

THE ROLE OF THE VENTROMEDIAL PREFRONTAL CORTEX IN VALUE-
BASED DECISION-MAKING

Linda Q. Yu

A DISSERTATION

in

Psychology

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2019

Supervisor of Dissertation

Joseph W. Kable, Baird Term Professor, Psychology

Graduate Group Chairperson

Sara Jaffee, Professor, Psychology

Dissertation Committee

Emily Falk, Associate Professor, Communication, Psychology and Marketing

Anjan Chatterjee, Elliott Professor, Neurology

ACKNOWLEDGMENTS

A sign propped up in our lab space reads, “If we knew what we were doing, it wouldn’t be called research. – A. Einstein.” If this quote is true, then I sure did a ton of research. I would like to thank the many people who helped me know slightly more what I am doing. First among them is my advisor in all senses of that word, Joe Kable. You admirably performed the duties of both supervisor and therapist, radiating calm and reassurance through my many neuroses. Thank you for your knowledge, support, guidance, and above all, maintaining good humour through the many pranks I have pulled.

I would also like to thank my committee, Anjan Chatterjee and Emily Falk, for being kind and supportive, and for providing valuable feedback on my research and helping my professional development over the years.

Thank you to Lesley Fellows and Christine Déry at McGill and Eileen Cardillo at Penn for making the patient part of my research possible, and *especially* to the patients themselves for donating their time to research.

Thanks to all Kable lab members, past and present, who made Penn and Philadelphia feel like home. I recognize how lucky I am to have hit the jackpot in finding so many good friends in the place I work. Thanks for the long talks, help with research, driving me to patient testing appointments, and being so down for happy hours, Bucket List adventures and Harry Potter themed events. A special thanks must go to Sangil “Arthur” Lee, who arrived at Penn the same day as I did, for being my co-author, lunch-eating partner, and trusty statistical guru for the last five-plus years.

Thank you from the bottom of my heart to my parents, Liyao Yu and Shiling Li, whose decision to bring me as a child from China to a wonderful land called Canada, where homework did not seem to exist, changed the course of my life. I owe everything I have to you, and I am so grateful every single day for your sacrifice and wisdom in giving me a life better than I could ever imagine. And thanks to Larry for making me the happiest, proudest big sister in the world.

Thank you to the Vaidya family for your kindness, innumerable hot meals, advice and laughter.

Finally, thank you to Avi, whom I am astonished at my good fortune to call my husband and best friend. Thank you for supporting my decision to go to grad school far away, lasting five years of long distance, editing papers, talking through research, making me laugh, cooking me dinner, and always encouraging me to be the best version of myself.

ABSTRACT

THE ROLE OF THE VENTROMEDIAL PREFRONTAL CORTEX IN VALUE-BASED DECISION-MAKING

Linda Q. Yu

Joseph W. Kable

The ventromedial prefrontal cortex (vmPFC) has been shown to correlate with the subjective value for options, across reward type and across hundreds of functional neuroimaging studies. Despite the prominence of its role in preference-based decision-making, its specific contributions to how decisions are made have not yet been well-characterised. Study 1 addresses *what* the vmPFC signal represents during decision-making. While the vmPFC signal has been shown to correlate highly with subjective value in past studies, this signal is also consistent with mental navigation through a conceptual attribute space using a grid-like code. We found that the mental navigation model lacked support in the evidence, and the subjective value model remains the best explanation for vmPFC signal during decision-making. After having established that the signal in vmPFC reflects subjective value, Study 2 addresses whether subjective value representations remain consistent for non-choice preference tasks, and *when* this representation comes online during the decision process. This study shows that the value network seen previously for choice tasks also is active during a matching bidding task, and that the vmPFC, interestingly, represents value only at the time of the final choice. Finally, in Study 3, I address the question of *how* the vmPFC is necessary for subjective value in my third chapter. Transitivity (the idea that if $A > B$, and $B > C$, then $A > C$) is a

key property of a value-based system. Individuals with ventromedial frontal lobe damage have been found to make more transitivity errors in the past, but it is not known whether vmPFC damage causes fundamentally intransitive choices (implying abolishment of value), or transitive but noisier choices (implying preservation of value but increased instability). We found strong evidence for the second case, demonstrating that vmPFC damage adds instability to valuation but does not abolish it. The evidence I present here is consistent with the theory that vmPFC is involved in a subjective value-based process during decision-making, yet that value is a distributed process over many brain regions where other regions may compensate for the loss of the vmPFC in calculating value.

TABLE OF CONTENTS

ACKNOWLEDGMENTS	II
ABSTRACT.....	IV
LIST OF TABLES	VIII
LIST OF ILLUSTRATIONS.....	IX
CHAPTER 1 -- INTRODUCTION.....	1
CHAPTER 2 - SUBJECTIVE VALUE, NOT A GRID-LIKE CODE, DESCRIBES NEURAL ACTIVITY IN VENTROMEDIAL PREFRONTAL CORTEX DURING DECISION- MAKING	12
Abstract	12
Introduction	13
Methods	15
Results.....	20
Discussion	26
Supplemental Materials	30
CHAPTER 3 -- TIMING OF VALUE REPRESENTATION IN VENTROMEDIAL PREFRONTAL CORTEX IN A COMPLEX AUCTION TASK	39
Abstract	39
Introduction	40
Materials and Methods	41
Results.....	46
Discussion	51
CHAPTER 4 -- INDIVIDUALS WITH VENTROMEDIAL FRONTAL DAMAGE HAVE MORE UNSTABLE BUT STILL FUNDAMENTALLY TRANSITIVE PREFERENCES .	55
Abstract	55

Introduction	56
Materials and Methods	58
Results.....	70
Discussion	84
CHAPTER 5 – GENERAL DISCUSSION.....	91
BIBLIOGRAPHY	98

LIST OF TABLES

Table 3-1: Peak foci of BOLD effects	47
Table 4-1. Demographics of participants.....	61
Table 4-2. Results of LOP analysis, by category.....	77

LIST OF ILLUSTRATIONS

Figure 2-1. Correlation between subjective value signal and its most similar hexagonal modulation signal.....	21
Figure 2-2. Regions with significant SV correlation in session 1 and session 2..	22
Figure 2-3. Z-transformed F-statistics of brain regions that are significantly explained by two hexagonal grid modulation regressors.	23
Figure 2-4. Effect of grid-angle consistency and SV in grid cell ROIs.....	25
Figure 2-5. Effect of grid-angle consistency and SV in SV-defined ROIs.....	26
Figure 3-1 .Whole brain effects correlated with participants' bid response on each trial at the four time points in the trial.....	50
Figure 3-2. Activity in the vmPFC ROI related to participant's bid is only significant at the final time point, that of bid response submission.....	51
Figure 4-1. Overlap images of the VMF and frontal control lesion groups..	60
Figure 4-2. Group average and individually plotted intransitive choices.....	71
Figure 4-3. Lesion tracings of the three individuals with VMF lesions who had significantly more intransitive choices compared to healthy control subjects.	75
Figure 4-4. DDM parameter fits: noise, drift rate, initial starting point, non-decision time, and alpha (exponent on rank distance).....	81
Figure 4-5. Value distance effect on RT, by group.....	83

Chapter 1 – Introduction

Our primate ancestors must be very jealous of us. We have a buffet of choices before us in the supermarket, endless aisles of protein, fat and sugar available with the single reach of the hand. The biggest problems we face (in rich, industrialized countries, anyway) these days is which restaurant we want to tap on our phone so they can deliver us dinner, and then what multi-million-dollar entertainment franchise to tap on next that we can watch while we eat. We have so many choices that, in fact, we are often paralyzed by them – a cacophony of multi-sensory influences competing for our hand. When you can have anything you desire, what do you desire?

The central question of this dissertation is to figure out how our brain deals with all this noise, and how we extract that precious thing called desire from them. This brain of ours originates not just from those unlucky primate ancestors; from the earliest days as a mere few cells strung together, it has had the task of keeping itself alive from when life began. To do so in the eons before smartphones, it had to figure out what actions to take to feed itself and procreate. How does it do so in a way that would maximize its chances of perpetuating itself? One way, economists thought, was that it should maximize *utility*, which is essentially the usefulness of something in advancing one's goals (Baron, 2008). Utility is like a summary of the attributes (either benefits or drawbacks) of an item or action, which can then be compared to those of another. Because utility is something that is abstracted from the attributes of the individual options, dissimilar items like apples and orange play-dough can be compared. Another important property of utility is that

preferences should be transitive (Samuelson, 1937). If a slime mold prefers agar patch A to agar patch B, and a.p. B to a.p. C, then it ought to prefer A to C. The reason to do so is to maximize utility towards its goal, which, for a slime mold, is to maximize metabolic energy and avoid light (Devi, Guttus, & Guttus, 1968).

An alternative to using an integrated value representation to adjudicate among options is using heuristics. For example, one could impose a rule that, when choosing among two different agar patches, if the light exposure is similar enough, one should choose the patch with more agar. Such a rule seems reasonable on its face, but would run into problems of violating transitivity under certain conditions (Tversky, 1967) and therefore would not maximize utility. Nevertheless, heuristics have been shown to be used when people's preferences are not well-formed, and/or if the decision-maker wishes to avoid making trade-offs (Tversky, Sattath, & Slovic, 1988).

A topic of study over the years has been if, and how, preferences are constructed. A utility signal that is invariant to context and mode of elicitation should give rise to the same preference for the same option no matter how the question is couched. This concept, called procedure invariance, has been clearly shown not to be the case, and people do change their preferences depending on how the question is asked, and even slime molds change preferences based on what other options are available (for review on humans see Payne, Bettman, and Johnson (1992); for slime mold see Latty and Beekman (2010)). However, it is debated as to how people construct their preferences – whether such a signal is more utility-like, but just influenced by context at the time, or something closer to heuristics, in which case it would be rule-based and not dependent on utility (Fischer & Hawkins, 1993; Mellers, Ordoñez, & Birnbaum, 1992).

Behavioral economic concepts of value were intended as *as-if* models (meaning people behave *as if* they are maximizing utility), and not as a way of describing how things actually worked in their heads (Kable & Glimcher, 2009). Similarly, behavioral studies are limited in their ability to assess the process by which people construct preferences. For organisms more complex than our single-celled mold friend, we can use neuroscience tools to investigate how this utility concept is actually implemented in the brain, and reveal the timing of such representations during the decision process. In humans, meta-analyses show a well-established network of regions known to be involved in representing reward and preference, including the ventral striatum, posterior cingulate cortex, and ventromedial prefrontal cortex, or vmPFC (Bartra, McGuire, & Kable, 2013; Levy & Glimcher, 2012; J. Peters & Büchel, 2010). The focus of this dissertation will be on the vmPFC (an area of the frontal lobes around eye-level), and its involvement in decision-making.

The Ventromedial Prefrontal Cortex

The definition of the vmPFC can vary based on the paper or method of investigation. In the fMRI literature, it generally refers to the anterior cingulate cortex (ACC) below the genu of the corpus callosum [ventral portions of areas 24, 25, 32, --as defined by Petrides and Pandya (1994)], the frontal pole (area 10), as well as medial orbitofrontal cortex (area 14). In the lesion literature, it often additionally includes central and lateral parts of the orbitofrontal cortex (areas 11, 13, and 12/47), and moreover the white matter as well as the grey matter. In this thesis, we will refer to this broader definition as “ventromedial frontal lobes”, in relation to discussions of studies of individuals with lesions.

The vmPFC and orbitofrontal cortex (OFC) contain two distinct patterns of connections, the medial network and the orbital network (Öngür & Price, 2000). The medial network, which includes the aforementioned ACC areas and medial OFC, is densely connected to limbic structures (e.g., amygdala, entorhinal cortex, hippocampus), as well as projecting to the ventral striatum. The orbital network, on the other hand, consists of the central and lateral OFC, and is characterised by connections to sensory inputs, primarily olfactory, gustatory, visual, and somatosensory, in addition to being connected to perirhinal cortex and more central parts of the striatum. The medial and orbital networks share connections with each other, particularly through intermediary areas like area 13 and posterior area 14.

The connections with limbic structures, then, puts the vmPFC at a particularly suitable position to represent preferences, gaining inputs from structures involved in emotion and memory, while projecting to reward-related centers like the nucleus accumbens. At the same time, the multiple sensory inputs in orbitofrontal cortex, and its connections to regions involved in recognition memory and action-learning, puts it in a position to represent learning and linking sensory inputs to outcomes (Haber, 2016).

Activity within overlapping regions of the vmPFC has been found to correlate with preferences for different categories of objects, suggesting that this region integrates value information for different types of stimuli into a common scale (Bartra et al., 2013; Chib, Rangel, Shimojo, & O'Doherty, 2009; Levy & Glimcher, 2012). The vmPFC's signal, moreover, describes a subjective assessment of the value of the options, which integrates dimensions of costs and benefits into a single representation that reflect individual differences (Hare, O'Doherty, Camerer, Schultz, & Rangel, 2008; Kable &

Glimcher, 2007; Plassmann, O'Doherty, & Rangel, 2007; Tom, Fox, Trepel, & Poldrack, 2007). More recently, it has been shown with multivariate methods that vmPFC represents a domain-general value signal, while central orbitofrontal cortex instead represents identity category-specific value (Howard, Gottfried, Tobler, & Kahnt, 2015; Howard & Kahnt, 2017; Pegors, Kable, Chatterjee, & Epstein, 2015). Individuals with damage to the vmPFC have been repeatedly found to be less consistent in preference-based choices between multiple kinds of goods (Fellows and Farah, 2007; Camille et al., 2011; Henri-Bhargava et al., 2012), suggesting a critical, general role for this region in valuation of stimuli.

Open Questions

Though the above evidence might show that value is robustly associated with the vmPFC, the opposite inference cannot be made. The vmPFC is involved in many other domains of cognition, a prominent one of which is schematic memory. fMRI studies have shown that vmPFC activity represents implicit relationships between elements within a task (Schuck, Cai, Wilson, & Niv, 2016; Zeithamova, Dominick, & Preston, 2012), and damage to the vmPFC has been shown to disrupt and weaken processing of semantic associations (Spalding, Jones, Duff, Tranel, & Warren, 2015; Spalding et al., 2018). One way of representing and inferring relationships between concepts is by situating them within a mental map, like one would buildings in a city – a theory known as the cognitive map (Tolman, 1948). Recently, fMRI studies have found that vmPFC appears to be involved in navigation through physical and conceptual space (Constantinescu, O'Reilly, & Behrens, 2016; Doeller, Barry, & Burgess, 2010), using the same mechanism of grid-like coding as the entorhinal cortex for spatial navigation (Hafting, Fyhn, Molden, Moser,

& Moser, 2005). This grid-like code has been proposed to underlie cognition broadly, including decision-making (Bellmund, Gärdenfors, Moser, & Doeller, 2018). In other words, making decisions between two movies would be like mentally traversing in a space made out of its attributes, e.g., genre and critical rating. Thus, though subjective value has been found to correlate with activity in the vmPFC, conceptual navigation represents a model of choice that could conceivably mimic such a signal that would not be value-based at all. The plausibility of such a model in decision-making has not yet been tested, and remains an open question.

Secondly, while there is an extensive literature on the vmPFC's role in preferences for choice tasks, it is not well-established as to how this process is conducted in the brain in preference tasks other than choice. That the same neural network would be active for a non-choice preference task is not to be taken for granted, given that people may change their preferences between choice and matching tasks (where one must come up with an amount for an attribute that would make one option equivalent to another), and it is thought that different strategies underlie choice in these different paradigms (Fischer & Hawkins, 1993; Tversky et al., 1988). In addition, there is little known about the temporal evolution of preferences in human vmPFC, and once again the available evidence is only on choice tasks (Harris, Adolphs, Camerer, & Rangel, 2011; Harris, Clithero, & Hutcherson, 2018). Thus, it is an open question as to how preferences are represented in non-choice preference tasks, and when this information arises during the decision process.

Finally, though individuals with vmPFC damage has been shown in the past to make more transitivity errors (Camille, Griffiths, Vo, Fellows, & Kable, 2011; Fellows &

Farah, 2007; Henri-Bhargava, Simioni, & Fellows, 2012), it is unknown if vmPFC damage causes fundamentally intransitive choices, or transitive but noisier choices. This distinction is important because the first possibility would imply that the vmPFC is necessary for value-based decisions *per se* (in which case decisions would be made without value, for example relying on heuristics), while the second would imply that the vmPFC is necessary for *an aspect* of value-based choice. This question has implications for whether the idea of value is a distributed system where each region of the frontal lobes computes attributes in a similar fashion, but contributes distinct aspects of value to the whole representation (Hunt & Hayden, 2017), or if the vmPFC represents the final common pathway for value representation.

My dissertation aims to describe vmPFC function during decision-making. It will address three questions: 1) does the vmPFC signal during decision-making truly reflect subjective value, or is it instead reflective of the conceptual navigation model? 2) how are preferences represented in the brain in a non-choice task, and when does it do so during the decision-process? 3) in what way is the VMF necessary to value-based choice?

General Methods

In this dissertation, I will use two cognitive neuroscience methods: analysis of neural data from functional magnetic resonance imaging (fMRI), and the study of the behavior of individuals who have sustained focal damage to specific brain areas. fMRI is an observational method in which participants perform cognitive tasks while lying in a scanner that measures the oxygenation level in blood flow (blood-oxygen-level

dependent, or BOLD, signal) throughout the brain. If the task elicits more activity in a certain brain region, then that region will consume more oxygen, therefore prompting more oxygenated blood flow there and increasing the BOLD signal. The evidence this method provides is *correlational*, meaning that the BOLD signal is correlated with neural activity — and hence an fMRI study tests the association of this activity with some task variable, or behavior. The advantage of this method is that it provides a relatively high resolution, whole-brain picture of activity. As such, it is possible to test various cognitive models and statistically assess how well they describe actual brain function. However, fMRI evidence does not tell us whether any brain region is necessary for the function in question.

My second method, study of individuals with focal lesions, provides evidence for necessity. A lesion study involves comparing the behavior of a group of individuals with damage to a target region (here vmPFC) against two control groups: an age and education matched healthy control group, and a group of individuals with damage to other parts of the brain (in this case, the frontal lobes) sparing the target region. The second control group is intended to rule out both that any observed behavioral differences are an effect of sustaining frontal lobe damage or having undergone a major medical event. If the vmPFC-damaged group performs worse than both control groups on a certain measure, it is taken as evidence that the vmPFC is *specifically* necessary for normal performance on that measure. This technique is inferentially powerful, but spatially imprecise (as there is no control over the extent or location of brain damage each individual sustains).

By combining the advantages of both methods in my dissertation, I will both be able to investigate models of representation and process in decision-making, as well as able to assess the *necessity* of the vmPFC to decision-making.

