Distinct representations of economic variables across regions of the frontal cortex



Antara Majumdar

Kellogg College

University of Oxford

A thesis submitted for the degree of

Doctorate of Philosophy

September 2024

This thesis is dedicated to my family

"*We live in the world when we love it.*"

Rabindranath Tagore

Acknowledgements

During my time at DPAG, I reflect on my scientific journey here and all the wonderful people who have inspired, taught and supported me throughout these years. Without this support, this thesis would not have been possible. I would like to thank everyone from the bottom of my heart.

Armin, thank you so much for your support and guidance since my first day in your lab. I have a photo of my first day in the lab and in it I am sitting near the Neuropixels rig and I remember feeling a quiet excitement forming that day. Since that day, you have so patiently taught me not only how to do electrophysiology experiments, design experiments and understand our dataset. All these moments always piqued my scientific curiosities and I only wish to bring forth with me all that you have taught me. I hope to mentor others with the patience and dedication that you have shown me. Since the moment I had my first conversation with you about science, I have not looked back and it was the easiest decision to join your lab and work on such exciting projects alongside you. You are an exceptional scientist and I am honored everyday to work with you on a scientific story and I hope to continue in this field and continue to grow my knowledge. You inspire me everyday and I am so thankful for your mentorship and kindness.

To Simon, thank you so much for being such an amazing supervisor throughout the last few years. I remember exiting all our meetings and feeling excited about the scientific questions we discussed. Your insight has always been incredibly valuable to me and has helped me grow so much as a scientist. Additionally, thank you for always being there to listen and being so willing to share your experiences throughout your scientific journey. I am incredibly thankful for your time and our conversations.

Matthias, thank you so much for all your endless support throughout the past few years. I remember the first time I ever had a meeting with you, I was struggling to write code and finding it hard to formulate different analyses questions. Throughout the years you have patiently guided me through not only code writing but provided incredible insights on the different analyses questions that I am interested in. I come out of each meeting with you feeling so excited to work on different ideas and most importantly to finish what I start. You have always reminded me to prioritize and I feel so lucky to have such awesome collaborations with you! I have learned so much from you and I

always feel so valued during our weekly discussions. You are an inspiration to me and I hope I can carry forward all that you have taught me throughout the years. Thank you for your mentorship!

My scientific journey in this lab would not have been complete without my collaborations. Caitlin, I don't have the words to express how much your friendship and our teamwork has meant to me in the past few years. We have worked together through deadlines and weeks of mice having streaks of learning or confusion with the task. Regardless of the moments being good or bad or even in the middle, you have always been a constant in my life. I am so grateful for all our coffee walks and conversations. I am so incredibly excited for all your medical school colleagues to see just how brilliant, compassionate and wonderful you are. Many hugs to you my dear friend!

Throughout my time in the Lak and Butt labs, I have been incredibly blessed to be among such wonderful people who I consider not only my colleagues but also friends. To Jess, Michael, Chiara, and Peter - I am so thankful for your friendship and I truly cherish all our adventures together. I am always so at home when I am around you all and I truly feel that I love being in science because of the people I meet. I always remember just how supportive you all have been throughout the years. I will cherish all our moments together and I hope to continue to build more such moments with you all. I will never forget our photowalks, gym classes, coffee walks, and our conversations. Thanks to all the present and past members of the Lak and Butt labs. You have all been the most supportive colleagues and I will fondly remember our lab retreat, pub trips, hiking in the Lake District, and many other moments.

I am also thankful for all my friendships - both the ones I was able to form in Oxford, the Cajal summer school, and friendships that have stayed with me throughout my undergrad, highschool and middle school years. Some of my favorite moments in the past years have been when I got to meet my friends in different cities and or they have been able to visit me here in Oxford. In particular, I would like to thank Andrada, Lauren, Daniel, Nanu, Maria, Alexandra, Erica, Rachel, Nisa and Priyanka for all the moments shared and I hope to have more adventures with you all. Thank you to everyone for your friendship. I am so incredibly lucky to have you all in my life. In addition, the Emotional Tasting of Hops (ETOH) society and the Super Fun Time Book Club (SFTBC) both have been incredibly fun to be a part of. I will always cherish the pub trips and sitting down to discuss books with you all!

I would not have been able to pursue my scientific career without the support of my family. In particular, I am incredibly blessed to have such supportive parents who I look up to and who gave me all the love, resources and time that allowed me to unravel my passions. To my mom and dad,

this thesis is for you both. Both of you have taught me to work hard and to create my space so that I can find the quiet for my creative pursuits. I will never forget how you helped me with my elementary school science projects, or how you drove me to the lab early in the morning so that I could do my experiments before the school day started. To my brother, Annesha didi and my little nephew - I am so thankful for all your support. Ma, Baba, Babi dada, Annesha didi and Manav - I love you.

Grigore, it is so difficult to sum up in a few sentences how much it means to me to have you in my life. I am so lucky to have a best friend in you - someone who I can share anything with and who supports me even during those moments when I did not feel complete. Thank you for reminding me that I am always enough. Life with you is so wonderful and I feel like the luckiest person in the world. In the time spent with you, you have amazed me with your brilliance and dedication to whatever you work on. You remind me to keep pushing through and to celebrate both the small and big achievements in life. I am incredibly excited to keep growing with you and building our careers and lives. Most importantly I am excited to have someone to share everyday with. I love you.

Abstract

We often make economic decisions in our daily lives which require us to evaluate and integrate different variables such as reward magnitude and risk to inform our choices. Previous studies have shown representations of such economic variables in distinct regions of the frontal cortex. However, these studies usually focused on select subpopulations of frontal cortical neurons. Therefore it remains unclear how representations of economic variables are structured within the frontal cortex and whether diverse frontal regions and their projections pathways encode distinct variables. To address this question we combined behavior and large-scale high-density electrophysiology to investigate how fine grained neural signals contribute to economic decisions. To this end, we devised a visual economic decision-making task in head-fixed mice, akin to those used in previous non-human primate studies. In each trial, mice chose between two simultaneously presented abstract visual stimuli that differed in their magnitude of associated water reward, and reward probability, resulting in options with different expected values and economic risks. Mice learned the task, consistently selecting stimuli with higher expected values. Population decoding of economic variables in the responses of frontal cortical neurons showed distinct gradients for different economic variables: expected value coding was stronger in dorsal compared to ventral frontal regions, while risk coding was stronger in medial compared to lateral regions. Lastly, we test the causal role of such representations with no conclusive evidence that inactivating frontal regions during our task disrupts economic choices. Our results shed light on the structure of the neural code underlying economic decision-making, demonstrating that frontal populations encode basic variables of economic decision-making.

Summary of contributions

I declare that this thesis is my own and that attribution has been given where information was derived from other sources. Aside from the contributions outlined below, I have performed the experiments, analyzed the data and constructed the figures. While the writing in this thesis is all my own, this project as a whole has relied on key contributions from others. Throughout the text, the pronoun 'we' is used indicating myself and the team below.

Caitlin Ashcroft: electrophysiological data acquisition, behavioral data acquisition and data analysis. **Matthias Fritsche:** kernel regression, population decoding of task variables. **Leah Mistry:** Data analysis of face videos. **Anita Maria-Rominto, Jake Swann:** spike sorting of neuronal data, and behavioral training. **Tolu Duckworth-Essilfie:** Histology and probe track localization. **Olena Didenko, Rathi Puliyadi:** behavioral training PV-Cre mice and histology. **Peter Zatka-Haas:** computer code for neuronal data pre-processing. **Simon Butt:** supervision throughout. **Armin Lak:** conceptualization, funding and supervision.

Contents

1 Introduction 1

1.1 Variables in economic choice 3

1.2 Neural structure of economic choice 8

1.3 Thesis outline 18

2 Economic decision-making in head-fixed mice 20

2.1 Introduction 20

2.2 Methods 24

2.3 Results 34

2.4 Discussion 52

2.5 Acknowledgements 54

3 Representations of economic variables across the frontal cortex 56

3.1 Introduction 56

3.2 Methods 62

3.3 Results 75

3.4 Discussion 95

3.4 Acknowledgements 104

4 Optogenetic inactivation during economic decision-making 106

4.1 Introduction 106

4.2 Methods 112

4.3 Results 116

4.4 Discussion 123

4.5 Acknowledgements 127

5 General Discussion 129

5.1 Future directions 139

References 147

List of figures

2.1 The 2AFC economic decision-making task setup
2.2 Cumulative probability of mice surpassing a given performance level as a function of trial number
2.3 Mice make economic decisions between abstract visual stimuli
2.4 Mice display faster response times for stimuli with higher EV
2.5 Choice accuracy split by stimulus side for single stimulus trials
2.6 Mice make choices based on economic value not visual properties of stimuli
2.7 Risk sensitivity evolves across sessions
2.8 Trial-by-trial updating happens in the space of actions
2.9 Facemap analysis of lick data
2.10 Correlating behavioural performance with response times and risk preference
3.1 Example frontal neuronal clusters' parameters on Phy GUI
3.2 Frontal neuronal responses reflect the economic value of stimuli
3.3 Residuals from the reduced rank kernel regression method show stimulus-locked responses
3.4 Kernel regression partially removes reaction-time related signals
3.5 Strength of representations are maintained in simultaneously recorded sessions
3.6 Representations of economic variables across frontal cortex exhibit distinct spatial gradients
3.7 Stronger contralateral value representations observed in frontal cortical neurons
3.8 Distinct representations of chosen and unchosen value in the frontal cortex
3.9 Decoding accuracies for mice with two stimuli conditions mirrors those in the general population
4.1 Histological confirmation of virus expression and fiber placement
4.2 Effects of unilateral inactivation of OFC and medial PFC areas
4.3 Effects of unilateral inactivation of OFC and medial PFC areas on response times
4.4 Effects of unilateral inactivation of OFC and medial PFC on individual single stimulus conditions

List of tables

1 Transitivity of choice preference and testing for weak or strong transitivity

1

Introduction

Contents

1.1 Variables in economic choice
1.2 Neural structure of economic choices
1.3 Thesis outline

Economic decision-making requires evaluating and integrating different economic variables such as reward magnitude and risk to inform our choices. Past studies have demonstrated representations of such economic decision variables across various regions of the frontal cortex. Early lesion studies in humans, non-human primates and other model organisms demonstrated that damage to the frontal cortex impairs decision-making (Bechara et al., 1994; Camille et al., 2011; Milner et al., 1963; Buckley et al., 2009; Rudebeck et al., 2006). These lesion studies were followed by a slew of electrophysiological studies that revealed that neurons across the frontal cortex can represent several economic variables across behaviorally relevant time-scales (Padoa-Schioppa et al., 2006; Kennerley et al., 2009a; O'Neill and Shultz, 2010; Tremblay and Schultz, 1999). Despite this progress, past studies, either measured neural signals with coarse spatial and temporal resolution (e.g. using

functional resonance imaging (fMRI)) or from small neural populations using electrophysiology. Therefore, it remains unclear how frontal neuronal signals represent economic decision variables.

These representations can appear at different levels in the frontal cortex. One level of representation is that different regions of the frontal cortex encode different economic variables. Indeed, previous studies have found evidence that neurons in specific frontal cortical regions represent different features of decision-making (Kuwabara et al., 2020; Hocker et al., 2021; Cai & Padoa-Schioppa, 2012; McCoy & Platt, 2005). However, such studies have typically focused on a small number of neurons in isolated regions, making it challenging to collectively compare representations across the entirety of the frontal cortex. Thereby it remains unclear how these representations of economic decision variables are structured within the frontal cortex.

To address this question, we trained mice in a visually-guided task to probe economic decision-making in head-fixed mice, akin to those used previously in studies of non-human primates, with different abstract visual stimuli predicting different reward magnitudes and probabilities. We show that mice make economic decisions between abstract visual stimuli, are sensitive to the expected value of the presented options and exhibit considerable variability in risk preference. We used high-density large-scale electrophysiological recordings (Jun et al., 2017) to measure neural signals across many frontal regions during the task. Investigating these representations of economic variables across the frontal cortex shows that population activity across the frontal cortex distinguishes between stimuli corresponding to different economic values. We next investigated whether neural responses reflected individual risk preferences and found that trials with the same expected value but different risk were decodable from neural responses of risk-sensitive mice, but were not decodable in risk-neutral mice. Furthermore, decoding of such economic variables from the neuronal population revealed distinct spatial gradients for the representation of expected value (EV) and risk. Having shown these functional spatial gradients for coding of EV and risk, we next asked whether the representations in these areas are causally linked to

economic decisions. Our inactivation results are thus far inconclusive and provide no concrete evidence that inactivating frontal regions disrupts economic evaluation.

Together, our results demonstrate that the encoding of economic variables assumes distinct spatial gradients of neural representation across the frontal cortex and provides a platform for investigating the neural basis of economic decision-making at a large scale with high spatial and temporal resolution. Furthermore, it provides a basis for further investigation of the causal role of these neural representations in guiding economic choices.

