

Neurocomputational Accounts of Choice Variability and Affect during Decision-making

Chew Jie Han Benjamin

Max Planck UCL Centre for Computational Psychiatry and Ageing Research
Wellcome Centre for Human Neuroimaging
UCL

September 2019

A dissertation submitted in partial fulfilment of the requirements for the degree
of Doctor of Philosophy at University College London

I, Benjamin Chew, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Benjamin Chew

Abstract

Humans exhibit surprising variability in behaviour, often making different choices under identical conditions. While the outcomes of these choices typically lead to explicit rewards that have been shown to influence subsequent affective states, less well understood is how the brain represents rewards that are intrinsically meaningful to an individual.

The first part of this thesis examines the contributions of endogenous fluctuations in brain activity to behaviour. Resting-state studies suggest that ongoing endogenous fluctuations in brain activity can influence low-level perceptual and motor processes but it remains unknown whether such fluctuations also influence high-level cognitive processes including decision making. Using a novel application of real-time functional magnetic resonance imaging, I find that low pre-stimulus brain activity lead to increased occurrences of risky choice. Using computational modeling, I show that greater risk taking is explained by enhanced phasic responses to offers in a decision network. These findings demonstrate that endogenous brain activity provides a physiological basis for variability in complex behaviour. I then examine how the neuroanatomy of the brain in the form of tissue microstructure relates to risk preferences by leveraging on *in vivo* histology using magnetic resonance imaging.

The second part of this thesis investigates how experienced events, such as rewards received following choice, are aggregated into affective states. Despite their relevance to ideas like goal-setting and well-being, little is known about the impact of intrinsic rewards on affective states and their representation in the brain. A reinforcement learning task incorporating a skilled performance component that did not influence payment was developed to examine this. Computational modeling revealed that momentary happiness depended on past extrinsic rewards and also intrinsic rewards related to the experience of successful skilled performance. Individuals for whom intrinsic rewards more strongly influence momentary happiness exhibit stronger ventromedial prefrontal cortex responses for successful skilled performance. These findings show that the ventromedial prefrontal cortex represents the subjective value of intrinsic rewards, and that computational models of mood dynamics provide a tool that can be used to measure implicit values of abstract goods and experiences.

Table of Contents

Abstract	3
Table of Contents	4
Impact Statement	7
Acknowledgements	8
Chapter 1: Introduction	10
1.1 <i>Conceptual Overview</i>	10
1.2 <i>Key Definitions</i>	11
1.2.1 <i>Rewards, Values, Risk, and Decisions</i>	11
1.2.2 <i>Subjective well-being</i>	13
1.3 <i>Thesis Outline</i>	13
References	15
Chapter 2: Literature Review	17
2.1 <i>Representation of Value</i>	17
2.2 <i>Valuation</i>	21
2.3 <i>Learning</i>	23
2.3.1 <i>Habitual or Model-free Systems</i>	24
2.3.2 <i>Goal-directed or Model-based Systems</i>	24
2.3.3 <i>Successor Representation</i>	25
2.3.4 <i>Dopamine: Learning and Decision-Making</i>	25
2.4 <i>Intrinsic Variability of Neural Signals and Behaviour</i>	29
References	33
Chapter 3: Non-invasive Methods for Investigating Activity and Tissue Microstructure in the Human Brain	42
3.1 <i>Blood Oxygen Level Dependent (BOLD) fMRI</i>	43
3.1.1 <i>MR Physics</i>	43
3.1.2 <i>Image Construction</i>	46
3.1.3 <i>What does the BOLD response reflect?</i>	47
3.2 <i>Real-time fMRI (rtfMRI) - Setup, Software, and Implementation</i>	48
3.2.1 <i>Setup and Software</i>	48
3.2.2 <i>Pre-processing of Images</i>	49
3.2.3 <i>Dealing with Additional Nuisance Regressors</i>	50
3.2.4 <i>Computation of Blood Oxygenation Level Dependent (BOLD) Signal and Level of Activity</i>	51
3.2.5 <i>Quantifying Level of BOLD Activity</i>	53

3.3 <i>Quantitative Magnetic Resonance Imaging</i>	55
3.4 <i>Drawing and Preparation of SN/VTA ROI for rtfMRI</i>	56
3.4.1 <i>Magnetization Transfer</i>	56
3.4.2 <i>Definition of ROI</i>	57
3.4.3 <i>Transformation of ROI</i>	58
3.5 <i>fMRI BOLD response and the SN/VTA complex</i>	58
References.....	63
Chapter 4: Endogenous fluctuations in the dopaminergic midbrain drive behavioural choice variability	70
4.1 <i>Methods</i>	72
4.1.2 <i>Procedure</i>	73
4.1.3 <i>Real-time fMRI</i>	75
4.1.4 <i>Offline Analyses</i>	77
4.2 <i>Results</i>	78
4.2.1 <i>Endogenous fluctuations in SN/VTA BOLD activity modulate risk taking</i>	78
4.2.2 <i>A computational mechanism for the effect of endogenous fluctuations on risk taking</i>	82
4.2.3 <i>Endogenous fluctuations affect phasic responses during choice</i>	86
4.3 <i>Discussion</i>	94
References.....	98
Chapter 5: Risk preferences correlate with tissue microstructure	104
5.1 <i>Methods</i>	105
5.1.1 <i>Participants</i>	105
5.1.2 <i>Study Design</i>	106
5.1.3 <i>Structural Data Acquisition and Processing</i>	107
5.2 <i>Results</i>	108
5.3 <i>Discussion</i>	112
References.....	115
Chapter 6: A neurocomputational model of mood and intrinsic rewards	120
6.1 <i>Methods</i>	122
6.1.1 <i>Participants</i>	122
6.1.2 <i>Procedure</i>	122
6.1.3 <i>Image Acquisition</i>	124
6.2 <i>Results</i>	125
6.2.1 <i>Happiness is modulated by reward outcomes and performance</i>	125
6.2.2 <i>Computational model of subjective well-being</i>	126

6.2.3 Neural correlates of extrinsic rewards and skilled performance	129
6.3 Discussion	132
References	135
Chapter 7: Conclusions and general discussion	139
7.1 Variability in the brain	139
7.2 Subjective well-being as a computational hedonometer	141
References	143

Impact Statement

The findings of this thesis provide further insight into the neural processes underpinning human decision making under risk. It details a brand new application of an existing technology that can be used to inform future work. It highlights an often overlooked role for endogenous brain activity when it comes to explaining why humans exhibit inconsistencies in their choice behaviour even when faced with identical options, underscoring the importance of taking time when making decisions as a different decision may be reached following a temporal delay. The findings demonstrate that higher-order cognition is influenced by fluctuations in internal brain states, providing a physiological basis for variability in complex human behaviour. Accounting for the influence of endogenous neural fluctuations on behaviour could be critical for future work looking to understand the neurobiological processes underlying cognition in health and disorder. Further, factoring the contribution of endogenous fluctuations in brain activity to cognition could be useful in the field of brain-machine interfaces to improve decodability of thoughts and actions.

Another avenue for impact based on the findings of this thesis is policy-making. High subjective well-being has been shown to convey many benefits to individuals and societies, yet it has been difficult to implement policies that reliably improve the well-being of people. Part of the problem stems from inaccuracies and biases when people are asked to estimate the value of hypothetical goods or experiences, or when they have to judge the impact of future events on their affective states. For example, when people are asked about the value of public goods like parks, there may be discrepancies between what they ought to report and whether such goods actually maximise their subjective well-being. The findings in this thesis outline an approach to estimate the value of such items within a framework grounded in computational modeling and neuroscience, providing a more reliable method of measuring the hedonic benefits of experiences and abstract goods that could then be used to inform policies.

Acknowledgements

The experiments detailed in this thesis would never have been possible without the wonderful support of all the scientists and support staff at the Max Planck UCL Centre and FIL. I have to begin by thanking my supervisors. I am extremely grateful to have been supported by Ray. His remarkable ability to find patterns in findings across domains is inspiring, as is his ability to quote James Joyce during lab meetings. The seminars in Ringberg were always jolly good fun and I am glad to have been invited to those. Robb has provided me with a constant stream of advice and pushed the boundaries of my scientific rigour. He taught me the important lessons of strategizing and keeping an eye on the bigger picture, and delivered encouragement that helped me push through the occasional rough patch. Tobias has been a delight to work with. I enjoyed the crazy ideas that pop up during our discussions and the social atmosphere that he fostered in his group. His high productivity despite his role as the social engine of the lab is nothing short of astounding.

I thank all my family – Dad, Mum, Marcus, and Charissa – whose love and support kept me going despite the large geographical distance between us. I enjoyed showing you guys the bits of London that I came to love over the course of my program.

The support staff at the FIL have been instrumental to the smooth running of everything behind the scenes. David and Peter were always at the top of their game when things went awry, and the radiographers Clive, Megan, Elisa, and Elaine made the long hours of scanning all the more bearable with their upbeat professionalism.

I also have to thank my labmates who are the most amazing bunch of people that I have had the privilege of working alongside. My desk buddies Sam and Rachel kept me grounded and were an incredible source of comfort and contagious enthusiasm when I was uncertain of my way forward. Johanna's cheerful disposition and ready smile helped brighten up my days more than she could ever imagine. Her breadth of knowledge was inspirational and I am grateful our paths crossed. Alex and I were already friends prior to the start of her PhD and it has been wonderful hanging out and watching her grow as a scientist. Magda possessed the most impressive tea collection outside of a commercial setting and was equally generous with it. Her curious nature often made for engaging conversations. Alisa's willingness to stand up

and support her opinions in the face of strong opposition was often a valuable lesson for me. My interactions with Yunzhe has almost always proved enlightening, and I can only hope that he found comfort in my presence as well. I am still keen to continue our discussions on optimal political systems one day. Toby introduced me to the world of insane cookies, and I will never forget our exploits in search of the best cookies New York and London had to offer. Bastien was supportive of my scientific efforts which I am thankful for, and I never cease to be amazed by the diversity of his organic supplements. Susan, the co-organiser of our informal lab dinners, was a picture of enthusiasm and efficiency which made organising social events so much easier. Hrvoje, Daniel, and Elliott have provided me with much food for thought throughout our impromptu conversations and I really appreciate how approachable they were. Rani and Mehdi have been great teachers from whom I have learnt a lot. I especially enjoyed watching Rani's progress as a stand-up comedian. Yuki has provided a breath of fresh air in the lab with her quirks and antics. Matilde and Jolanda have been a friendly source of encouragement and positivity throughout the time I have known them. Francesco took me in as a Masters student and paved the road for the scientist I am today. Akshay, Andrea, Andrei, Archy, Dorothea, Eran, Erie, Geert-Jan, Giles, Jessica, Jochen, Laurence, Liam, Lorenz, Lucy, Matthew, Max, Michael, Misun, Mona, Philip, Quentin, Ritwik, and Zeb have made life in the lab so much more pleasant and meaningful with their readiness to chat and their contagious enthusiasm for science. Toyah has also been a fantastic help whenever I needed any sort of assistance.

Last but not least, I am also grateful for my friends in Goodenough College and Singapore who provided me with respite from work. There are probably too many others to thank, and I hope the people that I have overlooked know that they have made a positive impact on my life.

Chapter 1: Introduction

This thesis addresses how the human brain supports value-based decision-making and moment-to-moment fluctuations in subjective well-being. It builds on a large body of work describing the functional neuroanatomy of choice under risk, and contributes to a growing field of research investigating the link between decision-making and changes in subjective feelings. In this chapter, I provide a conceptual overview of the work, define key terms that recur throughout the thesis, and present an outline of the thesis chapters.

We think, each of us, that we're much more rational than we are. And we think that we make our decisions because we have good reasons to make them. Even when it's the other way around – We believe in the reasons, because we've already made the decision.

Daniel Kahneman

1.1 Conceptual Overview

Decision-making is a pervasive part of everyday life, ranging from which breakfast item to have on any given day to more complicated decisions such as which cryptocurrency to invest in to increase one's likelihood of turning a profit. To enhance survivability in the face of environmental demands, it is not sufficient for the human brain to simply make a random choice when faced with a decision but it should instead evaluate available options and choose the best course of action. When the outcome of such an action is subsequently revealed, it is often accompanied by subjective feelings such as happiness or regret depending on whether the outcome was good or bad.

Many decisions contain elements of uncertainty that can be broadly attributed to the external world or internal states (D. Kahneman & Tversky, 1982), also known as external uncertainty and internal uncertainty respectively (Howell & Burnett, 1978). The former is associated with properties of the environment such as win-lose probabilities within a gamble, while the latter is associated with properties that reside with an agent such as a lack of knowledge (Volz, Schubotz, & von Cramon, 2004). While many studies have examined the contributions of external uncertainty to the

computation of subjective value in an agent's choice processes (Christopoulos, Tobler, Bossaerts, Dolan, & Schultz, 2009; Klein-Flügge, Barron, Brodersen, Dolan, & Behrens, 2013), few have examined the role of internal states such as subjective feelings (Rutledge, Skandali, Dayan, & Dolan, 2014), contextual representation (Louie, Khaw, & Glimcher, 2013; Rigoli, Rutledge, Dayan, & Dolan, 2016), and intrinsic variability of neural activity (Garrett et al., 2013) in relation to choice behaviour.

The experimental work in this thesis is concerned with how internal states either contribute to or arise as a result of decision-making processes. Through the use of human neuroimaging data, we will investigate a candidate mechanism that induce variability in economic choice when subjects have a probabilistic chance at obtaining rewards, look at the influence of brain microstructure on how risk is perceived, and finally examine how subjective well-being is altered as a function of implicit and explicit rewards. A common feature underlying the various analyses will be the use of computational modelling to provide mechanistic explanations for observed behaviour.

1.2 Key Definitions

We begin with some definitions of frequently encountered terms to reduce any uncertainty associated with them.

1.2.1 Rewards, Values, Risk, and Decisions

A central tenet of value-based decision-making is the concept of *reward*. Reward is an operational construct which refers to a something a subject finds desirable such as an item, action, or internal state (Schultz, Dayan, & Montague, 1997). A closely related term is *value* which refers to a scalar estimate or predictor of reward within frameworks like reinforcement learning, whereby an agent exclusively seeks to maximize long-run cumulative reward by seeking out states of highest value (Sutton & Barto, 2018). Ecologically, the reward signal is likely to be a vector rather than scalar and may be an abstract summary statistic derived from various brain circuits (Schultz et al., 1997; Sutton & Barto, 2018). In microeconomics and finance, value (or *utility*) can be further dissociated into *objective value* and *subjective value*. The former refers to an explicit measure of reward such as amount of cash or drops of juice, while the

latter refers to an internal valuation of the reward (Daniel Kahneman & Tversky, 1979; Rangel, Camerer, & Montague, 2008; von Neumann & Morgenstern, 1944). Subjective values cannot be measured directly but are instead revealed through choice behaviour as a proxy of “outward phenomena to which they [desires] give rise” (Marshall, 1920). *Risk* is present when a choice involves options associated with probabilistic outcomes, and can alter subjective values as a function of an individual’s risk attitude (Tobler, Christopoulos, O’Doherty, Dolan, & Schultz, 2009). For example, the subjective value of an option with an equal probability of receiving \$10 or nothing will be higher than that of a certain gain of \$5 for a risk-seeking compared to risk-averse individual, and identical for someone who is risk-neutral (Fig. 1).

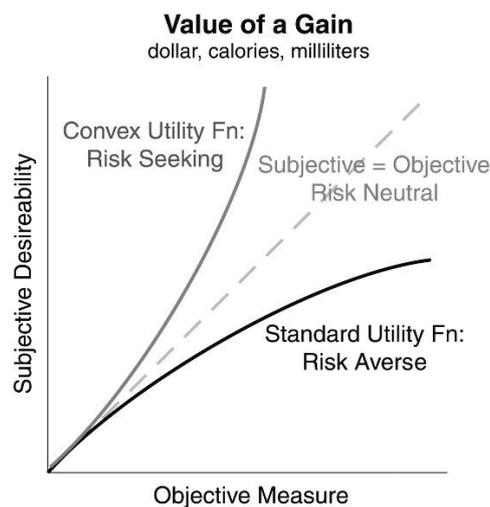


Fig. 1. Relationship between risk and utility. Solid black line: As the objective value of a gain increases, the subjective value similarly increases but at a slower pace. This accounts for the observation that people are generally averse to risk. Solid grey line: Conversely, risk seeking behaviour can be observed when the subjective value of a gain increases at a faster pace than the objective value. Dashed grey line: When the objective and subjective values of a gain grow at the same rate, risk neutral behaviour is predicted. Adapted from (Glimcher & Rustichini, 2004).

Finally, *decisions* refer to actions or choices selected by a subject when faced with options (Tobler et al., 2009). While actions can be sub-divided into categories such as Pavlovian, habitual, or goal-directed, these will be further touched on in **chapter 2**.

1.2.2 Subjective well-being

A key index of quality of life is subjective well-being (Oswald & Wu, 2010) which “refers to how people experience and evaluate their lives and specific domains and activities in their lives” (National Research Council, 2013). In 2009, the Report by the Commission on the Measurement of Economic Performance and Social Progress concluded that economic indicators such as Gross Domestic Product (GDP) alone were inadequate for assessing societal progress, and recommended that measures of well-being also be included to capture an extensive spectrum of nonmarket phenomena such as social connections and insecurities (Stiglitz, Sen, & Fitoussi, 2009). This makes intuitive sense for many of us, suggesting that the link between financial wealth and subjective well-being is a complex one (Easterlin, McVey, Switek, Sawangfa, & Zweig, 2010; Kahneman & Deaton, 2010). Recognizing the importance of well-being to public policies, an annual World Happiness Report has been published by the United Nations Sustainable Development Solutions Network since 2012 (Helliwell, Layard, & Sachs, 2018).

How can subjective well-being be measured? A popular method used in the fields of psychology and economics is the use of experience sampling whereby subjects are prompted to report their thoughts and feelings in real time, often over the course of several hours, days, or weeks (Kahneman & Krueger, 2006; Larson & Csikszentmihalyi, 1983). Similar to the United Nations report, the experiment described in **chapter 6** will use the terms “subjective well-being” and “happiness” interchangeably and draw upon experience sampling to derive individual measures of subjective well-being.

1.3 Thesis Outline

In the next chapter, I will review the literature on value-based decision-making with an emphasis on the involvement of the dopaminergic system in the processing of reward. I will also briefly touch on the distinction between modular and distributed processes underlying choice behavior. **Chapter 3** delves into the neuroimaging methods used in the empirical work of this thesis which is covered by **chapters 4, 5, and 6**.

Why are humans so often irrational? Economic theories struggle to account for the choice inconsistencies that humans exhibit. Recent advances in the study of the brain at rest suggest that spontaneous fluctuations are important for understanding the functional organization of the brain but the functional role of these fluctuations in behavior remains unknown. In **chapter 4**, I develop a novel real-time functional magnetic resonance imaging (rtfMRI) paradigm to examine the influence of such fluctuations in the substantia nigra and ventral tegmental areas on risk preferences, leading to some interesting findings. Drawing on recent technical developments that have enabled in vivo mapping of neuroimaging markers of biologically relevant quantities to be performed with high resolution, **chapter 5** investigates the relationship between risk preferences and brain microstructure. Together, these chapters provide novel insight into how risk preferences are influenced by the neurophysiology of the brain.

The neurotransmitter dopamine has been implicated in the processing of risk, rewards, and subjective-wellbeing. Recent animal studies have uncovered a dopaminergic signal in the ventral striatum that ramps up as animals approach a goal, yet the relationship between dopaminergic signals and subjective well-being are unknown when subjects are motivated by intrinsic as opposed to extrinsic rewards. With **Chapter 6**, I move on to the next part of the empirical work and ask how subjective well-being is influenced by rewards and the experience of skilled performance, providing an important contribution to the existing literature.

Finally, I will discuss the general implications of these works in **Chapter 7**.

References

- Christopoulos, G. I., Tobler, P. N., Bossaerts, P., Dolan, R. J., & Schultz, W. (2009). Neural correlates of value, risk, and risk aversion contributing to decision making under risk. *Journal of Neuroscience*, *29*(40), 12574–12583.
<https://doi.org/10.1523/JNEUROSCI.2614-09.2009>
- Easterlin, R. A., McVey, L. A., Switek, M., Sawangfa, O., & Zweig, J. S. (2010). The happiness–income paradox revisited. *PNAS*, *107*(52), 22463–22468.
<https://doi.org/10.1073/pnas.1015962107>
- Garrett, D. D., Samanez-Larkin, G. R., MacDonald, S. W. S., Lindenberger, U., McIntosh, A. R., & Grady, C. L. (2013). Moment-to-moment brain signal variability: A next frontier in human brain mapping? *Neuroscience & Biobehavioral Reviews*, *37*(4), 610–624.
<https://doi.org/10.1016/j.neubiorev.2013.02.015>
- Glimcher, P. W., & Rustichini, A. (2004). Neuroeconomics: The Consilience of Brain and Decision. *Science*, *306*(5695), 447–452. <https://doi.org/10.1126/science.1102566>
- Helliwell, J. F., Layard, R., & Sachs, J. D. (2018). *World Happiness Report 2018*. Retrieved from <http://worldhappiness.report/ed/2018/>
- Howell, W. C., & Burnett, S. A. (1978). Uncertainty measurement: A cognitive taxonomy. *Organizational Behavior and Human Performance*, *22*(1), 45–68.
[https://doi.org/10.1016/0030-5073\(78\)90004-1](https://doi.org/10.1016/0030-5073(78)90004-1)
- Kahneman, D., & Tversky, A. (1982). Variants of uncertainty. *Cognition*, *11*(2), 143–157.
- Kahneman, D., & Deaton, A. (2010). High income improves evaluation of life but not emotional well-being. *PNAS*, *107*(38), 16489–16493.
<https://doi.org/10.1073/pnas.1011492107>
- Kahneman, D., & Krueger, A. B. (2006). Developments in the measurement of subjective well-being. *Journal of Economic Perspectives*, *20*(1), 3–24.
<https://doi.org/10.1257/089533006776526030>
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291. <https://doi.org/10.2307/1914185>
- Klein-Flügge, M. C., Barron, H. C., Brodersen, K. H., Dolan, R. J., & Behrens, T. E. J. (2013). Segregated encoding of reward-identity and stimulus-reward associations in human orbitofrontal cortex. *Journal of Neuroscience*, *33*(7), 3202–3211.
<https://doi.org/10.1523/JNEUROSCI.2532-12.2013>
- Larson, R., & Csikszentmihalyi, M. (1983). The experience sampling method. *New Directions for Methodology of Social & Behavioral Science*, *15*, 41–56.

- Louie, K., Khaw, M. W., & Glimcher, P. W. (2013). Normalization is a general neural mechanism for context-dependent decision making. *PNAS*, 201217854. <https://doi.org/10.1073/pnas.1217854110>
- Marshall, A. (1920). *Principles of economics: An introductory volume* (8th ed.). London: Macmillan.
- National Research Council. (2013). *Subjective well-being: Measuring happiness, suffering, and other dimensions of experience*. <https://doi.org/10.17226/18548>
- Oswald, A. J., & Wu, S. (2010). Objective confirmation of subjective measures of human well-being: Evidence from the U.S.A. *Science*, 327(5965), 576–579. <https://doi.org/10.1126/science.1180606>
- Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, 9(7), 545. <https://doi.org/10.1038/nrn2357>
- Rigoli, F., Rutledge, R. B., Dayan, P., & Dolan, R. J. (2016). The influence of contextual reward statistics on risk preference. *NeuroImage*, 128, 74–84. <https://doi.org/10.1016/j.neuroimage.2015.12.016>
- Rutledge, R. B., Skandali, N., Dayan, P., & Dolan, R. J. (2014). A computational and neural model of momentary subjective well-being. *PNAS*, 111(33), 12252–12257. <https://doi.org/10.1073/pnas.1407535111>
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593–1599.
- Stiglitz, J., Sen, A., & Fitoussi, J.-P. (2009). Report by the commission on the measurement of economic performance and social progress. *The Commission on the Measurement of Economic Performance and Social Progress*.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning, An Introduction* (2nd ed.). Retrieved from <https://mitpress.mit.edu/books/reinforcement-learning-second-edition>
- Tobler, P. N., Christopoulos, G. I., O'Doherty, J. P., Dolan, R. J., & Schultz, W. (2009). Risk-dependent reward value signal in human prefrontal cortex. *PNAS*, 106(17), 7185–7190. <https://doi.org/10.1073/pnas.0809599106>
- Volz, K. G., Schubotz, R. I., & von Cramon, D. Y. (2004). Why am I unsure? Internal and external attributions of uncertainty dissociated by fMRI. *NeuroImage*, 21(3), 848–857. <https://doi.org/10.1016/j.neuroimage.2003.10.028>
- von Neumann, J., & Morgenstern, O. (1944). *Theory of Games and Economic Behavior* (60th Anniversary Commemorative). Retrieved from <https://press.princeton.edu/titles/7802.html>

Chapter 2: Literature Review

We are not living in a world where all roads are radii of a circle and where all, if followed long enough, will therefore draw gradually nearer and finally meet at the centre: rather in a world where every road, after a few miles, forks into two, and each of those into two again, and at each fork you must make a decision.

C.S. Lewis, *The Great Divorce*

In the course of everyday life, individuals are often presented with a multitude of choices that they have to make. At the time of choosing, outcomes for those choices may be certain or uncertain, immediate, or temporally delayed. Economic models such as Rational Choice Theory (RCT) casts individuals as rational agents whose goal is to maximize gains and minimize losses when faced with choices, and one characteristic of such models is the assumption of transitivity inherent to choices made by the agent. An agent that prefers mangoes to grapefruit, and grapefruit to pomegranates, should also show a preference for mangoes to pomegranates if their preferences are internally consistent. Although preferences are intuitive to think about, it is often handier to assign numbers to options presented to an agent, and have the agent select the option with the highest number. These numbers are known as *subjective values* as they are subjective to the agent. While subjective values are unobservable, there are methods to estimate them and it is often thought that the goal of an agent is to make decisions that maximize their long-run rewards. Research on value-based decision-making suggests that neural correlates of reward value, reward-based learning, and selection of actions can be observed across a wide distribution of brain regions.

2.1 Representation of Value

In the previous chapter, value has been defined as a scalar estimate or predictor of reward. The value of choosing an option thus involves an estimate of the reward obtained from making a decision. Many studies have used neuroimaging techniques like functional Magnetic Resonance Imaging (fMRI) to examine brain activity as people participated in a range of tasks involving either the presence or absence of choice. In a study investigating whether subjective values encoded in the brain correspond to individual preferences, participants were placed in an fMRI scanner and shown images of paintings and faces which they were asked to assign

pleasantness ratings to. Following that, they made pair-wise comparisons for the same set of images outside the scanner which allowed for their preferences to be established. It was observed that activity in parts of the ventromedial prefrontal cortex (vmPFC) and basal ganglia accurately reflected subjective values corresponding to individual preferences (Lebreton, Jorge, Michel, Thirion, & Pessiglione, 2009). Corroborating this finding, Levy, Lazzaro, Rutledge, and Glimcher (2011) found that despite the lack of choice while participants passively viewed consumer goods in the scanner, the magnitude of neural activity in the striatum and medial prefrontal cortex (mPFC) in response to the individual goods shown were predictive of subsequent choices, suggesting that the underlying mechanisms encoding subjective values were similar independent of whether an action was required.

The brain regions that encode subjective values vary to some extent with task measures as well as the types of goods presented. Activity in the medial orbitofrontal cortex (mOFC) appears to encode the value of appetitive reward as evident by a correlation with hungry participants' willingness-to-pay for different food items (Plassmann, O'Doherty, & Rangel, 2007). Single cell recordings of individual neurons in the mOFC of monkeys showed that these neurons were encoding both the offer and chosen value of various juice types (Padoa-Schioppa, 2007), and both humans and monkeys with OFC lesions displayed aberrant choice behaviour with failures to update appropriate values to stimuli following reversals in stimulus-reward associations (Fellows & Farah, 2007; Izquierdo, Suda, & Murray, 2004). From an ecological standpoint, many real world choices involve fairly complicated combinations of rewards and their associated subjective values, and one solution for comparison between different reward types or attributes is the idea of a common currency. In a demonstration of this, Levy and Glimcher (2011) had food- and water-deprived participants make risky choices for money, food, and water both inside and outside the scanner. They found that subjective values of food items and monetary rewards had distinct representations in the hypothalamic areas and posterior cingulate cortex respectively. Perhaps more importantly for choice, they also found that neural activations in sub-regions of ventromedial prefrontal cortex which represented subjective values of both food and monetary rewards lie on a common scale that could account for observed behavioural choices.

The activations of multiple brain regions in response to the encoding of subjective values need not necessarily mean that they utilize subjective values in a similar fashion. Camille, Tsuchida, and Fellows (2011) tested patients with damage to either the orbitofrontal cortex or dorsal anterior cingulate cortex on a learning task. The task comprised of value being attached to either 2 different-coloured piles of cards or 2 distinct hand actions. Reward contingencies were altered if patients chose the good deck for a certain amount of times, promoting reversal learning to occur. This paradigm demonstrated dissociation between the learning and acquisition of action- versus stimulus-value within the patient groups. Orbitofrontal cortex was implicated in value-based choices between stimuli that did not involve any motor action (i.e. no supination or pronation of the wrist, while dorsal anterior cingulate cortex was implicated only in choices involving motor action. For both groups of patients, feedback was absolutely crucial to the process of value encoding as observed in a trial-by-trial examination of performance. Integrated as a whole, any compromise to ventromedial frontal regions, or dorsal anterior cingulate cortex results in a disruption of behaviour reflective of utility maximization, and illustrates a chain of causation between the representations of subjective values in the listed regions and decision-making.

Up to this point, the computation and representation of subjective values have been shown to elicit activity in several brain regions. One way to systematically interpret these activity would be to group them according to function, and studies have found that those regions are separable into two functionally different groups. The first is involved in the computation of subjective values with areas implicated in this role being anterior ventral striatum, posterior cingulate cortex, and ventromedial prefrontal cortex. The second pertains to risk salience, and regions responding to this property include the insula, striatum, and dorsomedial prefrontal cortex (Bartra, McGuire, & Kable, 2013). In line with the encoding of subjective values, an intriguing question that might arise is whether the formation of subjective values is dependent on experience and retrieved from memory when needed, or dynamically computed. When rules learnt through experience result in expectations that change the subjective value of an item, conflicts may arise. For example, there is often a learnt association between the cost and quality of wine (i.e. more expensive wines tend to be higher in quality). Such associations have been found to interact with subjective pleasantness ratings, modulating activity in medial orbitofrontal cortex in a manner that was positively

correlated with pleasantness even though participants were tested on the exact same wine that bore different price tags (Plassmann, O'Doherty, Shiv, & Rangel, 2008). When students were shown objects like computer peripherals or candy and asked if any of them would buy one of those at a price randomly drawn from the digits of their social security number, Ariely, Loewenstein, and Prelec (2003) discovered that anchoring prices with any reference point – even an arbitrary one like a social security number – led to a deficit in value retrieval for common products. These studies suggest that in the absence of a readily-available reference point, people tend to fall back on values formed through experience and this could lead to sub-optimal decisions as bottlenecks in memory retrieval could introduce noise in decision-making processes (Giguère & Love, 2013).

Subjective values can be influenced by several factors such as risk and uncertainty and temporal discounting. There are two distinct categories of choices within the risk domain – certain, or uncertain. Uncertain choices are termed “risky” when the probability distributions of outcomes are known, and “ambiguous” if the reward probability distributions are not completely known (Christopoulos, Tobler, Bossaerts, Dolan, & Schultz, 2009). As such, levels of risk presented in a choice tend to be linked with the variance of outcomes (O'Neill & Schultz, 2013; Weber, Shafir, & Blais, 2004). That being said, operational definitions of ‘risk’ might sometimes differ across the literature and interpreted as the chance of obtaining an aversive outcome rather than the variability of potential rewards. If presented with a choice between £5 for sure and a 50% chance of obtaining either £0 and £10, someone who is risk-seeking would choose the gamble more often while someone who is risk-averse would prefer the £5. Notice that the expected value is the same for both options, and therefore someone who is risk-neutral would be indifferent to either option. Using a paradigm where participants were trained to associate cues with probabilistic food rewards, Tricomi and Lempert (2015) discovered that subregions of striatum were involved in reward processing with ventral striatum tracking value and dorsolateral and dorsomedial striatum tracking trialwise probabilities of receiving a reward. Under risk, the manner that options are presented also leads to decision biases that violate economic rationality, and the finding that this “framing effect” was associated with activity in the amygdala suggests an important role for emotional processes in the modulation of decision biases (De Martino, Kumaran, Seymour, & Dolan, 2006).

Relatedly, other studies have also found that the information content of negative and positive events is often skewed by affective biases as frequently observed in healthy and clinical populations, and normalising such biases could lead to more optimal decision-making (Pulcu & Browning, n.d.; Sharot & Garrett, 2016).

The magnitude of subjective values is also susceptible to temporal delays between the time of presentation and its receipt. For example, although a fixed monetary amount like £50 is objectively larger than £10, people might opt to receive £10 if it were offered immediately instead of £50 offered after 2 years. An individual discount function is able to capture this decrease in subjective value of a reward with respect to the increase in receipt delay, and this can be used in conjunction with neuroimaging to find brain regions that are positively correlated with objective value and negatively correlated with receipt delay. That is what Kable and Glimcher (2007) did and they found that medial prefrontal cortex, ventral striatum and posterior cingulate cortex displayed trade-offs between objective value and reward delay that mapped well onto the discount function, suggesting that subjective values of intertemporal choices were represented in those regions. One possible factor that could mediate this trade-off is the psychological construct of impulsivity, characterized by decreased levels of neural activity in nucleus accumbens to the magnitude of future rewards as well as greater deactivations in medial prefrontal cortex, dorsolateral prefrontal cortex, and posterior parietal cortex as delays of future rewards increased (Ballard & Knutson, 2009).

2.2 Valuation

Through a wealth of animal and human literature, there are thought to be several valuation systems in play such as the Pavlovian, habitual, and goal-directed systems (Daw & O'Doherty, 2014; Rangel, Camerer, & Montague, 2008). While it is likely that there are overlapping neural substrates between these systems (Doll, Simon, & Daw, 2012), they are nonetheless useful as frameworks within which valuation processes are carried out.

Some of the general ideas behind these systems can be traced back to scientists such as Edward Thorndike, Ivan Pavlov, Burrhus Frederic Skinner, and Edward Tolman in the late 19th and early 20th centuries. The “Law of Effect” proposed

by Thorndike states that any behaviour followed by pleasant consequences is likely to be repeated and any behaviour followed by unpleasant consequences is likely to cease, suggesting a learnt association between an action and outcome in the presence of reward. Evidence for this came from experiments where he placed hungry animals – often cats – in a series of puzzle boxes (Fig. 2.1) that they had to escape from in order to attain food. As the animals performed a repertoire of instinctive behaviours such as clawing, biting, or squeezing, they eventually chanced upon a specific action that released them from the box. Successive repetitions of this action and outcome pairing increased the strength of the association between them and decreased the time required by the animal to escape the box. Actions that were previously successful towards escape were repeated more often when animals were subsequently placed in a new box, and actions that were previously unsuccessful were reduced (Thorndike, 1898). This form of trial-and-error learning paved the way for Skinner’s work on Operant Conditioning which introduced *reinforcers* and *punishers*, referring to things in the environment that respectively increased or decreased the likelihood of a behaviour being repeated (Skinner, 1938).

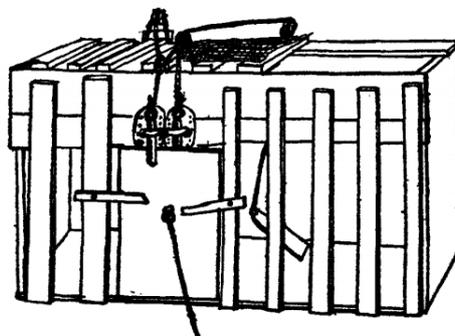


Fig. 2.1. An example of Thorndike’s puzzle boxes. These boxes were approximately 20 inches long, 15 inches wide, and 12 inches high. They often featured levers, latches, and strings attached to pulleys that were cleverly combined in a variety of ways so as to elicit different patterns of associations. Adapted from Thorndike (1898).

In contrast to operant conditioning, work by Pavlov (1927) on classical conditioning asserts that neutral stimuli can acquire or lose value based on repeated associations with a reinforcer that naturally produces a behaviour, even in the absence of action on the part of an animal. He demonstrated this in a classic experiment where the sound of a bell (“conditioned stimulus) gained a positive value and induced

salivation in a dog (“unconditioned response”) after being repeatedly followed by food delivery (“unconditioned stimulus”) despite having elicited no salivation response prior to its association with food. The Pavlovian system is related to this idea that values are assigned to some behaviours that naturally occur in response to environmental stimuli, and these behaviours are sometimes thought to be innate leading to conflicts with instrumental learning on some occasions. For example, Hershberger (1986) placed food-deprived chickens on a straight runway with a chicken feeder. The chicken feeder was set up to retreat from the chickens at twice the speed at which they ran at it, and approach the chickens at twice the speed at which they ran away from it. Thus, instead of approaching the feeder to get a reward as would normally be the case, the chickens had to run away from the feeder in order to get fed which was a behaviour they were unable to acquire. Similarly in humans, participants were found to be better at taking an action in anticipation of reward and withholding an action in anticipation of punishment as opposed to withholding an action in anticipation of reward and taking an action in anticipation of punishment (Guitart-Masip et al., 2012).

2.3 Learning

In order to make good decisions, people need to have reliable estimates of the values associated with an option or action given the state they are currently in. Unlike controlled environments such as a laboratory where values of stimuli are often explicitly available to participants, estimating values of items in the real world typically involve some form of learning. A central idea in learning is that of *surprise*, which often occurs as a consequence of deviations between *expectations* and observed outcomes. Surprise can be useful as it provides information that expectations need to be updated in order for future *predictions* to be more accurate given a non-stochastic environment, and I will briefly return to this as part of the work done in **Chapter 6**. However, rewards are often delayed in the future and progress towards obtaining a desired reward can involve multiple actions, leading to the problem of *credit assignment* where it becomes particularly important to appropriately assign value to actions that are crucial for reward.

2.3.1 Habitual or Model-free Systems

Habitual or model-free systems rely on past experience, using trial-and-error learning to estimate values of options across a wider range of actions compared to the Pavlovian system. As a consequence, learning the values of actions given the current state an agent is in takes place across a longer timescale. Over the course of experience, some goal-directed behaviours can also become habitual. One way of assessing whether behaviour is habitual or goal-directed involves reinforcer devaluation (Dickinson, 1985). In one such paradigm (Dickinson, Nicholas, & Adams, 1983), rats are trained to press a lever to obtain sucrose, after which sucrose is devalued by being paired with lithium chloride injections. When placed back into the operant chamber in the absence of reinforcers, rats that were trained on a ratio schedule made fewer lever presses, demonstrating an association between instrumental performance and value of the reinforcer that has been interpreted as goal-directed behaviour. Importantly, rats that were more extensively trained on the interval schedule did not appear to be responding less on the lever press, suggesting that they were responding habitually and not based on knowledge about the consequences of their actions. A similar pattern of behaviour has also been observed in humans under conditions of reward devaluation with activation in dorsolateral posterior putamen thought to be involved in the habitual control of instrumental actions (Tricomi, Balleine, & O'Doherty, 2009; Valentin, Dickinson, & O'Doherty, 2007).

2.3.2 Goal-directed or Model-based Systems

Unlike the habitual or model-free systems which contains cached estimates of long-run rewards for actions at each state, a goal-directed or model-based system contains representations of the environment and incorporates knowledge of transitions between states, which is the likelihood of transitioning to a certain state given the selected action (Dayan & Berridge, 2014). Such systematic organization of information has more recently been termed a cognitive map and is thought to support flexible behaviour by incorporating relationships between elements in the environment (Behrens et al., 2018; O'Keefe & Nadel, 19978). These relationships can be learnt even in the absence of reinforcers in what is known as *latent learning*. In a classic demonstration of this, rats that had been afforded an opportunity to explore a maze in

the absence of reward were quicker to navigate towards a subsequently rewarded goal state in the same maze compared to rats that had not been previously exposed to the maze (Tolman, 1948). In the earlier example where rats were trained to press a lever for sucrose, a reduction in lever presses after the sucrose was devalued has been interpreted as goal-directed behaviour because such a response requires knowledge about action-outcome and outcome-value contingencies (Rangel et al., 2008). Puzzlingly, the neural substrates associated with model-based or goal-directed influences appear to overlap with regions also involved in model-free learning (Daw, Gershman, Seymour, Dayan, & Dolan, 2011; van der Meer, Johnson, Schmitzer-Torbert, & Redish, 2010), which could be potentially related to how a goal-directed or model-based system is trained (Wang et al., 2018). One major drawback of the model-based system is that it is computationally expensive as it draws on the model of the environment as it looks ahead at each step of the decision tree.

2.3.3 Successor Representation

A recently resurrected idea that forms a neat compromise between flexibility and computational efficiency is the *successor representation* (Dayan, 1993; Momennejad et al., 2017). The key idea behind this is that the successor representation contains a summary of state transition statistics separate from reward statistics, meaning that changes in rewards for a given state do not necessitate a full recalculation of the value function before state values are updated (Gershman, 2018).

