i.e., when difficulty was high). Moreover, demands change with changing contexts (Shenhav et al., 2013). In our task, these changes take the form of congruency sequence effects wherein changing demands across trials (incongruent trials, e.g., a conflict trial preceded by a no-conflict trial) elicit increased need for control than consistent demands across trials (congruent trials, e.g., a conflict trial preceded by a conflict trial) (Egner, 2007). Indications of changing control demands correspond to decreased accuracy rates and increased reaction times in the presence of conflict relative to no-conflict as well as in the presence of incongruent relative to congruent consecutive trials (Egner, 2007; Shenhav et al., 2013). Accordingly, we hypothesized that accuracy rates and reaction times would index changing demands in our task by adopting these patterns.

Next, we considered whether neural components of decision making reflected conflict and/or surprise. Within our task, conflict was highest when the presented physiological reading was near the category border and decreased as a function of the distance from this border (see Figure 5). Conflict was highest at the category border because these readings almost equally corresponded to each diagnosis and was lowest away from this border because one diagnosis was increasingly more likely than the other. Correspondingly, conflict was low at the lowest extreme of the disease 1 range, increased to a maximum at the category border, and then decreased back to low at the highest extreme of the disease 2 range (see Figure 5c). Statistically, this pattern of conflict could be described as a quadratic polynomial (see Figure 6).

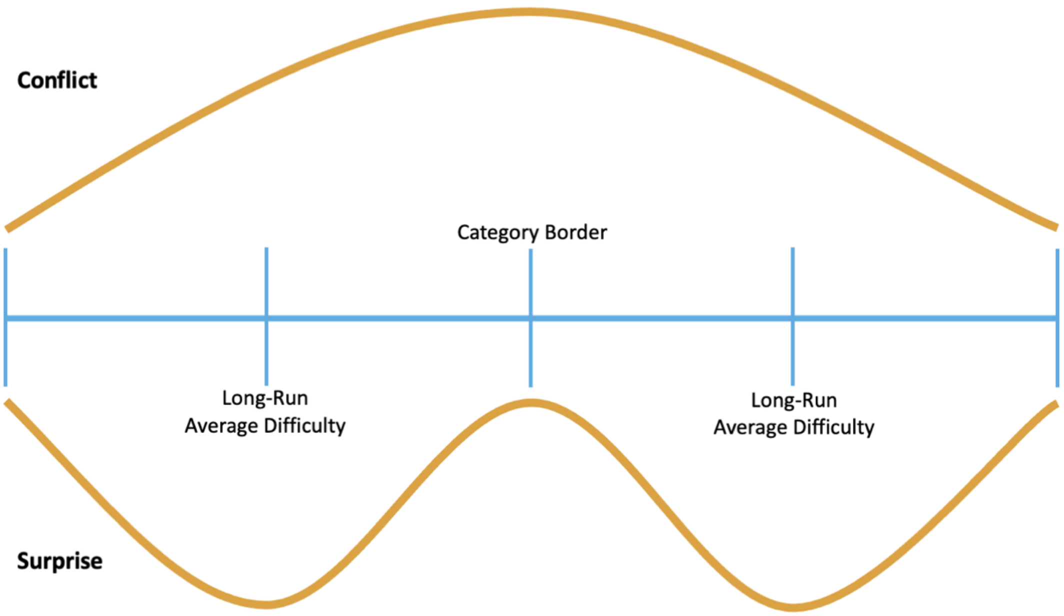

Figure 6: Experiment 2 (Figure 2)

Polynomials describing predictions of conflict (top orange line) and surprise (bottom orange line). The x-axis corresponds to task difficulty as depicted in Figure.

In addition, as the physiological readings were determined randomly from a uniform distribution on each trial, the long-run average difficulty, and thus the most expected and least surprising readings, for the diagnoses would be within the center of each diagnostic range (see Figure 5c). Surprise was lowest in the center of each range because participants learned to expect an average difficulty (or level of needed effort) of the task, and deviations from this expectations were unexpected and surprising. Thus, surprise was lowest within the middle of each disease's range and increased as the readings deviated from the center – i.e., towards the outer extremes and towards the category border. Statistically, this pattern of surprise could be described as a quartic polynomial (see Figure 6).

Here, we investigated oscillatory patterns in the brain corresponding to theta (3 – 8 Hz), alpha (8 – 14 Hz), and beta (14 – 20 Hz) rhythms while participants made decisions as modulations of these EEG rhythms have all been linked to the engagement of cognitive control (Cavanagh & Frank, 2014; DeLoof et al., 2019; Lin et al., 2018; Williams, Kappen, et al., 2019). Although there has been neural evidence for both conflict and surprise signals within effortful decision making (Cavanagh & Frank, 2014), they are rarely dissociated and investigated concurrently (Lin et al., 2018; Vassena et al., 2020). Our research here was exploratory thus we had no *a priori* hypothesis as to whether there would be individual or concurrent neural signals of conflict and surprise when making difficult, or else more effortful, decisions. Informally, however, in line with past research (Cavanagh, Zambrano-Vazquez, et al., 2012; DeLoof et al., 2019; Engel & Fries, 2010; Lin et al., 2018; Williams, Kappen, et al., 2019) we predicted that theta would be positively associated with conflict and surprise and that alpha and beta would be negatively associated with conflict but not associated with surprise.

Experimental Methods

Participants

Thirty-three undergraduate students from the University of Victoria participated in the experiment. Three participants were removed due to excessively noisy frontal data leaving us with thirty participants (21.47 years old [19.60, 23.33], 19 female, 10 male, 1 undisclosed). All participants had normal or corrected-to-normal vision and volunteered for extra course credit in a psychology course. Participants all provided informed consent approved by the Human Research Ethics Board at the University of Victoria.

Experimental Design

Participants were seated in a sound dampened room, viewed stimuli on a 19” LCD computer monitor, and responded using a handheld 5-button RESPONSEPixx controller (VPixx, Vision Science Solutions, Quebec, Canada). The task was written in MATLAB (version R2017b, Mathworks, Natick, U.S.A.) using the Psychophysics Toolbox extension (Version 3.0.14; (Brainard, 1997)).

Participants completed a simplified version of a decision making task used by Williams and colleagues (Williams et al., 2018) (see also (Bannister et al., 2016; Burak et al., 2015; Horrey et al., 2016; Kazoleas, 2016; Tang et al., 2016)). On each trial, they were presented with a simulated medical case including one physiological measure, Alkaline Phosphatase, and were to decide whether their virtual patient had general hepatocellular liver disease or cholestatic intrahepatic biliary disease (see Figure 7). Here, we use the term ‘medical case’ to be consistent with the aforementioned research. Although the current cases contained a physiological reading that is pertinent in diagnosing these diseases in real-world settings, we would like to note that the simplicity of what is presented in this research does not adequately reflect the variety of information used to diagnose real patients. Each disease had a unique range of readings (see Figure 5a) and on each trial a reading for the current disease was randomly determined from a uniform distribution of the corresponding disease range. The ranges were each 115 values wide and were shifted by a random number between 0 and 654 so that no two participants conducted the task with the same readings. Further, the diseases to which each range were representative of were counterbalanced across participants. Although participants were informed that each disease was represented by a range of numbers, they were not told the corresponding ranges, nor the range widths.

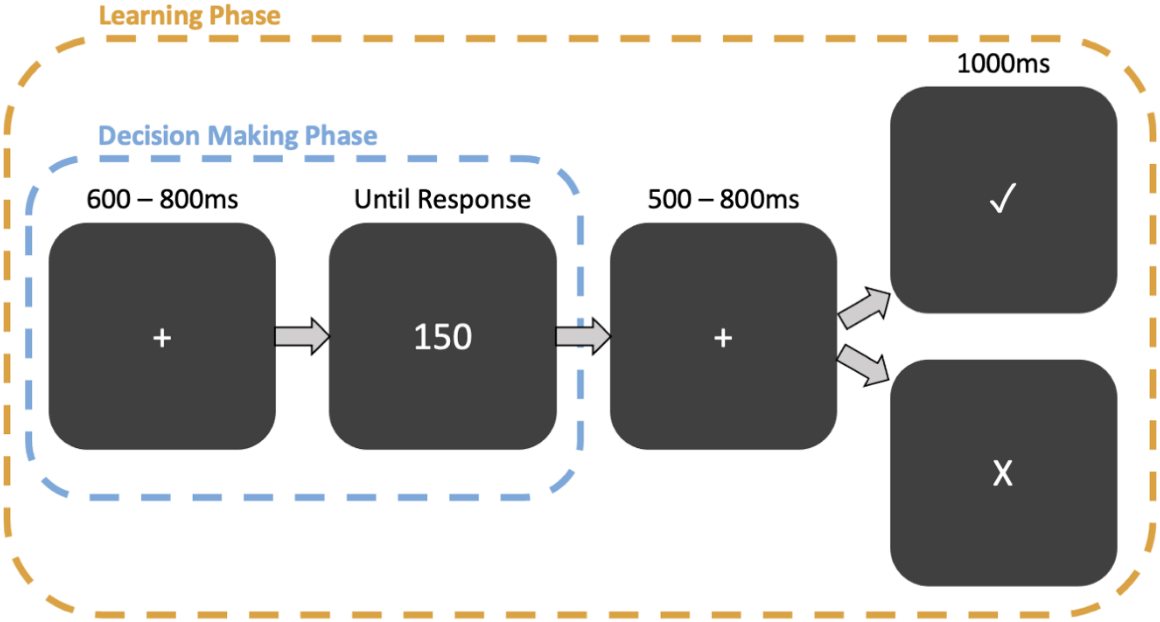

Figure 7: Experiment 2 (Figure 3)

Task paradigm demonstrating stimuli and timing for both the learning phase (orange) and the decision making phase (blue).

On each trial, participants saw a white fixation cross on a dark grey background for 600 to 800 ms which was followed by a white number representing the physiological measure of the patient. With no time limit, they were to then select one of the two diseases by pressing the left or right button on the response box. Which button represented which disease was counterbalanced and at the beginning of each block participants were reminded which button corresponded to which disease. The experiment was separated into two phases: a learning phase and a decision making phase (see Figure 7). In the learning phase participants would then see another white fixation cross for 500 to 800 ms which was followed by a one second presentation of a white ✓ or X indicating correct or incorrect, respectively. When participants achieved two consecutive blocks of twenty trials with an accuracy rate of 90% or higher, they moved into the decision making phase of the experiment.

Within the decision making phase, no feedback was presented and participants saw both typical and rare medical cases. Specifically, the readings presented to participants varied as a function of whether they were from the same ranges as in the learning phase (typical cases) or from a new range that had not previously been learned (rare cases), see Figure 5b. To avoid confusion, we want to emphasize that when using the terms typical and rare, we do not mean frequent and infrequent stimuli, but are rather using terminology that is coherent with the medical literature. The rare cases consisted of physiological readings that fell in between the two ranges that had been previously learned. Participants were told that this phase of the experiment was simply to practice what they had learned and were not told that the ranges of the diseases had been broadened nor that there would be unfamiliar cases. In this phase, participants completed ten blocks of twenty trials, half of which were typical cases and the other half rare cases. Although they were not presented feedback of their performance on each trial, they were informed of their performance after five blocks and at the end of the ten blocks.

Data Acquisition and Processing

Accuracy rates and reaction times were recorded using the response box. Behavioural responses were analyzed as markers of changing control demands. First, the typical cases and rare cases were analogous to conflict. In line with categorization literature, there exists a category border wherein readings equally activate each response option (Ratcliff et al., 2016). Rare cases straddle this border, resulting in conflict between the responses. In contrast, the typical cases encompass highly trained readings far from this border and result in no conflict. Thus, to confirm the manipulation of control demands in our paradigm, accuracy rates must decrease and reaction times increase for conflict relative to no conflict trials (Shenhav et al., 2013). Second, an additional marker of changing demands is the congruency sequence effect (Egner, 2007). Congruent trial types (i.e., typical – typical or rare – rare) result in higher accuracy rates and quicker reaction times than incongruent trial types (i.e., typical – rare or rare – typical). Thus, to further confirm changes of demand in our task would be to observe these behavioural patterns. Finally, we additionally analyzed these behavioural measures across task difficulty but describe these procedures below in section 2.4. Statistical Analyses.

EEG data were recorded from a 32 electrode EEG system (ActiCAP, Brain Products, GmbH, Munich, Germany) using Brain Vision Recorder (Version 1.10, Brain Products GmbH, Munich, Germany). Electrodes were all initially referenced to a common ground, impedances were on average kept below 20 kΩ, data were sampled at 500Hz, and filtered using an antialiasing low-pass filter of 245 Hz through an ActiCHamp amplifier (Revision 2, Brain Products GmbH, Munich, Germany). To ensure precise temporal resolution, we synced EEG markers and stimuli through a DataPixx processing box (VPixx, Vision Science Solutions, Quebec, Canada).

All EEG data were first processed in Brain Vision Analyzer (version 2.1.2.327, Brain Products GmbH, Munich, Germany). Excessively noisy or damaged electrodes were removed, and data was down-sampled to 250 Hz, re-referenced to an average mastoid, run through a dual-pass Butterworth filter (pass band: 0.1 Hz to 30 Hz, 4th order), and a notch filter of 60 Hz. To identify and remove blink artifacts, data were put through a restricted infomax independent component analysis (ICA) with classic PCA sphering, components were visually identified by component head maps and related factor loadings, and artifacts were removed via an ICA back transformation. Electrodes that had initially been removed were then interpolated using spherical splines.

All EEG data were then exported to a MATLAB format where the remainder of processing took place. Within MATLAB, data were segmented -500 to 1500 ms centered on markers of interest. The markers of interest coincided with the onset of medical cases. Next, artifact rejection with absolute difference of 200 μV and/or 20 μV/ms gradient criteria was applied. We then conducted wavelet analyses on individual trial data (Gaussian-windowed complex sine wave with a Morlet parameter of 6 for frequencies 1 to 30 in 30 linear steps, no baseline was used; script can be found at www.github.com/neuro-tools; also see (Brainard, 1997; Cohen et al., 2008)), and standardized the data within each participant. Standardization was completed for each electrode at each frequency within each participant by subtracting the mean of each frequency across conditions and dividing the output by the standard deviation of each frequency across conditions.

Within the time-frequency wavelet data, we first determined clusters of interest using the collapsed localizer approach (Luck & Gaspelin, 2017). This approach consists of averaging data across all conditions and identifying time-frequency clusters of interest. The frequency and time-width of these clusters were constrained to the contour lines of the plotted data (see Results). This approach resulted in two clusters of interest. The first cluster was at electrode FCz spanning frequencies 3 to 8 Hz within the time range of ~250 to 750 ms. Although this cluster's frequency range is in line with the theta band, we here caution from using frequency band terminology as it may constrain our ability to interpret time-frequency analyses (Haller et al., 2018). Hereafter, we will simply refer to this cluster as the frontal cluster. The second cluster was maximal at electrode POz and spanned frequencies 7 to 20 Hz within the time range of ~350 to 650 ms. This frequency range corresponds to frequency bands alpha and beta. Hereafter, this cluster will be referred to as the parietal cluster.

The task difficulty of each trial was then determined as the difference between the physiological reading presented and the category border, see Figure 5c. This measure corresponded to response conflict wherein readings near the category border would indicate both diseases as likely diagnoses. As the readings moved away from this border, one response became increasingly more favourable than the other, reducing task difficulty. This resulted in a measure of task difficulty on a continuous scale which we then segmented into 21 bins.

Statistical Analysis

All statistics were conducted in R (Version 4.0.0, the R Foundation, Vienna, Austria) (RStudioTeam, 2016) using RStudio (Version 1.1.463, RStudio Inc., Boston, U.S.A.) (RStudioTeam, 2016). All figures were created using the R package ggplot2 (Wickham, 2016) with the exception of the time-frequency wavelets and topographic maps which were created in MATLAB using EEGLab (Version 2019) (Delorme & Makeig, 2004).

For accuracy rates and reaction times, conflict and congruency sequence effects were determined with two-tailed, repeated-measures *t*-tests (alpha = .05), 95% within-subject credible intervals (Nathoo et al., 2018), and a Cohen’s *d* effect size. The assumption of normality was violated for both analyses of accuracy rates, and the assumption of homogeneity of variance was violated for accuracy analyses of conflict. For consistency, all behavioural analyses were then conducted using a Welch’s t-test.

Our manipulation of task difficulty provided opportunity to determine whether behavioural data confirmed that our task involved changing control demands and whether neural signals reflected conflict and/or surprise. For the former, changes in demand would be confirmed as any relationship between behavioural measures and task difficulty. For the latter, whereas conflict signals would decrease proportionally to the distance from the category border, or in other words via a second order quadratic polynomial, surprise signals would increase with deviations from the long-run average task difficulty, or else a fourth order quartic polynomial, see Figure 6. For accuracy rates, reaction times, frontal oscillatory clusters, and parietal oscillatory clusters, then, we employed linear mixed-effects modelling techniques (lme4 package) (Bates et al., 2015); lmerTest package (Kuznetsova et al., 2017) to determine quadratic fits:

$$Measure \sim \beta_0 + \beta_1 Bin + \beta_2 Bin^2 + (1|Participant) + \varepsilon \quad (1)$$

and quartic fits:

$$Measure \sim \beta_0 + \beta_1 Bin + \beta_2 Bin^2 + \beta_3 Bin^3 + \beta_4 Bin^4 + (1|Participant) + \varepsilon \quad (2)$$

Where *Measure* referred to accuracy, reaction time, frontal cluster activity, and parietal cluster activity and *Bin* referred to the binned task difficulty. To determine the fit of each model to neural data, we conducted a model comparison with a null model using chi-square difference tests:

$$Measure \sim \beta_0 + (1|Participant) + \varepsilon \quad (3)$$

For behavioural data, we determined that the presence of a significant quadratic and/or quartic fit would confirm that these measures are sensitive to task difficulty and thus provide additional evidence that our task elicited changing control demands. For neural data, we compared the quadratic model to the quartic model using chi-square difference tests to determine whether they reflected conflict or surprise signals. Although quadratic and quartic polynomials were the focus of our neural analyses due to their description of conflict and surprise, respectively, we also conducted linear mixed-effects modelling investigating linear and cubic polynomial fits. One note is that alpha and beta activity, as encompassed within the parietal cluster, are negatively associated with cognitive control (DeLoof et al., 2019; Engel & Fries, 2010; Williams, Kappen, et al., 2019) – increased activity within the parietal cluster signifies reduced control. Correspondingly, the patterns reflected in Figure 6 would then need to be interpreted inversely when considering the parietal cluster activity.

Data and Code Availability

In line with open science policies, all data and scripts (analysis, plotting, and statistics) used for this manuscript can be found at www.osf.io/a65sh/.

Results

Behavioural Findings Confirm the Recruitment of Cognitive Control

First, we sought to determine whether our paradigm engaged cognitive control in relation to changing control demands. Participants made a diagnostic decision between one of two diseases based on a physiological reading for a virtual patient. Each disease was characterized by a specific value range for the physiological reading and thus we manipulated task difficulty by varying the distance to which the presented readings was from the category border (see Figure 5). This manipulation of task difficulty afforded us the ability to investigate both conflict and congruency sequence effects. Specifically, readings near the category border would elicit conflict between the two diagnostic responses (conflict condition), and readings far from the border would not (non-conflict condition). Moreover, trials could be considered as congruent (e.g., conflict – conflict) or incongruent (e.g., non-conflict – conflict) depending on the demands of the preceding trial.

Effects of accuracy rate and reaction time when manipulating both conflict and trial-to-trial congruency, confirmed the manipulation of control demands (see Figure 8). Specifically with regard to conflict manipulations, accuracy rates were greater in the non-conflict condition (97% [96%, 98%]) relative to the conflict condition (80% [77%, 84%]), M_d = 17% [13%, 20%], *t*(29) = 7.94, *p* < .0001, *d* = 1.71 [1.03, 2.39]. Correspondingly, reaction times were quicker for the non-conflict condition (689 ms [646 ms, 733 ms]) relative to the conflict condition (783 ms [725 ms, 841 ms]), M_d = -93 ms [-123 ms, -63 ms], *t*(29) = -5.31, *p* < .0001, *d* = -0.51 [-0.71, -0.30]. Similarly, accuracy rates were higher in congruent trials (90% [88%, 92%]) relative to incongruent trials (88% [85%, 90%]), M_d = 2% [1%, 4%], *t*(29) = 2.73, *p* = .0105, *d* = 0.31 [0.08, 0.54]. Lastly, reaction times were quicker for congruent trials (716 ms [670 ms, 763 ms]) than for incongruent trials (755 ms [703 ms, 807 ms]), M_d = -39 ms [-58 ms, -19 ms], *t*(29) = -3.41, *p* = .0019, *d* = -0.23 [-0.37, -0.09].