1.1 Variables in economic choice

1.1.1 The development of neuroeconomics research

In our everyday lives, we often make decisions grounded in economic principles. When we encounter different offers in our environment, we make a choice among the available options. Among these options, typically there is no 'correct' choice as a lot of these economic choices are based on subjective preferences, which means that they are based on personal opinions and motivations rather than the objective computed value of the offer at hand. Economists and psychologists alike have been fascinated by these economic decisions for centuries and have established some core principles underlying such decisions. Work studying how humans make such economic choices has uncovered a series of phenomena that may impact economic behaviors observed in humans. Some of the phenomena described are hyperbolic discounting (a cognitive bias whereby people opt to choosing smaller more immediate rewards rather than wait for larger rewards that appear later), loss aversion (when a potential loss is perceived by individuals as more psychologically damaging than an equivalent gain), framing effects (where an individuals' decision behavior is influenced by the way that the information is presented to them), and lastly the Allais paradox (when individuals prefer a safe and certain outcome as opposed to a probabilistic and risky option) (Ainslie et al., 1992; Camerer et al., 2003; Kagel and Roth, 1999; Kahneman and Tversky,

2013). Many of these psychological phenomena have a basis in what economists uncovered from the 1930s, a decade during which the mathematical structure of consumer choice and behavior was investigated. Economists such as Samuelson, Arrow and Debreu developed models of economic choice that considered the impact of allocated resources and normative choices. One of the first frameworks introduced was the Weak Axiom of Revealed Preferences (WARP) developed by Samuelson in the 1930s (Samuelson et al., 1938). In this framework, an individual given a choice between an apple and an orange, choosing an apple, reveals a preference for an apple. This later evolved to the GARP ("generalized" axiom of revealed preference) in which an individual that has a preference for an apple over an orange, and a preference for a peach over an apple, has an "indirect" preference for a peach over an orange. The revealed preference theory was later extended and evolved as choices between safe and risky options were considered. These probabilistic versus guaranteed reward choice scenarios were characterized by von Neumann and Morgensterns's expected utility (EU) theory (von Neumann and Morgentern, 1944). The expected utility theory posits that when an individual is faced with a risky choice, they will make a choice that maximizes the utility (benefits minus the costs of making that choice). These theories will make the foundation for prospect theory - a descriptive model of decision-making under risk (Kahneman & Tversky, 1979) which deviates from the expected utility theory. Prospect theory integrates loss aversion and reference dependence (where outcomes such as wins and losses are viewed relative to an internal reference point) allowing it to model choice behavior accurately. Overall, these theories have sought to characterize how humans make decisions. Fundamental to these theories is that individuals integrate different economic variables to determine their choice. In the following sections, I will describe two variables integral to economic decisions: value and risk.

1.1.2 Value

One particular form of value that is quintessential in economic decision-making is expected value. Expected value is calculated by multiplying the magnitude of reward with its probability. One view of economic behavior is that in a goal-directed context, values are learned by experiencing the

outcome associated with an action and evaluating the reward magnitude (Skinner et al. 1953; 1981). In this associative view, values are purely learned through experience. Since this view postulates that values are learned through experience, scientists tested this by training animals within the context of an economic decision-making task. In a study by Padoa-Schioppa et al., 2005, they examined how a capuchin, when faced with making choices among two options, maps their choices and has a series of mental processes that enable effective decision-making among both pairs of familiar and novel stimuli. They found that when capuchins encountered unfamiliar stimuli with different values, they choose effectively from the very first exposure - suggesting that value is not based solely on a stimulus-action exposure effect, but values are also assigned to offers independently of experience. Once these values are assigned and learned, and an animal makes a choice between two offers such as in a 2 alternative forced choice task (2AFC), and they make decisions among different pairs and combinations of stimuli-reward pairs, one can map the preference relations. In other words, if a person prefers an 'apple' to an 'orange' and an 'orange' to a 'banana', then it can be assumed that the person will express a preference for the 'apple' in relation to the 'banana'. This is referred to as the axiom of transitivity and it is a cornerstone of economic decision-making paradigms with a numerical value assigned to the expected value variable (Luce et al., 1995; 1997; Regenwetter et al., 2010). Furthermore, theories of transitivity allow us to create a framework of preference-based decision-making. Despite the monotonic nature of value parameters in studies, preference-based decisions are often not perfect and we have to take into account the need to maximize on rewards acquired for survival. In a study by Polanía et al., 2019, the authors use a model to describe subjective value using a combination of inference of the values in the environment and maximizing information transmission by allocating resources to value-based decisions based on prior information about the distribution of value information. Finally, in order to study decision-making and the influence of value, we need a reliable way to collect a behavioral readout of how values influence decision-making. In several studies using non-human primates, this was developed by using different abstract visual stimuli, each associated with a different reward magnitude and or probability (Padoa-Schioppa & Assad, 2006; Tremblay & Schultz, 1999). In addition, a behavior

paradigm inspired by these non-human primate studies was created for mice in a study by Kuwabara et al., 2020. This paper used odor stimuli instead of visual stimuli to represent a particular offer type and the odor concentration to represent the reward magnitude. In both animal models of economic choice, animals were able to make choices scaled to the value associated with the stimuli, suggesting that they successfully compute and compare values of offers in their environment. Therefore all these studies highlight the importance of studying the representation of value.

1.1.3 Risk

Risky decisions are an ubiquitous part of our daily lives. Risk is defined as the "mathematical variance of binary symmetric probability distributions of reward magnitudes" (O' Neill and Shultz et al., 2010; Rothschild & Stiglitz, 1970), bringing to the forefront the variable of probability being heavily involved in economic decision-making. Risk sensitivity is often tested in animal models with behavioral paradigms that involve an animal choosing between two options: one with a guaranteed reward and the second option exhibiting a variance in probability or riskiness (O' Neill & Schultz, 2010). Animals can be categorized into three different risk attitudes from this assessment: animals that are insensitive to risk should be indifferent to either the guaranteed and the risky option, risk-averse animals show a preference for the safe option while risk-sensitive animals show a preference for the option with an associated risk. Having a descriptive study of animal models and risk attitudes is useful and prompts the question of what factors influence risk attitudes of different animal species. Some factors identified in animals that forage are: energetic status of the animal species (the balance between energy intake and energy usage to acquire resources) and it is found that animals with a positive energy status exhibit risk averse behavior while those with a negative energy budget are risk-seeking (Brito e Abreu & Kacelnik, 1999; Kacelnik & Bateson, 1996). Other factors influencing risk attitudes is whether the risk is introduced by varying the amount of food, typically promoting risk-aversion, or by varying the temporal delay to acquiring the offer, promoting risk-seeking tendencies. A third factor influencing risk attitudes involves the number of options presented to an animal in their environment (Hurly & Oseen, 1999; Bateson et al., 2002).

In addition to these cognitive variables, statistical variables impacting risk preference have been explored. Meta-analysis studies across animal species reveals that the coefficient of variation (CV), a measure of the amount of risk per unit of return, can predict how much risk an animal is willing to choose (Weber et al., 2004).

These cognitive and statistical factors have been shown to impact risk attitudes in studies that have set up formal economic decision-making tasks. For instance, the factor influencing risk attitude has been explored in a visual gambling task designed for rhesus macaques by Hayden and Platt, 2007, in which they asked whether risk sensitivity and temporal discounting (preference for immediate rewards relative to rewards presented with a longer delay) are closely tied to one another. Macaque monkeys in this study exhibited lower risk sensitivities as the delay to receiving a reward increased even while other factors in the task stayed constant. In addition to temporal discounting influencing the directionality of risk attitude, scientists have asked how individual wealth levels impact our economic choices in picking a risk. In a study by Yamada et al., 2013, the authors asked how the amount of food and water that an individual possesses impacts their risk attitudes, with a focus on rhesus macaques. They trained macaques in a gambling task where the reward was in a fluid form and assessed risk attitude over the course of a session as a function of blood osmolality (a measure of amount of water the animal possesses). When examining the risk attitude as a function of individual wealth defined by blood osmolality, they found that the more resource “rich” the individual, the more risk seeking they were - a finding that is incompatible with previous findings in the studies conducted by Bateson on foraging animals and their risk attitudes. Despite the differences with previous studies, this study allowed for a similarity to be found among human and monkey studies - both species display similar risk attitudes and economic rationality - suggesting that monkeys can serve as a model for human behavior. Beyond risk attitudes being determined dynamically as a function of individual wealth, the role of previous experiences, decisions and outcomes and their impact on risk attitudes was explored. Behavioral variability in rats choosing between a risky and a safe cue revealed the role of these sequential biases on current trial behavior (Constantinople et al.,

2019a). Specifically, the current trial probability of picking the risky cue was impacted by collective previous outcomes - for example, rats were more willing to take a risk following a risky win. Beyond the previous outcome shaping of current trial risk behavior, on a whole risk attitudes in rats show similarities to humans. Rats exhibit reference dependence (there is an internal reference point for each individual and outcomes are viewed as gains or losses compared to this) and probability distortion (where lower probabilities are weighed more strongly) coupled with trial history effects, pointing out how learning impacts economic choices (Constantinople et al., 2019). Studies from monkeys revealed a similar probability distortion effect similar to that observed in humans and rats, where monkeys use probability information to make decisions to obtain reward and specifically overweight low probability rewards and underweight high probability rewards (Stauffer et al., 2014). Overall, these studies on risk attitudes suggest that risk attitudes are influenced by different cognitive factors allowing us to further study brain regions that are involved in the dynamic representation of choices under risk. Additionally, it points out the importance of studying how risk attitudes differ across species to avoid ascribing the same phenomena observed in humans to other species. Finally, these behavioral studies on risk reveal the importance of studying the dynamic nature of risk attitude, as opposed to a static representation of this neurally and behaviorally.

1.2 Neural structure of economic choices

A fundamental question of neuroeconomics research is how our brains structure representations of distinct economic variables to enable economic decision-making. Past studies have revealed that neural activity in the frontal cortex encodes different economic variables, and that these signals are necessary for economic decision-making. The evidence supporting the crucial role of the frontal cortex comes from many studies, spanning from lesion studies, single-cell electrophysiological recordings, and non-invasive neuroimaging studies in humans. Such studies have collectively implicated the frontal cortex and other regions in decision making. In the next few subsections, we will briefly outline this evidence. Before we go into details about these studies, we would first like to give a brief overview of homologies to aid comparisons across species. Following this anatomical

overview, I then detail lesion studies across species followed by neuroimaging and electrophysiology literature focusing on reward magnitude and risk coding.

1.2.1 Anatomical differences and homologies among species

The mouse prefrontal cortical regions differ and have anatomical homologies to prefrontal cortical (PFC) regions defined in non-human primates and in humans. In mice as in rats, the prefrontal cortex is a portion of the agranular frontal lobe. Cytoarchitectonic staining has allowed for clear distinctions between areas in the mouse prefrontal areas (Van de Werd et al., 2010). Historically, the PFC was referred to as the areas in the brain targeted by projections from the mediodorsal thalamus (MD) (Rose and Woolsey, 1948). Even though the MD projects to portions of the mouse PFC, this draws into question the usefulness of relying on a definition of PFC that is dependent on projections from MD. It also draws into question: what are the anatomical similarities and dissimilarities between PFC regions across species? To answer this question, researchers studied the parcellation of the mouse prefrontal cortex as well as the connectivity patterns to and from these areas. In particular, researchers sought to characterize the corticocortical connectivity patterns among prefrontal cortical neurons, revealing a highly interconnected set of prefrontal cortical regions - what is referred to as a "prefrontal module in the mouse" (Harris et al., 2019). The mouse prefrontal cortical areas were identified to be prelimbic area (PL), infralimbic area (ILA), dorsal and ventral anterior cingulate areas (ACAd and ACAv), frontal pole (FRP), and orbital areas (lateral OFC, medial OFC and ventrolateral OFC) (Harris et al., 2019; Le Merre et al., 2019). In addition subregions of the secondary motor cortex (MOs) are included in the frontal cortical region as portions of the MOs are strongly connected to other frontal regions.

1.2.2 Lesion studies

1.2.2.1 Human lesion studies

One of the earliest observations connecting brain regions to decision-making involved lesion or neural damage studies. In 1847, when a railroad worker, Mr Phineas Gage, suffered an injury to his frontal cortex during an accident in which a 13-pound tamping iron propelled through his head - medical experts became interested in his case in the years to follow. In fact, people discovered reports written by Mr. Gage's physician detailing his patient's life post-accident and often note the changes in personality and attribute this to sustained damage of the frontal lobe (Harlow et al., 1848; Damasio et al., 1994; Neylan et al., 1999). These behavioral changes have been described as a deficit in social decision-making - often described in other patients with damage to the frontal cortex, specifically to the ventromedial areas (Macmillan et al., 2000). Physicians and scientists alike have sought to understand the effect on cognitive abilities when there is damage to frontal cortical areas. In 1963, Milner noted that patients with frontal lobe damage performed poorly on the Wisconsin Card Sorting test, a neuropsychological test which required them to classify cards and adapt to changing rules about which classification is correct or not. This is among one of the first cognitive tests employed to understand the extent of damage to the frontal cortex in patients - but more importantly it started to create a link between function in these areas and the ability to do cognitive tasks. In a clinical study by Bechara et al. (1994) asked whether they could formally quantify and investigate the behavioral deficits associated with damage to the frontal areas of the brain. To do so, they had patients with damage to the ventromedial portion of the prefrontal cortex taking part in a task which simulates real life decision-making scenarios with manipulation of uncertainty and outcomes. They found that patients with frontal damage compared to controls show reduced performance on this task and also are insensitive to the future consequences of their actions, suggesting an important role for the ventromedial regions in carrying out decisions that may maximize long-term payoffs. This insensitivity to the future consequences of one's actions is coupled with the increased likelihood of making choices that violate the Generalized Axiom of

Revealed Preference (Samuleson et al., 1938), a cornerstone in utility theory, as observed in another study with patients with ventromedial frontal lobe damage showing choices that are neither consistent nor transitive (Camille et al., 2011). This impairment is in contrast to findings in healthy control participants that show that using the same experimental paradigm that humans choose by utilizing value maximization (Harbaugh et al., 2001). Although human lesion studies pointed towards the role of frontal regions in cognitive tasks and social decision-making, the exact role of frontal regions was not found due to the nature of lesions in patients not always being localized at one area of interest. Therefore, more targeted lesion studies were required in order to have a more in depth understanding how frontal regions are involved in cognitive tasks and also how damage to these areas can disrupt crucial functions.