2.3.4 Dopamine: Learning and Decision-Making

A neurotransmitter that has often been implicated in the human and animal decision-making literature is dopamine. This, together with its receptor types (D1, D2, D3), make up the dopaminergic neurotransmitter system and exerts influence on choice behaviour and other processes through neuronal activity along the mesolimbic, mesocortical and nigrostriatal pathways (Everitt & Robbins, 2005). Schultz, Dayan, and Montague, (1997) devised an influential model of reinforcement learning based on work done with monkeys. Using single-cell recording, they showed that dopaminergic neurons in the midbrain responded to incongruences between expected and received rewards such that a reward that was not predicted lead to a spike in

neuronal activity, a reward that was predicted elicited no change in activity, and an absence of an expected reward lead to a depression in activity. The discrepancies between expected and actual outcomes are known as reward prediction errors (RPE). A simple reinforcement learning model can be illustrated by the equation:

$$V_t = V_{t-1} + \alpha (R_t - V_{t-1})$$

where V represents the value of an option, R represents the reward received, t refers to the current time point, and α represents an individual learning rate or extent to which previous values are updated by new information. The presence of reward prediction error signals in dopaminergic areas has also been observed in humans. Presenting deep-brain-surgery PD patients with separate decks of different-coloured cards, Zaghoul et al. (2009) informed them that one deck contained a greater chance of reward than the other and recorded neural activity during outcome presentation as patients selected cards from each deck. They discovered that activity in the dopaminergic midbrain resembled reward prediction errors, and suggested that damage to the dopaminergic system leads to abnormalities in the encoding of reward prediction errors resulting in PD patients displaying impaired decision-making as a function of incorrectly updated reward values.

As previously mentioned, the model-free system learns about the course of action that best maximises reward by using reward prediction errors to update values at each new point in time. This means that any action that leads to a successful reward may be reinforced, regardless of whether that action is actually instrumental to the task at hand. In contrast, model-based learning involves an internal task representation that is dependent on past experience, and optimal actions can be arrived at by searching through these representations. Since the model-free system is updated through reward prediction errors, one might expect that altering dopamine levels would disrupt the system and bias control towards the model-based system. Wunderlich, Smittenaar, and Dolan (2012) tested this hypothesis by administering L-DOPA to participants who played a two-stage Markov decision task and found that boosting dopamine biased behavioural control towards the model-based system.

A recent alternative to explain the role of dopamine comes from the concept of active inference (Friston et al., 2012). Active inference makes use of prior beliefs about

the world to guide behaviour, and the idea behind it is not unlike the model-based system. In fact, the very nature of active inference makes it model-based. Within active inference, the role of dopamine is thought to attenuate the precision or uncertainty regarding various Bayes-optimal representations of the world, and modelling changes to the postsynaptic gain of the same neurons that dopamine acts on appear to lead to similar deficits and pathologies observed in people with compromised dopaminergic neurotransmitter systems.

Several approaches can be adopted to understand the involvement of dopamine in reward (Berridge, 2007). First, one could examine what reward functions are lost when dopaminergic systems are compromised, for example in PD patients (Steeves et al., 2009; Voon et al., 2009; Weintraub et al., 2006; Zaghoul et al., 2009). Secondly, one could examine what reward functions are elevated by an increase in dopamine (Pessiglione, Seymour, Flandin, Dolan, & Frith, 2006; Sharot, Guitart-Masip, Korn, Chowdhury, & Dolan, 2012; Sharot, Shiner, Brown, Fan, & Dolan, 2009). Third, one could examine the coding of reward by dopaminergic neurons in the presence of stimulus values or outcome values (Rangel & Clithero, 2012). Using a combination of the first two approaches, Rutledge et al. (2009) tested PD patients with a foraging task where the distribution of reward probabilities and action values had to be learnt in the presence of feedback. They observed that the rates of learning were negatively correlated with the progression of the disease, meaning that patients that were more severely affected by PD displayed lower learning rates than patients that were less severely affected. This suggests that dopamine plays an active role in reinforcement learning, consistent with the idea that reward prediction errors are encoded by dopaminergic neurons. Apart from this, Rutledge et al. (2009) also found that PD patients demonstrated a perseveration of choice that did not depend on the history of rewards. One possible explanation for this could be that deficits in the updating of value by dopaminergic neurons interfere with the consolidation of past rewards, or perhaps dopamine is involved in the rate of reward decay with more weight placed on recent outcomes. Both learning rates and choice perseveration were improved following pharmacological administration of dopamine, reinforcing its role in the encoding and updating of value. Such impairments to dopaminergic neurons could also lead to abnormal activity in areas innervated by dopaminergic pathways such as the nucleus accumbens, leading to an increase in suboptimal choices in the presence

of risky options (Samanez-Larkin, Kuhnen, Yoo, & Knutson, 2010).

Administering dopamine to PD patients to treat their symptoms does not come without side effects, and observations that some PD patients on dopaminergic medication started to develop pathological gambling led to the investigation of dopamine on the distinct types of dopamine receptors. One potential explanation for this is that increased dopamine released through the use of dopamine agonists preferentially bind to D3-receptors littered throughout the limbic system – an area often associated with reward and hedonia (Dodd et al., 2005). Dopamine has often been implicated in hedonistic tendencies and rewards in part due to the propensity for addictive drugs and gambling behaviour to exert their influence on dopaminergic midbrain circuits, leading to the stipulation of dopamine as a ‘pleasure neurotransmitter’ (Clark, 2010; Everitt & Robbins, 2005). One major criticism of dopamine’s involvement in pleasure and hedonism – specifically the ‘high’ or pleasantness associated with rewarding items like drugs or food – is the fact many studies do not distinguish between ‘wanting’ and ‘liking’ (Berridge & Robinson, 2016). Excessive release of dopamine in the ventral striatum has been shown to be highly correlated with subjective ratings of ‘wanting’ and not ‘liking’ even for drugs like L-DOPA that do not induce euphoria (Evans et al., 2006). In response to this, Berridge (2007) proposed the incentive salience hypothesis, postulating that the notion of reward consists of attributes such as liking, learning, and wanting. The involvement of dopamine is restricted to the ‘wanting’ attribute, mediating the motivational aspect of attaining a reward by increasing or decreasing its incentive salience. Furthermore, the hypothesis also adds on to neural representations of learning signals by viewing incentive salience as a value that transforms predictions into the ‘wanting’ associated with reward stimuli. Manipulation of dopamine leads to behavioural changes by attaching ‘wanting’ to a conditioned stimulus predictive of rewards.

Role of Distinct Dopamine Receptors

To better inform us about the roles of dopamine and allow more detailed insight into processes underlying reward-related behavioural changes, it is important to examine the receptor sub-types that mediate effects of dopamine (Rogers, 2011). Glickstein, Desteno, Hof, and Schmauss (2005) used a visual discrimination task involving two-choices to test mice lacking in either D2- or D3-receptors and observed

that the former was impaired in the acquisition of task rules while the latter did not show such impairment. Both groups did not exhibit any compromise in ability to set-shift or attend to stimuli. Applying neuroimaging techniques to patients with Huntington's disease (HD) who displayed deficits in cognitive abilities like planning and attention, Lawrence et al. (1998) observed a robust correlation between performance on cognitive tasks and the levels of dopamine that were binding to striatal D1- and D2-receptors. In addition to this, Norbury, Manohar, Rogers, and Husain (2013) administered a D2- and D3- receptor agonist to healthy controls and observed an increase in sensitivity to information about positive reward probabilities and a decrease in sensitivity to the magnitude of potential losses. Furthermore, they found that the magnitude of the effect was dependent on a trait known as sensation-seeking.

In some instances, dopamine elicits similar behavioural changes regardless of receptor-types while at other times, differential effects of dopamine on reward-related behavioural changes are observed to be dependent on receptor-types. Incorporating the contributions of dopamine to action selection in humans, Frank and Hutchison (2009) described a model comprising of a direct and indirect pathway leading from the striatum up to cortex. The former was thought to facilitate a response appraised in the cortex ('Go signal') through the excitatory responses of D1-receptors while the latter kept a rein on competing responses ('No-Go signal') through inhibitory actions incited by D2-receptors. These results suggest that the action of dopamine is multi-faceted with complex interactions between receptor-types.

2.4 Intrinsic Variability of Neural Signals and Behaviour

Perhaps evident from previous sections, while a large portion of neuroscience studies have been dedicated to understanding neural responses to a task or stimulus, an often overlooked fact is that spontaneous fluctuations of Blood-Oxygen-Level-Dependent (BOLD) activity paint a picture of an active brain at rest despite the absence of external stimuli (Fox & Raichle, 2007). A key question of interest is whether such activity represent something meaningful rather than noise.

Mounting evidence from animal and human research suggest that BOLD activity at rest, known more commonly as resting-state BOLD, can be used to study

the functional organization of the brain by grouping brain regions with similar patterns of spontaneous activation, forming clusters that are spatially distributed but temporally correlated (Belloy et al., 2018; Damoiseaux et al., 2006; Power et al., 2011). These “intrinsic connectivity networks” are often functionally rather than structurally integrated (Yeo et al., 2011) and regional correlations in resting-state BOLD activity have been found to be fairly stable across short timescales, distinguishing them from momentary activations that can be attributed to cognitive processing (Laumann et al., 2017). By training a model to find relations between task-independent BOLD features during rest and subsequent activations in response to a task, Tavor et al. (2016) further demonstrated that differences in functional connectivity at rest was predictive of individual differences in task-evoked brain activity, suggesting that stable patterns within individuals could explain variability in activation patterns across individuals. Interestingly, the strength of intrinsic amygdala-cortical functional connectivity has even been found to be predictive of the size of one’s social network (Bickart, Hollenbeck, Barrett, & Dickerson, 2012). While stability represents a useful framework for understanding spontaneous fluctuations at the network level, nonetheless functional connectivity metrics have also been found to vary on some timescales under conditions such as sleep (Horowitz et al., 2009), learning (Bassett et al., 2011), and task demands (Fornito, Harrison, Zalesky, & Simons, 2012), leading to increased interest in dynamic functional connectivity (Hutchison et al., 2013). Such temporal fluctuations in functional connectivity are thought to reflect dynamic changes in brain organization and shifts of brain states (Liégeois, Laumann, Snyder, Zhou, & Yeo, 2017).

Many studies of spontaneous fluctuations have investigated regional relationships of infra-slow brain activity with respect to function, leaving a lacuna for the functional role of intrinsic activity *within* specific brain regions. At the neuronal level, intrinsic activity resembles distinct rather than stochastic patterns of firing (Ikegaya et al., 2004; Mao, Hamzei-Sichani, Aronov, Froemke, & Yuste, 2001; Tsodyks, Kenet, Grinvald, & Arieli, 1999), not unlike those elicited by electrophysiological stimulation or task events. In a fMRI study where human subjects pressed a button each time a stimulus changed colour on a monitor, Fox, Snyder, Vincent, and Raichle (2007) observed that intrinsic BOLD activity within the left somatomotor cortex was linked with changes of button press force from one trial to the next, establishing an association

between intrinsic brain activity and motor output. At the level of sensory perception, pre-stimulus fluctuations in medial thalamic BOLD activity have been found to influence conscious perception of a laser-induced somatosensory stimulus, and similar baseline fluctuations in the anterior cingulate cortex have been related to subsequent ratings of pain intensity when the stimulus was increased to nociceptive levels (Boly et al., 2007). In a similar vein of study, Hesselmann, Kell, & Kleinschmidt, 2008 presented subjects with visual stimuli featuring either coherent or random motion for extremely brief periods of time and asked them to report which of the two they had perceived. They found that intrinsic fluctuations in the right occipito-temporal cortex during baseline biased coherent percepts, suggesting a role for intrinsic fluctuations in perceptual performance that is also supported by studies conducted on other areas of the visual cortex (Schölvinck, Friston, & Rees, 2012). Within the context of perceptual decision-making, prescient activity has been thought to bias action selection through the use of outcome history (Hwang, Dahlen, Mukundan, & Komiyama, 2017), reflect the precision of prediction errors within a predictive coding framework (Hesselmann, Sadaghiani, Friston, & Kleinschmidt, 2010), or shift the initial starting point towards a decision boundary in an evidence accumulation model (Summerfield & de Lange, 2014).

In the domain of value-based decision-making, Maoz et al., 2013 recorded spiking activity from neurons in the striatum and dorsolateral prefrontal cortex of monkeys as they deliberated between receiving a small reward immediately and a larger reward later in time. Even before the options were presented, spiking activities of some neurons in the ventral striatum and dorsolateral prefrontal cortex were predictive of choice based on spatial location or reward magnitude respectively, constituting a form of pre-deliberation bias. A third group of neurons, found in the dorsolateral prefrontal cortex and caudate nucleus by Maoz et al., 2013 and in orbitofrontal cortex by Padoa-Schioppa (2013), had firing patterns that were modulated by choice. Fluctuations in pre-stimulus activity within this group of neurons were also predictive of forthcoming choice, especially when values of the options were close. Missing from the picture thus far is an important modulator of value – *risk*. In the real world, many decisions involve options with probabilistic outcomes such as purchasing a lottery ticket or placing a bet on a favourite sports team. Risk preferences have been associated with the dopaminergic system, and boosting dopamine levels

pharmacologically have been shown to increase the propensity to take risks (Rigoli et al., 2016), a phenomena also observed in some Parkinson's Disease patients treated with dopamine agonists (Driver-Dunckley, Samanta, & Stacy, 2003; Voon et al., 2007) . More recently, a large-scale study using a smartphone app found that older adults were less willing to take risks in order to win more points, a result consistent with age-related decline in dopamine levels (Rutledge, Skandali, Dayan, & Dolan, 2015).

To address the gap between intrinsic fluctuations of brain activity and risk preferences, I leverage on the observed relationship between dopamine and risk and hypothesized that spontaneous BOLD activity in the dopaminergic midbrain would exert an influence on risk preferences with higher activity leading to increased risk-taking. In **Chapter 3**, I describe a novel real-time fMRI method used to test this hypothesis. Briefly, resting BOLD activity was detected in real time and choice trials were presented to participants depending on whether this activity was lower or higher than the average. Full details and results of the experiment are presented in **Chapter 4**.

References

- Ariely, D., Loewenstein, G., & Prelec, D. (2003). "Coherent arbitrariness": Stable demand curves without stable preferences. *The Quarterly Journal of Economics*, *118*(1), 73–106. <https://doi.org/10.1162/00335530360535153>
- Ballard, K., & Knutson, B. (2009). Dissociable neural representations of future reward magnitude and delay during temporal discounting. *NeuroImage*, *45*(1), 143–150. <https://doi.org/10.1016/j.neuroimage.2008.11.004>
- Bartra, O., McGuire, J. T., & Kable, J. W. (2013). The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *NeuroImage*, *76*, 412–427. <https://doi.org/10.1016/j.neuroimage.2013.02.063>
- Bassett, D. S., Wymbs, N. F., Porter, M. A., Mucha, P. J., Carlson, J. M., & Grafton, S. T. (2011). Dynamic reconfiguration of human brain networks during learning. *PNAS*, *108*(18), 7641–7646. <https://doi.org/10.1073/pnas.1018985108>
- Behrens, T. E. J., Muller, T. H., Whittington, J. C. R., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron*, *100*(2), 490–509. <https://doi.org/10.1016/j.neuron.2018.10.002>
- Belloy, M. E., Naeyaert, M., Abbas, A., Shah, D., Vanreusel, V., van Audekerke, J., Verhoye, M. (2018). Dynamic resting state fMRI analysis in mice reveals a set of Quasi-Periodic Patterns and illustrates their relationship with the global signal. *NeuroImage*, *180*, 463–484. <https://doi.org/10.1016/j.neuroimage.2018.01.075>
- Berridge, K. C. (2007). The debate over dopamine's role in reward: The case for incentive salience. *Psychopharmacology*, *191*(3), 391–431. <https://doi.org/10.1007/s00213-006-0578-x>
- Berridge, K. C., & Robinson, T. E. (2016). Liking, Wanting and the Incentive-Sensitization Theory of Addiction. *The American Psychologist*, *71*(8), 670–679. <https://doi.org/10.1037/amp0000059>
- Bickart, K. C., Hollenbeck, M. C., Barrett, L. F., & Dickerson, B. C. (2012). Intrinsic amygdala–cortical functional connectivity predicts social network size in humans. *Journal of Neuroscience*, *32*(42), 14729–14741. <https://doi.org/10.1523/JNEUROSCI.1599-12.2012>
- Boly, M., Balteau, E., Schnakers, C., Degueldre, C., Moonen, G., Luxen, A., Laureys, S. (2007). Baseline brain activity fluctuations predict somatosensory perception in humans. *PNAS*, *104*(29), 12187–12192. <https://doi.org/10.1073/pnas.0611404104>

- Camille, N., Tsuchida, A., & Fellows, L. K. (2011). Double dissociation of stimulus-value and action-value learning in humans with orbitofrontal or anterior cingulate cortex damage. *Journal of Neuroscience*, *31*(42), 15048–15052.
<https://doi.org/10.1523/JNEUROSCI.3164-11.2011>
- Christopoulos, G. I., Tobler, P. N., Bossaerts, P., Dolan, R. J., & Schultz, W. (2009). Neural correlates of value, risk, and risk aversion contributing to decision making under risk. *Journal of Neuroscience*, *29*(40), 12574–12583.
<https://doi.org/10.1523/JNEUROSCI.2614-09.2009>
- Clark, L. (2010). Decision-making during gambling: An integration of cognitive and psychobiological approaches. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *365*(1538), 319–330. <https://doi.org/10.1098/rstb.2009.0147>
- Damoiseaux, J. S., Rombouts, S. a. R. B., Barkhof, F., Scheltens, P., Stam, C. J., Smith, S. M., & Beckmann, C. F. (2006). Consistent resting-state networks across healthy subjects. *PNAS*, *103*(37), 13848–13853. <https://doi.org/10.1073/pnas.0601417103>
- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, *69*(6), 1204–1215. <https://doi.org/10.1016/j.neuron.2011.02.027>
- Daw, N. D., & O'Doherty, J. P. (2014). Chapter 21—Multiple Systems for Value Learning. In P. W. Glimcher & E. Fehr (Eds.), *Neuroeconomics (Second Edition)* (pp. 393–410). <https://doi.org/10.1016/B978-0-12-416008-8.00021-8>
- Dayan, P. (1993). Improving generalization for temporal difference learning: The successor representation. *Neural Computation*, *5*(4), 613–624.
<https://doi.org/10.1162/neco.1993.5.4.613>
- Dayan, P., & Berridge, K. C. (2014). Model-based and model-free Pavlovian reward learning: Revaluation, revision, and revelation. *Cognitive, Affective & Behavioral Neuroscience*, *14*(2), 473–492. <https://doi.org/10.3758/s13415-014-0277-8>
- De Martino, B., Kumaran, D., Seymour, B., & Dolan, R. J. (2006). Frames, biases, and rational decision-making in the human brain. *Science*, *313*(5787), 684–687.
<https://doi.org/10.1126/science.1128356>
- Dickinson, A. (1985). Actions and habits: The development of behavioural autonomy. *Phil. Trans. R. Soc. Lond. B*, *308*(1135), 67–78. <https://doi.org/10.1098/rstb.1985.0010>
- Dickinson, Anthony, Nicholas, D. J., & Adams, C. D. (1983). The effect of the instrumental training contingency on susceptibility to reinforcer devaluation. *The Quarterly Journal of Experimental Psychology Section B*, *35*(1), 35–51.
<https://doi.org/10.1080/14640748308400912>

- Dodd, M. L., Klos, K. J., Bower, J. H., Geda, Y. E., Josephs, K. A., & Ahlskog, J. E. (2005). Pathological gambling caused by drugs used to treat Parkinson disease. *Archives of Neurology*, *62*(9), 1377–1381. <https://doi.org/10.1001/archneur.62.9.noc50009>
- Doll, B. B., Simon, D. A., & Daw, N. D. (2012). The ubiquity of model-based reinforcement learning. *Current Opinion in Neurobiology*, *22*(6), 1075–1081. <https://doi.org/10.1016/j.conb.2012.08.003>
- Driver-Dunckley, E., Samanta, J., & Stacy, M. (2003). Pathological gambling associated with dopamine agonist therapy in Parkinson's disease. *Neurology*, *61*(3), 422–423.
- Everitt, B. J., & Robbins, T. W. (2005). Neural systems of reinforcement for drug addiction: From actions to habits to compulsion. *Nature Neuroscience*, *8*(11), 1481–1489. <https://doi.org/10.1038/nn1579>
- Fellows, L. K., & Farah, M. J. (2007). The role of ventromedial prefrontal cortex in decision making: Judgment under uncertainty or judgment per se? *Cerebral Cortex*, *17*(11), 2669–2674. <https://doi.org/10.1093/cercor/bhl176>
- Fornito, A., Harrison, B. J., Zalesky, A., & Simons, J. S. (2012). Competitive and cooperative dynamics of large-scale brain functional networks supporting recollection. *PNAS*, *109*(31), 12788–12793. <https://doi.org/10.1073/pnas.1204185109>
- Fox, M. D., & Raichle, M. E. (2007). Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Reviews Neuroscience*, *8*(9), 700–711. <https://doi.org/10.1038/nrn2201>
- Fox, M. D., Snyder, A. Z., Vincent, J. L., & Raichle, M. E. (2007). Intrinsic Fluctuations within Cortical Systems Account for Intertrial Variability in Human Behavior. *Neuron*, *56*(1), 171–184. <https://doi.org/10.1016/j.neuron.2007.08.023>
- Friston, K. J., Shiner, T., FitzGerald, T., Galea, J. M., Adams, R., Brown, H., Bestmann, S. (2012). Dopamine, Affordance and Active Inference. *PLoS Computational Biology*, *8*(1), e1002327. <https://doi.org/10.1371/journal.pcbi.1002327>
- Gershman, S. J. (2018). The successor representation: Its computational logic and neural substrates. *Journal of Neuroscience*, 0151–18. <https://doi.org/10.1523/JNEUROSCI.0151-18.2018>
- Giguère, G., & Love, B. C. (2013). Limits in decision making arise from limits in memory retrieval. *PNAS*, *110*(19), 7613–7618. <https://doi.org/10.1073/pnas.1219674110>
- Glickstein, S. B., Desteno, D. A., Hof, P. R., & Schmauss, C. (2005). Mice lacking dopamine D2 and D3 receptors exhibit differential activation of prefrontal cortical neurons during tasks requiring attention. *Cerebral Cortex*, *15*(7), 1016–1024. <https://doi.org/10.1093/cercor/bhh202>

- Guitart-Masip, M., Huys, Q. J. M., Fuentemilla, L., Dayan, P., Duzel, E., & Dolan, R. J. (2012). Go and no-go learning in reward and punishment: Interactions between affect and effect. *Neuroimage*, *62–334*(1), 154–166.
<https://doi.org/10.1016/j.neuroimage.2012.04.024>
- Hershberger, W. A. (1986). An approach through the looking-glass. *Animal Learning & Behavior*, *14*(4), 443–451. <https://doi.org/10.3758/BF03200092>
- Hesselmann, G., Kell, C. A., & Kleinschmidt, A. (2008). Ongoing activity fluctuations in hMT+ bias the perception of coherent visual motion. *Journal of Neuroscience*, *28*(53), 14481–14485. <https://doi.org/10.1523/JNEUROSCI.4398-08.2008>
- Hesselmann, G., Sadaghiani, S., Friston, K. J., & Kleinschmidt, A. (2010). Predictive coding or evidence accumulation? False inference and neuronal fluctuations. *PLOS ONE*, *5*(3), e9926. <https://doi.org/10.1371/journal.pone.0009926>
- Horowitz, S. G., Braun, A. R., Carr, W. S., Picchioni, D., Balkin, T. J., Fukunaga, M., & Duyn, J. H. (2009). Decoupling of the brain's default mode network during deep sleep. *PNAS*, *106*(27), 11376–11381. <https://doi.org/10.1073/pnas.0901435106>
- Hwang, E. J., Dahlen, J. E., Mukundan, M., & Komiyama, T. (2017). History-based action selection bias in posterior parietal cortex. *Nature Communications*, *8*(1), 1242. <https://doi.org/10.1038/s41467-017-01356-z>
- Ikegaya, Y., Aaron, G., Cossart, R., Aronov, D., Lampl, I., Ferster, D., & Yuste, R. (2004). Synfire chains and cortical songs: Temporal modules of cortical activity. *Science*, *304*(5670), 559–564. <https://doi.org/10.1126/science.1093173>
- Izquierdo, A., Suda, R. K., & Murray, E. A. (2004). Bilateral orbital prefrontal cortex lesions in rhesus monkeys disrupt choices guided by both reward value and reward contingency. *Journal of Neuroscience*, *24*(34), 7540–7548. <https://doi.org/10.1523/JNEUROSCI.1921-04.2004>
- Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience*, *10*(12), 1625–1633. <https://doi.org/10.1038/nn2007>
- Laumann, T. O., Snyder, A. Z., Mitra, A., Gordon, E. M., Gratton, C., Adeyemo, B., Petersen, S. E. (2017). On the stability of BOLD fMRI correlations. *Cerebral Cortex*, *27*(10), 4719–4732. <https://doi.org/10.1093/cercor/bhw265>
- Lawrence, A. D., Weeks, R. A., Brooks, D. J., Andrews, T. C., Watkins, L. H., Harding, A. E., Sahakian, B. J. (1998). The relationship between striatal dopamine receptor binding and cognitive performance in Huntington's disease. *Brain: A Journal of Neurology*, *121* (Pt 7), 1343–1355.

- Lebreton, M., Jorge, S., Michel, V., Thirion, B., & Pessiglione, M. (2009). An automatic valuation system in the human brain: Evidence from functional neuroimaging. *Neuron*, *64*(3), 431–439. <https://doi.org/10.1016/j.neuron.2009.09.040>
- Levy, D. J., & Glimcher, P. W. (2011). Comparing apples and oranges: Using reward-specific and reward-general subjective value representation in the brain. *Journal of Neuroscience*, *31*(41), 14693–14707. <https://doi.org/10.1523/JNEUROSCI.2218-11.2011>
- Levy, I., Lazzaro, S. C., Rutledge, R. B., & Glimcher, P. W. (2011). Choice from non-choice: Predicting consumer preferences from blood oxygenation level-dependent signals obtained during passive viewing. *Journal of Neuroscience*, *31*(1), 118–125. <https://doi.org/10.1523/JNEUROSCI.3214-10.2011>
- Liégeois, R., Laumann, T. O., Snyder, A. Z., Zhou, J., & Yeo, B. T. T. (2017). Interpreting temporal fluctuations in resting-state functional connectivity MRI. *NeuroImage*, *163*, 437–455. <https://doi.org/10.1016/j.neuroimage.2017.09.012>
- Mao, B.-Q., Hamzei-Sichani, F., Aronov, D., Froemke, R. C., & Yuste, R. (2001). Dynamics of spontaneous activity in neocortical slices. *Neuron*, *32*(5), 883–898. [https://doi.org/10.1016/S0896-6273\(01\)00518-9](https://doi.org/10.1016/S0896-6273(01)00518-9)
- Maoz, U., Rutishauser, U., Kim, S., Cai, X., Lee, D., & Koch, C. (2013). Predeliberation activity in prefrontal cortex and striatum and the prediction of subsequent value judgment. *Frontiers in Neuroscience*, *7*. <https://doi.org/10.3389/fnins.2013.00225>
- Momennejad, I., Russek, E. M., Cheong, J. H., Botvinick, M. M., Daw, N. D., & Gershman, S. J. (2017). The successor representation in human reinforcement learning. *Nature Human Behaviour*, *1*(9), 680. <https://doi.org/10.1038/s41562-017-0180-8>
- Norbury, A., Manohar, S., Rogers, R. D., & Husain, M. (2013). Dopamine modulates risk-taking as a function of baseline sensation-seeking trait. *Journal of Neuroscience*, *33*(32), 12982–12986. <https://doi.org/10.1523/JNEUROSCI.5587-12.2013>
- O'Keefe, J., & Nadel, L. (1997). *The hippocampus as a cognitive map*. Oxford: Oxford University Press.
- O'Neill, M., & Schultz, W. (2013). Risk prediction error coding in orbitofrontal neurons. *Journal of Neuroscience*, *33*(40), 15810–15814. <https://doi.org/10.1523/JNEUROSCI.4236-12.2013>
- Padoa-Schioppa, C. (2007). Orbitofrontal cortex and the computation of economic value. *Annals of the New York Academy of Sciences*, *1121*, 232–253. <https://doi.org/10.1196/annals.1401.011>
- Padoa-Schioppa, C. (2013). Neuronal origins of choice variability in economic decisions. *Neuron*, *80*(5), 1322–1336. <https://doi.org/10.1016/j.neuron.2013.09.013>

- Pavlov, I. P. (1927). *Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex*. Oxford, England: Oxford Univ. Press.
- Pessiglione, M., Seymour, B., Flandin, G., Dolan, R. J., & Frith, C. D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*, *442*(7106), 1042–1045. <https://doi.org/10.1038/nature05051>
- Plassmann, H., O’Doherty, J., & Rangel, A. (2007). Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *Journal of Neuroscience*, *27*(37), 9984–9988. <https://doi.org/10.1523/JNEUROSCI.2131-07.2007>
- Plassmann, H., O’Doherty, J., Shiv, B., & Rangel, A. (2008). Marketing actions can modulate neural representations of experienced pleasantness. *PNAS*, *105*(3), 1050–1054. <https://doi.org/10.1073/pnas.0706929105>
- Power, J. D., Cohen, A. L., Nelson, S. M., Wig, G. S., Barnes, K. A., Church, J. A., ... Petersen, S. E. (2011). Functional network organization of the human brain. *Neuron*, *72*(4), 665–678. <https://doi.org/10.1016/j.neuron.2011.09.006>
- Pulcu, E., & Browning, M. (n.d.). Affective bias as a rational response to the statistics of rewards and punishments. *eLife*, *6*. <https://doi.org/10.7554/eLife.27879>
- Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, *9*(7), 545. <https://doi.org/10.1038/nrn2357>
- Rangel, A., & Clithero, J. A. (2012). Value normalization in decision making: Theory and evidence. *Current Opinion in Neurobiology*, *22*(6), 970–981. <https://doi.org/10.1016/j.conb.2012.07.011>
- Rigoli, F., Rutledge, R. B., Chew, B., Ousdal, O. T., Dayan, P., & Dolan, R. J. (2016). Dopamine increases a value-independent gambling propensity. *Neuropsychopharmacology*, *41*(11), 2658–2667. <https://doi.org/10.1038/npp.2016.68>
- Rogers, R. D. (2011). The roles of dopamine and serotonin in decision making: Evidence from pharmacological experiments in humans. *Neuropsychopharmacology*, *36*(1), 114–132. <https://doi.org/10.1038/npp.2010.165>
- Rutledge, R. B., Lazzaro, S. C., Lau, B., Myers, C. E., Gluck, M. A., & Glimcher, P. W. (2009). Dopaminergic drugs modulate learning rates and perseveration in parkinson’s patients in a dynamic foraging task. *Journal of Neuroscience*, *29*(48), 15104–15114. <https://doi.org/10.1523/JNEUROSCI.3524-09.2009>
- Rutledge, R. B., Skandali, N., Dayan, P., & Dolan, R. J. (2015). Dopaminergic modulation of decision making and subjective well-being. *Journal of Neuroscience*, *35*(27), 9811–9822. <https://doi.org/10.1523/JNEUROSCI.0702-15.2015>

- Samanez-Larkin, G. R., Kuhnen, C. M., Yoo, D. J., & Knutson, B. (2010). Variability in nucleus accumbens activity mediates age-related suboptimal financial risk taking. *Journal of Neuroscience*, *30*(4), 1426–1434.
<https://doi.org/10.1523/JNEUROSCI.4902-09.2010>
- Schölvinck, M. L., Friston, K. J., & Rees, G. (2012). The influence of spontaneous activity on stimulus processing in primary visual cortex. *NeuroImage*, *59*(3), 2700–2708.
<https://doi.org/10.1016/j.neuroimage.2011.10.066>
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A Neural Substrate of Prediction and Reward. *Science*, *275*(5306), 1593–1599.
<https://doi.org/10.1126/science.275.5306.1593>
- Sharot, T., & Garrett, N. (2016). Forming beliefs: Why valence matters. *Trends in Cognitive Sciences*, *20*(1), 25–33. <https://doi.org/10.1016/j.tics.2015.11.002>
- Sharot, T., Guitart-Masip, M., Korn, C. W., Chowdhury, R., & Dolan, R. J. (2012). How dopamine enhances an optimism bias in humans. *Current Biology*, *22*(16), 1477–1481. <https://doi.org/10.1016/j.cub.2012.05.053>
- Sharot, T., Shiner, T., Brown, A. C., Fan, J., & Dolan, R. J. (2009). Dopamine enhances expectation of pleasure in humans. *Current Biology*, *19*(24), 2077–2080.
<https://doi.org/10.1016/j.cub.2009.10.025>
- Skinner, B. F. (1938). *The behavior of organisms: An experimental analysis*. Oxford, England: Appleton-Century.
- Steeves, T. D. L., Miyasaki, J., Zurowski, M., Lang, A. E., Pellecchia, G., Van Eimeren, T., Strafella, A. P. (2009). Increased striatal dopamine release in Parkinsonian patients with pathological gambling: A [¹¹C] raclopride PET study. *Brain*, *132*(Pt 5), 1376–1385. <https://doi.org/10.1093/brain/awp054>
- Summerfield, C., & de Lange, F. P. (2014). Expectation in perceptual decision making: Neural and computational mechanisms. *Nature Reviews Neuroscience*, *15*(11), 745–756. <https://doi.org/10.1038/nrn3838>
- Tavor, I., Jones, O. P., Mars, R. B., Smith, S. M., Behrens, T. E., & Jbabdi, S. (2016). Task-free MRI predicts individual differences in brain activity during task performance. *Science*, *352*(6282), 216–220. <https://doi.org/10.1126/science.aad8127>
- Thorndike, E. L. (1898). Animal Intelligence: An Experimental Study of the Associative Processes in Animals. *Psychological Monographs: General and Applied*, *2*(4).
<https://doi.org/10.1037/h0092987>
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, *55*(4), 189–208.

- Tricomi, E., Balleine, B. W., & O'Doherty, J. P. (2009). A specific role for posterior dorsolateral striatum in human habit learning. *The European Journal of Neuroscience*, *29*(11), 2225–2232. <https://doi.org/10.1111/j.1460-9568.2009.06796.x>
- Tsodyks, M., Kenet, T., Grinvald, A., & Arieli, A. (1999). Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science*, *286*(5446), 1943–1946.
- Valentin, V. V., Dickinson, A., & O'Doherty, J. P. (2007). Determining the neural substrates of goal-directed learning in the human brain. *Journal of Neuroscience*, *27*(15), 4019–4026. <https://doi.org/10.1523/JNEUROSCI.0564-07.2007>
- van der Meer, M. A. A., Johnson, A., Schmitzer-Torbert, N. C., & Redish, A. D. (2010). Triple dissociation of information processing in dorsal striatum, ventral striatum, and hippocampus on a learned spatial decision task. *Neuron*, *67*(1), 25–32. <https://doi.org/10.1016/j.neuron.2010.06.023>
- Voon, V., Fernagut, P.-O., Wickens, J., Baunez, C., Rodriguez, M., Pavon, N., Bezaud, E. (2009). Chronic dopaminergic stimulation in Parkinson's disease: From dyskinesias to impulse control disorders. *The Lancet. Neurology*, *8*(12), 1140–1149. [https://doi.org/10.1016/S1474-4422\(09\)70287-X](https://doi.org/10.1016/S1474-4422(09)70287-X)
- Voon, V., Thomsen, T., Miyasaki, J. M., Souza, M. de, Shafro, A., Fox, S. H., Zurowski, M. (2007). Factors Associated With Dopaminergic Drug-Related Pathological Gambling in Parkinson Disease. *Archives of Neurology*, *64*(2), 212–216. <https://doi.org/10.1001/archneur.64.2.212>
- Weber, E. U., Shafir, S., & Blais, A.-R. (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review*, *111*(2), 430–445. <https://doi.org/10.1037/0033-295X.111.2.430>
- Weintraub, D., Siderowf, A. D., Potenza, M. N., Goveas, J., Morales, K. H., Duda, J. E., ... Stern, M. B. (2006). Association of dopamine agonist use with impulse control disorders in Parkinson disease. *Archives of Neurology*, *63*(7), 969–973. <https://doi.org/10.1001/archneur.63.7.969>
- Wunderlich, K., Smittenaar, P., & Dolan, R. J. (2012). Dopamine enhances model-based over model-free choice behavior. *Neuron*, *75*(3–4), 418–424. <https://doi.org/10.1016/j.neuron.2012.03.042>
- Yeo, B. T. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., Buckner, R. L. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, *106*(3), 1125–1165. <https://doi.org/10.1152/jn.00338.2011>

Zaghloul, K. A., Blanco, J. A., Weidemann, C. T., McGill, K., Jaggi, J. L., Baltuch, G. H., & Kahana, M. J. (2009). Human substantia nigra neurons encode unexpected financial rewards. *Science*, 323(5920), 1496–1499. <https://doi.org/10.1126/science.1167342>

Chapter 3: Non-invasive Methods for Investigating Activity and Tissue Microstructure in the Human Brain

The subject to be observed lay on a delicately balanced table which could tip downwards either at the head or at the foot if the weight of either end were increased. The moment emotional or intellectual activity began in the subject, down went the balance at the head-end, in consequence of the redistribution of blood in his system.

William James, *The Principles of Psychology*

William James's (James, 1890) account of an experimental setup by Italian scientist Angelo Mosso at the turn of the 19th century described an extraordinary endeavour. In patients with skull defects through which cerebral pulsations could be recorded, Mosso observed that cerebral blood flow appeared to increase at times of enhanced mental activity and with heightened emotional and sensory perceptions (Raichle & Shepherd, 2014). To test the relationship between cerebral circulation and mental activity in healthy people, Mosso built a "human circulation balance" which consisted of a wooden table rested on a fulcrum upon which a subject would be perfectly balanced (Fig. 3.1, Left). The theory was simple: If mental activity was accompanied by increased blood flow to the brain, the weight of the head would increase relative to the rest of the body resulting in a tilt of the balance. During sessions that spanned at least an hour, Mosso exposed subjects to a range of cognitive stimuli modulated by different difficulty levels, such as having them read a newspaper, novel, or mathematics manual. They were also presented with emotional stimuli, such as letters from a spouse or an agitated creditor. In these experiments, Mosso claimed that the balance tilted faster towards the head side as the stimuli increased in difficulty or emotional content (Sandrone et al., 2014). In a remarkable display of insight that would resonate with neuroscientists today, Mosso noted that the validity of blood flow analyses could be potentially confounded by physiological artefacts, leading him to include an array of sphygmographs and other measuring devices to record head movements, respiration, and pulse in a bid to differentiate between signal and noise (Sandrone et al., 2014); Fig. 3.1, Right).

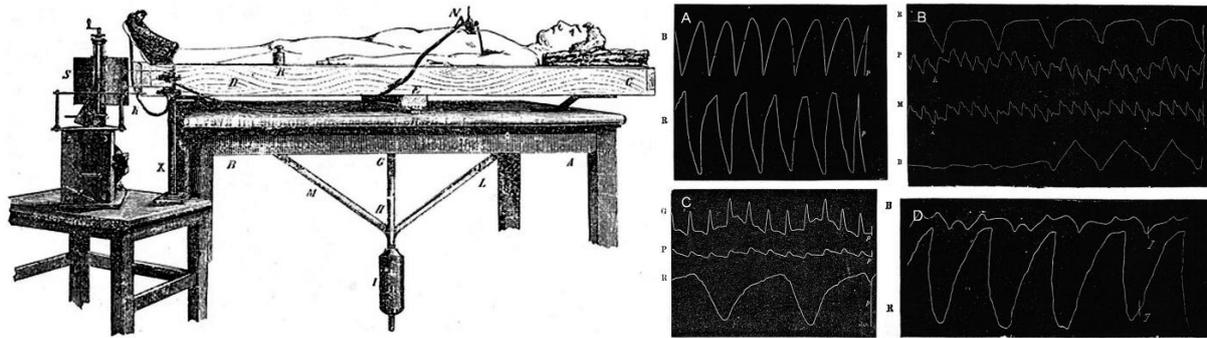


Fig. 3.1. The “human circulation balance”. (Left) Angelo Mosso’s “human circulation balance” consisted of a wooden table rested upon a fulcrum, resembling a seesaw. (Right) Paper tracings of balance movements and physiological artefacts such as respiration and pulse. Adapted from Sandrone et al., 2014.

While the findings by Mosso remain debatable, the suggestion put forth that cerebral circulation was coupled to mental activity proved astute, and work by (Roy & Sherrington, 1890) and Kety and Schmidt (1948) went on to show that changes in cerebral blood flow not only occurred at the local level, but also that these changes were modulated by the brain itself due to increased oxygen consumption by neurons leading to changes in vascular volume and blood flow.