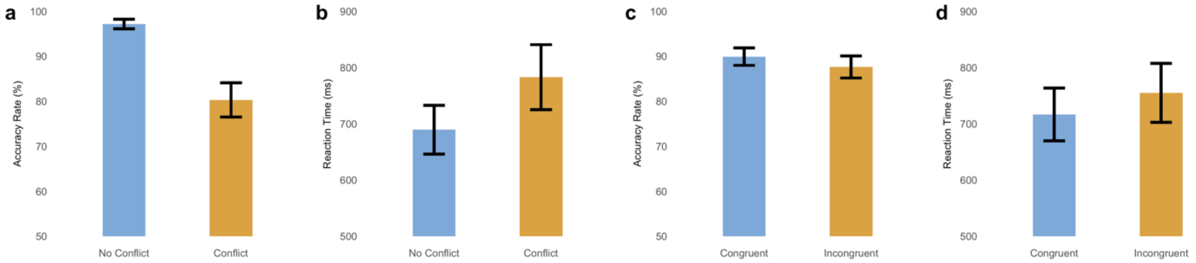

Figure 8: Experiment 2 (Figure 4)

Behavioural results reflect common findings of cognitive control – i.e., conflict effects (a and b) and congruency sequence effects (c and d). The behavioural effects confirm the recruitment of cognitive control within the task.

In addition to these analyses, we investigate accuracy rates and reaction times in a continuous manner using linear mixed-effects models. For these analyses the assumptions of linearity and homoskedasticity were met; however, the assumption of normality of residuals was violated. As linear mixed-effects models are robust to violations of residual normality (Winter, 2013), no corrections were made. These analyses present additional evidence that our task indeed elicited changing demands as demonstrated by both quadratic and quartic fits for accuracy rates (quadratic: $AIC = 52976$, $X^2(2, n = 30) = 356.61$, $p < .0001$; quartic: $AIC = 52753$, $X^2(4, n = 30) = 583.97$, $p < .0001$) and reaction times (quadratic: $AIC = 81304$, $X^2(2, n = 30) = 109.13$, $p < .0001$; quartic: $AIC = 81291$, $X^2(4, n = 30) = 126.08$, $p < .0001$), see Figure 9.

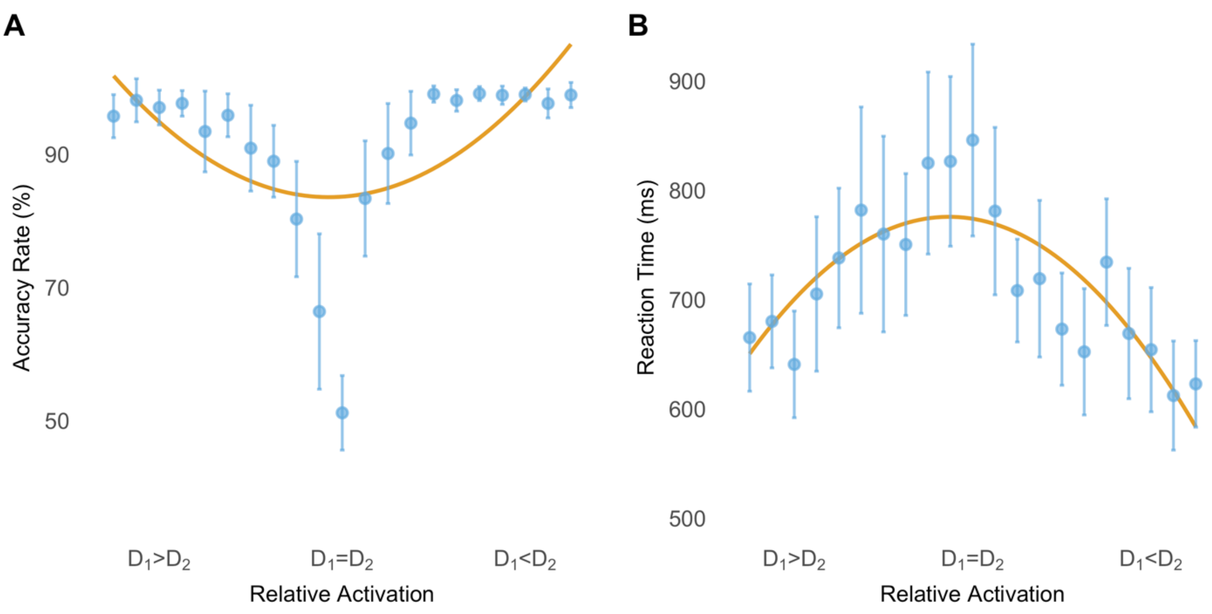

Figure 9: Experiment 2 (Figure 5)

Accuracy and reaction time reflect changing task difficulty. The x-axis corresponds to twenty-one bins of task difficulty as depicted in Figure 5c, ranging from favour for disease one to favour for disease two. Each point corresponds to a mean with 95% within-subject credible intervals. Orange lines reflect quadratic polynomials.

Frontal Activity Reflects Surprise and Parietal Activity Reflects Conflict

Our task affords us the opportunity to investigate task difficulty as a continuous function of the deviation between the readings and the category border with difficulty increasing as readings approached the border (see Figure 5c). Conflict signals could then be considered continuously in proportion to task difficulty and surprise signals could be considered as deviations from average task difficulty. Recall that conflict is highest at the category border because these readings almost equally indicate both diagnoses and decreases as the readings deviate from this border. Also recall that the best predictor of each diagnosis is the mean of their corresponding range, thus these values are most certain to participants and deviations from each range's center increases surprise. Statistically then, conflict would fit a quadratic pattern in the data while surprise would fit a quartic pattern (see Figure 6).

EEG time-frequency measures were fit to polynomials reflecting decision making model predictions using linear mixed-effect models. Specifically, frontal cluster power (reflective of theta band activity) and parietal cluster power (reflective of alpha and beta band activity) were identified (see Figure 10) and fit to quadratic and quartic polynomials to determine their correspondence with conflict and surprise, respectively (see Figure 11). The assumptions of linearity and homoskedasticity were met for all models, and the assumption of normality of residuals was violated for all models. As linear mixed-effects models are robust to violations of this latter assumption (Winter, 2013), no corrections were made.

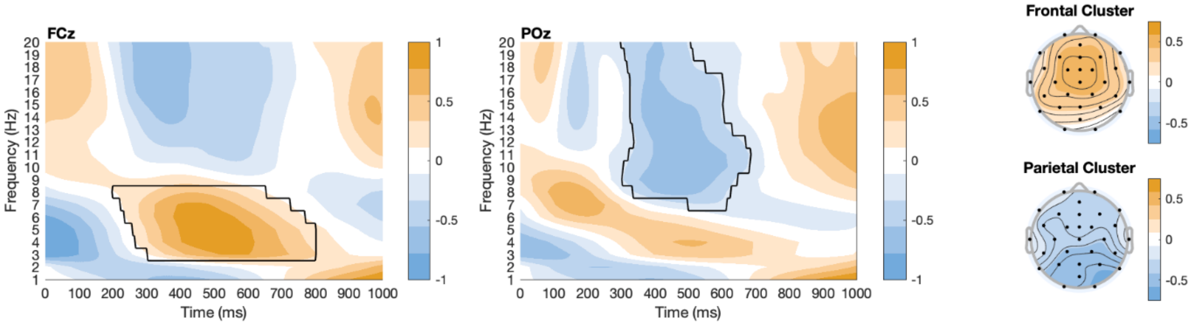

Figure 10: Experiment 2 (Figure 6)

Two clusters of interest at frontal (FCz) and parietal (POz) locations of the scalp. Standardized time-frequency wavelets correspond to data averaged across all trials. The frontal cluster reflects activity within the theta band (3-8 Hz) and the parietal cluster reflects activity within the alpha and beta bands (7-20 Hz). Black contour lines correspond to extracted regions for analysis. Topographic map values were extracted from time-frequency analyses in correspondence with the contour lines. Colour bars are in units of power (μV^2).

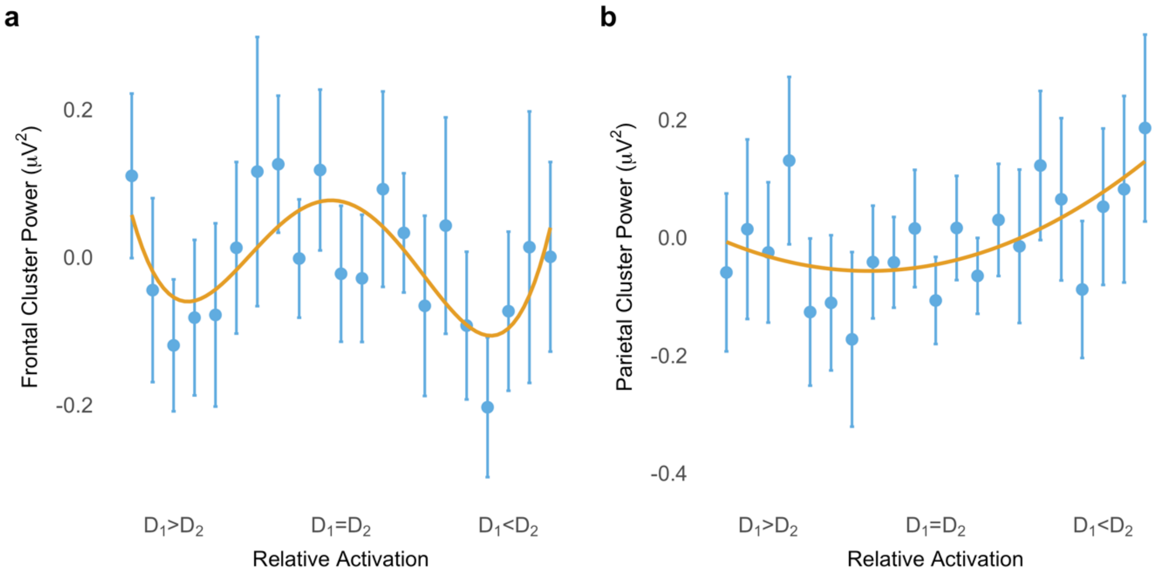

Figure 11: Experiment 2 (Figure 7)

Frontal cluster power (electrode FCz, frequencies 3-8 Hz) reflects surprise and parietal cluster power (electrode POz, frequencies 7-20 Hz) reflects conflict. Standardized power for the frontal and parietal cluster were determined by the clusters of Figure 20. The x-axis corresponds to twenty-one bins of task difficulty as depicted in Figure 5c, ranging from favour for disease one to favour for disease two. Each point corresponds to a mean with 95% within-subject credible intervals. Orange lines are best fit polynomials.

Frontal power was best fit by the quartic polynomial (quadratic: *AIC* = 15575, $X^2(2, n = 30) = 3.69$, *p* = .1581; quartic: *AIC* = 15568, $X^2(4, n = 30) = 14.71$, *p* = .0053; quartic versus quadratic: *ΔAIC* = -7.02, $X^2(2, n = 30) = 11.02$, *p* = .0040). Parietal power was best fit by the quadratic polynomial (quadratic: *AIC* = 15570, $X^2(2, n = 30) = 9.69$, *p* = .0079; quartic: *AIC* = 15574, $X^2(4, n = 30) = 9.91$, *p* = .0420; quartic vs quadratic: *ΔAIC* = 3.78, $X^2(2, n = 30) = 0.22$, *p* = .8939). We also provide a full complement of analyses comparing linear, quadratic, cubic, and quartic trends (1) to a null model and (2) to each other in Tables 1 and 2, respectively. Note that alpha and beta bands of the parietal cluster activity are inversely related to cognitive control, thus higher scores indicate lower control.

Table 1: Experiment 2 (Table 1)

Model fits of linear, quadratic, cubic, and quartic trends for frontal and parietal neural clusters. Chi-squared and p-values correspond to model comparison results with a null model. AIC = Akaike Information Criterion, BIC = Bayesian Information Criterion.

	AIC	**BIC**	X^2	***p-value***
Frontal Cluster				
Linear	15575	15601	1.58	.2092
Quadratic	15575	15608	3.69	.1581
Cubic	15576	15616	3.87	.2759
Quartic	15568	15614	14.71	.0053
Parietal Cluster				
Linear	15573	15599	5.05	.0247
Quadratic	15570	15603	9.69	.0079
Cubic	15572	15612	9.74	.0209
Quartic	15574	15620	9.91	.0420

Table 2: Experiment 2 (Table 2)

Model fits differences and statistical results when comparing across linear, quadratic, cubic, and quartic trends for accuracy, reaction time, the frontal neural cluster, and the parietal neural clusters. AIC = Akaike Information Criterion, BIC = Bayesian Information Criterion.

	ΔAIC	**ΔBIC**	**X^2**	***p-value***
Frontal Cluster				
Linear vs Quadratic	0	-7	2.11	.1461
Linear vs Cubic	-1	-15	2.29	.3178
Linear vs Quartic	7	-13	13.13	.0044
Quadratic vs Cubic	-1	-8	0.18	.6710
Quadratic vs Quartic	7	-6	11.02	.0040
Cubic vs Quartic	8	2	10.84	.0010
Parietal Cluster				
Linear vs Quadratic	3	-4	4.64	.0312
Linear vs Cubic	1	-13	4.69	.0957
Linear vs Quartic	-1	-21	4.87	.1820
Quadratic vs Cubic	-2	-9	0.05	.8186
Quadratic vs Quartic	-4	-17	0.22	.8939
Cubic vs Quartic	-2	-8	0.17	.6785

Discussion

Research has struggled to parsimoniously describe the factors that lead to effortful decision making (Alexander & Brown, 2011; J. W. Brown & Alexander, 2017; Kool & Botvinick, 2018; Shenhav et al., 2013, 2014; Vassena, Deraeve, et al., 2017; Vassena et al., 2020). Here, we had participants make diagnostic decisions based on a physiological reading of a series of virtual patients while we recorded EEG data. The proximity between the presented physiological reading and the category border served as a manipulation of task difficulty.

Specifically, when the presented reading was near the border that distinguished one diagnosis from the other, the likelihood of each diagnosis being correct was near equal and thus the decision was more difficult. In contrast, readings that were further from this border resulted in one diagnosis being more likely than the other thus making the decision easier. As such, our manipulation afforded us the ability to assess predictions of conflict and surprise in that conflict signals increased proportionally to task difficulty and surprise signals increased as a factor of deviation between the presented reading difficulty and the overall average task difficulty.

First, we determined that our paradigm indeed manipulated control demands and thus elicited effortful decision making by inspecting behavioural tendencies between the conflict and no-conflict conditions as well as between congruent and incongruent trials. Specifically, accuracy rates were lower and reaction times were higher for conflict trials relative to no-conflict trials, as well as, for incongruent trials relative to congruent trials. Changing control demands were additionally confirmed as accuracy rates and reaction times varied continuously across task difficulty. Next, we identified two oscillatory EEG patterns of interest: a frontal cluster of EEG activity within the theta range and a parietal cluster of EEG activity within the alpha and beta ranges. For each of these, we determined their reflection of conflict (as reflected by a quadratic polynomial) and surprise (as reflected by a quartic polynomial) via linear mixed-effects modelling. Frontal activity fit a quartic function rather than a quadratic function, indicating that it reflected signals of surprise rather than conflict. In contrast, parietal activity fit a quadratic function above a quartic function, determining that it reflected signals of conflict rather than surprise. Thus, our findings suggested that there existed a concurrent and dissociable account of conflict and surprise within effortful decision making.

Our findings demonstrate that neither conflict nor surprise may in themselves define what makes a decision difficult (i.e., more effortful). Further, they are neither able to independently account for the recruitment of cognitive control and the exertion of effort in decision making – indicating that cognitive control and effort are complex cognitive phenomena that are elicited by a range of factors. Both conflict and surprise need to be considered in parallel and with our findings we may now have dissociated neural indicators of each, pending replication.

Thus, we add to the neural evidence of concurrent contributing factors of cognitive control and effortful decision making. For example, Lin and colleagues (Lin et al., 2018) demonstrated both conflict and surprise signals in decision making. In their task, participants decided between an immediate reward and a delayed reward and task difficulty was manipulated as the subjective difference between the immediate and delayed reward magnitudes. They found that theta activity was highest when the subjective rewards of the options were similar (i.e., conflict was highest) and decreased with the degree to which the options diverged (i.e., decreasing conflict). They also presented participants with a small subset of surprising trials – i.e., no-brainer trials – which were designed so that the immediate reward was undoubtedly better than the delayed reward. They found heightened theta activity with no-brainer trials. As these trials were much easier than what the participants were used to, they concluded that theta was also indicative of surprise. As theta activity varied proportionally to conflict and was large to unexpected no-brainer trials, Lin and colleagues (Lin et al., 2018) findings demonstrate concurrent neural signals of conflict and surprise.

Our findings extend the work of Lin and colleagues by analyzing both conflict and surprise on a continuous scale – demonstrating patterns of both of these signals across task difficulty. With that said, one discrepancy between our findings and those of Lin and colleagues

(Lin et al., 2018) is that our frontal EEG activity exclusively reflected a surprise signal whereas Lin and colleagues' found that frontal theta reflected both conflict and surprise. These deviations, however, may simply be due to the different task demands between studies – Lin and colleagues (Lin et al., 2018) manipulated task difficulty as the relative subjective value of the immediate and delayed rewards and we manipulated task difficulty as the relative likelihood of each diagnosis.

In addition to our current findings and those of Lin and colleagues (Lin et al., 2018), Vassena and colleagues (Vassena et al., 2020) determined concurrent patterns of surprise and conflict while measuring brain activity via fMRI imaging. In their research, participants were to select one of two pairs of stimuli, each indicating a reward magnitude. Vassena and colleagues manipulated task difficulty, and thus conflict, as the discrepancy between the two option values wherein difficulty and conflict was highest when the magnitudes were similar. In addition, they signified that participants developed an expectation of presented rewards as the long-run averaged reward across trials and thus determined surprise to increase as values diverged from this expectation. They found that activity within the dorsal anterior cingulate cortex corresponded to surprise signals and activity within the ventral medial prefrontal cortex reflected control signals (similar to conflict in our findings). Curiously, despite finding the presence of both conflict-like and surprise signals, Vassena and colleagues (Vassena et al., 2020) concluded surprise as the sole contributor to the recruitment of cognitive control – a conclusion criticized by others who posited their findings to provide evidence for both conflict and surprise in effortful decision making (Shenhav et al., 2020).

Together, our findings with those of Lin (Lin et al., 2018), Vassena (Vassena et al., 2020), and their colleagues demonstrate dissociable and concurrent signals of conflict and surprise during effortful decision making. Although it is beyond the scope of the current manuscript, different computational models, such as the Expected Value of Control (EVC; (Grahek et al., 2020; Musslick et al., 2019; Shenhav et al., 2013)) and the Predicted Response Outcome (PRO; (Alexander & Brown, 2011; J. W. Brown & Alexander, 2017; Vassena et al., 2020)), may afford an explanation as to how our findings of dissociable conflict and surprise signals may affect decision making through the engagement of cognitive control and the employment of effort (Alexander & Brown, 2011; Shenhav et al., 2013). The EVC model posits that expected benefits and costs are utilized to determine whether to initiate cognitive control and the degree of effort to be enforced (Kool et al., 2017; Kool & Botvinick, 2018; Shenhav et al., 2013, 2016, 2020). Thus, control costs (e.g., to resolve conflict) as determined by this model explicitly function to influence action selection. Concurrently, the PRO model tracks expectations and determines alignment with the environment (Alexander & Brown, 2011; J. W. Brown & Alexander, 2017; Vassena, Deraeve, et al., 2017; Vassena et al., 2020). If discrepancies exist (a surprising event), control is initiated with the function of updating expectations to reduce future difficulty. Thus, conflict signals may recruit cognitive control to reactively address immediate action selection, while surprise signals may recruit cognitive control to proactively account for environmental demands via context-updating.

In conclusion, here we demonstrate that there may be dissociable and concurrent signals of conflict and surprise, indicating that neither alone may be able to explain the complexities of effortful decision making. Our findings indicate the complexity of recruiting factors and functions of cognitive control. Specifically, conflict may recruit cognitive control for action selection and surprise may recruit cognitive control to update expectations. We posit that future research may benefit by considering these signals as simultaneous contributors for the recruitment of cognitive control and the employment of effort in decision making as opposed to debating which of them should be considered as the lone contributor.