1.2.2.2 Non-human primate lesion studies

In order to have more targeted investigation into the role of different frontal regions in different aspects of decision-making, one possible way is to use targeted lesion studies in non-human primates. Previous macaque lesion studies have established the important role ventral PFC and OFC plays in rule-based decision-making (Walton and Behrens, 2010; Bussey et al., 2001; Murray et al., 2000; Petrides et al., 1982; Gaffan et al., 1989). These studies showed that macaques, like humans, can learn rule-based decision-making tasks, allowing for similar lesion and behavioral studies being used in primates for a closer examination of the neural correlates of rule-guided decisions. One particular study in non-human primates reported the effects of lesioning areas such as the macaque OFC, principal sulcus (PS), superior dorsolateral and ventrolateral PFC, and ACA (Buckley et al., 2009). In this study, they tested macaques ability to do a cognitive test previously used on human patients - the Wisconsin Card Sorting Test (WCST) - and found distinct functions for each area in this neuropsychological test. In the OFC lesion experiment, they noticed an impairment in updating the values of rewards while in the ACA lesion group. These impairments in the WCST analog test for macaques with lesions suggested that the OFC plays a crucial role in the representation of values of different rules in this task. The ACA impairments suggested that ACA

plays a complementary role in rule-based behavior, with a specific role in referring to the outcome in the rule selection and thereby playing an important role in referring to abstract rules to allow for optimal decision-making. This is one example study among a few others that attempts to bridge the gap between human neuropsychological clinical studies and rule-based behavior investigations in non-human primates.

1.2.2.3 Rodent lesion studies

Studies described above have shown that different lesion studies in humans and non-human primates allowed insight into the roles of frontal regions in rule-based decisions. Other factors such as potential costs and rewards may influence such decisions (Walton et al., 2002). In rodent studies, this extended role of frontal regions in contributing to the cost-benefit aspects of decision-making has been shown. In a particular study in rats, lesions in ACA impact the amount of effort rats put into attaining rewards and OFC lesions impacted the amount of time rats were willing to wait before attaining rewards (Rudebeck et al., 2006). Specifically, when OFC was lesioned they observed that rats made more impulsive decisions, and when ACA was lesioned, rats were biased in their effort-based decision-making meaning that they were prone to choosing smaller rewards with less effort over large rewards that required more effort. Therefore, rodent lesion studies point out separable roles for two frontal regions in decision-making. Lesion studies in non-human primates and rodents first introduced specific roles of specific regions in the frontal cortex in economic decision-making framework, spurring future studies to investigate this with a higher spatiotemporal profile.

1.2.3 Electrophysiological and neuroimaging studies

Since lesion studies may often impact multiple areas, it becomes difficult to disentangle how different variables are encoded by distinct regions. Several neuroimaging and electrophysiological studies have been conducted and have focused on characterizing the different neural representations

in different areas. In the following sections, we will detail such studies and highlight the neural representations of risk and value from such studies.

1.2.2.3 Representations of risk

As humans and animal models display a gradient of risk attitudes that are modulated by different cognitive variables, one question that arises is what are the neural correlates of these risk attitudes? A monkey study that shows that monkeys prefer the risky target in a visual gambling task revealed activations in a posterior cingulate cortex (CGp) when the monkey chose the risky option and this activation scaled with the degree of risk (McCoy & Platt, 2005). Other regions have been implicated in the representations of subjective value under risk such as the striatum, medial prefrontal cortex, and anterior cingulate (Levy et al., 2009; Christopoulos et al., 2009; Passecker et al., 2019; Maisson et al., 2021). These studies finding the neural correlates for decisions under risk are crucial for two reasons: first they provide validation for the parameters used in economic decision-making under risk studies and second they suggest that methods such as fMRI and blood oxygen-dependent (BOLD) are capable of finding the correlates of these decision variables and prompt even more detailed examination of the neural circuits underlying risk behavior. Imaging studies spearheaded further investigation into the neural underpinnings of risk attitude and electrophysiological studies found signatures of risk in a specific region of the frontal cortex: orbitofrontal cortex. In a study by O' Neill and Shultz, 2010, they found that a population of OFC neurons in monkeys exhibited risk responses that convaried with the variance in the reward distributions. Orbitofrontal cortical neurons were not only implicated coding of variance of risk but other factors as well. Interestingly, a risk study in rats revealed encoding of abstract biases (an increased propensity to choose the risky option following risky wins) in neural activity acquired from a lateral orbitofrontal cortex (lOFC) (Constantinople et al., 2019b). Another study by Passecker et al., 2019 in which they train rats to perform a two-arm bandit task and make choices between a small-certain and a large-gamble reward, uncovers neurons in the prelimbic cortex that predict future decisions during this gambling task.

Overall these studies collectively show the sheer spread of information about risk across brain regions as well as dynamic representations of risk preferences.

1.2.2.4 Representations of value

One of the regions studied quite closely in neuroeconomic studies is the orbitofrontal cortex (OFC). The OFC has been shown to encode different economic variables. In a primate study where subjects had to make choices between different types of juices, the activity of OFC is tied to encoding of variables such as the decision input (the value of the offer presented) and also the value of the chosen offer (Padoa-Schioppa & Assad, 2006). Importantly, these neural representations of economic variables were coded independently of visual properties of cues and the motor responses, suggesting that OFC has a crucial role in the abstract representation of economic variables. These neural representations being independent of visuospatial properties and motor factors separates OFC from other regions where value representations may be more closely tied to such factors. In addition to these representations, neurons in the OFC have been implicated in value comparisons (Ballesta et al., 2022) expanding the role of OFC in not only the computation of but also the comparison of values. Despite the plethora of evidence demonstrating the importance of OFC in value-based decision-making, it is less known whether OFC is involved in selecting the choice in order to make a decision. Furthermore, the OFC is known to connect weakly to motor areas (Carmichael & Price, 1996) and has a relatively weaker encoding of the choice compared to other variables (Kennerley et al., 2009; Cai & Padoa-Schioppa, 2014).

Given the overall weak representation of the choice response in OFC neurons, scientists examined how the value representations in the OFC may influence neural activity in regions such as ACA, a region that receives projections from OFC and is implicated in encoding the choices. In a study that simultaneously records from OFC and ACA neurons in nonhuman primates during an economic decision-making task, they show that the value representations in OFC interact with ACA - when OFC neurons are signaling a higher value cue, this in turn affects the ramping of ACA neurons

(Balewski et al., 2023). This suggests that not only are OFC neurons implicated in representing value independent of motor responses, but these representations appear to drive choice-related activity in other prefrontal regions. Studies that simultaneously record from two regions that represent values such as Balewski et al., (2023) and Sul et al., (2010) are instrumental to disentangling how value signals in different frontal regions interact with one another and enable complementary processes that are crucial to economic decisions. In particular, Sul et al., (2010) uses a two-armed bandit task for rats and records from OFC and medial prefrontal cortex. They found that both areas convey information about the past choices made and the outcomes received. The OFC distinguishes itself from the medial prefrontal cortex by having neural signals for the value of the chosen offer and having a reward prediction signal.

So far, we have discussed past OFC studies that highlight this area as crucial for signaling the values of offers in a current trial context. However, these value representations on a current trial are not made in isolation, but rather are a combination of past experiences (past successes and failures) and a collective reflection of previous outcomes (Sul et al., 2010; Shohamy and Daw, 2015). A question that arises from this is whether these value representations observed in OFC are modulated by previous experiences. Indeed, a study by Constantinople et al., (2019) where they trained rats in an economic task to choose between a safe and gamble rewards, revealed that OFC activity is reflective of past experiences (i.e. the history of outcomes). In particular, the authors found that OFC is important for modulating the propensity to choose a gamble following a previously rewarded gamble (i.e. a risky win). OFC activity is also reflective of adaptive coding, demonstrated by Kennerley et al., (2011), where OFC neurons encoded chosen value relative to the history of reward. These studies expand the role of OFC in not only encoding values of offers but integration of reward history.

While neuroeconomics studies have predominantly focused on value representations in OFC, there are several lines of evidence that show that ACA also plays important roles in coding value and risks

associated with offers (Cai & Padoa-Schioppa, 2012; Balewski et al., 2023; Christopoulos et al., 2009; Kennerley et al., 2011; Luk & Wallis, 2012). In a study by Amiez et al., (2006) where they vary the quantity and the probability of the juice offered, they show that individual ACA neurons are modulated by expected value. Given this evidence of ACA neurons encoding value signals, the next question was whether the ACA encodes pre-decision (offer value) or post-decision variables (chosen value) and whether these representations are closely tied to motor responses. Indeed, a study by Cai and Padoa-Schioppa (2012) showed that individual neurons in the ACA were modulated by the direction of the choice, the identity of the offer and the subjective value of the chosen offer but interestingly not by the value of individual offers. This study suggests a role of ACA in value representations that is seemingly distinct from representations in the OFC, while also showing that ACA may play an important role in transforming value-based choices into motor actions. Furthermore, this beautifully highlights how different regions may work in synchrony even if they have overlapping and non-overlapping representations of economic variables. ACA activity is not only correlated with value but also risk (Christopoulos et al., 2009) as demonstrated in a study with human participants choosing between a risky and a safe offer. Overall, ACA population activity is shown to be modulated by both economic variables and other variables that help the subject guide their choices. Interestingly, the same chosen value modulated ACA neurons are also encoding reward prediction errors (Kennerley et al., 2011), suggesting a role of ACA in learning, thereby promoting behavioral flexibility. This multiplexed role of ACA is not only shown in the context of encoding of reward prediction errors, but also in relation to other economic variables. For example, ACA neurons in a study by Kennerley et al., 2009) have multiplexed representations of three decision variables: probability, the potential payoff and the cost in terms of the effort needed to make a decision. Overall, ACA seems to play a crucial role in the neural circuits underlying decision-making. There are also observations in past studies that find encoding of decision variables in non-frontal areas (Bao et al., 2023; Tobler et al., 2007). In particular such studies have highlighted the role of the frontal orienting field in encoding of economic value during a risky decision-making task developed for rats (Bao et al., 2023).

1.2.4 Transformation of economic choice to an action output

Several studies in the past have implied that economic choices are made independent of the visuomotor properties of each task meaning that decisions are made in the space of goods. Given how decisions are independent of the spatial and visual properties of cues, it is crucial to understand how the value of the chosen offer is transformed into an action. The neural mechanisms underlying this transformation from chosen good to action have been shown to involve different frontal regions. Neural signals in lateral OFC regions - orbitofrontal cortex, ventrolateral prefrontal cortex (LPFCv) and dorsolateral prefrontal cortex (LPFCd) - encode the choice outcome in goods space and after a delay they represent the spatial location of the cues (Cai & Padoa-Schioppa, 2014). This neural signal transformation of chosen good to chosen value was observed in another value-based decision paradigm, this time a gambling task with recordings from the supplementary eye field (SEF) of monkeys (Chen & Stuphorn, 2015). In this study, they noted that the neural activity in the SEF first signaled the chosen good (chosen gamble option in this study) before signaling the action (a saccade) to obtain the reward. The ~100 ms delay between the chosen good and the action neural representation indicates that during economic choice, first the economic good is chosen in a 'goods space' prior to the action plan unfolding. This is important for understanding the several steps and complexity in economic decisions and suggests that even when both the value and the spatial information are provided simultaneously - that these two steps serve specialized and different functions. Regions in the brain such as the medial prefrontal cortex, striatum and OFC not only encode information about the value, but also encode spatiomotor variables in order complete the action needed to attain a reward (Matsumoto et al., 2003; Samejima et al., 2015; Feierstein et al., 2006) as groups of neurons that show responses to the direction of the choice made. Therefore, neurons do not just encode the value of an offer but also represent the properties of goals - thus making neurons that partake in both important for goal-directed behavior and the transformation of information from the goods space to an action.

1.3 Thesis outline

Economic decision-making requires the computation and comparison of information about available offers, such as their expected value and risk, to guide our actions. Previous studies have identified that neurons in the frontal cortex encode this information about the offers. However, despite extensive research on this topic, several questions remain. For example, do mice learn the economic value associated with abstract visual stimuli? Once they are experts on an economic task, how is the encoding of this information structured across regions of the frontal cortex? Are there distinct structures of representations for risk and EV across these regions? Do these frontal representations of economic variables play causal roles in decision making? In this thesis, I investigate these questions using a novel behavioral economic decision making task in mice, together with large-scale electrophysiological recordings, and optogenetic manipulation experiments. Overall, the goal of this thesis is to uncover the structure of neuronal representations of decisions guided by economic variables.

In **Chapter 2**, I will focus on the economic decision-making task designed for head-fixed mice. The main motivating question is to investigate how mice make economic choices. Specifically, I investigate how mice make decisions based on economic value and risk associated with different visual stimuli. Mice learn to perform the economic task successfully. Their choice accuracy is scaled with the expected value of the stimuli, whereas their risk preferences are best described as a gradient of risk sensitivities, ranging from risk-preferring to risk-averse mice.

It has been previously proposed that neurons in various frontal areas such as the OFC and ACA encode different economic decision variables. However, it is unclear how neuronal representations of economic variables are structured across the frontal cortex. To shed light on this question, in **Chapter 3**, I investigate the structure of neuronal representations of different economic decision variables. Specifically, I investigated how two economic variables - EV and risk - are structured across

the frontal cortex. Overall, I find that frontal cortical representations of EV and risk assume orthogonal spatial gradients to one another.

Building on evidence from Chapter 3, in **Chapter 4**, I investigate whether the representations in specific frontal regions play a causal role in economic choices. Specifically, I sought to test whether inactivation of either OFC or medial PFC impacts choice behavior in the economic task and whether inactivation of either area disrupts value-based computations. While I find preliminary evidence that inactivation of frontal areas impact choice behavior, the results pertaining to value-based computations remain inconclusive and will necessitate larger sample sizes.

Finally, in **Chapter 5**, I will summarize my findings outlined in this thesis and discuss the limitations and future directions both for this project and also in a broader context.

2

Economic decision-making in head-fixed mice

Contents

2.1 Introduction . . . . .
2.2 Methods . . . . .
2.3 Results . . . . .
2.4 Discussion . . . . .
2.5 Acknowledgments . . . . .

2.1 Introduction

We often make economic decisions in our everyday lives, usually with uncertain outcomes. Imagine you are looking to invest some money and are choosing one among two options. One of the options could be investing in a stock that offers a high return. The caveat to this high return offer is that there is a significantly high risk associated with this stock. Alternatively, you could look into investing in a stock that offers a smaller return of money associated with a promise of a safer and more certain outcome. This kind of economic decision-making is something we are faced with not only in our investments but also other contexts in our everyday lives. Human and animal behavioral studies posit that an individual chooses by assigning different "implicit" economic values to

different options, allowing them to compare these values, make a choice among them, and ultimately execute an action.