Research Overview

Chapter 2 addresses *what* the vmPFC signal represents during decision-making. While the vmPFC signal has been shown to correlate highly with subjective value in past studies, neural models inspired by spatial navigation have recently suggested that the vmPFC may be involved in navigation through conceptual space (Constantinescu et al., 2016). This has led to proposals that decision-making and other cognitive processes operate via mental navigation through a conceptual space made out of attributes (Behrens et al., 2018; Bellmund et al., 2018). Thus, in this chapter we assess both the theoretical plausibility and empirical fitness of this conceptual navigation model to vmPFC activity during decision-making. We find that the conceptual navigation model and subjective value model are highly correlated, so it is theoretically possible that the signal previously interpreted as subjective value could actually instead be reflective of mental navigation through conceptual space consisting of option attributes. We then sought to distinguish these two possibilities in a large dataset of 106 participants that performed an delay discounting task while going under functional neuroimaging. We found that the mental navigation model lacked support in the evidence, and the subjective value model remains the best explanation for vmPFC signal during decision-making.

After having established that the signal in vmPFC reflects subjective value, Chapter 3 addresses both how value is represented in a non-choice-based preference

paradigm, and *when* does this representation comes online during the decision process. We analysed the data of 37 participants who performed a bidding delay discounting task, where they were asked to decide the amount of money they would accept to receive immediately, in exchange for \$75 offered at varying delays. Participants entered their bid by scrolling through descending possible amounts of money and submitting their final decision. We found that first, that the brain does represent a similar pattern of value-related activity as found in choice tasks, and second, that the vmPFC represented the final bid only at the very end of the decision period, when the participants submitted their answers, and not any time prior. These results both support the idea of preference construction late into the decision, and demonstrates that neural representation between choice tasks and other types of preference-based tasks share common characteristics.

Finally, though evidence from both fMRI and past lesion studies show that vmPFC is critically involved in valuation, it is not known *how* it is necessary for subjective value. Specifically, is value abolished altogether in vmPFC damage, or does it remain but is altered in some way? I address this question in Chapter 4, by leveraging the idea that transitivity is a fundamental feature of value. Past studies have only presented choices between options once, so that violations of transitivity are considered in a deterministic way, which does not make clear whether the subject would always make such an error if given the same options again. To look at the question of whether VMF subjects are fundamentally transitive, it is necessary to present subjects with repeated pairs of choices, and look at whether they are consistent with transitivity in a probabilistic way, over the repetitions. We presented individuals with VMF damage, individuals with frontal damage sparing the VMF, and healthy controls with complete sets of pairwise

choices in three categories (art, chocolate brand, and gambles), where some of the pairs are repeated (to test probabilistic transitivity) and others were not (to test deterministic transitivity). We found strong evidence for the second possibility, where individuals with VMF damage were fundamentally transitive in tests of probabilistic transitivity, but still made more errors relative to control groups in deterministic definitions of transitivity. This finding is consistent with the theory that value is a distributed process over many brain regions, where other regions can compensate for the loss of VMF in calculating value. However, the VMF is necessary for stability in choices.

In conclusion, my three chapters advance the state of knowledge of the vmPFC's role in decision-making, by showing that 1) its signals are consistent with subjective value over an alternative model, 2) that value representations in vmPFC come online at the end of a decision process and not prior, and 3) that it is necessary for the stability and fidelity of decisions. The evidence I present here is consistent with the idea that vmPFC is involved in a value-based process during decision-making, yet that it cannot be the critical region solely responsible for value-based choice.

Chapter 2 -- Subjective value, not a grid-like code, describes neural activity in ventromedial prefrontal cortex during decision-making

Abstract

Recently, activity in the ventromedial frontal and entorhinal cortices has been shown to be modulated in a grid-like manner for navigation through conceptual as well as physical space. These findings have led to proposals that such a grid-like code could broadly underlie complex cognition, and specifically may be used for comparisons between multi-attribute objects in decision-making. We first assess the plausibility of this claim, showing theoretically that the activity correlated with subjective value observed in vmPFC in previous fMRI studies of decision-making could, in principle, reflect navigation through a conceptual space defined by the option attributes. We then empirically test for grid-like modulation in a large fMRI dataset of individuals making intertemporal choices. Here, though, we find that grid-like model fails confirmatory tests and does not provide the best description of the neural activity during decision-making. Our results constrain the type of tasks for which grid-like modulation is observed in vmPFC and further confirm that subjective value remains a good description of neural activity in vmPFC during decision-making.

Introduction

Neural activity in the ventromedial prefrontal cortex (vmPFC) has been shown to correlate with the subjective value of expected or experienced outcomes across a wide variety of decision-making tasks and categories of outcomes (Bartra et al., 2013; Clithero & Rangel, 2013; Levy & Glimcher, 2012). Neural correlates of subjective value have been found in vmPFC using both functional neuroimaging in humans as well as single cell recording in non-human animals (Kable & Glimcher, 2007; Strait, Blanchard, & Hayden, 2014; Yamada, Louie, Tymula, & Glimcher, 2018). One potential explanation for these correlates is that vmPFC encodes a representation of expected subjective value that could be used to make decisions or to guide learning (Kable & Glimcher, 2009).

However, recent studies have suggested that the similar area of vmPFC observed in human neuroimaging studies of decision-making encodes representations important for navigation through real and conceptual spaces (Constantinescu et al., 2016; Doeller et al., 2010; Jacobs et al., 2013). Both intracortical recordings and fMRI studies have revealed activity in medial prefrontal cortex while humans navigated virtual arenas that reflected a hexagonal spatial pattern characteristic of grid cells (Doeller et al., 2010; Jacobs et al., 2013). Grid cells were first discovered in entorhinal cortex during spatial navigation and can provide an efficient representation of two-dimensional space (Hafting et al., 2005). Furthermore, the same potential fMRI signature of grid cells has been observed in vmPFC during navigation in a purely conceptual space (Constantinescu et al., 2016). Specifically, stimuli in that study could be characterized by two dimensions, and when subjects imagined traversing the stimulus space defined by those two dimensions, activity

in vmPFC showed a similar response pattern as that observed during two-dimensional spatial navigation.

These results raise the intriguing question of whether the patterns of activity observed in vmPFC during conceptual navigation might explain the engagement of this region during decision-making tasks. Many decisions involve a choice between two options that differ along multiple dimensions or attributes. For example, choices can be between foods that differ in their taste and health or between monetary options that differ in their amount and probability or amount and delay. In the same way that spatial navigation involves moving through the two dimensions of longitude and latitude, and the conceptual navigation studied in Constantinescu et al. (2016) involves moving through a space defined by two stimulus attributes, might decisions involve navigating a conceptual space defined by the attribute dimensions of the choice options?

Here we aim to assess whether grid cell-like activity reflecting conceptual navigation through attribute space can account for the value correlates previously observed in vmPFC during decision making. If true, this would demonstrate that these signals do not reflect any encoding of value, and instead can be subsumed under an account of vmPFC function in terms of grid cell-like activity and more broadly in terms of encoding a cognitive map of decision space. We first show theoretically that grid cell-like activity could in principle explain subjective value correlates, as a plausible construction of grid-like regressors for a two-attribute choice task is highly correlated with subjective value. We then empirically test which of these two models better explains BOLD activity in vmPFC, using a large pre-existing dataset of subjects performing a standard intertemporal choice task (Kable et al., 2017). Our results unambiguously show

that vmPFC activity in this task is correlated with subjective value and cannot be explained by grid cell-like modulation.

Methods

Task.

We used the intertemporal task from Kable et al. (2017) as the decision-making task in the following analyses. Participant chose between a smaller immediate reward and a larger later reward. The smaller immediate reward was held constant at \$20 today while the larger later reward varied in amount (A : \$21 ~ \$85) and delay (D : 20 days ~ 180 days) from trial to trial. Each trial displayed the amount and delay of the delayed option; the immediate option as not displayed. Participants had 4 seconds to make their choice.

Theoretical correlation between subjective value and grid regressors.

Firstly, we sought to show that, in theory, hexagonal grid modulation could mimic or account for activity correlated with subjective value (SV) during decision making. To do this, we calculated the SV signal at a given discount rate for a given range of amounts and delays, and then estimated the best-fitting hexagonal grid modulation for this signal. The correlation between the SV signal and its best-fitting hexagonal grid modulation was taken as a measure of the highest potential similarity between the two signals. This correlation was examined across a range of discount rates.

For a given discount rate, the SVs of delayed monetary outcomes were calculated using the hyperbolic model:

$$SV \text{ of } \$A \text{ in } D \text{ days} = \frac{A}{1 + kD},$$

where k is the discount rate. The amount A varied from 20 to 80 in increments of 5 (13 levels) and the delay D varied from 20 to 180 in increments of 5 (33 levels) resulting in a total of 429 SVs for a given k . This 429-element vector of SV was then regressed against two hexagonal grid modulation regressors in order to estimate the best-fitting hexagonal modulation signal:

$$SV = \beta_0 + \beta_1 reg_1 + \beta_2 reg_2 + \epsilon$$

$$reg_1 = \cos(6\theta), \quad reg_2 = \sin(6\theta), \quad \theta = \text{atan}\left(\frac{A-20}{D}\right).$$

In words, the grid-like model contains a linear combination of sine and cosine of the direction of the trajectory angle θ with 60° periodicity. The trajectory angle θ is taken as the angle between the immediate option (\$20 now) and the delayed option on each trial. This model implicitly assumes that the direction of “navigation” is a straight line between the immediate and the delayed option.

The Pearson correlation between SV and the fitted signal (i.e., $\beta_0 + \beta_1 reg_1 + \beta_2 reg_2$) was assessed as the similarity statistic between the two signals. This procedure was repeated for 51 levels of k whose base-10 log ranged from -5 (i.e., $k = 0.00001$) to 0 (i.e., $k = 1$) in 0.1 increments.

Dataset.

In order to empirically test whether hexagonal grid modulation can better explain activity patterns previously attributed to SV during decision making, we used the intertemporal choice fMRI dataset from Kable et al. (2017). 107 participants (between the ages of 18-35) completed two sessions of scans 10 weeks apart. Each scan session contained 120 binary choices. 5 participants whose choices were entirely one-sided in a session (i.e., always choosing the immediate option or always choosing the delayed option) were removed from further analyses, making the total count of subjects 102.

Participants were scanned with a Siemens 3T Trio scanner with a 32-channel head coil. T1-weighted anatomical images were acquired using an MPRAGE sequence (T1 = 1100ms; 160 axial slices, 0.9375 x 0.9375 x 1.000 mm; 192 x 256 matrix). T2*-weighted functional images were acquired using an EPI sequence with 3mm isotropic voxels, (64 x 64 matrix, TR = 3,000ms, TE = 25ms; tilt angle = 30°) involving 53 axial slices with 104 volumes. B0 fieldmap images were also collected for distortion correction (TR = 1270ms, TE = 5 and 7.46ms). The datasets were preprocessed via FSL FEAT (FMRIB fMRI Expert Analysis Tool). Functional images were skull stripped with BET (FMRIB Brain Extraction Tool), motion corrected and aligned with MCFLIRT (FMRIB Linear Image Restoration Tool with Motion Correction), spatially smoothed with a FWHM 9mm Gaussian Kernel, and high pass filtered (104sec cutoff). Registration was performed with FNIRT with warp resolution of 20mm (FMRIB's Non-linear Image Registration Tool) to a 2mm MNI standard brain template.

Data analysis.

Empirical demonstration of hexagonal grid modulation involves two steps (Constantinescu et al., 2016). The first step is to identify regions in the brain where variance in activity is significantly explained by two grid-angle regressors (i.e., $\cos(6\theta)$, $\sin(6\theta)$). The second step is to show that a given person's unique grid angle, calculated from these two hexagonal regressors, is consistent across time, and that such consistency is only observed for 6-fold modulation and not for other numbers of folds (e.g., 4, 5, 7, 8).

For the first step, we ran a whole-brain GLM using FSL FEAT on the first session's data. Three regressors of interest were used: the event regressor that modeled average activity of all trials, and the two hexagonal grid angle regressors ($\cos(6\theta)$, $\sin(6\theta)$, $\theta = \text{atan}((A-20)/D)$). All three regressors were time-locked to the beginning of the trial with 0.1 duration. Standard 6-parameter motion regressors were also included as nuisance variables. We calculated the f-stat for the two hexagonal regressors and converted it to a z-stat¹. We performed group-level permutation testing using the 102 subjects' z-transformed f-stat map with threshold-free cluster enhancement to identify brain regions that are well explained by the two grid angle regressors.

For the second step, we identified ROIs in which we calculated each individual's grid angle and evaluated angle consistency across sessions. We created two spherical ROIs of 33 voxels (2mm voxels) from the peak activation coordinates reported by

¹ The conversion to a z-stat provided three benefits: 1) z-stats are more easily interpretable as they are from a standard normal distribution, 2) z-stats do not depend on degrees of freedom to calculate probabilities, and, most importantly, 3) the null distribution of z-stats is centered around 0 which allows for simple permutation testing by sign-flipping.

Constantinescu et al. (2016) in mPFC and ERC. To address any concerns that the loci of grid activation are different in our dataset from those of Constantinescu et al. (2016), we also defined two spherical ROIs of 33 voxels based on our peak activation coordinates from the GLM above, choosing the peaks in mPFC and ERC that were closest to the coordinates in Constantinescu et al. (2016). Finally, we also adopted the mPFC and ventral striatum (VS) ROI from Bartra et al. (2013), which identified consistent SV effects across hundreds of studies through meta-analysis.

In order to test for the consistency of hexagonal grid angles, we first needed to calculate the individual's unique grid angle. For a n -fold modulation, the individual's grid angle was calculated by first running a GLM on the first session's data with $\cos(n\theta)$ and $\sin(n\theta)$ as the modulators. Then we calculated the average coefficients for the cosine and sine regressors within the ROI ($\beta_{cosavg}, \beta_{sinavg}$). The n -fold grid angle was then calculated:

$$\phi_n = \frac{\text{atan}\left(\frac{\beta_{sinavg}}{\beta_{cosavg}}\right)}{n}.$$

Subsequently, the consistency of the grid angle was tested on the second session by calculating the ROI's z -stat (converted from t -stat) for the following GLM regressor:

$$\cos(n(\theta - \phi_n)), \theta = \text{atan}\left(\frac{A-20}{D}\right).$$

If there is significant hexagonal modulation in an ROI, the average coefficient for this regressor when $n = 6$ should be significantly positive, but not so when $n = 4, 5, 7$, or 8 .

Results

Theoretical correlation between subjective value and grid regressors

We first show via simulation that hexagonal grid modulation, of the form previously reported in medial prefrontal cortex, could account in theory for previously reported activity correlated with SV in this region. Using an intertemporal choice task as an example, we calculated the best-fitting hexagonal modulation for different theoretical SV signals that assumed different discount rates. **Figure 1** shows the correlation between an SV signal and its most similar hexagonal modulation signal at various discount rates.

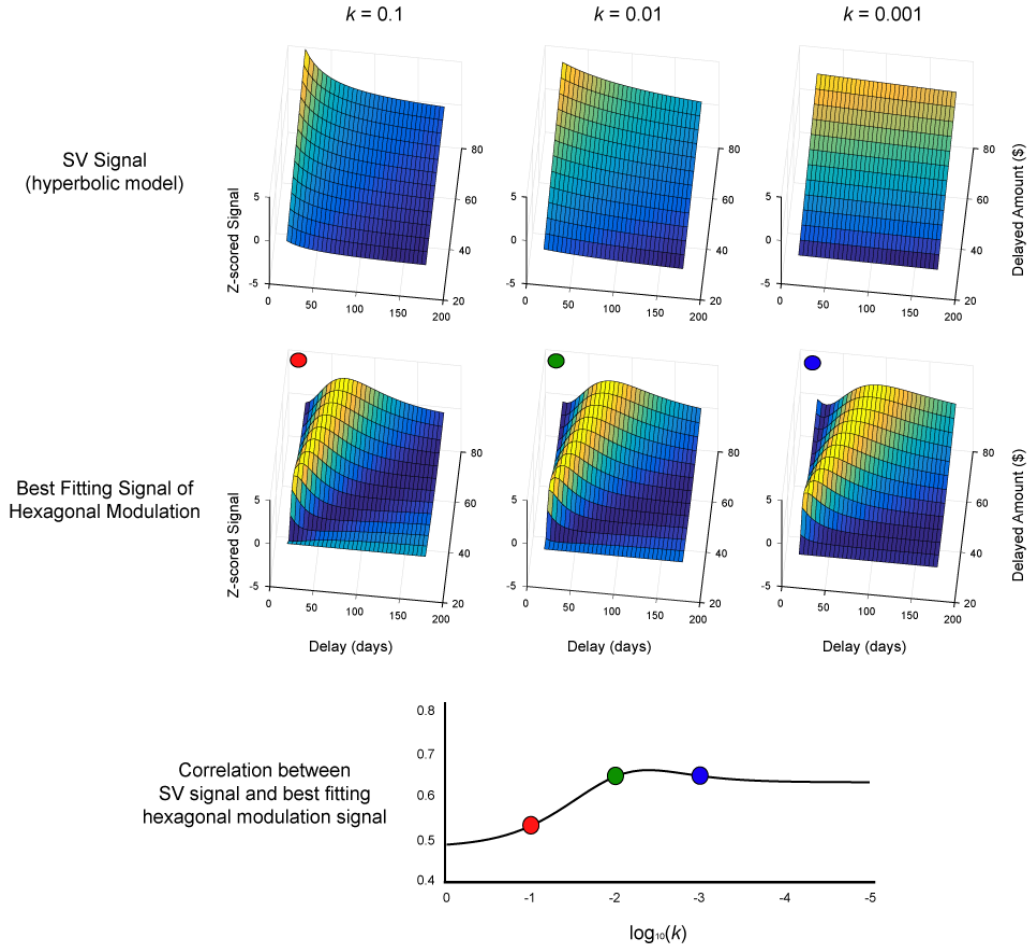


Figure 2-1. Correlation between subjective value signal and its most similar hexagonal modulation signal. The top three panels show simulated SV signal for various delayed amounts at various discount rates and the bottom three panels show the best fitting hexagonal grid modulations. The correlation between the two signals are provided below at various discount rates.

The correlation between the two signals ranges between $r = 0.5$ and $r = 0.7$ depending on the discount rate. These correlations show that it is possible for hexagonal grid modulation and a SV signal to be confused with each other, and therefore hexagonal grid

modulation provides a possible alternative theoretical account of previously reported activity correlated with SV.

Data.

We then examined whether hexagonal grid modulation does account for SV-related activity in actual data, using an intertemporal choice fMRI dataset from Kable et al. (2017). As previously reported, widespread regions show activity correlated with SV, calculated for each subject individually using their discount rate. These activations are consistent across two scanning sessions separated by 10 weeks (**Figure 2**).