3.1 Blood Oxygen Level Dependent (BOLD) functional Magnetic Resonance Imaging (fMRI)

In the early 1990s, independent groups from the AT&T Laboratories in New Jersey (Ogawa, Lee, Kay, & Tank, 1990) and Massachusetts General Hospital (Kwong et al., 1992) discovered that MRI was sensitive to changes in blood oxygenation with paramagnetic deoxyhemoglobin acting as an *endogenous* contrast agent. These findings were significant as they allowed for dynamic measurement of regional cerebral blood volume during resting and active cognitive states without the need for subjects to be injected with external contrast agents, paving the way for non-invasive BOLD fMRI used today.

3.1.1 MR Physics

Spin or *spin angular momentum* is a quantum property possessed by elementary particles such as protons. The method used to generate MR signals in BOLD fMRI is based on magnetization associated with nuclear spins of hydrogen

nuclei or protons found largely in water molecules in the brain. The spin of protons are randomly oriented in the absence of a magnetic field. The combination of spin and positive charge produces an electric current that also generates a small magnetic field. When these protons are placed in a magnetic field B_0 like that provided by an fMRI scanner, the magnetic moments of protons with positive spin tend to align in a parallel and antiparallel fashion with B_0 to create a net magnetization M_0 .

On top of aligning with B_0 , protons also precess or 'wobble' around the axis in a manner often likened to a spinning top with a rate ω termed the *Larmor frequency*. The relationship between B_0 and ω can be described with the equation:

$$\omega = \frac{\gamma}{2\pi} B_0$$

where γ refers to a gyromagnetic ratio that is constant for a given nuclei taking into consideration their size, mass, and spin. When protons precess together, they are said to be *in phase* and when they precess separately, they are said to be *out of phase*. γ for a proton is $2.675 \times 10^8 \text{ rad}\cdot\text{s}^{-1}\text{T}^{-1}$, so ω at a field strength of 3 Tesla used in this thesis results in a Larmor frequency of 128 MHz. Strength of the net magnetization M_0 along B_0 , typically parallel to the bore of the scanner or z-axis in three-dimensional Cartesian space, is largely dependent on the averaged sum of magnetic properties over individual protons that are aligned parallel rather than antiparallel to B_0 .

Experimental manipulation of the net magnetization vector involves the application of a radio frequency (RF) pulse B_1 perpendicular to the B_0 longitudinal field at the Larmor or *resonant* frequency. This serves to bring the protons in phase and tip the magnetization away from the z-axis through a transfer of energy. The extent of tipping relies on both the magnitude and duration of the RF pulse applied, and is described as the *flip angle*. For example, a flip angle of 90° means that magnetization was completely knocked from the longitudinal or z-axis onto the x-y plane which creates a transverse magnetization component. This excitation creates an oscillating net magnetic flux that generates an electric current in receiver coils, essentially a sine wave signal at the Larmor frequency ω . Due to atomic properties such as de-coherence of transverse spin magnetization, the precessing protons rapidly de-phase resulting in a loss of the transverse magnetization component in a process termed *transverse relaxation*. This signal loss can be described by an exponential decay function with time constant T_2 :

$$M_{xy} = M_0 e^{-t/T_2}$$

where M_{xy} refers to magnetization along the transverse x-y plane and t refers to time. Due to inhomogeneities in the main magnetic field such as field distortions produced by tissue, this decay occurs quicker than would be theoretically expected and the empirical decay is denoted with a time constant T_2^* where T_2^* is always less than or equal to T_2 . Since T_2^* depends on the de-phasing of precessing protons, its value can change with factors that influence the duration of de-phasing such as changes in concentration of paramagnetic deoxyhaemoglobin.

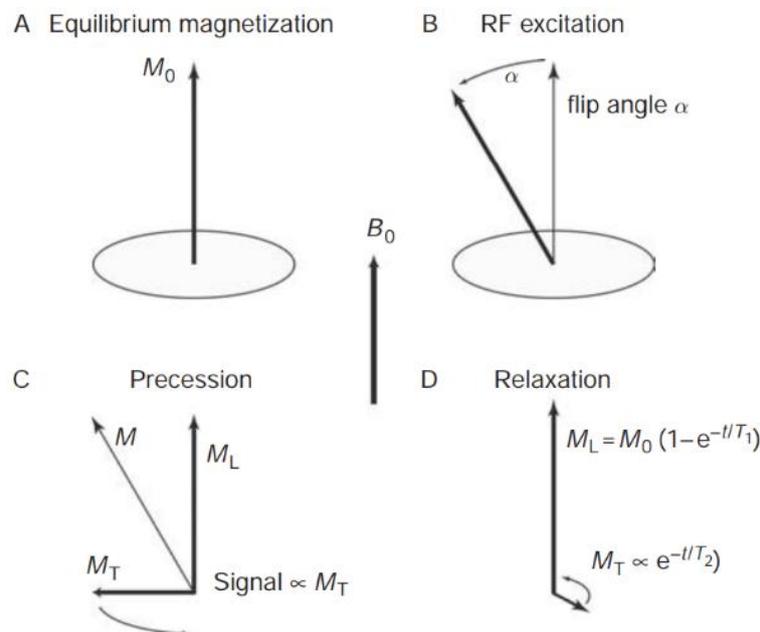


Fig. 3.2. General physics principles of fMRI. **A**, Net magnetization M_0 forms along B_0 due to the proportion of spin-up protons being larger than spin-down protons. **B**, Application of a radiofrequency (RF) pulse knocks M_0 over onto the transverse x-y plane, creating a transverse magnetization component termed M_T or M_{xy} . **C**, The protons precess in-phase around B_0 , generating a MR signal. **D**, The excited protons gradually lose their phase with M_T decaying back to zero with time constant T_2 and net magnetization M_L or M_z recovering to M_0 with time constant T_1 (Buxton, 2009).

In contrast, *longitudinal relaxation* refers to the process by which net magnetization M_z returns to its initial value M_0 aligned with B_0 with time constant T_1 :

$$M_z = M_0 (1 - e^{-t/T_1})$$

Since M_0 depends on the proportion of protons aligned parallel to B_0 in a lower energy spin-up state rather than higher energy spin-down state, regrowth of M_z partially depends on the transfer of energy from spins to the surrounding environment which can consist of neighbouring atoms and molecules. The value of $T1$ differs depending on the location of hydrogen nuclei in the brain, and can thus be used to distinguish between grey and white matter for example. These principles are summarized in Fig. 3.2.

3.1.2 Image Construction

A magnetic field *gradient* is an alteration of the magnetic field such that it varies linearly along a Cartesian axis, allowing the resonant frequency to also change linearly along that particular axis. These gradients are produced by orthogonal sets of gradient coils in the MRI scanner. MRI images typically consist of a series of sequentially-acquired two dimensional images, or *slices*, stitched together to form a three dimensional image or *volume*. The initial step of image acquisition leverages on gradients to restrict the range of on-resonance locations along B_0 or the z-axis when a RF pulse is applied at the Larmor frequency, allowing ω to be centred in the slice of interest. The thickness of these slices can be determined by changing the slope of the gradient or the bandwidth of the RF pulse. Essentially, the spatial distribution of transverse magnetization on the x-y plane at each point in time needs to be resolved for a MR image to be produced.

On the x-axis, a *frequency encoding gradient* provides a means to distinguish between signals on that axis as they would be spread across different frequencies depending on their positions along the gradient. On the y-axis, a *phase encoding gradient* allows protons to be differentiated based on the magnitude of phase shift (i.e. the gain or loss of phase) relative to a reference state. When gradient pulses are turned on followed by a data sample, a data matrix of spatial frequencies known as *k-space* (k_x, k_y) is incrementally produced. The direction of sampling in k-space depends on the repetition of a gradient echo pulse sequence, with a full distribution in the x-axis and a single sample along the y-axis obtained following each RF pulse. The MR image can then be constructed by conducting a Fourier transformation on k-space, and vice versa.

3.1.3 What does the BOLD response reflect?

The BOLD response is an indirect measure of neuronal activity that is contingent upon the variable ratio of deoxygenated to oxygenated blood (Huettel, Song, & McCarthy, 2009). Changes to this ratio are driven by neuronal dynamics occurring at the cellular and micro-circuitry level that involve synaptic or spiking activity (Friston, 2008). To resolve which of the two better accounts for the BOLD response, Logothetis et. al (2001) co-recorded local field potentials (LFP), multi-unit activity (MUA), and fMRI signals from monkeys. LFPs are made up of a weighted average of slow waveforms (< 200 Hz) such as post-synaptic voltage-gated membrane oscillations and somato-dendritic integrative processes, reflecting contributions from neuronal inputs into a regions and local intra-cortical processing respectively (Logothetis & Wandell, 2004). On the other hand, MUA consists of higher frequency waveforms (> 300-400 Hz) and represents the spiking of neuronal populations within a few hundred micrometres from the placement of the electrode (Viswanathan & Freeman, 2007). Logothetis et. al (2001) found that the BOLD response correlated with LFPs, suggesting that BOLD signals primarily reflect pre-synaptic firing rather than post-synaptic spikes. Further evidence for this comes from the finding that a similar correlation can also be observed in the absence of spiking, despite LFPs and spiking often being correlated (Logothetis & Wandell, 2004).

Associated with the BOLD response is the haemodynamic response, which refers to the idealised time course of the BOLD response and is modelled using a haemodynamic response function (HRF) in software packages such as Statistical Parametric Mapping. Often, this response encompasses elements such as blood volume and blood flow in addition to blood oxygenation (Heeger & Ress, 2002; Weiskopf, Hutton, Josephs, & Deichmann, 2006; Weiskopf et al., 2004). While there are several ways of modelling the HRF, the canonical HRF typically reaches peak intensity around 4 to 6 seconds following onset of neuronal activity, after which it gradually returns to baseline over the course of 10 to 20 seconds with a slight undershoot at around the 15-second mark (Lindquist, Loh, Atlas, & Wager, 2009).

fMRI relies on the tight coupling between neural activity and increases in blood flow which in turn changes $T2^*$ and the resulting signals. The neuroimaging studies

described later in **chapters 4** and **6** involve fMRI, with the former using a variant of fMRI termed real-time fMRI which can be defined as “any process that uses functional information from a MRI scanner where the analysis and display of the fMRI keep pace with data acquisition” that is described in the next section (Sulzer et al., 2013).

3.2 Real-time fMRI (rtfMRI) - Setup, Software, and Implementation

3.2.1 Setup and Software

One of the most widely used software packages for rtfMRI experiments is Turbo-Brainvoyager (TBV) (Brain Innovation, Maastricht, The Netherlands). This software is used to export functional images in real-time as Digital Imaging and Communications in Medicine (DICOM) mosaic images from the MRI console computer to a shared folder provided by the TBV computer (Fig. 3.3), resulting in a single file for each functional volume. This helps to avoid a scenario where multiple network queries are required to read the MRI images leading to a deceleration of real-time image export both from the MRI and TBV computers.

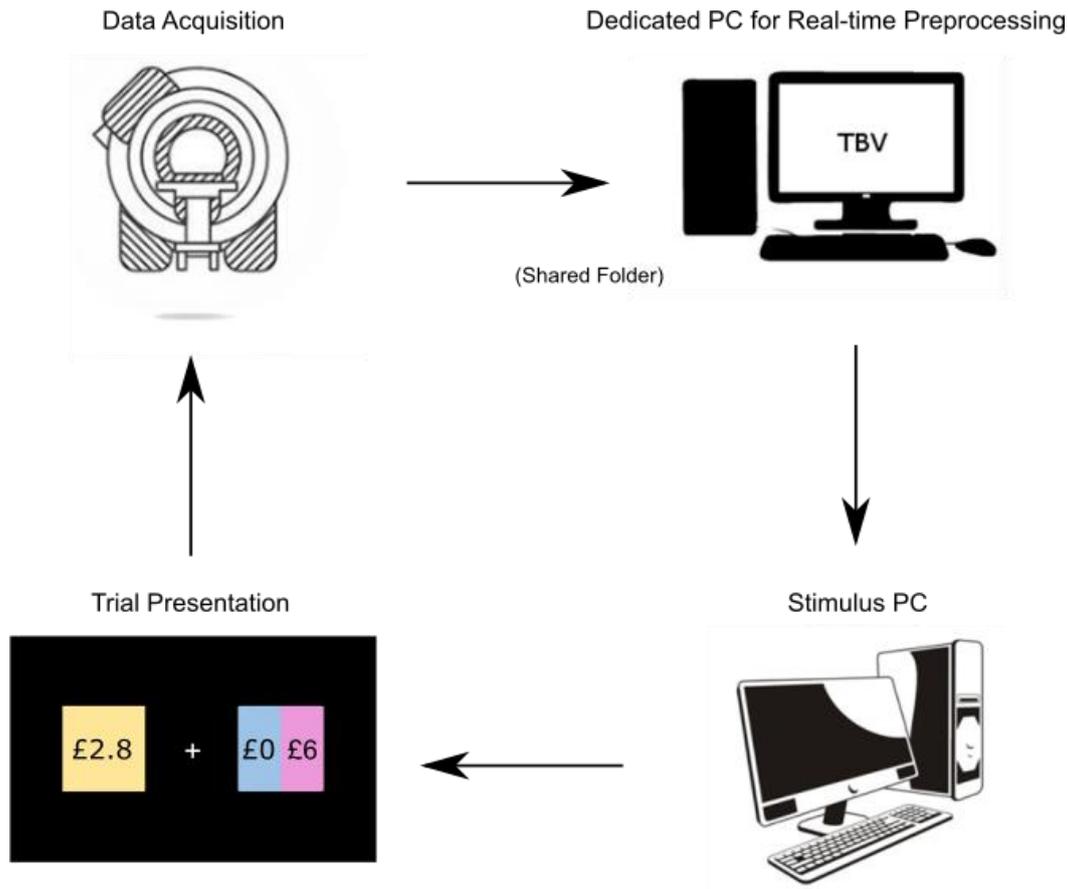


Fig. 3.3. General schematic for Real-Time pipeline. fMRI data is first acquired and exported in real time to a PC that is dedicated to the pre-processing of images using TBV. Note a non-dedicated PC might incur processing delays. The output from TBV is written to a shared folder accessible by the Stimulus PC, which computes changes in the BOLD signal and presents the signal for neurofeedback in a typical experiment.

3.2.2 Pre-processing of Images

Prior to analysis, functional data has to be pre-processed to ensure that certain statistical assumptions are met, such as timecourses coming from a single location and being uniformly spaced in time. Real-time pre-processing of the functional data was performed by TBV, which included realignment and spatial smoothing [6 mm Full Width at Half Maximum (FWHM)]. In the experiment presented in **chapter 4**, I was predominantly interested in the substantia nigra / ventral tegmental area (SN/VTA) complex as my main region-of-interest (ROI). Consequently, time courses for every voxel within this ROI were extracted, averaged, and exported by TBV.

3.2.3 Dealing with Additional Nuisance Regressors

To remove physiological noise arising from respiratory phasicity and cardiac pulsatility, participants were fitted with a pneumatic respiratory belt and a pulse oximeter. Physiological measurements from these devices are modelled using a Fourier expansion of physiological phases based on the RETROICOR model (Glover, Li, & Ress, 2000), and the subsequent regressors can be categorised into ‘phase’ regressors and ‘rate’ regressors. The former set of regressors are used to model cyclic fluctuations in the respiratory and cardiac cycles, while the latter are used to account for dependency of the signal on the rate of physiological processes. These were incrementally regressed out in real time from the exported time courses (Fig. 3.4) using a custom-made MATLAB (MathWorks, Natick, USA) toolbox. Essentially, this amounts to removal of lower frequency drifts through the use of a high-pass filter. Linear detrending of the signal was also performed to correct for signal drift caused by fluctuations in the superconducting magnetic field over time. A failure to correct for this can lead to errors when it comes to the estimation of quantitative or diffusion parameters in neuroimaging. The ensuing ‘cleaned’ time courses were then used in the main experiment.

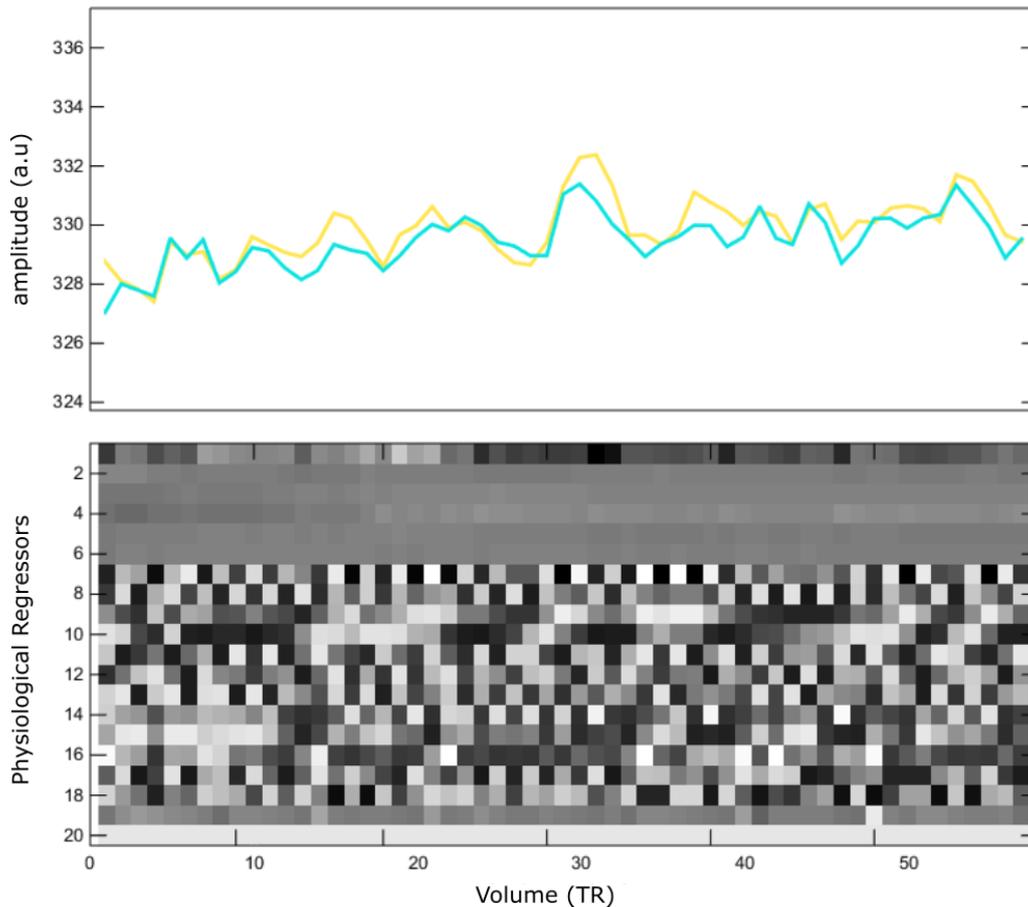


Fig. 3.4. Cartoon example of an exported TBV signal with a General Linear Model (GLM) containing motion and physiological regressors. (Top) The raw BOLD signal is represented here in yellow, while the filtered BOLD signal is coloured blue. (Bottom) Design Matrix containing movement and physiological regressors. Using an incremental GLM approach, these regressors are removed from the raw BOLD signal and the cleaned signal is calculated from the residuals.

3.2.4 Computation of Blood Oxygenation Level Dependent (BOLD) Signal and Level of Activity

From a methodological standpoint, any rtfMRI experiment necessitates that brain states are detectable and can be reliably transformed into a signal used either for neurofeedback or in a novel manner described with more detail in **chapter 4**. rtfMRI approaches are to some degree constrained by the spatial and temporal resolution at which the BOLD response can be measured. Nevertheless, the exact limits are flexible and some studies report early signal changes less than 2 seconds following neuronal activity (Yacoub & Hu, 1999), although a caveat here is that this early component scales with magnetic field strength. This could have implications for rtfMRI as it suggests that the resolution afforded by higher magnetic field strengths (e.g. 7 Tesla)

may be advantageous when it comes to improving not only the spatial resolution of the ROIs, but also the temporal resolution of the BOLD signal exploited in neurofeedback.

In typical rtfMRI studies involving neurofeedback, the presence of haemodynamic delay may not be as crucial since there is evidence demonstrating that delayed feedback as long as 60 seconds can still be used for learning, as long as this delay is consistent and predictable (Miall, Weir, Wolpert, & Stein, 1993). However, this delay was an important consideration for me and advancement in computing power, network speed, and the relative efficiency of TBV have enabled delays associated with acquisition and pre-processing of functional images to be reduced to approximately 1 second, as was achieved in my main experiment.

There are a variety of methods used to derive the BOLD signal for use in rtfMRI experiments, and they are often based either on the General Linear Model (GLM) or multivariate pattern analysis (MVPA). GLM methods typically involve defining a ROI through the use of a functional localizer, such as for primary motor cortex (M1) (Chiew, LaConte, & Graham, 2012; Yoo & Jolesz, 2002), but can also be defined anatomically with the aid of anatomical landmarks or brain atlases. The choice of whether to define a ROI functionally or anatomically depends on 1) the availability of a reliable functional localizer and 2) whether the target area is anatomically well demarcated (Sulzer et al., 2013). While the target signal is usually the average BOLD response from a single ROI, differential activity between two ROIs can also be used in certain cases, adding on another layer of complexity (Chiew et al., 2012). By applying an incremental GLM to the time series from each ROI as a new measurement becomes available, the fit of the model is updated and nuisance variables are regressed out. The residual of the GLM is scaled by the standard deviation of the fit to derive activation estimates for each voxel, which can then be averaged across the ROI for use in the rtfMRI experiment (Hinds et al., 2011). An alternative to the incremental GLM method is the use of functional connectivity measures over a short sliding window of a few seconds (Zilverstand, Sorger, Zimmermann, Kaas, & Goebel, 2014). This approach helps increase temporal resolution while limiting spatial coverage (Posse et al., 2003), and the computed signal derived from this method is based on functional correlations from one time point to the next rather than the level of activation in voxels.

When network activity is anticipated over the whole brain, or a pre-selected set of regions, rather than specific ROIs, or when participants are afforded some degree of freedom regarding cognitive strategies for the same task, MVPA might be more advantageous than a GLM. MVPA can be construed as a supervised machine learning problem which means that techniques such as support vector machines (LaConte, 2011), neural networks (Eklund et al., 2009), logistic regressions (deBettencourt, Cohen, Lee, Norman, & Turk-Browne, 2015), and many others can be used to derive optimal weights for the combination of BOLD signal across voxels. MVPA has the capability to adapt to the individual or task, and is often used to predict brain states – physiological/behavioural events, or mental processes for which neural correlates can be obtained – by examining the distributed (or multivariate) relationships between task measures and fMRI images (LaConte, 2011).

3.2.5 Quantifying Level of BOLD Activity

BOLD Signal change due to functional activation is typically around 0.5% to 1.5% in many brain areas (Brühl, 2015). While the experiment described in **chapter 4** does not involve feedback of the signal to participants, it nevertheless requires the ability to detect changes in the signal that is compatible with methods used for neurofeedback studies.

Several factors have to be taken into consideration when deciding on a method to obtain a measure of baseline activity. Firstly, slow drifts and changes in variance of the BOLD signal have to be accounted for within the baseline measure. An intuitive approach to quantify the change in signal is to calculate the mean or median percentage signal change. The advantage of the former is that it is less susceptible to noise, but can be affected by outlier activations from the voxels recorded. Comparing the activity at each incoming time point without correcting for drifts in the signal or temporal fluctuations in the amplitude of the signal means that the summary statistics of the signal may not be sufficiently sensitive to changes of the signal over time. Take for example a case where the variance of the signal is not constant over time. Two data points with the same amplitude can be significantly different from the local mean in one case but not another (Fig. 3.5), yet may be labelled in a similar manner by an algorithm that simply averages the baseline across the entire timeseries. In a similar

vein, even if the variance were to remain constant, the presence of a slow drift would change the global mean over time such that activity of the most recent data point may appear significantly different from the mean even if its amplitude was exactly the same as a previous data point that was not.

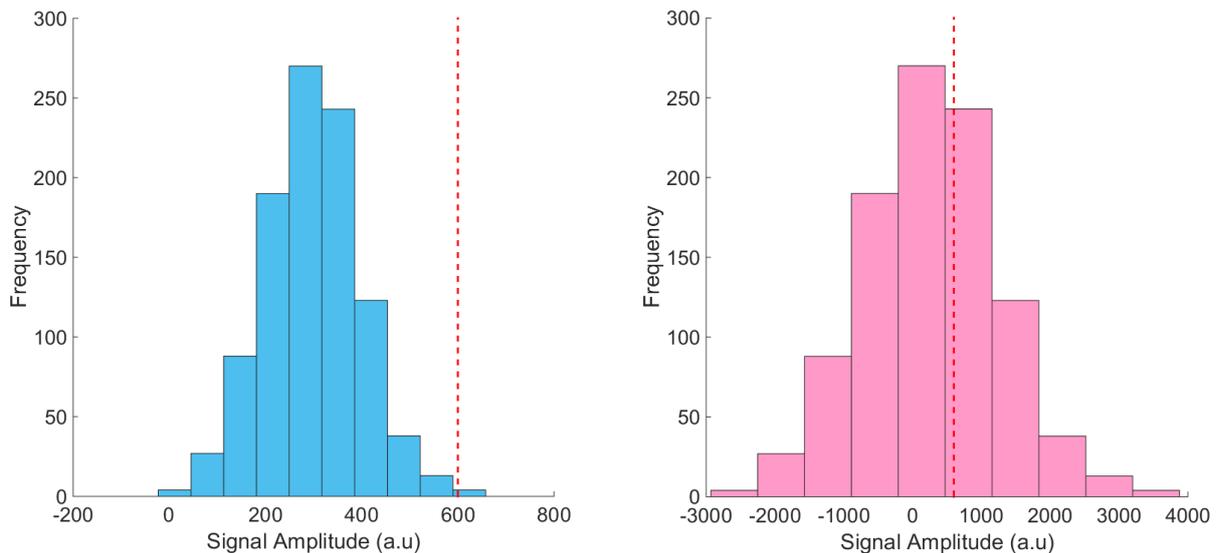


Fig. 3.5. Simulated distribution of signal amplitudes. A peak amplitude of 600 would be distant from the mean in the distribution on the *left* but close to the mean in the distribution on the *right*, yet would be labelled in a similar manner by an algorithm that only relies on the change of amplitude from the global mean without taking the variance of the signal into consideration.

To overcome these potential problems, I took a sliding window approach in my experiment. By using a temporal window of 2 minutes and normalizing the signal within this timeframe, I could utilize a normal cumulative distribution function to quantify the distribution of signals in terms of percentiles, and detect changes in the signal based on this. Activity was averaged for the most recent 2 volumes acquired and compared against the preceding normalised baseline in order to further reduce susceptibility to noise. This approach is more robust in the face of activation outliers over the course of each run and also accounts for changes in the variance of the signal over time. However, as with most methods of quantifying BOLD activity, this method is unable to correct for large head movements beyond a threshold of 3 mm. Turning to the future, more sophisticated methods to account for motion are currently being developed such as the prospective motion correction (PMC) system. The PMC system uses an optical

camera in the bore of the scanner to track a marker attached to the head of the subject in real time, allowing the imaging parameters to be updated dynamically to account for motion (Callaghan et al., 2015).

3.3 Quantitative Magnetic Resonance Imaging

Recent developments in neuroimaging have enabled *in vivo* mapping of neuroimaging markers of biologically relevant quantities to be performed with high resolution in the form of *Quantitative MRI*. Anatomical studies relying on morphometric analyses of grey and white matter volumes using voxel-based morphometry typically utilize T1-weighted images, leading to difficulties comparing anatomical data between imaging sites and sessions in multi-center studies due to the arbitrary units of signal intensities. On the other hand, Quantitative MRI is sensitive to tissue microstructure and provides voxel-wise absolute measures of myelination, water concentration, and iron levels that are comparable across imaging sites and time points (Weiskopf et al., 2013). The main approach used in Quantitative MRI is known as Multi-Parameter Mapping (MPM), which uses biophysical models to produce quantitative maps of magnetisation transfer saturation (MT), proton density (PD*) and relaxometry measures R_1 and R_2^* where $R_1 = 1/T_1$ and $R_2^* = 1/T_2^*$ that depend on underlying tissue microstructure, increasing sensitivity of Quantitative MRI to specific microstructural features such as myelination (Weiskopf, Mohammadi, Lutti, & Callaghan, 2015). While whole-brain MPMs can be acquired, parameters such as the field of view and isotropic resolution can be altered to examine deep brain structures that have been difficult to distinguish using traditional T1-weighted images, such as the locus coeruleus, substantia nigra, and ventral tegmental areas. In **chapter 4**, I leverage on the excellent contrast provided by the magnetisation transfer saturation (MT) map to define the substantia nigra / ventral tegmental area complex using a procedure described in the next section.

The ability to estimate tissue microstructure *in vivo* and non-invasively has opened up new research possibilities for neuroscience and clinical applications. To disentangle normal age-related changes in the brain from pathological neurodegeneration, Callaghan et al. (2014) acquired multi-parameter maps for a healthy cohort between 19-75 years of age to establish a quantitative baseline against

which similar measures from a clinical population can be compared. Leveraging on the fact that quantitative measures can be compared across different time, Ziegler et al. (2018) examined a longitudinal dataset where MT maps were acquired for subjects over multiple sessions during adolescent development. They found that psychiatric risk factors like compulsivity and impulsivity were tightly linked with aberrant myelin-related growth in the cingulate and ventral striatum for the former, and dorsal striatum and lateral prefrontal regions for the latter. Quantitative MRI can also be used to link microstructure to cognitive tasks and abilities. For example, myeloarchitecture in the right anterior prefrontal cortex appears to be predictive of metacognitive ability in the perceptual domain (Allen et al., 2017), while iron and myelin levels in the ventral striatum were predictive of performance on verbal learning memory tests in the ageing brain (Steiger, Weiskopf, & Bunzeck, 2016).

In **chapter 5**, I use Quantitative MRI to investigate the relationship between tissue microstructure and risk preferences.

3.4 Drawing and Preparation of SN/VTA ROI for rtfMRI

3.4.1 Magnetization Transfer

Protons involved in the generation of the MR signal are found predominantly in three pools when it comes to non-fatty tissue: 1) Free Water, 2) Bound Water, and 3) Macromolecules. Protons found in free water contain rapid rotations and are largely unstructured, making them inefficient in the production of MR signal. Protons present in bound water - found on the surface of macromolecules – and macromolecules are moderately- to highly-structured resulting in a restricted range of motion which leads to different relaxation properties when a RF pulse is applied than when protons are able to rotate freely (Henkelman, Stanisz, & Graham, 2001).

Magnetization transfer typically refers to the transfer of energy between protons found in free water, bound water, and macromolecules such as those present in myelin. Measures of magnetization transfer have been demonstrated to correlate with myelin content as assessed histologically (Schmierer, Scaravilli, Altmann, Barker, & Miller, 2004). By using a specially designed RF pulse to saturate the macromolecular protons, net magnetization of areas with higher levels of macromolecules will be reduced as

energy gets transferred to the surrounding unaffected pools in a bid to return the energy state of the macromolecular protons to equilibrium (Wolff & Balaban, 1989). In the case of the SN, this results in greater delineation of SN from the surrounding white matter compared to T1-weighted images where little contrast is observed (Helms, Draganski, Frackowiak, Ashburner, & Weiskopf, 2009) (Fig. 3.6).

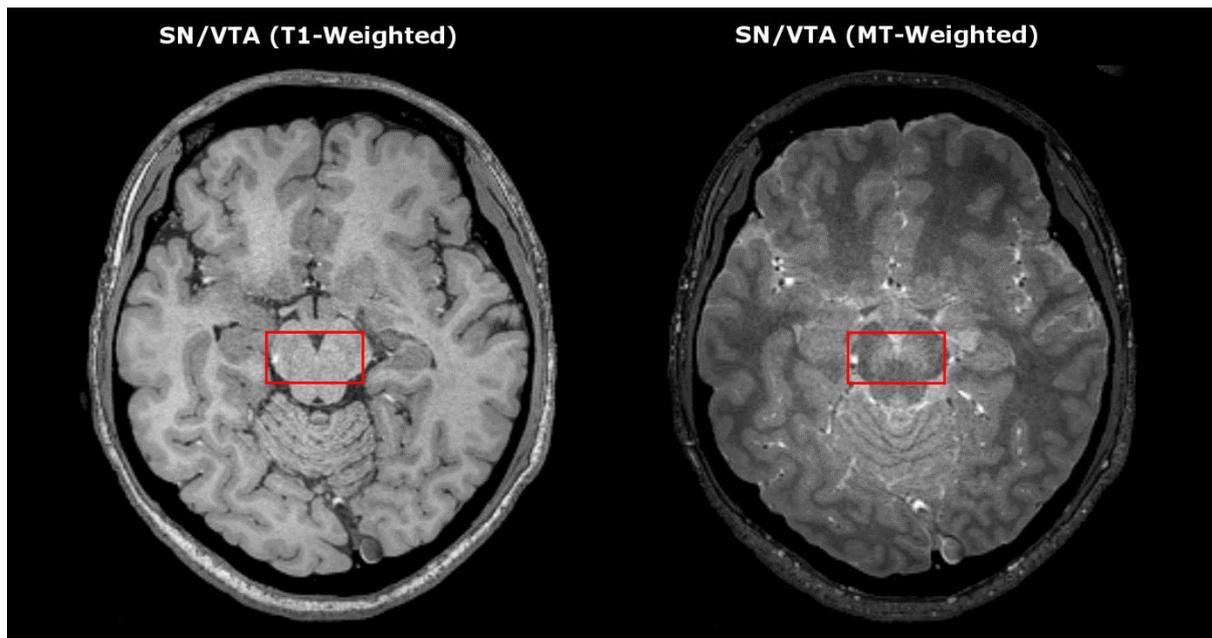


Fig. 3.6. Illustration of contrast around the region of interest. (Left) In the T1-Weighted image, SN/VTA is indistinguishable from the surrounding white matter tract. (Right) In the MT-Weighted image, SN/VTA shows up as a hyperintense region amidst the darker grey of the surrounding tissue.

3.4.2 Definition of ROI

Bright areas in MT-contrast images have been shown to be coextensive with the SN as delineated histologically by tyrosine hydroxylase immunohistochemistry, which stains dopaminergic neurons (Bolding et al., 2013) that are the key component of SN/VTA. Leveraging upon this, SN/VTA ROIs can be hand-drawn for each individual participant in MRICron (Rorden & Brett, 2000) using their MT-weighted structural images.

In accordance with procedures previously outlined by Chowdhury, Guitart-Masip, Lambert, Dolan, and Düzel (2013) and Düzel et al. (2008), medial and lateral boundaries of the SN/VTA ROI were defined based on the change in contrast between

its bright grey colour and the dark grey colour of the adjacent cerebral peduncle and interpeduncular fossa. Lower and upper boundaries of the ROI were selected as the slices preceding the ones where the intensity of SN/VTA was indistinguishable from surrounding tissue, totalling between 6 to 9 slices contingent on individual SN/VTA size differences.

3.4.3 Transformation of ROI

To prepare the hand-drawn SN/VTA ROI for use in TBV, it needs first be co-registered and transformed to the space and resolution of the EPIs. Co-registration is carried out using a single EPI volume (see Image Acquisition) as the reference image, and the individual-specific T1-weighted image as the source image. Following this, the EPI voxels corresponding to each ROI voxel is indexed based on Euclidean distance calculated in native space. Since the coordinate space used in Turbo-Brainvoyager differs from common ones such as the Montreal Neurological Institute (MNI) space, coordinates for the ROI must also be transformed before it can be accurately positioned in Turbo-Brainvoyager. This series of co-registration and transformations is executed using custom MATLAB scripts that have been made available on [Github](#) (Chew & Hauser, 2017). Following these steps, the SN/VTA ROI can now be used in Turbo-Brainvoyager.

3.5 fMRI BOLD response and the SN/VTA complex

Before we embark on the experimental portion of the thesis from the next chapter, I will briefly review the SN/VTA complex and the interpretations of BOLD responses within the region as these will be relevant for the study described in **chapter 4**.

Dopamine cells in the human midbrain largely resides within three different groups: retrobulbar, SN, and VTA (Volkow et al., 1996). While evidence supporting the functional and anatomical segregation of SN and VTA in primates is generally weak, unlike in rodents (E. Düzel et al., 2009), recent studies have attempted to parcellate the SN/VTA into sub-regions based on anatomical connectivity to other brain areas like striatum, as assessed by a combination of diffusion tensor imaging (DTI) and quantitative MRI (qMRI) (Chowdhury, Lambert, Dolan, & Düzel, 2013).

Although the proportion of dopaminergic neurons in relation to non-dopaminergic neurons appears higher in SN than VTA, taken as a whole, the SN/VTA complex contains the highest concentration of dopamine cells in the human brain. Dopamine neurons from the SN mainly project to the dorsal striatum while those from the VTA form broad reciprocal connections with cortical and sub-cortical brain regions, making the SN/VTA complex an important site of information integration (Oades & Halliday, 1987).

Through the mesolimbic pathway, the SN/VTA complex forms reciprocal connections with parts of the limbic system including nucleus accumbens (nAcc), hippocampus, amygdala, and cingulate cortex (Adell & Artigas, 2004). This complex projects to, and receives, afferent inputs from cortex - especially the prefrontal cortex (PFC) - via the mesocortical pathway. In addition, it also has connections with the thalamus, hypothalamus, locus coeruleus (LC), raphe nuclei, ventral pallidum, lateral habenula, superior colliculus, and reticular formation periaqueductal grey (Ferrucci, Giorgi, Bartalucci, Busceti, & Fornai, 2013; Haber & Fudge, 1997). In terms of afferent-type, the SN/VTA complex receives glutamatergic inputs from the prefrontal cortex, glutamatergic, cholinergic, and GABAergic inputs from the pedunculo-pontine tegmental nucleus (Nestler, Hyman, & Malenka, 2008; Watabe-Uchida, Zhu, Ogawa, Vamanrao, & Uchida, 2012), noradrenergic afferents from the LC (Sara, 2009), serotonergic innervations from the raphe nuclei (Nakamura, 2013), and GABAergic projections from the nAcc (Russo & Nestler, 2013). Activity of SN/VTA dopamine neurons is suppressed by activation of the lateral habenula, which projects to GABAergic neurons within subregions of the VTA via glutamatergic inputs and its inputs inhibits firing of dopamine neurons (Ji & Shepard, 2007).

As mentioned in **chapter 2**, dopamine is involved in modulatory aspects of behaviour and cognition, and has been found to govern neural mechanisms underlying reward-seeking (Pessiglione, Seymour, Flandin, Dolan, & Frith, 2006), addiction (Hyman, Malenka, & Nestler, 2006), reinforcement learning (Schultz, Dayan, & Montague, 1997), motivation (Westbrook & Braver, 2016), vigour (Niv, Daw, Joel, & Dayan, 2007; Rigoli, Chew, Dayan, & Dolan, 2016), risk-seeking (Rigoli et al., 2016), subjective well-being (Rutledge, Skandali, Dayan, & Dolan, 2015), working memory (Williams & Goldman-Rakic, 1995), and action selection and movement (Parker et al., 2016). Despite the abundance of dopamine neurons in the SN/VTA complex, it is

important to note when considering the meaning of the BOLD response in SN/VTA that there are also glutamatergic and GABAergic efferents from SN/VTA to areas such as nAcc (Qi et al., 2016) as well as an abundance of reciprocal connections present in SN/VTA. Although it is tempting on this basis to conclude that the BOLD response in SN/VTA reflects input from another brain area, such as projections from the lateral habenula, this is not necessarily the case as the firing of pre-synaptic terminals can also be driven by recurrent excitatory-inhibitory loops of a microcircuit within the same brain area (Goense & Logothetis, 2008).

Duzel et al. (2009) propose that the BOLD signal in SN/VTA may arise from several physiological candidates such as local field potentials produced by glutamatergic inputs onto dopamine neurons or the inhibition of GABAergic inputs onto dopamine neurons (Fig. 3.7).

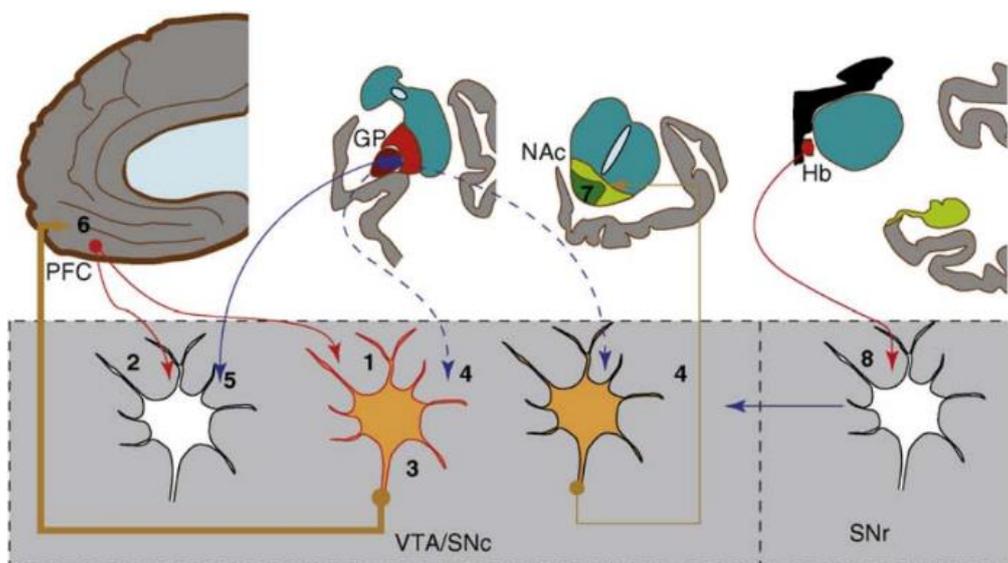


Fig. 3.7. Possible origins of the BOLD response in SN/VTA. 1) Local field potentials by glutamatergic afferents onto active dopamine neurons; 2) Local field potentials by glutamatergic afferents onto silent dopamine neurons; 3) spiking dopamine neurons; 4) Local field potentials by inhibition of GABAergic afferents onto dopamine neurons; 5) Local field potentials by GABAergic afferents onto dopamine neurons; 6) Spiking dopamine neurons releasing dopamine; 7) Tonic dopamine neurons releasing dopamine (Duzel et al., 2009).