Making economic decisions such as investing in stocks requires us to evaluate and integrate different economic variables such as reward magnitude and risk to inform our choices. One of the key questions that interests neuroscientists is: how are economic variables represented in the brain? In order to study this, past research has focused on creating behavioral paradigms to study how individuals make economic decisions, assign values to the options that are available and make a decision among these options. Initial studies in neuroeconomics stem from lesion or neurological damage studies in humans that show impairments in decision-making. One of the impairments observed is that individuals with damage to frontal regions have an increased probability of making choices that violate the generalized axiom of revealed preference, an economic theory that ensures that preferences among offers remain constant (Camille et al., 2011). Such an impairment was characterized in a study with patients with damage to the ventromedial frontal lobe, an area in the frontal cortex that plays a role in encoding the subjective value of offers. Past studies have overall helped create experimental economic paradigms in humans that allow a careful and quantitative assessment of violations of utility theory (Kahneman and Tversky, 2013) and also how subjects make economic decisions (Levy et al., 2010), with a focus on variables crucial for shaping economic decisions.

Studies in nonhuman primates have further expanded the framework for studying economic choices, using tasks that involve choosing between juice or food rewards varying in magnitude or probability (Tremblay and Schultz 1999; Padoa-Schioppa et al., 2006a). Key economic variables of economic risk and value were manipulated in a study by O'Neill and Shultz (2010) permitting two main observations: (a) nonhuman primates displayed a graded preference for the risky option as opposed to the safe option across individuals and (b) they showed an increased preference for stimuli associated with higher value. Overall studies in primates where the monkeys have to choose

between different food types have shown that monkeys can make economic choices and such tasks integrate several components of the decision-making process allowing future studies that further scrutinize neural activity during different behavioral events and allowing for furthering our understanding of economic choice (Padoa-Schioppa & Assad, 2006; Padoa-Schioppa & Assad, 2008). Primate studies of economic decision-making prompted further investigations into the neural representations of different aspects of behavior. This is readily investigated in animal models where genetic tools are readily available and manipulations of projections and specific cell types are more easily done. Thereby rodent models of economic choices have become an integral part of maximizing our knowledge of economic decision making and the neural underpinnings of different variables in decision making.

Rodent models of economic choice developed in the last twenty years have developed comparable tasks to humans and nonhuman primates and have characterized a species-specific view of economic choice. Studies that have developed a mouse model of economic choice have integrated stimulus-reward associations and show that mice make choices with patterns that closely resemble those in primates and their responses are scaled with reward magnitude (Kuwabara et al., 2020). In particular, this study utilizes olfactory stimuli to denote different magnitudes of reward. Other studies developed a risky choice task for rats and showed that rats can make cue-based decisions between a probabilistic and a safe stimulus, and this was further characterized using model-based observations to examine different risk preferences in rats (Bao & Zhu, 2023). Rodent studies using a behavioral paradigm studying risk sensitivity found that rats display probability distortion, loss aversion and reference dependence (Constantinople et al., 2019a). Apart from characterizing heterogeneous risk attitudes across rats, studies such as this one bridge the gap between human and rodent studies looking at risk-sensitivities and value-based decision-making. The plethora of cross-species development of economic choice paradigms lay a strong foundation for the use of animal models in studying the different steps of this rather complex decision-making process and allowing for observations that can further highlight the similarities and differences in evaluation of

economic variables such as risk between species. These studies motivate our investigation of head-fixed visually-guided economic decision-making behavior.

Previous rodent economic studies have sought to bridge the gap to human and non-human primate studies. However, one big caveat is that unlike human and non-human primate studies, most rodent studies associate economic variables either with fixed choices, or with scaled physical stimulus attributes such as intensity or salience of a stimulus, making it hard to dissociate the representation of economic values from choice or stimulus representations. Our task additionally utilizes distinct visual stimuli, used previously in economic tasks designed for monkeys (Padoa-Schioppa & Assad, 2006; O' Neill & Shultz, 2010), allowing for interpretation of the economic value rather than choices guided by perceptual intensity or salience of a stimulus. Our stimuli differ from other stimuli used in rodent studies mainly when we consider the mapping of stimulus features to the economic variables. In particular, the visual stimuli we use in our task have low level features which do not directly map onto economic variables. In contrast, previous studies such as Constantinople et al., 2019 have continuous changes in low level features of stimuli which correspond to continuous changes in economic variables. In particular, the Constantinople et al., 2019 study used the frequency of sound to signal the value of reward (higher frequency, higher reward) and the number of light flashes conveyed the probability of reward (higher number of flashes, higher probability). Overall, our study utilizes stimuli that are different from stimuli used in other rodent studies with the main defining point being the mapping of stimulus features to the economic variables.

We designed a visually-guided head-fixed economic task to characterize how mice make economic choices by evaluating choices that vary in two dimensions: reward magnitude and risk. Our study integrated different levels of both aforementioned dimensions allowing us to characterize how mice may make choices scaled to the magnitude and how they treat a risky option compared to a safe option. Overall, our task takes inspiration from previous work in non-human primates and

ultimately seeks to create a standardized task for evaluating economic choices in mice. Once we show that this task design works, we can further implement it in the context of investigating how neural representations across brain areas enable economic decision-making guided by expected value (EV) and risk.

2.2 Methods

2.2.1 Animals

The dataset comprises data from 34 male mice aged between 10-24 weeks. Wildtype (C57/BL6J) mice from Charles River were used. The data for the 'single stimulus' task design was collected from 34 mice. Of these mice, 14 mice were additionally trained in the 'two stimulus' task design. Furthermore, we conducted a control experiment using 5 mice. Mice were housed in a facility that maintained a 12 hour light/dark cycle, a room ambient temperature of 20-24℃ and a controlled humidity level of around 40%. All experiments were conducted under the UK Animals Scientific Procedures Act (1986) under specific personal and project licenses for this project.

2.2.2 Surgeries

Mice were implanted with custom-made steel head plates to allow for head fixation during behavioral experiments. These surgeries were carried out in accordance with specific personal and project licenses for this project. A headplate implantation surgery was conducted in order to implant the custom-made steel headplate onto the skull. At the time of the headplate surgery, mice were generally 8-9 weeks of age. Mice were induced into an anesthetized state using 5% isoflurane (1-2% in O_2) and then positioned onto a stereotaxic frame prior to intraperitoneal injection of 5 mg/kg meloxicam and 0.1 mg/kg buprenorphine. Following these injections, mice were maintained on 1.5% isoflurane for the duration of the surgical procedure. Throughout the surgery, mice were kept on a heating pad to ensure proper maintenance of body temperature.

In the beginning of the surgery the hair over the skull was shaved and the muscles covering the central part of the skull were surgically removed. Any connective tissue was removed using forceps. Thorough cleaning of the skull surface was done by washing with sterile saline. Repeated saline washes ensured that no tissue or hair remained on the surface of the skull. We used lambda and bregma dorsal-ventral (DV) axis measurements to ensure the skull was level and recorded distance between these landmarks to mark locations for future surgeries which involve craniotomies over frontal regions. The custom-made headplate was then attached to the skull surface using dental cement (Super Bond C&B; Sun Medical). It was attached specifically to the bone posterior to the bregma. The portion anterior to bregma was covered using a layer of clear cement. Following the clear cement drying, we re-marked locations for bregma, lambda along with locations over both hemispheres for targeting frontal regions. The mouse was left to recover in a post-recovery cage placed over a heating pad. The mouse was provided with soft food, water, and jelly mixed with oral Metacam. The mouse was monitored until locomotion and grooming behavior were normal and the mouse was fully recovered.

Following surgery, mice recovered for 5 days prior to initiating water restriction. Once water restriction was started, mice were handled regularly for 5-7 consecutive days for 20 minutes each day. Following these 5-7 days of handling, mice were acclimated to head-fixation prior to the start of training. The head-fixation procedure involves starting with 10 minutes of head-fixation on the first day, followed by increasing the amount of time of head-fixation by adding 5 minutes to each day. From the first day of head-fixation mice were trained on the visual economic decision-making task starting at the initial shaping stage (described in further detail below).

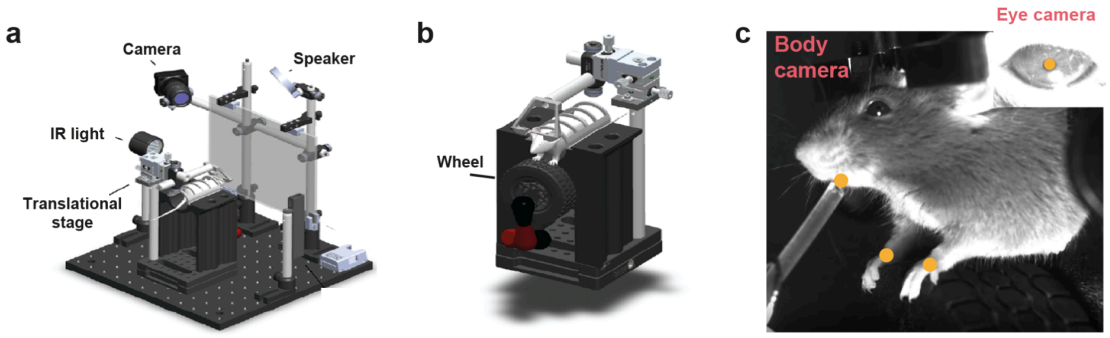


Fig 2.1 | The 2AFC economic decision-making task setup

a, A basic behavioral rig setup with different components such as an LCD computer screen, a 3D printed mouse holder, and infrared light source to allow for illumination on the face and eye for capturing videos. **b,** A demonstration of a head-fixed mouse in this behavioral rig setup. The mouse is head fixed using the headplate attached to the surface of the skull using dental cement. The headplate is slotted into a notch on the head plate holder, which is itself adjoined to a rod that spans along the setup of the mouse holder. The mouse's forepaws are resting on the Lego wheel below. Once a mouse is head fixed, a 3d printed mouse cover is used to cover the mouse in the mouse holder setup. **c,** An example frame from a face and eye video (top right) from a recording session. In the face video frame, the orange dots are marking key points for tracking mouth, right and left paw motion. In the eye video the center dot is tracking the motion of the pupil. Overall this setup allows the mouse to make decisions based on visual stimuli presented on the screen, and allows us to track the choices the mice make as the wheel is attached to a wheel encoder that tracks the movement of the wheel. Additionally, the face and eye camera allows us to capture videos that we can align to different behavioral events. This figure was adapted from the International Brain Laboratory et al., 2021 manuscript.

2.2.3 Materials and apparatus

Mice were head-fixed and trained on a standardized behavioral rig which has a LCD computer screen (9.7 inch diagonal), a 3-D printed plastic mouse holder, a head-fixation bar clamp to hold the mouse in a position and allow their forepaws to rest on a Lego wheel (Kubler 05.2400.1122.0360) (**Fig. 2.1 a,b**; Burgess et al., 2017; International Brain Laboratory et al., 2021). The angular movement of the Lego wheel along with the wheel velocity and direction was measured using a rotary encoder attached to the center, and a detailed breakdown of the hardware can be found here (https://www.ucl.ac.uk/cortexlab/tools/wheel). The LCD computer screen was positioned centrally at a distance of 11 cm from the mouse's eyes and was refreshed at 60 Hz. Water reward was delivered through a silicon plastic tube, positioned near the mouse's mouth, and controlled by a solenoid pinch valve. Water calibration was done every 3 months on rigs in order to ensure that the

range of water delivery allowed for reliable delivery of different water volumes. Pupil size, face and body movements were monitored using two cameras (right eye camera: FLIR, MVL50M23 - 50 mm EFL, f/2.8, for 2/3" C-Mount Format Camera, with Lock; face and body camera: FLIR, lens: MVL16M23 - 16 mm EFL, f/1.4, for 2/3" C-Mount Formal Cameras, with Lock). Face and eye videos were collected during the neuronal recording sessions). The face and eye of the mouse were illuminated using an infrared light source (BW 48 LED Illuminator Light CCTV IR Infrared Lamp). Experiments were run using custom-made Rigbox software written in MATLAB 2020 (Bhagat et al., 2020).

2.2.4 Visual economic decision-making task

The economic decision-making task was developed by myself and colleagues in the Lak Lab. This task was based on an established visually-guided behavior task (Burgess et al., 2017, International Brain Laboratory, 2020). We adapted the International Brain Laboratory (IBL) task in order to examine how mice make economic decisions guided by different abstract visual stimuli, each corresponding to different reward magnitudes and reward probabilities. Prior to the presentation of the stimulus, there was a variable time window between 0.2 and 0.5 seconds during which no movement of the wheel was permitted. The wheel was not physically blocked during this period but mice had to stop moving. After this pre-stimulus wait period, mice were presented with a single visual stimulus, appearing on either the right or left side of the screen (±35° azimuth, 0° altitude). After stimulus onset, a short delay of 0.2 - 0.3 seconds was introduced in order to keep the stimulus in its original position and to not allow any premature choices. Following this delay, mice were able to move the wheel under their forepaws to move the stimuli on the screen. Mice learned to move the stimulus to the center of the screen in order to obtain a water reward. The reward delivery occurred 0.1 seconds after the time when the choice was completed. If the mouse incorrectly moved the stimulus 37 degrees to the edge of the screen then they do not receive a water reward. The inter-trial interval was sampled randomly from a uniform distribution of 1 and 3 seconds.

Visual stimuli presented were abstract to control for any impact of visual properties on behavioral and neuronal readouts in this study (**Fig. 2.3c**). There were three abstract visual stimuli presented, each corresponding to a different combination of reward magnitude and probability: a large, guaranteed reward condition (3 μL, 100% of correct trials rewarded, EV = 3 μL), a small, guaranteed reward condition (1.5 μL, 100% of correct trials rewarded, EV = 1.5 μL), and a gamble condition (3 ul reward, 50% of correct trials rewarded, EV = 1.5 μL). These three visual stimuli are referred to as 'Large', 'Small' and 'Gamble' going forward. Each of the abstract visual stimuli are randomly paired with trial conditions across mice. Stimuli were presented in pseudo randomized order in each session and were shown randomly with the same probability during each session.