Chapter 4: Experiment 3

Stabilizing Expectations when Shifting from Analytical to Intuitive Reasoning:

The Role of Prediction Errors in Reasoning

Chad C. Williams*, Cameron D. Hassall, and Olave E. Krigolson

The Theoretical and Applied Neuroscience Laboratory, University of Victoria, Victoria, British Columbia, V8W 2Y2, Canada

Under Review in Cortex, 2022

Williams, C.C., Hassall, C.D., & Krigolson, O.E. (2022). Stabilizing Expectations when Shifting from Analytical to Intuitive Reasoning: The Role of Prediction Errors in Reasoning. Cortex. https://doi.org/10.31234/osf.io/gc6u9

Abstract

As humans, we rely on intuitive reasoning for most of our decisions. However, when there is a novel or atypical decision to be made, we must rely on a slower and more deliberative thought process - analytical reasoning. As we gain experience with these novel or atypical decisions, our reasoning shifts from analytical to intuitive which parallels a reduction in the need for cognitive control. Here, we sought to confirm this claim by employing electroencephalographic (EEG) measures of cognitive control as participants performed a simple perceptual decision-making task. Specifically, we had participants categorize "*blobs*" into families based on their visual attributes so we could examine how their reasoning changed with learning. In a key manipulation, halfway through the experiment we introduced novel blob families to categorize, thus temporarily increasing the need for analytical reasoning (i.e., cognitive control). Congruent with past research, we focused our EEG analyses on frontal theta activity as it is linked to cognitive control and analytical thinking. As hypothesized, we found a transition from analytical to intuitive decision-making systems with learning as indexed by a decrease in frontal theta power. Further, when the novel blobs were introduced at the midpoint of the experiment, we found that decisions about these stimuli recruited analytical reasoning as indicated by increased theta power in comparison to decisions about well-practiced stimuli. We propose our findings to reflect prediction errors to decision demands – a monitoring process that determines whether our expectations of demands are met. Shifting from analytical to intuitive reasoning thus may reflect the stabilization of our expectations of decision demands, which can be violated with unexpected demands when encountering novel stimuli.

Introduction

As humans we are inherently lazy – we avoid putting out effort whenever we can (Kool et al., 2010). As a result, we rely on gut-hunches and heuristics to guide us in everyday decisions and only put forth effort when something important comes along; we reluctantly take some extra time to ensure we are making the best possible decision (DeNeys & Pennycook, 2019; Evans & Stanovich, 2013; Kahneman, 2011). Dual-process decision-making theory describes this as the trade-off between intuitive and analytical reasoning (DeNeys & Pennycook, 2019; Evans & Stanovich, 2013; Kahneman, 2011). Within this framework, the majority of our decisions are governed by heuristics. For example, we instinctively slam on our brakes when a child chases a ball into the road. But sometimes we must think something through and spend effort in doing so. For instance, when we deliberate whether to move to a new city for a job. In general, intuitive reasoning is thought to be faster but less accurate than analytical reasoning – but this is an acceptable trade-off because most of our intuitive decisions work out (Evans & Stanovich, 2013; Norman et al., 2017). That is, intuitive reasoning is effective when everything is familiar and our responses are well-practiced, a situation that defines most of our daily decisions.

If intuition is so great, why do we ever need to be analytical? As noted above, intuitive reasoning fails when a situation is unfamiliar or uncertain and we do not have a rehearsed response. Another way of putting this is that intuitive reasoning fails us when we encounter something that is novel or atypical (Cavanagh & Frank, 2014; Williams, Ferguson, Hassall, Wright, et al., 2021). Let us say, on a walk to get coffee you are texting on your cellphone, and you take the wrong turn. You end up lost in an unknown neighbourhood. Because intuitions rely on our past knowledge(Croskerry, 2009, 2017; Evans & Stanovich, 2013; Lin et al., 2017), in this novel neighbourhood you are going to have to start thinking and looking for signs that will help you find your way out. This is, of course, effortful and analytical. However, just finding your way out does not mean that the next time you end up in this same neighbourhood you will be able to rely on intuitions because it is still an atypical place for you to be (Ericsson et al., 1993; Williams, Ferguson, Hassall, Wright, et al., 2021); considerable experience is needed for something to become intuitive (Ericsson et al., 1993).

Indeed, experience allows us to shift from making analytical to intuitive decisions (Ericsson et al., 1993). We have all heard that *practice makes perfect* and that is exactly what we are talking about here. Let us think about learning a new skill, such as driving. The first time you sat behind the wheel of a car, you likely had more to think about than your brain could handle. You were confronted with a load of gauges, the pedals, and of course the steering wheel. Moreover, you do not know what to pay attention to, so you try and pay attention to all of it. Every decision you make – for example pulling onto the street from your driveway – is a stressful process requiring cognitive effort. However, driving gets easier with practice and now as an expert driver you can hold a conversation with your passenger while listening to the radio and drinking a coffee all as you commute along a busy highway.

So, what dissociates intuitive and analytical thinking – or in other words, what changes as you practice making decisions? Emerging research suggests that analytical decision making requires cognitive control (Kool et al., 2017; Pennycook et al., 2015; Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Kappen, et al., 2019; Williams, VanOorschot, & Krigolson, 2021) – a mechanism that functions to coordinate cognitive processes. Think of cognitive control as the manager of a store where the manager supervises their employees and directs them to make sure the store is functioning as best as possible. This means that analytical decision making not only relies on a brain region to act as the manager, but also relies on the coordination of a brain network, or of multiple brain networks (Williams, Kappen, et al., 2019).

Neuroimaging research has described cognitive control to coordinate brain networks and recent electroencephalographic (EEG) research has explicitly linked neural signals of cognitive control to analytical thinking (Cavanagh, Figueroa, et al., 2012; Cavanagh & Frank, 2014; Cavanagh & Shackman, 2015; Eisma et al., 2021; Umemoto et al., 2019; Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Kappen, et al., 2019; Williams, VanOorschot, & Krigolson, 2021). According to Cavanagh and Frank, a neural brain oscillation, specifically frontal theta band activity (4 to 7 Hz), reflects cognitive control (Cavanagh & Frank, 2014; Cavanagh & Shackman, 2015; Eisma et al., 2021) and has been linked to control demands (Cavanagh & Shackman, 2015; Eisma et al., 2021). Work by our laboratory has elaborated on these findings by linking theta oscillations (i.e., cognitive control) to analytical thinking. For example, in 2019 (Williams, Kappen, et al., 2019) we showed increased frontal theta power corresponded to analytical rather than intuitive reasoning. In 2021 (Williams, Ferguson, Hassall, Wright, et al., 2021), we further demonstrated that, during reasoning, increased theta power was an indicator of surprise, which reflected a novel or atypical environment and thus also the need to recruit cognitive control to reason analytically. Further in 2021 (Williams, VanOorschot, & Krigolson, 2021), we linked theta power to a person's tendency to respond analytically when completing word problems. Altogether, frontal theta oscillations seem to be related to analytical thinking. However, our previous work contrasted the neural activity of people when placed in situations where they either needed to reason analytically or not. What is yet to be investigated is how cognitive control, as reflected by theta oscillations, changes with learning.

Here we here sought to investigate how neural signals of cognitive control change as someone shifts from reasoning analytically to reasoning intuitively. To examine this, we had participants classify complex shapes – which we call *blobs* – into different families. We investigated how cognitive control, as indexed by theta power, changed with experience. We hypothesized that theta power, and thus cognitive control, would decrease with learning, implying the transfer from analytical to intuitive reasoning. In a critical manipulation, halfway through the experiment we introduced two new families of blobs to determine whether these novel stimuli would elicit increased cognitive control relative to stimuli which the participants had already learned. Here, we hypothesized that the introduction of novel stimuli at the halfway point of the study would be associated with increased theta power when making decisions about these new stimuli relative to decisions about previously learned stimuli.

Materials and Methods

Participants

Fifteen undergraduate students ($\bar{X}_{age}$ = 19.60 years [SD: 18.02 years, 21.28 years], 14 female, 1 male) from Dalhousie University participated in this study for course credit via the online sign-up system. All participants had normal or corrected-to-normal vision and provided informed consent approved by the Research Ethics Board at Dalhousie University.

Apparatus and Procedure

Participants sat in a dark, sound attenuated room in front of a 19” LCD computer monitor and responded to the task using a Logitech game controller. The task (written in MATLAB version, 8.3, Mathworks, Natick, MA, using the Psychophysics Toolbox (Brainard, 1997)) consisted of participants classifying different families of polygons, or more casually *blobs*. Six families of blobs were constructed in the same way as Krigolson and colleagues (Krigolson et al., 2009) wherein six prototype blobs were created by dividing a circle into 20 vertices, randomly placing the vertices within 30-70% of the original circles radius and interconnecting them to form a closed polygon. Each family prototype was then modulated into one hundred exemplars by randomly adjusting the radius of each vertex by ±20%, thus creating six families of blobs.

On each trial, participants first viewed a fixation cross for 500-1000ms, followed by a blob presented in the center of the screen (see Figure 12). After a 1000-1500ms delay, a family label (families were named from A to F) was presented underneath the blob and the participants were to report whether the blob fit into the family. The label was a correct match 50% of the time. Following this, a second fixation cross appeared for 500-1000ms and correct or incorrect feedback was presented as a checkmark or an X, respectively, for 1000ms. There were two phases in this experiment in that the first half of the experiment presented four families of blobs for participants to learn and the second half of the experiment presented two ‘old’ families of blobs that were present in the first half of the experiment and two ‘new’ families of blobs that participants had not yet seen. Each half of the experiment consisted of three blocks of 100 trials and thus there were a total of 600 trials in this experiment. Our analyses will focus on early trials of each half of the experiment (when learning was highest) and categorize three conditions. The ‘Early’ condition reflects the initial four families of blobs in the first half of the experiment, the ‘Old’ condition reflects the two families of blobs within the second half of the experiment that were extensively learned in the first half, and the ‘New’ condition reflects the two families of blobs introduced in the second half of the experiment. Thus, we are able to contrast practice effects on decision making by comparing the Early and Old conditions and by comparing the New and Old conditions.

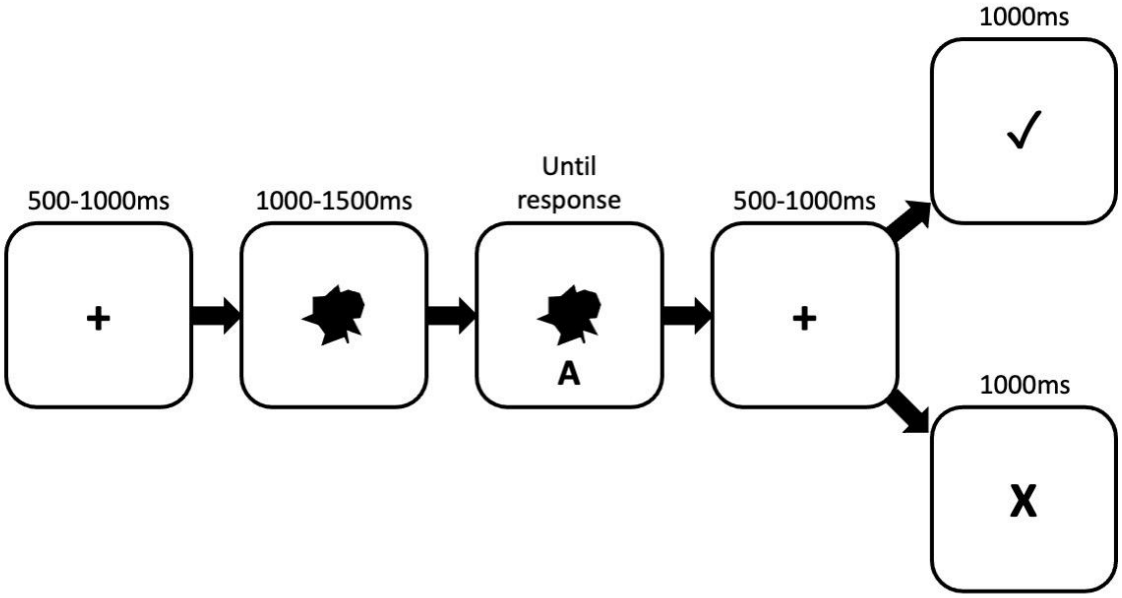

Figure 12: Experiment 3 (Figure 1)

Task presentations with timing. On each trial, participants see a fixation cross followed by a blob stimulus. After a delay, a blob family label is presented (here, the ‘A’ family) until the participant indicates whether the family and blob match. Upon their response, a fixation cross precedes feedback (checkmark = correct, X = incorrect).

Data Acquisition and Processing

Behavioural data were collected using a Logitech game controller and were recorded using MATLAB. EEG data were collected using a 64-electrode system with a 10-20 layout (ActiCAP, Brainproducts GmbH, Munich, Germany) and recorded using Brain Vision Recorder (Brainproducts GmbH, Munich, Germany). Data were recorded with a sampling rate of 500Hz and filtered with an antialiasing low-pass filter of 245Hz via an ActiCHamp amplifier (Brainproducts GmbH, Germany). Electrode impedances were, on average, below 10 kΩ.

Although the data were collected in 2014, they were analyzed in 2021 using MATLAB (version, 9.9, Mathworks, Natick, MA) and custom MATLAB scripts and toolboxes (github.com/neuro-tools: MATLAB-EEG-fileIO toolbox, MATLAB-EEG-preProcessing toolbox, and MATLAB-EEG-timeFrequencyAnalysis toolbox), which rely on EEGLAB (Delorme & Makeig, 2004). All data were referenced to averaged mastoids, except for two participants whose mastoid electrodes were faulty and as such data were referenced to an averaged electrode, filtered using a Butterworth passband filter (0.1 to 30Hz with a 4th order) and a notch filter (60Hz, 2nd order). Next, eye blinks were corrected using independent component analyses (ICA) wherein ICA blink components were visually identified by factor loadings and scalp topographies and EEG data were reconstructed after removing these components. Data were then segmented from -500 to 1500ms respective of events of interests. The events of interests were the onset of blob stimuli in the Early, Old, and New conditions. Data were then baseline corrected using 200ms prior to event onset and run through an artifact rejection algorithm with a 150μV max-min criteria. Any electrode that exceeded rejection rates of 20% were tagged as noisy or faulty for removal. Thus far, we have described the first pass of data processing, which we use to identify faulty electrodes. Our second pass of data processing was exactly as just described except with the removal of faulty electrodes at the beginning and the introduction of re-constructing faulty electrodes using topographic interpolation with spherical splines prior to segmentation.

Data were then transformed into time-frequency representations using a Gaussian-windowed complex sine wave with a normalized 6 cycle Morlet parameter to result in frequencies from 1 to 30Hz. As time-frequency transformations create edge artifacts we then resegmented data to be -200 to 1000ms respective of events of interest. Further, we only here present data from 1 to 20Hz and focus our analyses within theta frequencies (roughly 4 to 8 Hz).

Data Analysis

There were three conditions of interest considered in this manuscript. The first is named the ‘Early’ condition wherein we investigate the four blob families presented in the first half of the experiment. Two of these four blob families are re-investigated within the second half of the experiment, and this is the ‘Old’ condition. Finally, two new families of blobs are introduced within the second half of the experiment and investigated as the ‘New’ condition. Although the families share lots of features and as such are difficult to categorize, participants performed better than we expected from pilot research. The new blob families in the second half of the experiment were learned at a very quick pace, likely due to generalizability effects of what was learned in the first half of the experiment. As such, we had to focus our analyses to early trials of each experimental half and thus behavioural and neural data for each condition and participant reflected the average of the first five trials of the first/second half of the experiment.

Our paradigm has participants make difficult classifications of complex stimuli and thus would recruit effortful decision making systems of the brain to do so early in learning. As participants become well practiced with learning, they rely less heavily on effortful decision making systems in favour of recruiting effortless decision making systems. As such, we can investigate the effects of learning in two ways. First, we can determine how decision making for a family of blobs changes from the first half (Early condition) to the second half (Old condition) of the experiment, and second we can compare how decision making differs between learned blob families (Old condition) and unlearned blob families introduced in the second half of the experiment (New condition). In other words, we will use the Old condition as an indicator of learned and effortless decision making and compare this to the two unlearned and effortful decision making conditions, namely the Early and New conditions.

Behavioural measures of accuracy rates and reaction times were determined for each condition (see Figure 13A&B) and repeated-measures, two-tailed, t-tests (alpha = 0.05) were conducted to compare the Early and the New conditions to the Old condition. Additionally, we provide effect sizes in the form of 95% confidence intervals and Cohen's d measures.

Neural measures of decision making required the extraction of frontal theta components. First, we created grand averaged time-frequency representations for each condition (see Figure 14A-C). For each comparison (i.e., Early versus Old and New versus Old), we created difference time-frequency representations by subtracting the Old condition from each of the Early and New conditions (see Figure 14D&E). At electrode FCz (Williams, Ferguson, Hassall, Abimbola, et al., 2021; Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Kappen, et al., 2019), we then extracted a cluster of theta activity for each of the comparisons. For the Early-Old contrast, this resulted in a cluster that ranged from 4 to 8Hz and from ~0 to ~400ms post-blob stimulus onset (see Figure 14D). For the New-Old contrast, this ranged from 5 to 8Hz and from ~-50 to ~600ms post-blob stimulus onset (see Figure 14E). Note that both components peak at similar frequencies and times and that the latter extends prior to the onset of the blob stimulus (i.e.,Interest 0ms) due to temporal smearing, an artifact of time-frequency transforms (Cohen, 2014).

We then extracted theta activity for each participant and for the Early and New conditions by using their respective clusters. We also extracted theta activity for each participant for the Old condition by using the combination of both clusters to ensure that the Old condition was consistent across both contrasts (see Figure 13C). We conducted repeated-measures, two-tailed, t-tests (alpha = 0.05) for each contrast and also provide 95% confidence intervals and Cohen's d effect sizes.

Results

For accuracy rates, reaction times, and frontal theta activity, we contrasted the Early and New conditions with the Old condition, see Figure 13. Early accuracy rates (M = 48% [35%, 61%]) were statistically lower than Old accuracy rates (M = 72% [60%, 84%]), M_d = -24% [-43%, -5%], $t(14) = -2.67$, $p = .0183$, $d = -0.98$, see Figure 13A. In contrast, New accuracy rates (M = 65% [57%, 73%]) were not statistically different from Old accuracy rates, M_d = -7% [-22%, 8%], $t(14) = -0.96$, $p = .3535$, $d = -0.35$, see Figure 13A. Early reaction times (M = 909ms [745ms, 1,073ms]) were not different from Old reaction times (M = 903ms [781ms, 1,026ms]), M_d = 6ms [-143ms, 155ms], $t(14) = 0.08$, $p = .9361$, $d = 0.03$, see Figure 13B. Similarly, New reaction times (M = 954ms [800ms, 1,107ms]) were not statistically different from Old accuracy rates, M_d = 50ms [-80ms, 181ms], $t(14) = 0.83$, $p = .4215$, $d = 0.30$, see Figure 13B.

Neural analyses first required the identification of frontal theta clusters for each contrast, see Figure 14. Indeed, we found a theta cluster for each contrast with similar frequency, timing, and scalp topography, see Figure 14D-G. Generally, these clusters were from 4 to 8Hz, 0 to 500ms post-stimulus onset, and frontal-central. Each cluster was used to extract theta activity for each participant from the corresponding effortful conditions (Early, New). As the old condition serves as a non-effortful contrast for each of the effortful conditions, we extracted data using the combination of both clusters. Indeed, the Early theta activity ($M = 5.03\mu V^2$ [$3.62\ \mu V^2$, $6.45\ \mu V^2$]) was statistically larger than theta activity in the Old condition ($M = 4.08\mu V^2$ [$3.24\ \mu V^2$, $4.93\ \mu V^2$]), $M_d = 0.95\mu V^2$ [$0.04\mu V^2$, $1.86\mu V^2$], $t(14) = 2.24$, $p = .0419$, $d = 0.82$, see Figure 13C. The New theta activity ($M = 5.48\mu V^2$ [$3.99\ \mu V^2$, $6.97\ \mu V^2$]) was also statistically larger than the Old condition's theta activity, $M_d = 1.40\mu V^2$ [$0.26\mu V^2$, $2.54\mu V^2$], $t(14) = 2.63$, $p = .0198$, $d = 0.96$, see Figure 13C.

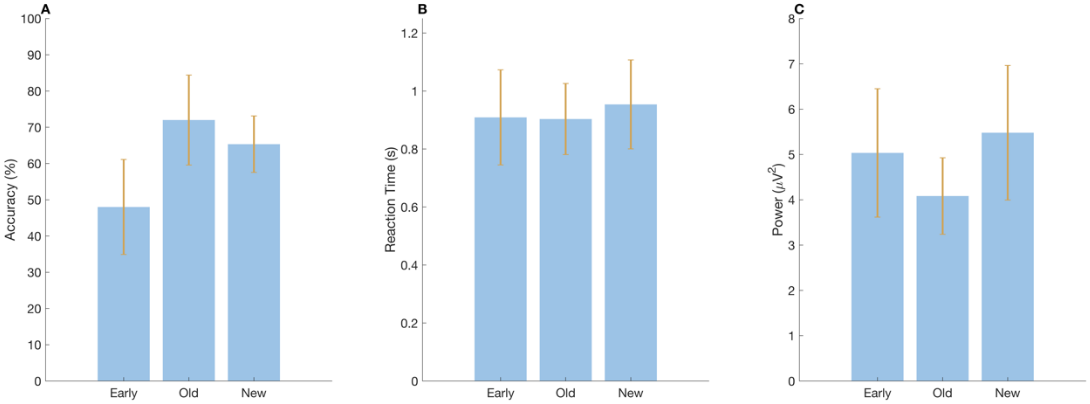

Figure 13: Experiment 3 (Figure 2)

Accuracy rates (A), reaction times (B), and frontal theta power (C) across conditions. Bars indicate means with 95% confidence intervals.