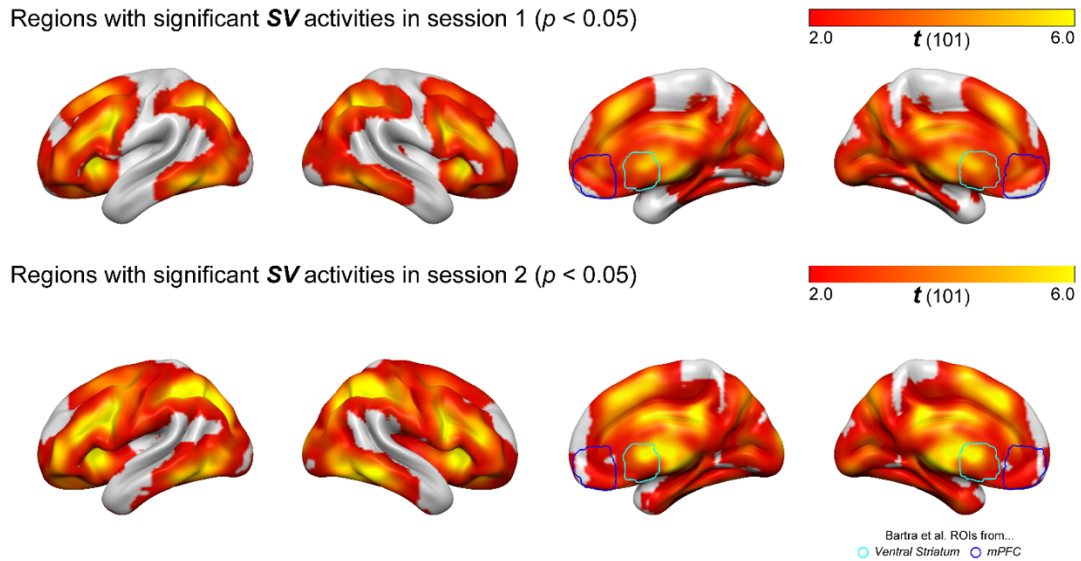


Figure 2-2. Regions with significant SV correlation in session 1 and session 2. Top panel shows the t-statistics (estimated from permutation) of regions that show significant correlation with SV signal ($p < 0.05$ with permutation testing). Bottom panel shows the t-statistics (estimated from permutation) of regions that show significant correlation with

SV signal ($p < 0.05$ with permutation testing). Bottom right brains show overlays of the two ROIs from Bartra et al. (2013).

Though hexagonal grid modulation could in theory account for this SV-correlated activity, we do not find that it does so in this dataset. Testing the hexagonal grid modulation hypothesis requires two steps: (1) showing that activity in a region is significantly explained by hexagonal grid regressors, and (2) showing that the grid angle implied by those regressors is consistent across time. There are many brain regions in this dataset where activity was significantly explained by the combination of two grid angle regressors ($\sin(6\theta)$, $\cos(6\theta)$). In fact, almost the entire brain reaches statistical significance for the F-test of these two regressors (Fig. 2). The peaks in this map include mPFC and ERC, as well as several other regions including somatosensory cortex, right TPJ, and dmPFC.




Figure 2-3. Z-transformed F-statistics of brain regions that are significantly explained by two hexagonal grid modulation regressors. Most brain regions were significant at $p < 0.05$ level. The right panel shows overlays of grid-cell ROIs: two spherical ROIs defined from peak coordinates of Constantinescu et al. (2016) and two

spherical ROIs defined from peak F-statistic activation that was closest to the coordinates of Constantinescu et al. (2016).

However, though the two grid angle regressors account for significant variance in activity in widespread regions including MPFC and ERC, the grid angle implied by these regressors is not consistent across time as required by grid cell hypothesis. We evaluated grid angle consistency in 4 ROIs in MPFC or ERC. These ROIs were defined based on the peaks in the above analyses (**Figure 3**) or previous reports of grid cell like activity (Constantinescu et al., 2016). We calculated the average grid angle in the first session's data in a given ROI, and then tested whether the activity in the second session was well aligned with the first session's grid angle. The grid angle consistency effect for hexagonal modulation was not significant in any ROI, nor was this effect larger in size for the 6-fold regressor than that observed for modulation at other folds (4 fold, 5-fold, 7-fold and 8-fold). (**Figure 4**). Note that the pattern observed in MPFC, of decreasing grid angle consistency from 4-8 fold modulation, is predicted of a signal that is correlated with SV (**Supplementary figure 2**).

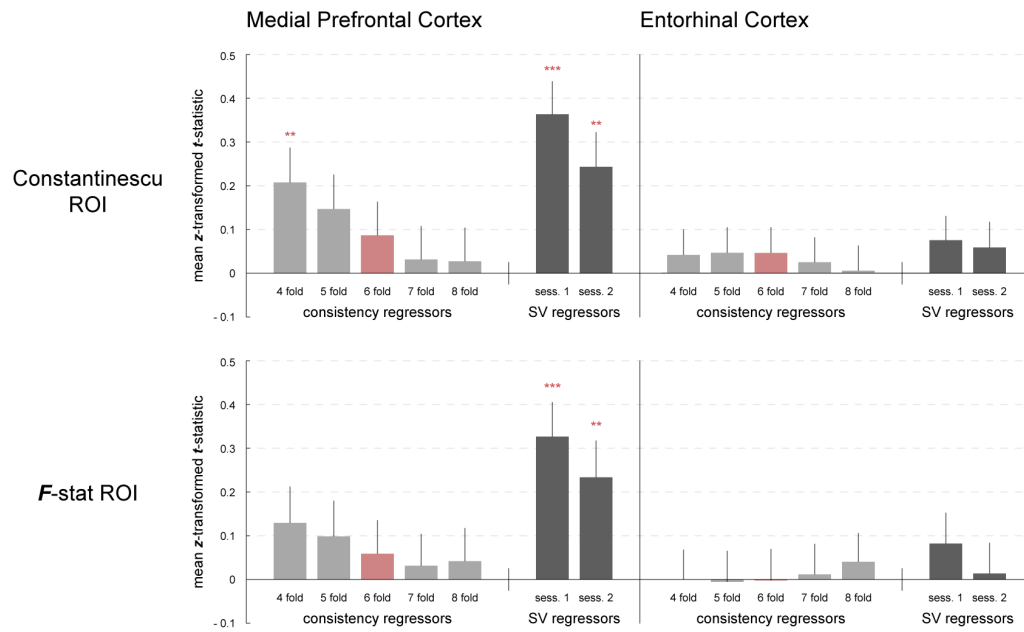


Figure 2-4. Effect of grid-angle consistency and SV in grid cell ROIs. The top two panels show consistency effects and SV effects in ROIs defined by a previous study by Constantinescu et al. (2016), the middle two show consistency effects and SV effects in ROIs defined by the closest F-statistic peaks to coordinates of Constantinescu et al. (2016). The left panels show ROIs in mPFC and the right two panels show ROIs in ERC. ** $p < 0.01$, *** $p < 0.001$ (uncorrected t-test against zero).

We also checked for grid angle consistency in regions previously associated with SV via meta analysis by Bartra et al. (2013). In both mPFC and VS, there is no evidence of hexagonal grid consistency and rather a similar pattern as in other mPFC ROIs: decreasing grid angle consistency from 4-8 fold manipulation (Fig. 5).




Figure 2-5. Effect of grid-angle consistency and SV in SV-defined ROIs. The left panel shows consistency effects and SV effects in mPFC ROI from Bartra et al. (2013) and the right panel shows them in VS ROI from Bartra et al. (2013). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ (uncorrected t-test against zero).

Discussion

In this paper, we assessed the idea that vmPFC activity during decision-making, commonly interpreted as subjective value, could be explained by a grid-like signal reflecting conceptual navigation through option attribute space. We showed that this idea

is theoretically plausible, in that regressors from a subjective value model are highly correlated with those of a grid-like navigational model, implying that past findings of subjective value-related activity could in principle be consistent with the conceptual navigational explanation. Thus, it was necessary to empirically test whether the conceptual navigation model can, in fact, appropriately describe activity during decision-making. We assessed the fit of this model, replicating the methods of Constantinescu et al. (2016), to the fMRI data of a large dataset of subjects performing a standard two-attribute decision task (intertemporal choice) across two sessions. In the critical analysis, we assessed the consistency of subject-specific navigational angle within a region of interest between two days of data collection, as well as whether 6-fold modulation (a characteristic of grid cell activity) can describe activity better than modulations at other frequencies. We found neither to be the case: the subject-specific navigational angle was not consistent between two days of data collection, and 6-fold modulation did not outperform control models with modulations at other frequencies. Thus, there is no evidence that the grid-like conceptual navigational model explains vmPFC activity during decision-making, and subjective value remains the best description of this activity.

Our results constrain the implications of Constantinescu et al. (2016), by limiting the conditions under which grid-like activity is observed in the vmPFC. Our standard intertemporal choice task is obviously very different from the sort of mental navigation demanded in Constantinescu et al. (2016). Their study involved extensive learning of a novel conceptual space by the participants prior to the navigational task; ours takes advantage of a spontaneous one created by the option attributes as they are presented. Their navigational task involved a period of mental simulation where the participants are

asked to “imagine” the progress of the stimulus, similar to spatial navigation tasks, whereas we simply ask participants to make a choice. Nevertheless, the two tasks share important similarities in the sense that both operate within a two-dimensional space where grid-like representation of the task structure might be expected (Behrens et al., 2018). Our finding means that the task space required when making decisions between options is one scenario that does *not* provoke grid-like representations, and thus grid-like coding cannot explain all cognitive activity in the vmPFC during such tasks.

One potential objection to our results might be that we do not see grid-like representations because we are not considering either the appropriate two-dimensional mental space or method of mental travel. One assumption made in our analyses was linear spacing between each unit in amount and delay. As we know that human beings do not generally perceive numbers in a linear way, but rather loglinearly as numbers become larger (Zauberman, Kim, Malkoc, & Bettman, 2009), perhaps loglinear spacing would be more appropriate. Indeed, in loglinear space, the grid-like model correlates even more highly with the subjective value model (see Supplementary materials). However, when tested empirically, the grid-like model with loglinear spacing also fails to account for BOLD activity. Another assumption we made in our grid-like model was that the subjects would “navigate” in a straight line between the immediate and delayed options. There are other conceivable assumptions such an analysis could make; however, we chose these assumptions precisely because they resulted in high correlations between the subjective value and grid-like regressors, meaning that the latter could plausibly mimic past results seen with the former. Thus, we have shown that the most plausible forms of grid-like model did not describe decision-making activity in the brain. Finally, an inherent

limitation of the intertemporal choice paradigm is that choices with tradeoffs necessarily preclude “no-brainer” pairs (i.e., strictly dominated options), and therefore many of the possible angles of conceptual travel in the attribute space. However, we could have sufficiently detected evidence of sixfold modulation in the quadrant of angles we were able to test.

The vmPFC is important for a wide variety of functions, from learning and decision-making to schematic memory and social cognition. A possible explanation for its diverse function is that different subregions of the vmPFC serve different functions. Though we used the same ROI as in Constantinescu et al. (2016) for our analysis, fMRI does not allow for the resolution necessary to differentiate between populations of neurons that may have different functions. Another possibility that has been raised is that the vmPFC represents a cognitive map of the inferred, hidden structure of the current task (Behrens et al., 2018; Wilson, Takahashi, Schoenbaum, & Niv, 2014). We have presented here evidence against the cognitive map based on a grid-like code during decision-making. However, the nature of the coding scheme in vmPFC may instead depend on the demands of the task at hand. Subjective value is a representation that has long been known to afford useful features for decision-making; for example, using such a representation of multi-attribute options and choosing the maximal valued option is guaranteed to result in transitive, non-cyclical choices (Samuelson, 1937). Subjective value may therefore be the most efficient representation of the option space for the kind of decision tasks we studied here.

Supplemental Materials

Methods.

Theoretical correlation between subjective value and grid regressors in logarithmic space.

In addition to grid angle analyses in the manuscript, we also considered the possibility that the cognitive map may be represented in logarithmic scale such that the hexagonal grid angle only manifests when one scales both axes of the space logarithmically. Hence, we re-did all the analyses in the manuscript with logarithmic grid angle calculations. All methods are as described in the manuscript except for the calculation of angles which is now $\theta = \text{atan}((\ln(A) - \ln(20))/\ln(D))$ instead of $\theta = \text{atan}((A-20)/D)$.

We first assessed the theoretical correlation between SV signal and hexagonal grid modulation signal in logarithmic space. This was done by calculating a range of SVs for a given discount rate and then fitting the best hexagonal modulation signal (methods in main manuscript). The Pearson correlation between the SV and the fitted signal was assessed as the similarity between the two signals.

Simulation of angle consistency in grid space assuming subjective value signal

We were interested in finding out how the grid-like regressor would behave in within-subject, inter-session angle consistency analyses if the true underlying signal actually represented subjective value. This simulation would be useful to compare against our analyses in real data, to better assess how subjective value describes the pattern of results. Thus, an additional simulation was performed to predict the pattern of grid angle

consistency analyses if we assume the true signal actually represented SV, in both regular and logarithmic space. First, we calculated the SVs of 102 people in our dataset for both session 1 and session 2. A Gaussian noise with $\sigma = 10$ was added on to the SV to simulate fMRI noise in signal. Then, we estimated each individual's grid angle based on the first session's simulated SV signal by running the following regression:

$$SV_{sess1} + N(\mathbf{0}, 10^2 \mathbf{I}) = \beta_0 + \beta_1 reg_1 + \beta_2 reg_2 + \epsilon,$$

$$reg_1 = \cos(n\theta), \quad reg_2 = \sin(n\theta),$$

$$\theta = \text{atan}\left(\frac{A-20}{D}\right) \quad \text{or} \quad \theta = \text{atan}\left(\frac{\ln A - \ln 20}{\ln D}\right).$$

Based on the grid angle calculated from session 1 (ϕ_n), we used the following regression to assess the consistency effect:

$$SV_{sess2} + N(\mathbf{0}, 10^2 \mathbf{I}) = \beta_0 + \beta \cos(n(\theta - \phi_n)) + \epsilon.$$

The resulting t -statistic of β was converted to a z -statistic and then compared across different number of folds ($n = 4 \sim 8$).

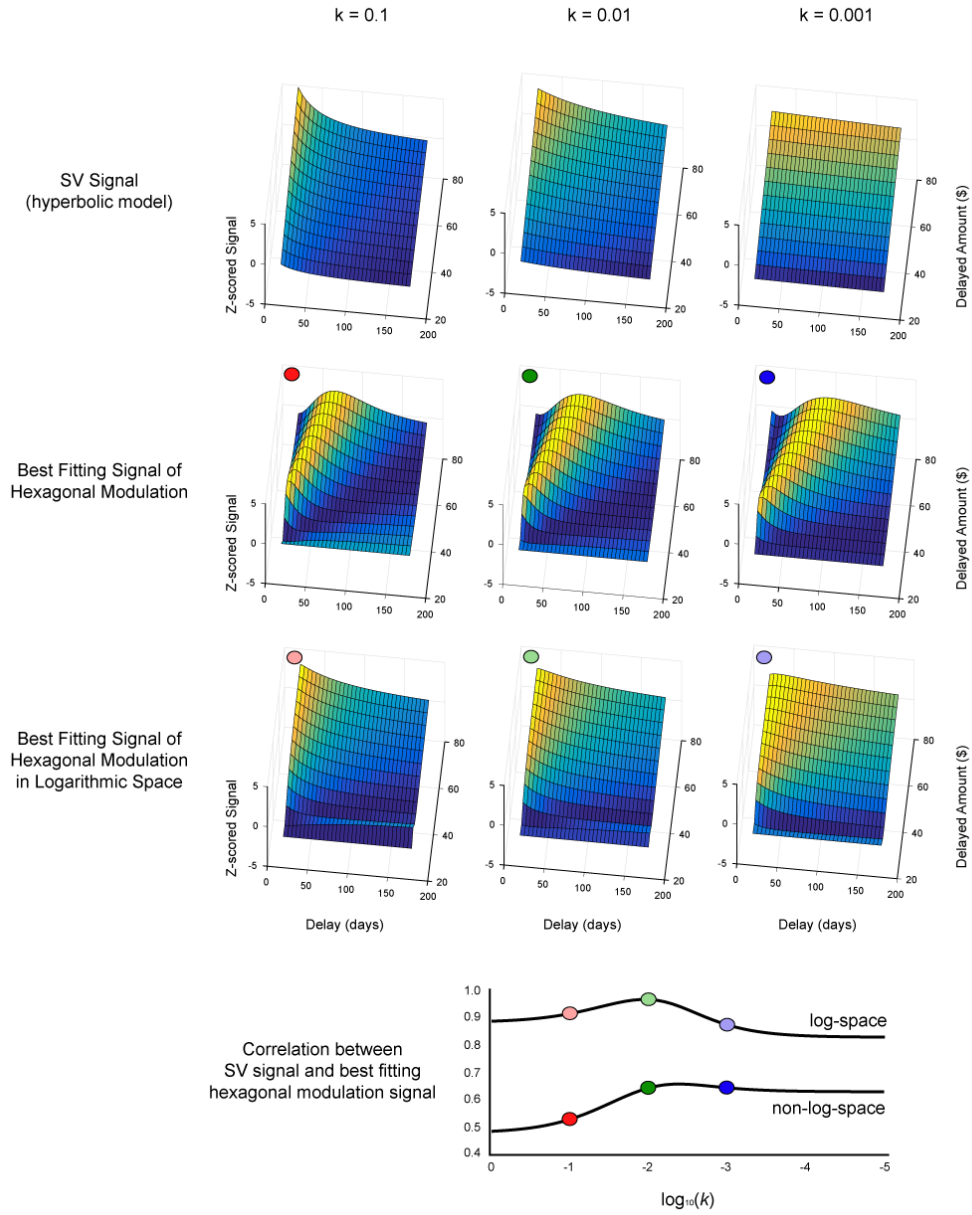
Analysis of data

We proceeded largely in the same manner as the main manuscript. The only difference being the calculation of grid angles (described above).

Results.

Theoretical correlation between subjective value and grid regressors in logarithmic space.

Firstly, our simulations show that it is quite possible for logarithmic hexagonal grid activity to mimic SV signal, even more so than non-logarithmic hexagonal grid activity. Across different discount rates, we found that the correlation between logarithmic grid angle activity and SV activity was very high ($r = 0.8\sim 1.0$) and always higher than the correlation between non-logarithmic grid angle activity and SV (**Suppl. Figure 1**).




Supplementary figure 1. Correlation between SV signal and its most similar hexagonal modulation signal in regular and logarithmic space. The top three panels show simulated SV signal for various delayed amounts at various discount rates and the middle three panels show the best fitting hexagonal grid modulations in regular space. The bottom three panels show the best fitting hexagonal grid modulations in logarithmic

space. The correlation between the two signals are provided below at various discount rates.

Simulation of angle consistency in grid space assuming true subjective value signal.