Although there is no conclusive physiological evidence that points to one single mechanism eliciting the BOLD response, the most straightforward method of triggering a BOLD response – given the correlation between BOLD and local field potentials –

would be via glutamatergic afferents from the prefrontal cortex (PFC) or mesopontine nuclei, at least for motivationally-salient or reward-related stimuli. For example, in one of the first human fMRI studies focused exclusively on the VTA, D'Ardenne, McClure, Nystrom, and Cohen (2008) leveraged upon the reward prediction error hypothesis of dopamine function (Schultz, Dayan, & Montague, 1997) to show that BOLD response in the VTA reflected positive reward prediction errors which occurs when the reward received by the participant exceeded expectations. Excitatory glutamatergic direct afferents to the VTA could account for the increased BOLD response during a positive RPE, and although D'Ardenne et al. (2008) did not find a significant BOLD response corresponding to a negative RPE, it is not inconceivable that glutamatergic afferents onto inhibitory interneurons in the SN/VTA inhibit the firing of dopamine neurons despite elevated local field potentials. A reason that this might not have translated to a BOLD signal if local field potentials did change is that the proportion of inhibitory interneurons in the SN/VTA is relatively small (~30%), limiting the sensitivity of fMRI to detect their engagement (Düzel et al., 2009).

Recent studies in rodents suggest a causal role for SN/VTA GABAergic neurons in the calculation of RPEs, which suggests the plausibility of recurrent activity within the same area generating local field potentials that modulate the BOLD response as opposed to just glutamatergic afferents onto these neurons (Eshel et al., 2015). Thus, it is highly likely that the origin of the BOLD signal in the SN/VTA complex is heterogeneous and dependent – at least partially – on the task at hand since direct and recurrent afferents into the SN/VTA complex arise from many different brain regions and the SN/VTA itself (Lammel, Ion, Roeper, & Malenka, 2011). Given that many of these connections with other brain regions are reciprocal in nature, it has recently been demonstrated that activating and inhibiting dopamine neurons in the SN/VTA using optogenetics lead to downstream modulation of reward-related BOLD signal in striatum – and indeed also global BOLD signal changes across the brain (Lohani, Poplawsky, Kim, & Moghaddam, 2017) – and this relationship can be attenuated by a top-down input from PFC that is in turn correlated with predictable changes in reward-related behaviours (Ferenczi et al., 2016). Understanding the origin of the BOLD signal in SN/VTA is crucial to the interpretation of studies such as this.

Although a causal link between firing of dopamine neurons and the BOLD signal has yet to be established physiologically in humans, there is nevertheless converging

evidence that BOLD response in the SN/VTA is strongly linked to dopamine neurotransmission. In rats, optogenetic stimulation of midbrain dopamine neurons increased BOLD activity in downstream regions such as the striatum (Ferenczi et al., 2016). Following a neurotoxic lesion that render rhesus macaques hemiparkinsonian, Zhang et al., (2006) found that BOLD response elicited by the dopamine agonist amphetamine was correlated with the amount of residual dopamine neurons that remained intact in the lesioned site of SN/VTA. Using a delayed monetary incentive delay task in humans, Schott et al. (2008) demonstrated that reward-related DA release in the nAcc, as measured by [¹¹C]raclopride positron emission tomography (PET), was correlated with BOLD response in the SN/VTA during reward anticipation in the same set of participants, providing evidence for a quantitative relationship between SN/VTA BOLD and dopamine neurotransmission.

In the preceding chapters, I have reviewed the current literature on value-based decision-making, discussed the evidence for a relationship between dopamine and risk preferences, and provided a primer on non-invasive measures of neural activity and how the BOLD response can be interpreted. In the following chapter, I bring together these components in a study where I use a real-time fMRI protocol to investigate the influence of endogenous fluctuations of BOLD activity in the SN/VTA on risk preferences.

References

- Adell, A., & Artigas, F. (2004). The somatodendritic release of dopamine in the ventral tegmental area and its regulation by afferent transmitter systems. *Neuroscience and Biobehavioral Reviews*, 28(4), 415–431.
<https://doi.org/10.1016/j.neubiorev.2004.05.001>
- Allen, M., Glen, J. C., Müllensiefen, D., Schwarzkopf, D. S., Fardo, F., Frank, D., Rees, G. (2017). Metacognitive ability correlates with hippocampal and prefrontal microstructure. *NeuroImage*, 149, 415–423.
<https://doi.org/10.1016/j.neuroimage.2017.02.008>
- Bolding, M. S., Reid, M. A., Avsar, K. B., Roberts, R. C., Gamlin, P. D., Gawne, T. J., ... Lahti, A. C. (2013). Magnetic transfer contrast accurately localizes substantia nigra confirmed by histology. *Biological Psychiatry*, 73(3), 289–294.
<https://doi.org/10.1016/j.biopsych.2012.07.035>
- Brühl, A. B. (2015). Making sense of real-time functional magnetic resonance imaging (rtfMRI) and rtfMRI neurofeedback. *The International Journal of Neuropsychopharmacology*, 18(6). <https://doi.org/10.1093/ijnp/pyv020>
- Buxton, R. B. (2009). *Introduction to Functional Magnetic Resonance Imaging: Principles and Techniques* (Vol. 9780521899956). Cambridge: Cambridge University Press.
- Callaghan, M. F., Freund, P., Draganski, B., Anderson, E., Cappelletti, M., Chowdhury, R., Weiskopf, N. (2014). Widespread age-related differences in the human brain microstructure revealed by quantitative magnetic resonance imaging. *Neurobiology of Aging*, 35(8), 1862–1872. <https://doi.org/10.1016/j.neurobiolaging.2014.02.008>
- Callaghan, M. F., Josephs, O., Herbst, M., Zaitsev, M., Todd, N., & Weiskopf, N. (2015). An evaluation of prospective motion correction (PMC) for high resolution quantitative MRI. *Frontiers in Neuroscience*, 9. <https://doi.org/10.3389/fnins.2015.00097>
- Chew, B., & Hauser, T. U. (2017). *Functions that are helpful when working with real-time fMRI & TurboBrainVoyager* [Matlab]. Retrieved from <https://github.com/tuhauser/rtfMRI> (Original work published 2017)
- Chiew, M., LaConte, S. M., & Graham, S. J. (2012). Investigation of fMRI neurofeedback of differential primary motor cortex activity using kinesthetic motor imagery. *NeuroImage*, 61(1), 21–31. <https://doi.org/10.1016/j.neuroimage.2012.02.053>
- Chowdhury, R., Guitart-Masip, M., Lambert, C., Dolan, R. J., & Düzel, E. (2013). Structural integrity of the substantia nigra and subthalamic nucleus predicts flexibility of instrumental learning in older-age individuals. *Neurobiology of Aging*, 34(10), 2261–2270. <https://doi.org/10.1016/j.neurobiolaging.2013.03.030>

- Chowdhury, R., Lambert, C., Dolan, R. J., & Düzel, E. (2013). Parcellation of the human substantia nigra based on anatomical connectivity to the striatum. *NeuroImage*, *81*, 191–198. <https://doi.org/10.1016/j.neuroimage.2013.05.043>
- D'Ardenne, K., McClure, S. M., Nystrom, L. E., & Cohen, J. D. (2008). BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. *Science*, *319*(5867), 1264–1267. <https://doi.org/10.1126/science.1150605>
- deBettencourt, M. T., Cohen, J. D., Lee, R. F., Norman, K. A., & Turk-Browne, N. B. (2015). Closed-loop training of attention with real-time brain imaging. *Nature Neuroscience*, *18*(3), 470–475. <https://doi.org/10.1038/nn.3940>
- Düzel, E., Bunzeck, N., Guitart-Masip, M., Wittmann, B., Schott, B. H., & Tobler, P. N. (2009). Functional imaging of the human dopaminergic midbrain. *Trends in Neurosciences*, *32*(6), 321–328. <https://doi.org/10.1016/j.tins.2009.02.005>
- Düzel, S., Schütze, H., Stallforth, S., Kaufmann, J., Bodammer, N., Bunzeck, N., Düzel, E. (2008). A close relationship between verbal memory and SN/VTA integrity in young and older adults. *Neuropsychologia*, *46*(13), 3042–3052. <https://doi.org/10.1016/j.neuropsychologia.2008.06.001>
- Eklund, A., Ohlsson, H., Andersson, M., Rydell, J., Ynnerman, A., & Knutsson, H. (2009). Using real-time fMRI to control a dynamical system by brain activity classification. *Medical Image Computing and Computer-Assisted Intervention: MICCAI ... International Conference on Medical Image Computing and Computer-Assisted Intervention*, *12*(Pt 1), 1000–1008.
- Eshel, N., Bukwich, M., Rao, V., Hemmelder, V., Tian, J., & Uchida, N. (2015). Arithmetic and local circuitry underlying dopamine prediction errors. *Nature*, *525*(7568), 243–246. <https://doi.org/10.1038/nature14855>
- Ferenczi, E. A., Zalocusky, K. A., Liston, C., Grosenick, L., Warden, M. R., Amatya, D., Deisseroth, K. (2016). Prefrontal cortical regulation of brainwide circuit dynamics and reward-related behavior. *Science*, *351*(6268), aac9698. <https://doi.org/10.1126/science.aac9698>
- Ferrucci, M., Giorgi, F. S., Bartalucci, A., Busceti, C. L., & Fornai, F. (2013). The effects of locus coeruleus and norepinephrine in methamphetamine toxicity. *Current Neuropharmacology*, *11*(1), 80–94. <https://doi.org/10.2174/157015913804999522>
- Friston, K. (2008). Neurophysiology: The brain at work. *Current Biology*, *18*(10), R418–R420. <https://doi.org/10.1016/j.cub.2008.03.042>
- Glover, G. H., Li, T. Q., & Ress, D. (2000). Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. *Magnetic Resonance in Medicine*, *44*(1), 162–167.

- Goense, J. B. M., & Logothetis, N. K. (2008). Neurophysiology of the BOLD fMRI signal in awake monkeys. *Current Biology*, *18*(9), 631–640.
<https://doi.org/10.1016/j.cub.2008.03.054>
- Haber, S. N., & Fudge, J. L. (1997). The primate substantia nigra and VTA: Integrative circuitry and function. *Critical Reviews in Neurobiology*, *11*(4), 323–342.
- Heeger, D. J., & Ress, D. (2002). What does fMRI tell us about neuronal activity? *Nature Reviews Neuroscience*, *3*(2), 142–151. <https://doi.org/10.1038/nrn730>
- Helms, G., Draganski, B., Frackowiak, R., Ashburner, J., & Weiskopf, N. (2009). Improved segmentation of deep brain grey matter structures using magnetization transfer (MT) parameter maps. *Neuroimage*, *47*(1), 194–198.
<https://doi.org/10.1016/j.neuroimage.2009.03.053>
- Henkelman, R. M., Stanisz, G. J., & Graham, S. J. (2001). Magnetization transfer in MRI: A review. *NMR in Biomedicine*, *14*(2), 57–64. <https://doi.org/10.1002/nbm.683>
- Hinds, O., Ghosh, S., Thompson, T. W., Yoo, J. J., Whitfield-Gabrieli, S., Triantafyllou, C., & Gabrieli, J. D. E. (2011). Computing moment-to-moment BOLD activation for real-time neurofeedback. *NeuroImage*, *54*(1), 361–368.
<https://doi.org/10.1016/j.neuroimage.2010.07.060>
- Huettel, S. A., Song, A. W., & McCarthy, G. (2009). *Functional magnetic resonance imaging* (2nd ed.). Sunderland, Mass.: Sinauer Associates.
- Hyman, S. E., Malenka, R. C., & Nestler, E. J. (2006). Neural mechanisms of addiction: The role of reward-related learning and memory. *Annual Review of Neuroscience*, *29*, 565–598. <https://doi.org/10.1146/annurev.neuro.29.051605.113009>
- James, W. (1890). *The principles of psychology*. London: Macmillan.
- Ji, H., & Shepard, P. D. (2007). Lateral habenula stimulation inhibits rat midbrain dopamine neurons through a GABA(A) receptor-mediated mechanism. *Journal of Neuroscience*, *27*(26), 6923–6930. <https://doi.org/10.1523/JNEUROSCI.0958-07.2007>
- Kety, S. S., & Schmidt, C. F. (1948). The nitrous oxide method for the quantitative determination of cerebral blood flow in man: Theory, procedure and normal values 1. *Journal of Clinical Investigation*, *27*(4), 476–483.
- Kwong, K. K., Belliveau, J. W., Chesler, D. A., Goldberg, I. E., Weisskoff, R. M., Poncelet, B. P., Turner, R. (1992). Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *PNAS*, *89*(12), 5675–5679.
- LaConte, S. M. (2011). Decoding fMRI brain states in real-time. *NeuroImage*, *56*(2), 440–454. <https://doi.org/10.1016/j.neuroimage.2010.06.052>

- Lammel, S., Ion, D. I., Roeper, J., & Malenka, R. C. (2011). Projection-specific modulation of dopamine neuron synapses by aversive and rewarding stimuli. *Neuron*, *70*(5), 855–862. <https://doi.org/10.1016/j.neuron.2011.03.025>
- Lindquist, M. A., Loh, J. M., Atlas, L. Y., & Wager, T. D. (2009). Modeling the hemodynamic response function in fMRI: Efficiency, bias and mis-modeling. *NeuroImage*, *45*(1), S187-198. <https://doi.org/10.1016/j.neuroimage.2008.10.065>
- Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, *412*(6843), 150–157. <https://doi.org/10.1038/35084005>
- Logothetis, N. K., & Wandell, B. A. (2004). Interpreting the BOLD Signal. *Annual Review of Physiology*, *66*(1), 735–769. <https://doi.org/10.1146/annurev.physiol.66.082602.092845>
- Lohani, S., Poplawsky, A. J., Kim, S.-G., & Moghaddam, B. (2017). Unexpected global impact of VTA dopamine neuron activation as measured by opto-fMRI. *Molecular Psychiatry*, *22*(4), 585–594. <https://doi.org/10.1038/mp.2016.102>
- Miall, R. C., Weir, D. J., Wolpert, D. M., & Stein, J. F. (1993). Is the cerebellum a smith predictor? *Journal of Motor Behavior*, *25*(3), 203–216. <https://doi.org/10.1080/00222895.1993.9942050>
- Nakamura, K. (2013). The role of the dorsal raphé nucleus in reward-seeking behavior. *Frontiers in Integrative Neuroscience*, *7*, 60. <https://doi.org/10.3389/fnint.2013.00060>
- Nestler, E. J., Hyman, S. E., & Malenka, R. C. (2008). *Molecular neuropharmacology: A foundation for clinical neuroscience* (2nd ed.). McGraw-Hill Professional.
- Niv, Y., Daw, N. D., Joel, D., & Dayan, P. (2007). Tonic dopamine: Opportunity costs and the control of response vigor. *Psychopharmacology*, *191*(3), 507–520. <https://doi.org/10.1007/s00213-006-0502-4>
- Oades, R. D., & Halliday, G. M. (1987). Ventral tegmental (A10) system: Neurobiology. 1. Anatomy and connectivity. *Brain Research*, *434*(2), 117–165.
- Ogawa, S., Lee, T. M., Kay, A. R., & Tank, D. W. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *PNAS*, *87*(24), 9868–9872.
- Parker, N. F., Cameron, C. M., Taliaferro, J. P., Lee, J., Choi, J. Y., Davidson, T. J., Witten, I. B. (2016). Reward and choice encoding in terminals of midbrain dopamine neurons depends on striatal target. *Nature Neuroscience*, *19*(6), 845–854. <https://doi.org/10.1038/nn.4287>
- Pessiglione, M., Seymour, B., Flandin, G., Dolan, R. J., & Frith, C. D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*, *442*(7106), 1042–1045. <https://doi.org/10.1038/nature05051>

- Posse, S., Fitzgerald, D., Gao, K., Habel, U., Rosenberg, D., Moore, G. J., & Schneider, F. (2003). Real-time fMRI of temporolimbic regions detects amygdala activation during single-trial self-induced sadness. *NeuroImage*, *18*(3), 760–768.
- Qi, J., Zhang, S., Wang, H.-L., Barker, D. J., Miranda-Barrientos, J., & Morales, M. (2016). VTA glutamatergic inputs to nucleus accumbens drive aversion by acting on GABAergic interneurons. *Nature Neuroscience*, *19*(5), 725–733. <https://doi.org/10.1038/nn.4281>
- Raichle, M. E., & Shepherd, G. M. (2014). *Angelo Mosso's Circulation of Blood in the Human Brain*. Retrieved from <http://oxfordmedicine.com/view/10.1093/med/9780199358984.001.0001/med-9780199358984>
- Rigoli, F., Chew, B., Dayan, P., & Dolan, R. J. (2016). The dopaminergic midbrain mediates an effect of average reward on pavlovian vigor. *Journal of Cognitive Neuroscience*, *28*(9), 1303–1317. https://doi.org/10.1162/jocn_a_00972
- Rigoli, F., Rutledge, R. B., Chew, B., Ousdal, O. T., Dayan, P., & Dolan, R. J. (2016). Dopamine increases a value-independent gambling propensity. *Neuropsychopharmacology*, *41*(11), 2658–2667. <https://doi.org/10.1038/npp.2016.68>
- Rorden, C., & Brett, M. (2000). Stereotaxic display of brain lesions. *Behavioural Neurology*, *12*(4), 191–200.
- Roy, C. S., & Sherrington, C. S. (1890). On the regulation of the blood-supply of the brain. *The Journal of Physiology*, *11*(1–2), 85-158.17.
- Russo, S. J., & Nestler, E. J. (2013). The brain reward circuitry in mood disorders. *Nature Reviews Neuroscience*, *14*(9), 609–625. <https://doi.org/10.1038/nrn3381>
- Rutledge, R. B., Skandali, N., Dayan, P., & Dolan, R. J. (2015). Dopaminergic modulation of decision making and subjective well-being. *Journal of Neuroscience*, *35*(27), 9811–9822. <https://doi.org/10.1523/JNEUROSCI.0702-15.2015>
- Sandrone, S., Bacigaluppi, M., Galloni, M. R., Cappa, S. F., Moro, A., Catani, M., Martino, G. (2014). Weighing brain activity with the balance: Angelo Mosso's original manuscripts come to light. *Brain*, *137*(2), 621–633. <https://doi.org/10.1093/brain/awt091>
- Sara, S. J. (2009). The locus coeruleus and noradrenergic modulation of cognition. *Nature Reviews Neuroscience*, *10*(3), 211–223. <https://doi.org/10.1038/nrn2573>
- Schmierer, K., Scaravilli, F., Altmann, D. R., Barker, G. J., & Miller, D. H. (2004). Magnetization transfer ratio and myelin in postmortem multiple sclerosis brain. *Annals of Neurology*, *56*(3), 407–415. <https://doi.org/10.1002/ana.20202>

- Schott, B. H., Minuzzi, L., Krebs, R. M., Elmenhorst, D., Lang, M., Winz, O. H., Bauer, A. (2008). Mesolimbic functional magnetic resonance imaging activations during reward anticipation correlate with reward-related ventral striatal dopamine release. *Journal of Neuroscience*, *28*(52), 14311–14319. <https://doi.org/10.1523/JNEUROSCI.2058-08.2008>
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*(5306), 1593–1599.
- Steiger, T. K., Weiskopf, N., & Bunzeck, N. (2016). Iron level and myelin content in the ventral striatum predict memory performance in the aging brain. *Journal of Neuroscience*, *36*(12), 3552–3558. <https://doi.org/10.1523/JNEUROSCI.3617-15.2016>
- Sulzer, J., Haller, S., Scharnowski, F., Weiskopf, N., Birbaumer, N., Blefari, M. L., Sitaram, R. (2013). Real-time fMRI neurofeedback: Progress and challenges. *NeuroImage*, *76*, 386–399. <https://doi.org/10.1016/j.neuroimage.2013.03.033>
- Viswanathan, A., & Freeman, R. D. (2007). Neurometabolic coupling in cerebral cortex reflects synaptic more than spiking activity. *Nature Neuroscience*, *10*(10), 1308–1312. <https://doi.org/10.1038/nn1977>
- Volkow, N. D., Fowler, J. S., Gatley, S. J., Logan, J., Wang, G. J., Ding, Y. S., & Dewey, S. (1996). PET evaluation of the dopamine system of the human brain. *Journal of Nuclear Medicine*, *37*(7), 1242–1256.
- Watabe-Uchida, M., Zhu, L., Ogawa, S. K., Vamanrao, A., & Uchida, N. (2012). Whole-brain mapping of direct inputs to midbrain dopamine neurons. *Neuron*, *74*(5), 858–873. <https://doi.org/10.1016/j.neuron.2012.03.017>
- Weiskopf, N., Hutton, C., Josephs, O., & Deichmann, R. (2006). Optimal EPI parameters for reduction of susceptibility-induced BOLD sensitivity losses: A whole-brain analysis at 3 T and 1.5 T. *NeuroImage*, *33*(2), 493–504. <https://doi.org/10.1016/j.neuroimage.2006.07.029>
- Weiskopf, N., Mathiak, K., Bock, S. W., Scharnowski, F., Veit, R., Grodd, W., Birbaumer, N. (2004). Principles of a brain-computer interface (BCI) based on real-time functional magnetic resonance imaging (fMRI). *IEEE Transactions on Bio-Medical Engineering*, *51*(6), 966–970. <https://doi.org/10.1109/TBME.2004.827063>
- Weiskopf, N., Mohammadi, S., Lutti, A., & Callaghan, M. F. (2015). Advances in MRI-based computational neuroanatomy: From morphometry to in-vivo histology. *Current Opinion in Neurology*, *28*(4), 313–322. <https://doi.org/10.1097/WCO.0000000000000222>

- Weiskopf, N., Suckling, J., Williams, G., Correia, M. M., Inkster, B., Tait, R., Lutti, A. (2013). Quantitative multi-parameter mapping of R1, PD*, MT, and R2* at 3T: A multi-center validation. *Frontiers in Neuroscience*, 7. <https://doi.org/10.3389/fnins.2013.00095>
- Westbrook, A., & Braver, T. S. (2016). Dopamine does double duty in motivating cognitive effort. *Neuron*, 89(4), 695–710. <https://doi.org/10.1016/j.neuron.2015.12.029>
- Williams, G. V., & Goldman-Rakic, P. S. (1995). Modulation of memory fields by dopamine D1 receptors in prefrontal cortex. *Nature*, 376(6541), 572–575. <https://doi.org/10.1038/376572a0>
- Wolff, S. D., & Balaban, R. S. (1989). Magnetization transfer contrast (MTC) and tissue water proton relaxation in vivo. *Magnetic Resonance in Medicine*, 10(1), 135–144.
- Yacoub, E., & Hu, X. (1999). Detection of the early negative response in fMRI at 1.5 Tesla. *Magnetic Resonance in Medicine*, 41(6), 1088–1092.
- Yoo, S.-S., & Jolesz, F. A. (2002). Functional MRI for neurofeedback: Feasibility study on a hand motor task. *Neuroreport*, 13(11), 1377–1381.
- Zhang, Z., Andersen, A. H., Ai, Y., Loveland, A., Hardy, P. A., Gerhardt, G. A., & Gash, D. M. (2006). Assessing nigrostriatal dysfunctions by pharmacological MRI in parkinsonian rhesus macaques. *NeuroImage*, 33(2), 636–643. <https://doi.org/10.1016/j.neuroimage.2006.07.004>
- Ziegler, G., Hauser, T., Moutoussis, M., Bullmore, E. T., Goodyer, I. M., Fonagy, P., Dolan, R. J. (2018). Compulsivity and impulsivity are linked to distinct aberrant developmental trajectories of fronto-striatal myelination. *BioRxiv*, 328146. <https://doi.org/10.1101/328146>
- Zilverstand, A., Sorger, B., Zimmermann, J., Kaas, A., & Goebel, R. (2014). Windowed correlation: A suitable tool for providing dynamic fMRI-based functional connectivity neurofeedback on task difficulty. *PLoS ONE*, 9(1), e85929. <https://doi.org/10.1371/journal.pone.0085929>

Chapter 4: Endogenous fluctuations in the dopaminergic midbrain drive behavioural choice variability

To know one's own state is not a simple matter. One cannot look directly at one's own face with one's own eyes, for example. One has no choice but to look at one's reflection in the mirror. Through experience, we come to believe that the image is correct, but that is all.

Haruki Murakami, *The Wind-Up Bird Chronicle*

Human behaviour is inherently variable. Even when facing the same task repeatedly, humans often act in inconsistent ways. This observation led the English poet Horace Smith to suggest that, "inconsistency is the only thing in which men are consistent." Inconsistencies in value-based decision making often violate the tenets of rational economic theory. Many economic models explain this variability by injecting stochasticity into subjective preferences (Harless & Camerer, 1994).

The human brain shows substantial regional activity fluctuations in the absence of external stimulation (i.e., resting-state) (Fox & Raichle, 2007; Tavor et al., 2016). The functional role of these fluctuations is not well understood. Endogenous fluctuations endure when participants perform externally imposed tasks and can explain neural variability in task-evoked responses (Fox, Snyder, Zacks, & Raichle, 2006). Studies investigating low-level cognitive processes have shown that endogenous fluctuations also influence how stimuli are processed. Endogenous fluctuations in task-relevant areas influence perception of auditory (Sadaghiani, Poline, Kleinschmidt, & D'Esposito, 2015) and somatosensory stimuli (Boly et al., 2007) and can influence the force exerted during simple motor actions, such as button presses (Fox, Snyder, Vincent, & Raichle, 2007). However, it remains unknown whether intrinsic fluctuations also affect complex cognitive processes, such as decision making, and whether variability in pre-stimulus brain activity can predict future decisions.

In this study, we hypothesized that endogenous fluctuations in areas implicated in decision making would explain variability in choice. In particular, we hypothesized that endogenous fluctuations in the dopaminergic midbrain, encompassing substantia nigra and ventral tegmental area (SN/VTA), play a key role in decision making under risk. SN/VTA contains the largest assembly of dopamine neurons in the human brain and is centrally involved in decision making (Starkweather, Babayan, Uchida, &

Gershman, 2017; Steinberg et al., 2013). Modulating dopamine neurotransmission can increase risk taking (Burke et al., 2018; Rigoli et al., 2016; Rutledge, Skandali, Dayan, & Dolan, 2015) and dopamine dysfunction is strongly linked to problem gambling and impulsive behaviors (Buckholtz et al., 2010). Although it is not possible to directly assess dopaminergic activity using fMRI, dopamine-related quantities such as reward prediction errors (Lak, Stauffer, & Schultz, 2014) are observed in BOLD activity within the SN/VTA (D'Ardenne, McClure, Nystrom, & Cohen, 2008; Hauser, Eldar, & Dolan, 2017).

To test our hypothesis, we developed a novel real-time fMRI framework to trigger presentation of options based on intrinsic fluctuations of blood oxygenation level-dependent (BOLD) activity in the SN/VTA (Fig. 4.1). We developed an algorithm that detected epochs of very high and very low activity, providing a trigger to probe participants with a matched set of choices between a safe and a risky option in these two background brain states (Rigoli et al., 2016; Rutledge et al., 2015). The risky option comprised equal probabilities of a prize (£6, £9, or £12) or £0. The value of the safe option was always lower than the potential prize from the risky option and varied systematically around each subject's economic indifference point, the offer for which a subject chooses safe and risky options in equal proportion. Safe option values were determined from pre-scanning decisions from an extensive choice set (see Methods).

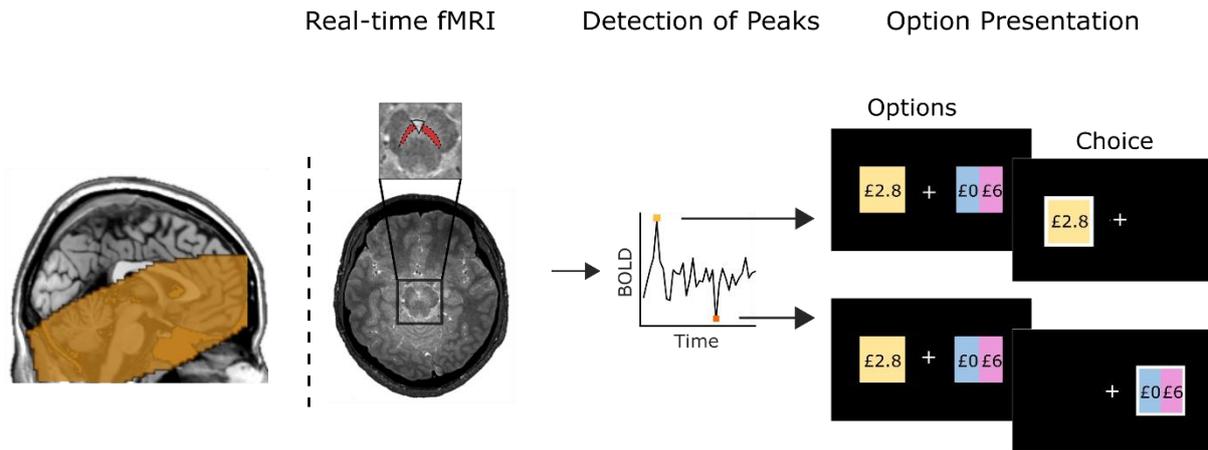


Fig. 4.1. Schematic of real-time fMRI setup. BOLD activity from anatomically defined SN/VTA is extracted and denoised (removing movement, breathing, and pulsatile artifacts) in real-time. The overlay on the sagittal image indicates intersecting coverage across all subjects in the study. Endogenous activity reflecting a low/high background activity state (exceeding a 15th/85th percentile cut-off) triggered presentation of a trial with a choice between a safe option (here, £2.8 guaranteed reward) and a risky option (here, £0 or £6 with equal probability). To ensure similar rates of risk taking across individuals, safe options varied around each subject's indifference point which was determined prior to scanning. This design allowed us to efficiently and selectively probe subjects with identical options during very low and very high endogenous SN/VTA activity. Any difference in behaviour can therefore be attributed to endogenous SN/VTA activity.

4.1 Methods

4.1.1 Participants

49 healthy, young adults (age: 25.2 ± 4.2 , mean \pm SD) were recruited through the University College London (UCL) Psychology Subject Database. Participants were screened to ensure no history of neurological or psychiatric disorders. Six participants were excluded from analyses: 3 participants because of excessive number of missed trials (>20) and 3 due to frequent large head movements (>3mm). A total of 43 participants (Group 1: 10 females, 2 males; Group 2: 21 females, 10 males) were included. Participants in both groups went through identical procedures with the only difference being that the range of values for the safe options, drawn around each subject's indifference points, was wider for Group 2 than Group 1, allowing us to better distinguish between competing computational models (see **Procedure**). The study was approved by the UCL research ethics committee, and all participants gave written informed consent.

4.1.2 Procedure

Our study protocol spanned two sessions approximately 24 hours apart. On the first day, we assessed gambling behaviour and collected structural brain scans. These scans were used to define individualized anatomical masks of the dopaminergic midbrain for use in the following session. On the second day, decision making was reassessed before participants participated in the real-time fMRI experiment reported.

Day 1

Probabilistic Gambling Task. Participants first played a probabilistic gambling task consisting of 180 trials. On each trial, participants chose between a certain monetary amount and a gamble with equal probabilities of two outcomes. There were three gamble options available: £0 and £6, £9, or £12. The certain amounts were determined using 12 divisors (0.82, 0.87, 0.93, 1, 1.1, 1.23, 1.4, 1.6, 1.9, 2.25, 2.75, 3.5) on the expected value of the gambles, chosen to accommodate a wide range of risk sensitivity. Take for example a fraction of 3.5 and a gamble between £0 and £6. The expected value of a £0 or £6 gamble is £3 ($0.5 \times £0 + 0.5 \times £6$), which divided by 3.5 gives a certain amount of £0.86. There were 12 certain amounts for each gamble option in total, and each trial was repeated 5 times in a randomized sequence.

Structural Scans. Multi-Parameter Maps were acquired for each subject (Weiskopf et al., 2013). The magnetization-transfer (MT) saturation image was used for the drawing of the region-of-interest (SN/VTA) due to its ability to delineate grey and white matter in subcortical/brainstem regions, in line with preceding studies (Hauser et al., 2017; Koster, Guitart-Masip, Dolan, & Düzel, 2015).

Day 2

Prior to the real-time fMRI session, participants completed a shorter version of the probabilistic gambling task consisting of 108 trials to recalibrate the participants' indifference points. The only difference between this task and the task on Day 1 was that each trial was repeated 3 instead of 5 times.

Probabilistic Gambling Task Inside the MRI Scanner. Choice behaviour across both days was fitted to a prospect theory-based parametric decision model that has been used in past studies (Rutledge, Skandali, Dayan, & Dolan, 2014; Sokol-Hessner

et al., 2009) to describe decision-making under risk. The expected utility of the certain options and gambles were determined using the following equations:

$$U_{\text{gamble}} = 0.5(V_{\text{gain}})^{\alpha}$$

$$U_{\text{certain}} = (V_{\text{certain}})^{\alpha}$$

where V_{gain} is the value of the potential gain from a gamble and V_{certain} is the value of the certain option. α alters the degree of curvature of the utility function and represents the degree of risk aversion. When presented with an option where the expected values for the certain gain and the gamble are equal, a subject with $\alpha = 1$ would be risk-neutral and indifferent between the two, a risk-seeking individual with $\alpha > 1$ would choose the gamble more often, and a risk-averse individual with $\alpha < 1$ would choose the certain gain more often. The probability of selecting a gamble was determined by the following softmax rule:

$$P_{\text{gamble}} = \frac{1}{1 + e^{-\mu(U_{\text{gamble}} - U_{\text{certain}})}}$$

where the degree of stochasticity in choice behaviour is captured by the inverse temperature parameter μ . When μ is low, participants are more likely to choose randomly between safe and risky options irrespective of their subjective values. When μ is high, participants increasingly choose the action leading to the highest expected reward. Expected utilities for the certain option were sampled evenly (5 bins) between $P_{\text{gamble}} = 0.3$ and 0.7 for each gamble level for the first group of participants, and $P_{\text{gamble}} = 0.1$ and 0.9 for the second group of participants. These utilities were then converted back to objective values and used as the safe options in the real-time fMRI session. The real-time fMRI task consisted of 90 trials in total with 30 trials for each gamble level (£0 and £6, £9, or £12) of which 15 trials were allocated to the low baseline condition and 15 trials allocated to the high baseline condition according to criteria defined in the following section.

4.1.3 Real-time fMRI

Software and Preprocessing of Images. Real-time preprocessing of the functional data was performed using Turbo-BrainVoyager (TBV) (Brain Innovation, Maastricht, The Netherlands) and custom scripts. Time courses for every voxel within the SN/VTA ROI were extracted from smoothed and realigned images [6mm Full Width at Half Maximum (FWHM)] and exported using TBV. Exported data were then corrected for additional noise sources (movement and physiological noise; cf below). Physiological noise arising from breathing and pulsatile artifacts were modeled using a Fourier expansion of physiological phases based on the RETROICOR model (Glover, Li, & Ress, 2000) and respiratory volume (Birn, Diamond, Smith, & Bandettini, 2006) were incrementally regressed out in real time from the exported time courses using a custom-made MATLAB (MathWorks, Natick, USA) toolbox. The ensuing filtered time courses were then analysed to detect endogenous fluctuations.

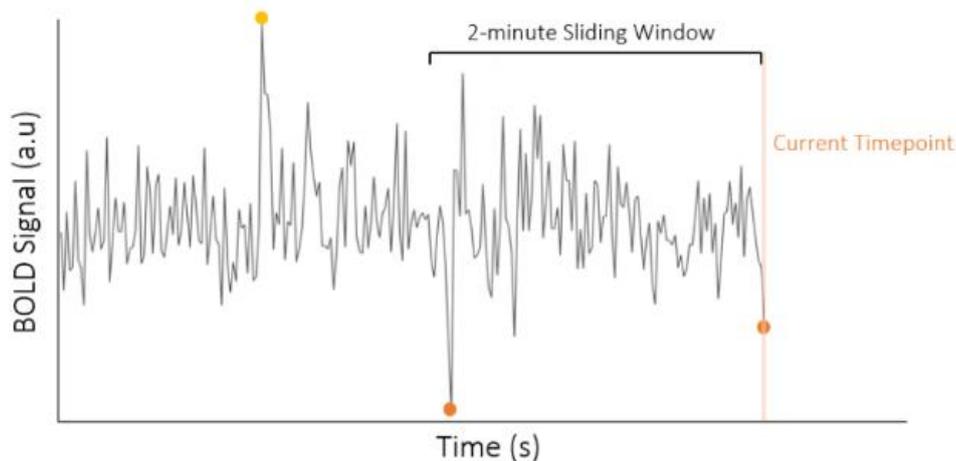


Fig 4.2. Sliding window approach to quantifying BOLD activity levels. Endogenous activity reflecting a low/high background BOLD activity state (exceeding a 15th/85th percentile cut-off based on an incremental sliding window of 2 minutes) triggered presentation of a trial involving a choice between a safe option and a risky option. Trials presented in both low/high background activity states were matched and thus any difference in behaviour can be attributed to distinct levels of endogenous SN/VTA BOLD activity.

Quantifying the Level of BOLD Activity

We used a sliding window approach to quantify endogenous activation of the SN/VTA over the course of the experiment (Fig. 4.2). This measure not only takes

scanner-induced and other slow signal drifts (e.g., due to a warming of the gradient coils) into consideration but is also robust to outlier activations and can account for changes in the variance of the signal over time. A normal cumulative distribution function was used to quantify the distribution of BOLD signal within an ongoing sliding window consisting of 69 volumes (approximately 2 minutes). The mean of the most recent 2 volumes was compared to the previous 69 volumes over the progression of the entire experiment. The distribution of the sliding window was updated with each new volume acquired. Thresholds for the trials were set below the 15th percentile for low baseline trials and above the 85th percentile for high baseline trials. When BOLD activity reaching the thresholds was detected, a trial was immediately presented. There was a minimum inter-trial interval of 20s to allow the hemodynamic response for each trial to return close to baseline. If threshold criteria were not met by 55s, a trial was presented and categorized as low or high depending on whether it was lower or higher than the mean of the preceding baseline, respectively. This procedure was applied to the $15.1 \pm 5.8\%$ (mean \pm SD) of trials that did not reach the threshold criteria.

Image Acquisition.

MRI data were acquired at the Wellcome Centre for Human Neuroimaging at UCL, using a Siemens Trio 3-Tesla scanner equipped with a 32-channel head coil. A partial-volume 2D echo-planar imaging (EPI) sequence that was optimized for striatal, medial prefrontal, and brainstem regions was selected for the functional images. Each volume consisted of 25 slices with 2.5mm isotropic voxels [repetition time (TR): 1.75s; echo time (TE): 30ms; slice tilt: -30°]. At the beginning of each functional session, 10 EPI volumes were acquired with the 10th volume selected as the template used to co-register the ROI. Field maps with 3mm isotropic voxels (whole brain coverage) were also acquired to correct the EPIs for any inhomogeneity in magnetic field strength. Subsequently, the first 6 volumes of each run were discarded to allow for T1 saturation effects. Sequence settings were identical across participants (e.g., no variation in tilt angle) and no slices were discarded. Overlapping coverage across all participants is indicated in orange in Figure 4.1.

Structural images consisted of 3 spoiled multi-echo 3D fast low angle shot (FLASH) acquisitions at 0.8mm isotropic resolution with T1 (TR: 18.7ms; flip angle: 20°), proton density (PD) (TR: 23.7ms; flip angle: 6°), and magnetization transfer (MT)

(TR: 23.7ms; flip angle: 6°; excitation preceded by a 2kHz off-resonance Gaussian radiofrequency (RF) pulse with 4ms duration and 200° nominal flip angle) weightings. Additional B1 mapping and field maps were acquired to get calibration data measuring the spatial distribution of the B1+ transmit field in order to detect the spatial variation in flip angle.

ROI Definition and Transformation.