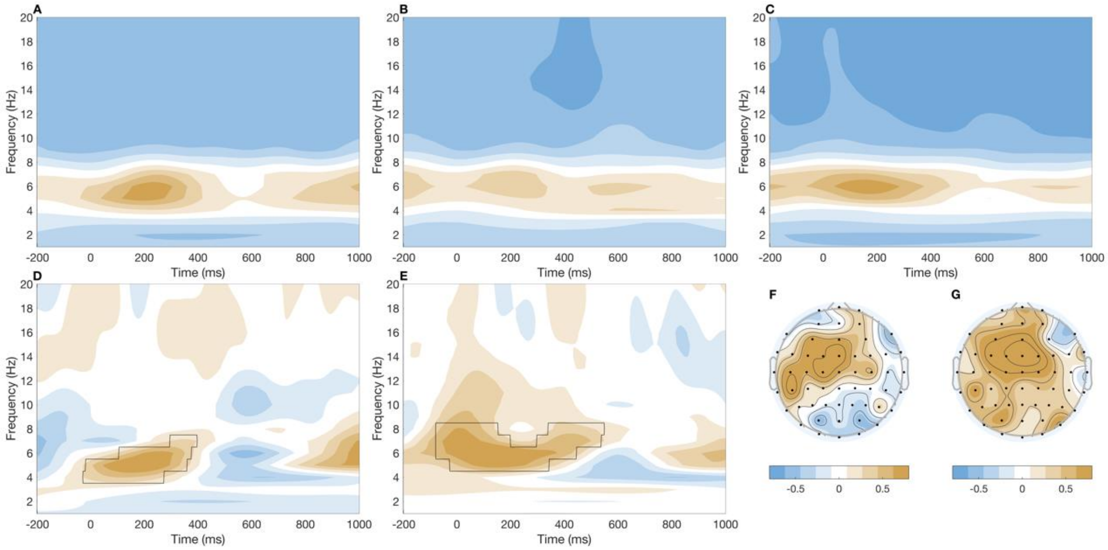

Figure 14: Experiment 3 (Figure 3)

Time-frequency plots of the Early (A), Old (B), and New (C) conditions, as well as of the difference between Early minus Old (D) and the difference between New minus Old (E) – contour lines demonstrate theta clusters of interest. To exemplify theta effects of interest, we applied an asymmetric Hamming window which peaked at 5 and 6Hz. In addition, topographic plots of the theta clusters of the difference between Early and Old (F) and the difference between New and Old (G) across electrodes.

Discussion

As expected, we found neural evidence of a transfer from analytical to intuitive reasoning systems over practice and further of a dissociation between novel and practiced decisions. As a reminder, novel and atypical decisions rely on the engagement of cognitive control (Evans & Stanovich, 2013; Pennycook et al., 2015; Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Kappen, et al., 2019; Williams, VanOorschot, & Krigolson, 2021) – a neural mechanism that facilitates decision-making and increases the coordination of different brain regions. With practice of the same decisions, we see a shift of effortful analytical reasoning to effortless intuitive reasoning – or in other words, a reduction in the need for cognitive control

(DeNeys & Pennycook, 2019; Pennycook, 2017; Pennycook et al., 2015). Here we demonstrated this shift by measuring frontal theta activity – an EEG signal thought to reflect the engagement of cognitive control (Cavanagh & Frank, 2014; Cavanagh & Shackman, 2015; Lin et al., 2020; Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Kappen, et al., 2019; Williams, VanOorschot, & Krigolson, 2021). Specifically, we found that frontal theta oscillations decreased with learning. We also introduced novel stimuli halfway through the experiment and found that making classification decisions about them required cognitive control (i.e., analytical reasoning) as evidenced by increased theta activity relative to the well-practiced stimuli, which relied on effortless intuitive thinking.

These findings are even more interesting when considering accuracy rates. As might be expected, we observed an increase in accuracy with learning, yet we did not see reduced accuracy for the novel stimuli introduced in the second half of the experiment. To explain this, we propose that participants were able to generalize their knowledge from the first half of the experiment to the decisions made with these novel stimuli. But what is interesting is that there may still be a need for cognitive control and analytical thinking, as evidenced by increased theta oscillations, for the novel stimuli. Together, this means that decisions about novel stimuli recruited cognitive control to ensure optimal functioning and accurate decisions – both core descriptors of analytical decision-making (Evans & Stanovich, 2013; Kahneman, 2011).

Earlier we discussed getting lost in an unfamiliar neighbourhood as an example of when you would need to employ analytical decision-making to find your way out. Now, let us say you have just moved and the neighbourhood you were previously lost in is your new neighbourhood. At first, you would need to employ analytical decision-making to navigate your new environment, but over time you would learn the streets and landmarks and eventually you could rely on intuitive heuristics. As we are only human, we would most likely only learn a single optimal route out of our neighbourhood and stick to it. But what if one day there is unexpected construction, and your normal route is blocked off? Well, you would again need to rely on your analytical reasoning system to find a novel path. If the construction lasted long enough, then the need for analytical reasoning would diminish and you would eventually know that novel path like the back of your hand, navigating it intuitively.

In the present study we have shown evidence that we can quantify this shift from analytical to intuitive reasoning, and subsequent re-engagement of analytical reasoning when it is needed. Here, these shifts were quantified as frontal theta oscillations and associated with a person's previous experience with an item or event. We propose that, during reasoning, theta oscillations may reflect a prediction error – a monitoring process wherein expectations are compared to actual events (Alexander & Brown, 2011; Krigolson, 2018; Proudfit, 2015; Williams, Ferguson, Hassall, Abimbola, et al., 2021). Specifically, prediction errors emerge in response to unexpected decision demands – i.e., if a decision is more or less demanding than expected. For example, walking along your normal neighbourhood path has become your expectation, yet the disruption of this path by construction elicits a prediction error, as your path is now going to be more demanding than usual.

Here, when we introduced novel families of blobs, demands of the task became more difficult than expected and prediction errors emerged. As such, the increased frontal theta activity to novel stimuli during reasoning reflected prediction errors concerning the demands of a decision. Indeed, our findings are in line with research describing frontal theta activity to be a signal of surprise in decision making (Lin et al., 2018; Williams, Ferguson, Hassall, Wright, et al., 2021). Here, any unexpected event elicits a prediction error and thus is surprising. Further, the Predicted Response Outcome (PRO) model posits that these surprising events are the consequence of a monitoring process that tracks the likelihood of events, such as stimuli, responses, and their outcomes (Alexander & Brown, 2011; J. W. Brown & Alexander, 2017; Vassena et al., 2020; Williams, Ferguson, Hassall, Wright, et al., 2021). In line with the PRO model, we provide evidence that these prediction errors may monitor decision making demands. Here, prediction errors emerge when decision demands are not as expected – e.g., when participants encountered unexpected novel blob families to learn. As such, increased theta activity may be an indication of a prediction error that recruits cognitive control to update one's expectations of decision demands for future reasoning (Alexander & Brown, 2011; J. W. Brown & Alexander, 2017). Consequently, the shift from analytical to intuitive reasoning then involves this monitoring process that stabilizes our expectations of decision demands. Further, if a well-practiced decision presents unexpected characteristics, then a prediction error emerges, and cognitive control is employed to update our expectations.

To summarize, in the present study we found that frontal theta activity was a predictor of cognitive control and analytical reasoning, in congruence with past research (Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Kappen, et al., 2019; Williams, VanOorschot, & Krigolson, 2021). What is novel here is that we demonstrated that changes in theta oscillations reflected a transition from analytical to intuitive reasoning with practice. Additionally, we saw that introducing novel decisions in a well-practiced environment also recruited cognitive control (i.e., increased theta oscillations) to reason analytically and achieve accuracy. We propose that the transition from analytical to intuitive reasoning may involve a process wherein our expectations of decision demands are stabilized. Further, if an intuitive decision presents atypical characteristics, we recognize this in the form of a prediction error to recruit cognitive control and update our expectations.

Chapter 5: Experiment 4

Expectations of Decision Demands Influence Cost-Benefit Computations During Reasoning

Chad C. Williams[1]*, Mathew R. Hammerstrom[1], Michael Inzlicht[2], and Olave E. Krigolson[1]

[1]The Theoretical and Applied Neuroscience Laboratory, University of Victoria, Victoria, Canada
[2]Department of Psychology, University of Toronto, Toronto, Canada

*Corresponding author: Chad C. Williams, The Theoretical and Applied Neuroscience Laboratory, University of Victoria, P.O. Box 17000 STN CSC, Victoria, British Columbia, Canada, V8W 2Y2

Tel: 1-250-858-4682; Fax: 1-250-721-6601; E-mail address: ccwillia@uvic.ca

Abstract

The decision to exert effort in a judgment relies on a comparison of the expected costs and benefits. Recently, it has been proposed that this process involves the building of cost-benefit expectations that allow us to proactively determine the degree of effort to exert in upcoming judgments. In this framework, any deviation from our expectations elicits a prediction error and a subsequent reassessment of the effort being exerted. Here, we leveraged electroencephalography to investigate whether prediction errors play a role in reasoning alongside cost-benefit computations by having participants make a series of decisions to exert or withhold effort for different monetary rewards. Participants encountered judgments with equal costs and benefits and judgments that were overvalued and undervalued – these latter two eliciting prediction errors. Indeed, neural evidence expressed as frontal theta activity corroborated the influence of prediction errors on determining the degree of effort to exert in a decision.

Introduction

"Thirty white horses on a red hill,
First they champ,
Then they stamp,
Then they stand still"(Tolkien, 1937)
What are they?

Well, did you solve the riddle? Better yet – did you *try* to solve the riddle, or did you rather decide to move on? Either choice is fine – sure, it might be fun, and you might feel accomplished by solving the riddle, but it might also just be better for you to save your time and energy and get to the paper that you are here to read. The choice to exert effort or not is surprisingly complicated and that is the focus of the work presented here. Oh, and by the way – the answer is teeth.

When presented with a decision, such as a riddle, we first determine how we will approach the problem, and then we engage with it to reach a solution (Alexander & Brown, 2011; J. W. Brown & Alexander, 2017; Collins & Shenhav, 2022; Frömer & Shenhav, 2021; Kool et al., 2017; Kool & Botvinick, 2018; Shenhav et al., 2013, 2017, 2021; Vassena, Holroyd, et al., 2017; Vassena et al., 2020; Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Hassall, et al., 2022). In terms of the riddle, you would have first determined whether you would bother trying to solve it at all and then either attempt to solve it or skip ahead. Here we will discuss these two concepts – the strategy selection phase and the engagement phase of a decision – in reverse order. With regard to the engagement phase, we will discuss effort and its limitations; in terms of the strategy selection phase, we will focus on hypotheses that ground strategy selection as a cost-benefit computation.

The Engagement Phase: Exerting and Withholding Effort

Exerting or withholding effort in a decision is a process that we are all very familiar with, and whether you spent the effort to solve the riddle or not has very different consequences. Exerting effort requires more time and mental resources, but with the advantage of being more accurate (DeNeys, 2017; DeNeys & Pennycook, 2019; Evans & Stanovich, 2013; Kahneman, 2011; Pennycook, 2017). In contrast, withholding effort is quick and requires little-to-no mental resources, yet is prone to error (DeNeys, 2017; DeNeys & Pennycook, 2019; Evans & Stanovich, 2013; Kahneman, 2011; Pennycook, 2017). Despite its propensity for mistakes, the majority of our daily decisions are done so effortlessly (Kahneman, 2011). We can often rely on effortless reasoning because it only fails when encountering atypical or novel rather than highly familiarized decisions (Cavanagh & Frank, 2014; Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Hassall, et al., 2022). Just as there are times where effortless reasoning is optimal, there are times where exerting effort is necessary (Williams, Hassall, et al., 2022) – for example, when deciding whether to move to a new city for a job. Indeed, it is unlikely that we are going to make such a life-changing decision on a whim without much thought. Instead, we will contrast the pros and the cons of this move to eventually reach a solution.

These examples highlight the two extremes of the effort spectrum: cases where we would not fathom spending time and effort to reach a solution, and other cases where we would be foolish to decide without a thought (Alexander & Brown, 2011; Botvinick & Cohen, 2014; Egner, 2017; Nigg, 2017; Shenhav et al., 2013). But there are a lot of decisions in-between these two extremes that could go either way depending on who is making the decision and how they were feeling in the moment of making that decision (Alexander & Brown, 2011; J. W. Brown & Alexander, 2017; Collins & Shenhav, 2022; Frömer & Shenhav, 2021; Kool et al., 2017; Kool & Botvinick, 2018; Shenhav et al., 2013, 2017, 2021; Vassena, Holroyd, et al., 2017; Vassena et al., 2020; Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Hassall, et al., 2022). For example, what's for dinner? Answering with "anything is fine" versus "what are our options?" have very different demands. Whereas the former relinquishes the need to effortfully decide what to eat, it also comes with the risk of getting something that you do not want. In contrast, the latter requires the effortful consideration of a lot of factors to reach a solution yet ensures that you get the meal you are craving. Why, then, don't we always just put forth effort to make these types of decisions? Afterall, wouldn't we always want to satisfy our cravings? It turns out that we are limited in how many effortful decisions we can make in a given amount of time and so we must choose carefully what decisions we will effortfully engage with (Inzlicht & Friese, 2019; Inzlicht & Schmeichel, 2012; Shenhav et al., 2017).

The depletion of mental energy during effortful reasoning is similar to the depletion of physical energy during a workout (Baumeister et al., 1998; Baumeister & Heatherton, 1996; Inzlicht & Friese, 2019; Inzlicht & Schmeichel, 2012; Muraven et al., 1998; Shenhav et al., 2017). In both cases, the exertion of energy results in a lessened likelihood of again exerting energy. In the physical domain, we know that one major contributor to physical depletion is glucose and that we can quickly replenish our physical capacity by consuming sugar (Baumeister et al., 1998; Baumeister & Heatherton, 1996; Muraven et al., 1998). In the mental domain, it is not this simple. There is a rich history of research that considered this same substance, blood glucose, to be proportional to our mental capacity (Gailliot et al., 2007), yet this claim has since been criticized and is no longer valid (Dang, 2016; Finley et al., 2019; Kurzban, 2010; Kurzban et al., 2013; Vadillo et al., 2016). That is not to say that the concept of a metabolic constraint in itself has been refuted, but simply that glucose is not the substance underlying our mental capacity. In fact, others have proposed astrocytic glycogen to be proportional (Christie & Schrater, 2015), and amyloid-Beta neurotoxin build up to be inversely proportional (Holroyd, 2015), to mental capacity. None-the-less, these latter substances are still under early investigation and so we do not yet know the extent to which they influence our capacity to engage in effortful reasoning (Shenhav et al., 2017).

Some have rather proposed our mental capacity to be limited by our motivation (Inzlicht et al., 2014; Inzlicht & Berkman, 2015; Inzlicht & Friese, 2019; Inzlicht & Schmeichel, 2012; Kool & Botvinick, 2014). In other words, it is not an inability to continuously engage in effortful reasoning, but rather a lack of motivation to do so. Even another account has suggested that there rather exists a structural capacity – that is, effort is guided by a single mechanism that has computational limitations (Cowan, 2012; Luck & Vogel, 1997; Ma et al., 2014; Oberauer et al., 2016). Although this hypothesis draws on the well-studied limitations of working memory, it proposes a single mechanism as the controller for all of human effort and has been criticized to be biologically implausible (Cowan, 2012; Luck & Vogel, 1997; Ma et al., 2014; Oberauer et al., 2016). Indeed, others have reframed this concept to be system-wide in that exerting effort places strain on information-processing pathways that persists and interferes with our ability to re-engage in another effortful decision (Feng et al., 2014; Musslick, Jang, et al., 2018; Musslick & Cohen, 2019, 2021). Together, it seems that engaging in an effortful decision results in a reduced chance of engaging in another effortful decision but it is still unclear why this happens biologically.

The Strategy Selection Phase: The Expected Costs and Benefits of a Decision

Regardless of the biological basis, we know that we must be conservative when deciding to engage effortfully with a decision. This makes the strategy selection phase of decision-making very important because if we robustly engage effortfully with our decisions, we will run out of steam by lunch and struggle for the rest of the day. Balancing our expectations of the costs and benefits of a decision is how the Expected Value of Control (EVC) model (Collins & Shenhav, 2022; Frömer & Shenhav, 2021; Kool et al., 2017; Shenhav et al., 2013, 2017, 2021) frames strategy selection in reasoning. It poses that we compare our *expected* benefits and our *expected* costs of engaging effortfully and if the benefits outweigh the costs, then we exert effort, but if the costs outweigh the benefits, then we withhold effort. This claim turns out to be complicated because benefits and costs are not objective truths of the world, but subjective estimations that we compute (Collins & Shenhav, 2022; Frömer & Shenhav, 2021; Kool et al., 2017; Shenhav et al., 2013, 2017, 2021).

Sometimes benefits can seem objective and easy to predict – for example, your paycheque. There is often no ambiguity in what your paycheque will be, but that is not the only benefit for doing a good job is it? Alongside a paycheque, we could get praise and respect, but also a promotion. Similarly, doing a good job avoids getting fired – another benefit. But again, all of these benefits are subjective and vary across people and across time (Frömer & Shenhav, 2021; Kool & Botvinick, 2018; Shenhav et al., 2013, 2021; Westbrook et al., 2013; Yee et al., 2021). Money is more important for some than others and at some times than other times, and so the expected benefit of that paycheque varies accordingly. Likewise, some may be aiming for that promotion, while others could not care less if they get fired. And those are only a few examples as to how our expectations of benefits can be pulled this way and that.

It should be no surprise that our expectations of costs work in the same way. We have so far considered costs in terms of the mental resources drained when making an effortful decision, and although a major factor, this is far from the whole picture (Frömer & Shenhav, 2021; Kool & Botvinick, 2018; Shenhav et al., 2013, 2017, 2021; Westbrook et al., 2013; Yee et al., 2021). For instance, the time it takes to make an effortful decision is also a cost. Spending extra time to effortfully consider a problem may lead to the lost opportunity of engaging with another decision that would require effort. If emergency doctors, for example, effortfully engage with their patients that portray typical symptoms of a flu, they would be taking time away from seeing other patients that may have life-threatening symptoms. These opportunity costs become an important factor in deciding whether we should exert effort (Frömer & Shenhav, 2021; Kool & Botvinick, 2018; Kurzban et al., 2013; Lin et al., 2022; Shenhav et al., 2013, 2017, 2021; Westbrook et al., 2013; Yee et al., 2021).

Altogether, in the strategy selection phase of reasoning, we weigh the expected costs and benefits of exerting effort, both of which are shrouded in complexity. Yet, the story still does not end there – costs and benefits also interact with each other. When deciding whether to exert effort in a decision, we must also consider the efficacy of exerting that effort. First, performance efficacy refers to whether exerting effort will have any influence on attaining more benefits (Frömer et al., 2021; Frömer & Shenhav, 2021; Schevernels et al., 2014; Shenhav et al., 2013, 2021). When performance efficacy exists, the more effort exerted, the more benefits we may expect. If there is no performance efficacy, then the likelihood of achieving the benefits stays the same, no matter how much effort we exert. In these cases, we are not motivated to exert effort. Second, control efficacy regards whether exerting effort will have any influence on our performance (Frömer et al., 2021; Frömer & Shenhav, 2021; Schevernels et al., 2014; Shenhav et al., 2013, 2021). When control efficacy exists, exerting more effort leads to increases in performance, and thus higher likelihood of attaining the benefits. In other words, our beliefs about whether we are even able to solve the decision at all matters. If we have no control efficacy, there is no benefit of exerting effort as there is no possibility of actually attaining it.

We do have one final consideration – our proactive expectations of decision demands. Here, decision demands correspond to a decision's expected costs and benefits, and we build our expectations to facilitate future decision making (Alexander & Brown, 2011; J. W. Brown & Alexander, 2017; Vassena et al., 2020; Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Hassall, et al., 2022). These expectations are the *long-run average* of decision demands and are proactive, meaning that they are determined before encountering a decision. When presented with a decision, if the demands (costs and benefits) match our expectations, we continue with our predisposed reasoning strategy (Frömer & Shenhav, 2021; Shenhav et al., 2021; Williams, Hassall, et al., 2022). However, if the decision demands deviate from our expectations then a prediction error occurs. Prediction errors occur when our expectations are violated and signify that we must re-assess our predisposed strategy (Alexander & Brown, 2011; J. W. Brown & Alexander, 2017; Holroyd & Krigolson, 2007; Krigolson, 2018; Proudfit, 2015; Vassena et al., 2020; Williams et al., 2017; Williams, Ferguson, Hassall, Abimbola, et al., 2021; Williams, Hassall, et al., 2022). In reasoning, prediction errors may signify a violation of our expectations of decision demands and may occur both when a decision is more or less demanding than expected (Williams, Hassall, et al., 2022). In other words, in consistent environments, we settle into a strategy, whether that is to employ effort or to withhold effort. We continue to use this strategy until demands do not meet our expectations, and then we re-assess to determine whether we should change our strategy.