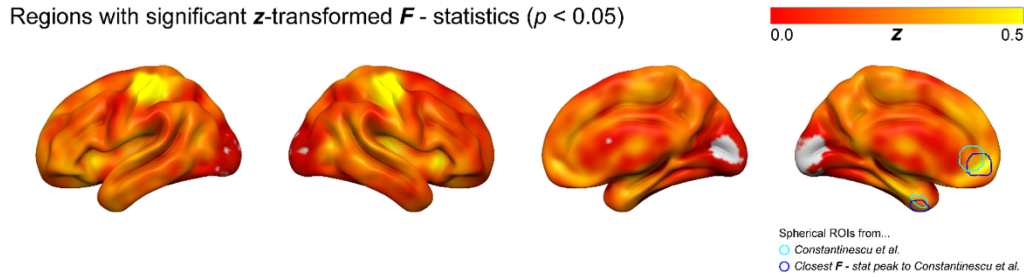
Our simulations of fMRI data that posit a true underlying SV signal shows that, in this scenario, the grid regressors in regular, non-logarithmic space would show a pattern of decreasing inter-session angle consistency as the folds of modulation of the angles grow. The consistency is highest for the four-fold regressor, which is the regressor that most closely resembles the pattern of subjective value (i.e., highest for low delays and high amounts, then decreases as the delays grow). The six-fold regressor, in this case, would not be the model with the highest angle consistency. In logarithmic space, the inter-session angle consistency would be flat across modulations of folds, and the six-fold regressor would once again not be the model with the highest consistency (**Suppl. Figure 2**).



Supplementary figure 2. Angle consistency in simulated fMRI data positing subjective value in non-logarithmic and logarithmic space. The left panel shows results of the inter-session angle consistency analysis for simulated fMRI data that assume underlying subjective value in non-logarithmic space, for the 6-fold grid regressor and the control folds. The right panel shows the same in logarithmic space. In both cases, the 6-fold regressor does not outperform the other folds for high angle consistency.

Analysis of data.


Just like non-logarithmic grid angle regressors, the logarithmic grid angle regressors ($\cos(6\theta)$ and $\sin(6\theta)$) together were able to significantly explain variance in most of the brain regions (**Suppl. Figure 3**).



Supplementary figure 3. Z-transformed F-statistics of Brain regions that are significantly explained by two logarithmic hexagonal grid modulation regressors.


Most brain regions were significant at $p < 0.05$ level. The right panel shows overlays of grid-cell ROIs: two spherical ROIs defined from peak coordinates of Constantinescu et al. (2016) and two spherical ROIs defined from peak F-statistic activation that was closest to the coordinates of Constantinescu et al. (2016).

We then repeated our analyses for inter-session, within-subject angle consistency analyses for each ROI, but this time in logarithmic space (MPFC and ERC ROIs defined from peak activation in Constantinescu et al., 2016 and from peak activation in our own GLM analysis above; the MPFC and VS ROIs from Bartra et al., 2013). Just like in analyses from non-logarithmic space, the grid angle consistency effect for hexagonal modulation was not significant in any ROI for logarithmic space, nor was this effect larger in size for the 6-fold regressor than that observed for modulation at other modulations (**Suppl. Figure 4; Suppl. Figure 5**).



Supplementary figure 4. Effect of grid-angle consistency and SV in grid cell ROIs.

The top two panels show consistency effects and SV effects in ROIs defined by a previous study by Constantinescu et al. (2016), the middle two show consistency effects and SV effects in ROIs defined by the closest F-statistic peaks to coordinates of Constantinescu et al. (2016). The left panels show ROIs in mPFC and the right two panels show ROIs in ERC.



Supplementary figure 5. Effect of grid-angle consistency and SV in SV-defined ROIs. The left panel shows consistency effects and SV effects in mPFC ROI from Bartra et al. (2013) and the right panel shows them in VS ROI from Bartra et al. (2013).

Chapter 3 -- Timing of value representation in ventromedial prefrontal cortex in a complex auction task

Abstract

The ventromedial prefrontal cortex (vmPFC) has been shown to be involved in value representations in many studies of decision-making. However, these past studies have mostly focused on choices between options, with very little work concerning preference tasks other than choice (such as ratings or auction tasks). As psychological theories have posited that those types of tasks involve different cognitive mechanisms compared to choice tasks, it is necessary to investigate whether the neural response to decisions in auction tasks would reflect the same value-related representations as have been found in choice tasks, and when those representations might occur. We conduct a study where participants play an auction task in which they selected the amount they would be willing to receive in exchange for a fixed amount of \$75 at a variable delay. We find that vmPFC and other regions in the value network correlated with the participant's bid, and that this representation occurred only at the end of the bidding period – i.e., the time of the bid submission. Our study shows that vmPFC activity reflects the same value-based process in a complex auction task as in choice tasks, and supports the view that value is constructed.

Introduction

Previous fMRI studies have shown that activity in the ventromedial prefrontal cortex is correlated with value across many different domains. Typically, participants in these experiments are given a choice, usually between two options (Levy & Glimcher, 2012). These studies have shown that the ventromedial prefrontal cortex is a region that appears to overlap in reward representation for various categories of rewards including foods, trinkets, attractive faces, and monetary gambles, with support for domain-general valuation across both univariate studies and multivariate studies (Chib et al., 2009; Howard et al., 2015; Pegors et al., 2015). This region thus may make common comparisons between dissimilar goods (Bartra et al., 2013), and this representation scales with a participant's subjective preferences (Kable & Glimcher, 2007; Tom et al., 2007).

Choices are, of course, not the only way people make decisions. We haggle over treasures at a flea market, put a bid on a house, and rate products or services online. These other forms of decisions have been studied in the laboratory using different kind of auction, rating or matching tasks. Such studies have demonstrated that different cognitive mechanisms can be engaged depending on the form of the decision. For instance, participants assign greater weight to an attribute when it is compatible with the response scale, or use a lexicographic step in choice but not matching tasks (Fischer & Hawkins, 1993; Tversky et al., 1988). Though this literature has made a clear case that decisions can depend on the manner in which they are elicited, little is known about how preference is expressed in the brain in tasks other than choice. Though several fMRI studies technically involve ratings (Hare, Camerer, & Rangel, 2009) or auctions (Plassmann et al., 2007), in these tasks there were only a few discrete responses (e.g., 5 point scale, or

\$0-\$4) making them similar to a choice between four or five options. Whether the same neural mechanisms are active in more complex and cognitively demanding auction or matching tasks therefore remains an open question. Answering this question would inform our understanding of how decisions may be constructed depending on the task demands.

In this study, we ask participants to bid on monetary rewards at different delays while measuring functional activity with BOLD fMRI. This task has two features of interest that are novel. First, in this complex auction task there were a large array of possible responses (from \$0-\$75; Cooper, Kable, Kim, and Zauberman (2013)), making the task similar to bidding and matching tasks studied behaviorally. Second, the response period for up to 10 seconds, so that the timing of neural signals within the task can be reasonably studied given the temporal resolution of fMRI. Thus, we have a valuable paradigm that is set to reveal the a) manner of value representation in a complex bidding task and b) its onset during the decision process. With this task, we seek to answer two questions: does the valuation network in the brain respond in a similar manner in this auction task as it does during choice? And, if so, at which point in time during the task does this signal occur?

Materials and Methods

Subjects.

Forty participants (16 male; 88% right handed) were recruited for this study from University of Pennsylvania and surrounding community. Participants had a mean age of 21.75 years ($SD = 3.27$ years). Participants were compensated for their time on two testing sessions; they also received incentive payments based on their decisions in auction task. All participants provided consent in accordance with the Institutional Review Board of the University of Pennsylvania.

Task.

All subjects completed two sessions involving different tasks: a subjective time estimation task (first session) and intertemporal decision-making task (second session). The two sessions were separated by an average of 10 d (SD 5 d, range, 4-21 d). Only the analysis of the latter task is discussed in this paper. For discussion of the subjective time estimation task, as well as behavioral results from the intertemporal decision task, see Cooper et al. (2013).

In the intertemporal decision task, participants bid on delayed monetary rewards. They selected the amount they would be willing to receive immediately in exchange for a fixed amount of \$75 at a variable delay (e.g., “I feel indifferent between receiving \$75 in 28d and receiving \$? now”). Bids were entered using an interface that allowed subjects to scroll through possible values. The bid amount always began at the maximum amount of \$75 now. Participant used two buttons to increase or decrease their bid and a third button to submit their bid.

The intertrial interval was variable, between 0.5 and 13.5s. In the “question period”, lasting 3s to 5s in duration, participants saw the delayed option they were asked

to evaluate. They were not able to respond at this time. After this period terminated, participants saw “\$75 now” on the screen and could begin to adjust this amount and select their bid. This “bidding period” lasted up to 10s. The amount on the screen when the participants pressed the ‘submit’ key or when the 10s timed out was registered as the participant’s response.

Participants knew that there was no penalty for not pressing the ‘submit’ key. Because of this, several participants adopted the strategy of not pressing the ‘submit’ key on some trials (22% of trials overall). Because the imaging analysis depended on having accurate response time data, we excluded the trials in which no ‘submit’ key was pressed. Three participants were excluded from analysis entirely due to a lack of sufficient valid trials (two participants had below 13% of trials in which the ‘submit key’ was pressed; the third participant had the same response of “\$75 now” for every single trial).

Payments.

In addition to a flat \$15 fee for participation, participants were paid according to their bidding decisions using an incentive-compatible Becker-DeGroot-Marschak protocol (Becker, DeGroot, & Marschak, 1964). A random trial was selected, and a random “counteroffer” (between \$0 and \$75) was generated for that trial. If the participant’s bid on that trial was greater than the counteroffer, they received \$75 at the specified delay. If the participant’s bid was below the counteroffer, they received the counteroffer amount immediately. Using such a protocol, participants are incentivized to bid their true valuations, as their bid affects the likelihood they will receive the item but not the price

they will pay for it. All payments were made using prepaid debit cards (described in Kable & Glimcher, 2007; Cooper et al., 2013).

MRI image acquisition.

Functional and anatomical images were collected using a 3T Siemens Trio scanner equipped with an eight-channel head coil. T2*-weighted functional images were collected using an EPI sequence (TR = 3s; TE = 30ms; 45 axial slices, 3 x 3 x 3mm; 64 x 64 matrix). Each scan consists of 150-152 images. All participants completed four scans in each session. High resolution T1 weighted anatomical images were collected using an MPAGE sequence (TI = 1100ms; 160 axial slices, 0.9375 x 0.9375 x 1.000mm; 192 x 256 matrix).

Imaging data analysis.

Functional images were analyzed using FSL (Jenkinson, Beckmann, Behrens, Woolrich, & Smith, 2012). Functional images were first interpolated in time to correct for staggered slice acquisitions, corrected for head motion using a six-parameter rigid-body transformations, and detrended and high-pass filtered (cutoff of 3 cycles/scan, or 0.0066 Hz) to remove low frequency drift in fMRI signal. Images were coregistered with each subject's high-resolution anatomical scan, and normalized into MNI space. Normalized data were spatially smoothed (kernel FWHM = 5mm) and thresholded.

To further control for excessive motion, we discarded any functional run in which > 5% of TRs exhibited > 0.5mm image-to-image displacement, and we discarded any

subjects with fewer than 2 functional runs passing this quality control criterion. Four additional subjects were excluded in this manner.

The general linear model for the intertemporal decision task contained 8 covariates of interest. These covariates modeled activity at the following four different time points in the trial: 1) onset of the question period (*Qon*), the onset of the response period (*Ron*), mid-point of the participant's response (*Rmid*), and the time at which the participant submitted the response (*Roff*). The first four regressors modeled the average activity at each of these time points, while the next four included the participant's bid on each trial as a parametric modulator at these time points. The bid values were mean-centered, and all regressors were convolved with a hemodynamic function included with the FSL package. All of these regressors only included trials in which the participant pressed the submit button. The model also included an additional 8 regressors of no interest that duplicated the regressors named above, but for trials in which the participant did not press the submit button.

Whole brain group analyses were assessed for significance using permutation testing implemented by FSL's *randomise* function. Corrected p-values were calculated by sign-flipping the entire map with 5000 iterations, and the threshold-free-cluster-enhancement (TFCE) method was used to form clusters without the need for an arbitrarily defined cluster forming threshold (Smith & Nichols, 2009). Results were then thresholded at the $p = 0.05$ level (two-tailed).

Region of interest analysis.

We used regions of interest (ROIs) from Bartra et al. (2013), a quantitative meta-analysis reporting value-related neural signals during decision-making. Our region of interest was ventromedial prefrontal cortex (609 voxels at 3 x 3 x 3mm, centred on MNI coordinates -2, 40, and -8). We extracted the average z-value from each participant within the ROI, and performed a one-sample t-test against 0.

Results

Whole brain analysis.

We first performed a whole-brain analysis testing for activity correlated with the participant's bid across four different time points within the auction task. We only found neural correlates of value similar to those observed during choice tasks at the end of an auction trial. At the first time point, that of question onset, we did not find any significant activity. At the second time point, that of onset of the bid response period, we found activity negatively correlated with the participant's eventual bid in the occipital cortex. At the third time point, midway through the bid response period, we found effects in both visual and motor regions. Only at the final time point, that of response submission, did we find widespread neural activity that was correlated with the participant's bid, including positive effects in classical areas of the valuation network, vmPFC and posterior cingulate cortex (PCC) (**Figure 1**).

Table 3-1: Peak foci of BOLD**effects**

<i>MNI coordinates</i>					<i>peak value</i>
<i>Region</i>	<i>x</i>	<i>y</i>	<i>z</i>	<i># voxels</i>	<i>(z)</i>
<i>Response onset</i>					
<i>Negative effects</i>					
R. inf. Occipital	34	-90	-6	644	1.04
L. inf. Occipital	-28	94	-6	374	1.07
L. mid. Occipital	-44	-74	2	43	0.81
R. mid. temporal	48	-70	0	173	0.95
R. mid. Temporal pole	48	10	-32	7	0.75
<i>Response midperiod</i>					
<i>Positive effects</i>					
L. Precentral gyrus	-36	-18	50	2473	0.88
L. sup. frontal gyrus	18	-10	70	404	0.85

L. Rolandic operculum	-44	-2	10	88	0.75
L. mid cingulate	-8	14	34	24	0.62
L. inf. Operculum	-40	10	18	23	0.7
<i>Negative effects</i>					
L. inf. Occipital	-28	-96	-6	570	1.38
R. inf. Occipital	30	-94	-4	340	1.44
R. mid. temporal	50	-72	4	38	1.04
<i>Response offset</i>					
<i>Positive effects</i>					
VMPFC	2	50	16	6911	1.00
Cuneus	2	-84	34	1948	1.12
L. angular gyrus	52	-60	42	1209	0.96
R. angular gyrus	-54	-54	32	1115	1.13
R. inf. orbital gyrus	54	38	-6	554	0.86
R. mid temporal gyrus	60	-34	-6	531	0.94

L. mid temporal gyrus	-62	-34	-8	268	0.85
PCC	-2	-44	34	238	0.93
L. Caudate	-18	18	18	63	0.73
L. superior temporal pole	-16	-42	18	30	0.82
R. SMA	8	18	68	9	0.65
<i>Negative effects</i>					
L. Precentral gyrus	-58	8	24	20070	1.64
R. Occipital	26	-92	-6	2038	1.49
L. Occipital	-32	-92	-6	1013	1.49

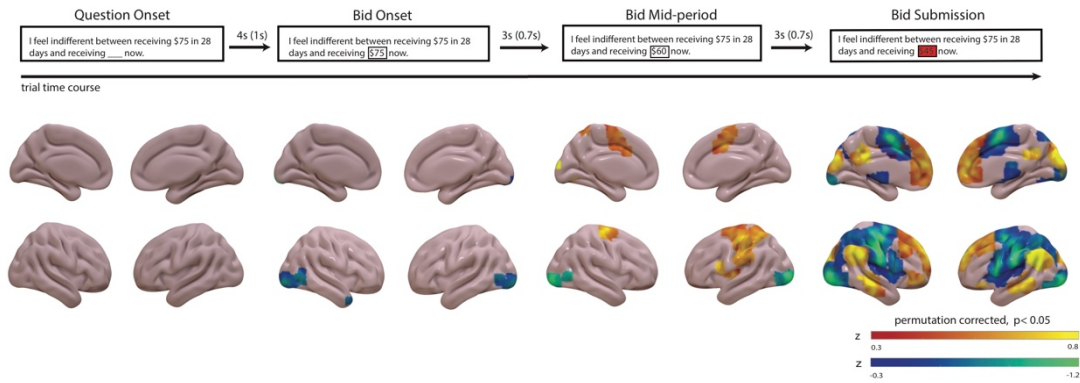


Figure 3-1 .Whole brain effects correlated with participants' bid response on each trial at the four time points in the trial. The boxes below each time point constitute example screens of the event the participant sees at that time point, and the arrows in between each event represent the average (with standard deviation in brackets) of the intervals subjects experienced between each time point. The “Bid mid-period” event represents the middle of the subject trial reaction time.

ROI results.

Region-of-interest analyses recapitulated these results, as activity in ventromedial prefrontal cortex tracked the participant's bid, but only near the completion of the trial when the decision was registered. Across subjects, activity in ventromedial prefrontal cortex was significantly correlated with the subject's bid only at the time of bid submission (R_{off} ; $t(32) = 3.86$; $p = 0.0005$), and not at any of the earlier time points in the trial (Q_{on} , ($p = 0.16$), R_{on} ($p = 0.15$), or R_{mid} ($p=0.78$)).




Figure 3-2. Activity in the vmPFC ROI related to participant's bid is only significant at the final time point, that of bid response submission. Asterisks denote significance at $p < 0.001$.

Discussion

In this paper, we measured neural activity during a complex auction task. Our goals were twofold: first, to see whether the neural correlates of value routinely observed in simpler choice tasks were also present during this form of decision; and, if so, secondly to test when during the decision task such value correlates were observed. We found that activity in the valuation network, including the ventromedial prefrontal cortex (vmPFC) and posterior cingulate cortex (PCC), was correlated with the participant's bid, but only late in the trial at the time of response submission. This result shows that the brain's valuation network does respond in a similar manner during an auction task as during choice, but only does so near the conclusion of the decision process.

Intriguingly, our results regarding the timing of value correlates at a longer timescale during auctions mirrors those observed on a shorter timescale during choices. Harris and colleagues have recorded EEG in several studies while participants made value-based choices and found that value signals, putatively linked to the vmPFC, typically appear late in the trial, around 500ms-800ms after stimulus onset (Harris et al., 2011; Harris et al., 2018). Similarly, studies that have argued that vmPFC arrives at choice through mutual inhibition between options have shown that activity in this region reflects the comparison of values only at the end of the trial (Hunt et al., 2012; Hunt, Woolrich, Rushworth, & Behrens, 2013; Jocham, Hunt, Near, & Behrens, 2012; Strait et al., 2014). Thus, though our task involves a very different paradigm and takes place over a longer period, our results are consistent with this body of evidence that signals related to value are only manifest in vmPFC near the end of the decision process. To date there are no neural models that would address the valuation mechanism in matching tasks. A model from computational psychology, decision field theory (Johnson & Busemeyer, 2005), posits that, in matching tasks, participants mentally sample bid amounts and then adjust them either downward (if the generated bid is too high) or upward (if bid is too low) until an indifference threshold is reached. In the future, more work should be done to investigate the neural evidence behind this theory.