In addition to *single-stimulus* trials described above, a subset of mice experienced a second trial type termed *two-stimulus* trial. In the two stimuli trials, mice were presented with two different abstract visual stimuli simultaneously. The two simultaneously presented stimuli differed either in their reward magnitudes or probabilities . In the two-stimuli trials, mice were required to choose between the two stimuli, by moving one of the stimuli to the screen center and consequently obtained the water reward associated with the chosen stimulus. The two-stimuli task was designed to further demonstrate that mice make choices that are sensitive to value or probability differences among pairs of stimulus conditions. Additionally, these experiments provided a readout of the risk sensitivity of each mouse in the task, by calculating the choice preference between two stimuli with the same EV, but different levels of risk.

Mice started behavioral training 5 days post surgery allowing for adequate time for recovery. Initially, mice are acclimatized and handled for 4-5 days prior to training on the visual economic decision-making task. During each behavioral session, we stopped the task when we observed that the mouse lost motivation, indicated by slower reaction times, and or the mouse had received enough water to maintain its body weight of above 85% of the pre-surgery weight. After training each day on the visual task, mice received extra water if necessary in order to receive a minimum

daily amount that maintained the body weight relative to the pre-surgery weight. During each day of water restriction and behavioral training, body weight and general health were monitored in order to ascertain if there are any signs of dehydration.

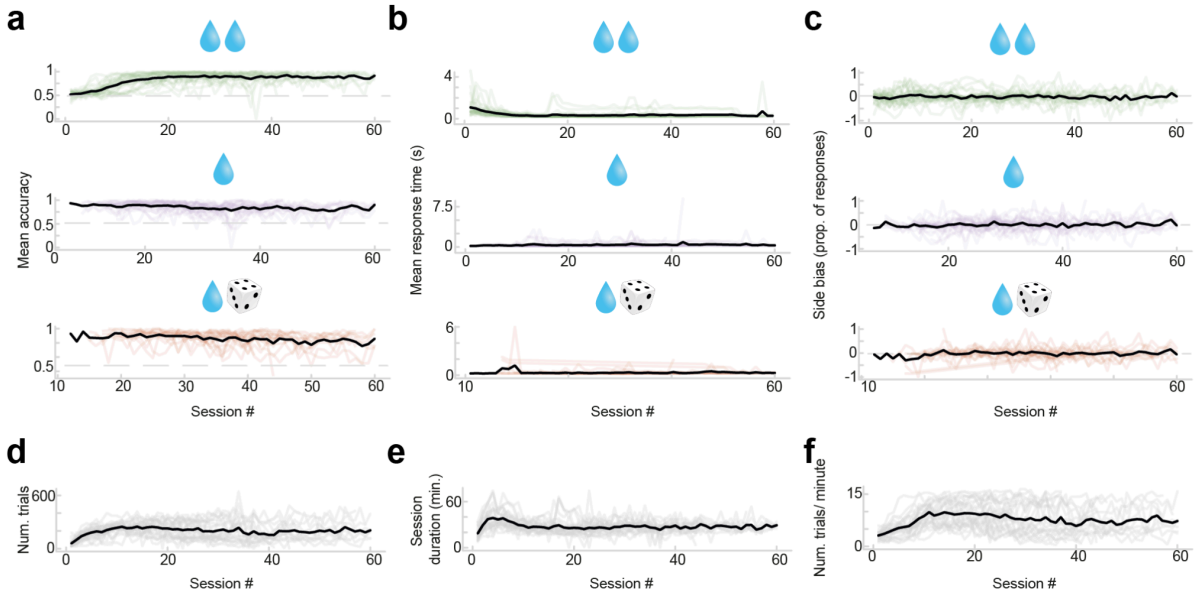


Fig 2.2 | Cumulative probability of mice surpassing a given performance level as a function of trial number

a, The mean accuracy of the mouse successfully moving the 'Large' (top), 'Small' (middle) and 'Gamble' (bottom) to the center of the screen per session and plotted across 60 sessions. The dotted gray line denotes chance level performance. Light colored lines indicate average accuracy for each mouse pertaining to each condition. Dark lines indicate the average across all mice. **b,** The mean response time, calculated by taking the difference between the choice complete time and the stimulus onset time, plotted by taking the average for each session and calculated for each condition separately. **c,** Side bias calculated for each session by taking the difference between right and left side choices and dividing by the total number of trials within a session. **d,** The total number of trials within each session for each mouse. As the number of training sessions progressed - the average number of trials within a session also steadily increased. **e,** Average session duration across 60 sessions. **f,** The rate of trials completed per minute across each session.

2.2.5 Training protocol

Mice were trained in several stages, by gradually introducing new stimuli.. In the initial stage of training, the mice were presented with only the large safe stimulus with a corresponding reward magnitude of 3 μL. Once mice achieved a sufficient level of accurate performance on the 'large' single-stimulus trials (>= 70% correct) and maintained this for 5-7 days, a stimulus associated with a

lower reward magnitude (‘small’ single-stimulus) was introduced. The last of the three single-stimulus conditions, with 3μL reward magnitude present with 50% probability - referred to as the ‘gamble’ condition - was introduced after the mouse reached 70% correct performance on the ‘small’ condition and maintained criterion performance for 5-7 days.

Following the introduction of all three single-stimulus conditions and reaching sufficient accuracy across all three, the two-stimuli conditions with different EV but same reward probability were introduced (‘large’ versus ‘small’). Mice had to display above a certain level of performance (>= 65% correct) on this different EV two-stimuli condition and maintain such a performance, before being introduced to the next different EV two-stimuli condition (‘large’ versus ‘gamble’). Similarly mice had to reach and maintain above a certain level of performance on this second two-stimuli condition before being introduced to the last two-stimuli condition which is a same EV two-stimuli condition where the mouse has to make a choice between a ‘small’ and a ‘gamble’ stimulus. Mice were trained at this stage for at least 1-2 weeks as this was the final stage of the training protocol. No threshold performance was expected from choices between stimuli with the same EV as we expected that the choice preference between these stimuli depended on individual risk sensitivities.

During the training period for all stimuli with the exception of the same EV two stimuli condition, mice were required to repeat trials if there was an incorrect response on the current trial. During these ‘repeat incorrect’ trials, the previous stimulus location for either single- or two-stimuli trials was repeated. The full task included six trial types (‘Large’, ‘Small’, ‘ Gamble’, ‘Large versus Small’, ‘Large versus Gamble’, ‘Small’ versus Gamble’). Once the training protocol was completed, the behavioral paradigm was combined with large-scale electrophysiology experiments to probe the neural correlates of the decision-making process.

Experimental sessions during which mice displayed disengagement from the task were removed from future analysis pipelines. General disengagement was characterized based on the following

criteria. We excluded trials with repeat numbers greater than 1 and reaction times longer than 10 seconds. In the case that a mouse is switched from one rig to a new one, the first two days on the new rig were excluded from analysis, along with the first two sessions with introduction of a new stimulus condition. Both these exclusion criteria allowed for adjusting the mouse's behavior to a novel environment and also adjusting to learning a new stimulus condition adequately. We further excluded sessions in which the choice accuracy on all three single stimulus conditions was lower than 80%. This led to an exclusion of 171 out of 618 sessions (27.67%).

2.2.6 Control behavioral experiment

We trained mice (n = 5) in a visual economic decision-making task control experiment, in order to provide further evidence that mice make decisions based on the economic value associated with a stimulus and not visual properties of the stimulus. The design of the experiment in this control experiment was similar to the main visual decision-making task with the exception of a few modifications. One of the modifications was the addition of two different abstract high contrast stimuli with the same reward magnitude (3 μL) and probability (100%), resulting in two high EV stimuli. We reasoned that if the mice were basing their decisions on the economic value of each stimuli, regardless of their visual properties, they would choose either high EV stimulus over a lower EV stimulus. We trained mice in various stages with different stages representing introducing new stimuli in sequential order in order to gradually increase the difficulty of the task and to ensure that the mice learn the task and the stimulus-economic value associations adequately. In the initial stage of the shaping process, mice were introduced to only the large safe stimulus with a corresponding reward output of 3 μL with a 100% probability. This portion is the same as the main experiment. However, we will refer to this stimulus condition as 'large 1' for future reference. Once mice reach the threshold level of performance on this stage of the task (>= 70% correct) and maintain this for 5-7 days, we introduce a second stimulus ('small' single-stimulus) similar to the main experiment. Importantly, after the mice reach criterion performance on the 'small' condition, we introduce the

'large 2' single-stimulus condition, a third stimulus with 3 μL reward magnitude present with 100% probability.

Following the introduction of all three single-stimulus conditions and reaching threshold accuracy across all three, the two stimuli conditions with different EV but same reward probability were introduced ('large 1' versus 'small'). Mice had to display above threshold performance on this different EV two stimuli condition (>= 65% correct) and maintain such a performance, before being introduced to the next different EV two stimuli condition ('large 2' versus 'small'). Similarly mice had to reach and maintain criterion performance on this second two stimuli condition before being introduced to the last two stimuli condition which is a same EV two stimuli condition where the mouse has to make a choice between a 'large 1' and a 'large 2' stimuli. Mice were trained at this stage for at least 1-2 weeks as this was the final stage of the training protocol. No threshold performance was expected from choices between stimuli with the same EV as slight variability in preferences for either 'large 1' or 'large 2' stimulus was anticipated.

2.2.7 Video monitoring

Face and eye videos were collected during the neuronal recording days. The eye and face movements of the mice were illuminated using an infrared light source (BW 48 LED Illuminator Light CCTV IR Infrared Lamp). The right eye is monitored using a camera (FLIR, lens: MVL50M23 - 50 mm EFL, f/2.8, for 2/3" C-Mount Format Camera, with Lock) (**Fig. 2.1c**). The face and body movements were monitored using a camera (FLIR, lens: MVL16M23 - 16 mm EFL, f/1.4, for 2/3" C-Mount Formal Cameras, with Lock) (**Fig. 2.1c**). The sampling rate of both videos was 20 Hz. The eye camera was adjusted and focused to pick up the pupil size and eye movements. The face camera was used to pick up orofacial and paw movements.

2.2.8 Facemap analysis

To investigate whether mice showed different licking behavior for large and small rewards, and whether they engaged in anticipatory licking before reward delivery, we extracted lick rates from face videos. Using Facemap (Syeda et al., 2024), lick rates were extracted from face videos recorded on Neuropixels recording days. In the analysis pipeline, we manually tracked the movement of two keypoints for the upper and lower lip, which was diagnostic of licking behavior. Refinement of the base model in Facemap used 50 training frames which were randomly sampled from one recording session. Our analysis pipeline for face videos implemented an exclusion criteria in order to avoid noise: we only included sessions in which the distance between the spout for water delivery and the mouth was pronounced and we could notice clearly defined licks. We excluded sessions from this analysis in which the water spout was positioned too close to the mouth, making it difficult to extract licks from the video data (39 of 114 sessions, 34.21%). The frame-by-frame cartesian coordinates of the key points were used to determine “mouth open” and “mouth closed” states. In particular, the Euclidean distance between keypoints was calculated, and distances exceeding 1 standard deviation were classified as “mouth open”. Licks were defined as a transition between “mouth closed” and “mouth open” states across successive frames.

We next analyzed the lick data aligned to different task events. In each trial, we only included events that occurred within a specific time window of -1 to 2 relative to outcome in order to analyze the lick rate during a pre- and post-outcome time window. To assess licking activity before and after outcomes, lick rates (licks/s) were calculated in a sliding window (200ms width, 4 frames) and aligned to outcome times. When comparing average lick rates between different trial types such as rewarded versus unrewarded trials, cluster permutation tests were conducted at the group-level (i.e. at the level of mice). The cluster-forming threshold was set to $\alpha = 0.05$. Permutations were computed by shuffling trial labels (e.g. rewarded and unrewarded) per mouse, and recording the maximum cluster statistic (summed t-values). This permutation procedure was repeated 1000 times. P-values were calculated as the proportion of permutations in which the maximum cluster

statistic exceeded that of the largest empirical cluster. The alpha-level was set to 0.025 (two-sided test). Further analysis focused on the mean lick rate within a time window of -0.1 to 0 seconds relative to outcome, which was taken as the 'pre-outcome' time period. Paired t-tests were carried out to determine if there were statistically significant differences in mean lick rate between the three single stimulus conditions during this pre-outcome time period.

2.3 Results

2.3.1 Mice make economic decisions between abstract visual stimuli

To study economic decisions guided by visual stimuli in head-fixed mice we developed a novel behavioral task, akin to those previously used for non-human primates (Padoa-Schioppa & Assad, 2006; O'Neill & Shultz, 2010). In our task, mice made choices among visual stimuli that varied in their associated reward magnitude and/or probability, thereby signaling different expected value (EV) and economic risk. Mice were head-fixed with a water spout close to their mouth, a rotating wheel beneath their forepaws, and a computer screen in front of them (**Fig. 2.1 a,b, 2.3a**). In each trial of the task, mice were presented with one or two visual stimuli on the screen (**Fig. 2.3b**). In 'single stimulus' trials, a single stimulus was presented on either the right or left side of the screen. The mouse had to bring the stimulus to the center by rotating the wheel, promoting delivery of the reward that the stimulus was associated with. In these trials, wheel movement to the opposite side moved the stimulus off screen and resulted in no reward. In 'two-stimuli' trials, mice were simultaneously presented with two distinct visual stimuli on opposite sides of the screen and had to make a choice between them by moving one of the stimuli to the center of the screen, resulting in the delivery of water reward associated with the chosen stimulus (**Fig. 2.3b**). The task included three visual stimuli (**Fig. 2.3c**). One stimulus predicted a large reward (3 µL) delivered with a probability of 100%, which we denote as a 'Large' stimulus with EV = 3 µL. A second stimulus predicted a small (1.5 µL) reward with a probability of 100% (EV = 1.5 µL), which we term a 'Small' stimulus. Lastly a third stimulus predicted a large reward (3 µL) delivered with a probability of 50%,

yielding a stimulus with EV = 1.5 μL. This stimulus thus had the same EV as the 'Small' stimulus but entailed risk, thereby we will refer to this as a 'Gamble' stimulus (**Fig. 2.3c**). The mapping between stimuli and associated reward was counterbalanced across mice. A subset of mice encountered both single stimulus and two stimuli trials randomly interleaved in each session, while other mice only encountered single stimulus trials.