The proactive expectation of decision demands is in line with the Predicted Response Outcome (PRO) model (Alexander & Brown, 2011, 2014; Vassena, Deraeve, et al., 2017; Vassena et al., 2020), which poses a monitoring process that tracks the likelihood of events. We have recently proposed that both the PRO and EVC hypotheses may work in tandem during reasoning (Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Hassall, et al., 2022). Specifically, the PRO model acts to proactively account for decision demands while the EVC model acts to reactively determine whether to exert effort in a decision. Although we now know these two mechanisms are involved in reasoning, it is still unclear how they interact.

The Current Study

Here, we sought to determine neural signals that reflect prediction errors of decision demands and/or cost-benefit computations. We hypothesize that we develop expectations of decision demands and that this leads to the adoption of a predisposed amount of effort to exert for upcoming decisions. As long as the decision demands remain constant for these new decisions, we continue to exert the same level of effort. Yet, if we detect a change in decision demands when encountering a decision, we tune our cost-benefit computation to determine the level of effort to exert. Deviations from our expectations of decision demands are identified as prediction errors, which could then lead to a change in the degree of effort being exerted as determined by a cost-benefit computation.

Here, we had participants decide between two upcoming judgments that would require them to exert or withhold effort, respectively. Each judgment was tied to different monetary rewards – specifically, the effortful decision would always provide more monetary incentive than the effortless decision, yet the discrepancy between the two varied as a function of participant's past choices. This manipulation resulted in decisions with equal expected costs and benefits. We also presented participants with scenarios where decision demands changed in that effortless decisions were either overvalued or undervalued. If they were overvalued, there would be more incentive to withhold effort and if they were undervalued, there would be more incentive to exert effort. In other words, the benefits of the overvalued judgments were better than expected, but the benefits of the undervalued judgments were worse than expected – both of which should elicit a prediction error.

We measured neural signals of reasoning by leveraging electroencephalographic (EEG) imaging techniques. Specifically, we investigated whether a neural oscillatory signal of the brain, frontal theta activity, reflected an expected value of control signal and/or a prediction error of decision demands signal (Cavanagh & Frank, 2014; Lin et al., 2018, 2022; Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Hassall, et al., 2022; Williams, Kappen, et al., 2019; Williams, VanOorschot, & Krigolson, 2021). If the former is true we predicted we would see a linear decrease in frontal theta activity from the overvalued to the undervalued judgments, whereas if the latter is true we predicted that overvalued and undervalued judgments would demonstrate equal frontal theta activity to each other but different from the equally valued judgments.

Methods

Participants

Sixty-eight students from the University of Victoria participated in the study; however, current analyses focused on 57 of these. Eight participants were removed from analyses due to *a priori* exclusion criteria based upon behavioural responses. Specifically, our analyses focused on a decision-making process wherein participants selected between an easy and hard judgement for varying degrees of monetary reward. The design of the paradigm tunes monetary reward for each participant so that participants select each judgment type equally across the experiment. Seven participants rather responded towards one of the two judgments 85% of the time or more, signifying that they were not following instructions to base their decisions on the monetary values of each judgment. As such, we removed them from analyses. Three other participants were removed from analyses due to technical issues – specifically, the EEG recording computer shut down during testing and participants did not encounter one of the conditions. As such, our sample size of analysis consisted of 57 participants ($\bar{x}_{age}$ = 21.11 years old [SD: 15.11 years old, 27.10 years old]; 40 female, 15 male, 1 non-binary, 1 other).

Our participants were recruited using an online sign-up system, where they earned credit for psychology courses. Additionally, participants earned an average monetary reward of $7.60 CAD [SD: $6.92 CAD, $8.27 CAD] dependent on their selections and performance on the task (see section 2.2. Paradigm and Procedures for details). All participants had normal or corrected-to-normal vision and provided informed consent that was approved by the Human Research Ethics Board at the University of Victoria (Ethics Protocol Number: 16-428).

Paradigm and Procedures

The experiment took place in a dark, sound-dampened room. Participants viewed stimuli on a 19” LCD computer monitor and responded using a ResponsePixx Controller (VPixx, Vision Science Solutions, Quebec, Canada). The task was written in MATLAB 2021a (Mathworks, Natick, MA, using the Psychophysics Toolbox Extension version 3.0.17 (Brainard, 1997).

The Reasoning Strategy Selection Paradigm

Here, we developed a novel paradigm called the *Reasoning Strategy Selection Paradigm* (RSSP; see Figure 15). This paradigm was designed to test assumptions of the expected value of control (EVC) theory (Shenhav et al., 2013). As described in the introduction, this theory posits that the decision to put forth effort in a judgment hinges on a cost-benefit computation – if the benefits outweigh the costs then effort will be exerted, but if the costs outweigh the benefits then effort will be withheld. If this theory is correct, then we would be able to find neural activity in the brain that 1) tracks the benefits of judgments, 2) tracks the costs of judgments, and 3) tracks the expected value of control – i.e., the difference between costs and benefits. It also affords us to track prediction errors when decision demands do not match our expectations. The RSSP was developed to manipulate benefits while holding costs constant and thus can investigate a range of neural signatures, such as the expected value of control as proportional to the difference between the costs and benefits of a decision. Furthermore, it can be used to investigate prediction errors of reasoning – when decision demands do not match our long-run expectations.

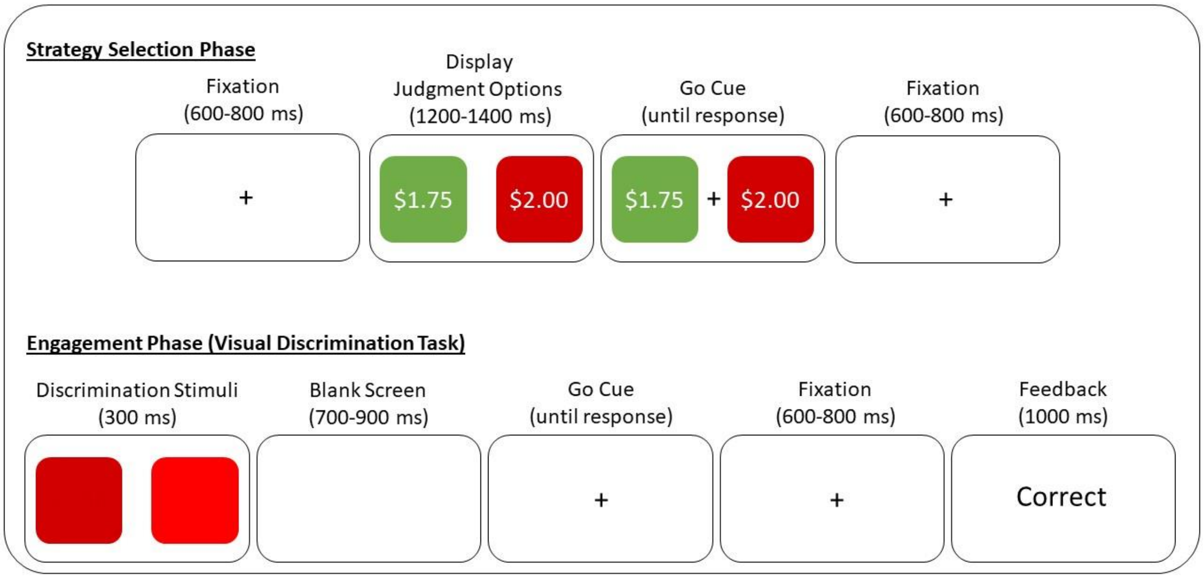

Figure 15: Experiment 4 (Figure 1)

The Reasoning Strategy Selection Paradigm (RSSP). This paradigm has two phases – the strategy selection phase (top) and the engagement phase (bottom). The former has participants decide whether to exert effort in a judgment for varying monetary gains and the latter has participants complete those judgments (which here were visual discriminations).

Each trial of the RSSP had two phases: *the strategy selection phase* and *the engagement phase*. The strategy selection phase corresponded to determining which strategies a person will use to complete a judgment. The EVC model theorizes that this involves a computation of costs and benefits to determine whether effort will be exerted in a judgment (Shenhav et al., 2013). As such, this phase had participants deciding whether to withhold effort in an easy judgment or exert effort in a hard judgment for varying monetary rewards. To tie these decisions with real-world outcomes, they then performed the selected judgment in the engagement phase. The judgment task may be any standard psychological task that varies effort and here we opted to use a visual discrimination task wherein participants identified which of two squares were brighter. Altogether, participants selected to complete either an easy or hard judgment for varying monetary rewards and then completed the selected judgment.

The Strategy Selection Phase

The strategy selection phase is novel and the foundation of the RSSP. Again, this phase had participants select their decision making strategy (e.g., whether to exert effort) through a cost-benefit computation. After a white fixation was displayed on a dark grey background for 600-800ms, two squares (100px in height and width) appeared each with a dollar value within it. The squares were green and red corresponding to the easy and hard judgments, respectively. The side to which each condition was presented was randomly determined on each trial. These squares with attached dollar values were presented for 1200-1400ms and then a white fixation cross appeared between them indicating to participants that they may now select whether they would like to engage in an easy or hard judgment for the corresponding reward. Once a response was made, a white fixation appeared for 600-800ms and then the trial moved onto the engagement phase as described in the next section.

The critical manipulation of this phase involved the dollar values presented to participants. The goal of the RSSP was to control how often participants encountered a decision where both the easy and hard judgments were equally subjectively valued. Again, this phase attached dollar values to both the easy and hard judgment options. In line with similar research (Lin et al., 2018), the value of the hard judgment stayed the same at \$2.00[1]. The value of the easy

---

[1] Note that the dollar values described here are, in fact, the values that participants saw; however, participants knew that they were only gaining a proportion of these values on each trial (although they did not know the proportion itself, they knew they could make up to \$10 CAD). The total money earned was dependent both on the dollar values selected by participants and their performance. The money earned was as follows: $\frac{rewards\ chosen}{total\ possible\ rewards} * Accuracy * \$10$, where rewards chosen was the sum of values that the participant selected, total possible rewards corresponds to always selecting the hard condition for \$2.00 across all six blocks with 60 trials per block (\$720),

judgment varied dependent on participant responses (see Figure 16). For each block, the reward for the easy judgment began at $1.75 – thus, the difference between the easy ($1.75) and hard ($2.00) judgment was ¢25. In this scenario, the participant must decide whether ¢25 is enough gain to make exerting effort in the hard judgment worth it. If they believed the extra effort to be worth the ¢25 then they would select the hard judgment, but if they did not believe it was worth it then they would select the easy judgment.

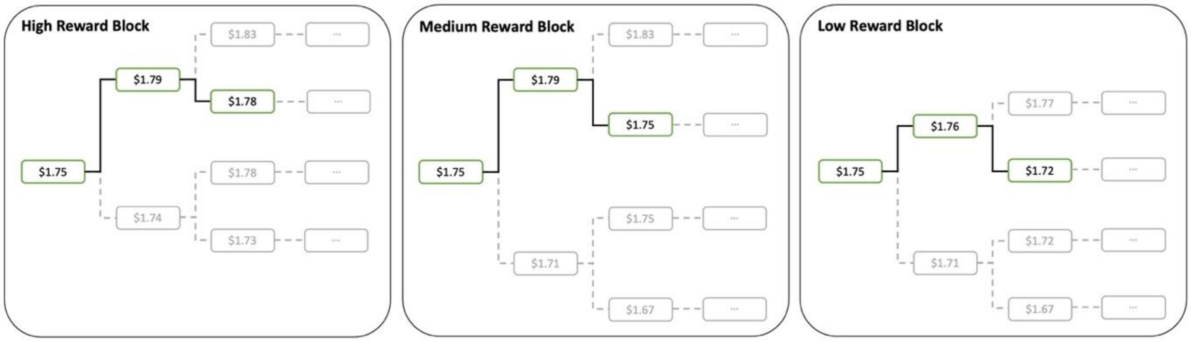

Figure 16: Experiment 4 (Figure 2)

Illustration of the reward manipulations across the three blocks. The value of the easy judgment (i.e., the dollar values here presented) changes based upon the previously selected judgment. If a hard judgment was selected, then the easy judgment's reward increased in value as depicted in the first choice of each figure panel. If an easy judgment is rather selected, the easy judgment's reward decreases in value, as depicted in the second choice of each figure panel. The degree to which the easy judgment changes is dependent on the block-wise manipulation.

After each decision, the value of the easy judgment either increased or decreased by ¢4 depending on their selection. If the participant chose the hard judgment, then the easy judgment

---

accuracy was the proportion of discriminations made correctly, and $10 indicated the total possible earnings for participating in the study. Participants received these dollar values rounded to the nearest whole dollar.

increased by ¢4 and in the next trial would be $1.79. This then reduced the gain for selecting the hard judgment (now, ¢21) and thus made it less likely to be selected. Conversely, if the participant rather selected the easy judgment, then its value decreased by ¢4 cents on the next trial and would be $1.71. This increased the gain (now, ¢29) for, and thus the likelihood of, selecting the hard judgment on future trials. The participant would eventually reach an equilibrium wherein the value of the easy condition made it equally likely to be selected as the hard condition – and this is what we found (see section 3. Results)[2].

This was the middle of three conditions: the high reward condition, the medium reward condition, and the low reward condition (see Figure 16). The difference in the high and low reward conditions was how much the values changed in light of participant responses. The high reward condition increased the value of the easy judgment after selecting a hard judgment by ¢4 (e.g., $1.75 changed to $1.79, same as the medium reward condition) but only decreased the value of the easy judgment when it was selected by ¢1 (e.g., $1.75 changed to $1.74). As such, after selecting a hard judgment, it took four easy judgments to return to the equilibrium value. This, on average, overvalued the easy judgments and thus would mathematically result in participants selecting the easy judgments 75% of the time. In contrast, the low reward condition flipped these changes in that selecting a hard judgment only increased the easy judgment’s value by ¢1 (e.g., $1.75 changed to $1.76) and selecting an easy judgment reduced its value by ¢4 (e.g., $1.75 changed to $1.71, same as the medium reward condition). As such, participants would mathematically select easy judgments 25% of the time as it was, on average, undervalued. The paradigm presented the high reward, medium reward, and low reward conditions in a

---

[2] One deviation from this pattern is that the value of the easy judgment was asymptotic when approaching either $1.00 or $2.00 to ensure that it did not surpass or equal these values (i.e., did not go lower than $1.00 or higher than $2.00). The asymptote function only affected values in the ranges of $1.00 to $1.08 and of $1.92 to $2.00.

blocked design where each condition was presented twice, all in a random order. Thus, this paradigm had 6 blocks of 60 trials, resulting in a total of 360 trials, or 120 trials per condition.

The Engagement Phase

After participants determined whether to exert effort in the strategy selection phase, they needed to make that judgment. The engagement phase can be any psychological task that varies effort and here we used a visual discrimination task where participants were to determine which of two red squares were brighter. Specifically, on a dark grey background, two red squares (100px height and width) were presented next to each other for 300ms (the side of the brighter square was randomly determined each trial), which was followed by a blank screen for 700-900ms. A white fixation cross then appeared in the center of the screen to indicate that participants were to judge whether the left or right square was brighter (with no time limit). After a response was selected, the fixation remained for 600-800ms and then deterministic feedback of performance was presented as white text (‘Correct’, ‘Incorrect’) for 1000ms. This was then the end of the trial.

The visual discrimination task varied effort requirements depending on how closely the two squares aligned. Specifically, the redness of the squares was defined by RGB values wherein easy judgments differed by 41 to 50 degrees (randomly determined on each trial – e.g., red values 150 versus 200 indicated a difference of 50), and hard judgments differed by 1 to 10 degrees (randomly determined on each trial – e.g., red values 150 and 160 indicated a difference of 10). Red RGB values were constrained so that they never fell below 150 and so they never surpassed the upper limit of 255, while green and blue RGB values were always set to 1.

Data Acquisition and Processing

We recorded behavioural data in MATLAB with a ResponsePixx Controller (VPixx, Vision Science Solutions, Quebec, Canada). Data from a standard 10-20 layout, 64-electrode ActiCHamp EEG system (Bainproducts GmbH, Munich, Germany) were recorded using Brain Vision Recorder (Version 2.1., Bainproducts GmbH, Munich, Germany). Data were sampled at 500 Hz with an antialiasing low-pass filter of 245 Hz. For optimal temporal precision, marker and stimuli synchronization was managed using a DATAPixx stimulus timing unit (VPixx, Vision Science Solutions, Quebec, Canada).

EEG data was processed using custom MATLAB scripts (github.com/neuro-tools: MATLAB-EEG-fileIO toolbox, MATLAB-EEG-preProcessing toolbox, and MATLAB-EEG-timeFrequencyAnalysis toolbox), which depended on EEGLab (Delorme & Makeig, 2004). We employed the following processing steps:

1. Removed excessively noisy and faulty electrodes
2. Re-referenced data to an averaged mastoid
3. Applied a Butterworth filter (passband: 0.1 to 30 Hz, 4th order) and a notch filter (60 Hz, 2nd order)
4. Corrected eye blinks using an infomax independent component analysis, determining components to be removed through visual inspection and factor loadings
5. Interpolated removed electrodes using spherical splines
6. Segmented data from -500ms to 1500ms relative to markers of interest (specifically, the onset of judgment options with attached reward dollar values)
7. Baseline corrected data from -200ms to 0ms

8. Removed trials using artifact rejection with a 150µV max-min criterion
9. Conducted time-frequency transforms using a Guassian-windowed complex sine wave with a Morlet parameter of 6 cycles (Cohen, 2014)
10. Averaged across trials for each condition and participant

Statistical Analysis

Behavioural Manipulation Checks

We begin our analyses with two behavioural manipulation checks to ensure that the task elicited the manipulations that we had planned. First, we sought to identify whether we correctly manipulated how often participants selected each judgment type (easy, hard) across the three conditions (high reward, medium reward, low reward). Yet, one caveat discovered in pilot testing was that there would be a series of trials at the beginning of each block that occurred before participant responses reached an equilibrium. As such, we first plotted participant responses across trials to visually determine the location at which responses stabilized, see Figure 17A. Through visual inspection, we determined that it generally took the first fifteen trials of each block for the values to stabilize. As such, we removed these trials from analyses for the following behavioural manipulation checks and EEG analyses.

With these trials removed, we determined whether response rates increased linearly across conditions from high reward to low reward (see Figure 17B). To do so, we employed linear mixed effects modelling (using R (RCore, 2021; RStudioTeam, 2016) packages lme4 (Bates et al., 2015) and lmerTest (Kuznetsova et al., 2017)) with a random intercept to determine whether there existed a linear increase of response rates across the three conditions, where response rates refers to the percentage of times the hard judgment was selected:

$$Response\_Rates = \beta_0 + \beta_1 Block\_Type + (1|Participant) + \varepsilon$$

Next, we needed to ensure that the rewards linearly decreased across conditions from high reward to low reward, see Figure 17C and 17D. We again employed linear mixed effects modelling with a random intercept to determine whether there existed a linear decrease across conditions:

$$Reward\_Values = \beta_0 + \beta_1 Block\_Type + (1|Participant) + \varepsilon$$

EEG Analyses

For each participant, preprocessing resulted in time-frequency data for each condition (high reward, medium reward, low reward) for the strategy selection phase (onset of judgment options with attached dollar values), see Figure 18. In line with past research (Cavanagh & Frank, 2014; Lin et al., 2018, 2022; Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Hassall, et al., 2022; Williams, Kappen, et al., 2019; Williams, VanOorschot, & Krigolson, 2021), our analyses focused on theta frequencies (e.g., 4–8 Hz). First, we created two difference time-frequency data for each participant, comparing each of the high reward and low reward condition to the medium reward condition. Specifically, we contrasted 1) high reward condition minus medium reward condition (see Figure 5A), and 2) low reward condition minus medium reward condition (see Figure 5B). As such, the medium reward can be considered a referent condition, in line with the linear and quadratic modelling that we will highlight momentarily. We then determined a theta cluster of interest using the collapsed localizer method (Luck, 2014; Luck & Gaspelin, 2017) by averaging the two difference time-frequency data and finding a cluster that exceeded a $0.40\mu V^2$ threshold (see Figure 5C). We then extracted conditional theta activity by averaging data within the cluster for each condition and participant. This resulted in a single theta power score for each condition for each participant, which we used for statistical investigations, see Figure 21.