Bright areas in MT-contrast images have been shown to be coextensive with the SN as delineated histologically by tyrosine hydroxylase immunohistochemistry, which stains dopaminergic neurons (Bolding et al., 2013) that are the key component of SN/VTA. Leveraging upon this, SN/VTA ROIs were hand-drawn for each individual in MRIcron (Rorden & Brett, 2000) using MT-weighted structural images. In accordance with procedures outlined previously (Chowdhury, Guitart-Masip, Lambert, Dolan, & Düzel, 2013), medial and lateral boundaries of the SN/VTA ROI were defined based on the change in contrast between its bright grey colour and the dark grey colour of the adjacent cerebral peduncle and interpeduncular fossa. Lower and upper boundaries of the ROI were selected as the slices preceding the ones where the intensity of SN/VTA was indistinguishable from surrounding tissue, totalling between 6 to 9 slices contingent on individual SN/VTA size differences. To prepare the hand-drawn SN/VTA ROI for use in TBV, it needs to be co-registered and transformed to the space and resolution of the EPIs. Co-registration was carried out using a single EPI volume as the reference image, and the individual-specific T1-weighted image as the source image. Following this, the EPI voxels corresponding to each ROI voxel were indexed based on Euclidean distance calculated in native space. Since the coordinate space in TBV differs from more common ones such as the Montreal Neurological Institute (MNI) space, coordinates for the ROI were transformed before use in TBV. This series of co-registration and transformations was executed using custom MATLAB scripts available on Github (<https://github.com/tuhauser/rffMRI>).

4.1.4 Offline Analyses

Images were preprocessed using standard procedures in SPM 12 (Wellcome Centre for Human Neuroimaging, UCL). This consisted of unwarping EPIs using field

maps, motion correction, spatial transformation to the MNI template, and spatial smoothing with a 6-mm full-width at half-maximum Gaussian kernel.

4.2 Results

4.2.1 Endogenous fluctuations in SN/VTA BOLD activity modulate risk taking

We first asked how the two modes of endogenous SN/VTA activity (low and high) influenced choice behaviour. On average, participants chose the risky option more when pre-stimulus SN/VTA activity was low compared to when it was high (low activity: $59.6 \pm 1.5\%$ (mean \pm SEM), high activity: $56.2 \pm 1.8\%$, $t_{42} = 3.83$, $P < 0.001$, Fig. 4.3A). This effect of greater risk taking following low compared to high activity was present in 30 of 43 participants (Fig. 4.3B).

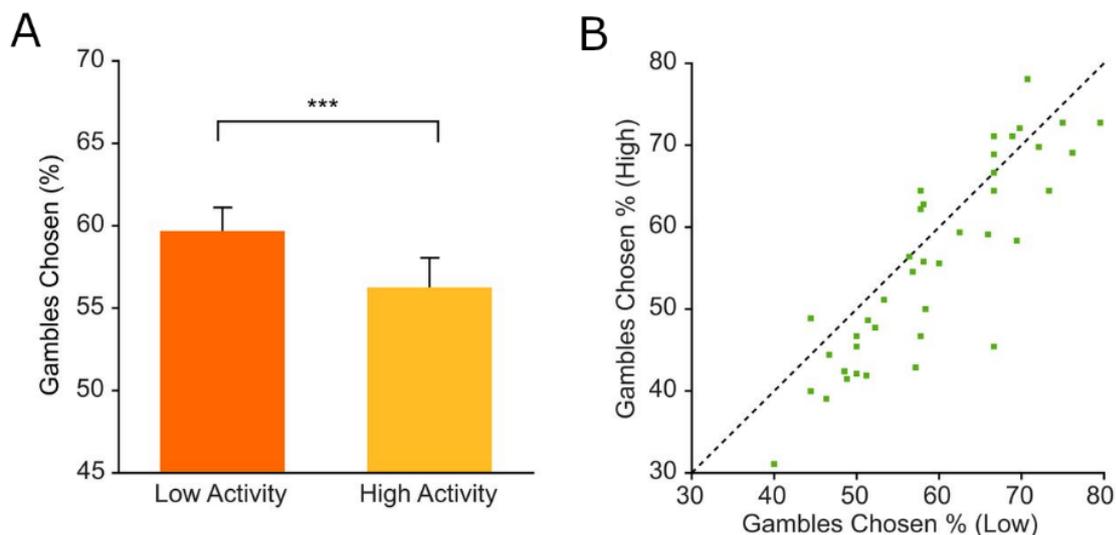


Fig. 4.3. Endogenous fluctuations in SN/VTA BOLD activity bias behavioural choice. A, Subjects ($n = 43$) gambled more when options were presented against a background of low compared to high endogenous SN/VTA activity. **B,** This effect of greater risk taking for low than high activity was consistent across subjects. *** $P < 0.001$. Data are mean \pm SEM.

Control Analysis

To validate the results obtained from our online procedure and to examine whether the effect of endogenous BOLD activity on risky choice behaviour was a

general property across the brain, we sampled activity from multiple regions. The ROI for vmPFC was derived from www.neurosynth.org, the VS ROI was bilateral 8-mm spheres at MNI coordinates derived from a previous study (Rutledge et al., 2014), a group anatomical mask from a previous study (Hauser et al., 2017) was used for SN/VTA ROI, and the primary auditory cortex (A1) was Brodmann Areas 41 from the Wake Forest University PickAtlas toolbox for SPM (Maldjian, Laurienti, & Burdette, 2004).

BOLD time courses for these ROIs were extracted and filtered using an incremental GLM with the same motion and physiological regressors as in the real-time fMRI experiment. Based on our real-time procedure, BOLD activity for each region was averaged for the 2 most recent TRs prior to trial presentation and compared against each preceding baseline window of 2 minutes. As our design was optimized to detect activity fluctuations in the SN/VTA, the threshold used to categorize trials as low or high activity in the SN/VTA would be overly conservative when applied to other brain regions. This would lead to many trials being left uncategorized. To ensure that all trials were categorized, we relaxed the threshold and categorized each trial as low or high depending on whether pre-trial BOLD activity for each of these regions was lower or higher than the mean of the preceding baseline period.

To test whether our main effect of risk preference change is specific to SN/VTA BOLD activity, we investigated the relationship between endogenous fluctuations of BOLD activity in other brain regions and risky choice. We conducted offline analyses on A1 as a control area, as well as VS and vmPFC, which are regions strongly implicated in value-based decision making (Bartra, McGuire, & Kable, 2013). We used independent ROIs for all areas including SN/VTA and re-categorized trials based on endogenous activity in each of these ROIs.

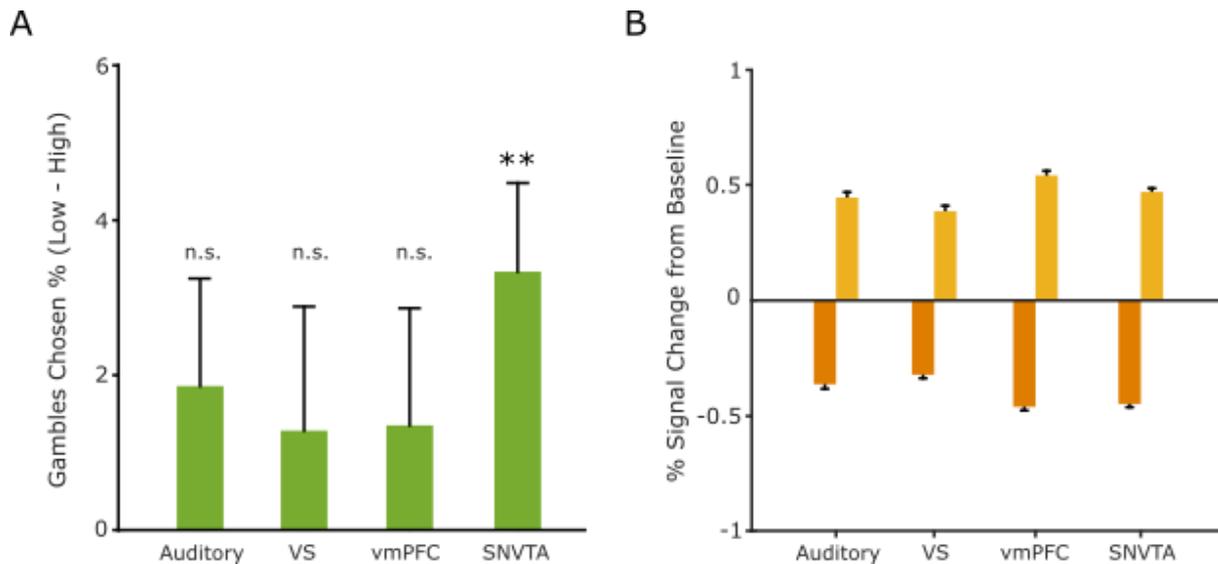


Fig 4.4. Control analyses. **A**, Offline re-classification of endogenous activity using independent ROIs (see Methods) revealed that only endogenous dopaminergic midbrain fluctuations were significantly associated with choice variability. **B**, Differences in signal change between low and high activity conditions were largest in vmPFC and smallest in VS (yellow: trials categorised as high activity, orange: signals categorised as low activity). This suggests that endogenous fluctuations in SN/VTA were not more extreme than other regions. ** $P < 0.01$. Data are mean \pm SEM.

Risky choice behaviour was significantly greater for low compared to high baseline activity in the independent SN/VTA ROI (low baseline activity: $59.7 \pm 1.5\%$ (mean \pm SEM), high baseline activity: $56.0 \pm 1.9\%$, $t_{42} = 2.92$, $P = 0.003$). There was no significant relationship between risk taking and endogenous activity in any of the control and decision-related areas tested (Fig. 4.4A). Risk taking was similar for low and high baseline activity in VS (low: $58.6 \pm 1.7\%$, high: $57.0 \pm 1.8\%$, $t_{42} = 0.95$, $P = 0.35$), vmPFC (low: $58.6 \pm 1.6\%$, high: $57.1 \pm 1.9\%$, $t_{42} = 1.01$, $P = 0.32$), and A1 (low: $58.9 \pm 1.6\%$, high: $57.0 \pm 1.8\%$, $t_{42} = 1.43$, $P = 0.16$). To further verify that the effects we observe are driven by local rather than global fluctuations, we tested whether SN/VTA activity was still predictive of risk taking even after controlling for activity in control area A1 ($t_{42} = 2.34$, $P = 0.02$). These findings suggest that the effect is not a general effect of low and high BOLD activity modes across the brain, but specific to local fluctuations in the dopaminergic midbrain that explain variability in risk taking.

A caveat of the above analysis is that the absence of any effect in a control area could be due to reduced endogenous signal variability. To rule out this alternative explanation, we calculated the signal change of epochs used to trigger each trial

relative to their preceding baselines. Differences in signal change between low and high activity conditions were largest in vmPFC and smallest in VS, suggesting that activity used to trigger trials in SN/VTA was no more extreme than that observed in other regions, supporting our finding of a specific effect of SN/VTA endogenous fluctuations on risk taking (Fig. 4.4B).

To determine whether the VS results were affected by partial volume effects due to its location and the image acquisition parameters, we re-ran the preprocessing steps and re-analysed the data after discarding the top and bottom slices of the partial volumes. We found that risk taking was still similar for low and high baseline activity in VS (low: $58.7 \pm 1.7\%$, high: $57.0 \pm 1.9\%$, $t_{42} = 1.02$, $P = 0.32$), suggesting that the absence of an association between VS BOLD activity and risk taking was not due to partial volume effects.

As SN/VTA BOLD signals recorded in real-time may be contaminated by signals from surrounding structures due to smoothing, we also performed offline analyses on unsmoothed functional images using the same algorithm to reclassify pre-stimulus activity and found consistent results in unsmoothed data. Risk taking was higher for trials presented against a background of low compared to high SN/VTA BOLD activity (low activity: $59.9 \pm 1.8\%$, high activity: $55.6 \pm 2.1\%$, $t_{42} = 3.2$, $P = 0.003$).

To test how sensitive the effect we observe is to the timing of pre-stimulus activity, we reanalysed the data, reclassifying activity levels as high or low based on volumes t-2 and t-3 before trial onset (instead of t-1 and t-2). Discarding the final volume of SN/VTA signal before trial onset did not affect the relationship between pre-stimulus activity and risk taking ($t_{42} = 2.95$, $P = 0.005$), suggesting that the effect we observe does not depend on the precise timing of option presentation.

These control analyses suggest that the effect of greater risk taking following low compared to high activity is unaffected by the precise timing of option presentation or degree of smoothing and is specific to SN/VTA with no effect in other decision or control areas.

4.2.2 A computational mechanism for the effect of endogenous fluctuations on risk taking

We next examined how endogenous fluctuations in SN/VTA BOLD activity influenced risk taking and tested whether the effect was specific to a certain set of offers. We computed the difference between the average return for risky and safe options and identified a main effect of this value difference, indicating that increased value for risky relative to safe options was associated with an increased propensity to choose the risky option ($F(2.822, 118.505) = 107.580, P < 0.001$, Fig. 4.5A). We found a main effect of endogenous fluctuations in SN/VTA activity on risk taking ($F(1,42) = 14.356, P < 0.001$, Fig. 4.5A) but no interaction with value difference ($F(3.113, 130.749) = 0.127, P = 0.95$), indicating that low SN/VTA activity is associated with greater risk taking irrespective of how much risky and safe options differed in value. We also found no interaction between risky option value (i.e., £6, £9, or £12) and activity ($F(1.957, 82.208) = 0.493, P = 0.61$), further supporting an association between low endogenous SN/VTA activity and a value-independent increase in risk taking.

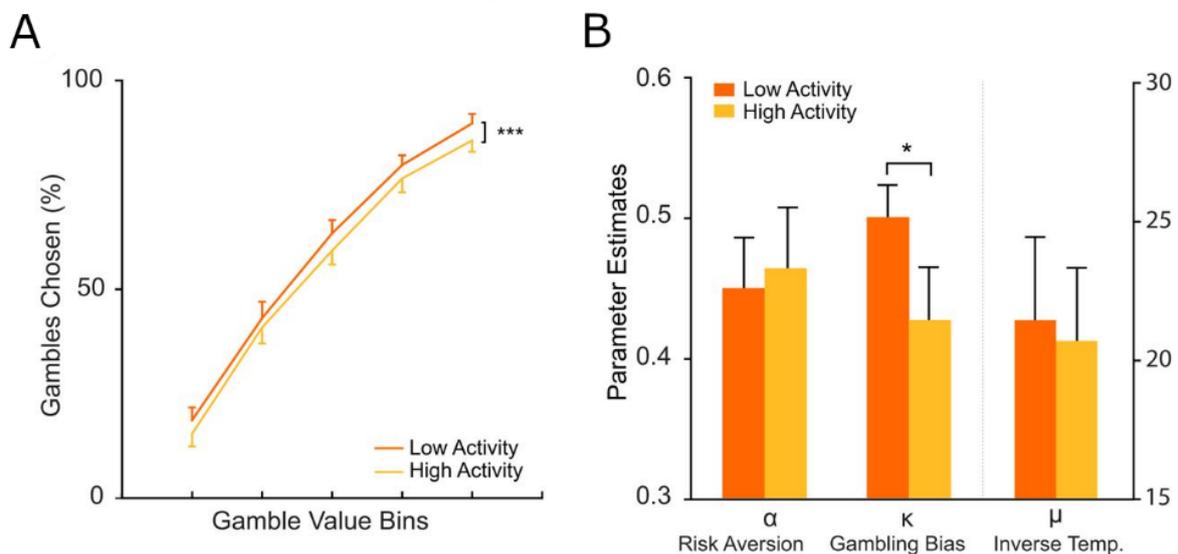


Fig 4.5. Endogenous fluctuations in SN/VTA BOLD activity modulate value-independent influences on choice. **A**, The activity-induced shift in risk taking was independent of value with low endogenous activity leading to increased risk taking irrespective of option value. Differences in objective value between risky and safe options were divided into bins of equal sizes for each subject. **B**, Choices were fitted to a parametric decision model based on prospect theory with the best-fitting model including a gambling bias parameter that was higher when endogenous activity was low. Positive gambling bias parameters reflects a tendency to take risks irrespective of option value. *** $P < 0.001$. Data are mean \pm SEM.

We next asked whether endogenous SN/VTA BOLD activity influenced option valuation in a manner consistent with standard economic models. We tested the following four models of choice.

Parametric Decision Model Based on Prospect Theory

Details of this model are provided in the previous section (*Procedure, Day 2*). This model provided a good fit for choice behaviour in both low- and high-activity conditions with an average pseudo- R^2 of 0.44 (SD: 0.15).

Parametric Approach-Avoidance Decision Model

A recent model (Rutledge et al., 2015) that was developed to account for value-independent tendencies to choose gambles is the approach-avoidance model, which allows choice probabilities to differ from 0 or 1 in the limit when a softmax rule is used. Expected utilities were determined using equations in the prospect theory model described earlier.

The main difference lies in the softmax rule where the probability of gambling depended on a new parameter, β , determined by the following equations:

$$P_{\text{gamble}} = \frac{1 - \beta}{1 + e^{-\mu(U_{\text{gamble}} - U_{\text{certain}})}} + \beta \text{ if } \beta \geq 0$$

$$P_{\text{gamble}} = \frac{1 + \beta}{1 + e^{-\mu(U_{\text{gamble}} - U_{\text{certain}})}} \text{ if } \beta < 0$$

If β is positive, choice probabilities are mapped from $(\beta, 1)$. If β is negative, choice probabilities are mapped from $(0, 1 + \beta)$. This model also provided a good fit of behaviour with an average pseudo- R^2 of 0.47 (SD: 0.14).

Parametric Decision Model Based on Prospect Theory with Gambling Bias

To account for the possibility of a shift in indifference points leading to a difference in tendencies to choose gambles, the softmax rule in the parametric prospect theory model included an additional parameter, κ , such that:

$$P_{\text{gamble}} = \frac{1}{1 + e^{-\mu(U_{\text{gamble}} - U_{\text{certain}} + \kappa)}}$$

κ here represents a gambling bias that is additive to the expected utilities. This model provided the best fit out of all the models tested with a pseudo- R^2 of 0.55 (SD: 0.12). Model comparison based on Bayesian Information Criterion (BIC) confirmed the fit and revealed that this model fitted the data best (Table 4.1). Larger effects of endogenous SN/VTA fluctuations on task-evoked SN/VTA responses (as measured using the average of an epoch corresponding to 5.25s to 10.5s in Fig. 4.8A later) correlated with larger increases in gambling bias parameter.

Parametric Decision Model Using Expected Values

The final model tested was one that used the expected values of the gamble (E_{gamble}) and certain gain (E_{certain}) and passed through the following softmax with the same gambling bias term β as before:

$$P_{\text{gamble}} = \frac{1}{1 + e^{-\mu(E_{\text{gamble}} - E_{\text{certain}} + \kappa)}}$$

This model had the lowest fit with a pseudo- R^2 of 0.36 (SD: 0.17), which suggests that more of the variance could be accounted for by the inclusion of a risk aversion parameter to convert objective values into subjective values.

Model comparison revealed that a parametric model based on prospect theory (Sokol-Hessner et al., 2009) provided a good description of behaviour (pseudo- $R^2 = 0.44 \pm 0.15$) but was outperformed by a model (Brown et al., 2014; Timmer, Sescousse, Esselink, Piray, & Cools, 2017) that included a gambling bias parameter (pseudo- $R^2 = 0.55 \pm 0.12$; Table 4.1). Changes in this gambling bias parameter κ shift the sigmoidal decision function in standard models, capturing a propensity to take risks irrespective of offer value. This gambling bias parameter was significantly higher in low compared to high activity conditions ($t_{30} = 2.21$, $P = 0.04$). No differences were observed for other model parameters (risk aversion α : $t_{30} = -0.5$, $P = 0.62$, inverse temperature μ : $t_{30} = 0.13$, $P = 0.9$; Fig. 4.5B). This finding suggests that endogenous SN/VTA activity does not impact the valuation process in a value-dependent way, but instead influences a more general decision process that does not depend on the relative values of available options.

Table 4.1. Model Comparison Results.

Model	Parameters	Mean R ²	BIC	ΔBIC
Prospect Theory	4	0.44	2636	130
Approach-Avoidance	6	0.47	2804	298
Prospect Theory with Gambling Bias	6	0.55	2506	0
Expected Values with Gambling Bias	4	0.36	2927	421

BIC measures are summed across 31 subjects. The winning model (lowest BIC) here was the parametric decision model based on prospect theory with the addition of a gambling bias. ΔBIC refers to the difference in BIC scores between each model and the winning model.

Variability (i.e., standard deviation) in SN/VTA BOLD activity was uncorrelated with the difference in risk taking between low and high activity across participants (Spearman $\rho = -0.19$, $P = 0.29$). By design, all participants were offered a set of options in the real-time fMRI task such that each should gamble half of the time on average. However, the percentage of risky choices was negatively correlated with the difference in risk taking between low and high SN/VTA activity (Spearman $\rho = -0.46$, $P = 0.002$). This means that the decisions of people who gamble less than predicted by prospect theory are more susceptible to endogenous SN/VTA fluctuations.

Dopamine activity is known to influence behaviour in multiple ways. For example, high tonic dopamine is proposed to mediate an enhanced motivational vigour (Hamid et al., 2016; Niv, Daw, Joel, & Dayan, 2007), leading to faster reaction times. We reasoned that if endogenous SN/VTA BOLD fluctuations reflect changes in tonic dopamine, participants should choose more quickly when endogenous activity is high. Matching this prediction, we found faster reaction times in high (1.67 ± 0.05 s) compared to low (1.72 ± 0.05 s) activity conditions ($t_{42} = 3.13$, $P = 0.003$; Fig. 4.6A), consistent with an influence of tonic dopamine on endogenous SN/VTA BOLD activity. This effect was present in 31 of 43 participants (Fig. 4.6B). We conducted an additional multiple linear regression and predicted reaction times based on SN/VTA BOLD activity, choice to safe or risky option, and the absolute value of the difference in option subjective values, an index of choice difficulty. Even when controlling for these variables, reaction times were still significantly related to SN/VTA BOLD activity ($t_{42} = 3.08$, $P = 0.004$).

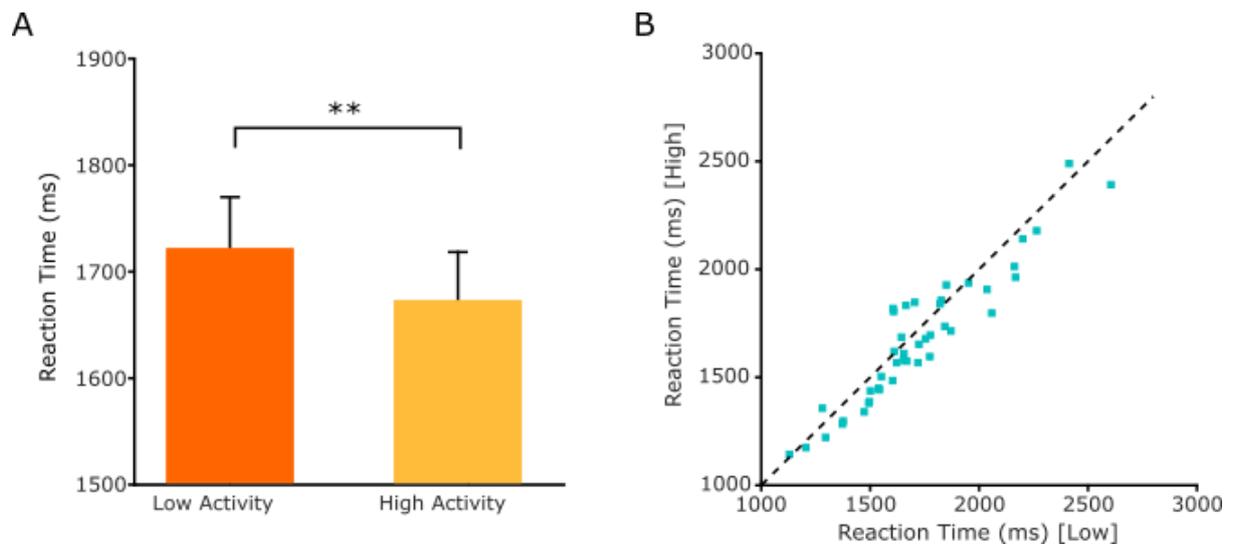


Fig. 4.6. Endogenous fluctuations in dopaminergic midbrain modulate vigour. A, Endogenous activity in dopaminergic midbrain modulated vigour as captured by response speed. Subjects ($n = 43$) were faster ($P < 0.01$) to make choices for options presented on a background of high compared to low endogenous activity. **B,** This effect of faster response speeds for high than low activity was consistent across subjects. ** $P < 0.01$.

4.2.3 Endogenous fluctuations affect phasic responses during choice

We next asked how endogenous SN/VTA BOLD fluctuations shifts preferences in risky decision making as described in our computational model. Given a known association between baseline activity and task-evoked responses (Fox et al., 2006), we hypothesized that endogenous SN/VTA BOLD fluctuations impact risk taking through an influence on the expression of phasic task-evoked activity known to represent choice-relevant information (Lak et al., 2014). Apart from the SN/VTA, other candidate sites that could be involved are the ventral striatum (VS) and ventromedial prefrontal cortex (vmPFC), regions that receive dense dopaminergic innervation (Ferenczi et al., 2016; Williams & Goldman-Rakic, 1998) and express strong functional connectivity with the SN/VTA (Fig. 4.7).

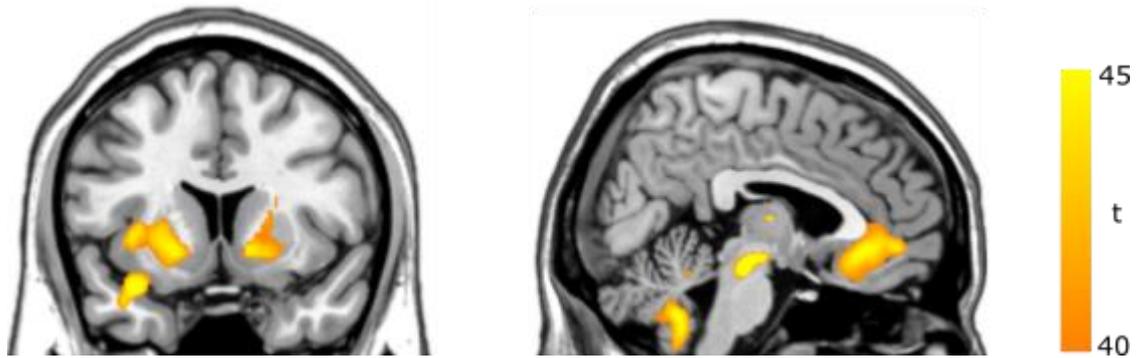


Fig. 4.7. Intrinsic fluctuations in dopaminergic midbrain co-activate a decision network. To identify a network of brain areas that the SN/VTA was embedded in and whose activity co-varied with endogenous fluctuations in the SN/VTA, we first extracted the BOLD time course from each subject's SN/VTA using the same independent ROI as in the offline analyses. This was then included as an additional regressor in a GLM at the 1st level analysis in SPM. T-Contrasts on this SN/VTA regressor were used in the 2nd level group analysis, revealing that activity in dopaminergic midbrain co-activates a decision-related network of areas including VS and vmPFC (both $P < 0.05$, FWE-corrected).

We examined task-evoked SN/VTA responses and found that phasic responses to offer presentation were significantly increased in low compared to high pre-stimulus activity (Fig. 4.8A; $P < 0.01$, cluster-extent permutation test, height threshold $t = 2$, 5000 permutations). We next examined task-evoked responses in VS and vmPFC and found the same effect as in SN/VTA with low endogenous SN/VTA BOLD activity leading to larger phasic task-evoked responses in both VS ($P < 0.01$, cluster-extent permutation, Fig. 4.8B) and vmPFC ($P < 0.01$; Fig. 4.8C). Consistent with previous studies (Bartra et al., 2013), we also found that phasic BOLD responses in VS and vmPFC reflected the subjective values of options.

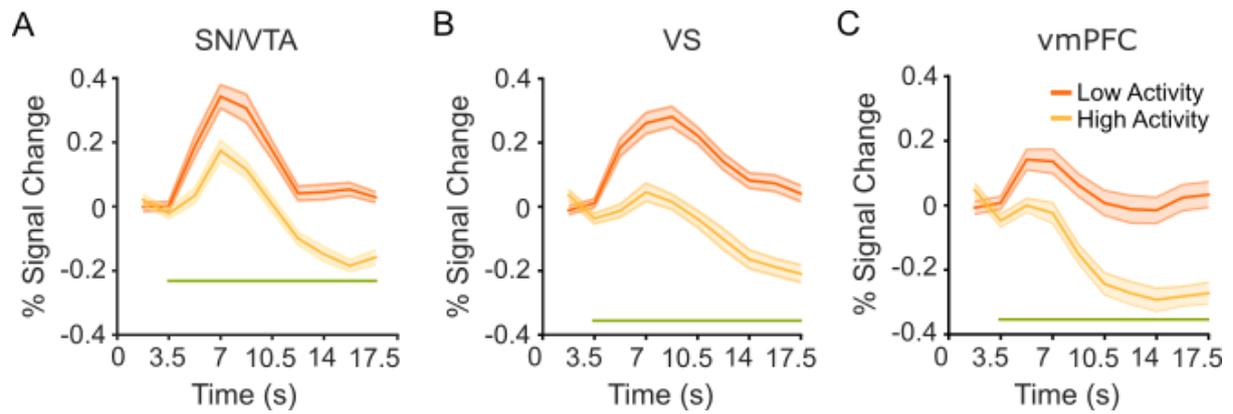


Fig. 4.8. Task-evoked responses are influenced by endogenous fluctuations. A, B, C, Endogenous fluctuations lead to distinct task-evoked response patterns with greater BOLD responses in SN/VTA, VS, and vmPFC when offers are presented against a background of low endogenous SN/VTA activity. Percent signal change was calculated relative to the two volumes following stimulus onset to correct for differences in starting baseline. The green horizontal line indicates statistical significance ($P < 0.01$).

We reasoned that if task-evoked responses play a critical role in translating endogenous fluctuations into risky choice, then participants with stronger effects of endogenous SN/VTA BOLD fluctuations on task-evoked responses should show a greater difference in risk taking in low compared to high activity conditions. We found this to be the case, with larger effects of endogenous SN/VTA fluctuations on task-evoked SN/VTA responses predicting larger increases in the gambling bias parameter κ ($r = 0.39$, $P = 0.03$; Fig. 4.9A). This effect was specific to phasic SN/VTA responses as there was no such effect in decision and control areas (all $P > 0.1$).

Lastly, we investigated the relative contributions of both endogenous SN/VTA BOLD fluctuations and task-evoked responses to risk taking using multi-level mediation analyses (Atlas, Lindquist, Bolger, & Wager, 2014). These analyses were conducted using the Mediation Toolbox (<http://wagerlab.colorado.edu/tools>) (Wager, Davidson, Hughes, Lindquist, & Ochsner, 2008). Evoked responses in the SN/VTA, VS, and vmPFC were determined as the maximum percentage change in BOLD signal within a 10s epoch following trial onset, while baseline SN/VTA was determined as the percentile that each trial was triggered off. Distribution of path coefficients were estimated by drawing 10,000 random samples and significance estimates were computed through bootstrapping.

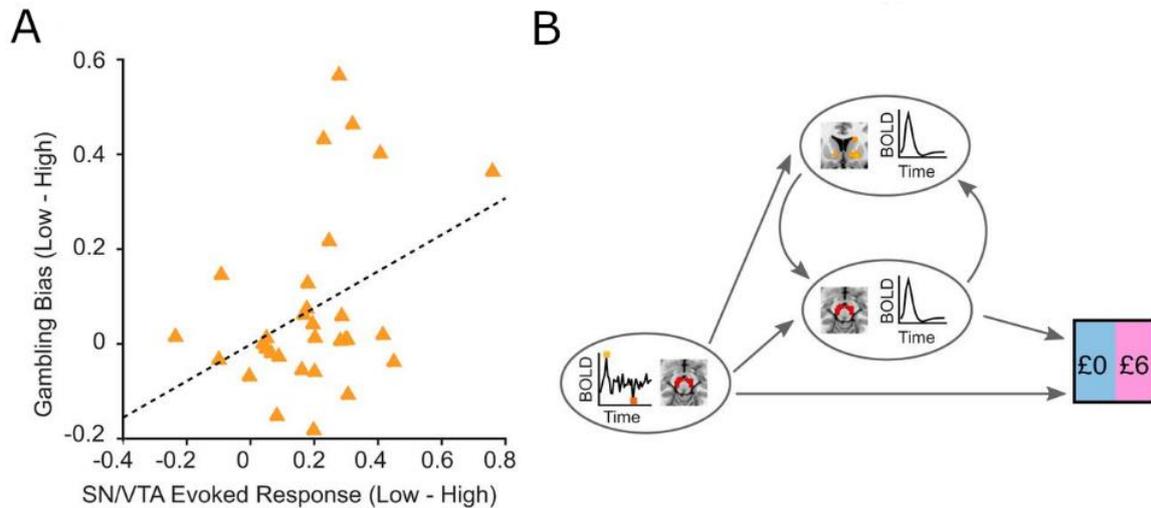


Fig. 4.9. Endogenous fluctuations modulate risk taking via task-evoked responses. **A**, The effect of endogenous SN/VTA activity on risk taking is associated with phasic task-evoked SN/VTA responses. Subjects with a larger difference in task-evoked responses between low and high activity conditions had larger differences in gambling bias parameter κ ($r = 0.39$, $P = 0.03$). **B**, Mediation analysis shows task-evoked VS responses mediate the influence of endogenous SN/VTA fluctuations on risk taking through their influence on task-evoked SN/VTA responses, indicating the effect of endogenous SN/VTA fluctuations on behaviour is under the influence of reciprocal dynamics between SN/VTA and VS.

The mediation analysis tests whether baseline SN/VTA BOLD activity influences the magnitude of task-evoked responses in SN/VTA [VS / vmPFC] (path a), whether task-evoked responses in SN/VTA [VS / vmPFC] are correlated with choice controlling for SN/VTA baseline (path b), whether the relationship between SN/VTA baseline and risk taking is reduced after controlling for task-evoked responses (path c'), and finally a test of mediation. A mediator can be interpreted as an indirect pathway through a brain region that links endogenous fluctuations in SN/VTA baseline activity with choice, whereby this relationship would be reduced or abolished if the mediator is disrupted. To further understand how task-evoked responses in VS mediates baseline SN/VTA BOLD activity and choice despite the absence of a direct link between VS and choice, we conducted an additional analysis using task-evoked responses in VS to predict choice using task-evoked responses in SN/VTA as a mediating variable.

We found that SN/VTA BOLD activity significantly impacted task-evoked response in both SN/VTA and VS (Fig. 4.9B; Table 4.2). Task-evoked SN/VTA responses modulated risk taking, but task-evoked VS responses influenced risk taking

only indirectly through their impact on task-evoked SN/VTA responses (Fig. 4.10; Table 4.2). These results show that endogenous SN/VTA BOLD fluctuations shape decision making through their influence on task-evoked responses to offers in a decision network.

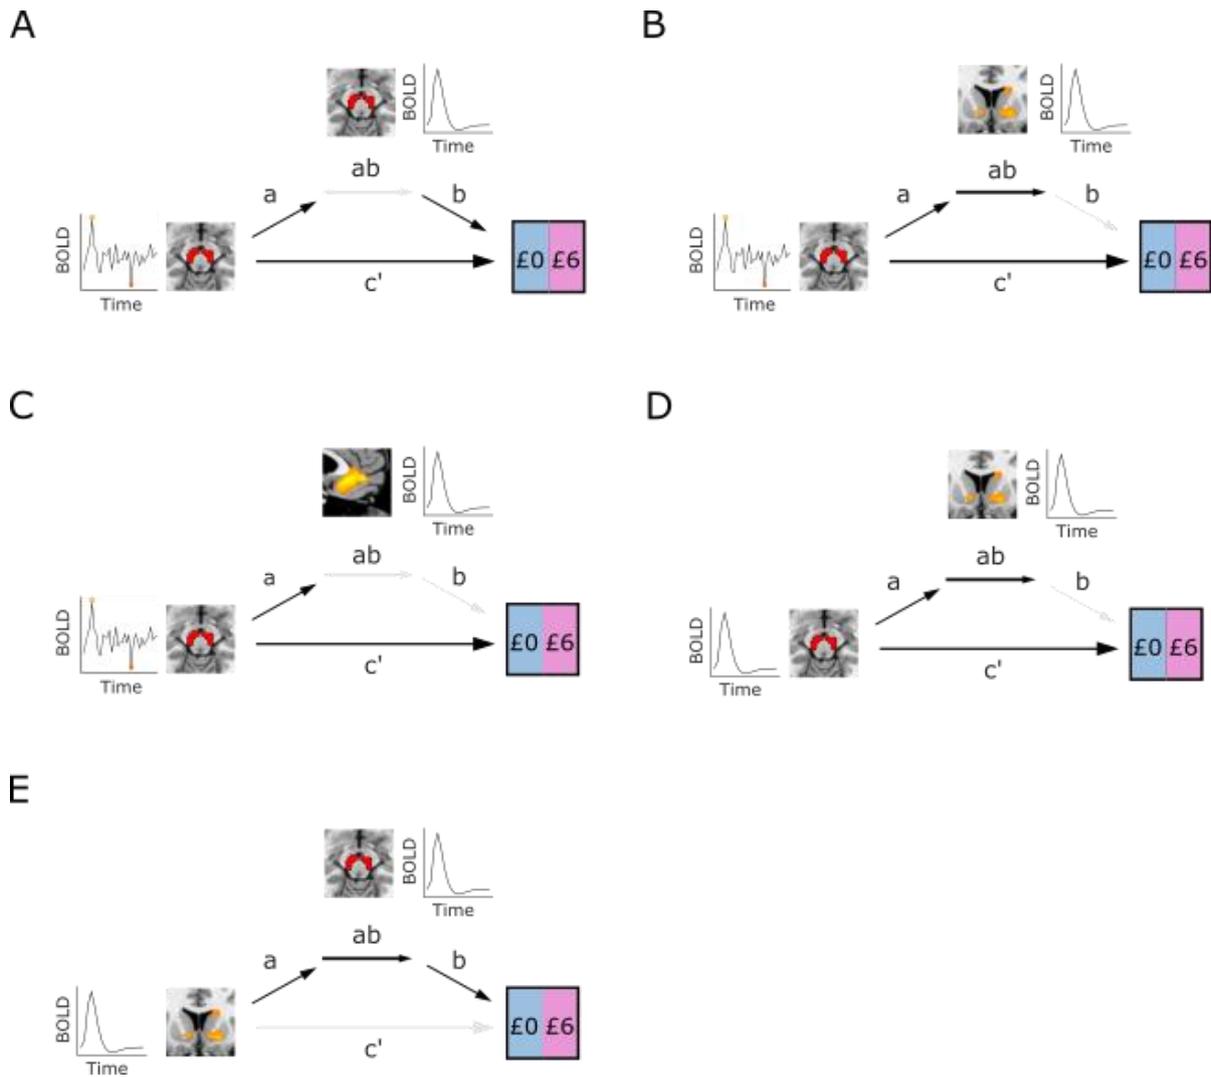


Fig. 4.10. Endogenous SN/VTA activity modulates task-evoked responses in a decision network. **A**, Endogenous SN/VTA activity influences task-evoked responses in SN/VTA, which in turn modulate decision making. **B**, **C**, Endogenous SN/VTA activity also influences task-evoked responses in VS in **B** and vmPFC in **C**. VS thereby mediates the effect of endogenous SN/VTA activity on risk taking. **D**, **E**, Further analysis revealed that task-evoked responses in SN/VTA and VS dynamically interact with each other. Of these three decision areas, only SN/VTA directly influences risk taking, with task-evoked VS responses indirectly influencing behaviour through their influence on task-evoked SN/VTA responses. This result indicates that endogenous SN/VTA activity induces differential phasic responses in SN/VTA and VS, which in turn interact with and dynamically influence risk taking.

Table 4.2. ROI-based Mediation Results.

		SN/VTA	VS	VMPFC
Mediation of endogenous SN/VTA fluctuations on choice	Path a	-0.16*** (0.03) P < 0.001	-0.21*** (0.04) P < 0.001	-0.19*** (0.05) P < 0.001
	Path b	-0.16* (0.08) P = 0.03	-0.0011 (0.05) P = 0.98	-0.04 (0.05) P = 0.31
	Direct c'	-0.14* (0.06) P = 0.01	-0.10* (0.05) P = 0.04	-0.12* (0.05) P = 0.02
	Mediation a x b	-0.13 (0.15) P = 0.40	-0.26* (0.10) P = 0.02	0.01 (0.09) P = 0.96
Mediation of task-evoked VS responses on choice	Path a	0.52*** (0.03) P < 0.001		
	Path b	-0.24** (0.09) P = 0.006		
	Direct c'	0.13 (0.07) P = 0.08		
	Mediation a x b	-2.41** (0.89) P = 0.005		

	SN/VTA	VS	VMPFC
Mediation of task-evoked SN/VTA responses on choice	Path a	0.63*** (0.04) P < 0.001	
	Path b	0.13 (0.07) P = 0.07	
	Direct c'	-0.24** (0.09) P = 0.006	
	Mediation a x b	1.58* (0.78) P = 0.04	

Coefficients, standard errors, and p-values for the different paths in the mediation analyses (n=43). *P < 0.05, **P < 0.01, ***P < 0.001

4.3 Discussion

The brain expresses substantial ongoing activity in the absence of external stimulation. Although many studies capitalize on this fact and have described this 'resting state' (Fox & Raichle, 2007), little is known about the function of spontaneous fluctuations and whether it carries relevance for higher-order cognition. We show that endogenous fluctuations in the dopaminergic midbrain have direct behavioural relevance in modulating a preference for risky decision making in humans. Using a novel framework to study the influence of intrinsic fluctuations on behaviour, we find greater risk taking when choice options are presented against a background of low compared to high SN/VTA BOLD activity. Our findings highlight that the endogenous state of a network relevant for behaviour is critical for determining which actions are taken.