Such a conclusion is broadly consistent with theories that value is computed or constructed online during decision-making, instead of being “read off some master list” (Tversky et al., 1988). However, Tversky et al. (1988) and much other work on “constructed preferences” has also focused on how different modes of preference elicitation – for example, choices versus auctions – occasion different types of cognitive

processes (Payne et al., 1992). Even more dramatically, this work has shown violations of procedural invariance, the principle that one's preference between two options ought to be the same no matter how you assess this. Instead, decision-makers can sometimes express different contradictory preferences in different response modes. One limitation of our study is that we do not directly compare choice and auction tasks in the same participants, so we cannot assess whether there are any differences in preferences across the two response modes, nor can we link any such behavioral differences to changes in activation in different brain regions. Clearly such investigations should be a priority for future work.

Contrasting the results we observed in the vmPFC, we did not find value-correlated activity in ventral striatum in our auction task, which we would have expected based on its status in the valuation network (Bartra et al., 2013; J. Peters & Büchel, 2010). The lack of signals in the ventral striatum may be due to the temporally extended nature of our task. For instance, the persistence task in McGuire and Kable (2015) also involves a prolonged period of waiting, and did not elicit value-correlated activity in ventral striatum. These results add to a growing list of findings suggesting some dissociations between vmPFC and ventral striatum during decision-making tasks (Hare et al., 2008; Knutson, Rick, Wimmer, Prelec, & Loewenstein, 2007).

Critically, though, our results do convincingly show that there are some neural processes in common across choice and auction tasks. We found that much of the same core valuation network that exhibits value-correlated activity in choice tasks (Kable & Glimcher, 2007), including regions like vmPFC and PCC, also exhibits activity correlated

with the subject's bid in this auction task. These results further establish the generality of this network's role in decision-making across a variety of tasks.

Chapter 4 -- Individuals with ventromedial frontal damage have more unstable but still fundamentally transitive preferences

Yu, L.Q., Dana, J., & Kable, J.W. (2018). *bioRxiv*, 384024.

Abstract

Transitivity of preferences (i.e., if one prefers A over B, and B over C, one should prefer A over C) is a hallmark of making rational, value-based decisions. Damage to the ventromedial frontal lobes (VMF) has been shown in previous studies to increase intransitive choice cycles (i.e., choosing A over B and B over C, but C over A). However, past studies have examined transitivity by treating preferences as deterministic rather than probabilistic, which could mask an important distinction in the critical role of the VMF in value-based choices: are individuals with VMF damage prone to choosing irrationally, or are they transitive, but simply more variable in what they prefer? We present individuals with focal VMF damage, controls with other frontal damage, and healthy controls with incentive compatible stimuli (artwork, brands of chocolate, and gambles) and have them make repeated choices between all possible pairs. Using cutting edge tests of a model of stochastic transitivity, and replicating previous analyses of transitivity that treat preferences as deterministic, we find that individuals with VMF damage made decisions consistent with stochastic transitivity. We also replicate previous findings that these individuals more frequently violate deterministic notions of transitivity. Our results are consistent with the hypothesis that individuals with VMF damage are not, in fact, more irrational, but do have noisier preferences. The implication is that the VMF is critical to maintaining the stability of preferences across time and

context during decision-making, rather than for the ability for choices to reflect preferences at all.

Introduction

A central assumption of many theories of choice is that decision-makers compare different options on a single dimension of subjective value and choose the highest valued option. Satisfying this assumption is equivalent to the observed choices being transitive (Samuelson, 1937). An example of transitivity is the following: if you choose to listen to Adele (A) over Britney Spears (B), and Britney over Celine Dion (C), then you would also choose Adele (A) over Celine (C). There is a strong argument that choices ought to be transitive, as an intransitive chooser could be exploited (e.g., as a “money pump”) and would get caught in choice cycles that do not advance towards any goal. Given this, one might expect that organisms develop internal representations of subjective value to ensure transitivity. Key studies in neuroeconomics have identified neural signals in the ventromedial frontal lobe (VMF) that scale with subjective value across different goods, in the firing rate of single neurons in the orbitofrontal cortex in monkeys (Padoa-Schioppa & Assad, 2006) and in the BOLD signal of ventromedial prefrontal cortex in humans (Bartra et al., 2013; Levy & Glimcher, 2012).

Consistent with the idea that neural signals in the VMF support value maximization, inconsistency has long been recognized as a hallmark of VMF damage: Phineas Gage was “capricious and vacillating” (Harlow, 1868) and EVR would drive on a single street for hours trying to decide on a restaurant (Eslinger & Damasio, 1985). More recently, individuals with VMF damage have been shown to make more

intransitive choices than healthy controls or individuals with damage elsewhere in the frontal lobe (Camille et al., 2011; Fellows & Farah, 2007; Henri-Bhargava et al., 2012). In the above example with the songstresses, an individual with VMF damage would be more likely to choose C(eline) over A(dele).

Axioms of rational choice, like transitivity, are usually stated deterministically. In contrast, behavior in experiments is probabilistic, because people can make different choices given the same pair of options over time (Luce, 1959; McFadden, 1980; Regenwetter, Dana, & Davis-Stober, 2011; Tversky, 1969). There are different ways to recast transitivity in probabilistic terms (Regenwetter et al., 2011; Tversky, 1969); however, testing any probabilistic model requires observing repeated choices over many instances of the same stimulus pairs. Noting a cycle (e.g., choosing C over A when one has chosen A over B and B over C) is not sufficient to disentangle whether one has fundamentally intransitive preferences versus variable preferences.

Previous studies have only asked individuals with VMF damage about their preferences between each pair of stimuli a single time. Therefore, the greater tendency of individuals with VMF damage to make intransitive choices in these experiments is consistent with two very different possibilities from a probabilistic perspective. One possibility is that the choices of individuals with VMF damage are fundamentally intransitive. In this case, their choices would not satisfy probabilistic notions of transitivity (e.g., by consistently and reliably choosing $C > A$ above). This could occur if individuals with VMF damage chose according to stimulus-response associations or rules that lack any higher order transitive structure, such as the lexicographic semiorder heuristic (Tversky, 1969). A second possibility is that the choices of individuals with

VMF damage are fundamentally transitive, but noisier. In this case, their choices would satisfy probabilistic notions of transitivity despite violating deterministic notions more often (e.g., they might choose A over C above with greater than 50% probability, but not 100% of the time). This could occur if individuals with VMF damage chose according to underlying values, but did so less reliably.

Here we test which of these two possibilities holds. The answer is both clinically relevant, as it sheds light on the nature of “capricious and vacillating” behavior after VMF damage, and theoretically relevant, as it determines whether VMF is necessary for choices to reliably reflect underlying values or for choices to be value-based at all.

Materials and Methods

Experimental design.

Participants. Fourteen individuals with focal damage to the frontal lobes were recruited from the Focal Lesion Database (FoLD) at the University of Pennsylvania, and ten were recruited from the Cognitive Neuroscience Research Registry at McGill University (Fellows, Stark, Berg, & Chatterjee, 2008). Individuals were eligible to participate if they had a lesion primarily affecting the frontal lobes. One individual was excluded due to incomplete data collection (the individual completed one session and was not able to be scheduled for the second). Fourteen females and 9 males were included in the final sample. Participants were tested a minimum of 5 months after injury (median = 10.29 years, range: 5 months to 17.75 years).

Participants were divided into two groups *a priori* based on location of damage, assessed with MR or computed tomography images by a neurologist blind to task performance. The ventromedial frontal lobe (VMF) group consisted of individuals who sustained damage to the VMF, while the frontal control group (FC) consisted of individuals who sustained damage to the frontal lobe sparing the VMF. Lesions were drawn on a common space [Montreal Neurological Institute (MNI) brain] by neurologists at the research sites blind to task performance. The overlap images for the groups are found in **Figure 1**. Damage in the VMF group was caused by aneurysm or subarachnoid hemorrhage in 5 cases, stroke in 2 cases, tumor resection in 3 cases, glioma in one case, and meningioma in 2 cases. Damage in the FC group was caused by hemorrhage, stroke or infarct in 7 cases, glioma in 2 cases, and meningioma in one case.

Age and education matched healthy controls (HC) were recruited from the corresponding Normal Control Databases of the University of Pennsylvania (N = 14) and McGill University (N = 6), including 15 females and 5 males (**Table 1**). They were free of neurological and psychiatric disorders. All subjects provided informed consent and were compensated for their time. The study protocol was approved by the institutional review boards of both the University of Pennsylvania and McGill University.

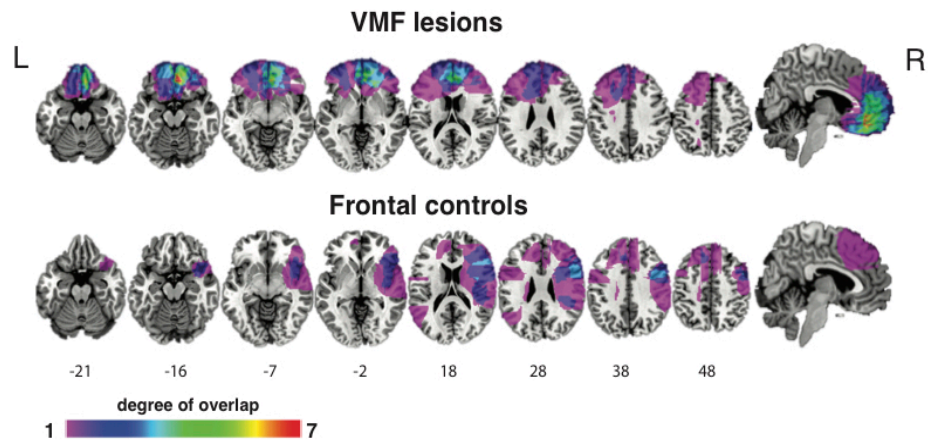


Figure 4-1. Overlap images of the VMF and frontal control lesion groups. Numbers below slices indicate the MNI z-coordinates. Colors indicate extent of overlap. L = left; R = right.

Table 4-1. Demographics of participants.

Group (n)	Gender	Mean age (sd)	Education in yrs
VMF (13)	7F:6M	59 (15)	14
FC (10)	7F:3M	66 (8)	14
HC (12)	15F:5M	62 (8)	15

Apparatus. All tasks were programmed using EPrime 2.0 (Psychology Software Tools). Participants were tested at the Hospital of the University of Pennsylvania, at the MNI, or at their own home in the greater Philadelphia or Montreal area. Participants saw stimuli on a laptop monitor and responded using the 1 and 0 keys of the keyboard.

Stimuli. Stimuli consisted of images of artwork, chocolate bars, and pie charts representing gambles. There were two sets of stimuli: 10-11 stimuli for each of the categories (10 for chocolate bar brands, 11 for art and gambles) used in non-repeated choices that allow deterministic tests of transitivity (set A), and 5 stimuli for each of the categories (art, chocolate bar brands, gambles) used in repeated choices that allow probabilistic tests of transitivity (set B). Choices constructed using set A and set B stimuli were intermingled in each block. For each category, we strove to design option sets in

which the options were close in preference, as intransitive choices are less likely between items that have widely different values.

The artwork stimuli were paintings that were rated highly by participants in Vaidya and Fellows (2015a). The set B stimuli consisted of 5 paintings by Monet, which were all within the top 20 most highly rated paintings by those subjects. We selected Monet as he was the artist that occurred most frequently in the top 20 rated paintings of Vaidya and Fellows (2015a). The 5 selected paintings were roughly similarly preferred (i.e., chosen with close to the same frequency in pair-wise choices across the whole sample) in a sample of 107 participants recruited from Amazon Mechanical Turk. Set A consisted of paintings of the similar style/era (Impressionist, Romantic periods) in the top 40 ranked paintings of the Vaidya and Fellows (2015a) stimuli set.

The chocolate bars were from five brands (Lindt, Godiva, Ghirardelli, Dove, and Cadbury). We selected five brands that were roughly similarly preferred across the population. These brands were being sold for similar prices, were rated similarly on a seven-point scale by a sample of 103 participants from Amazon Mechanical Turk (mean rating = 5.76), and were selected at roughly similar frequencies in pair-wise choices across another sample of 101 Mechanical Turk participants. Milk chocolate bars from each of the 5 brands were in set B, while dark chocolate and dark chocolate almond bars from each brand were in set A. The stimuli consisted of publicly available pictures of the front side of the chocolate bar packaging.

We used sixteen gambles of equal expected value (\$8.80). The stimuli consisted of a pie chart showing the probability of winning, with text on top indicating both the

cash amount to be won and the probability of winning. The five set B gambles were the “Cash II” set in Regenwetter et al. (2011), which used contemporary monetary equivalents of the Tversky (1969) five gamble set. The probabilities were 28%, 32%, 36%, 40%, and 44%. Set A consisted of 11 other gambles with the same expected value (probabilities of 8%, 17%, 25%, 33%, 42%, 50%, 58%, 67%, 75%, 83%, 92%).

Procedure. Participants completed a binary forced choice task. On each trial, participants first saw a central fixation point for 1s, then a screen with two choice stimuli (placed to the left and the right of the center). Participants indicated which stimulus they preferred, by pressing buttons for left or right. Participants had as much time as they needed to make their selection. Following their selection, there was an inter-trial interval of 1s where a black screen was presented.

For set A stimuli, participants faced all possible pairings of either 10 (for brands) or 11 (for art and gambles) options, constituting 45 and 55 pairs in total, respectively. Each pair was faced once. For set B stimuli, participants faced all possible pairings of 5 options, constituting 10 pairs, and each pair was repeated 15 times. Therefore, there were 195 (for brands) or 205 (for art and gambles) total choices in each category across the entire experiment.

Choice trials were presented in blocks, in which participants made choices between items within a single category (art, brands, gambles). There were five blocks of choices for each category, containing 39 (for brands) or 41 (for art and gambles) trials each. Each block contained 9 or 11 choices composed from set A and 30 choices composed from set B. Choices from set A and set B were intermingled with each other

within a block, with the set A stimuli inserted into a block of B stimuli in positions randomly selected from a uniform distribution.

We took a number steps to reduce any potential memory effects for choices constructed with set B stimuli. We designed the sequence of trials so that: (1) the same pairing was not repeated within a minimum of 3 trials; (2) the same stimulus rarely appeared on immediately adjacent trials (no more than 9 times throughout the entire experiment); and (3) when the same pairing was repeated the choices immediately preceding and following that pairing differed from its previous occurrence (to minimize contextual memory).

Furthermore, the side on which stimuli were presented was counterbalanced across repetitions. Finally, we divided the experiment into two sessions, held on separate days for every subject except two (due to scheduling constraints). The two sessions were held on average 8.09 (sd = 11.73) days apart (excepting the two who were tested on the same day, the sessions ranged from 1 day to 57 days apart). We did not observe a significant correlation between total number of intransitive choices made across all participants (see explanation of measure below) and days between the two sessions ($r = 0.24$, $p = 0.12$).

Statistical analysis.

Deterministic tests of transitivity. All data was analyzed with MATLAB (Mathworks). We used the set A choices to perform deterministic tests of transitivity, replicating previous studies. We first determined the preference ordering within each category for each subject. The 10 or 11 options within each category were ranked according to how many times each was chosen by that subject. Then, for each trial, a choice was counted as intransitive if a lower-ranked item was chosen over a higher-ranked item. Following Henri-Bhargava et al. (2012), ties were maintained in the

rankings (i.e., more than option could have the same rank) to provide a more conservative definition of intransitive choices. Because the intransitive choice counts are not normally distributed, we used non-parametric statistics to test for group differences. We used Kruskal-Wallis tests to detect effects between groups, followed by one-tailed Wilcoxon ranked sum *post hoc* pairwise tests as appropriate (as several previous studies have found increased intransitive choices after VMF damage, we had strong hypotheses about the direction of the results). To test for within-subject effects, we used repeated measures analysis of variance (ANOVA) on rank-transformed data for the omnibus test and Wilcoxon signed-rank *post hoc* tests as appropriate.

Probabilistic tests of transitivity. We used the set B choices to perform probabilistic tests of transitivity, extending on previous studies. We first obtained the proportion of choices (out of a possible total of 15 choices) for each of the 10 choice pairs afforded by all possible pairings of the 5 options in each category. We then tested the random mixture model of preference by noting whether the choices violated the linear ordering polytope (LOP) (Regenwetter et al., 2011). The random mixture model states that a person's response comes from a probability distribution over all possible orderings of the stimuli. Thus, at any one time, preferences are transitive, but the transitive state that one is in can vary. The probability of a person choosing one option (X) over another (Y) in a binary choice is the sum of all the preference states in which X is preferred to Y. In a two alternative forced choice task, this is constrained by the triangle inequalities. For every distinct X, Y, and Z in a choice set:

$$P_{xy} + P_{yz} - P_{xz} \leq 1$$

Where P_{xy} denotes the probability of choosing X over Y, etc. For up to 5 options in a 2AFC task, satisfying the triangle inequalities, which together define the LOP, is necessary and sufficient for a set of choices to be consistent with the random mixture model.

For choice probabilities that did not satisfy the triangle inequalities, we used the Q-test (Regenwetter et al., 2014) software to determine whether the data were significantly outside of the LOP. Q-test uses maximum likelihood estimation to find the goodness of fit of the data at each vertex in the polytope, using a chi-squared bar distribution with simulated weights (Regenwetter, Dana, & Davis-Stober, 2010; Regenwetter et al., 2014). Any subject with choices in a category that produced $p < 0.05$ in this test were considered as significantly violating the LOP and thus, the random mixture model of preference.

Sensitivity of probabilistic tests. We performed several simulations to determine the sensitivity of the probabilistic test of transitivity, i.e., the rate at which this test would declare different forms of random or heuristic-based choice to be transitive. First, following Regenwetter et al. (2011), we randomly picked a choice probability for every pair from a uniform distribution (from 0 to 100%). As previously shown in Regenwetter et al. (2011), only about 5% of the choice datasets simulated in this manner satisfy the triangle inequalities. That is, only 5% of the possible set of choice proportions for 10 pairs/5 stimuli satisfy the random mixture model.

Second, we simulated an intransitive chooser who has an entirely consistent preference within each pair (i.e., choosing A 100% of time when it is paired with B) that

is unconstrained by any higher order transitive structure (i.e., the preference in each pair is independent from that of all other pairs). This type of intransitive chooser only satisfies the triangle inequalities about 12% of the time for choice proportions for 10 pairs/5 stimuli as in our dataset.

Third, we simulated an intransitive chooser using the lexicographic semiorder heuristic (LS; Tversky, 1969). The LS heuristic is easiest to demonstrate with the gambles stimulus set. Following Tversky (1969), we defined our LS rule as follows: if two gambles are adjacent (i.e., next to each other in the set in terms of probabilities/payouts), always choose the gamble with the higher payout (amount); for all other (non-adjacent) gamble pairs, always select the gamble with the higher probability. Such a chooser would never satisfy the triangle inequalities in our dataset. Together, the first three sets of simulations show that our probabilistic test is very sensitive to different forms of intransitive choice.