As discussed in the introduction, we were here investigating whether neural signatures of reasoning tracked expected value of control and/or prediction error signals. If neural signals scaled to the expected value of control, then they would follow a linear increase from the low to high reward conditions. However, if they reflected prediction errors then we would expect to see a quadratic function in that the high and low reward conditions would have equal power to each other but different power from the medium reward condition. To analyze the expected value of control and prediction error patterns, we employed linear mixed effect modelling with a random intercept and slope in both a linear model:

$$Theta\_Power = \beta_0 + \beta_1 Block\_Type + (Block\_Type|Participant) + \varepsilon$$

and a quadratic model:

$$Theta\_Power = \beta_0 + \beta_1 Block\_Type + \beta_2 Block\_Type^2 + (Block\_Type|Participant) + \varepsilon$$

Results

Behavioural Manipulation Checks

We assessed behavioural measures to determine whether the task elicited the effects of interest. First, we needed to determine whether we properly manipulated response rates (i.e., the percentage of times that the hard judgment was selected) across the three conditions (high reward, medium reward, low reward). If our manipulation were successful, we would see a linear increase in selected hard judgments from the high reward to low reward conditions, and indeed this is what we found, $\beta_1 = 19.30$, $F(1, 113) = 285.22$, $p < .0001$, see Figure 17A & 17B.

Next, we needed to confirm that the averaged dollar value of the easy judgment was also properly manipulated across the conditions. A successful manipulation would find a linear decrease in values from the high reward to low reward conditions. Indeed, we found this relationship, $\beta_1 = -0.15$, $F(1, 113) = 128.11$, $p < .0001$, see Figure 17C & 17D. As such, our manipulation checks confirmed our block-wise manipulation, indicated by a linear increase in selecting the hard judgments, and a linear decrease in dollar values, from the high reward to the low reward conditions.

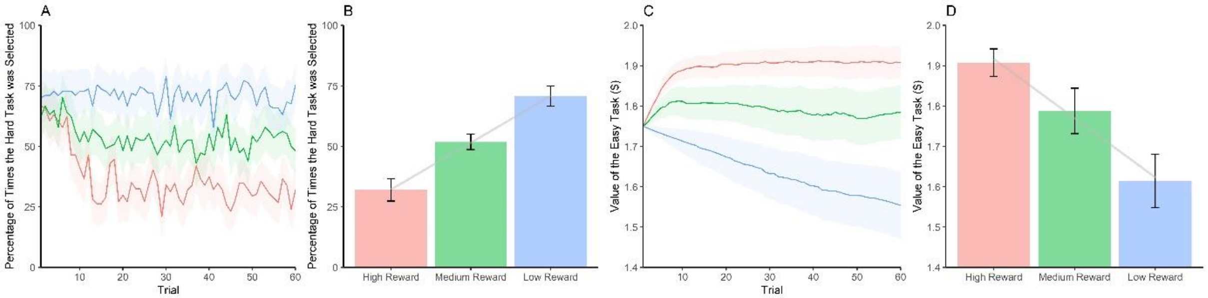

Figure 17: Experiment 4 (Figure 3)

Behavioural manipulation checks demonstrating averaged trial-by-trial changes with 95% confidence interval ribbons of hard judgments selected as a function of the three block types (Panel A), the averaged percentage of hard judgments selected for the three block types with 95% confidence intervals and a grey linear trend line (Panel B), the averaged trial-by-trial changes with 95% confidence interval ribbons of the easy judgment value as a function of the block types (Panel C), and the averaged value of the easy judgment for the three block types with 95% confidence intervals and a grey linear trend line (Panel D).

Electroencephalographic Findings

As described in the methods, we needed to determine a cluster of theta activity for analysis – see Figure 18 for conditional time-frequency plots, Figure 19 for difference time-frequency plots, and Figure 20 for topographic maps of conditional and difference theta activity. We determined theta clusters by creating two difference time-frequency transforms with the medium reward condition as the referent (i.e., high reward condition – medium reward condition, see Figure 19A & low reward condition – medium reward condition, see Figure 19B), averaging these difference time-frequency data (see Figure 19C), and using a threshold criterion of $0.40\mu V^2$

to isolate theta activity (see contour lines in Figure 19). The theta cluster encompassed 4-7Hz and 194-552ms, see Figure 19C. We then averaged the cluster for each condition and participant.

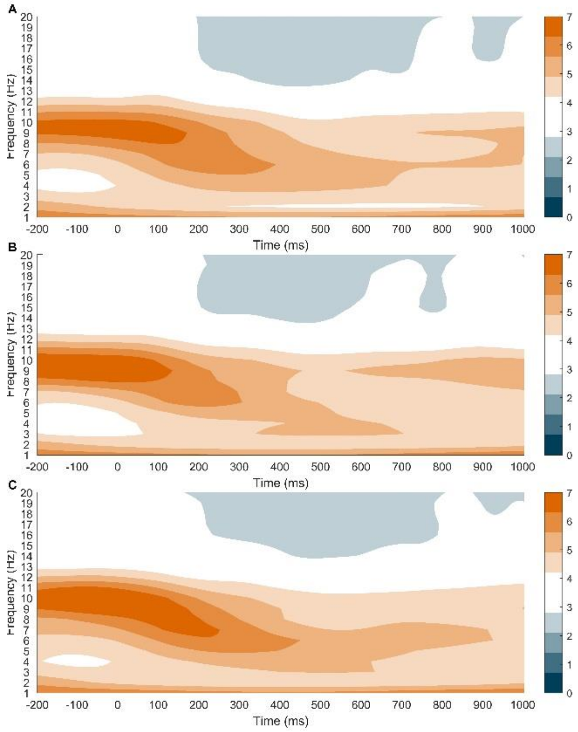

Figure 18: Experiment 4 (Figure 4)

Grand averaged conditional time-frequency plots of the High Reward (Panel A), Medium Reward (Panel B), and Low Reward (Panel C) block types at electrode FCz. Time zero indicates the onset of judgment options with attached reward dollar values. Our analyses focused on the theta range of these transforms, specifically within the 4-8Hz band.

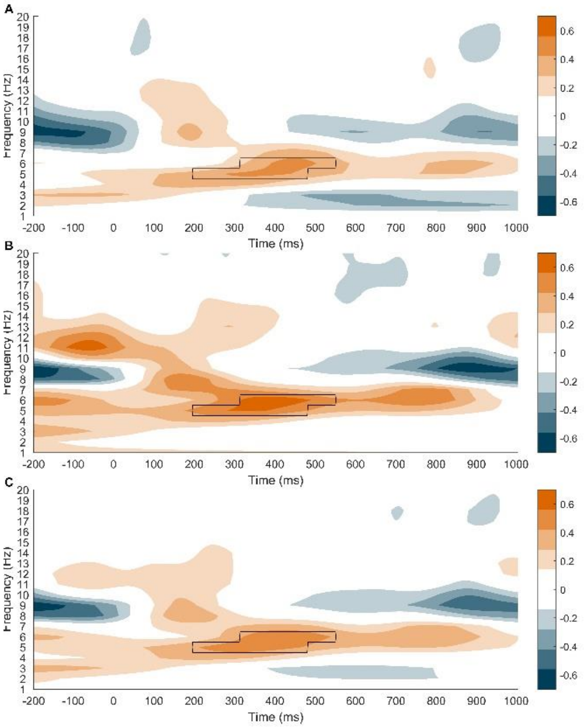

Figure 19: Experiment 4 (Figure 5)

Grand averaged difference time-frequency plots of High Reward - Medium Reward block types (Panel A), Low Reward - Medium Reward block types (Panel B), and the average of these two differences (Panel C), all at electrode FCz. Time zero indicates the onset of judgment options with attached reward dollar values. The theta cluster used for analysis was determined from the averaged difference time-frequency data (presented in Panel C), albeit here projected on all plots as a black contour line for comparison. As such, our analyses focused on a theta cluster with a frequency range of 4-7Hz and a time range of 194-552ms.

Using these data, we investigated whether theta activity corresponded to an expected value of control signal or a prediction error signal in that the former would conform to a linear trend and the latter to a quadratic trend. We found theta activity to follow a quadratic pattern, $\beta_1 = 0.07$, $\beta_2 = 0.48$, $F(1, 56) = 5.67$, $p = .0206$, rather than a linear pattern, $\beta_1 = 0.07$, $F(1, 56) = 0.19$, $p = .6606$, see Figure 21. Comparing the two models determined the quadratic trend as a better fit than the linear trend, $AIC_{linear} = 759.30$, $AIC_{quadratic} = 755.80$, $BIC_{linear} = 778.15$, $BIC_{quadratic} = 777.79$, $\chi^2(1, n = 57) = 5.50$, $p = .0190$. As such, findings indicated that theta activity reflected prediction errors of decision demands rather than the expected value of control.

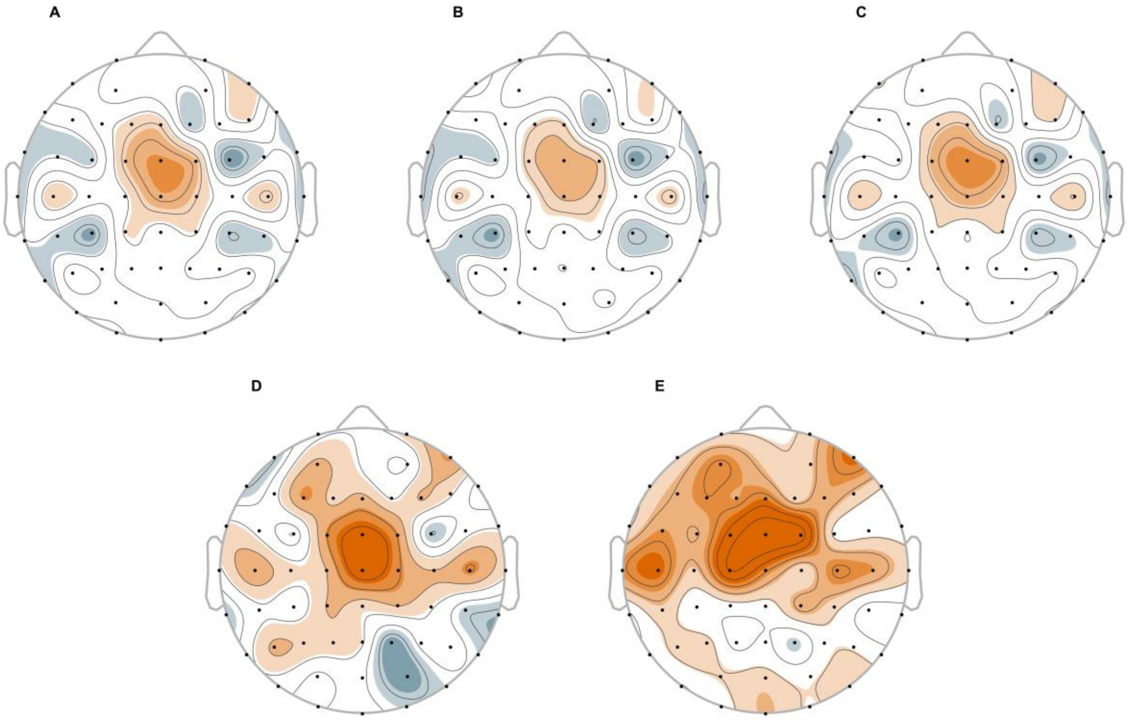

Figure 20: Experiment 4 (Figure 6)

Topographic plots of the theta cluster across the scalp for the High Reward (Panel A), Medium Reward (Panel B), and Low Reward (Panel C) block types. In addition, we present topographic plots of the High Reward - Medium Reward block types difference data (Panel D) and the Low Reward – Medium Reward block types difference data (Panel E).

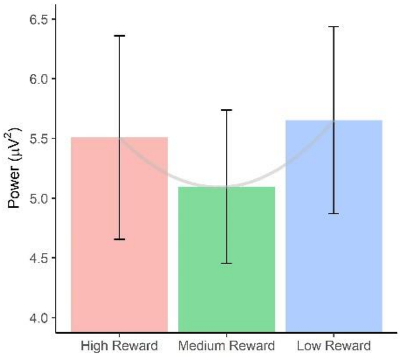

Figure 21: Experiment 4 (Figure 7)

The averaged theta power for the three block types with 95% confidence intervals and a grey quadratic trend line. The theta power was extracted from each conditional block type (High Reward, Medium Reward, Low Reward) using the determined theta cluster (4-7Hz, 192-552ms) at electrode FCz.

Discussion

We here provide neural evidence that our expectations are involved in deciding the level of effort to put into a decision. Specifically, we had participants make a series of decisions as to whether they would exert or withhold effort for different monetary rewards. This resulted in them building long-run expectations of decision demands (Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Hassall, et al., 2022). We additionally provided them problems where the effortless decisions were overvalued and undervalued – each leading to a deviation from their expectations and thus a prediction error (Alexander & Brown, 2011; J. W. Brown & Alexander, 2017; Holroyd & Krigolson, 2007; Krigolson, 2018; Proudfit, 2015; Vassena et al., 2020;

Williams et al., 2017; Williams, Ferguson, Hassall, Abimbola, et al., 2021; Williams, Hassall, et al., 2022). We observed these prediction errors as increased frontal theta activity and thus corroborated the involvement of expectations of decision demands in deciding how much effort to exert in a given decision.

Our findings support our proposal that expectations proactively account for decision demands while cost-benefit computations reactively determine the degree of effort to exert in a decision (Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Hassall, et al., 2022). Although we speculate that our expectations of decision demands function to relieve the burden of cost-benefit computations (Bustamante et al., 2021; Grahek et al., 2020; Lieder et al., 2018; Musslick, Cohen, et al., 2018; Musslick et al., 2015; Musslick & Cohen, 2019), we caution that we still do not know how this works. Perhaps it is that cost-benefit computations draw from our expectations to better fine tune our reasoning strategies. Expectations then provide our cost-benefit computations with a starting point to be adjusted if necessary (Musslick, Cohen, et al., 2018; Musslick et al., 2015). Another possibility could be that our expectations act as a gating mechanism for cost-benefit computations (Frömer & Shenhav, 2021; Shenhav et al., 2021). Specifically, if our expectations of decision demands are met when encountering a decision, no prediction error occurs and there is no need for a cost-benefit computation as we will exert our predisposed amount of effort. Yet if our expectations of decision demands are violated, a prediction error occurs, and a cost-benefit assessment is computed to determine the appropriate strategy for the decision at hand. In this hypothesis, cost-benefit computations are only conducted when decision demands do not meet our expectations. If everything is going as we predicted, then we simply adopt the same strategy as always.

In summary, we have provided neural evidence that our expectations of decision demands influence cost-benefit computations when determining the degree of effort to exert in a decision. This was determined by having participants make a series of decisions wherein the expected costs and benefits were equal, and additionally providing them with decisions that violated this expectation. The violation of these expectations resulted in a prediction error and was reflected as frontal theta activity. We propose that our expectations of decision demands serve to relieve the burden of cost-benefit computations, but caution that it is still unclear how this occurs. Finally, we speculate that our expectations could either tune each cost-benefit computation or act as a gating mechanism of these computations.

Chapter 6: General Discussion

I propose that prediction errors of decision demands influence cost-benefit computations, which determine the reasoning strategy to be employed. Specifically, repeated encounters of the same or similar decisions provide an opportunity to develop expectations of the prospective costs and benefits and this facilitates the reasoning process (Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Hammerstrom, et al., 2022; Williams, Hassall, et al., 2022). For each decision we are to make, we must first determine the degree of effort that we are going to exert, and this can range from no effort to full effort (Alexander & Brown, 2011; J. W. Brown & Alexander, 2017; Collins & Shenhav, 2022; Frömer & Shenhav, 2021; Kool et al., 2017; Kool & Botvinick, 2018; Shenhav et al., 2013, 2017, 2021; Vassena, Holroyd, et al., 2017; Vassena et al., 2020; Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Hassall, et al., 2022). Over time, we can learn the optimal degree of effort for familiar decisions and use this level of effort as the default – this process is considered stabilizing our expectations of decision demands (Williams, Hassall, et al., 2022). As such, we settle into a reasoning strategy as long as the decisions to be made are as expected and this lessens or releases the burden of cost-benefit computations for each judgment (Williams, Hammerstrom, et al., 2022). However, when the decision demands do not meet our expectations – i.e., they are either better or worse than expected – our predisposed reasoning strategy may no longer be optimal. These violations of our expectations are considered prediction errors of decision demands and signal the need to rely on cost-benefit computations more forcefully (Williams, Hammerstrom, et al., 2022). As such, the predisposed reasoning strategy is re-assessed through cost-benefit computations. If the strategy is still deemed optimal, it remains; however, if it is deemed sub-optimal then a new strategy is employed.

The findings across my four experiments here presented provide evidence for my proposition of the influence of prediction errors on cost-benefit computations during reasoning. These experiments demonstrate the building process of this proposition in that the first experiment provided evidence for the influence of theta oscillations within effortful reasoning (Williams, VanOorschot, & Krigolson, 2021), the second experiment revealed theta oscillations may reflect prediction errors within reasoning (Williams, Ferguson, Hassall, Wright, et al., 2021), the third experiment supported the presence of prediction errors within reasoning (Williams, Hassall, et al., 2022), and the fourth experiment tied the prediction errors to cost-benefit computations during the strategy selection phase of reasoning (Williams, Hammerstrom, et al., 2022). We will now consider each of these experiments in detail and how the findings relate to the proposition here presented.

The first experiment was inspired by classic reasoning research that utilizes word problems (Williams, VanOorschot, & Krigolson, 2021). A famous example from Frederick's (Frederick, 2005) cognitive reflection test is the bat-and-ball problem: "*A bat and a ball cost $1.10 in total. The bat costs $1.00 more than the ball. How much does the ball cost?*" In this problem, an intuitive incorrect response of ¢10 comes to mind immediately, but it requires effortful contemplation to achieve the correct response of ¢5. Another example of word problems, and those of interest in this research, is the base-rate problem (Kahneman & Tversky, 1973; Pennycook et al., 2014): "*There are clowns and accountants. Person A is funny. There are 30 clowns and 70 accountants. Is person A a clown or an accountant?*" Base-rate problems have two responses, one that focuses on the stereotypic information – here that Person A is funny – and the other that relies on the base-rate information – here that there are more accountants than clowns. Whereas responding in line with the stereotypic information is thought to rely on effortless intuitions, responding in line with the base-rate information is thought to rely on effortful contemplation (Pennycook et al., 2014). When presented with a series of these questions, people do not always respond in the same way and thus I was able to derive a continuum proportional to effortful processing as the percentage of times that individuals responded in line with the base-rate information (Williams, VanOorschot, & Krigolson, 2021). I found theta activity to strongly correlate with this effortful continuum and thus posited that theta activity was an indication of employing effortful reasoning strategies. At the time of this experiment, I was not considering prediction errors within reasoning, but these findings led to the conception of my next experiment that began to consider the decisional factors that lead to increased difficulty, and thus costs, in a decision.

The first experiment indicated that theta activity is enhanced during contemplative reasoning, but it was unclear what cognitive mechanism it reflected during this process. It could reflect several things such as the degree of effort being exerted or the tracking of decision difficulty. As such, I designed my next experiment to determine whether theta activity reflected the perceived costs of a decision as reflected by its difficulty (Botvinick & Cohen, 2014; Egner, 2011, 2017; Nigg, 2017). This experiment was inspired by research on accumulation models, which are computational frameworks that describe how responses are derived (S. D. Brown & Heathcote, 2008; Ratcliff et al., 2016; Rouder & Ratcliff, 2004). These models rely on tasks that present response conflicts – or in other words, situations wherein multiple responses may be similarly likely. I had participants diagnose a series of simulated patients into one of two diagnoses dependent on a single metric. This metric ranged on a continuum that contained an inflection point wherein any value below that point characterized one disease and any value above that point characterized the other disease. As the participants were blinded to what value described this inflection point, the proximity between the presented metric and this point was proportional to the difficulty of the question and as such the expected costs of exerting effort. If theta oscillations indeed reflected costs, it would too match this pattern. To my surprise, it did not conform to cost signals but rather demonstrated patterns in line with prediction errors (Alexander & Brown, 2011; Krigolson, 2018; Proudfit, 2015; Williams, Ferguson, Hassall, Abimbola, et al., 2021). Specifically, theta activity was lowest at the average of each disease's diagnostic range and increased with any deviations from these points. As such, I found evidence that theta oscillations reflected surprise, and so prediction errors, within reasoning. Further, I found another oscillatory signal, alpha and beta activity, to reflect decision difficulty and this implied that these prediction error signals may function in parallel to the assessment of decision difficulty and thus cost-benefit computations.