We show that endogenous SN/VTA BOLD activity influences risky decisions via modulation of phasic task-evoked responses to potential rewards. Our results are consistent with findings that impulsive behaviour is linked to phasic dopamine release (Buckholz et al., 2010; Joutsa et al., 2012) and to levodopa administration (Rigoli et al., 2016; Rutledge et al., 2015), assumed to increase phasic dopamine (Pessiglione, Seymour, Flandin, Dolan, & Frith, 2006). Low pre-stimulus activity and levodopa administration may both exert their effects on risk taking by boosting task-evoked phasic responses, which in turn promote risk-taking behaviour.

Reward-predicting cues elicit phasic responses in midbrain dopamine neurons (Morris, Nevet, Arkadir, Vaadia, & Bergman, 2006). In rodents, optogenetic manipulation of SN dopamine neurons boosts striatal dopamine release and biases action selection (Howard, Li, Geddes, & Jin, 2017). Optogenetic stimulation of striatal D2-receptor neurons modulates risk preferences (Zalocusky et al., 2016). Attenuation of pre-choice phasic dopamine via electrical stimulation of the lateral habenula reduces preference for risk in rodents (Stopper, Tse, Montes, Wiedman, & Floresco, 2014). Our study builds on these results by identifying a possible link between pre-stimulus brain activity, phasic responses to stimuli, and subsequent risky choice.

The functional role of these endogenous fluctuations remains unclear, but they might form a reference point relative to which potential offers are evaluated. While standard models of economic decision making often treat preferences as independent

of a rational agent's current state, real-world behaviour often reflects comparison against a reference point that can change over time (Tversky & Kahneman, 1991), sometimes substantially changing the subjective value of an offer (Kahneman & Tversky, 1979). If endogenous SN/VTA BOLD activity reflects slow dopaminergic fluctuations, proposed to index environmental reward rate (Hamid et al., 2016; Niv et al., 2007) or reward anticipation (MacInnes, Dickerson, Chen, & Adcock, 2016), then these fluctuations could represent a reference point against which potential rewards are compared during decision making (Kőszegi & Rabin, 2006; Louie & Glimcher, 2012). Potential rewards presented on a background of low activity could lead to enhanced task-evoked responses linked to greater risk taking (Stopper et al., 2014).

Endogenous fluctuations may constitute an evolutionarily conserved principle that enables the brain to introduce variability across a wide variety of processes including perception (Boly et al., 2007; Sadaghiani et al., 2015) and motor action (Fox et al., 2007). Neural variability has been hypothesized to reflect the dynamic range of potential responses to environmental stimuli, allowing the brain to flexibly transition between states in response to changing task demands (Garrett, Kovacevic, McIntosh, & Grady, 2013). It could also reduce susceptibility to becoming entrenched in specific behavioural repertoires (van Leeuwen, 2008) and promote exploration in dynamic environments that are a common feature of the natural world (Wilson, Geana, White, Ludvig, & Cohen, 2014).

One possible source of variability that could relate to our results is D2/D3 autoreceptor availability in the SN/VTA. Lower autoreceptor availability is associated with greater dopamine release following amphetamine administration and greater trait impulsivity (Buckholtz et al., 2010). The link we find between risk taking and phasic responses is also consistent with the finding that phasic dopamine during gambling tasks is greater in pathological gamblers (Joutsa et al., 2012).

One limitation of our study is that BOLD activity is an indirect measure of local neuronal activity thought to consist of an ensemble of signals including afferent and recurrent inputs (Logothetis & Wandell, 2004). Phasic and tonic dopamine release may contribute to fluctuations in SN/VTA BOLD activity, while optogenetic stimulation of dopamine neurons in VTA is sufficient to elicit BOLD activity in VTA (Brocka et al., 2018). However, SN/VTA BOLD activity may also reflect activity in other cell types

including glutamatergic (Sesack & Grace, 2010) and GABAergic neurons that act to inhibit dopamine neurons when reward is expected (Eshel et al., 2015). Reduced GABAergic activity could also be associated with greater phasic dopamine release and provide an alternate explanation for greater risk taking when pre-stimulus SN/VTA BOLD activity is low.

Previous studies have shown that dopamine release in VS, as measured using PET, is linked to reward-related SN/VTA BOLD activity (Schott et al., 2008). While we have focused on the SN/VTA, the extent to which BOLD activity in the downstream VS responds to a cue is tightly coupled to SN/VTA BOLD activity. Future studies might extend these findings with direct striatal recordings that assess the relationship between spontaneous fluctuations and risk taking.

Our effect is consistent across individuals, albeit modest in terms of effect size (on average, a 3.4% increase in the number of risky options chosen). We would predict a larger effect size with direct electrophysiological recordings, since fMRI measurements are inherently noisy at several levels (Webb, Levy, Lazzaro, Rutledge, & Glimcher, 2019). However, given the many factors that contribute to risky decision making, it would be surprising if the state of the brain when options are presented had a large effect on the probability of risky decisions, especially in the absence of any environmental changes.

The effect size is comparable in size to previous studies. For example, a standard clinical dose of 150mg of levodopa increased risk taking by only 5% on average (Rutledge et al., 2015) and natural aging leads to a comparable decrease in risk taking, which we surmised may reflect age-related dopaminergic decline estimated at 5-10% per decade (Rutledge et al., 2016). It is also noteworthy that the effect of low pre-stimulus SN/VTA activity on risk taking is particularly large in relative terms for unattractive gambles. The probability that individuals choose the least attractive gambles (chosen less than 20% of the time) is much greater under low than high pre-stimulus activity (18.7% vs 15.4%), a 21% relative increase. In contrast, the probability that individuals choose the most attractive gambles (chosen more than 80% of the time) is only 5% greater under low than high pre-stimulus activity in relative terms. Our findings may be particularly relevant to understanding pathological gamblers, who may take risks that others would generally avoid.

Our key finding is that variability in higher-order cognition can emerge out of a neurophysiologically well-defined process. While risk preferences are thought of as personality traits determined partly by genetic variation (Frydman, Camerer, Bossaerts, & Rangel, 2010), we show that the expression of risk preferences reflects in part individual susceptibility to endogenous fluctuations. Neural variability may change with task experience, consistent with reductions in neural variability during skill learning (Santos, Oliveira, Jin, & Costa, 2015) and the impact of endogenous fluctuations may be largest in novel environments. Aberrant endogenous fluctuations might also play a role in disorders where there is excessive behavioural variability or risk taking, such as ADHD (Hauser, Fiore, Moutoussis, & Dolan, 2016) and pathological gambling (Joutsa et al., 2012). Accounting for the influence of endogenous neural fluctuations on behaviour is critical for understanding the neurobiological processes underlying cognition in health and disorder.

References

- Atlas, L. Y., Lindquist, M. A., Bolger, N., & Wager, T. D. (2014). Brain mediators of the effects of noxious heat on pain. *Pain*, *155*(8), 1632–1648.
<https://doi.org/10.1016/j.pain.2014.05.015>
- Bartra, O., McGuire, J. T., & Kable, J. W. (2013). The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *NeuroImage*, *76*, 412–427. <https://doi.org/10.1016/j.neuroimage.2013.02.063>
- Birn, R. M., Diamond, J. B., Smith, M. A., & Bandettini, P. A. (2006). Separating respiratory-variation-related fluctuations from neuronal-activity-related fluctuations in fMRI. *NeuroImage*, *31*(4), 1536–1548. <https://doi.org/10.1016/j.neuroimage.2006.02.048>
- Bolding, M. S., Reid, M. A., Avsar, K. B., Roberts, R. C., Gamlin, P. D., Gawne, T. J., Lahti, A. C. (2013). Magnetic transfer contrast accurately localizes substantia nigra confirmed by histology. *Biological Psychiatry*, *73*(3), 289–294.
<https://doi.org/10.1016/j.biopsych.2012.07.035>
- Boly, M., Balteau, E., Schnakers, C., Degueldre, C., Moonen, G., Luxen, A., Laureys, S. (2007). Baseline brain activity fluctuations predict somatosensory perception in humans. *PNAS*, *104*(29), 12187–12192. <https://doi.org/10.1073/pnas.0611404104>
- Brocka, M., Helbing, C., Vincenz, D., Scherf, T., Montag, D., Goldschmidt, J., Lippert, M. (2018). Contributions of dopaminergic and non-dopaminergic neurons to VTA-stimulation induced neurovascular responses in brain reward circuits. *NeuroImage*, *177*, 88–97. <https://doi.org/10.1016/j.neuroimage.2018.04.059>
- Brown, H. R., Zeidman, P., Smittenaar, P., Adams, R. A., McNab, F., Rutledge, R. B., & Dolan, R. J. (2014). Crowdsourcing for cognitive science – the utility of smartphones. *PLOS ONE*, *9*(7), e100662. <https://doi.org/10.1371/journal.pone.0100662>
- Buckholz, J. W., Treadway, M. T., Cowan, R. L., Woodward, N. D., Li, R., Ansari, M. S., Zald, D. H. (2010). Dopaminergic network differences in human impulsivity. *Science*, *329*(5991), 532–532. <https://doi.org/10.1126/science.1185778>
- Burke, C. J., Soutschek, A., Weber, S., Beharelle, A. R., Fehr, E., Haker, H., & Tobler, P. N. (2018). Dopamine receptor-specific contributions to the computation of value. *Neuropsychopharmacology*, *43*(6), 1415–1424. <https://doi.org/10.1038/npp.2017.302>
- Chowdhury, R., Guitart-Masip, M., Lambert, C., Dolan, R. J., & Düzel, E. (2013). Structural integrity of the substantia nigra and subthalamic nucleus predicts flexibility of instrumental learning in older-age individuals. *Neurobiology of Aging*, *34*(10), 2261–2270. <https://doi.org/10.1016/j.neurobiolaging.2013.03.030>

- D'Ardenne, K., McClure, S. M., Nystrom, L. E., & Cohen, J. D. (2008). BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. *Science*, *319*(5867), 1264–1267. <https://doi.org/10.1126/science.1150605>
- Eshel, N., Bukwich, M., Rao, V., Hemmelder, V., Tian, J., & Uchida, N. (2015). Arithmetic and local circuitry underlying dopamine prediction errors. *Nature*, *525*(7568), 243–246. <https://doi.org/10.1038/nature14855>
- Ferenczi, E. A., Zalocusky, K. A., Liston, C., Grosenick, L., Warden, M. R., Amatya, D., Deisseroth, K. (2016). Prefrontal cortical regulation of brainwide circuit dynamics and reward-related behavior. *Science*, *351*(6268), aac9698. <https://doi.org/10.1126/science.aac9698>
- Fox, M. D., & Raichle, M. E. (2007). Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Reviews Neuroscience*, *8*(9), 700–711. <https://doi.org/10.1038/nrn2201>
- Fox, M. D., Snyder, A. Z., Vincent, J. L., & Raichle, M. E. (2007). Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior. *Neuron*, *56*(1), 171–184. <https://doi.org/10.1016/j.neuron.2007.08.023>
- Fox, M. D., Snyder, A. Z., Zacks, J. M., & Raichle, M. E. (2006). Coherent spontaneous activity accounts for trial-to-trial variability in human evoked brain responses. *Nature Neuroscience*, *9*(1), 23–25. <https://doi.org/10.1038/nn1616>
- Frydman, C., Camerer, C., Bossaerts, P., & Rangel, A. (2010). MAOA-L carriers are better at making optimal financial decisions under risk. *Proceedings of the Royal Society B: Biological Sciences*, *278*, 2053–2059. <https://doi.org/10.1098/rspb.2010.2304>
- Garrett, D. D., Kovacevic, N., McIntosh, A. R., & Grady, C. L. (2013). The modulation of BOLD variability between cognitive states varies by age and processing speed. *Cerebral Cortex*, *23*(3), 684–693. <https://doi.org/10.1093/cercor/bhs055>
- Glover, G. H., Li, T. Q., & Ress, D. (2000). Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. *Magnetic Resonance in Medicine*, *44*(1), 162–167.
- Hamid, A. A., Pettibone, J. R., Mabrouk, O. S., Hetrick, V. L., Schmidt, R., Weele, C. M. V., Berke, J. D. (2016). Mesolimbic dopamine signals the value of work. *Nature Neuroscience*, *19*(1), 117–126. <https://doi.org/10.1038/nn.4173>
- Harless, D. W., & Camerer, C. F. (1994). The predictive utility of generalized expected utility theories. *Econometrica*, *62*(6), 1251–1289. <https://doi.org/10.2307/2951749>
- Hauser, T. U., Eldar, E., & Dolan, R. J. (2017). Separate mesocortical and mesolimbic pathways encode effort and reward learning signals. *PNAS*, 7395–7404. <https://doi.org/10.1073/pnas.1705643114>

- Hauser, T. U., Fiore, V. G., Moutoussis, M., & Dolan, R. J. (2016). Computational psychiatry of ADHD: Neural gain impairments across Marrian levels of analysis. *Trends in Neurosciences*, 39(2), 63–73. <https://doi.org/10.1016/j.tins.2015.12.009>
- Howard, C. D., Li, H., Geddes, C. E., & Jin, X. (2017). Dynamic nigrostriatal dopamine biases action selection. *Neuron*, 93(6), 1436-1450.e8. <https://doi.org/10.1016/j.neuron.2017.02.029>
- Joutsa, J., Johansson, J., Niemelä, S., Ollikainen, A., Hirvonen, M. M., Piepponen, P., Kaasinen, V. (2012). Mesolimbic dopamine release is linked to symptom severity in pathological gambling. *NeuroImage*, 60(4), 1992–1999. <https://doi.org/10.1016/j.neuroimage.2012.02.006>
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–291. <https://doi.org/10.2307/1914185>
- Koster, R., Guitart-Masip, M., Dolan, R. J., & Düzel, E. (2015). Basal ganglia activity mirrors a benefit of action and reward on long-lasting event memory. *Cerebral Cortex*, 25(12), 4908–4917. <https://doi.org/10.1093/cercor/bhv216>
- Kőszegi, B., & Rabin, M. (2006). A model of reference-dependent preferences. *The Quarterly Journal of Economics*, 121(4), 1133–1165.
- Lak, A., Stauffer, W. R., & Schultz, W. (2014). Dopamine prediction error responses integrate subjective value from different reward dimensions. *PNAS*, 111(6), 2343–2348. <https://doi.org/10.1073/pnas.1321596111>
- Logothetis, N. K., & Wandell, B. A. (2004). Interpreting the BOLD Signal. *Annual Review of Physiology*, 66(1), 735–769. <https://doi.org/10.1146/annurev.physiol.66.082602.092845>
- Louie, K., & Glimcher, P. W. (2012). Efficient coding and the neural representation of value. *Annals of the New York Academy of Sciences*, 1251(1), 13–32. <https://doi.org/10.1111/j.1749-6632.2012.06496.x>
- MacInnes, J. J., Dickerson, K. C., Chen, N., & Adcock, R. A. (2016). Cognitive neurostimulation: Learning to volitionally sustain ventral tegmental area activation. *Neuron*, 89(6), 1331–1342. <https://doi.org/10.1016/j.neuron.2016.02.002>
- Maldjian, J. A., Laurienti, P. J., & Burdette, J. H. (2004). Precentral gyrus discrepancy in electronic versions of the Talairach atlas. *NeuroImage*, 21(1), 450–455.
- Morris, G., Nevet, A., Arkadir, D., Vaadia, E., & Bergman, H. (2006). Midbrain dopamine neurons encode decisions for future action. *Nature Neuroscience*, 9(8), 1057. <https://doi.org/10.1038/nn1743>

- Niv, Y., Daw, N. D., Joel, D., & Dayan, P. (2007). Tonic dopamine: Opportunity costs and the control of response vigor. *Psychopharmacology*, *191*(3), 507–520.
<https://doi.org/10.1007/s00213-006-0502-4>
- Pessiglione, M., Seymour, B., Flandin, G., Dolan, R. J., & Frith, C. D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*, *442*(7106), 1042–1045. <https://doi.org/10.1038/nature05051>
- Rigoli, F., Rutledge, R. B., Chew, B., Ousdal, O. T., Dayan, P., & Dolan, R. J. (2016). Dopamine increases a value-independent gambling propensity. *Neuropsychopharmacology*, *41*(11), 2658–2667. <https://doi.org/10.1038/npp.2016.68>
- Rorden, C., & Brett, M. (2000). Stereotaxic display of brain lesions. *Behavioural Neurology*, *12*(4), 191–200.
- Rutledge, R. B., Skandali, N., Dayan, P., & Dolan, R. J. (2014). A computational and neural model of momentary subjective well-being. *PNAS*, *111*(33), 12252–12257.
<https://doi.org/10.1073/pnas.1407535111>
- Rutledge, R. B., Skandali, N., Dayan, P., & Dolan, R. J. (2015). Dopaminergic modulation of decision making and subjective well-being. *Journal of Neuroscience*, *35*(27), 9811–9822. <https://doi.org/10.1523/JNEUROSCI.0702-15.2015>
- Rutledge, R. B., Smittenaar, P., Zeidman, P., Brown, H. R., Adams, R. A., Lindenberger, U., Dolan, R. J. (2016). Risk taking for potential reward decreases across the lifespan. *Current Biology*, *26*(12), 1634–1639. <https://doi.org/10.1016/j.cub.2016.05.017>
- Sadaghiani, S., Poline, J.-B., Kleinschmidt, A., & D'Esposito, M. (2015). Ongoing dynamics in large-scale functional connectivity predict perception, *PNAS*, *112*(27), 8463–8468.
<https://doi.org/10.1073/pnas.1420687112>
- Santos, F. J., Oliveira, R. F., Jin, X., & Costa, R. M. (2015). Corticostriatal dynamics encode the refinement of specific behavioral variability during skill learning. *eLife*, *4*, e09423.
<https://doi.org/10.7554/eLife.09423>
- Schott, B. H., Minuzzi, L., Krebs, R. M., Elmenhorst, D., Lang, M., Winz, O. H., Bauer, A. (2008). Mesolimbic functional magnetic resonance imaging activations during reward anticipation correlate with reward-related ventral striatal dopamine release. *Journal of Neuroscience*, *28*(52), 14311–14319. <https://doi.org/10.1523/JNEUROSCI.2058-08.2008>
- Sesack, S. R., & Grace, A. A. (2010). Cortico-basal ganglia reward network: Microcircuitry. *Neuropsychopharmacology*, *35*(1), 27–47. <https://doi.org/10.1038/npp.2009.93>
- Sokol-Hessner, P., Hsu, M., Curley, N. G., Delgado, M. R., Camerer, C. F., & Phelps, E. A. (2009). Thinking like a trader selectively reduces individuals' loss aversion. *PNAS*, *106*(13), 5035–5040. <https://doi.org/10.1073/pnas.0806761106>

- Starkweather, C. K., Babayan, B. M., Uchida, N., & Gershman, S. J. (2017). Dopamine reward prediction errors reflect hidden-state inference across time. *Nature Neuroscience*, *20*(4), 581–589. <https://doi.org/10.1038/nn.4520>
- Steinberg, E. E., Keiflin, R., Boivin, J. R., Witten, I. B., Deisseroth, K., & Janak, P. H. (2013). A causal link between prediction errors, dopamine neurons and learning. *Nature Neuroscience*, *16*(7), 966–973. <https://doi.org/10.1038/nn.3413>
- Stopper, C. M., Tse, M. T. L., Montes, D. R., Wiedman, C. R., & Floresco, S. B. (2014). Overriding phasic dopamine signals redirects action selection during risk/reward decision making. *Neuron*, *84*(1), 177–189. <https://doi.org/10.1016/j.neuron.2014.08.033>
- Tavor, I., Jones, O. P., Mars, R. B., Smith, S. M., Behrens, T. E., & Jbabdi, S. (2016). Task-free MRI predicts individual differences in brain activity during task performance. *Science*, *352*(6282), 216–220. <https://doi.org/10.1126/science.aad8127>
- Timmer, M. H. M., Sescousse, G., Esselink, R. A. J., Piray, P., & Cools, R. (2017). Mechanisms underlying dopamine-induced risky choice in Parkinson's disease with and without depression. *Computational Psychiatry*, *2*, 11–27. https://doi.org/10.1162/CPSY_a_00011
- Tversky, A., & Kahneman, D. (1991). Loss aversion in riskless choice: A reference-dependent model. *The Quarterly Journal of Economics*, *106*(4), 1039–1061. <https://doi.org/10.2307/2937956>
- van Leeuwen, C. (2008). Chaos breeds autonomy: Connectionist design between bias and baby-sitting. *Cognitive Processing*, *9*(2), 83–92. <https://doi.org/10.1007/s10339-007-0193-8>
- Wager, T. D., Davidson, M. L., Hughes, B. L., Lindquist, M. A., & Ochsner, K. N. (2008). Prefrontal-subcortical pathways mediating successful emotion regulation. *Neuron*, *59*(6), 1037–1050. <https://doi.org/10.1016/j.neuron.2008.09.006>
- Webb, R., Levy, I., Lazzaro, S. C., Rutledge, R. B., & Glimcher, P. W. (2019). Neural random utility: Relating cardinal neural observables to stochastic choice behavior. *Journal of Neuroscience, Psychology, and Economics*, *12*(1), 45–72. <https://doi.org/10.1037/npe0000101>
- Weiskopf, N., Suckling, J., Williams, G., Correia, M. M., Inkster, B., Tait, R., Lutti, A. (2013). Quantitative multi-parameter mapping of R1, PD*, MT, and R2* at 3T: A multi-center validation. *Frontiers in Neuroscience*, *7*. <https://doi.org/10.3389/fnins.2013.00095>
- Williams, S. M., & Goldman-Rakic, P. S. (1998). Widespread origin of the primate mesofrontal dopamine system. *Cerebral Cortex*, *8*(4), 321–345.

Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology*, *143*(6), 2074–2081. <https://doi.org/10.1037/a0038199>

Zalocusky, K. A., Ramakrishnan, C., Lerner, T. N., Davidson, T. J., Knutson, B., & Deisseroth, K. (2016). Nucleus accumbens D2R cells signal prior outcomes and control risky decision-making. *Nature*, *531*(7596), 642–646. <https://doi.org/10.1038/nature17400>

Chapter 5: Risk preferences correlate with tissue microstructure

Like the entomologist in search of colourful butterflies, my attention has chased in the gardens of the grey matter cells with delicate and elegant shapes, the mysterious butterflies of the soul, whose beating of wings may one day reveal to us the secrets of the mind.

Santiago Ramón y Cajal, *Recollections of My Life*

As discussed in **Chapter 2**, decisions encountered in the real world generally depend on the values of options available to an organism. These values are *subjective* quantities that are internally generated by the organism rather than objective attributes of the options. Imagine that a decision has to be made between a single sumptuous chocolate chip cookie and two scoops of espresso biscotti ice-cream. Someone who dislikes the taste of coffee might opt for the former while someone who enjoys the flavour of coffee might opt for the latter, resulting in different decision outcomes despite identical options being presented to both people. Such decisions are afforded an additional layer of complexity known as *risk* when the expected outcomes of options are *probabilistic* rather than deterministic, and the influence of risk on subjective values may vary according to the risk attitude of each individual. Individuals who are averse to risk assign lower subjective values to options with low probabilities of receiving a desired outcome, while individuals who are risk-seeking show the opposite preference. A sure gain of £5 has a higher subjective value than an option with an equal probability of winning £10 or nothing for a risk-averse individual while the opposite is true for a risk-seeking individual.

Many studies have previously investigated the functional anatomy of risky decision-making, with areas such as ventromedial prefrontal cortex (vmPFC) (Bartra, McGuire, & Kable, 2013; Levy & Glimcher, 2012), orbitofrontal cortex (OFC) (Padoa-Schioppa, 2013; Padoa-Schioppa & Cai, 2011), ventral striatum (VS) (Bartra et al., 2013), and substantia nigra and ventral tegmental area (SN/VTA) (D'Ardenne, McClure, Nystrom, & Cohen, 2008) being implicated in the decision process. Despite this, the neurobiology driving the relationship between neuroanatomy and risk preferences remains unclear. An underlying source of uncertainty is that volumetric measures of brain structure are often qualitative in nature and lack specificity in the form of properties like axon, myelin, iron, and water concentrations (Tabelow et al.,

2019). In the vein of Cajal and his quest to study the structure of the cerebral cortex, the emerging field of *in vivo* histology seeks to address the preceding lacunae by using a combination of biophysical models, quantitative magnetic resonance imaging (qMRI), and voxel-based quantification (VBQ) using the multi-parameter mapping (MPM) approach. MPM quantifies the longitudinal relaxation rate, R_1 , effective transverse relaxation rate, R_2^* , percent saturation due to magnetization transfer (MT) and effective proton density (PD^*). An advantage of these measures is that they are largely impervious to inter-site biases and are also comparable across different time points due to their quantitative nature as an absolute measure (Weiskopf et al., 2013), facilitating reproducibility and increasing neurobiological specificity of neuroanatomy findings. Collectively, this approach is coined *in vivo* histology using MRI (hMRI).

In the present study, I use hMRI to estimate parameters sensitive to tissue microstructure in order to better understand the neuroanatomical correlates of risky decision-making. High resolution MPMs (800 μm isotropic) were acquired for 46 healthy participants who also engaged in a risky decision-making task involving a choice between a safe and a risky option. The risky option comprised equal probabilities of a prize (£6, £9, or £12) or £0. The value of the safe option was always lower than the potential prize of the risky option and was obtained using divisors on the expected value of the gamble (see Methods).

5.1 Methods

5.1.1 Participants

46 healthy, young adults (age: 25.0 ± 4.2 , mean \pm SD) were recruited through the University College London Psychology Subject Database. Participants were screened to ensure no history of neurological or psychiatric disorders. One subject was excluded due to excessive movements during scanning as measured by the standard deviation of R_2^* maps (Balteau et al., 2018; Castella et al., 2018), leaving a total of 45 participants included in the study. The study was approved by the University College London (UCL) research ethics committee, and all participants gave written informed consent.

5.1.2 Study Design

Participants completed the experiment at the Wellcome Centre for Human Neuroimaging, University College London, in an appointment that lasted approximately an hour. Each appointment consisted of a probabilistic gambling task that lasted 25 minutes, and a 35-minute neuroimaging session where multi-parameter maps (Callaghan et al., 2014; Draganski et al., 2011) were acquired while participants viewed a muted nature documentary.

Probabilistic Gambling Task

Participants played a probabilistic gambling task consisting of 180 trials (Fig. 5.1). On each trial, participants chose between a certain monetary amount and a gamble with equal probabilities of two outcomes. There were three gamble options available: £0 or £6, £0 or £9, and £0 or £12. The certain amounts were determined using 12 divisors (0.82, 0.87, 0.93, 1, 1.1, 1.23, 1.4, 1.6, 1.9, 2.25, 2.75, and 3.5) on the expected value of the gambles, chosen to accommodate a range of risk sensitivity. Take for example a fraction of 3.5 and a gamble between £0 and £6. The expected value of a £0 or £6 gamble would be $(0.5 \times £0) + (0.5 \times £6)$ which gives £3, and this is then divided by 3.5 to give a certain amount of £0.86. There were 12 certain amounts for each gamble option in total, and each trial was repeated 5 times in a randomized sequence. Participants had a maximum time of 3 seconds to make a choice or receive £0 automatically upon exceeding that. The chosen option was displayed for a further 700 milliseconds before an intertrial interval (4.2 to 5.8 seconds) commenced. The outcome of a randomly selected trial was added to a starting pay of £10.

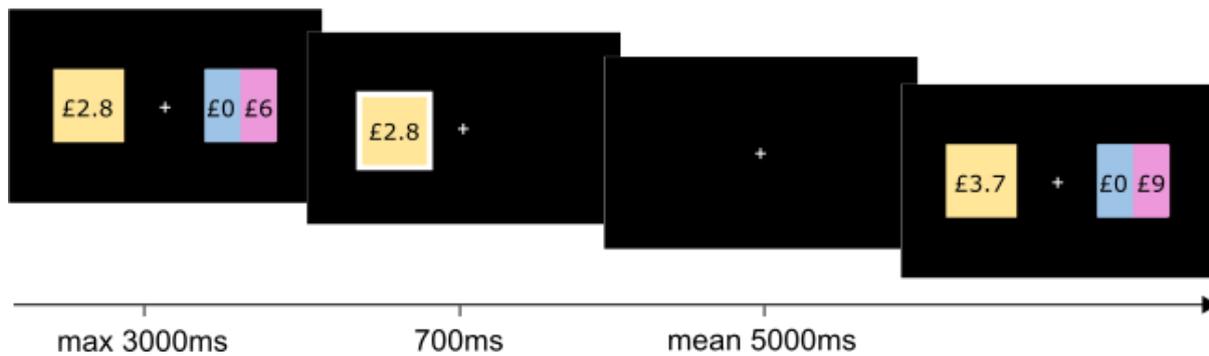


Fig. 5.1. Experimental design. On each trial, participants chose between a safe option (here, £2.8 guaranteed reward) and a risky option (here, £0 or £6 with equal probability). Their chosen option remained on the screen for a further 700ms before the commencement of an intertrial interval with a mean of 5000ms. Outcomes of chosen gambles were not revealed to the participants.

5.1.3 Structural Data Acquisition and Processing

Imaging data consisted of 3 spoiled multi-echo 3D fast low angle shot (FLASH) acquisitions at 0.8 mm isotropic resolution with T1 [TR: 18.7 ms; flip angle: 20°], proton density (PD) [TR: 23.7 ms; flip angle: 6°], and magnetization transfer (MT) [TR: 23.7 ms; flip angle: 6°; excitation preceded by a 2kHz off-resonance Gaussian radiofrequency (RF) pulse with 4 ms duration and 200° nominal flip angle] weightings. One RF sensitivity map was acquired for each contrast to address RF sensitivity bias during processing. Additional B1 Mapping and Field Maps were also acquired to get calibration data measuring the spatial distribution of the B1+ transmit field in order to detect the spatial variation in flip angle. These data can be combined to generate multi-parameter maps (Callaghan et al., 2014).

All images were processed using SPM 12 (Wellcome Centre for Human Neuroimaging, <http://www.fil.ion.ucl.ac.uk/spm/>). Maps of R_2^* were estimated from T1-, PD-, and MT-weighted RF gradient echos using an ordinary least squares ESTATICS approach (Weiskopf, Callaghan, Josephs, Lutti, & Mohammadi, 2014) that provides a more robust estimation of R_2^* with a higher signal-to-noise ratio. The data acquired for each contrast (T1w, PDw, MTw) were then averaged to increase their signal-to-noise ratio. The 3 resulting volumes were then used to calculate MT and R1 maps (Helms, Dathe, Kallenberg, & Dechent, 2008; Weiskopf et al., 2013). Subsequent processing steps followed the standard pipeline within the hMRI toolbox (<https://hmri->

group.github.io/hMRI-toolbox/) which included co-registration, segmentation, diffeomorphic image registration (DARTEL), and smoothing [6mm Full Width at Half Maximum (FWHM)]. Additionally, an enhanced tissue probability map was used to improve segmentation for subcortical areas (Lorio et al., 2016).

5.2 Results

5.2.1 Average gambling is correlated with tissue microstructure in SN/VTA, VS, and Right Hippocampus

I first asked whether decision areas involved in risk were extended to neuroanatomy (Bartra et al., 2013; D'Ardenne et al., 2008; Rangel, Camerer, & Montague, 2008). To do so, I conducted region of interest (ROI) analyses using multiple regression. vmPFC ROI was derived from a meta-analysis of experiments investigating neural correlates of subjective value (Bartra et al., 2013), OFC ROI was a single 8-mm sphere at MNI coordinates [x = 26, y = 18, z = -16] derived from a previous study (Peters & Büchel, 2009), VS ROI was bilateral 8-mm spheres at MNI coordinates [left: x = -10, y = 12, z = -8; right: x = 10, y = 12, z = -8] derived from a previous study (Robb B. Rutledge, Skandali, Dayan, & Dolan, 2014), a group anatomical mask from a previous study (Hauser, Eldar, & Dolan, 2017) was used for SN/VTA ROI, and to test an *a priori* prediction of the model-based approach detailed later, LC was included and defined as 4 × 6 × 10 mm cuboids centered at x = -5, y = -34, z = -21 on the left and x = 7, y = -34, z = -21 on the right (Bär et al., 2016), placing the LC in the floor of the fourth ventricle and in the rostral pons. As no results were significant within the vmPFC and OFC ROIs, they have been dropped from further discussion.

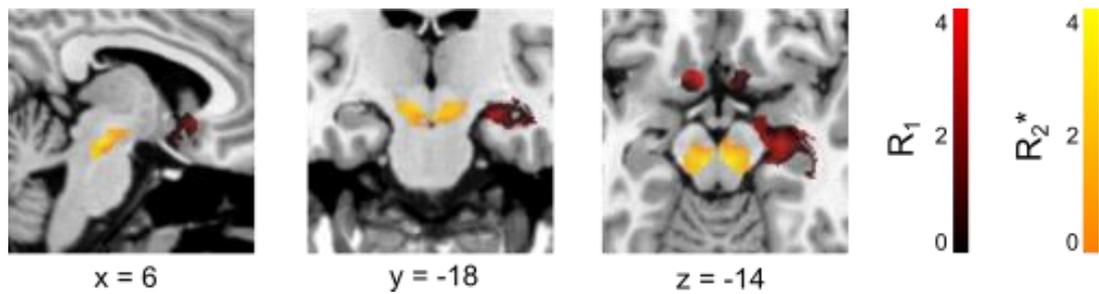


Figure 5.2. VBQ region of interest (ROI) analysis and average gambling. Myelin increases in right hippocampus and ventral striatum with higher average gambling are estimated from increased R_1 (red) while reductions in iron in SN/VTA with higher average gambling are estimated from decreased R_2^* (orange). Figure thresholded at $P < 0.001$ uncorrected level for display purposes.

The percentage of gambles chosen, or *average gambling*, by each subject was modelled as the key dependent variable with age, gender, and total intracranial volume included as nuisance covariates (Ridgway et al., 2008). Age in particular contributes greatly to global patterns of microstructure variation observed using VBQ measures (Callaghan et al., 2014). The main effects of average gambling was analysed within each ROI corrected for small-volume as well as multiple comparisons using a Family-wise error (FWE) cluster alpha of 0.05. Often used as an index of cortical myelination (Lutti, Dick, Sereno, & Weiskopf, 2014), R_1 in the ventral striatum (Fig. 5.2) was found to be positively related to average gambling. Anterior right hippocampus was also positively related to average gambling at the whole-brain FWE-cluster level. R_2^* , a marker of hepatic iron content (Daugherty & Raz, 2013; Wood et al., 2005), in the SN/VTA was negatively related to average gambling (Fig. 5.2) and comparable in quantity to other published R_2^* values (VBQ measure: 23.7 ± 1.6 , Literature (Callaghan et al., 2014): 26.7 ± 4.1 , mean \pm SD). See Table 5.1 for summary of these results.

Table 5.1. Results for average gambling.

Map	Region	P_{FWE}	k_E	t	z	x	y	z
R ₁ +	R Hippo.	< 0.001	206	3.97	3.63	26	-8	-23
R ₁ +	VS	0.001	59	4.37	3.92	-7	15	-13
R ₂ *-	SN/VTA	0.02	15	3.62	3.34	6	-22	-15

+ and - refer to positive and negative t-contrasts for the main effect of average gambling. Hippo = hippocampus, VS = ventral striatum, SN/VTA = substantia nigra and ventral tegmental area.

5.2.2 Iron in SN/VTA and LC is associated with risk aversion and choice stochasticity

Having established a relationship between R1 and R2* maps and a model-free measure of risk, I next turned to a model-based approach which has the added advantage of showing how tissue microstructure is related to *specific processes* underlying risky choice (O'Doherty, Hampton, & Kim, 2007). Choice behaviour was fitted to a prospect theory-based parametric decision model that has been used in previous studies (Rutledge, Skandali, Dayan, & Dolan, 2015; Sokol-Hessner et al., 2009) to describe decision-making under risk. The expected utility of the certain options and gambles were determined using the following equations:

$$U_{\text{gamble}} = 0.5(V_{\text{gain}})^{\alpha}$$

$$U_{\text{certain}} = (V_{\text{certain}})^{\alpha}$$

where V_{gain} is the value of the potential gain from a gamble and V_{certain} is the value of the certain option. α alters the degree of curvature of the utility function and represents the degree of risk aversion. When presented with an option where the expected values for the certain gain and the gamble are equal, a subject with $\alpha = 1$ would be risk-neutral and indifferent between the two, a risk-seeking individual with $\alpha > 1$ would choose the gamble more often, and a risk-averse individual with $\alpha < 1$ would choose the certain gain more often. The probability of selecting a gamble was determined by the following softmax rule:

$$P_{\text{gamble}} = \frac{1}{1 + e^{-\beta(U_{\text{gamble}} - U_{\text{certain}})}}$$

where the degree of stochasticity in choice behavior is captured by the inverse temperature parameter β . When β is low, participants are more likely to choose randomly between safe and risky options irrespective of their subjective values. When β is high, participants increasingly choose the action leading to the highest expected reward. This model provided a good fit for choice behaviour with an average pseudo- R^2 of 0.45 (SD: 0.16).

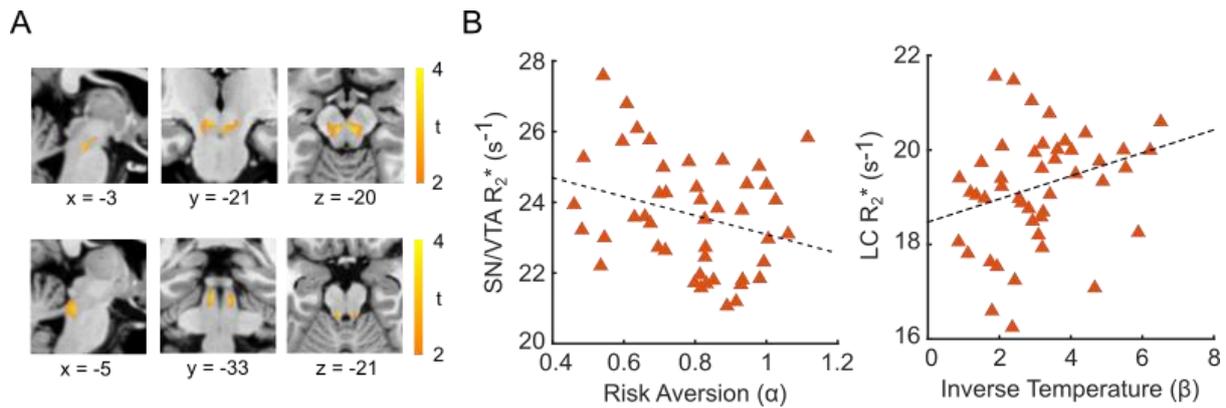


Fig. 5.3. VBQ ROI analysis and model parameters. **A**, Iron reduction in SN/VTA with higher risk aversion parameter is estimated from decreased R_2^* (Top). Increased iron in LC with higher inverse temperature parameter is estimated from increased R_2^* (Bottom). **B**, Correlation between mean R_2^* within the SN/VTA ROI and risk aversion parameter was just above the significance threshold ($\rho = -0.28$, $P = 0.058$, *Left*) while mean R_2^* within the LC ROI was positively correlated with the inverse temperature parameter (Spearman $\rho = 0.34$, $P = 0.02$, Fig 5.3B, *Right*).

Similar to the model-free analysis, risk aversion (α) and inverse temperature (β) parameters were placed in a multiple regression analysis with age, gender, and total intracranial volume included as nuisance covariates. Higher risk aversion parameters related to decreased R_2^* in the SN/VTA (Fig. 5.3A, *Top*) whereas higher inverse temperature parameters related to increased R_2^* in the LC (Fig. 5.3A, *Bottom*). VBQ measures extracted using the same SN/VTA ROI as before revealed that the negative correlation between SN/VTA R_2^* and risk aversion parameter was just above the significance threshold ($\rho = -0.28$, $P = 0.058$, Fig. 5.3B, *Left*) suggesting that a sub-region of the SN/VTA may be more important for driving this association. LC R_2^* and the inverse temperature parameter was significantly correlated (Spearman $\rho = 0.34$, $P = 0.02$, Fig. 5.3B, *Right*) suggesting that individuals who are more deterministic in their choices possess higher levels of iron in the LC. See Table 5.2 for summary of these results.

Table 5.2. ROI results for risk aversion (α) and inverse temperature (β).

Parameter	Map	Region	P_{FWE}	k_E	t	z	x	y	z
α	R_2^* -	R SN/VTA	0.03	14	3.63	3.35	5	-22	-16
α	R_2^* -	L SN/VTA	0.01	18	3.62	3.34	-9	-25	-18
β	R_2^* +	L LC	0.002	23	3.77	3.47	-3	-32	-19

+ and - refer to positive and negative t-contrasts for the main effect of each computational parameter. SN/VTA = substantia nigra and ventral tegmental area, LC = locus coeruleus.