Finally, we simulated a completely random chooser (i.e., someone who flips a coin on every single trial). The choice proportions for such a random chooser are given by the binomial probabilities with $p=0.5$. Such a chooser satisfies the triangle inequalities 80% of the time in our dataset (5 stimuli, 10 choice pairs repeated 15 times). This high percentage is not unexpected, as 50% choice probabilities across all pairs is consistent with the random mixture model (i.e., $0.5 + 0.5 - 0.5 < 1$). We use this rate below to assess whether the behavior of VMF subjects is consistent with completely random choice.

Drift diffusion modelling and analysis of reaction times. We calculated ranks of options similar to the method we used in the set A (deterministic transitivity) above,

where the option that was chosen most often overall was ranked first, and the option chosen second-most was ranked second, etc., and broke ties by looking at which options were more often chosen more than half of the time in every pair (Henri-Bhargava et al., 2012). It was necessary to break ties here for the purposes of calculating the effect of value distance on reaction times (RTs). Three subjects still had tied ranks after this process, in one category each: two are HC subjects in the gambles domain, the other is a VMF patient from the Art domain. These subjects in these categories only are dropped from the ANOVA analysis and drift diffusion modelling below.

We fit a drift diffusion model (Ratcliff, 1978) to the choices and RTs from all set B choices for every other subject and category in our experiment. We modelled the decision process as a decision variable (DV) that increased linearly with a slope $d*v^\alpha$, where d was the drift rate, v was the value difference of the options (expressed as the absolute rank difference between the two items for that individual), and α was an exponent accounting for potential non-linearities in the effect of rank difference. We also assume that at each time step there is Gaussian noise added to the DV, with a standard deviation of ε . We assumed 10ms time steps. We also assume there is a non-decision time (ndt) before accumulation begins, and an initial value (int) of the DV that is constant across trials. Choices are made when the DV crosses a threshold.

Thus there are five free parameters: d , α , ε , int and ndt . Note that the threshold was a fixed parameter across subjects, as one of the threshold, d , or ε must be fixed for the other two parameters to be estimable. We chose to fix threshold after a model-comparison process showed that option to provide the best model fits. Threshold was held constant at

(+/-) 0.15. Values for d are sampled between 0 and 1, for ε are sampled between 0 and 1, for α are sampled between 0 and 3, for int are sampled between the threshold bounds, and for ndt are sampled between 0 and the minimum RT minus 10ms for that subject.

To fit these free parameters, we first calculated the cumulative probability that the DV crossed the threshold for the subject's choice (T_{correct} or $T_{\text{incorrect}}$, where “correct” was defined as choosing the option of higher rank) across all time steps. For each trial, we then calculated the joint likelihood of the subject's choice at the time which they made that choice (their trial RT, minus ndt), by taking the derivative of this cumulative probability at the timestep of the subject's choice (every 10ms to the maximum RT for the subject). The model was then fit using the MATLAB function *fmincon*, where the cost function was defined as the sum of the negative log likelihoods of the instantaneous probabilities of the subject's choices and RTs in all trials. The fitting procedure was repeated 10 times for each subject, with each iteration varying in randomly sampled starting values for the free parameters as specified above; the parameters with the lowest log likelihood out of the 10 was taken for that subject. The model was fit individually to each of the three reward categories (art, brands, gambles) for each subject.

To look at differences in DDM parameters between groups across categories, we performed a mixed ANOVA on each of the free parameters, with group as the cross-subject factor and reward category as the within-subject factor.

Finally, we performed a mixed ANOVA with group and value distance as factors to look for the effect of value distance on RTs across groups.

Results

Deterministic tests of transitivity.

Individuals with frontal damage exhibit more choice cycles. A subset of the choices in our experiment, Set A, consists of a single instance of all pairwise choices from a total of nine or ten items within a category, which allows us to first replicate two previous studies of transitivity (Fellows and Farah, 2007; Henri-Bhargava et al., 2012). Combining all three categories (art, brands, gambles) in our experiment, we replicate the finding that individuals with VMF damage make more intransitive choices, though we do not replicate that this effect is selective to VMF damage in the frontal lobe (**Figure 2**). There was a moderate difference in intransitive choices in set A summed across all three categories (Kruskal-Wallis $H = 5.05$, $p = 0.08$). Because three previous studies have found increased intransitive choices after VMF damage (Fellows and Farah, 2007; Henri-Bhargava et al., 2012; Camille et al., 2011), we conducted planned comparisons between groups. Similar to previous studies, our VMF group (mean = 9.93%, $sd = 6.65$) made more intransitive choices than the HC group (mean = 5.71%, $sd = 4.05$; Wilcoxon ranked sums $Z = 1.64$, $p = 0.05$). Unlike previous studies though, our FC group (mean = 9.09%, $sd = 3.74$) also made more intransitive choices than the HC group ($Z = 2.05$, $p = 0.02$) and the difference between VMF and FC was not significant ($Z = 0.12$, $p = 0.45$).

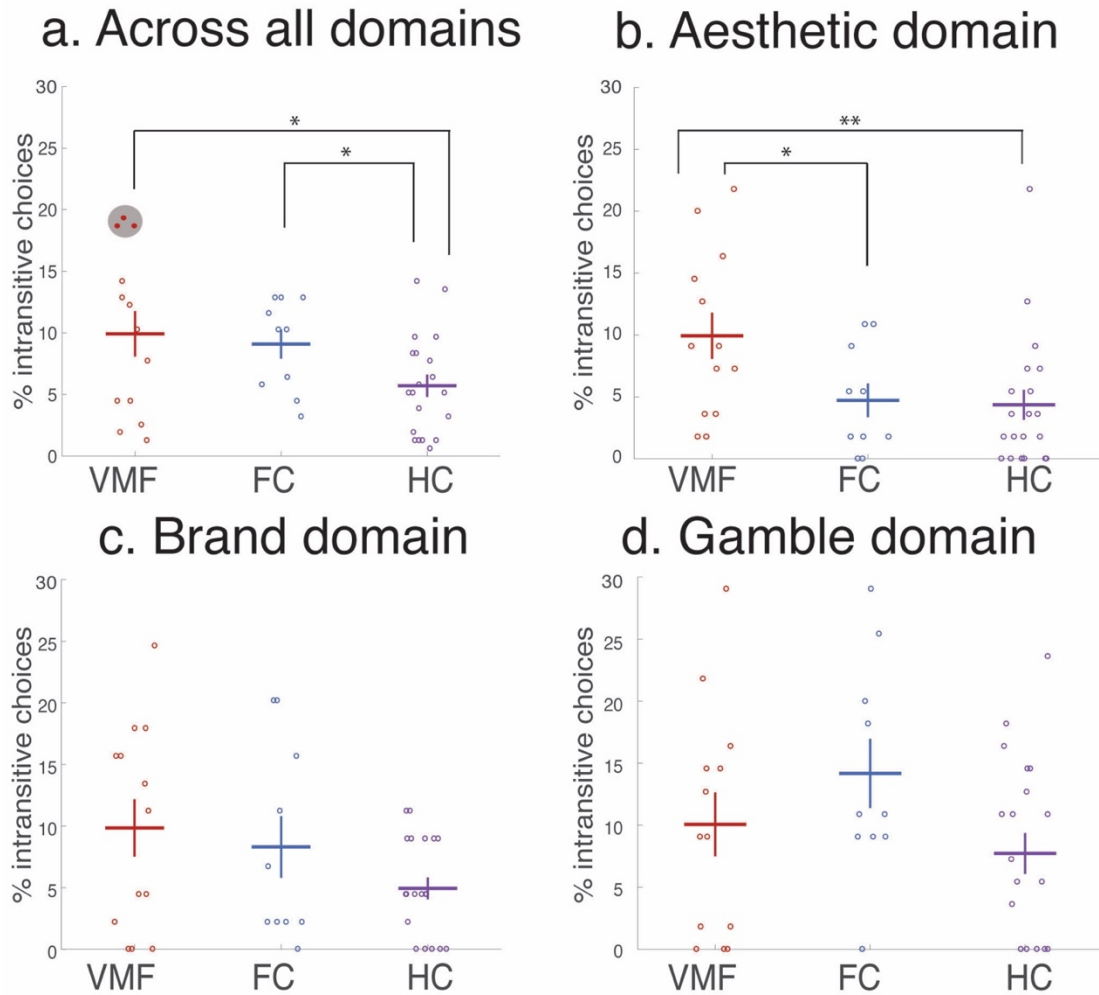


Figure 4-2. Group average and individually plotted intransitive choices in (deterministic) set A across a) all domains, and b-c) in each reward domain. Filled-in dots encircled in gray denote the VMF subjects whose errors were significantly higher compared to the HC group, and whose lesion extents are depicted in Figure 3. Error bars are standard errors of the mean.

Differences among reward categories. However, the analysis above obscures differences across individuals and choice categories that point to more specific effects of VMF damage. We first examined how intransitive choices in set A differ across choice categories. In the one choice category used in previous studies of transitivity, art, there was significant difference in intransitive choices across groups (Kruskal-Wallis $H = 7.62$, $p = 0.02$), which replicated the previously reported pattern of selective VMF deficit. The VMF group (mean = 9.93%, sd = 1.86) made significantly more intransitive choices in the art category than both the FC group (mean = 4.73%, sd = 1.36; Wilcoxon ranked sum $Z = 1.91$, $p = 0.03$) and the HC group (mean = 3.64%, sd = 0.97; Wilcoxon ranked sum $Z = 2.62$, $p = 0.004$). In contrast, in the two categories that have not been used in previous studies, brands and gambles, we did not find significant differences between the three groups (brands, $H = 2.42$, $p = 0.29$; gambles, $H = 3.01$, $p = .22$ respectively).

In **Figure 2b-d**, it appears that number of intransitive choices is relatively stable across categories in the VMF and HC groups, but variable across categories in the FC group. Indeed, the effect of reward category is significant for the FC group ($F(2,18) = 3.88$, $p = 0.04$), but not for the VMF ($p = 0.92$) or the HC group ($p = 0.27$). In the FC group, the number of intransitive choices in the gamble category was significantly greater than in the art category ($Z = 2.40$, $p = 0.02$), while the differences between gambles and brands ($p = 0.19$) and art and brands ($p = 0.18$) were not significant.

Differences among individuals. We then examined how intransitive choices in set A differ across individuals. To do this, we considered each individual with a VMF or FC lesion as a single case, and compared their total number of intransitive choices (i.e., across all three categories) against healthy controls. We made this comparison using case-control t-tests (Crawford & Howell, 1998) which are modified to compare an individual against a normative group when the sample size is small. In the VMF group, four individuals made significantly more intransitive choices than healthy controls, before corrections for multiple comparisons (Subject 350: $t(19) = 2.04, p = 0.03$; Subject 10403: $t(19) = 3.28, p = 0.002$, Subject 12402: $t(19) = 3.13, p = 0.003$; Subject 775: $t(19) = 3.13, p = 0.003$). These differences remained significant in the latter three individuals after correcting for multiple comparisons using FDR (corrected $p = 0.023$ for all three individuals). Lesion extent of these three subjects are shown in **Figure 3**. In contrast, in the FC group, none of the individuals made significantly more intransitive choices than healthy controls (all $p \geq 0.05$ before multiple comparison correction).

This result suggests that a subset of individuals with VMF damage show the most pronounced increase in intransitive choices. However, we did not find evidence to support any particular account of this heterogeneity. The total number of intransitive choices (i.e., across all three categories) was not significantly correlated with lesion size (in cc's), whether considering all subjects with lesions (Spearman's $\rho = -0.14, p = 0.51$) or only those with VMF damage ($\rho = -0.13, p = 0.67$). Within the VMF group, the total number of intransitive choices was also not significantly correlated with lesion volume within a vmPFC mask defined based on value effects in fMRI studies (Bartra, McGuire & Kable, 2013; $\rho = -0.06, p = 0.83$). Finally, across all subjects, the total

number of intransitive choices was not significantly correlated with any of the demographic variables (gender, point biserial $r = 0.13$, $p = 0.39$; age, rho = 0.14, $p = 0.35$; education, rho = 0.24, $p = 0.11$).




Figure 4-3. Lesion tracings of the three individuals with VMF lesions who had significantly more intransitive choices compared to healthy control subjects, as determined by case-control t-tests. Red denotes areas where at least one of these subjects had a lesion; yellow denotes the areas where at least one of these subjects had lesions *outside* of all other lesion subjects. There was very little overlap in lesions within the three subjects (only maximally two out of three on only in a small number of voxels). Numbers below axial slices indicate the MNI z-coordinates.

Probabilistic tests of transitivity. Individuals with VMF damage make choices consistent with probabilistic models of transitivity. After replicating the finding that individuals with VMF damage make an increased number of intransitive choices, we next turned to the central question motivating our study, which is whether or not the choices of these individuals violate probabilistic notions of transitivity. To do this, we examined the subset of choices in our experiment, Set B, which involve 15 repetitions each of 10 different binary choices in each of the three categories. Set B provides sufficient data for evaluating whether the choices each participant made are consistent with the random mixture model, a probabilistic model of transitive choice. None of the individuals with VMF damage violated the random mixture model in any of the three domains (a total of 39 tests, see **Table 2**). Similarly, none of the individuals with frontal damage outside the VMF violated the random mixture model in any of the three domains (a total of 36 tests).

Table 4-2. Results of LOP analysis, by category

	Art	Brands	Gambles
<i>Respondent</i>	<i>p-value</i>	<i>p-value</i>	<i>p-value</i>
<i>Individuals with VMF lesions</i>			
1	✓	✓	✓
2	✓	✓	0.64
3	✓	✓	0.83
4	✓	✓	✓
5	✓	✓	✓
6	✓	✓	✓
7	✓	0.57	✓
8	✓	✓	✓
9	✓	✓	0.92
10	✓	✓	✓
11	✓	✓	✓
12	✓	✓	✓
<i>Frontal controls</i>			
1	0.2	✓	✓
2	✓	✓	✓
3	✓	✓	✓
4	✓	✓	✓
5	✓	✓	0.57
6	✓	✓	✓
7	✓	✓	0.48

8	✓	✓	0.36
9	✓	✓	0.14
10	✓	✓	✓
<i>Healthy controls</i>			
1	✓	0.71	0.0016
2	✓	✓	0.9
3	✓	✓	✓
4	✓	✓	✓
5	✓	✓	✓
6	✓	✓	✓
7	✓	✓	✓
8	✓	✓	0.95
9	✓	✓	0.01
10	✓	✓	0.55
11	✓	✓	✓
12	✓	✓	0.09
13	✓	✓	✓
14	✓	✓	✓
15	0.24	✓	✓
16	✓	0.87	0.27
17	✓	✓	✓
18	✓	0.36	✓
19	✓	✓	0.26
20	✓	✓	✓

Note: Each participant participated in choices for all three categories.

Checkmark indicates subject fulfilled triangle inequalities for that category. Significant violations of linear ordering polytope are marked in bold.

Interestingly, two healthy controls significantly violated the random mixture model in the gambles domain ($p = 0.002$ and $p = 0.01$, respectively). One of these individuals followed Tversky's (1969) lexicographic semiorder heuristic exactly and the other followed this heuristic partially. Their results demonstrate the sensitivity of our test to detect individuals choosing on the basis of attribute-based heuristics that lack higher order transitive structure.

Individuals with VMF damage are not choosing randomly. One possible explanation for why individuals with VMF damage conform to probabilistic models of transitivity despite making a greater number of individual intransitive choices is that they are simply choosing randomly, as completely random choices fulfill the random mixture model 80% of the time in our experimental design (see methods). However, individuals with VMF damage are not simply choosing randomly. First, the probability that a group of random choosers the size of the VMF group ($N=13$) would all make choices consistent with the random mixture model in all three domains is extremely low, $p = 1.66e-04$. Second, we can evaluate directly the likelihood that an individual is choosing randomly by comparing their choice proportions ($N=10$ in each category) against those expected under the binomial distribution. For every single individual with VMF damage, and in all three domains, the likelihood that their choice proportions arose from completely random choice was extremely low (all $p < 1e-06$).

Individuals with VMF damage do not have systematically different preferences. A second possible explanation for why individuals with VMF damage conform to probabilistic models of transitivity despite making a greater number of individual intransitive choices is that they have systematically different preferences. For example, we might expect that a risk-neutral chooser would be more likely to make occasional intransitive choices in our gambles category than a strongly risk averse chooser. However, individuals in the VMF group did not make systematically different types of choices than individuals in the other groups. In a MANOVA on the choice proportions for each of the 10 binary choices the participants faced in each category, there were no significant differences between groups in the art category [Wilks' Lambda = 0.64, $F(18,64) = 0.9$, $p = 0.58$], the brand category [Wilks = 0.64, $F(18,64) = 0.90$, $p = 0.58$], or the gambles category [Wilks = 0.46, $F(18,64) = 1.67$, $p = 0.07$].

Individuals with VMF damage have noisier preferences. A third possible explanation for why individuals with VMF damage conform to probabilistic models of transitivity despite making a greater number of individual intransitive choices is that they are noisier choosers. That is, their choices reflect underlying transitive preference orderings, but they vacillate among preference orderings more than other choosers. To further test this possibility, we fit each individual's choices and RTs in Set B to a drift diffusion model (DDM), which assumed that choices and RTs were a probabilistic function of the rank distance in preference ordering between the two options. These fits revealed that individuals with VMF damage were noisier choosers. The only parameter of the DDM that was significantly different across groups was the noise parameter ε [$F(2,37) = 6.25$, $p = 0.005$]. Specifically, the VMF group (mean = 0.12, sd = 0.03) had

significantly higher ε than HC (mean = 0.09, sd = 0.04)[t(28) = 2.08, $p = 0.047$] and FC (mean = 0.07, sd = 0.02) [t(20) = 3.94, $p < 0.001$]. No other parameters differed between the three groups (**Figure 4**).

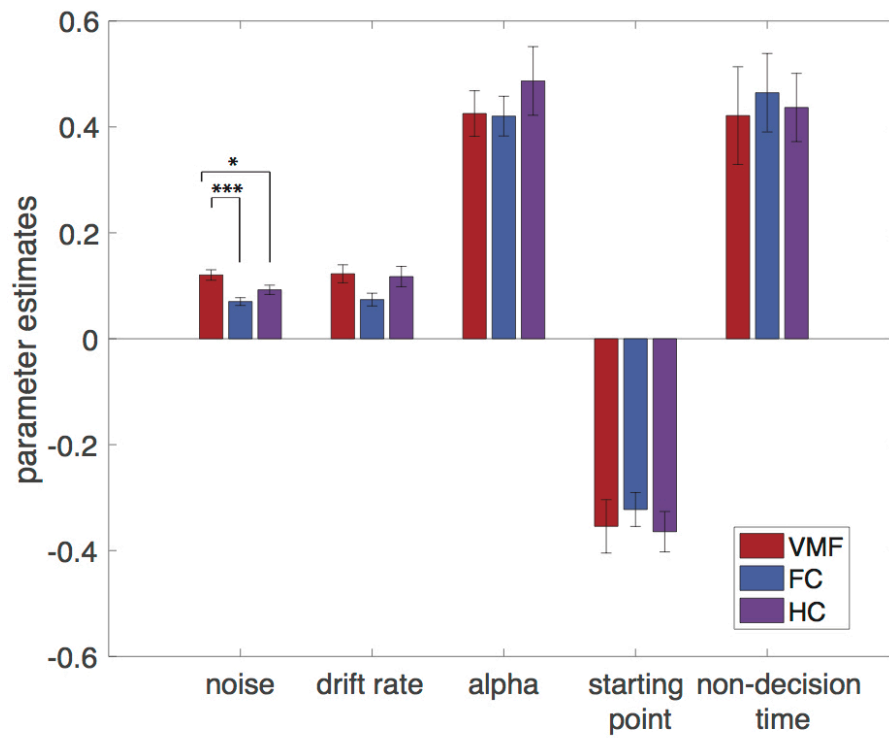


Figure 4-4. DDM parameter fits: noise, drift rate, initial starting point, non-decision time, and alpha (exponent on rank distance). Error bars are standard errors of the mean.