Together, the first experiment indicated theta oscillations may reflect mechanisms involved within effortful contemplation and the second experiment described the mechanism to be prediction errors. For my third experiment, I wanted to support the position that theta oscillations indeed reflected prediction errors during reasoning by drawing from a task that included a manipulation known to influence prediction errors. Specifically, we know that prediction errors diminish with learning as our expectations stabilize and so if theta oscillations indeed reflected prediction errors they too would diminish with learning (Krigolson et al., 2009, 2014; Williams, Hassall, et al., 2019). Participants were to categorize a series of complex polygons into their respective families. Participants familiarized themselves with each polygon family through trial-and-error, which consequently stabilized their expectations of the shapes being presented. Indeed, I found theta oscillations to diminish with learning, thus indicating that they could be reflecting prediction errors in reasoning. These findings alone, however, did not provide strong enough evidence for this conclusion because perhaps it was not the familiarity with the decisions that led to the diminishment of prediction errors and theta oscillations, but simply the time that has passed while performing the task. Critically, half-way through the experiment, participants were presented with two new families of polygons to categorize alongside two learned families of polygons. This manipulation allowed me to determine whether the diminished neural response reflected familiarity with the decisions or else the time taken to perform the task. If the former were true, prediction errors and thus theta oscillations, would remain diminished with the learned decisions but would be elicited with the novel decisions – or in other words, there would be a difference in activity between the learned and novel decisions (Krigolson et al., 2009, 2014; Williams, Hassall, et al., 2019). In contrast, if the diminished signal instead reflected the time on task, both the learned and novel decisions would present diminished activity. Indeed, theta oscillations presented as diminished for learned decisions but not for novel decisions, suggesting that they followed assumptions of, and therefore reflected the occurrence of, prediction error signals within reasoning.

When designing my final experiment, my research had thus far indicated that prediction errors of decision demands may have manifested during reasoning in parallel to cost-benefit computations. This final experiment was when I explicitly began to consider the reasoning process to involve two phases, specifically the strategy selection phase and the engagement phase. As such, I recognized that my first three experiments did not consider this dissociation and thus could not describe whether these prediction errors were in response to selecting a strategy and thus an influential factor in determining the degree of effort to exert via cost-benefit computations. As such, I designed a task that explicitly separated these two phases in that participants were first to determine whether they would engage with a decision that would require effort or not, each for different monetary rewards, and then engage with the selected decision. This task kept the costs of the decisions stable but systematically manipulated the benefit of each judgment to achieve decisions that had subjectively equal costs and benefits. Moreover, I designed the task so that some judgments were overvalued – i.e., the benefits outweighed the costs – and other judgments were undervalued – i.e., the costs outweighed the benefits. As such, I had participants select judgments that spanned from benefits outweighing costs to costs outweighing benefits (Alexander & Brown, 2011; J. W. Brown & Alexander, 2017; Collins & Shenhav, 2022; Frömer & Shenhav, 2021; Kool et al., 2017; Kool & Botvinick, 2018; Shenhav et al., 2013, 2017, 2021; Vassena et al., 2020; Vassena, Holroyd, et al., 2017; Williams, Ferguson, Hassall, Wright, et al., 2021; Williams, Hassall, et al., 2022). This approach elicited cost-benefit computations and so I investigated the influence of prediction errors within this strategy selection framework. Here, prediction errors would be expressed with decisions whose demands deviated from the long-run averaged demands – or in other words, prediction errors would be elicited only for the overvalued and undervalued decisions (Alexander & Brown, 2011; Krigolson, 2018; Proudfit, 2015; Williams, Ferguson, Hassall, Abimbola, et al., 2021). As such, if theta oscillations truly reflected the consideration of prediction errors alongside cost-benefit computations during the strategy selection phase of reasoning, it would follow this same pattern. Indeed, theta activity followed this pattern of results, supporting the influence of prediction errors of decision demands on cost-benefit computations.

Altogether, I have provided a progression of evidence that ultimately led to my proposition that prediction errors of decision demands influence cost-benefit computations, yet my research has still not revealed how this influence is expressed. I speculate two possible methods for the interaction of these mechanisms, each in need of future research for testing. The first is that stabilizing our expectations of decision demands results in a lessened burden of cost-benefit computations. Specifically, if each decision requires a cost-benefit computation, our expectations may provide a starting point within the cost-benefit space that facilitates the strategy selection process. In terms of the EVC model, the cost-benefit space can range from benefits far outweighing costs to costs far outweighing benefits. As such, this explanation anchors the expected decision demands by drawing on our past experiences. This process draws from computational expressions of the EVC model, which assume the starting point of cost-benefit computational space to be uninformed or else randomly determined for each computation (Musslick, Cohen, et al., 2018; Musslick et al., 2015). As such, the inclusion of prediction errors in this framework would cause a foundational change by using past experience to inform cost-benefit computations.

My second explanation for the interaction of prediction errors and cost-benefit computations is that prediction errors act as a gating mechanism in strategy selection not to lessen the burden of cost-benefit computations but rather to remove the need for them altogether (Frömer & Shenhav, 2021; Shenhav et al., 2021). With familiar decisions, we have an established expectation of the costs and benefits of a judgment. This consequently means that we have an established degree of effort necessary for these judgments, which we adopt as our default. As such, when encountering similar decisions, we can rely on our adopted strategy rather than engaging cost-benefit computations to determine the degree of effort to exert. However, we can only rely on our adopted strategy as long as the decision at hand prescribes to our expectations. If the demands of a decision violates our expectations, then a prediction error emerges to signal the need to re-assess our strategy. This results in a cost-benefit computation to determine the degree of effort to exert. This assessment could lead to the same strategy being employed or else a new strategy is adopted.

Here, the critical difference between these two explanations is whether cost-benefit computations are conducted for each decision (as in the first explanation) or else only when our predisposed strategy needs re-assessment (as in the second explanation). To me, the second explanation seems more intuitively likely, as it seems inefficient to continuously compute cost-benefit calculations with every decision, especially when we are well-familiarized with the judgments. However, we know that our intuitions can be faulty and as such I will lean on future research to continue probing the relationship between prediction errors of decision demands and cost-benefit computations in reasoning.

In summary, I sought to determine the mechanisms underlying selecting a reasoning strategy by leveraging neural imaging techniques. My investigations began by exploring neural correlates of effortful contemplation and evolved to test assumptions of prediction errors as it seemed that they were an influential factor. Once providing evidence for the presence of prediction errors in reasoning, I tied this mechanism to the strategy selection phase of reasoning and cost-benefit computations. From these findings, I proposed that prediction errors of decision demands function to lessen or remove the burden of cost-benefit computations by either informing these computations to ease their computational needs or by acting as a gating mechanism to ensure that these computations are only recruited when necessary. Although there is still much work to do to test my propositions and further understand strategy selection, I hope that my contributions will serve towards understanding how we reason in our ever-growing complicated world.

References

Alexander, W. H., & Brown, J. W. (2011). Medial prefrontal cortex as an action-outcome predictor. *Nature Publishing Group, 14*(10), 1338–1344. https://doi.org/10.1038/nn.2921

Alexander, W. H., & Brown, J. W. (2014). A general role for medial prefrontal cortex in event prediction. *Frontiers in Computational Neuroscience*, *8*, 69. https://doi.org/10.3389/fncom.2014.00069

Alexander, W. H., & Vassena, E. (2020). Selective application of theory does not replace formal model tests : a reply to Shenhav et al . *PsyArXiv*, 1–7.

Apps, M. A. J., Grima, L. L., Manohar, S., & Husain, M. (2015). The role of cognitive effort in subjective reward devaluation and risky decision-making. *Nature Publishing Group*, *5*(1), 1–11. https://doi.org/10.1038/srep16880

Bago, B., Frey, D., Vidal, J., Houdé, O., Borst, G., & Neys, W. D. (2018). Fast and Slow Thinking: Electrophysiological evidence for early conflict sensitivity. *Neuropsychologia*, *117*, 483–490. https://doi.org/10.1016/j.neuropsychologia.2018.07.017

Bago, B., & Neys, W. D. (2018). The Intuitive Greater Good: Testing the Corrective Dual Process Model of Moral Cognition. *Journal of Experimental Psychology: General*, *148*(10), 1782–1801. https://doi.org/10.1037/xge0000533

Bago, B., & Neys, W. D. (2019). The Smart System 1: evidence for the intuitive nature of correct responding on the bat-and-ball problem. *Thinking & Reasoning*, *25*(3), 1–43. https://doi.org/10.1080/13546783.2018.1507949

Banks, A. P. (2017). Comparing dual process theories: Evidence from event - related potential. In W. D. Neys (Ed.), *Dual Process Theory 2.0* (pp. 66–81).

Bannister, S. L., Au, H., Forbes, K. L., Zucker, M., & Weiler, G. (2016). A Card Is Worth aThousand Questions: Leveraging the Power of Key Features and Semantic Qualifiersin Writing Clinical Questions. *Council on Medical Student Education in Pediatrics 2016 Annual Meeting*.

Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, *67*(1), 1–48.

Baumeister, R. F., Bratslavsky, E., Muraven, M., & Tice, D. M. (1998). Ego Depletion: Is the Active Self a Limited Resource? *Journal of Personality and Social Psychology*, *74*(5), 1252–1265. https://doi.org/10.1037/0022-3514.74.5.1252

Baumeister, R. F., & Heatherton, T. F. (1996). Self-Regulation Failure: An Overview. *Psychological Inquiry*, *7*(1), 1–15. https://doi.org/10.1207/s15327965pli0701_1

Bernat, E. M., Nelson, L. D., & Baskin-Sommers, A. R. (2015). Time-frequency theta and delta measures index separable components of feedback processing in a gambling task. *Psychophysiology*, *52*(5), 626–637. https://doi.org/10.1111/psyp.12390

Bernat, E. M., Nelson, L. D., Holroyd, C. B., Gehring, W. J., & Patrick, C. J. (2008). *Separating cognitive processes with principal components analysis of EEG time-frequency distributions*. 70740S-70740S – 10. https://doi.org/10.1117/12.801362

Bernat, E. M., Nelson, L. D., Steele, V. R., Gehring, W. J., & Patrick, C. J. (2011). Externalizing Psychopathology and Gain–Loss Feedback in a Simulated Gambling Task: Dissociable Components of Brain Response Revealed by Time-Frequency Analysis. *Journal of Abnormal Psychology*, *120*(2), 352–364. https://doi.org/10.1037/a0022124

Botvinick, M. M., & Cohen, J. D. (2014). The computational and neural basis of cognitive control: Charted territory and new frontiers. *Cognitive Science*, *38*(6), 1249–1285. https://doi.org/10.1111/cogs.12126

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*(4), 433–436. https://doi.org/10.1163/156856897x00357

Braver, T. S., Krug, M. K., Chiew, K. S., Kool, W., Westbrook, J. A., Clement, N. J., Adcock, R. A., Barch, D. M., Botvinick, M. M., Carver, C. S., Cools, R., Custers, R., Dickinson, A., Dweck, C. S., Fishbach, A., Gollwitzer, P. M., Hess, T. M., Isaacowitz, D. M., Mather, M., … Somerville, L. H. (2014). Mechanisms of motivation–cognition interaction: challenges and opportunities. *Cognitive, Affective, & Behavioral Neuroscience*, *14*(2), 443–472. https://doi.org/10.3758/s13415-014-0300-0

Brown, J. W. (2013). Beyond Conflict Monitoring : Cognitive Control and the Neural Basis of Thinking Before You Act. *Current Directions in Psychological Science*, *22*(3), 179–185. https://doi.org/10.1177/0963721412470685

Brown, J. W., & Alexander, W. H. (2017). Foraging Value , Risk Avoidance , and Multiple Control Signals : How the Anterior Cingulate Cortex Controls. *Jordan Journal of Civil Engineering*, *10*(4), 1656–1673. https://doi.org/10.1162/jocn

Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, *57*(3), 153–178. https://doi.org/10.1016/j.cogpsych.2007.12.002

Burak, K. W., McLaughlin, K. J., Coderre, S. P., Busche, K. D., & Raman, M. (2015). The FlippedClassroom Improves Exam Performance in Undergraduate Medical Education. *The Office of Health and Medical Education Scholarship Symposium, Calgary, AB.*

Bustamante, L., Lieder, F., Musslick, S., Shenhav, A., & Cohen, J. (2021). Learning to Overexert Cognitive Control in a Stroop Task. *Cognitive, Affective, & Behavioral Neuroscience*, *21*(3), 453–471. https://doi.org/10.3758/s13415-020-00845-x

Cavanagh, J. F., Figueroa, C. M., Cohen, M. X., & Frank, M. J. (2012). Frontal theta reflects uncertainty and unexpectedness during exploration and exploitation. *Cerebral Cortex*, *22*(11), 2575–2586. https://doi.org/10.1093/cercor/bhr332

Cavanagh, J. F., & Frank, M. J. (2014). Frontal theta as a mechanism for cognitive control. *Trends in Cognitive Sciences*, *18*(8), 414–421. https://doi.org/10.1016/j.tics.2014.04.012

Cavanagh, J. F., & Shackman, A. J. (2015). Frontal midline theta reflects anxiety and cognitive control: Meta-analytic evidence. *Journal of Physiology Paris*, *109*(1–3), 3–15. https://doi.org/10.1016/j.jphysparis.2014.04.003

Cavanagh, J. F., Zambrano-Vazquez, L., & Allen, J. J. B. (2012). Theta lingua franca: A common mid-frontal substrate for action monitoring processes. *Psychophysiology*, *49*(2), 220–238. https://doi.org/10.1111/j.1469-8986.2011.01293.x

Christie, S. T., & Schrater, P. (2015). Cognitive cost as dynamic allocation of energetic resources. *Frontiers in Neuroscience*, *9*, 289. https://doi.org/10.3389/fnins.2015.00289

Cohen, M. X. (2014). *Analyzing neural time series data: theory and practice: Vol. MIT Press*.

Cohen, M. X. (2017). Neurophysiological Oscillations and Action Monitoring. In T. Egner, *The Wiley Handbook of Cognitive Control* (pp. 242–258). John Wiley & Sons, Incorporated.

Cohen, M. X., & Donner, T. H. (2013). Midfrontal conflict-related theta-band power reflects neural oscillations that predict behavior. *Journal of Neurophysiology*, *110*(12), 2752–2763. https://doi.org/10.1152/jn.00479.2013

Cohen, M. X., Ridderinkhof, K. R., Haupt, S., Elger, C. E., & Fell, J. (2008). Medial frontal cortex and response conflict: Evidence from human intracranial EEG and medial frontal cortex lesion. *Brain Research*, *1238*(Cereb. Cortex. 12 2002), 127–142. https://doi.org/10.1016/j.brainres.2008.07.114

Collins, A. G. E., & Shenhav, A. (2022). Advances in modeling learning and decision-making in neuroscience. *Neuropsychopharmacology*, *47*(1), 104–118. https://doi.org/10.1038/s41386-021-01126-y

Cowan, N. (2012). *Working Memory Capacity*. Psychology press.

Croskerry, P. (2009). A universal model of diagnostic reasoning. *Acad Med*, *84*(8), 1022–1028. https://doi.org/10.1097/acm.0b013e3181ace703

Croskerry, P. (2017). A Model for Clinical Decision-Making in Medicine. *Medical Science Educator*, *27*(Suppl 1), 9–13. https://doi.org/10.1007/s40670-017-0499-9

Dai, Z., Souza, J. de, Lim, J., Ho, P. M., Chen, Y., Li, J., Thakor, N., Bezerianos, A., & Sun, Y. (2017). EEG Cortical Connectivity Analysis of Working Memory Reveals Topological Reorganization in Theta and Alpha Bands. *Frontiers in Human Neuroscience*, *11*, 237. https://doi.org/10.3389/fnhum.2017.00237

Dang, J. (2016). Testing the role of glucose in self-control: A meta-analysis. *Appetite*, *107*, 222–230. https://doi.org/10.1016/j.appet.2016.07.021

DeLoof, E., Vassena, E., Janssens, C., Taeye, L. D., Meurs, A., Roost, D. V., Boon, P., Raedt, R., & Verguts, T. (2019). Preparing for hard times : Scalp and intracranial physiological signatures of proactive cognitive control. *Psychophysiology*, *56*(10), 1–14. https://doi.org/10.1111/psyp.13417

Delorme, A., & Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics. *Journal of Neuroscience Methods*, *134*, 9–21.

DeNeys, W. (2017). Bias, conflict, and fast logic: Towards a hybrid dual process future? In *Dual Process Theory 2.0* (pp. 1–35). Routledge. https://doi.org/10.1080/13546783.2013.854725

DeNeys, W., & Pennycook, G. (2019). Logic, Fast and Slow: Advances in Dual-Process Theorizing. *Current Directions in Psychological Science*, *28*(5), 503–509. https://doi.org/10.1177/0963721419855658

DeNeys, W., Vartanian, O., & Goel, V. (2008). Smarter Than We Think. When Our Brains Detect That We Are Biased. *Psychological Science*, *19*(5), 483–489. https://doi.org/10.1111/j.1467-9280.2008.02113.x

Egner, T. (2007). Congruency sequence effects. *Cognitive, Affective, & Behavioral Neuroscience*, *7*(4), 380–390.

Egner, T. (2011). Right Ventrolateral Prefrontal Cortex Mediates Individual Differences in Conflict-driven Cognitive Control. *Journal of Cognitive Neuroscience, 23*(12), 3903–3913. https://doi.org/10.1162/jocn_a_00064

Egner, T. (2017). *The Wiley handbook of cognitive control* (["Egner & T."], Eds.). John Wiley & Sons.

Eisma, J., Rawls, E., Long, S., Mach, R., & Lamm, C. (2021). Frontal midline theta differentiates separate cognitive control strategies while still generalizing the need for cognitive control. *Scientific Reports*, *11*(1), 14641. https://doi.org/10.1038/s41598-021-94162-z

Engel, A. K., & Fries, P. (2010). Beta-band oscillations-signalling the status quo? *Current Opinion in Neurobiology*, *20*(2), 156–165. https://doi.org/10.1016/j.conb.2010.02.015

Ericsson, K. A., Krampe, R. Th., & Tesch-Römer, C. (1993). The Role of Deliberate Practice in the Acquisition of Expert Performance. *Psychological Review*, *100*(3), 363–406. https://doi.org/10.1037/0033-295x.100.3.363

Evans, J. S. B. T. (2011). Dual-process theories of reasoning: Contemporary issues and developmental applications. *Developmental Review*, *31*(2–3), 86–102. https://doi.org/10.1016/j.dr.2011.07.007

Evans, J. S. B. T., & Stanovich, K. E. (2013). Dual-Process Theories of Higher Cognition: Advancing the Debate. *Perspectives on Psychological Science*, *8*(3), 223–241. https://doi.org/10.1177/1745691612460685

Feng, S. F., Schwemmer, M., Gershman, S. J., & Cohen, J. D. (2014). Multitasking versus multiplexing: Toward a normative account of limitations in the simultaneous execution of control-demanding behaviors. *Cognitive, Affective, & Behavioral Neuroscience*, *14*(1), 129–146. https://doi.org/10.3758/s13415-013-0236-9

Finley, A. J., Tang, D., & Schmeichel, B. J. (2019). Sweet Nothings. *Social Psychology*, *50*(5–6), 322–331. https://doi.org/10.1027/1864-9335/a000376

Foti, D., Weinberg, A., Bernat, E. M., & Proudfit, G. H. (2015). Anterior cingulate activity to monetary loss and basal ganglia activity to monetary gain uniquely contribute to the feedback negativity. *Clinical Neurophysiology*, *126*(7), 1338–1347. https://doi.org/10.1016/j.clinph.2014.08.025

Frederick, S. (2005). Cognitive Reflection and Decision Making. *Journal of Economic Perspectives*, *19*(4), 25–42. https://doi.org/10.1257/089533005775196732

Frömer, R., Lin, H., Wolf, C. K. D., Inzlicht, M., & Shenhav, A. (2021). Expectations of reward and efficacy guide cognitive control allocation. *Nature Communications*, *12*(1), 1030. https://doi.org/10.1038/s41467-021-21315-z

Frömer, R., & Shenhav, A. (2021). Filling the gaps: Cognitive control as a critical lens for understanding mechanisms of value-based decision-making. *Neuroscience & Biobehavioral Reviews*, *134*, 104483. https://doi.org/10.1016/j.neubiorev.2021.12.006

Gailliot, M. T., Baumeister, R. F., DeWall, C. N., Maner, J. K., Plant, E. A., Tice, D. M., Brewer, L. E., & Schmeichel, B. J. (2007). Self-Control Relies on Glucose as a Limited Energy Source: Willpower Is More Than a Metaphor. *Journal of Personality and Social Psychology*, *92*(2), 325–336. https://doi.org/10.1037/0022-3514.92.2.325

Grahek, I., Musslick, S., & Shenhav, A. (2020). A computational perspective on the roles of affect in cognitive control. *International Journal of Psychophysiology*, *151*(February), 25–34. https://doi.org/10.1016/j.ijpsycho.2020.02.001

Haller, M., Donoghue, T., Peterson, E., Varma, P., & Sebastian, P. (2018). Parameterizing neural power spectra. *Biorxiv*, 1–16.

Holroyd, C. B. (2015). The Waste Disposal Problem of Effortful Control. In T. S. Braver (Ed.), *The Waste Disposal Problem of Effortful Control* (p. 26). Routledge.