5.3 Discussion

Recent developments in the field of quantitative MRI have permitted in vivo mapping of biologically relevant measures to be executed with high resolution. The findings here demonstrate that individual differences in risky decision-making is related to quantitative measures of myelination and iron in the brain. Previous studies have identified several key areas involved in risky decision-making such as VS, vmPFC, OFC, and SN/VTA (Bartra et al., 2013; D'Ardenne et al., 2008; Levy & Glimcher, 2012; Nieuwenhuis, Aston-Jones, & Cohen, 2005; Padoa-Schioppa & Cai, 2011; Polezzi, Sartori, Rumiati, Vidotto, & Daum, 2010; Rangel et al., 2008). I extend these findings using quantitative magnetic resonance imaging to demonstrate that risk preference is related to myelination in the VS and right hippocampus, as well as iron levels in SN/VTA. The peak R_1 observed in left VS is located in close proximity to coordinates reported in a different study (Peters & Büchel, 2009) that also found subjective value coding of both delayed and probabilistic rewards in the left VS, suggesting that levels of myelination may be involved in the coding of subjective value. That myelination in right hippocampus is related to risky decision-making here is interesting especially in the absence of outcomes and future work may build on these results by investigating its functional role in risk.

The model-based results suggest that markers of microstructural iron in LC is linked to individual variation in choice stochasticity with more deterministic individuals displaying higher R_2^* . The LC is a key component of the noradrenaline system and has been implicated in attention (Aston-Jones, Rajkowski, & Cohen, 2000; Nieuwenhuis et al., 2005), surprise (Clewett, Schoeke, & Mather, 2014; Preuschoff, 't Hart, & Einhauser, 2011), emotional memory (Hämmerer et al., 2018, 2017),

modulation of neural gain (Eldar, Cohen, & Niv, 2013), and metacognitive performance (Hauser, Allen, et al., 2017). While I can only speculate as to the exact process underlying the observed relationship here, attention and gain appear to be viable candidates with propranolol (a β -adrenergic blocker) administration in previous studies impairing attention that could result in choices that are more stochastic (De Martino, Strange, & Dolan, 2008), and lowered neural gain as a function of malfunctioning catecholamine systems resulting in increased choice variability (Hauser, Fiore, Moutoussis, & Dolan, 2016).

While just above the significance threshold, the relation between SN/VTA R_2^* and risk aversion parameter is interesting to speculate on. The main iron compound found in SN/VTA dopamine neurons and LC noradrenaline neurons is the neuromelanin-iron complex (Zucca et al., 2017). Although the interplay between iron, dopamine, and neuromelanin is a complicated process and equilibrium between the three is essential for homeostasis in cells, increased iron deposition has been observed in the brains of Parkinson's disease patients (Stankiewicz et al., 2007; Ulla et al., 2013) which is a disease characterized by the degeneration and loss of dopamine neurons in the SN (Damier, Hirsch, Agid, & Graybiel, 1999; Sofic et al., 1988). Although the link between iron levels and quantity of dopamine neurons is not entirely straightforward, the negative correlation between SN/VTA R_2^* with both average gambling and the risk aversion parameter is at least consistent with findings that pharmacologically increasing dopamine levels leads to increased gambling (Rigoli et al., 2016; Rutledge et al., 2015). Furthermore, iron accumulation is observed in ageing which is also related to a decline in the dopamine system and reduced risk taking for potential rewards (Rutledge et al., 2016; Stankiewicz et al., 2007).

One limitation of the study is that the R_1 and R_2^* quantitative maps may also be sensitive to factors other than the ones facilitating straightforward interpretations like myelin and hepatic iron content. For example, R_1 values could be influenced by axonal density (Gouw et al., 2008) and other variations in microstructure (Harkins et al., 2016) with iron and other macromolecules also playing a potential role (Callaghan, Helms, Lutti, Mohammadi, & Weiskopf, 2015) while some variation in R_2^* could be due to fibre orientation related to the position of participants within MRI scanner bore (Bender & Klose, 2010; Wharton & Bowtell, 2012). However, the close correspondence between measures of R_2^* in the SN/VTA here and a previous study (Callaghan et al., 2014) is

reassuring and suggests that differences due to such factors may not be too large a confound for interpretation. A second limitation of the study is the use of ROI analyses on the basis of regions thought to be involved in the decision network. Peak locations of effects of interest, such as representation of subjective value in the vmPFC, often vary from study to study and although significant effects were found in VS, SN/VTA, LC, and right hippocampus here, these were missing in the OFC and vmPFC which raises questions for future work regarding the biological substrates of certain quantities assumed to play a role in risky decision-making.

In this study, I used hMRI to reveal the association between risk preferences and tissue microstructure. The model-free results suggest that myelin and iron in the VS, SN/VTA, and right hippocampus are key predictors of average gambling while the model-based results suggest that levels of iron in the LC is predictive of how deterministic an individual will be in their choice behaviour. As a whole, these results suggest a link between processes underlying decision-making and existing microstructure of each individual's neuroanatomy beyond grey and white matter volume. Future work may use a richer task design to better isolate components of risk and reward to extend these findings, especially since the quantitative nature of MPMs can be used for multi-site comparisons and facilitate reproducible research.

References

- Aston-Jones, G., Rajkowski, J., & Cohen, J. (2000). Locus coeruleus and regulation of behavioral flexibility and attention. *Progress in Brain Research*, 126(2000), 165–182. [https://doi.org/10.1016/S0079-6123\(00\)26013-5](https://doi.org/10.1016/S0079-6123(00)26013-5)
- Balteau, E., Tabelow, K., Aeshburner, Callaghan, M. F., Draganski, B., Helms, G., Mohammadi, S. (2018). HMRI -- A toolbox for using quantitative MRI in neuroscience and clinical research. *Weierstrass Institute for Applied Analysis and Stochastics: Preprint 2527*. Retrieved from <http://www.wias-berlin.de/publications/wias-publ/run.jsp?template=abstract&type=Preprint&year=2018&number=2527>
- Bär, K.-J., de la Cruz, F., Schumann, A., Koehler, S., Sauer, H., Critchley, H., & Wagner, G. (2016). Functional connectivity and network analysis of midbrain and brainstem nuclei. *NeuroImage*, 134, 53–63. <https://doi.org/10.1016/j.neuroimage.2016.03.071>
- Bartra, O., McGuire, J. T., & Kable, J. W. (2013). The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *NeuroImage*, 76, 412–427. <https://doi.org/10.1016/j.neuroimage.2013.02.063>
- Bender, B., & Klose, U. (2010). The in vivo influence of white matter fiber orientation towards B(0) on T2* in the human brain. *NMR in Biomedicine*, 23(9), 1071–1076. <https://doi.org/10.1002/nbm.1534>
- Callaghan, M. F., Freund, P., Draganski, B., Anderson, E., Cappelletti, M., Chowdhury, R., Weiskopf, N. (2014). Widespread age-related differences in the human brain microstructure revealed by quantitative magnetic resonance imaging. *Neurobiology of Aging*, 35(8), 1862–1872. <https://doi.org/10.1016/j.neurobiolaging.2014.02.008>
- Callaghan, M. F., Helms, G., Lutti, A., Mohammadi, S., & Weiskopf, N. (2015). A general linear relaxometry model of R1 using imaging data. *Magnetic Resonance in Medicine*, 73(3), 1309–1314. <https://doi.org/10.1002/mrm.25210>
- Castella, R., Arn, L., Dupuis, E., Callaghan, M. F., Draganski, B., & Lutti, A. (2018). Controlling motion artefact levels in MR images by suspending data acquisition during periods of head motion. *Magnetic Resonance in Medicine*, 80(6), 2415–2426. <https://doi.org/10.1002/mrm.27214>
- Clewett, D., Schoeke, A., & Mather, M. (2014). Locus coeruleus neuromodulation of memories encoded during negative or unexpected action outcomes. *Neurobiology of Learning and Memory*, 111, 65–70. <https://doi.org/10.1016/j.nlm.2014.03.006>
- Damier, P., Hirsch, E. C., Agid, Y., & Graybiel, A. M. (1999). The substantia nigra of the human brain II. Patterns of loss of dopamine-containing neurons in Parkinson's disease. *Brain*, 122(8), 1437–1448. <https://doi.org/10.1093/brain/122.8.1437>

- D'Ardenne, K., McClure, S. M., Nystrom, L. E., & Cohen, J. D. (2008). BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. *Science*, *319*(5867), 1264–1267. <https://doi.org/10.1126/science.1150605>
- Daugherty, A., & Raz, N. (2013). Age-related differences in iron content of subcortical nuclei observed in vivo: A meta-analysis. *NeuroImage*, *70*, 113–121. <https://doi.org/10.1016/j.neuroimage.2012.12.040>
- De Martino, B., Strange, B. A., & Dolan, R. J. (2008). Noradrenergic neuromodulation of human attention for emotional and neutral stimuli. *Psychopharmacology*, *197*(1), 127–136. <https://doi.org/10.1007/s00213-007-1015-5>
- Draganski, B., Ashburner, J., Hutton, C., Kherif, F., Frackowiak, R. S. J., Helms, G., & Weiskopf, N. (2011). Regional specificity of MRI contrast parameter changes in normal ageing revealed by voxel-based quantification (VBQ). *NeuroImage*, *55*(4), 1423–1434. <https://doi.org/10.1016/j.neuroimage.2011.01.052>
- Eldar, E., Cohen, J. D., & Niv, Y. (2013). The effects of neural gain on attention and learning. *Nature Neuroscience*, *16*(8), 1146–1153. <https://doi.org/10.1038/nn.3428>
- Gouw, A. A., Seewann, A., Vrenken, H., van der Flier, W. M., Rozemuller, J. M., Barkhof, F., Geurts, J. J. G. (2008). Heterogeneity of white matter hyperintensities in Alzheimer's disease: Post-mortem quantitative MRI and neuropathology. *Brain*, *131*(12), 3286–3298. <https://doi.org/10.1093/brain/awn265>
- Hämmerer, D., Callaghan, M. F., Hopkins, A., Kosciessa, J., Betts, M., Cardenas-Blanco, A., Düzel, E. (2018). Locus coeruleus integrity in old age is selectively related to memories linked with salient negative events. *PNAS*, *115*(9), 2228–2233. <https://doi.org/10.1073/pnas.1712268115>
- Hämmerer, D., Hopkins, A., Betts, M. J., Maaß, A., Dolan, R. J., & Düzel, E. (2017). Emotional arousal and recognition memory are differentially reflected in pupil diameter responses during emotional memory for negative events in younger and older adults. *Neurobiology of Aging*, *58*, 129–139. <https://doi.org/10.1016/j.neurobiolaging.2017.06.021>
- Harkins, K. D., Xu, J., Dula, A. N., Li, K., Valentine, W. M., Gochberg, D. F., Does, M. D. (2016). The microstructural correlates of T1 in white matter. *Magnetic Resonance in Medicine*, *75*(3), 1341–1345. <https://doi.org/10.1002/mrm.25709>
- Hauser, T. U., Allen, M., Purg, N., Moutoussis, M., Rees, G., & Dolan, R. J. (2017). Noradrenaline blockade specifically enhances metacognitive performance. *ELife*, *6*, e24901. <https://doi.org/10.7554/eLife.24901>

- Hauser, T. U., Eldar, E., & Dolan, R. J. (2017). Separate mesocortical and mesolimbic pathways encode effort and reward learning signals. *PNAS*, 7395–7404. <https://doi.org/10.1073/pnas.1705643114>
- Hauser, T. U., Fiore, V. G., Moutoussis, M., & Dolan, R. J. (2016). Computational psychiatry of ADHD: Neural gain impairments across marrian levels of analysis. *Trends in Neurosciences*, 39(2), 63–73. <https://doi.org/10.1016/j.tins.2015.12.009>
- Helms, G., Dathe, H., Kallenberg, K., & Dechent, P. (2008). High-resolution maps of magnetization transfer with inherent correction for RF inhomogeneity and T1 relaxation obtained from 3D FLASH MRI. *Magnetic Resonance in Medicine*, 60(6), 1396–1407. <https://doi.org/10.1002/mrm.21732>
- Levy, D. J., & Glimcher, P. W. (2012). The root of all value: A neural common currency for choice. *Current Opinion in Neurobiology*, 22(6), 1027–1038. <https://doi.org/10.1016/j.conb.2012.06.001>
- Lorio, S., Fresard, S., Adaszewski, S., Kherif, F., Chowdhury, R., Frackowiak, R. S., Draganski, B. (2016). New tissue priors for improved automated classification of subcortical brain structures on MRI. *NeuroImage*, 130, 157–166. <https://doi.org/10.1016/j.neuroimage.2016.01.062>
- Lutti, A., Dick, F., Sereno, M. I., & Weiskopf, N. (2014). Using high-resolution quantitative mapping of R1 as an index of cortical myelination. *NeuroImage*, 93 Pt 2, 176–188. <https://doi.org/10.1016/j.neuroimage.2013.06.005>
- Nieuwenhuis, S., Aston-Jones, G., & Cohen, J. D. (2005). Decision making, the P3, and the locus coeruleus—Norepinephrine system. *Psychological Bulletin*, 131(4), 510–532. <https://doi.org/10.1037/0033-2909.131.4.510>
- O’Doherty, J. P., Hampton, A., & Kim, H. (2007). Model-based fMRI and its application to reward learning and decision making. *Annals of the New York Academy of Sciences*, 1104, 35–53. <https://doi.org/10.1196/annals.1390.022>
- Padoa-Schioppa, C. (2013). Neuronal origins of choice variability in economic decisions. *Neuron*, 80(5), 1322–1336. <https://doi.org/10.1016/j.neuron.2013.09.013>
- Padoa-Schioppa, C., & Cai, X. (2011). Orbitofrontal cortex and the computation of subjective value: Consolidated concepts and new perspectives. *Annals of the New York Academy of Sciences*, 1239, 130–137. <https://doi.org/10.1111/j.1749-6632.2011.06262.x>
- Peters, J., & Büchel, C. (2009). Overlapping and distinct neural systems code for subjective value during intertemporal and risky decision making. *Journal of Neuroscience*, 29(50), 15727–15734. <https://doi.org/10.1523/JNEUROSCI.3489-09.2009>

- Polezzi, D., Sartori, G., Rumiati, R., Vidotto, G., & Daum, I. (2010). Brain correlates of risky decision-making. *NeuroImage*, *49*(2), 1886–1894.
<https://doi.org/10.1016/j.neuroimage.2009.08.068>
- Preuschoff, K., 't Hart, B. M., & Einhauser, W. (2011). Pupil dilation signals surprise: Evidence for noradrenaline's role in decision making. *Frontiers in Neuroscience*, *5*.
<https://doi.org/10.3389/fnins.2011.00115>
- Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, *9*(7), 545. <https://doi.org/10.1038/nrn2357>
- Ridgway, G. R., Henley, S. M. D., Rohrer, J. D., Scahill, R. I., Warren, J. D., & Fox, N. C. (2008). Ten simple rules for reporting voxel-based morphometry studies. *NeuroImage*, *40*(4), 1429–1435. <https://doi.org/10.1016/j.neuroimage.2008.01.003>
- Rigoli, F., Rutledge, R. B., Chew, B., Ousdal, O. T., Dayan, P., & Dolan, R. J. (2016). Dopamine increases a value-independent gambling propensity. *Neuropsychopharmacology*, *41*(11), 2658–2667. <https://doi.org/10.1038/npp.2016.68>
- Rutledge, R. B., Skandali, N., Dayan, P., & Dolan, R. J. (2014). A computational and neural model of momentary subjective well-being. *PNAS*, *111*(33), 12252–12257.
<https://doi.org/10.1073/pnas.1407535111>
- Rutledge, R. B., Skandali, N., Dayan, P., & Dolan, R. J. (2015). Dopaminergic modulation of decision making and subjective well-being. *Journal of Neuroscience*, *35*(27), 9811–9822. <https://doi.org/10.1523/JNEUROSCI.0702-15.2015>
- Rutledge, R. B., Smittenaar, P., Zeidman, P., Brown, H. R., Adams, R. A., Lindenberger, U., Dolan, R. J. (2016). Risk taking for potential reward decreases across the lifespan. *Current Biology*, *26*(12), 1634–1639.
<https://doi.org/10.1016/j.cub.2016.05.017>
- Sofic, E., Riederer, P., Heinsen, H., Beckmann, H., Reynolds, G. P., Hebenstreit, G., & Youdim, M. B. H. (1988). Increased iron (III) and total iron content in post mortem substantia nigra of parkinsonian brain. *Journal of Neural Transmission*, *74*(3), 199–205. <https://doi.org/10.1007/BF01244786>
- Sokol-Hessner, P., Hsu, M., Curley, N. G., Delgado, M. R., Camerer, C. F., & Phelps, E. A. (2009). Thinking like a trader selectively reduces individuals' loss aversion. *PNAS*, *106*(13), 5035–5040. <https://doi.org/10.1073/pnas.0806761106>
- Stankiewicz, J., Panter, S. S., Neema, M., Arora, A., Batt, C., & Bakshi, R. (2007). Iron in chronic brain disorders: Imaging and neurotherapeutic implications. *Neurotherapeutics: The Journal of the American Society for Experimental NeuroTherapeutics*, *4*(3), 371–386. <https://doi.org/10.1016/j.nurt.2007.05.006>

- Tabelow, K., Balteau, E., Ashburner, J., Callaghan, M. F., Draganski, B., Helms, G., Mohammadi, S. (2019). HMRI – A toolbox for quantitative MRI in neuroscience and clinical research. *NeuroImage*, *194*, 191–210.
<https://doi.org/10.1016/j.neuroimage.2019.01.029>
- Ulla, M., Bonny, J. M., Ouchchane, L., Rieu, I., Claise, B., & Durif, F. (2013). Is $r2^*$ a new MRI biomarker for the progression of parkinson's disease? A longitudinal follow-up. *PLOS ONE*, *8*(3), e57904. <https://doi.org/10.1371/journal.pone.0057904>
- Weiskopf, N., Callaghan, M. F., Josephs, O., Lutti, A., & Mohammadi, S. (2014). Estimating the apparent transverse relaxation time ($R2^*$) from images with different contrasts (ESTATICS) reduces motion artifacts. *Frontiers in Neuroscience*, *8*.
<https://doi.org/10.3389/fnins.2014.00278>
- Weiskopf, N., Suckling, J., Williams, G., Correia, M. M., Inkster, B., Tait, R., Lutti, A. (2013). Quantitative multi-parameter mapping of $R1$, PD^* , MT , and $R2^*$ at 3T: A multi-center validation. *Frontiers in Neuroscience*, *7*. <https://doi.org/10.3389/fnins.2013.00095>
- Wharton, S., & Bowtell, R. (2012). Fiber orientation-dependent white matter contrast in gradient echo MRI. *PNAS*, *109*(45), 18559–18564.
<https://doi.org/10.1073/pnas.1211075109>
- Wood, J. C., Enriquez, C., Ghugre, N., Tyzka, J. M., Carson, S., Nelson, M. D., & Coates, T. D. (2005). MRI $R2$ and $R2^*$ mapping accurately estimates hepatic iron concentration in transfusion-dependent thalassemia and sickle cell disease patients. *Blood*, *106*(4), 1460–1465. <https://doi.org/10.1182/blood-2004-10-3982>
- Zucca, F. A., Segura-Aguilar, J., Ferrari, E., Muñoz, P., Paris, I., Sulzer, D., Zecca, L. (2017). Interactions of iron, dopamine and neuromelanin pathways in brain aging and Parkinson's disease. *Progress in Neurobiology*, *155*, 96–119.
<https://doi.org/10.1016/j.pneurobio.2015.09.012>

Chapter 6: A neurocomputational model of mood and intrinsic rewards

All men seek happiness. This is without exception. Whatever different means they employ, they all tend to this end. The cause of some going to war, and of others avoiding it, is the same desire in both, attended with different views. The will never takes the least step but to this object. This is the motive of every action of every man, even of those who hang themselves.

Blaise Pascal

Previous chapters have touched on the contributions of neural variability and neuroanatomy to risk. Common across many tasks involving the assessment of risk preferences is the use of rewards like food or money. However, there also exist rewards that are intrinsically experienced by an individual and these are more difficult to measure due to the lack of economic markets for such goods. This chapter examines the influence of intrinsic rewards on mood and discusses a framework that can be used to estimate values of intangible goods or experiences.

A key index of quality of life is subjective well-being which “refers to how people experience and evaluate their lives and specific domains and activities in their lives” (Oswald & Wu, 2010). Individuals with higher subjective well-being display lower mortality rates (Chida & Steptoe, 2008; Steptoe, Deaton, & Stone, 2015) and are less at risk of disease (Davidson, Mostofsky, & Whang, 2010). In the workplace, employees who possess higher subjective well-being demonstrate increased productivity without a loss of output quality (Oswald, Proto, & Sgroi, 2015), have reduced rates of absenteeism (Pelled & Xin, 1999), are rated more positively by their supervisors, and also produce greater financial performance (Peterson, Luthans, Avolio, Walumbwa, & Zhang, 2011). Maximising subjective well-being should therefore be of interest to individuals, organisations, and governments, and be considered a target for health and economic policies (Dolan & White, 2007).

The problem arises when it comes to designing measures that are likely to increase well-being. When contemplating the future, people routinely engage in *affective forecasting* where they make predictions about what it would feel like to experience specific events like winning the lottery or meeting their favourite celebrity. However, studies on affective forecasting have revealed that people habitually

misjudge how future events would impact on their emotional states, occasionally leading them to perform actions that may be detrimental to the maximization of their subjective well-being (Meyvis, Ratner, & Levav, 2010; Wilson & Gilbert, 2005). This is known as the impact bias where people overestimate both the intensities and durations of their hedonic responses to future events (Gilbert & Wilson, 2007; Morewedge & Buechel, 2013). Whereas the value of tangible goods can be assessed by their net economic prices or willingness-to-pay (Plassmann, O'Doherty, & Rangel, 2007), the value of intangible goods (e.g. autotelic experiences) are more difficult to define or elicit accurately due to biases (Nisbet & Zelenski, 2011; Van de Mortel, 2008) and mixed findings regarding the predictive validity of implicit measures (Keatley, Clarke, & Hagger, 2013; Levesque, Copeland, & Sutcliffe, 2008). Neuroscience-informed methods involving self-reported values of intrinsic rewards – such as the experience of mastering a musical composition for its own sake – and neural measurements pertaining to those values may provide some insight as to how the true valuation of intrinsic rewards may be obtained relative to money (Krajbich, Camerer, Ledyard, & Rangel, 2009).

Drawing on these ideas, here we hypothesized that the experience of skilled performance relating to intrinsic rewards would influence the momentary happiness of participants in a similar vein as extrinsic rewards. In particular, we hypothesized that the extent to which momentary happiness was influenced by both types of reward would be reflected in the brain, representing the subjective value of rewards. Recent studies (Rutledge, Skandali, Dayan, & Dolan, 2014, 2015; Vinckier, Rigoux, Oudiette, & Pessiglione, 2018) have demonstrated a use for experience sampling as an effective approach to relating rewards and subjective feelings when affective and motivational responses are elicited by extrinsic rewards, and here we extend these methods to investigate how momentary happiness is influenced by intrinsic rewards.

To test our hypothesis, we developed a reinforcement learning task incorporating both an explicit reward component and a skilled performance component that did not affect payment. On each trial, participants chose between two boxes, one of which was more rewarding than the other, before having to navigate a cursor through a series of barriers. Over the course of the experiment, the box containing the higher reward switched and when the end of the track was reached, they received

some points from their chosen box. Performance was titrated using a continuous staircase method to keep performance at a similar level across participants.

6.1 Methods

6.1.1 Participants

37 healthy, young adults (age: 25.8 ± 4.7 , mean \pm SD) were recruited through the University College London Psychology Subject Database. Subjects were screened to ensure no history of neurological or psychiatric disorders. 4 participants were excluded due to excessive head movements during scanning, leaving a total of 33 participants (age: 26.1 ± 4.9) included in the study. The study was approved by the University College London (UCL) research ethics committee, and all participants gave written informed consent.

6.1.2 Procedure

Participants completed the experiment at the Wellcome Centre for Human Neuroimaging, University College London, in an appointment that lasted approximately 90 minutes.

Experimental Task

Stimuli were presented in MATLAB (MathWorks, Inc.) using Cogent 2000. The layout of each trial resembled a T-Maze (Howe, Tierney, Sandberg, Phillips, & Graybiel, 2013). On each trial, participants first chose between a blue or magenta box, one of which contains a higher number of points (Mean: 50 ± 10 , mean \pm SD) than the other average (Mean: 25 ± 10) on. Points from the better box were drawn from a Gaussian with a mean of 50 and standard deviation of 10, while points from the worse box were drawn from a Gaussian with a mean of 25 and standard deviation of 10. Every 19 to 23 trials, a reversal occurred where the box that previously contained the higher number of points on average now contained a lower number of points and vice versa for the other box. On half the trials, participants were afforded a free choice and for the remaining half, participants had were only presented with a single option.

After a choice was made, the chosen option was highlighted and 4 barriers appeared on the screen along with a small square at the bottom of the screen.

Following a 1s delay, the small square automatically moved up at the T-Maze towards the chosen option. Participants were able to control the horizontal position of the moving square to avoid crashing into the barriers. If they passed a barrier without crashing, the barrier turned green. Otherwise, the barrier turned red providing immediate feedback about performance. Crucially, the participants' final payment depended only on the number of points accumulated across the experiment and not their performance on each trial. After the cursor had entered the chosen box, the outcomes were displayed for 800ms following a 1.5s delay. Total cumulated points was displayed to the participants on the top right of the screen throughout the experiment. Participants were presented with the question, "How happy are you at this moment?" after every two to three trials. After a 1s delay period, a rating line appeared and participants had 4s to move a cursor, which always started at the midpoint, along the scale with button presses. The left end of the line was labelled "very unhappy" and the right end of the line was labelled "very happy".

Staircase Procedure

To ensure that differences in mood-related responses were not due to skill-related differences in how often each participant crashed into the barriers, we used a staircase procedure - Parametric Estimation by Sequential Testing (PEST) (Taylor & Creelman, 1967) - to calibrate the speed at which the cursor moved for every participant such that they would be able to avoid crashing into the barriers on approximately 70% of the trials. This calibration was carried out over 60 trials prior to the start of the task in the scanner.

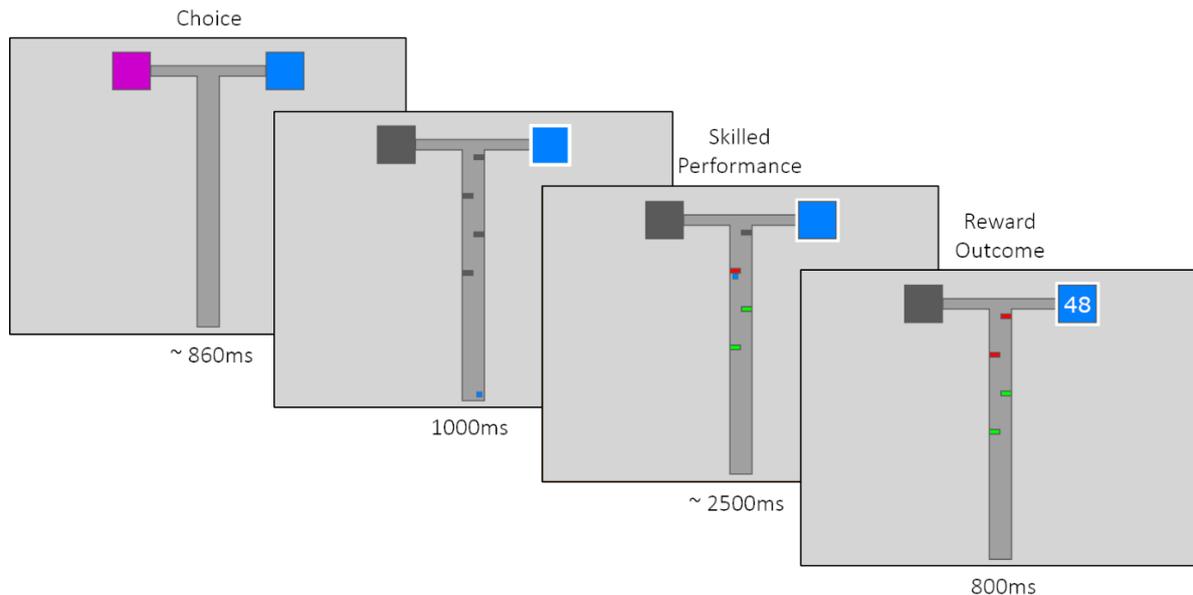


Fig. 6.1. Experiment paradigm. On each trial, subjects choose between two options, one of which leads to a higher number of points than the other on average. The high-reward option reverses every 18-21 trials. After choices are made, selected options are immediately outlined and a cursor appears at the bottom of the T-Maze along with 4 barriers. After a 1000ms delay, the cursor automatically moves up the maze and participants navigate around the barriers with button presses constituting a form of skilled performance that may be intrinsically rewarding for some participants despite the lack of instrumental value. Contacted barriers turn red and avoided barriers turn green. The performance component typically lasted for 2500ms. After the cursor reaches the selected option, the reward outcome (i.e., extrinsic rewards) is displayed for 800ms. A 1500ms delay inter-trial interval follows each trial. After every 2-3 trials, subjects are asked to rate their current happiness by moving a cursor on a line with a maximum time limit of 4000ms.

6.1.3 Image Acquisition

MRI was acquired at the Wellcome Centre for Human Neuroimaging, University College London, using a Siemens Prisma 3-Tesla scanner equipped with a 64-channel head coil. Functional images were acquired with a gradient echo T2*-weighted echo-planar sequence with whole brain coverage and each volume consisted of 48 slices with 3mm isotropic voxels [repetition time (TR): 3.36s; echo time (TE): 30ms; slice tilt: 0°] in ascending order. A field map [double-echo FLASH, TE1 = 10ms, TE2 = 12.46ms] with 3mm isotropic voxels (whole brain coverage) was also acquired for each participant to correct the functional images for any inhomogeneity in magnetic field strength. Subsequently, the first 6 volumes of each run were discarded to allow for T1 saturation effects. Structural images were T1-weighted (1 x 1 x 1 mm resolution) images acquired using a MPRAGE sequence.

6.2 Results

6.2.1 Happiness is modulated by reward outcomes and performance

We first asked whether participants were able to track which box contained higher value outcomes and found that participants were able to learn the reward contingencies (Fig. 6.2A), choosing the option that contained a higher value $84.8\% \pm 5.6$ (mean \pm SD) of the time when afforded a free choice. We next examined the determinants of momentary happiness and found that participants reported greater average happiness after receiving higher than lower outcomes ($t_{32} = 8.4$, $P < 0.001$, Fig. 6.2B), consistent with previous research (Rutledge, Skandali, Dayan, & Dolan, 2015; Rutledge, Skandali, Dayan, & Dolan, 2014). Participants were not penalized or rewarded with additional points whether they had crashed into any of the barriers or successfully avoided all of them, meaning that performance was non-instrumental to the receipt of eventual monetary reward. Despite this, participants also reported being happier when they navigated through the maze without crashing into any barriers compared to when they crashed into at least one barrier ($t_{32} = 6.4$, $P < 0.001$, Fig. 6.2B), suggesting that a determinant of momentary happiness may be intrinsically derived. Accordingly, there was considerable variation across participants in terms of how much extrinsic rewards and skilled performance contributed to momentary happiness (Fig. 6.2C).

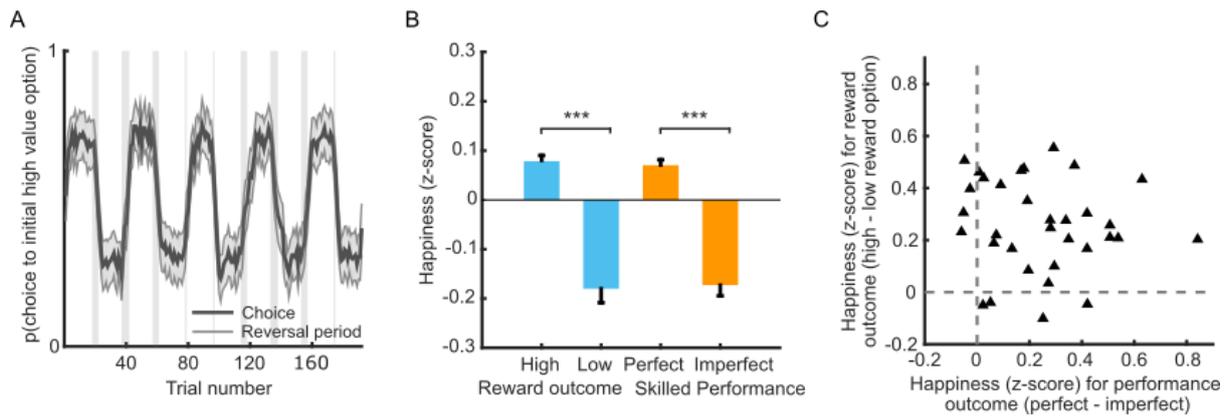


Fig. 6.2. Task behaviour and modulation of happiness by reward outcomes and skilled performance. **A**, Average choice behaviour across participants ($n = 33$) in black with the shaded area corresponding to SEM. Grey vertical bands represent intervals where reversals can occur with timing jittered across participants. High-value options were chosen more often than low-value options, indicating that subjects tracked changing reward contingencies. In half of trials, only one of the two options was available. **B**, Subjects were happier when they received an outcome from the high-value option compared to the low-value option ($t_{32} = 8.4$, $P < 0.001$, in blue). Subjects were happier on average when they navigated through the T-maze without hitting a barrier compared to when they hit at least one barrier ($t_{32} = 6.4$, $P < 0.001$, in orange). **C**, The majority (29 of 33) were happier after receiving a reward from a high-value compared to low-value option and the majority (29 of 33) were happier after achieving perfect compared to imperfect performance. There was no relationship between happiness for reward outcomes compared to happiness for skilled performance ($R = -0.11$, $P = 0.55$). *** $P < 0.001$.

6.2.2 Computational model of subjective well-being

Next, we investigated the relationship between outcomes, performance, and happiness using a computational model of mood. In line with previous literature (Rutledge et al., 2014, 2015), we accounted for the decay of influences over time:

$$\text{Happiness}(t) = w_1 \sum_{j=1}^t \gamma^{t-j} \text{Reward}_j + w_2 \sum_{j=1}^t \gamma^{t-j} \text{Performance}_j$$

where t is the trial number, weights w capture the influence of different task events, $0 \leq \gamma \leq 1$ represents a discount factor that reduces the impact of distal relative to recent events, reward is the z-scored outcome of the chosen box on each trial, and performance is the z-scored result of whether a barrier was hit on each trial where a 1 is assigned if no barriers were hit, and 0 if at least one barrier was hit. Parameters were fit to happiness ratings in each individual subject. This model explained a decent amount of fluctuations in happiness with $r^2 = 0.27 \pm 15.3$ (mean \pm SD, Fig. 6.3A), and

we found that weights for both performance ($t_{32} = 5.79$, $P < 0.001$, Fig. 6.3B) and reward ($t_{32} = 8.27$, $P < 0.001$, Fig. 6.3B) were positive on average. The discount factor γ was 0.44 ± 0.31 (mean \pm SD). Model comparison revealed that this model performed better than models containing individual terms for reward ($r^2 = 0.19 \pm 13.6$) and performance ($r^2 = 0.08 \pm 12.2$), suggesting that the happiness of participants are not solely dependent on the receipt of explicit points, but is also influenced by the non-instrumental experience of skilled performance.

Table 6.1. Model Comparison Results.

Model	Parameters	Mean R^2	BIC	Δ BIC
Reward	2	0.19	-6747	507
Performance	2	0.08	-5907	1347
Reward and Performance	3	0.27	-7254	0
Reward and Performance (Separate discount factors γ)	4	0.28	-7131	123

BIC measures are summed across 33 subjects. The winning model (lowest BIC) here was the model with both reward and performance having identical discount factors. Δ BIC refers to the difference in BIC scores between each model and the winning model.

Despite performance being held constant (Percent perfect: 69 ± 2.4 , mean \pm SD, Fig. 6.3C), we found considerable variation across individuals in how much performance contributed to happiness. Furthermore, these measures were uncorrelated (Pearson's $\rho = 0.26$, $P = 0.15$). Interestingly, the median speed of the cursor and the weight for performance were positively correlated ($R = 0.42$, $P = 0.01$) suggesting that individuals who may have been more proficient at executing a sequence of motor actions or who were more motivated on the task had their happiness impacted by performance by a greater amount. Similar to performance, there was a good amount of variation across individuals in how much explicit outcomes contributed to happiness even as most participants learnt the reward contingencies to a similar extent in terms of the percentage of times they chose the higher value option (Fig. 6.3D).

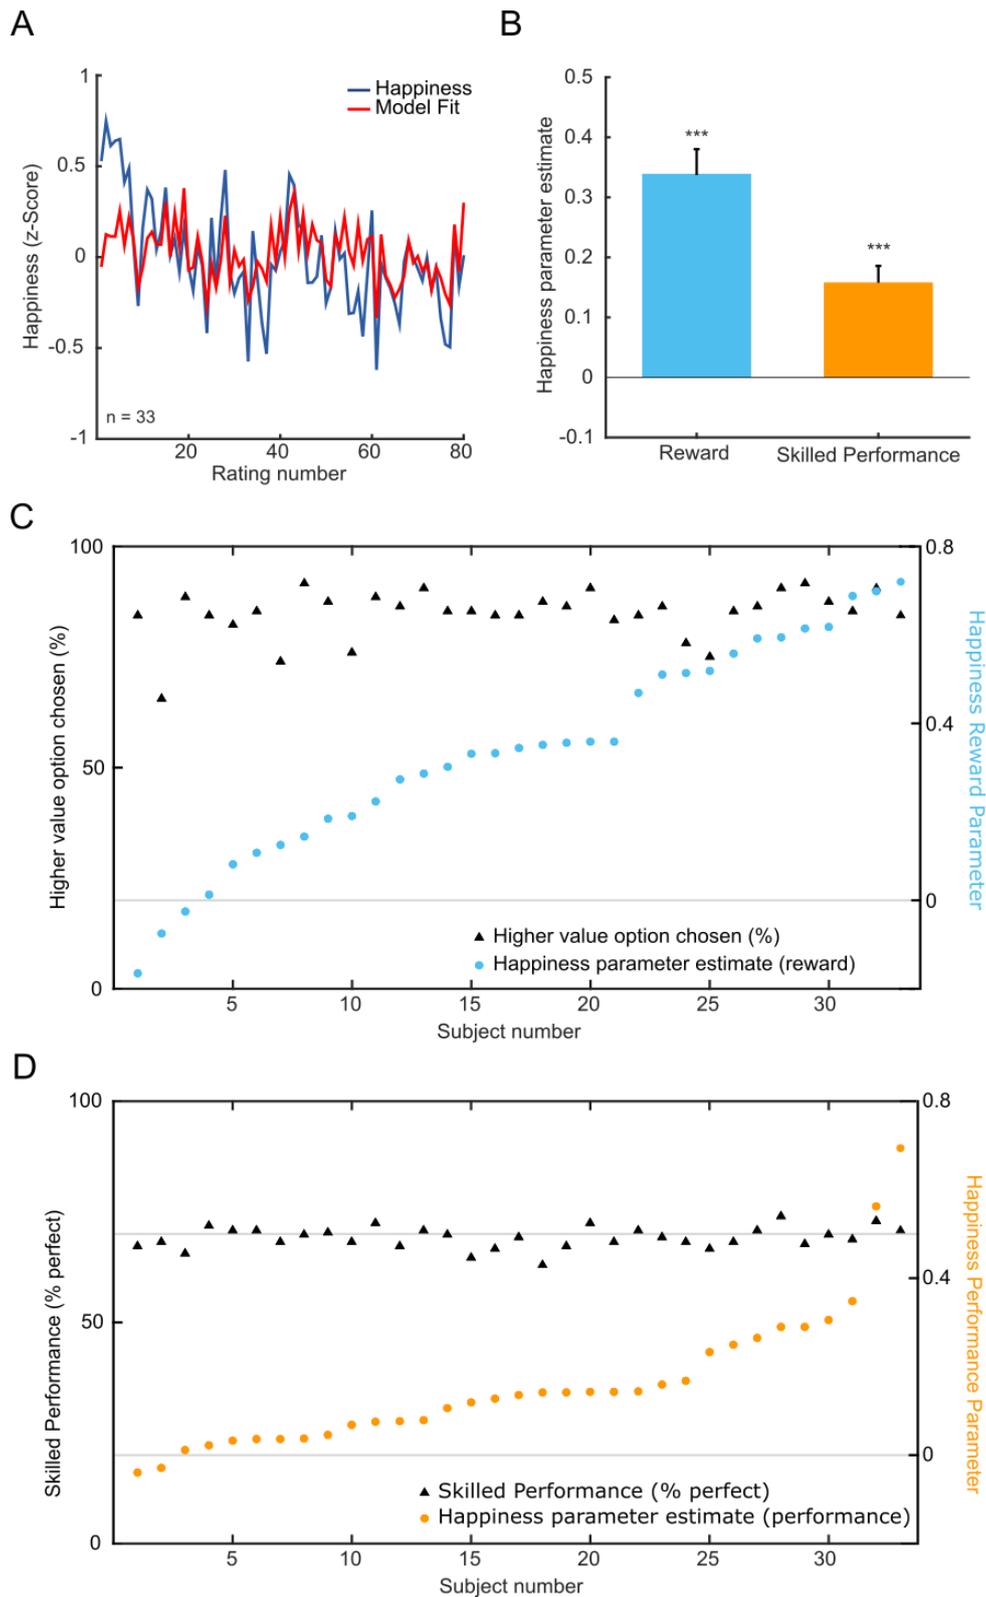


Fig. 6.3. Computational modelling of mood. **A**, Average happiness across the task and model fit is displayed for the computational model of mood ($n=33$, mean $r^2 = 0.27$). **B**, Happiness was significantly related to the history of extrinsic rewards in the form of points ($t_{32} = 8.27$, $P < 0.001$) and also to the history of skilled performance, a proxy for intrinsic rewards ($t_{32} = 5.79$, $P < 0.001$). **C**, **D**, The contribution of reward to happiness varied across subjects despite a similar high choice accuracy across subjects. Despite titrating difficulty to match performance around 70%, subjects displayed considerable variation in weights for performance impacts on happiness in the computational model. *** $P < 0.001$.