Individuals with VMF damage show a less pronounced effect of value on reaction times. RTs in individuals with VMF damage also showed a less pronounced effect of ranked value distance, consistent with the increased noise parameter observed in the DDM fits. We performed a mixed ANOVA on median RTs with value distance and group as factors. We found a significant main effect of value distance [$F(3, 111) = 28.63$, $p < 0.0001$], a significant main effect of group [$F(2,37) = 4.93$; $p = 0.01$], and a significant interaction between the two [$F(6,111) = 3.76$; $p = 0.002$].

The significant effect of value distance reflected the expected decrease in RTs as the distance in preference ordering rank gets larger. The average median RT for a rank difference of 1 (mean = 2800ms, sd = 1458) was significantly slower than a rank difference of 2 (mean = 2500ms, sd = 1211) [$Z = 4.86$, $p < 0.0001$], which in turn was slower than the rank difference of 3 (mean = 2300ms, sd = 1166) [$Z = 3.59$, $p < 0.001$], which in turn was slower than a rank difference of 4 (mean = 2180ms, sd = 1045) [$Z = 2.78$, $p = 0.005$].

The effect of group reflected longer RTs in the FC group. RTs in the FC group (mean = 3380ms, sd = 1596 ms) were significantly slower than in VMF group (mean = 1883ms, sd = 469ms) ($Z = 3.13$, $p = 0.002$), and a similar slowing relative to the HC group (mean = 2439ms, sd = 1116ms) exhibited a non-significant trend ($Z = 1.70$, $p = 0.09$). RTs in the VMF and HC groups were not significantly different ($Z = 1.16$, $p = 0.24$).

The interaction between value distance and group reflected a reduced effect of value distance on RTs in the VMF group. We took the Spearman correlation between RT and the difference in preference ordering rank as an index of the value distance effect. The VMF group (mean $\rho = -0.16$) exhibited a flatter value distance-RT relationship than the HC group (mean $\rho = -0.22$) [$t(28) = 2.20$; $p = 0.04$]. The value distance-RT relationship in the FC group (mean $\rho = -0.19$) was intermediate and not significantly different from the VMF ($p = 0.42$) or HC ($p = 0.61$) groups. As shown in **Figure 5**, these differences can be accounted for by the DDM fits described above.




Figure 4-5. Value distance effect on RT, by group. Dotted line are simulated RTs from DDM parameter fits. Error bars are standard errors of the mean.

Discussion

Individuals with damage to the ventromedial frontal lobes (VMF) have been shown previously to be more inconsistent in their choices (Camille et al., 2011; Fellows & Farah, 2007; Henri-Bhargava et al., 2012). These previous findings, however, are consistent with two possible patterns of behavior, with very different implications for the function of the VMF. One possibility is that individuals with VMF damage are fundamentally intransitive: that they reliably choose in an intransitive manner when given the same choice between the same options repeatedly. A second possibility is that individuals with VMF damage are more variable in their choices, yet still fundamentally transitive. Here we distinguished between these two possibilities by testing whether the choices of individuals with VMF damage satisfy probabilistic notions of transitivity, as the first possibility predicts they do not and second predicts they do. We overwhelmingly find evidence for the second possibility, as all individuals with VMF damage make choices in all domains that are consistent with probabilistic models of transitivity.

The first possibility, that individuals with VMF damage are fundamentally intransitive choosers, implies that the VMF is necessary for choices to be value-based, as transitivity is the key hallmark of a value-based choice (Samuelson, 1937; Von Neumann & Morgenstern, 1945). According to this view, individuals with VMF damage would only be able to choose in a non-value-based manner, for example, according to rules or heuristics. Our data, however, provide strong evidence against this possibility. This result is difficult to reconcile with the view that VMF is *the* critical substrate for value-based choice.

In contrast, we found strong evidence for the second possibility, that individuals with VMF damage are fundamentally transitive, that is their choices satisfy probabilistic models of transitivity, even though they make more intransitive choices according to deterministic notions of transitivity. Furthermore, we showed that this pattern was not due to individuals with VMF damage choosing in an entirely random manner, nor was it due to these individuals having preferences that were systematically different from those of the other groups. Rather, this pattern was due to individuals with VMF damage being noisier or more variable choosers. This is consistent with the suggestion of Henri-Bhargava et al. (2012), that “values are unstable, fluctuating from trial to trial in those with VMF damage.” We illustrated this by fitting a drift diffusion model (Ratcliff, 1978) to each individual’s choices. In this model, the VMF group had a significantly higher noise term, i.e., more variance around the decision variable, than healthy individuals or those with frontal damage outside the VMF. Importantly, the VMF group did not differ on the value of any other parameters. Reaction times in the VMF group were also similar to healthy controls, arguing against accounts of their behavior based on impulsivity (faster RTs) or indecision (slower RTs). Overall our modeling further strengthens the conclusion that the VMF serves to make preferences more stable, so that individuals would be less likely to select an option that is typically less preferred.

These results are easier to reconcile with a framework in which valuation and value-based choice are distributed processes, to which multiple regions of the brain contribute in some respect (Hunt & Hayden, 2017). This framework would predict that others regions can compensate for damage to the VMF, so that such damage does not fundamentally abolish the transitivity of preferences. The modest effect size in

deterministic tests, which is typically an increase of around 5% in the number of intransitive choices in the VMF group relative to control groups in our study and previous ones (Fellows & Farah, 2007; Henri-Bhargava et al., 2012), is also more consistent with this view. As making transitive choices that maximize value is incredibly important to the survival of an organism, it would make sense that value is a highly conserved process that is not abolished by damage to one part of the cortex. Future studies could more directly test hypotheses about compensation by examining activity in inconsistent individuals with fMRI, as it is also possible that regions that compensate are in the still intact parts of VMF rather than in other regions entirely.

Our results do not speak to how exactly the VMF supports choice stability. One possibility is that VMF contributes some part of the composition of subjective value. If subjective value is computed through the interaction of several brain regions, the loss of VMF may make this computation noisier and less reliable, akin to the greater noise we see in our DDM results. Alternatively, as a flattening of the value distance-RT relationship is consistent with greater indifference between options, the VMF could amplify or enhance the differences in value between different options (Henri-Bhargava et al., 2012). It is also possible that the VMF contributes a unique, specific component to valuation. For example, it has been suggested that the VMF contributes emotional content when making aesthetic judgments (Vaidya, Sefranek, & Fellows, 2017), and in other contexts that it contributes motivational salience that can distinguish close options from one another more clearly (Manohar & Husain, 2016; Pujara, Philippi, Motzkin, Baskaya, & Koenigs, 2016; Vaidya & Fellows, 2015b).

Another broad set of possibilities can be generated by considering the nature of the random mixture model that individuals with VMF damage satisfy. In this model, choosers are allowed to have different preference orderings in different contexts or at different points in time. It is possible, therefore, that VMF somehow contributes to the same preference ordering being repeated reliably. For example, individuals might use episodic memories of their previous choices (e.g., “I remember choosing A over B before”) to guide their decisions. Although we tried to reduce the influence of such memories, it is difficult to eliminate their influence entirely (Birnbbaum, 2011) and VMF has been implicated in episodic memory processes (Bertossi, Tesini, Cappelli, & Ciaramelli, 2016). Alternatively, VMF could support a representation of the context of the experiment that in turn activates a specific set of preferences, such as in a schematic network. Consistent with this idea, previous work has shown VMF involvement in schema formation (Schlichting & Preston, 2016; Spalding et al., 2018).

Finally, we extended previous studies that considered only deterministic notions of transitivity by identifying heterogeneity in these effects both across individuals and across domains. There was considerable heterogeneity within the VMF group, where some participants made as few intransitive choices as healthy controls, while other participants made significantly more intransitive choices. We did not find any systematic differences in lesion location or size that accounted for this heterogeneity. The lesions of the three individuals in the VMF group who made significantly more intransitive choices overall did not overlap much in their location, and the overlap areas were in the same location where other individuals had sustained lesions. The lesions of the three most inconsistent individuals in the VMF group did tend to extend more posteriorly towards

the basal forebrain and ventral striatum, though given the sample size in our study this potential explanation will need to be rigorously evaluated in future work with a larger number of subjects. Future studies could also test alternative explanations that we were unable to assess by using more advanced imaging to test whether damage to specific white matter tracts or disruptions in specific connectivity networks are linked to making more intransitive choices.

There was also considerable heterogeneity across domains, with the pattern of intransitive choices being most consistent with previous studies (i.e., showing a deficit selective to VMF damage) in the one the domain, art, that had been used in those studies. The greatest heterogeneity across domains, though, was in the frontal control group. This group looked similar to healthy controls in the art domain but made the most number of intransitive choices in the gamble domain. The frontal control group includes individuals with damage to the dorsomedial or dorsolateral prefrontal cortex, and both of these regions have been previously shown to be involved in decisions about risk (Christopoulos, Tobler, Bossaerts, Dolan, & Schultz, 2009; Hsu, Krajbich, Zhao, & Camerer, 2009). Previous studies have started to consider how the brain regions necessary for preference consistency may vary across domains (Fellows & Farah, 2007; Henri-Bhargava et al., 2012), and our results further highlight the need to examine a variety of domains in future work.

In conclusion, we found that individuals with VMF damage make choices that are noisier, but still fundamentally transitive. This result both characterizes how erratic choices manifest after damage to the VMF (Eslinger & Damasio, 1985; Harlow, 1868), as well as potentially explains why studies using similar decision-making paradigms in

individuals with VMF-damage can yield different results (Fellows, 2011). In addition, our findings further clarify and define the necessary role the VMF plays in value-based decision-making. Specifically, though each choice still reflects some subjective preference ordering after VMF damage, an intact VMF is necessary for preference orderings to remain stable and reliable across time and contexts.

Chapter 5 – General Discussion

Whether you are swiping between potential mates on a dating app, rating a pair of winter boots you just bought, or putting a bid on a house, there are many factors in your valuation of those options. However, we seem to be able to, most of the time, successfully and quickly resolve these factors into a preference. This dissertation sought to explore how the brain arrives at preferences across several different types of tasks. In Chapter 2, we assessed a model of decision-making that posited that choosing between options is akin to navigating a mental space made of option attributes. We found that this model did not have support in neural responses to a standard choice task, and that subjective value instead remains a better explanation for the data. In Chapter 3, we looked for the neural correlates of valuation in a different type of preference task, the matching task, which featured a longer response period. We found that the valuation network in the brain responded in much the same way as in choice tasks, and late into the response period. In Chapter 4, we tested whether the vmPFC is necessary for valuation by using a probabilistic model of transitivity. We found that vmPFC damage increased noise in the consistency of preferences, but did not make preferences fundamentally intransitive. This result shows that though vmPFC is critical for a component of valuation, multiple neural regions are likely needed to contribute to the valuation process.

This dissertation affirms subjective value as a model of neural representation in vmPFC during decision-making, that it is robust to the type of preference elicitation and represents value late into the decision process, and it is likely part of a distributed network of regions necessary for value construction. These studies confirm some pre-existing theories and challenges others. It supports the prevalent subjective value theory

of vmPFC function (Kable & Glimcher, 2009), but it limits the role of the vmPFC to being necessary for an aspect of value, rather than valuation as a whole. This dissertation is also the first to assess, and challenge, a conceptual navigation account of decision-making (Behrens et al., 2018; Bellmund et al., 2018).

Open Questions

This dissertation raises two important questions. In the following section, I will discuss these questions, and some future directions that could address them. The first question is, *what role does vmPFC play in subjective value?* Chapter 4 showed it is necessary for some aspect of value that reduces noise in valuation – so what could those aspects be? Relatedly, what models could account for the representations of value at the end of consideration period? The second question is, *what role does the vmPFC play in other cognitive tasks that could relate to its role in decision-making?* For instance, can our studies that support a view of vmPFC’s involvement in subjective value in decision-making, be reconciled with the studies that show that it is involved in spatial or conceptual navigation (Constantinescu et al., 2016; Doeller et al., 2010)? These two questions are intertwined, and feed into the main question of *what does the vmPFC do?*

First, what is the role of the vmPFC in valuation? The results of Chapter 4 support a view of value as a distributed network. Hunt and Hayden (2017) posit that multiple regions of the prefrontal cortex compute similar variables (e.g., the attributes of the options) simultaneously, and feedback into each other in a recurrent, hierarchical network. Each region of the prefrontal network (such as anterior cingulate cortex, orbitofrontal cortex, and dorsolateral cortex) computes similar calculations, but receive

unique inputs to contribute to the overall representation (for example, the OFC receives sensory inputs, while the vmPFC receives limbic inputs, etc.). Thus, no one region conducts one part of the valuation process; rather, all regions do so simultaneously. This theory is supported by single cell literature showing that virtually all regions of the frontal cortex compute similar decision variables (Hosokawa, Kennerley, Sloan, & Wallis, 2013; Hunt, Behrens, Hosokawa, Wallis, & Kennerley, 2015), and the dense reciprocal interconnections within regions of the frontal cortex (Felleman & Van Essen, 1991; Jbabdi, Sotiropoulos, Haber, Van Essen, & Behrens, 2015). Another prevalent distributed model of decision-making, the affordance competition hypothesis (Cisek, 2007), posits that goals and action plans for multiple options are prepared simultaneously and compete against each other. In this framework, competition between different action plans take place across the frontal-parietal dorsal regions, which is influenced by biasing factors from prefrontal cortex and basal ganglia. Both of these theories provide mechanisms by which the valuation representations are presented later into the decision (e.g., mutual inhibition between options, or competition between action plans) (Hunt et al., 2012; Jocham et al., 2012; Pastor-Bernier & Cisek, 2011). However, neither of them address scenarios where valuation is required, but there are no competing options, as in matching or rating tasks. Though there is a cognitive model from decision field theory which addresses matching tasks (Johnson & Busemeyer, 2005), there are to date no neural evidence or mechanisms proposed for this process. This is a gap in the literature that should be addressed.

Secondly, can vmPFC's role in subjective value be reconciled with its role in non-value-based cognitive domains, from memory, emotional and social cognition, and

valuation (Roy, Shohamy, & Wager, 2012)? One possibility is that different neural populations of vmPFC subserve these different functions, beyond the resolution that fMRI can provide. Conversely, however, there could be a broad functionality underlying these different domains. vmPFC and OFC is thought to represent the underlying structure of the task (Stalnaker, Cooch, & Schoenbaum, 2015; Wilson et al., 2014). For example, vmPFC could represent subjective value in decision-making tasks because subjective value is the most efficient and useful representation for preference tasks, and it could represent grid-like coding in navigation tasks because that form is the most efficient for representing a two-dimensional space for that purpose (Behrens et al., 2018). The reason knowing underlying task structure is useful is because when you are able to represent different states of the world (beyond just simple response-outcome associations), you can simulate future events and make predictions. Hippocampus, an area involved in prospective thinking, projects to the vmPFC, and lesions of the vmPFC causes deficits in both prospective memory and planning (Bertossi et al., 2016; Fellows & Farah, 2005; S. L. Peters, Fellows, & Sheldon, 2017). Additionally, both fMRI and lesion evidence have shown that vmPFC is involved in inferring unseen associations between objects, and for representing the underlying context of a task, and likely works with the hippocampus to do so (Zeithamova et al., 2012; Schuck et al., 2016; Spalding et al., 2018). However, though there have been computational models proposed of how these hidden attributes can be computed, as well as a role for vmPFC in this function (Gershman, 2018; Momennejad, Otto, Daw, & Norman, 2017), much work still needs to be done to uncover the processes behind them. I will discuss some of the potential future avenues of research now.

Future Directions

As Chapter 4 suggests that vmPFC is critical for some aspect of valuation, one main avenue of research would be to figure out precisely those aspects are. The perspective in Hunt and Hayden (2017) suggests that vmPFC would contribute components related to its limbic connections (e.g., memory or emotional aspects) to an overall value representation. There has been some work done that show that individuals with VMF damage use emotional or social information less in preference judgments of art and politicians (Vaidya et al., 2017; Xia, Stolle, Gidengil, & Fellows, 2015). These studies hint at the vmPFC contributing an emotional or contextual component to valuation, which could be further investigated with more systematic examination of the components of options. For instance, in studies of consumer products, art, or social stimuli, it would be a good idea to have the participants rate these items on various attributes of those items, and then subsequently ask about their preferences among those items (preferably in a way so that transitivity can be assessed, to check for inconsistencies). Having both the attribute rating and the subsequent preference assessment is important because it would show whether the differences in the VMF group come from a deficit in assessing certain attributes, or if they come from the failure to take attributes into account during preference assessments.

Secondly, another major research direction would be to test theories of vmPFC function that can account for its role in both value and non-value based tasks. More work should be done on both the nature of the vmPFC's involvement in tasks that require structure learning, as well as its necessity in these tasks. Specifically, the vmPFC and OFC have been shown to be critical in devaluation, as well as being involved in

reevaluation tasks in humans (Momennejad et al., 2017; Reber et al., 2017). These tasks require the participant to learn about direct associations between conditioned stimuli and reinforcers (e.g., different types of food, or monetary outcomes), learning that the reinforcer's value change in some way separately from conditioned stimulus (either devalued through satiation, or changed in monetary amounts), and then querying the value of the conditioned stimuli to see if those have changed along with the values of their associated reinforcers. Neural evidence has suggested that offline replay, that is, hippocampal activity during rest periods in the middle of the task, supported later reevaluation (Momennejad et al., 2017). Thus, it would be useful to look at the functional connectivity between hippocampus and vmPFC/OFC regions during periods of structure learning. Additionally, the OFC has been shown to encode the changes in the latent state of a task with probabilistic outcomes (Nassar, McGuire, Ritz, & Kable, 2018). This type of task is analogous to our anthropoid ancestors learning about the weather shifts which determine the probabilistic flowering of individual fruit trees. Successful performance is predicated on determining the true source of the outcomes, and when this source changes. Future directions for this research include finding out whether this latent state representation includes recognition of previously learnt states when the environment returns to them, as well as whether the OFC is necessary for this type of learning in lesion experiments.