Holroyd, C. B., & Krigolson, O. E. (2007). Reward prediction error signals associated with a modified time estimation task. *Psychophysiology*, *44*(6), 913–917. https://doi.org/10.1111/j.1469-8986.2007.00561.x

Horrey, K., Keegan, D., Paget, M. K., & Tan, A. (2016). Creating open access FM micro-cases for online medical student learning. *2016 Conference on Medical Student Education. The Society of Teachers of Family Medicine, Phoenix, AZ.*

Hsieh, L.-T., & Ranganath, C. (2014). Frontal midline theta oscillations during working memory maintenance and episodic encoding and retrieval. *NeuroImage*, *85*, 721–729. https://doi.org/10.1016/j.neuroimage.2013.08.003

Inzlicht, M., & Berkman, E. (2015). Six Questions for the Resource Model of Control (and Some Answers). *Social and Personality Psychology Compass*, *9*(10), 511–524. https://doi.org/10.1111/spc3.12200

Inzlicht, M., & Friese, M. (2019). The Past, Present, and Future of Ego Depletion. *Social Psychology*, *50*(5–6), 370–378. https://doi.org/10.1027/1864-9335/a000398

Inzlicht, M., & Schmeichel, B. J. (2012). What Is Ego Depletion ? Toward a Mechanistic Revision of the Resource Model of Self-Control Self-Control Exertion at Time 1 Self-Control Failure at Time 2. *Perspectives on Psychological Science*, *7*(5), 450–463. https://doi.org/10.1177/1745691612454134

Inzlicht, M., Schmeichel, B. J., & Macrae, C. N. (2014). Why self-control seems (but may not be) limited. *Trends in Cognitive Sciences*, *18*(3), 127–133. https://doi.org/10.1016/j.tics.2013.12.009

Inzlicht, M., Shenhav, A., & Olivola, C. Y. (2018). The Effort Paradox: Effort Is Both Costly and Valued. *Trends in Cognitive Sciences*, *22*(4), 337–349. https://doi.org/10.1016/j.tics.2018.01.007

Jiang, J., Beck, J., Heller, K., & Egner, T. (2015). An insula-frontostriatal network mediates flexible cognitive control by adaptively predicting changing control demands. *Nature Communications*, *6*(May), 8165. https://doi.org/10.1038/ncomms9165

Kahneman, D. (2011). *Thinking, Fast and Slow*. Macmillian.

Kahneman, D., & Frederick, S. (2002). *Representativeness revisited: Attribute substitution in intuitive judgment. 49*, 81.

Kahneman, D., Peavler, S., & Onuska, L. (1968). Effects of verbalization and incentive on the pupil response to mental activity. *Canadian Journal of Psychology/Revue Canadienne de Psychologie*, *22*(3), 186–196. https://doi.org/10.1037/h0082759

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, *80*(4), 237–251. https://doi.org/10.1037/h0034747

Kazoleas, K. (2016). Med students practice clinical problem solving by playing cards. *UCalgary Medicine Magazine Winte*.

Kool, W., & Botvinick, M. (2014). A Labor/Leisure Tradeoff in Cognitive Control. *Journal of Experimental Psychology: General*, *143*(1), 131–141. https://doi.org/10.1037/a0031048

Kool, W., & Botvinick, M. (2018). Mental labour. *Nature Human Behaviour*, *2*(December), 899–908. https://doi.org/10.1038/s41562-018-0401-9

Kool, W., McGuire, J. T., Rosen, Z. B., & Botvinick, M. M. (2010). Decision Making and the Avoidance of Cognitive Demand. *Journal of Experimental Psychology: General*, *139*(4), 665–682. https://doi.org/10.1037/a0020198

Kool, W., Mcguire, J. T., Wang, G. J., & Botvinick, M. M. (2013). Neural and Behavioral Evidence for an Intrinsic Cost of. *PLoS ONE*, *8*(8), 1–6. https://doi.org/10.1371/journal.pone.0072626

Kool, W., Shenhav, A., & Botvinick, M. M. (2017). Cognitive Control as Cost-Benefit Decision Making. In *The Wiley Handbook of Cognitive Control* (pp. 167–189). https://doi.org/10.1002/9781118920497.ch10

Krajbich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, *13*(10), 1292–1298. https://doi.org/10.1038/nn.2635

Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(33), 13852–13857. https://doi.org/10.1073/pnas.1101328108

Krigolson, O. E. (2018). Event-related brain potentials and the study of reward processing: Methodological considerations. *International Journal of Psychophysiology*, *132*(November 2017), 175–183. https://doi.org/10.1016/j.ijpsycho.2017.11.007

Krigolson, O. E., Hassall, C. D., & Handy, T. C. (2014). How We Learn to Make Decisions: Rapid Propagation of Reinforcement Learning Prediction Errors in Humans. *Journal of Cognitive Neuroscience*, *26*(3), 1–10. https://doi.org/10.1162/jocn_a_00509

Krigolson, O. E., Pierce, L. J., Holroyd, C. B., & Tanaka, J. W. (2009). Learning to Become an Expert: Reinforcement Learning and the Acquisition of Perceptual Expertise. *Journal of Cognitive Neuroscience*, *21*(9), 1833–1840. https://doi.org/10.1162/jocn.2009.21128

Kurzban, R. (2010). Does the Brain Consume Additional Glucose during Self-Control Tasks? *Evolutionary Psychology*, *8*(2), 147470491000800200. https://doi.org/10.1177/147470491000800208

Kurzban, R., Duckworth, A., Kable, J. W., & Myers, J. (2013). An opportunity cost model of subjective effort and task performance. *Behavioral and Brain Sciences*, *36*(6), 661–679. https://doi.org/10.1017/s0140525x12003196

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest Package: Tests in Linear Mixed Effects Models. *Journal of Statistical Software*, *82*(13). https://doi.org/10.18637/jss.v082.i13

Lieder, F., Shenhav, A., Musslick, S., & Griffiths, T. L. (2018). Rational metareasoning and the plasticity of cognitive control. *PLoS Computational Biology*, *14*(4), e1006043. https://doi.org/10.1371/journal.pcbi.1006043

Lin, H., Ristic, J., Inzlicht, M., & Otto, R. (2022). The opportunity cost of time modulates behavioral and neural indices of effortful control allocation. *PsyarXiv*.

Lin, H., Saunders, B., Friese, M., Evans, N. J., & Inzlicht, M. (2020). Strong Effort Manipulations Reduce Response Caution: A Preregistered Reinvention of the Ego-Depletion Paradigm. *Psychological Science, 31*(5), 531–547. https://doi.org/10.1177/0956797620904990

Lin, H., Saunders, B., Hutcherson, C. A., & Inzlicht, M. (2017). Midfrontal theta and pupil dilation parametrically track subjective conflict (but also surprise) during intertemporal choice. *BioRxiv*, 172122. https://doi.org/10.1101/172122

Lin, H., Saunders, B., Hutcherson, C., & Inzlicht, M. (2018). Midfrontal theta and pupil dilation parametrically track subjective conflict (but also surprise) during intertemporal choice. *NeuroImage*, *172,* 838–852.

Luck, S. J. (2014). *An introduction to the event-related potential technique*. MIT Press.

Luck, S. J., & Gaspelin, N. (2017). How to get statistically significant effects in any ERP experiment (and why you shouldn ’ t). *Psychophysiology*, *54*(1), 146–157. https://doi.org/10.1111/psyp.12639

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, *390*(6657), 279–281. https://doi.org/10.1038/36846

Luft, C. D. B. (2014). Learning from feedback: The neural mechanisms of feedback processing facilitating better performance. *Behavioural Brain Research*, *261*, 356–368. https://doi.org/10.1016/j.bbr.2013.12.043

Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nature Neuroscience*, *17*(3), 347–356. https://doi.org/10.1038/nn.3655

Monahan, J. S. (2001). Coloring Single Stroop Elements: Reducing Automaticity or Slowing Color Processing? *The Journal of General Psychology*, *128*(1), 98–112. https://doi.org/10.1080/00221300109598901

Müller, T., & Apps, M. A. J. (2018). Motivational fatigue: A neurocognitive framework for the impact of effortful exertion on subsequent motivation. *Neuropsychologia*, *123*, 141–151. https://doi.org/10.1016/j.neuropsychologia.2018.04.030

Muraven, M., Tice, D. M., & Baumeister, R. F. (1998). Self-Control as Limited Resource: Regulatory Depletion Patterns. *Journal of Personality and Social Psychology*, *74*(3), 774–789. https://doi.org/10.1037/0022-3514.74.3.774

Musslick, S., & Cohen, J. D. (2019). A Mechanistic Account of Constraints on Control-Dependent Processing: Shared Representation, Conflict and Persistence. *CogSci*.

Musslick, S., & Cohen, J. D. (2021). Rationalizing constraints on the capacity for cognitive control. *Trends in Cognitive Sciences*, *25*(9), 757–775. https://doi.org/10.1016/j.tics.2021.06.001

Musslick, S., Cohen, J. D., & Shenhav, A. (2018). Estimating the costs of cognitive control from task performance: theoretical validation and potential pitfalls. *CogSci*, 798–803.

Musslick, S., Cohen, J. D., & Shenhav, A. (2019). Decomposing individual differences in cognitive control: A model-based approach. *CogSci*, 2427–2433.

Musslick, S., Jang, S. J., Shvartsman, M., Shenhav, A., & Cohen, J. D. (2018). Constraints associated with cognitive control and the stability-flexibility dilemma. *CogSci*.

Musslick, S., Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2015). A computational model of control allocation based on the Expected Value of Control. *RLDM*, 1–4.

Nathoo, F. S., Kilshaw, R. E., & Masson, M. E. J. (2018). A better (Bayesian) interval estimate for within-subject designs. *Journal of Mathematical Psychology*, *86*, 1–9. https://doi.org/10.1016/j.jmp.2018.07.005

Nelson, L. D., Patrick, C. J., Collins, P., Lang, A. R., & Bernat, E. M. (2011). Alcohol impairs brain reactivity to explicit loss feedback. *Psychopharmacology*, *218*(2), 419. https://doi.org/10.1007/s00213-011-2323-3

Nigg, J. T. (2017). Annual Research Review: On the relations among self-regulation, self-control, executive functioning, effortful control, cognitive control, impulsivity, risk-taking, and inhibition for developmental psychopathology. *Journal of Child Psychology and Psychiatry*, *58*(4), 361–383. https://doi.org/10.1111/jcpp.12675

Norman, G. R. (2000). The Epistemology of Clinical Reasoning. *Academic Medicine*, *75*(10), S127–S133. https://doi.org/10.1097/00001888-200010001-00041

Norman, G. R., Monteiro, S. D., Sherbino, J., Ilgen, J. S., Schmidt, H. G., & Mamede, S. (2017). The Causes of Errors in Clinical Reasoning. *Academic Medicine*, *92*(1), 23–30. https://doi.org/10.1097/acm.0000000000001421

Oberauer, K., Farrell, S., Jarrold, C., & Lewandowsky, S. (2016). What Limits Working Memory Capacity? *Psychological Bulletin*, *142*(7), 758–799. https://doi.org/10.1037/bul0000046

Pennycook, G. (2017). *A perspective on the theoretical foundation of dual - process models: Vol. Dual Process Theory 2.0* (pp. 5–27). Psychology Press. https://doi.org/10.4324/9781315204550-2

Pennycook, G., Fugelsang, J. A., & Koehler, D. J. (2015). What makes us think? A three-stage dual-process model of analytic engagement. *Cognitive Psychology*, *80*, 34–72. https://doi.org/10.1016/j.cogpsych.2015.05.001

Pennycook, G., Neys, W. D., Evans, J. St. B. T., Stanovich, K. E., & Thompson, V. A. (2018). The Mythical Dual-Process Typology. *Trends in Cognitive Sciences*, *22*(8), 667–668. https://doi.org/10.1016/j.tics.2018.04.008

Pennycook, G., Trippas, D., Handley, S. J., & Thompson, V. A. (2014). Base Rates: Both Neglected and Intuitive. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*(2), 544–554. https://doi.org/10.1037/a0034887

Proudfit, G. H. (2015). The reward positivity: From basic research on reward to a biomarker for depression. *Psychophysiology*, *52*(4), 449–459. https://doi.org/10.1111/psyp.12370

Ratcliff, R., Smith, P. L., Brown, S. D., & Mckoon, G. (2016). Diffusion Decision Model : Current Issues and History. *Trends in Cognitive Sciences*, *20*(4), 260–281. https://doi.org/10.1016/j.tics.2016.01.007

RCore. (2021). *R: A Language and Environment for Statistical Computing*. https://www.r-project.org/

Rouder, J. N., & Ratcliff, R. (2004). Comparing categorization models. *Journal of Experimental Psychology: General*, *133*(1), 63–82. https://doi.org/10.1037/0096-3445.133.1.63

RStudioTeam. (2016). *RStudio: Integrated Development Environment for R*. http://www.rstudio.com/

Sambrook, T. D., & Goslin, J. (2015). A Neural Reward Prediction Error Revealed by a Meta-Analysis of ERPs Using Great Grand Averages. *Psychological Bulletin*, *141*(1), 213–235. https://doi.org/10.1037/bul0000006

Schevernels, H., Krebs, R. M., Santens, P., Woldorff, M. G., & Boehler, C. N. (2014). Task preparation processes related to reward prediction precede those related to task-difficulty expectation. *NeuroImage*, *84*, 639–647. https://doi.org/10.1016/j.neuroimage.2013.09.039

Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). Review The Expected Value of Control : An Integrative Theory of Anterior Cingulate Cortex Function. *Neuron*, *79*(2), 217–240. https://doi.org/10.1016/j.neuron.2013.07.007

Shenhav, A., Cohen, J. D., & Botvinick, M. M. (2016). Dorsal anterior cingulate cortex and the value of control. *Nature Neuroscience*, *19*(10), 4–6. https://doi.org/10.1038/nn.4382

Shenhav, A., Fahey, M. P., & Grahek, I. (2021). Decomposing the Motivation to Exert Mental Effort. *Current Directions in Psychological Science*, *30*(4), 307–314. https://doi.org/10.1177/09637214211009510

Shenhav, A., Musslick, S., Botvinick, M. M., & Cohen, J. D. (2020). *Misdirected vigor: Differentiating the control of value from the value of control*. https://doi.org/10.31234/osf.io/5bhwe

Shenhav, A., Musslick, S., Lieder, F., Kool, W., Griffiths, T. L., Cohen, J. D., & Botvinick, M. M. (2017). Toward a Rational and Mechanistic Account of Mental Effort. *Annual Review of Neuroscience*, *40*(1), 1–26. https://doi.org/10.1146/annurev-neuro-072116-031526

Shenhav, A., Straccia, M. A., Cohen, J. D., & Botvinick, M. M. (2014). Anterior cingulate engagement in a foraging context reflects choice difficulty , not foraging value. *Nature Publishing Group, 17*(9), 1249–1254. https://doi.org/10.1038/nn.3771

Tajima, S., Drugowitsch, J., & Pouget, A. (2016). Optimal policy for value-based decision-making. *Nature Communications*, *7*(1), 1–12. https://doi.org/10.1038/ncomms12400

Tang, A., Kwan, E., Paget, M., Coderre, S., Burak, K., & McLaughlin, K. (2016). Cards: A Novel, Case-based Method for Undergraduate Medical Students to Learn Key Concepts in Geriatric. *Presented at the 36th Annual Scientific Meeting of the Canadian GeriatricsSociety. (Vancouver, BC).*

Tolkien, J. R. R. (1937). *The Hobbit*. George Allen & Unwin.

Umemoto, A., Inzlicht, M., & Holroyd, C. B. (2019). Neuropsychologia Electrophysiological indices of anterior cingulate cortex function reveal changing levels of cognitive e ff ort and reward valuation that sustain task performance. *Neuropsychologia*, *123*(October 2017), 67–76. https://doi.org/10.1016/j.neuropsychologia.2018.06.010

Vadillo, M. A., Gold, N., & Osman, M. (2016). The Bitter Truth About Sugar and Willpower. *Psychological Science*, *27*(9), 1207–1214. https://doi.org/10.1177/0956797616654911

Vartanian, O., Beatty, E. L., Smith, I., Blackler, K., Lam, Q., Forbes, S., & Neys, W. D. (2018). The Reflective Mind: Examining Individual Differences in Susceptibility to Base Rate Neglect with fMRI. *Journal of Cognitive Neuroscience*, *30*(7), 1011–1022. https://doi.org/10.1162/jocn_a_01264

Vassena, E., Deraeve, J., & Alexander, W. H. (2017). Predicting Motivation: Computational Models of PFC Can Explain Neural Coding of Motivation and Effort-based Decision-making in Health and Disease. *Journal of Cognitive Neuroscience*, *29*(10), 1633–1645. https://doi.org/10.1162/jocn_a_01160

Vassena, E., Deraeve, J., & Alexander, W. H. (2020). Surprise, value and control in anterior cingulate cortex during speeded decision-making. *Nature Human Behaviour*, *4*(4), 412–422. https://doi.org/10.1038/s41562-019-0801-5

Vassena, E., Holroyd, C. B., & Alexander, W. H. (2017). Computational Models of Anterior Cingulate Cortex : At the Crossroads between Prediction and Effort. *Frontiers in Neuroscience*, *11*(June), 1–9. https://doi.org/10.3389/fnins.2017.00316

Walsh, M. M., & Anderson, J. R. (2012). Learning from experience: Event-related potential correlates of reward processing, neural adaptation, and behavioral choice. *Neuroscience & Biobehavioral Reviews*, *36*(8), 1870–1884. https://doi.org/10.1016/j.neubiorev.2012.05.008

Westbrook, A., Kester, D., & Braver, T. S. (2013). What Is the Subjective Cost of Cognitive Effort ? Load , Trait , and Aging Effects Revealed by Economic Preference. *PLoS ONE*, *8*(7), 1–8. https://doi.org/10.1371/journal.pone.0068210

Wickham, H. (2016). *ggplot2: Elegant Graphics for Data Analysis.* Springer-Verlag, New York.

Williams, C. C., Ferguson, T. D., Hassall, C. D., Abimbola, W., & Krigolson, O. E. (2021). The ERP, frequency, and time–frequency correlates of feedback processing: Insights from a large sample study. *Psychophysiology*, *58*(2), e13722. https://doi.org/10.1111/psyp.13722

Williams, C. C., Ferguson, T. D., Hassall, C. D., Wright, B., & Krigolson, O. E. (2021). Dissociated neural signals of conflict and surprise in effortful decision Making: Theta activity reflects surprise while alpha and beta activity reflect conflict. *Neuropsychologia*, *155*, 107793. https://doi.org/10.1016/j.neuropsychologia.2021.107793

Williams, C. C., Hammerstrom, M. R., Inzlicht, M., & Krigolson, O. E. (2022). Expectations of Decision Demands Influence Cost-Benefit Computations During Reasoning. *PsyArXiv*.

Williams, C. C., Hassall, C. D., & Krigolson, O. E. (2022). Stabilizing Expectations when Shifting from Analytical to Intuitive Reasoning: The Role of Prediction Errors in Reasoning. *Psyarxiv*.

Williams, C. C., Hassall, C. D., Lindenbach, T., & Krigolson, O. E. (2019). Reward Prediction Errors Reflect an Underlying Learning Process That Parallels Behavioural Adaptations: A Trial-to-Trial Analysis. *Computational Brain & Behavior*, 1–11. https://doi.org/10.1007/s42113-019-00069-4

Williams, C. C., Hassall, C. D., Trska, R., Holroyd, C. B., & Krigolson, O. E. (2017). When theory and biology differ: The relationship between reward prediction errors and expectancy. *Biological Psychology*, *129*, 265–272. https://doi.org/10.1016/j.biopsycho.2017.09.007

Williams, C. C., Hecker, K. G., Paget, M. K., Coderre, S. P., Burak, K. W., Wright, B., & Krigolson, O. E. (2018). The application of reward learning in the real world: Changes in the reward positivity amplitude reflect learning in a medical education context. *International Journal of Psychophysiology*, *132*(Med. Sci. Educ. 26 1 2016), 236–242. https://doi.org/10.1016/j.ijpsycho.2017.10.010

Williams, C. C., Kappen, M., Hassall, C. D., Wright, B., & Krigolson, O. E. (2019). Thinking theta and alpha: Mechanisms of intuitive and analytical reasoning. *NeuroImage*, *189*(December 2018), 574–580. https://doi.org/10.1016/j.neuroimage.2019.01.048

Williams, C. C., VanOorschot, F., & Krigolson, O. E. (2021). A Window into the Rational Mind: The Neural Underpinnings of Human Reasoning. *PsyArXiv*, 1–21. https://doi.org/10.31234/osf.io/gc6u9

Winter, B. (2013). What These Models Are Doing łdots and Then We'Ll Spend Most of the Time. *ArXiv*, *Tutorial 2*, 1–22. https://doi.org/http://arxiv.org/pdf/1308.5499.pdf

Yee, D. M., Leng, X., Shenhav, A., & Braver, T. S. (2021). Aversive Motivation and Cognitive Control. *Neuroscience & Biobehavioral Reviews*, *133*, 104493. https://doi.org/10.1016/j.neubiorev.2021.12.016