Mice learned to perform the task successfully. Examining the performance of animals in single stimulus trials confirmed that mice were able to successfully move the stimuli to the center of the screen, as signified by high levels of accuracy across all three conditions (**Fig. 2.3d;** proportion correct trials: 0.92 $\pm$ 0.01, 0.88 $\pm$ 0.01 and 0.87 $\pm$ 0.01, mean $\pm$ SEM for Large, Gamble and Small stimuli, respectively). These high performances were largely independent of the stimulus position on the screen (**Fig. 2.3b**; 0.89 $\pm$ 0.003, 0.90 $\pm$ 0.003, mean $\pm$ SEM for left side accuracy, right side accuracy respectively). When examining the performance of animals in different single stimulus trials split by stimulus position on the screen, it revealed a bias in one of the three conditions but on average mice display balanced performance on both stimulus sides (**Fig. 2.5c;** 0.92 $\pm$ 0.01, 0.92 $\pm$ 0.01, accuracy difference: -6.2 x 10-4 $\pm$ 0.02; 0.83 $\pm$ 0.02, 0.88 $\pm$ 0.01, accuracy difference: 0.04 $\pm$ 0.02; 0.87 $\pm$ 0.01, 0.88 $\pm$ 0.01, accuracy difference: 0.01 $\pm$ 0.02; mean $\pm$ SEM and Right minus accuracy difference for Large- left side, Large - right side, Small - left side, Small - right side, Gamble - left side, Gamble - right side respectively).

Crucially, even with a single stimulus to choose from, animals' performance scaled with the EV: mice displayed the highest accuracy on trials with the highest EV stimulus (**Fig. 2.3e,** Large minus Small: 0.06 $\pm$ 0.01; Large minus Gamble: 0.04 $\pm$ 7.9 x 10-3; Small minus Gamble: -0.02 $\pm$ 9.8 x 10-3; accuracy difference - ΔLarge - Small: $t(34) = 5.50$, $p = 3.86 \times 10\text{-}6$; large safe minus gamble: $t(34) = 5.14$, $p = 2.24 \times 10\text{-}5$) and comparable accuracy for trials with the same EV but different risk (small safe minus gamble: $t(34) = 1.67$, $p = 0.11$). Examining the response times (computed by taking the difference of the stimulus onset time from choice start time for each trial) (**Fig. 2.4a**) of

animals in single stimulus trials, revealed fast reaction times for each condition (**Fig. 2.4b;** 0.29 ± 0.02, 0.31 ± 0.03 and 0.31 ± 0.02, mean reaction time in seconds ± SEM for large safe, gamble and small safe stimuli, respectively). We also examined the response times averaged across all single stimulus conditions and the overall performance on single stimulus trials and found that overall as mice overall exhibit higher accuracy and they also exhibit faster response times (**Fig. 2.10a;** $r(33) = -0.65$, $p = 2.44 \times 10\text{-}5$ (Pearson's correlation)). Consistent with choice accuracy, response times were also faster for the stimulus associated with a higher EV (**Fig. 2.4c;** Large minus Small: -0.03 ± 7.2 x 10-3; Large minus Gamble: -0.02 ± 9.1 x 10-3; Small minus Gamble: 9.8 x 10-3 ± 0.01; reaction time difference ΔLarge - Small: $t(34) = -4.05$, $p = 8.46 \times 10\text{-}4$; ΔLarge - Gamble: $t(34) = -2.13$, $p = 8.1 \times 10\text{-}2$; Δ Small - Gamble: $t(34) = 0.884$, $p = 0.38$). Inspecting choices in two stimuli trials revealed that mice were able to successfully choose the stimulus with higher EV on a majority of trials and thus make rational choices between two simultaneously presented choice options (**Fig. 2.3f;** P('Choose higher EV'); Large versus Small: 0.72 ± 0.03, $t(12) = 7.71$, $p = 5.5 \times 10\text{-}6$; Large versus Gamble: 0.659 ± 0.035, $t(12) = 4.47$, $p = 7.64 \times 10\text{-}4$). Mice were slightly but significantly faster choosing between Small and Large stimuli, differing in EV, compared to Small and Gamble stimuli with the same EV (**Fig. 2.4d;** Small versus Gamble: 0.31 ± 0.02; Large versus Small: 0.29 ± 0.02; reaction time difference: ΔLarge versus Small - Small versus Gamble: $t(12) = -3.23$, $p = 0.02$). Response times on Large versus Gamble trials were not significantly different from the other two-stimulus trials (**Fig. 2.4d;** Large versus Small: 0.29 ± 0.02; Large versus Gamble: 0.299 ± 0.02; reaction time difference ΔLarge versus Small - Large versus Gamble: $t(12) = -1.48$, $p = 0.165$). Overall, examining performance and reaction times on single and two stimuli trials revealed that mice are able to successfully make economic choices.

We also examined whether choice preferences on two-stimulus trials satisfied choice transitivity. In particular, we tested whether choices violated weak and strong axioms of revealed preference. For instance, the axiom of weak transitivity posits that a mouse preferring 'Large' over 'Gamble' and Gamble' over 'Small' stimuli, should also prefer 'Large' over 'Small' stimuli. We determined that the

choices made by a majority of mice satisfied choice transitivity (**Table 1;** 12/13 mice showed strong transitivity and 13/13 showed weak transitivity). Strong and weak transitivity are measurements for choice preferences among different pairwise comparisons with strong transitivity implying a higher strength of choice preferences among pairs of offers. Overall, examining choice transitivity allowed us to determine that mice exhibit consistent choice preferences when making economic choices guided by abstract visual stimuli.

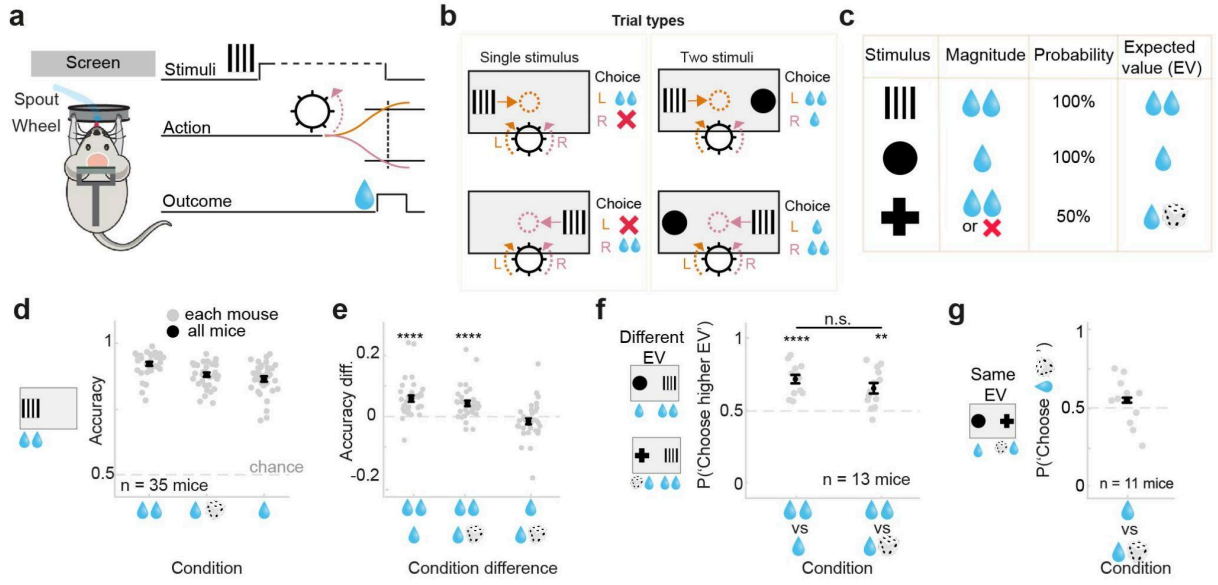


Fig. 2.3 | Mice make economic decisions between abstract visual stimuli

a, Schematic of the two alternative forced choice economic decision-making task. Head-fixed mice report decisions guided by visual stimuli on a computer screen by steering a wheel with their forepaws, to receive water reward delivered through a spout. **b,** Mice can encounter two randomly interleaved trial types. In single stimulus trials (**left**) one stimulus appears on the left or right side of the screen and the mouse has to move the stimulus to the center to receive the associated reward. Moving the stimulus outside of the screen is considered an incorrect trial and does not yield reward. In two stimuli trials (**right**), the mouse can choose between two simultaneously presented stimuli by moving one of the stimuli to the center of the screen. **c,** Each mouse is trained on three abstract, high contrast visual stimuli. The first two stimuli are associated with different reward magnitudes but the same guaranteed probability, yielding different EVs (termed 'Large' and 'Small' stimuli). The third stimulus has the same EV as the small safe stimulus but different risk (termed gamble stimulus). **d,** Behavioral choice accuracy for all three single stimulus conditions (n = 35 mice). Gray dots denote individual mice, and black dots show group averages. Error bars in all panels depict SEMs. Mice overall exhibit high accuracy across all three conditions, successfully moving a single stimulus to the screen center. **e,** Difference in accuracy between single stimulus conditions for n = 35 mice (one-sample t-test). Mice are more accurate in the high EV condition (large safe, EV = 3uL), compared to the low EV conditions (small safe and gamble, both EV = 1.5 uL, first and second columns). Choice accuracy is similar between conditions

with the same EV (small safe and gamble, third column). **f,** Two stimulus trial performance defined as the proportion of choosing the stimulus with the higher EV. Mice (n = 13, one-sample and paired t-tests) made rational choices, choosing the higher EV large safe stimulus over a lower EV small safe or gamble stimulus. **g,** Proportion of choosing the gamble stimulus over the small safe stimulus. Despite both stimuli having the same EV, mice (n = 11, one-sample t-test) exhibited diverse choice preferences ranging from risk-seeking (preferring the gamble stimulus) to risk-aversion (preferring the safe stimulus), but no consistent choice preference at the group level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, **** $p < 0.0001$, one-sample or paired t-tests.

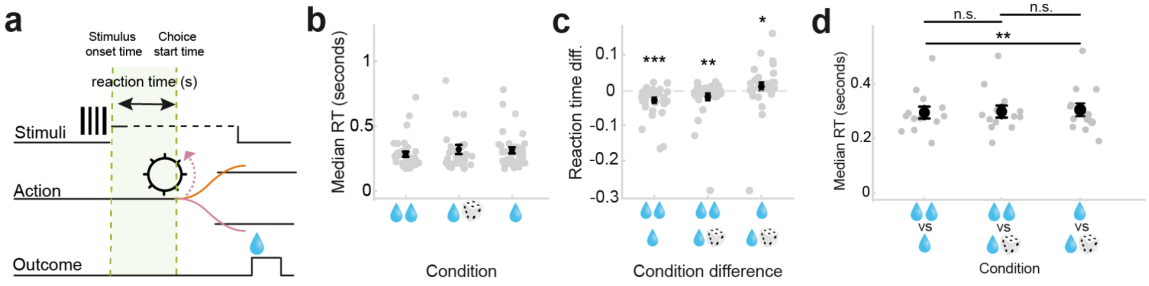


Fig. 2.4 | Mice display faster response times for stimuli with higher EV

a, Diagram showing the time course of a single trial and response time calculation by taking the difference between the choice start time and the stimulus presentation time. **b,** Median reaction times in seconds for all three single stimulus conditions (n = 35 mice). **c,** Difference in response times between single stimulus conditions. Mice (n = 35 mice, one-sample t-tests) overall exhibit faster reaction times for single stimulus conditions with a higher EV compared to the stimuli with lower EV. **d,** Median reaction times in seconds for the different and same EV two stimuli trials. Mice (n = 13 mice, one-sample and paired t-tests) overall do not exhibit a statistically significant difference in median RT between the two different EV conditions. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, one-sample or paired t-tests.