6.2.3 Neural correlates of extrinsic rewards and skilled performance

Having established inter-individual variability in the influence of outcomes and performance on happiness, we then asked whether this variability was also predictive of neural responses to extrinsic rewards and performance. First, we regressed event-related activity on parametrically-modulated task events to find areas of the brain that could reflect these quantities. We found an effect of extrinsic rewards at time of outcome display in vmPFC (Fig. 6.4A, top: -3, 38, -1; $t_{32} = 5.92$, $P < 0.05$ Family-Wise-Error (FWE) cluster-corrected), and an effect of skilled performance in a region of the vmPFC anterior to that (Fig. 6.4A, bottom: -3, 50, -1; $t_{32} = 4.24$, $P < 0.05$ FWE cluster-corrected). Next, we extracted the weights for reward and skilled performance using an independent vmPFC mask (Bartra, McGuire, & Kable, 2013) and found that similar to our happiness model, the vmPFC weights in this unbiased mask for both extrinsic rewards ($t_{32} = 3.36$, $P = 0.002$) and skilled performance ($t_{32} = 2.90$, $P = 0.007$) were significantly positive (Fig. 6.4B) suggesting an association between vmPFC activity and these quantities.

After determining that neural responses in the vmPFC were associated with both extrinsic rewards and skilled performance, we next examined whether these responses differed in individuals whose momentary happiness was more strongly influenced by one quantity over the other. We first split participants into two groups depending on whether they had a higher happiness-model-derived weight for performance compared to reward and vice versa. We found that the group with larger happiness weights for performance had a significant positive vmPFC weight for performance (signrank: $P = 0.004$) but not extrinsic rewards ($P = 0.13$), while the group with larger happiness weights for extrinsic rewards had a significant positive vmPFC weight for extrinsic rewards ($P = 0.01$) but not performance ($P = 0.2$). The former group also showed greater responses in vmPFC for skilled performance compared to the group with larger happiness weights for explicit rewards ($P = 0.003$, Fig. 6.4C), suggesting that we can distinguish between the two groups based on neural responses to intrinsic rewards (experience of skilled performance).

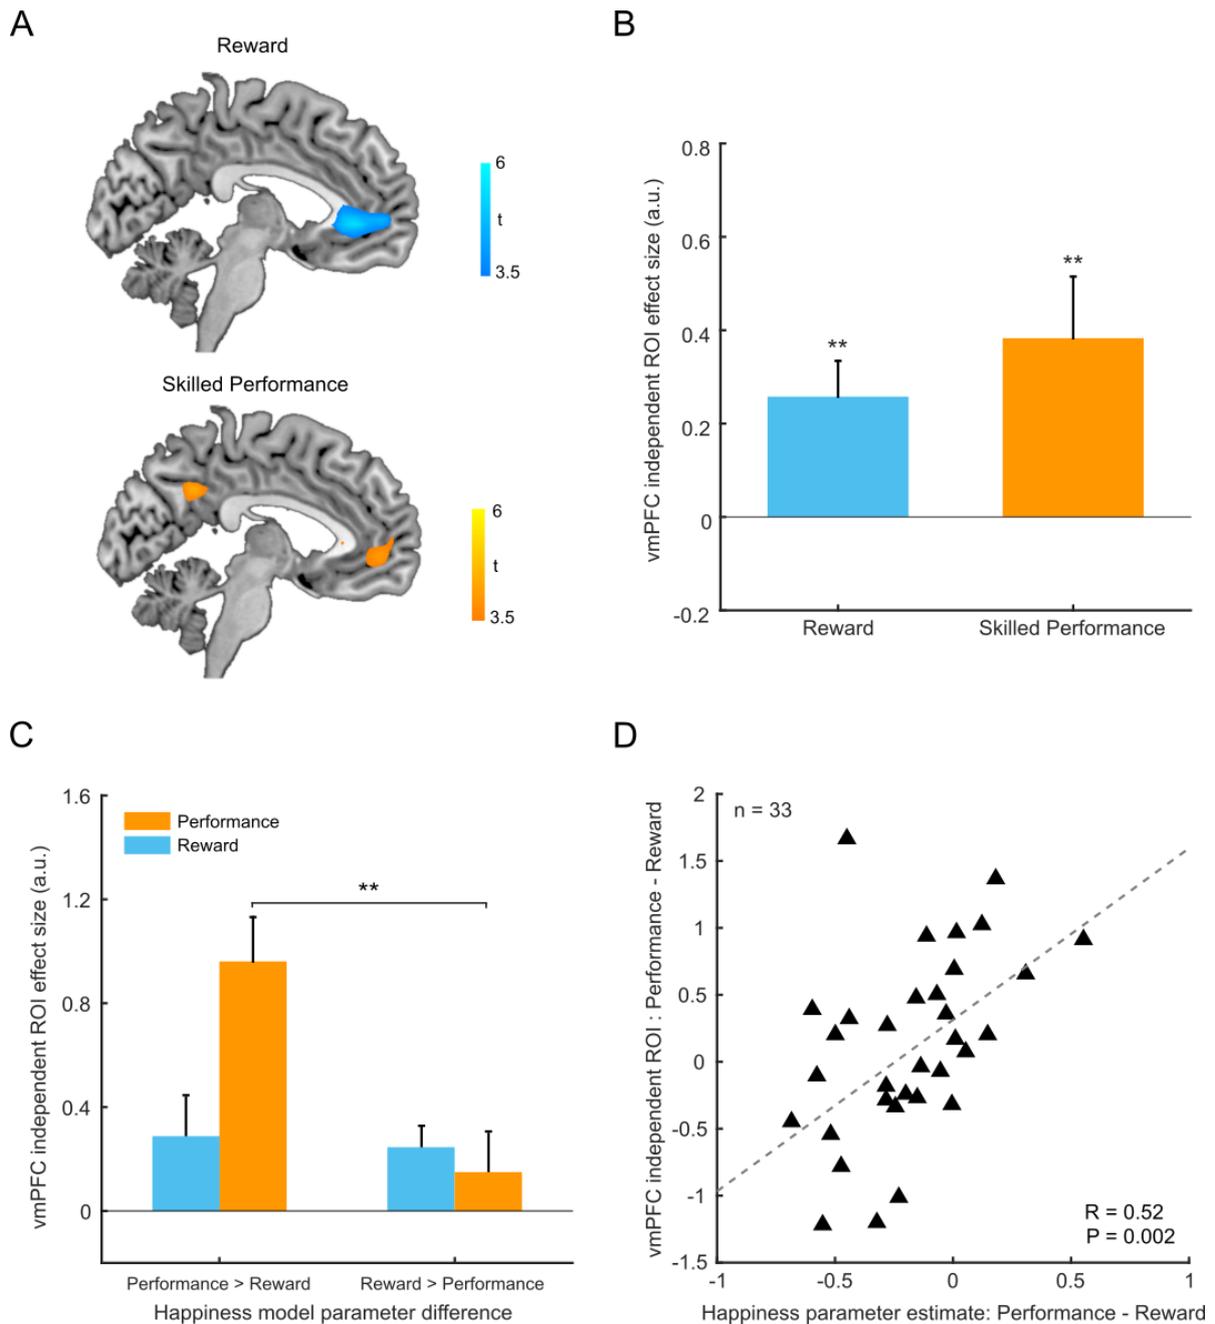


Fig. 6.4. Relative impacts of reward and performance on mood predict vmPFC responses. **A**, *Top*. BOLD activity in vmPFC was parametrically modulated by reward outcome (Peak: -3,38,-1). *Bottom*. Bold activity in an overlapping region of vmPFC was modulated by trial-by-trial performance (Peak: -3,50,-1). **B**, An independent vmPFC region-of-interest (ROI) shows the expected significant modulation by both reward outcome and performance (both $P < 0.01$). **C**, In the independent vmPFC ROI, subjects with higher performance than reward weights in the computational analysis of mood displayed stronger neural responses in the vmPFC for performance than subjects with higher reward than performance weights ($P = 0.01$). **D**, The difference between performance and reward weights in the happiness model, a measure potentially reflecting the relative subjective values of intrinsic and extrinsic rewards, predicted the difference in neural responses for successful performance relative to reward in the independent vmPFC ROI ($\rho = 0.52$, $P = 0.002$). * $P < 0.05$, ** $P < 0.01$.

Across participants, we found a significant positive relationship (Spearman's $\rho = 0.52$, $P = 0.002$, Fig. 6.4D) between the relative weights for extrinsic rewards and performance in our mood model, a measure potentially reflecting the subjective values for extrinsic and intrinsic rewards respectively, and the pattern of neural responses to both quantities within the independent vmPFC mask. This finding suggests that the extent to which an individual's mood is influenced by extrinsic or intrinsic rewards is reflected in the brain. As previous studies (Rutledge et al., 2014) have found a positive association between BOLD activity in the right anterior insula and z-scored happiness ratings, we conducted a region-of-interest (ROI) analysis using an 8mm sphere (coordinates 42, 5, -14) and similarly found increased BOLD activity in the right anterior insula for positive versus negative z-scored happiness ratings after regressing out outcomes for explicit rewards and skilled performance (Fig. 6.5).

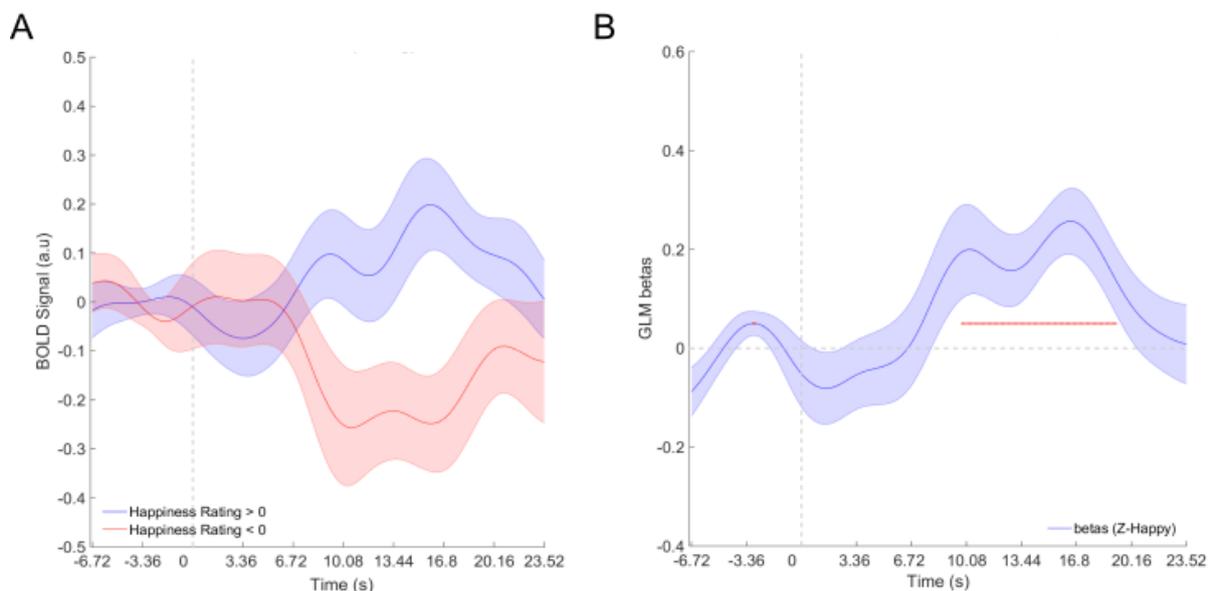


Fig. 6.5. Effect of happiness question on right anterior insula. (A) After controlling for extrinsic outcomes and skilled performance, BOLD activity in the right anterior insula was higher when participants made positive ratings of happiness than negative ratings. Dashed vertical grey line indicates onset of happiness question. (B) A generalized linear model was constructed by regressing BOLD activity in right anterior insula against z-scored happiness ratings. Timepoints where happiness ratings were significantly represented in the BOLD activity are indicated in red.

6.3 Discussion

Using experience sampling (Kahneman, Krueger, Schkade, Schwarz, & Stone, 2004; Reis & Gable, 2000) combined with functional neuroimaging, we investigated the contributions of intrinsic and extrinsic rewards to subjective well-being or happiness using a task designed to disentangle affective responses to both reward types. Recent studies have used a similar approach to examine subjective feelings within the context of value-based decision-making (Eldar, Rutledge, Dolan, & Niv, 2016; Rutledge et al., 2014; Vinckier et al., 2018) and social interactions (Rutledge, de Berker, Espenhahn, Dayan, & Dolan, 2016). On top of extending those findings to the reinforcement learning domain, our findings here highlight that fluctuations in momentary happiness is not only influenced by rewards that are externally quantifiable but also those that are intrinsically experienced. Furthermore, vmPFC BOLD activity reflects the extent to which an individual's happiness is driven by intrinsic over extrinsic rewards.

While improvements in skilled performance can be enhanced by rewarding individuals for performance (Sugawara, Tanaka, Okazaki, Watanabe, & Sadato, 2012), holding performance constant across participants here allowed us to investigate how happiness varied independently of how skilled an individual was in the task. We show that individuals whose happiness were largely influenced by intrinsic rewards exhibited increased vmPFC BOLD responses for successful versus unsuccessful skilled performance relative to individuals whose happiness were largely influenced by extrinsic rewards. The positive correlation between the performance parameter in our mood model and the median speed of the cursor suggests that individuals who may have been more proficient at the task were also those whose momentary happiness were greater impacted by their performance.

The vmPFC, together with orbitofrontal cortex, codes for the value of different types of goods and anticipatory outcomes like food or juice (Hare, Malmaud, & Rangel, 2011; Padoa-Schioppa, 2007), money (Martino, Kumaran, Seymour, & Dolan, 2006), and even aesthetic judgments (Jacobsen, Schubotz, Höfel, & Cramon, 2006; Kawabata & Zeki, 2004). An interpretation of this is that the vmPFC plays a principle role in representing qualitatively different types of goods on a common scale which would facilitate decisions to be made between otherwise incommensurable goods

(Levy & Glimcher, 2011, 2012). Our study builds on these results by identifying an association between vmPFC BOLD activity and rewards that are intrinsically valued, such as the experience of performing a skilled task flawlessly.

The vmPFC has further been demonstrated to play a role in affect with subjective emotional experiences elicited by images and pleasurable music leading to changes in vmPFC BOLD activity and regional cerebral blood flow (Blood & Zatorre, 2001; Winecoff et al., 2013; Zald, Mattson, & Pardo, 2002). Damage to the vmPFC has resulted in abnormal emotional responses (Hiser & Koenigs, 2018; Koenigs et al., 2007; Zald & Andreotti, 2010) and maladaptive decision-making in environments where emotional regulation may be useful (Grossman et al., 2010; Spaniol, Di Muro, & Ciaramelli, 2019). Our finding that the extent to which individuals care more about intrinsic or extrinsic rewards insofar as their momentary happiness is influenced by them is reflected in a similar manner in the vmPFC suggests that the strong association observed between measures of happiness and vmPFC BOLD activity could allow subjective values of intrinsic rewards to be estimated from an individual's mood dynamics.

People exhibit biases when it comes to predicting how future events would impact on their emotional states and are also prone to making sub-optimal decisions by misjudging the hedonic consequences of pursuing an option (Meyvis et al., 2010; Nisbet & Zelenski, 2011; Wilson & Gilbert, 2005), posing a difficulty for enacting policies that objectively increased subjective well-being. Additional factors such as the social desirability bias (Van de Mortel, 2008) further decrease the reliability of self-reported value when an individual's assessment of a hypothetical experience or good, such as the availability of public parks, differs from what society perceives it ought to be. Our results suggest that the combination of computational modelling and measures of momentary happiness can be used as a tool to obtain accurate value estimates of intangible goods like performance or experiences. An advantage of our method is that it can be applied to any experience without a need to probe people explicitly about the content of those experiences, reducing biases associated with social desirability. Furthermore, the use of computational modelling and experience sampling in a previous study have found that parameter estimates obtained from a similar approach were not confined to a controlled setting but also replicated outside of the laboratory (Rutledge et al., 2014). Future work could expand on our findings in

more naturalistic settings such as a corporate workplace, and build on them using experiences that could be important for well-being.

Consistent with previous results, we also observed differential BOLD activity in right anterior insula during onset of the rating question for positive and negative happiness ratings, providing evidence that it plays a role in the interoceptive awareness of affective states (Craig, 2009; Damasio & Carvalho, 2013; Rutledge et al., 2014). Our main finding is that vmPFC BOLD activity in response to intrinsic and extrinsic rewards can be predicted by measures of momentary happiness, suggesting that such measures can be used as a convenient tool to approximate implicit values of abstract goods and experiences that may be otherwise challenging to quantify. Such a tool would prove useful for policies targeting the maximisation of subjective well-being by providing a computational and neuroscientific framework for measuring the implicit values of abstract goods and experiences.

References

- Bartra, O., McGuire, J. T., & Kable, J. W. (2013). The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *NeuroImage*, *76*, 412–427. <https://doi.org/10.1016/j.neuroimage.2013.02.063>
- Blood, A. J., & Zatorre, R. J. (2001). Intensely pleasurable responses to music correlate with activity in brain regions implicated in reward and emotion. *PNAS*, *98*(20), 11818–11823. <https://doi.org/10.1073/pnas.191355898>
- Chida, Y., & Steptoe, A. (2008). Positive psychological well-being and mortality: A quantitative review of prospective observational studies. *Psychosomatic Medicine*, *70*(7), 741–756. <https://doi.org/10.1097/PSY.0b013e31818105ba>
- Craig, A. D. (2009). How do you feel — now? The anterior insula and human awareness. *Nature Reviews Neuroscience*, *10*(1), 59–70. <https://doi.org/10.1038/nrn2555>
- Damasio, A., & Carvalho, G. B. (2013). The nature of feelings: Evolutionary and neurobiological origins. *Nature Reviews Neuroscience*, *14*(2), 143–152. <https://doi.org/10.1038/nrn3403>
- Davidson, K. W., Mostofsky, E., & Whang, W. (2010). Don't worry, be happy: Positive affect and reduced 10-year incident coronary heart disease: the Canadian Nova Scotia Health Survey. *European Heart Journal*, *31*(9), 1065–1070. <https://doi.org/10.1093/eurheartj/ehp603>
- Dolan, P., & White, M. P. (2007). How can measures of subjective well-being be used to inform public policy? *Perspectives on Psychological Science*, *2*(1), 71–85. <https://doi.org/10.1111/j.1745-6916.2007.00030.x>
- Eldar, E., Rutledge, R. B., Dolan, R. J., & Niv, Y. (2016). Mood as Representation of Momentum. *Trends in Cognitive Sciences*, *20*(1), 15–24. <https://doi.org/10.1016/j.tics.2015.07.010>
- Gilbert, D. T., & Wilson, T. D. (2007). Propection: Experiencing the future. *Science*, *317*(5843), 1351–1354. <https://doi.org/10.1126/science.1144161>
- Grossman, M., Eslinger, P. J., Troiani, V., Anderson, C., Avants, B., Gee, J. C., Antani, S. (2010). The role of ventral medial prefrontal cortex in social decisions: Converging evidence from fMRI and frontotemporal lobar degeneration. *Neuropsychologia*, *48*(12), 3505–3512. <https://doi.org/10.1016/j.neuropsychologia.2010.07.036>
- Hare, T. A., Malmaud, J., & Rangel, A. (2011). Focusing attention on the health aspects of foods changes value signals in vmPFC and improves dietary choice. *Journal of Neuroscience*, *31*(30), 11077–11087. <https://doi.org/10.1523/JNEUROSCI.6383-10.2011>

- Hiser, J., & Koenigs, M. (2018). The multifaceted role of ventromedial prefrontal cortex in emotion, decision-making, social cognition, and psychopathology. *Biological Psychiatry*, *83*(8), 638–647. <https://doi.org/10.1016/j.biopsych.2017.10.030>
- Howe, M. W., Tierney, P. L., Sandberg, S. G., Phillips, P. E. M., & Graybiel, A. M. (2013). Prolonged dopamine signalling in striatum signals proximity and value of distant rewards. *Nature*, *500*(7464), 575–579. <https://doi.org/10.1038/nature12475>
- Jacobsen, T., Schubotz, R. I., Höfel, L., & Cramon, D. Y. v. (2006). Brain correlates of aesthetic judgment of beauty. *NeuroImage*, *29*(1), 276–285. <https://doi.org/10.1016/j.neuroimage.2005.07.010>
- Kahneman, D., Krueger, A. B., Schkade, D. A., Schwarz, N., & Stone, A. A. (2004). A survey method for characterizing daily life experience: The day reconstruction method. *Science*, *306*(5702), 1776–1780. <https://doi.org/10.1126/science.1103572>
- Kawabata, H., & Zeki, S. (2004). Neural correlates of beauty. *Journal of Neurophysiology*, *91*(4), 1699–1705. <https://doi.org/10.1152/jn.00696.2003>
- Keatley, D., Clarke, D. D., & Hagger, M. S. (2013). The predictive validity of implicit measures of self-determined motivation across health-related behaviours. *British Journal of Health Psychology*, *18*(1), 2–17. <https://doi.org/10.1111/j.2044-8287.2011.02063.x>
- Koenigs, M., Young, L., Adolphs, R., Tranel, D., Cushman, F., Hauser, M., & Damasio, A. (2007). Damage to the prefrontal cortex increases utilitarian moral judgements. *Nature*, *446*(7138), 908–911. <https://doi.org/10.1038/nature05631>
- Krajchich, I., Camerer, C., Ledyard, J., & Rangel, A. (2009). Using neural measures of economic value to solve the public goods free-rider problem. *Science*, *326*(5952), 596–599. <https://doi.org/10.1126/science.1177302>
- Levesque, C., Copeland, K. J., & Sutcliffe, R. A. (2008). Conscious and nonconscious processes: Implications for self-determination theory. *Canadian Psychology*, *49*(3), 218–224. <https://doi.org/10.1037/a0012756>
- Levy, D. J., & Glimcher, P. W. (2011). Comparing apples and oranges: Using reward-specific and reward-general subjective value representation in the brain. *Journal of Neuroscience*, *31*(41), 14693–14707. <https://doi.org/10.1523/JNEUROSCI.2218-11.2011>
- Levy, D. J., & Glimcher, P. W. (2012). The root of all value: A neural common currency for choice. *Current Opinion in Neurobiology*, *22*(6), 1027–1038. <https://doi.org/10.1016/j.conb.2012.06.001>

- Martino, B. D., Kumaran, D., Seymour, B., & Dolan, R. J. (2006). Frames, biases, and rational decision-making in the human brain. *Science*, *313*(5787), 684–687. <https://doi.org/10.1126/science.1128356>
- Meyvis, T., Ratner, R. K., & Levav, J. (2010). Why don't we learn to accurately forecast feelings? How misremembering our predictions blinds us to past forecasting errors. *Journal of Experimental Psychology*, *139*(4), 579–589. <https://doi.org/10.1037/a0020285>
- Morewedge, C. K., & Buechel, E. C. (2013). Motivated underpinnings of the impact bias in affective forecasts. *Emotion*, *13*(6), 1023–1029. <https://doi.org/10.1037/a0033797>
- Nisbet, E. K., & Zelenski, J. M. (2011). Underestimating nearby nature: Affective forecasting errors obscure the happy path to sustainability. *Psychological Science*, *22*(9), 1101–1106. <https://doi.org/10.1177/0956797611418527>
- Oswald, A. J., Proto, E., & Sgroi, D. (2015). Happiness and productivity. *Journal of Labor Economics*, *33*(4), 789–822. <https://doi.org/10.1086/681096>
- Oswald, A. J., & Wu, S. (2010). Objective confirmation of subjective measures of human well-being: Evidence from the U.S.A. *Science*, *327*(5965), 576–579. <https://doi.org/10.1126/science.1180606>
- Padoa-Schioppa, C. (2007). Orbitofrontal cortex and the computation of economic value. *Annals of the New York Academy of Sciences*, *1121*, 232–253. <https://doi.org/10.1196/annals.1401.011>
- Pelled, L. H., & Xin, K. R. (1999). Down and out: An investigation of the relationship between mood and employee withdrawal behavior. *Journal of Management*, *25*(6), 875–895. [https://doi.org/10.1016/S0149-2063\(99\)00027-6](https://doi.org/10.1016/S0149-2063(99)00027-6)
- Peterson, S. J., Luthans, F., Avolio, B. J., Walumbwa, F. O., & Zhang, Z. (2011). Psychological capital and employee performance: A latent growth modeling approach. *Personnel Psychology*, *64*(2), 427–450. <https://doi.org/10.1111/j.1744-6570.2011.01215.x>
- Plassmann, H., O'Doherty, J., & Rangel, A. (2007). Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *Journal of Neuroscience*, *27*(37), 9984–9988. <https://doi.org/10.1523/JNEUROSCI.2131-07.2007>
- Reis, H. T., & Gable, S. L. (2000). Event-sampling and other methods for studying everyday experience. In *Handbook of research methods in social and personality psychology* (190–222). New York, NY, US: Cambridge University Press.
- Rutledge, R. B., de Berker, A. O., Espenhahn, S., Dayan, P., & Dolan, R. J. (2016). The social contingency of momentary subjective well-being. *Nature Communications*, *7*, 11825. <https://doi.org/10.1038/ncomms11825>

- Rutledge, R. B., Skandali, N., Dayan, P., & Dolan, R. J. (2014). A computational and neural model of momentary subjective well-being. *PNAS*, *111*(33), 12252–12257. <https://doi.org/10.1073/pnas.1407535111>
- Rutledge, R. B., Skandali, N., Dayan, P., & Dolan, R. J. (2015). Dopaminergic modulation of decision making and subjective well-being. *Journal of Neuroscience*, *35*(27), 9811–9822. <https://doi.org/10.1523/JNEUROSCI.0702-15.2015>
- Spaniol, J., Di Muro, F., & Ciaramelli, E. (2019). Differential impact of ventromedial prefrontal cortex damage on “hot” and “cold” decisions under risk. *Cognitive, Affective, & Behavioral Neuroscience*, *19*(3), 477–489. <https://doi.org/10.3758/s13415-018-00680-1>
- Stephoe, A., Deaton, A., & Stone, A. A. (2015). Subjective wellbeing, health, and ageing. *The Lancet*, *385*(9968), 640–648. [https://doi.org/10.1016/S0140-6736\(13\)61489-0](https://doi.org/10.1016/S0140-6736(13)61489-0)
- Sugawara, S. K., Tanaka, S., Okazaki, S., Watanabe, K., & Sadato, N. (2012). Social rewards enhance offline improvements in motor skill. *PLOS ONE*, *7*(11), e48174. <https://doi.org/10.1371/journal.pone.0048174>
- Taylor, M. M., & Creelman, C. D. (1967). PEST: Efficient estimates on probability functions. *The Journal of the Acoustical Society of America*, *41*(4A), 782–787. <https://doi.org/10.1121/1.1910407>
- Van de Mortel, T. F. (2008). Faking it: Social desirability response bias in self-report research. *Australian Journal of Advanced Nursing*, *The*, *25*(4), 40.
- Vinckier, F., Rigoux, L., Oudiette, D., & Pessiglione, M. (2018). Neuro-computational account of how mood fluctuations arise and affect decision making. *Nature Communications*, *9*(1). <https://doi.org/10.1038/s41467-018-03774-z>
- Wilson, T. D., & Gilbert, D. T. (2005). Affective forecasting: Knowing what to want. *Current Directions in Psychological Science*, *14*(3), 131–134. <https://doi.org/10.1111/j.0963-7214.2005.00355.x>
- Winecoff, A., Clithero, J. A., Carter, R. M., Bergman, S. R., Wang, L., & Huettel, S. A. (2013). Ventromedial prefrontal cortex encodes emotional value. *Journal of Neuroscience*, *33*(27), 11032–11039. <https://doi.org/10.1523/JNEUROSCI.4317-12.2013>
- Zald, D. H., & Andreotti, C. (2010). Neuropsychological assessment of the orbital and ventromedial prefrontal cortex. *Neuropsychologia*, *48*(12), 3377–3391. <https://doi.org/10.1016/j.neuropsychologia.2010.08.012>
- Zald, D. H., Mattson, D. L., & Pardo, J. V. (2002). Brain activity in ventromedial prefrontal cortex correlates with individual differences in negative affect. *PNAS*, *99*(4), 2450–2454. <https://doi.org/10.1073/pnas.042457199>

Chapter 7: Conclusions and general discussion

It is good to have an end to journey toward; but it is the journey that matters, in the end.

Ursula K Le Guin, *The Left Hand of Darkness*

This thesis has used a mixture of neuroimaging and computational modelling approaches to investigate how neurophysiological processes and neuroanatomical features influence risky choice, and the relationship between the implicit value of an experience and subjective well-being. We have drawn upon ideas discussed in **chapter 2** and methods summarised in **chapter 3** to form the basis of predictions made in later empirical chapters. In **Chapter 4**, we used a novel application of real-time functional magnetic resonance imaging (rt-fMRI) to test the prediction that endogenous fluctuations in the dopaminergic midbrain influence behavioural choice variability. **Chapter 5** details how risk preferences were related to the content of brain tissue microstructure by leveraging on the development of *in vivo* histology using MRI (hMRI). In **Chapter 6**, we used fMRI to investigate the neural correlates of intrinsic and extrinsic rewards and their relation to subjective well-being.

In this final chapter, I discuss some general principles that can be drawn from these studies.

7.1 Variability in the brain

In **chapter 2**, we discussed how brain imaging signals that co-varied with experimental variables of interest were thought to be functionally relevant to decision processes through representations of choice features (e.g. subjective value) or computations. Indeed, studies have shown that people with lesions to areas implicated in the decision network like vmPFC and OFC demonstrated impairments in preference judgements and value maximization when afforded choices between multiple options (Camille, Griffiths, Vo, Fellows, & Kable, 2011; Fellows, 2011; Gläscher et al., 2012; Rushworth, Noonan, Boorman, Walton, & Behrens, 2011). However, such signals may not always be explicitly used to guide the behaviour of an organism. In a group of depressed participants who were administered amisulpride (a D₂/D₃ receptor antagonist), researchers found acute enhancement of reward-related striatal

activation and corticostriatal functional connectivity relative to depressed participants who received placebos as participants played a monetary incentive delay task in the MRI scanner (Admon et al., 2017). Despite this, reward-motivated behaviour of depressed participants in the experimental group remained significantly different from that of the healthy controls, suggesting that normalised neural signals need not necessarily translate to overt behavioural differences and that there could be other factors at play when representations are transformed to behaviour.

One such factor could be neural variability. A source of variability is noise which permeates the nervous system at sensory, cellular, and motor levels. Noise is an inevitable consequence of brains operating across multiple scales from molecular components to complex systems of neurons that produce behaviour. Another source of variability is the initial state of a neural circuit that could lead to divergent neuronal and behavioural responses (Faisal, Selen, & Wolpert, 2008). A system with dynamics that are highly sensitive to the initial conditions when an event commences will demonstrate greater response variability. While much work in the field of decision-making have focused on the representation and computation of value in response to options and outcomes presented to participants, often overlooked are the contributions of endogenous fluctuations in brain imaging signals to choice even in the absence external stimulation. In **chapter 4**, we provided evidence that endogenous fluctuations in the dopaminergic midbrain is able to drive variability in risky choice behaviour, suggesting a systematic relationship between variability at the neural level and complex behaviours.

Why might variability be observed across a broad spectrum of processes like perception (Boly et al., 2007; de Gee et al., 2017; Sadaghiani, Poline, Kleinschmidt, & D'Esposito, 2015), motor action (Fox, Snyder, Vincent, & Raichle, 2007), and higher cognition (Chew et al., 2019)? One compelling hypothesis is that neural variability reflects the dynamic range of potential responses to environmental stimuli, allowing the brain to flexibly transition between states in response to changing task demands (Garrett, Kovacevic, McIntosh, & Grady, 2013). This could help reduce susceptibility to becoming entrenched in specific behavioural repertoires (van Leeuwen, 2008) and promote exploration in dynamic environments that are often a feature of the natural world (Wilson, Geana, White, Ludvig, & Cohen, 2014). This does not necessarily mean that behaviour is completely variable from moment to moment – recall that other

factors such as the value of options still dominated the decision to take a risk for the experiment described in **chapter 4**. Rather, the initial states of a circuit could provide a frame of reference used in computations performed by downstream regions. Within the context of noise, systems that use thresholds for signal detection demonstrate strategies like stochastic resonance that rely on changes in noise levels for the periodic detection sub-threshold signals (Russell, Wilkens, & Moss, 1999). Further, neuronal networks formed in the presence of noise display reduced generalization errors and are less prone to overfitting (Krogh & Hertz, 1992).

The results described in **chapter 5** suggest that risk preferences can also be predicted to some extent by individual differences in the makeup of tissue microstructure. Together with the finding that neural variability is translated to behavioural variability in some people more than others, future work might expand on these ideas using a multi-modal approach to paint a unifying picture of the relationship between neuroanatomy and variability within the context of risk. Such an avenue of research could be potentially useful within the field of computational psychiatry where illnesses can manifest as extremes of behaviours ranging from repetitive actions to highly unpredictable ones.

7.2 Subjective well-being as a computational hedonometer

In **chapter 6**, we touched on the idea of subjective well-being and how it should be a target for policies to maximise due to the benefits it conveys to individuals and societies. A difficulty with designing potential measures to improve well-being presents itself in the form of biases when it comes to *affective forecasting* and self-reports that are sometimes used to inform change. For example, despite the future boost in mood associated with outdoor walks in nearby nature, people failed to anticipate the hedonic benefits of choosing that option (Nisbet & Zelenski, 2011).

When it comes to reporting the value of a good or experience, a common economic assumption is that the smallest amount an individual would accept to give up the good or largest amount an individual would pay to obtain the good reflects its economic value for the individual. However, in the absence of markets for hypothetical experiences or public goods, willingness to pay may not necessarily represent economic value but could instead be an expression of an individual's *attitudes* towards

the good (Kahneman, Ritov, Jacowitz, & Grant, 1993). An example of this was demonstrated in a recent study where the willingness to pay was used to investigate the value of receiving thoughts and prayers following hardship (Thunström & Noy, 2019). The authors found that atheistic and religiously agnostic participants were not only willing to pay to avoid receiving prayers, they also negatively valued thoughts from Christians compared to other secular groups. Willingness to pay may thus reflect a measure of attitude on an arbitrary scale for goods that lack economic markets.

The experimental results presented in **chapter 6** provide evidence that the values of both intrinsic and extrinsic rewards are encoded in the ventromedial prefrontal cortex, supporting the idea of a common currency in which subjective values of different goods are represented (Levy & Glimcher, 2012). That the extent to which an individual's happiness is influenced by intrinsic over extrinsic reward - as revealed by computational modelling - is reflected in this region suggests that our approach could be useful in estimating the subjective value of intrinsic rewards. Rather than probing participants explicitly about the contents of their experiences which may be susceptible to biases, our computational hedonometer is able to disentangle the contributions of distinct goods or events to measures of momentary happiness. This allows us to derive value estimates for abstract goods or experiences relative to money and could prove useful for informing policies that aim to improve subjective well-being. Future studies could expand on these results by applying our approach to other naturalistic settings and goods (e.g. public goods like parks) to examine whether changes inspired by such measures of intrinsic rewards generate larger hedonic benefits and greater improvement in well-being compared to traditional methods.

References

- Admon, R., Kaiser, R. H., Dillon, D. G., Beltzer, M., Goer, F., Olson, D. P., Pizzagalli, D. A. (2017). Dopaminergic enhancement of striatal response to reward in major depression. *The American Journal of Psychiatry*, *174*(4), 378–386. <https://doi.org/10.1176/appi.ajp.2016.16010111>
- Boly, M., Balteau, E., Schnakers, C., Degueldre, C., Moonen, G., Luxen, A., Laureys, S. (2007). Baseline brain activity fluctuations predict somatosensory perception in humans. *PNAS*, *104*(29), 12187–12192. <https://doi.org/10.1073/pnas.0611404104>
- Camille, N., Griffiths, C. A., Vo, K., Fellows, L. K., & Kable, J. W. (2011). Ventromedial frontal lobe damage disrupts value maximization in humans. *Journal of Neuroscience*, *31*(20), 7527–7532. <https://doi.org/10.1523/JNEUROSCI.6527-10.2011>
- Chew, B., Hauser, T. U., Papoutsi, M., Magerkurth, J., Dolan, R. J., & Rutledge, R. B. (2019). Endogenous fluctuations in the dopaminergic midbrain drive behavioral choice variability. *PNAS*, *116*(37), 18732–18737. <https://doi.org/10.1073/pnas.1900872116>
- de Gee, J. W., Colizoli, O., Kloosterman, N. A., Knapen, T., Nieuwenhuis, S., & Donner, T. H. (2017). Dynamic modulation of decision biases by brainstem arousal systems. *ELife*, *6*, e23232. <https://doi.org/10.7554/eLife.23232>
- Faisal, A. A., Selen, L. P. J., & Wolpert, D. M. (2008). Noise in the nervous system. *Nature Reviews Neuroscience*, *9*(4), 292–303. <https://doi.org/10.1038/nrn2258>
- Fellows, L. K. (2011). Orbitofrontal contributions to value-based decision making: Evidence from humans with frontal lobe damage. *Annals of the New York Academy of Sciences*, *1239*(1), 51–58. <https://doi.org/10.1111/j.1749-6632.2011.06229.x>
- Fox, M. D., Snyder, A. Z., Vincent, J. L., & Raichle, M. E. (2007). Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior. *Neuron*, *56*(1), 171–184. <https://doi.org/10.1016/j.neuron.2007.08.023>
- Garrett, D. D., Kovacevic, N., McIntosh, A. R., & Grady, C. L. (2013). The modulation of BOLD variability between cognitive states varies by age and processing speed. *Cerebral Cortex*, *23*(3), 684–693. <https://doi.org/10.1093/cercor/bhs055>
- Gläscher, J., Adolphs, R., Damasio, H., Bechara, A., Rudrauf, D., Calamia, M., Tranel, D. (2012). Lesion mapping of cognitive control and value-based decision making in the prefrontal cortex. *PNAS*, *109*(36), 14681–14686. <https://doi.org/10.1073/pnas.1206608109>

- Kahneman, D., Ritov, I., Jacowitz, K. E., & Grant, P. (1993). Stated willingness to pay for public goods: A psychological perspective. *Psychological Science, 4*(5), 310–315. <https://doi.org/10.1111/j.1467-9280.1993.tb00570.x>
- Krogh, A., & Hertz, J. A. (1992). Generalization in a linear perceptron in the presence of noise. *Journal of Physics A: Mathematical and General, 25*(5), 1135–1147. <https://doi.org/10.1088/0305-4470/25/5/020>
- Levy, D. J., & Glimcher, P. W. (2012). The root of all value: A neural common currency for choice. *Current Opinion in Neurobiology, 22*(6), 1027–1038. <https://doi.org/10.1016/j.conb.2012.06.001>
- Nisbet, E. K., & Zelenski, J. M. (2011). Underestimating nearby nature: Affective forecasting errors obscure the happy path to sustainability. *Psychological Science, 22*(9), 1101–1106. <https://doi.org/10.1177/0956797611418527>
- Rushworth, M. F. S., Noonan, M. P., Boorman, E. D., Walton, M. E., & Behrens, T. E. (2011). Frontal cortex and reward-guided learning and decision-making. *Neuron, 70*(6), 1054–1069. <https://doi.org/10.1016/j.neuron.2011.05.014>
- Russell, D. F., Wilkens, L. A., & Moss, F. (1999). Use of behavioural stochastic resonance by paddle fish for feeding. *Nature, 402*(6759), 291–294. <https://doi.org/10.1038/46279>
- Sadaghiani, S., Poline, J.-B., Kleinschmidt, A., & D'Esposito, M. (2015). Ongoing dynamics in large-scale functional connectivity predict perception. *PNAS, 112*(27), 8463–8468. <https://doi.org/10.1073/pnas.1420687112>
- Thunström, L., & Noy, S. (2019). The value of thoughts and prayers. *PNAS, 201908268*. <https://doi.org/10.1073/pnas.1908268116>
- van Leeuwen, C. (2008). Chaos breeds autonomy: Connectionist design between bias and baby-sitting. *Cognitive Processing, 9*(2), 83–92. <https://doi.org/10.1007/s10339-007-0193-8>
- Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology, 143*(6), 2074–2081. <https://doi.org/10.1037/a0038199>