Finally, as alluded to above, very little work has been done on modelling the mechanisms by which valuation occurs in the absence of competition between options, as in matching or rating tasks. One obvious direction is to find neural correlates for the decision field theory model proposed by Johnson and Busemeyer (2005), which posits

that sampled values are adjusted up or down until an indifference decision threshold is crossed. Similar accumulation-to-bound models for choice tasks (e.g., the drift diffusion model) has support in the firing in areas like the lateral intraparietal area in primates (Shadlen & Newsome, 1996), which may be a candidate for this modified choice model.

In summary, this dissertation refines our understanding of the vmPFC's role during decision-making. It shows that subjective value remains the best explanation of its role in determining preferences, whether for choice or for matching tasks, but also limits it to a role later in the decision-process. It furthermore clarifies that the vmPFC is necessary for a component of the valuation process that maintains the stability of preferences, but a lesion of the vmPFC does not abolish valuation. This work points to ways that our understanding of the vmPFC can be further advanced, not just for decision-making, but for its role in cognition more broadly as well.

BIBLIOGRAPHY

- Baron, J. (2008). Thinking and deciding 4th ed: New York: Cambridge University Press.
- 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.
- Becker, G. M., DeGroot, M. H., & Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral Science*, 9(3), 226-232.
- Behrens, T. E., Muller, T. H., Whittington, J. C., 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.
- Bellmund, J. L., Gärdenfors, P., Moser, E. I., & Doeller, C. F. (2018). Navigating cognition: Spatial codes for human thinking. *Science*, 362(6415), eaat6766.
- Bertossi, E., Tesini, C., Cappelli, A., & Ciaramelli, E. (2016). Ventromedial prefrontal damage causes a pervasive impairment of episodic memory and future thinking. *Neuropsychologia*. doi:<http://dx.doi.org/10.1016/j.neuropsychologia.2016.01.034>
- Birnbaum, M. H. (2011). Testing mixture models of transitive preference: Comment on Regenwetter, Dana, and Davis-Stober (2011).
- Camille, N., Griffiths, C. A., Vo, K., Fellows, L. K., & Kable, J. W. (2011). Ventromedial frontal lobe damage disrupts value maximization in humans. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 31(20), 7527-7532. doi:10.1523/JNEUROSCI.6527-10.2011

- Chib, V. S., Rangel, A., Shimojo, S., & O'Doherty, J. P. (2009). Evidence for a common representation of decision values for dissimilar goods in human ventromedial prefrontal cortex. *The Journal of Neuroscience*, 29(39), 12315-12320.
- 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.
- Cisek, P. (2007). Cortical mechanisms of action selection: the affordance competition hypothesis. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 362(1485), 1585-1599.
- Clithero, J. A., & Rangel, A. (2013). Informatic parcellation of the network involved in the computation of subjective value. *Social cognitive and affective neuroscience*, 9(9), 1289-1302.
- Constantinescu, A. O., O'Reilly, J. X., & Behrens, T. E. (2016). Organizing conceptual knowledge in humans with a gridlike code. *Science*, 352(6292), 1464-1468.
- Cooper, N., Kable, J. W., Kim, B. K., & Zauberman, G. (2013). Brain activity in valuation regions while thinking about the future predicts individual discount rates. *Journal of Neuroscience*, 33(32), 13150-13156.
- Crawford, J. R., & Howell, D. C. (1998). Comparing an individual's test score against norms derived from small samples. *The Clinical Neuropsychologist*, 12(4), 482-486.
- Devi, V. R., Guttes, E., & Guttes, S. (1968). Effects of ultraviolet light on mitosis in *Physarum polycephalum*. *Experimental cell research*, 50(3), 589-598.

- Doeller, C. F., Barry, C., & Burgess, N. (2010). Evidence for grid cells in a human memory network. *Nature*, 463(7281), 657.
- Eslinger, P. J., & Damasio, A. R. (1985). Severe disturbance of higher cognition after bilateral frontal lobe ablation. *Patient EVR*, 35(12), 1731-1731.
doi:10.1212/wnl.35.12.1731
- Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral cortex (New York, NY: 1991)*, 1(1), 1-47.
- Fellows, L. K. (2011). 16 The Neurology of Value. *Neurobiology of sensation and reward*, 351.
- Fellows, L. K., & Farah, M. J. (2005). Dissociable elements of human foresight: a role for the ventromedial frontal lobes in framing the future, but not in discounting future rewards. *Neuropsychologia*, 43(8), 1214-1221.
doi:10.1016/j.neuropsychologia.2004.07.018
- 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 (New York, N.Y. : 1991)*, 17(11), 2669-2674. doi:10.1093/cercor/bhl176
- Fellows, L. K., Stark, M., Berg, A., & Chatterjee, A. (2008). Patient registries in cognitive neuroscience research: Advantages, challenges, and practical advice. *Journal of cognitive neuroscience*, 20(6), 1107-1113.
- Fischer, G. W., & Hawkins, S. A. (1993). Strategy compatibility, scale compatibility, and the prominence effect. *Journal of Experimental Psychology: Human Perception and Performance*, 19(3), 580.

- Gershman, S. J. (2018). The successor representation: Its computational logic and neural substrates. *Journal of Neuroscience*, 38(33), 7193-7200.
- Haber, S. N. (2016). Corticostriatal circuitry. *Neuroscience in the 21st Century*, 1-21.
- Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052), 801-806.
- Hare, T. A., Camerer, C. F., & Rangel, A. (2009). Self-control in decision-making involves modulation of the vmPFC valuation system. *Science*, 324(5927), 646-648.
- Hare, T. A., O'Doherty, J., Camerer, C. F., Schultz, W., & Rangel, A. (2008). Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors. *Journal of Neuroscience*, 28(22), 5623-5630.
- Harlow, J. M. (1868). Recovery from the passage of an iron bar through the head. *Publications of the Massachusetts Medical Society*, 2(3), 274-281.
- Harris, A., Adolphs, R., Camerer, C., & Rangel, A. (2011). Dynamic construction of stimulus values in the ventromedial prefrontal cortex. *PLoS ONE*, 6(6), e21074.
- Harris, A., Clithero, J. A., & Hutcherson, C. A. (2018). Accounting for taste: A multi-attribute neurocomputational model explains the neural dynamics of choices for self and others. *Journal of Neuroscience*, 38(37), 7952-7968.
- Henri-Bhargava, A., Simioni, A., & Fellows, L. K. (2012). Ventromedial frontal lobe damage disrupts the accuracy, but not the speed, of value-based preference judgments. *Neuropsychologia*, 50(7), 1536-1542.
doi:10.1016/j.neuropsychologia.2012.03.006

- Hosokawa, T., Kennerley, S. W., Sloan, J., & Wallis, J. D. (2013). Single-neuron mechanisms underlying cost-benefit analysis in frontal cortex. *The Journal of Neuroscience*, 33(44), 17385-17397.
- Howard, J. D., Gottfried, J. A., Tobler, P. N., & Kahnt, T. (2015). Identity-specific coding of future rewards in the human orbitofrontal cortex. *Proceedings of the National Academy of Sciences*, 112(16), 5195-5200.
- Howard, J. D., & Kahnt, T. (2017). Identity-specific reward representations in orbitofrontal cortex are modulated by selective devaluation. *Journal of Neuroscience*, 3473-3416.
- Hsu, M., Krajbich, I., Zhao, C., & Camerer, C. F. (2009). Neural response to reward anticipation under risk is nonlinear in probabilities. *Journal of Neuroscience*, 29(7), 2231-2237.
- Hunt, L. T., Behrens, T. E., Hosokawa, T., Wallis, J. D., & Kennerley, S. W. (2015). Capturing the temporal evolution of choice across prefrontal cortex. *eLife*, 4, e11945.
- Hunt, L. T., & Hayden, B. Y. (2017). A distributed, hierarchical and recurrent framework for reward-based choice. *Nature Reviews Neuroscience*, 18(3), 172.
- Hunt, L. T., Kolling, N., Soltani, A., Woolrich, M. W., Rushworth, M. F., & Behrens, T. E. (2012). Mechanisms underlying cortical activity during value-guided choice. *Nature Neuroscience*, 15(3), 470.
- Hunt, L. T., Woolrich, M. W., Rushworth, M. F., & Behrens, T. E. (2013). Trial-type dependent frames of reference for value comparison. *PLoS computational biology*, 9(9), e1003225.

- Jacobs, J., Weidemann, C. T., Miller, J. F., Solway, A., Burke, J. F., Wei, X.-X., . . .
- Fried, I. (2013). Direct recordings of grid-like neuronal activity in human spatial navigation. *Nature Neuroscience*, *16*(9), 1188-1190.
- Jbabdi, S., Sotiropoulos, S. N., Haber, S. N., Van Essen, D. C., & Behrens, T. E. (2015). Measuring macroscopic brain connections in vivo. *Nature Neuroscience*, *18*(11), 1546.
- Jenkinson, M., Beckmann, C. F., Behrens, T. E., Woolrich, M. W., & Smith, S. M. (2012). Fsl. *NeuroImage*, *62*(2), 782-790.
- Jocham, G., Hunt, L. T., Near, J., & Behrens, T. E. (2012). A mechanism for value-guided choice based on the excitation-inhibition balance in prefrontal cortex. *Nature Neuroscience*, *15*(7), 960.
- Johnson, J. G., & Busemeyer, J. R. (2005). A dynamic, stochastic, computational model of preference reversal phenomena. *Psychological Review*, *112*(4), 841.
- Kable, J. W., Caulfield, M. K., Falcone, M., McConnell, M., Bernardo, L., Parthasarathi, T., . . . Lerman, C. (2017). No effect of commercial cognitive training on brain activity, choice behavior, or cognitive performance. *Journal of Neuroscience*, *37*(31), 7390-7402.
- Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience*, *10*(12), 1625-1633.
doi:10.1038/nn2007
- Kable, J. W., & Glimcher, P. W. (2009). The neurobiology of decision: consensus and controversy. *Neuron*, *63*(6), 733-745.

- Knutson, B., Rick, S., Wimmer, G. E., Prelec, D., & Loewenstein, G. (2007). Neural predictors of purchases. *Neuron*, 53(1), 147-156.
- Latty, T., & Beekman, M. (2010). Food quality and the risk of light exposure affect patch-choice decisions in the slime mold *Physarum polycephalum*. *Ecology*, 91(1), 22-27.
- 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.
doi:<https://doi.org/10.1016/j.conb.2012.06.001>
- Luce, R. (1959). Individual Choice Behaviour. A Theoretical Analysis: New York. Wiley.
- Manohar, S. G., & Husain, M. (2016). Human ventromedial prefrontal lesions alter incentivisation by reward. *Cortex: A Journal Devoted to the Study of the Nervous System and Behavior*, 76, 104-120.
doi:<http://dx.doi.org/10.1016/j.cortex.2016.01.005>
- McFadden, D. (1980). Econometric Models for Probabilistic Choice Among Products. *The Journal of Business*, 53(3), S13-S29.
- McGuire, J. T., & Kable, J. W. (2015). Medial prefrontal cortical activity reflects dynamic re-evaluation during voluntary persistence. *Nature Neuroscience*, 18(5), 760.
- Mellers, B. A., Ordoñez, L. D., & Birnbaum, M. H. (1992). A change-of-process theory for contextual effects and preference reversals in risky decision making. *Organizational Behavior and Human Decision Processes*, 52(3), 331-369.
- Momennejad, I., Otto, A. R., Daw, N. D., & Norman, K. A. (2017). Offline Replay Supports Planning: fMRI Evidence from Reward Revaluation. *bioRxiv*, 196758.

- Nassar, M. R., McGuire, J. T., Ritz, H., & Kable, J. (2018). Dissociable forms of uncertainty-driven representational change across the human brain. *bioRxiv*, 364638.
- Öngür, D., & Price, J. L. (2000). The Organization of Networks within the Orbital and Medial Prefrontal Cortex of Rats, Monkeys and Humans. *Cerebral Cortex*, 10(3), 206-219. doi:10.1093/cercor/10.3.206
- Padoa-Schioppa, C., & Assad, J. A. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature*, 441(7090), 223-226.
- Pastor-Bernier, A., & Cisek, P. (2011). Neural correlates of biased competition in premotor cortex. *Journal of Neuroscience*, 31(19), 7083-7088.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1992). Behavioral decision research: A constructive processing perspective. *Annual review of psychology*, 43(1), 87-131.
- Pegors, T. K., Kable, J. W., Chatterjee, A., & Epstein, R. A. (2015). Common and unique representations in pFC for face and place attractiveness. *Journal of cognitive neuroscience*.
- Peters, J., & Büchel, C. (2010). Neural representations of subjective reward value. *Behavioural brain research*, 213(2), 135-141.
- Peters, S. L., Fellows, L. K., & Sheldon, S. (2017). The ventromedial frontal lobe contributes to forming effective solutions to real-world problems. *Journal of cognitive neuroscience*, 29(6), 991-1001.
- Petrides, M., & Pandya, D. (1994). Comparative architectonic analysis of the human and the macaque frontal cortex. In F. Boller & J. Grafman (Eds.), *Handbook of neuropsychology* (Vol. 9). Amsterdam: Elsevier.

- 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.
- Pujara, M. S., Philippi, C. L., Motzkin, J. C., Baskaya, M. K., & Koenigs, M. (2016). Ventromedial prefrontal cortex damage is associated with decreased ventral striatum volume and response to reward. *Journal of Neuroscience*, 36(18), 5047-5054.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59.
- Reber, J., Feinstein, J. S., O'doherty, J. P., Liljeholm, M., Adolphs, R., & Tranel, D. (2017). Selective impairment of goal-directed decision-making following lesions to the human ventromedial prefrontal cortex. *Brain*, 140(6), 1743-1756.
- Regenwetter, M., Dana, J., & Davis-Stober, C. P. (2010). Testing transitivity of preferences on two-alternative forced choice data. *Frontiers in psychology*, 1, 148.
- Regenwetter, M., Dana, J., & Davis-Stober, C. P. (2011). Transitivity of preferences. *Psychological Review*, 118(1), 42-56. doi:10.1037/a0021150
- Regenwetter, M., Davis-Stober, C. P., Lim, S. H., Guo, Y., Popova, A., Zwilling, C., . . . Messner, W. (2014). QTest: Quantitative testing of theories of binary choice. *Decision*, 1(1), 2.
- Roy, M., Shohamy, D., & Wager, T. D. (2012). Ventromedial prefrontal-subcortical systems and the generation of affective meaning. *Trends in Cognitive Sciences*, 16(3), 147-156.

- Samuelson, P. A. (1937). A note on measurement of utility. *The review of economic studies*, 4(2), 155-161.
- Schlichting, M. L., & Preston, A. R. (2016). Hippocampal–medial prefrontal circuit supports memory updating during learning and post-encoding rest. *Neurobiology of Learning and Memory*, 134, 91-106.
- Schuck, N. W., Cai, M. B., Wilson, R. C., & Niv, Y. (2016). Human orbitofrontal cortex represents a cognitive map of state space. *Neuron*, 91(6), 1402-1412.
- Shadlen, M. N., & Newsome, W. T. (1996). Motion perception: seeing and deciding. *Proceedings of the National Academy of Sciences*, 93(2), 628-633.
- Smith, S. M., & Nichols, T. E. (2009). Threshold-free cluster enhancement: addressing problems of smoothing, threshold dependence and localisation in cluster inference. *NeuroImage*, 44(1), 83-98.
- Spalding, K. N., Jones, S. H., Duff, M. C., Tranel, D., & Warren, D. E. (2015). Investigating the Neural Correlates of Schemas: Ventromedial Prefrontal Cortex Is Necessary for Normal Schematic Influence on Memory. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 35(47), 15746-15751. doi:10.1523/JNEUROSCI.2767-15.2015
- Spalding, K. N., Schlichting, M. L., Zeithamova, D., Preston, A. R., Tranel, D., Duff, M. C., & Warren, D. E. (2018). Ventromedial prefrontal cortex is necessary for normal associative inference and memory integration. *Journal of Neuroscience*, 38(15), 3767-3775.
- Stalnaker, T. A., Cooch, N. K., & Schoenbaum, G. (2015). What the orbitofrontal cortex does not do. *Nature Neuroscience*, 18(5), 620.

- Strait, C. E., Blanchard, T. C., & Hayden, B. Y. (2014). Reward Value Comparison via Mutual Inhibition in Ventromedial Prefrontal Cortex. *Neuron*, 82(6), 1357-1366. doi:10.1016/j.neuron.2014.04.032
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, 55(4), 189.
- Tom, S. M., Fox, C. R., Trepel, C., & Poldrack, R. A. (2007). The neural basis of loss aversion in decision-making under risk. *Science*, 315(5811), 515-518.
- Tversky, A. (1969). Intransitivity of preferences. *Psychological Review*.
- Tversky, A., Sattath, S., & Slovic, P. (1988). Contingent weighting in judgment and choice. *Psychological Review*, 95(3), 371.
- Vaidya, A. R., & Fellows, L. K. (2015a). Testing necessary regional frontal contributions to value assessment and fixation-based updating. *Nature communications*, 6, 10120. doi:10.1038/ncomms10120
- Vaidya, A. R., & Fellows, L. K. (2015b). Ventromedial frontal cortex is critical for guiding attention to reward-predictive visual features in humans. *The Journal of Neuroscience*, 35(37), 12813-12823. doi:<http://dx.doi.org/10.1523/JNEUROSCI.1607-15.2015>
- Vaidya, A. R., Sefranek, M., & Fellows, L. K. (2017). Ventromedial Frontal Lobe Damage Alters how Specific Attributes are Weighed in Subjective Valuation. *Cerebral Cortex*, 1-11.
- Von Neumann, J., & Morgenstern, O. (1945). *Theory of games and economic behavior*: Princeton University Press Princeton, NJ.
- Wilson, R. C., Takahashi, Y. K., Schoenbaum, G., & Niv, Y. (2014). Orbitofrontal cortex as a cognitive map of task space. *Neuron*, 81(2), 267-279.

- Xia, C., Stolle, D., Gidengil, E., & Fellows, L. K. (2015). Lateral orbitofrontal cortex links social impressions to political choices. *The Journal of Neuroscience*, 35(22), 8507-8514. doi:<http://dx.doi.org/10.1523/JNEUROSCI.0526-15.2015>
- Yamada, H., Louie, K., Tymula, A., & Glimcher, P. W. (2018). Free choice shapes normalized value signals in medial orbitofrontal cortex. *Nature communications*, 9(1), 162.
- Zauberman, G., Kim, B. K., Malkoc, S. A., & Bettman, J. R. (2009). Discounting time and time discounting: Subjective time perception and intertemporal preferences. *Journal of Marketing Research*, 46(4), 543-556.
- Zeithamova, D., Dominick, A. L., & Preston, A. R. (2012). Hippocampal and ventral medial prefrontal activation during retrieval-mediated learning supports novel inference. *Neuron*, 75(1), 168-179.