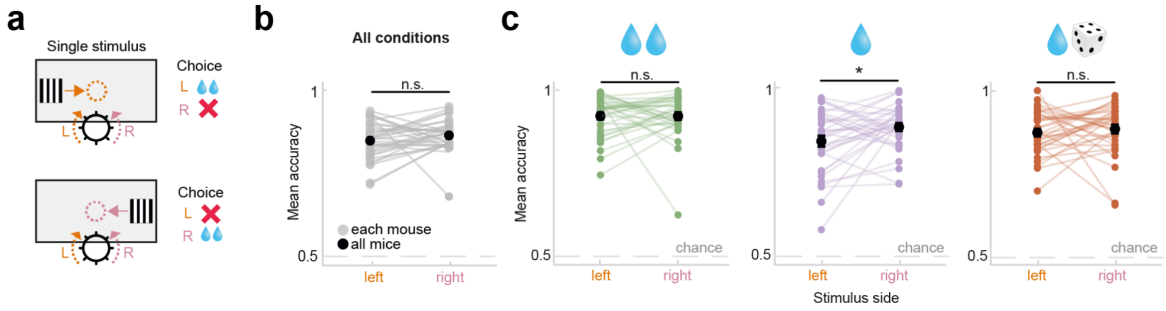


Fig. 2.5 | Choice accuracy split by stimulus side for single stimulus trials

a, Schematic of different stimulus positions in the single stimulus trials in the two-alternative choice economic decision-making task. The top schematic shows the single stimulus trials with a single stimulus presented on the left side of the screen. The schematic on the bottom shows a stimulus position on the right side. **b,** Behavioral choice accuracy (n = 35 mice, paired t-test) split by stimulus presentation position (either right or left side). Light gray dots denote individual mice, and black dots show group averages. Error bars in all panels depict SEMs. Mice overall exhibit high accuracy irrespective of the stimulus side. **c,** Same as in **b** but for individual single stimulus conditions (left: large, middle: small, right: gamble; n = 35 mice, paired t-test). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, paired t-tests.

mouse ID	P(💧💧>💧⚪)	P(💧⚪ > 💧)	P(💧💧 > 💧)	P(💧⚪>💧💧)	P(💧 > 💧⚪)	P(💧 > 💧💧)	weak transitivity	strong transitivity
ALK112	0.8444	0.5809	0.8562	0.1556	0.4191	0.1438	True	True
AMR001	0.5538	0.5377	0.5575	0.4463	0.4623	0.4424	True	True
AMR007	0.6136	0.7705	0.8861	0.3864	0.2295	0.1139	True	True
AMR008	0.6450	0.5859	0.7565	0.3555	0.4141	0.2435	True	True
AMR009	0.5208	0.4828	0.5714	0.4791	0.5172	0.4286	True	False
AMR011	0.4315	0.7430	0.6420	0.5685	0.2570	0.3580	True	True
AMR013	0.6847	0.6436	0.7967	0.3153	0.3564	0.2033	True	True
AMR014	0.5382	0.7069	0.7410	0.4618	0.2931	0.2590	True	True
AMR017	0.5887	0.5778	0.7461	0.4113	0.4221	0.2539	True	True
AMR019	0.8038	0.4252	0.7276	0.1962	0.5748	0.2724	True	True
AMR021	0.7687	0.3694	0.6473	0.2313	0.6306	0.3527	True	True
AMR028	0.7070	0.5867	0.8098	0.2930	0.4133	0.1902	True	True
AMR031	0.8103	0.2635	0.6185	0.1897	0.7365	0.3815	True	True

Table 1 | Transitivity of choice preference and testing for weak or strong transitivity

We examined whether choice preferences on two-stimulus trials ('Large versus Small', 'Large versus Gamble' and 'Small versus Gamble') satisfied choice transitivity. In particular we tested whether choices violated weak or strong axioms of revealed preference. An example scenario for when 'weak choice transitivity' is satisfied is when the mouse prefers 'Large over Small' at least 50% of the time, and 'Small over Gamble' at least 50% of the time, then 'Large' should be chosen over 'Gamble' at least 50% of the time. An example scenario for when 'strong choice transitivity' is satisfied is when the mouse prefers 'Large over Small' at least 50% of the time, and 'Small over Gamble' at least 50% of time, then 'Large' should be chosen over 'Gamble' more than the maximum preference between the two previous scenarios. In the table, we detail all six choice preference scenarios and use these to calculate whether the animal's choices satisfy 'weak' or 'strong' or both.

2.3.2 Mice make choices based on economic value and not visual properties

Two other lines of evidence indicate that mice make choices based on the economic value of the stimuli, rather than their visual properties. First, in a control experiment in which mice (n = 5) were trained to make choices among two 'Large' ('Large 1' and 'Large 2') and one 'Small' stimulus, mice were able to successfully choose the stimulus with higher EV on a majority of trials and thus make rational choices between two simultaneously presented choice options (**Fig. 2.6a;** 0.69 $\pm$ 0.04, 0.59 $\pm$ 0.03, mean $\pm$ SEM for: 'Large 1 versus Small' and 'Large 2 versus Small' respectively). Mice preferred both 'Large' stimuli over the 'Small' stimulus significantly above chance level (**Fig. 2.6a**; 'Large 1 versus Small' one-sample t-test against 0.5: $t(4) = 4.15$, $p = 1.43 \times 10\text{-}2$; 'Large 2 versus Small' one-sample t-test against 0.5: $t(4) = 2.96$, $p = 4.17 \times 10\text{-}2$). This was despite the Large 2 stimulus being introduced last. These results suggest that irrespective of differences in visual properties and order of introduction, mice are able to make choices based on the economic value associated with each stimulus.

Second, choice accuracy in single stimulus trials was not dependent on the visual properties of the stimuli as stimulus-reward pairings were counterbalanced across mice. Examining the performance associated with each visual object, regardless of its assigned economic value, confirmed that mice were able to successfully move the stimuli to the center of the screen irrespective of the stimulus identity, as signified by high levels of accuracy across all visual objects (**Fig. 2.6b (left)**; 0.88 $\pm$ 9.9 x 10-3, 0.90 $\pm$ 8.52 x 10-3, 0.89 $\pm$ 1.58 x 10-2 and 0.89 $\pm$ 1.16 x 10-2; mean +/- SEM for Circle, Cross, Horizontal Grating and Vertical Grating, respectively; **Fig. 2.6b (right)**). These results demonstrate that mice are capable of making choices guided by the economic value of visual stimuli.

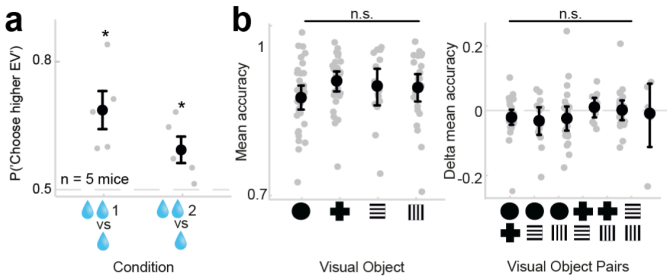


Fig. 2.6 | Mice make choices based on economic value not visual properties of stimuli

a, Behavioral choice accuracy on two stimuli conditions in a control experiment (n = 5 mice, one-sample t-test). The two conditions presented had 'Small' versus 'Large' in both conditions but with different visual cues denoting the large safe in each condition. Overall mice prefer both 'Large' stimuli over the 'Small' stimuli. **b (left),** Behavioral accuracy for different abstract visual objects across all mice (n = 35, paired t-tests). Gray dots denote individual mice, and black dots show group averages. Mice on average exhibit high accuracy across all visual objects. **b (right),** Difference in behavioral accuracy between different pairs of visual objects. On average, mice do not exhibit a preference for one visual object over another as the difference in mean accuracy is not statistically significant (n = 35 mice, 6 one-sample t-tests against 0 for each pair of visual objects difference in mean accuracy). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, one-sample or paired t-tests.

2.3.3 Mice exhibit diverse risk attitudes that evolve over extended periods of time

In this behavioral paradigm we assessed the preference for value and risk by introducing different trial types. The formal interpretation of risk, wherein the expected value is kept constant but there is introduction of variation through manipulation of the probability, was used in this study (Rothschild and Stiglitz, 1970). In this part of the task, only one level of risk was introduced - through a condition referred to as 'Gamble' - that had two reward magnitudes (3 and 0 μL) with equal probability (p = 0.5) and the expected value was identical to that of the 'Small' condition. At the group level, mice did not prefer one of these stimuli to the other, i.e. on average they were indifferent to the 'Gamble' or the 'Small' stimulus (**Fig. 2.3g**; P('choose gamble') = 0.55 $\pm$ 0.02), mean $\pm$ SEM. However, examining the preference of individual mice revealed a gradient of risk sensitivities from substantial risk-seeking to risk-aversive behavior. We used statistical tests, to assess whether a given mouse's mean per-day risk preference is statistically different from 50%. A statistically significant result implies that the mouse's risk preference is consistently different from 50% across days. Across all recording sessions, 6/11 mice were consistently risk-seeking (all $p < 0.05$, one-sample t-test of session-wise risk preference), whereas 3/11 mice were consistently risk-averse ($p < 0.05$, one-sample t-test of session-wise risk preference). Two mice were risk-neutral with no consistent risk preference.

We next asked whether individual risk attitudes remained stable across sessions of expert performance. Despite stable high performance of animals in choosing the stimuli with larger EV over sessions, mice exhibited a gradual decrease in their preference for the Gamble stimulus (**Fig. 2.7c,** P("choose Large over lower EV stimulus") one-sample t-test against 0: $t(9) = 1.16$, $p = 0.28$; P("choose Gamble over Small stimulus") one-sample t-test against 0: $t(9) = -3.70$, $p = 0.005$).

In spite of this gradual change in risk attitude over the long time scale (i.e. across sessions), the risk attitude of mice on a shorter (i.e. trial-by-trial) time scale did not change significantly; mice did not show a higher probability of choosing the gamble depending on whether they got a rewarded gamble outcome on the previous trial (**Fig. 2.7e**). Having shown that the risk attitude is stable within a session, we next asked whether this stability in risk attitude on a short-term timescale is uniform across several trials. In order to examine this, we analyze the probability of the mouse choosing the Gamble stimulus over the Small stimulus as a function of a previous Gamble win on different n-back trial numbers (**Fig. 2.7f**). Specifically we asked whether there is a particular previously rewarded Gamble trial that might be influencing the current trial risk attitude. We found a similar risk attitude irrespective of the n-back rewarded Gamble trial. Finally, we wanted to see whether a previously unrewarded Gamble trial impacts the current trial risk attitude. To examine this, we analyze the probability of the mouse choosing the Gamble stimulus over the Small stimulus as a function of a previous Gamble loss on different n-back trial numbers (**Fig. 2.7g**). We found a similar risk attitude irrespective of the n-back unrewarded Gamble trial. These results indicate that mice exhibit diverse risk attitudes and gradually decrease their risk seeking behavior.

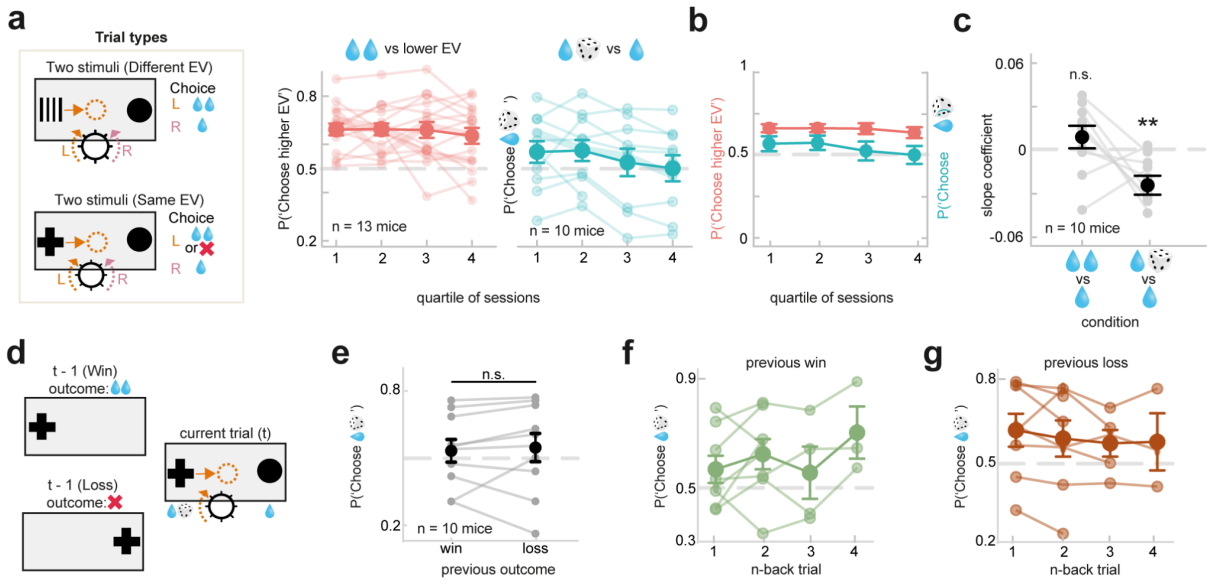


Fig. 2.7 | Risk sensitivity evolves across sessions

a, Schematic of the different two stimuli trials types. On the top is an example trial type with a 'large' stimuli presented on the left and a 'small' presented on the right. These two stimuli when presented simultaneously have a difference in EV and therefore these are labeled as a 'different EV' trial type. In the schematic on the bottom, is an example trial with a 'gamble' stimuli presented on the left and a 'small' stimuli presented on the right. **(left side graph)** Behavioral choice accuracy for both different EV two stimuli trial conditions split by all sessions for a mouse being split into different quartiles (n = 13 mice). Lighter lines denote individual mice, and dark and thicker lines show group averages. Error bars in all panels depict SEMs. Mice overall exhibit a stable and high accuracy across both different EV two stimuli conditions, successfully choosing the stimuli with the higher EV ('large') and moving it to the center of the screen. **(right side graph)** Behavioral choice metric for risk sensitivity (probability of choosing the 'gamble' stimulus over the 'small safe') for the same EV two stimuli condition across all sessions for each mouse divided into different quartiles (n = 10 mice). Lighter lines denote individual mice and darker and thicker lines show group averages. Mice overall exhibit a risk seeking behavior that shifts towards neutrality (equal probability of picking 'gamble' or the 'small' stimulus). **b,** Comparison of averaged risk sensitivity and behavioral accuracy on 'different EV two stimuli' trials across all sessions for all mice divided into different quartiles. On a group level, as the behavioral accuracy on different EV two stimuli trials is stable across different quartiles, the risk sensitivity on same EV two stimuli trials shifts across quartiles of sessions. **c,** Slope coefficient for linear regressions describing the change in either preference for the stimulus with the higher EV or for the mouse choosing the 'Gamble' stimulus over the 'Small' in the risky two stimuli condition (n = 10 mice, one-sample t-tests). **d,** Schematic of a current trial and the trial prior having one of two different outcomes: either the 'Gamble' was chosen and the mouse received a large reward or the 'Gamble' was chosen and the mouse did not receive any reward. **e,** When selecting between stimuli with the same EV but different risk and asking whether the probability of choosing the Gamble stimulus over the Small stimulus is influenced by a previous Gamble win or loss, mice (n = 10 mice, paired t-test) did not exhibit a win-stay or lose-switch strategy. This shows that the risk attitude of mice on a shorter (i.e. trial-by-trial) time scale did not change significantly. **f,** The probability of the mouse choosing the Gamble stimulus over the Small stimulus as a function of a previous Gamble win on different n-back trial numbers. We show that there is a similar impact on gamble attitude irregardless of how many

trials back the mouse got a reward for picking a Gamble stimulus (what we refer to as a 'Win'). **g,** The probability of the mouse choosing the Gamble stimulus over the Small stimulus as a function of a previous Gamble loss on different n-back trial numbers. We show that there is a similar impact on gamble attitude irregardless of how many trials back the mouse did not receive a water reward for picking a Gamble stimulus (what we refer to as a 'Loss'). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, one-sample or paired t-tests.

2.3.4 Trial-by-trial updating occurs in the space of actions, not stimuli

Early models investigating economic choice have tried to examine how different animal species make choices. One of the main frameworks that emerged from early work is that economic decisions occur based on action selection, an action-based hypothesis (Lau & Glimcher et al., 2007). Typically when making decisions between different stimuli, it is believed that the brain assigns values to each option and then compares them to ultimately make a choice. One of the main questions that remains in neuroscience is how the brain makes these decisions - is it by comparing the value of the goods in a goods space or by first assigning the values to the associated actions and then making a choice over actions. Previous work in primates shows that monkeys tend to choose the same cue chosen in a previous trial, what is referred to as choice hysteresis – or stimulus based updating. However despite work in primates suggesting that the formation of economic decisions begins in the goods, i.e. stimulus space (Padoa-Schioppa et al., 2011; Cai and Padoa-Schioppa, 2014), it remains unknown whether economic decisions can happen as a result of an organism selecting their preferred economic offer or by selecting an action that will obtain the good. It is therefore imperative to investigate how in our economic decision-making framework, value-based decisions take place in either the goods or the action space.

We examined how choices in our economic decision-making task may be influenced by either actions, goods or a combination of the two. Here we examined the two possible sources of choice biases related to the outcome of the previous trial: either action-based updating (where the animal is biased to repeat the previously rewarded action) or stimulus-based updating (where the animal is biased to choose a previously rewarded stimulus regardless of the action that was taken to obtain it).

In order to examine these possible choice biases, we analyzed whether the current trial stimulus side chosen is guided by a previously rewarded action. Mice presented a consistent action-based updating bias, repeating previously rewarded actions (**Fig. 2.8 a-c**: P('choose right with previous rewarded left') = 0.44 ± 0.04; P('choose right with previous rewarded right') = 0.54 ± 0.04; , mean ± SEM; delta-P("choose right") - Δprevious rewarded right - previous rewarded left: $t(15) = 6.62$, $p = 8.2 \times 10\text{-}6$). Having shown that mice exhibit a bias of repeating previously rewarded actions, we next examined whether animals showed a bias linked to past outcomes linked to particular abstract visual stimuli.

We show that animals showed relatively no bias by past outcomes linked to abstract visual stimuli, independent of action. We first examined whether this stimulus-based bias is guided by a previously rewarded Large stimulus or Small stimulus (**Fig. 2.8 d-f**: P('choose Large with previous rewarded chosen Large') = 0.71 ± 0.03; P('choose Large with previous rewarded chosen Small') = 0.68 ± 0.04; mean ± SEM; delta-P("choose Large") - Δprevious rewarded Large - previous rewarded Small: $t(15) = 0.82$, $p = 0.43$) and found no significant effect. We next examined whether a stimulus-based bias is guided by a previously rewarded Large stimulus or Gamble stimulus (**Fig. 2.8 g-i**: P('choose Large with previous rewarded chosen Large') = 0.61 ± 0.04; P('choose Large with previous rewarded chosen Gamble) = 0.68 ± 0.04; , mean ± SEM; delta-P("choose Large") - Δprevious rewarded Large - previous rewarded Gamble: $t(15) = -1.71$, $p = 0.11$) and once again found no significant effect of either previously rewarded cue on current choice behavior. Lastly, we asked whether mice exhibit a bias linked to past outcomes linked to either a Small or Gamble stimuli (**Fig. 2.8 j-l**: P('choose Small with previous rewarded chosen Small') = 0.47 ± 0.03; P('choose Small with previous rewarded chosen Gamble) = 0.45 ± 0.04; , mean ± SEM; delta-P("choose Small") - Δprevious rewarded Small - previous rewarded Gamble: $t(15) = 0.56$, $p = 0.58$) and similarly did not find an effect. These results indicate that in our experiments trial-by-trial updating happens in the space of actions and not stimuli. Overall, we identify a bias to repeat previous actions but no stimulus-based bias which

suggests that the mouse's understanding of value is relatively stable over the course of a short timescale.

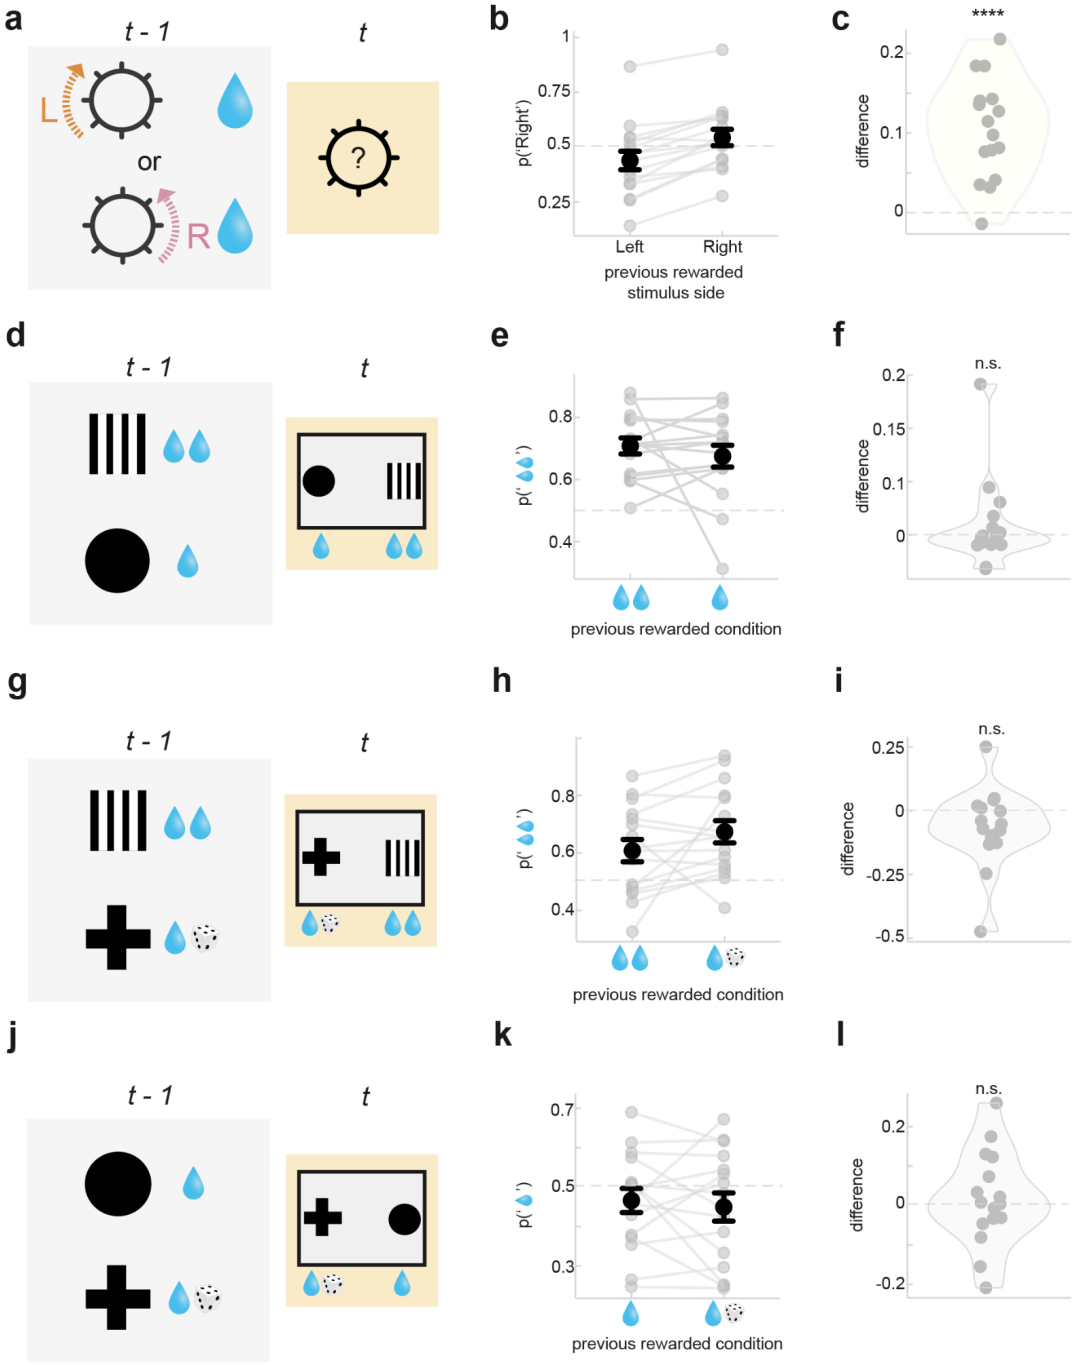


Fig. 2.8 | Trial-by-trial updating happens in the space of actions

a, Schematic of the previous trial (t-1) and the current trial (t). In the previous trial, we depict one of two different rewarded choices that a mouse can make: either picking the stimulus on the left by moving the stimulus from the left side of the screen to the center, or picking the stimulus presented on the right side.

Given that either of these choices was rewarded, we are interested in the current trial choice: will the mouse pick the stimulus on the right or the left side in a two stimuli trial context? **b,** Probability of the mouse picking the right side stimulus in a current two stimuli trial split by previous rewarded left or right side trial. Light gray dots denote individual mice, and black dots show group averages. Error bars in all panels depict SEMs. Mice (n = 16 mice) overall exhibit an increased propensity to choose the side that was rewarded in a previous trial. **c,** Difference in the probability of picking Right given a previously rewarded Right side choice minus the probability of picking Right given a previously rewarded Leftward choice. On average there is a significant difference (one-sample t-test) between these two probabilities, suggesting that the choices mice make are guided by an action updating strategy. **d,** Schematic of the previous trial (t-1) and the current trial (t). In the previous trial, we depict one of two different rewarded choices that a mouse can make: either picking the cue associated with a large reward and 100% probability ('Large'), or picking the cue with a small reward and 100% probability ('Small'). Given that either of these choices was rewarded in either trial type, we are interested to see whether in a current two stimuli ('Large versus Small') trial, the mouse exhibits increased propensity to pick the previously rewarded cue. **e,** Same as (b) but plotting the probability of picking the Large over the Small in a current two stimuli trial split by a previously rewarded chosen Large or a previously rewarded Small (n = 16 mice). **f,** Same as (c) but with the difference in the probability of picking Large given a previously rewarded Large choice minus the probability of picking Large given a previously rewarded Small (n = 16 mice, one-sample t-test). On average there is no significant difference between these two probabilities, suggesting that the mice's choices are not guided by a cue-updating strategy. **g,** Schematic of the previous trial (t-1) and the current trial (t). In the previous trial, we depict one of two different rewarded choices that a mouse can make: either picking the cue associated with a large reward and 100% probability ('Large', or picking the cue with a large reward and 50% probability ('Gamble'). Given that either of these choices was rewarded in either trial type, we are interested to see whether in a current two stimuli ('Large versus Gamble') trial, the mouse exhibits increased propensity to pick the previously rewarded cue. **h,** Same as (b) but plotting the probability of picking the Large over the Gamble in a current two stimuli trial split by a previously rewarded chosen Large or a previously rewarded Gamble (n = 16 mice). **i,** Same as (c) but with the difference in the probability of picking Large given a previously rewarded Large choice minus the probability of picking Large given a previously rewarded Gamble (n = 16 mice, one-sample t-test). On average there is no significant difference between these two probabilities, suggesting that the choices mice make are not guided by a cue-updating strategy. **j,** Schematic of the previous trial (t-1) and the current trial (t). In the previous trial, we depict one of two different rewarded choices that a mouse can make: either picking the cue associated with a small reward and 100% probability (Small), or picking the cue with a large reward and 50% probability ('Gamble'). Given that either of these choices was rewarded in either trial type, we are interested to see whether in a current two stimuli ('Small versus Gamble') trial, the mouse exhibits increased propensity to pick the previously rewarded cue. **k,** Same as (b) but plotting the probability of picking the Small over the Gamble in a current two stimuli trial split by a previously rewarded chosen Small or a previously rewarded Gamble (n = 16 mice). **l,** Same as (c) but with the difference in the probability of picking Small given a previously rewarded Small choice minus the probability of picking Small given a previously rewarded Gamble (n = 16 mice, one-sample t-test). On average there is no significant difference between these two probabilities, suggesting that the choices mice make are not guided by a cue-updating strategy. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, **** $p < 0.0001$, one-sample t-tests.

2.3.5 Post outcome lick rates are reflective of reward magnitude

To investigate whether mice showed different licking behavior for large and small rewards, and whether they engaged in anticipatory licking before reward delivery, we extracted lick rates from face videos. We used Facemap to track the positions of two points on the face of each mouse performing the visual decision-making task. These two points are the upper lip and the lower lip and using the Euclidean distance between these two points, we were able to examine the licking activity aligned to different behavioral events. One of the main questions to address is whether the licking rate is scaled to the delivered reward magnitude post-outcome period when the mouse receives the water reward. In our task, mice received water rewards that vary in magnitude, for example, the reward associated with a successful picking of the 'Large' single stimulus trial is 3 ul of water. Using Facemap we were able to reliably track from individual frames of each recorded videos' periods of either 'No lick' or when there is a 'Lick'. Furthermore by implementing a threshold for the z-scored distance between the upper and lower lip keypoints, we were able to further examine the lick rate as we estimated licks by thresholding the keypoint distance to determine 'mouse-open' and 'mouth-closed' states (**Fig.2.9 a-c**).

We first asked whether the lick rate aligned to outcome time is different between trials where the mouse receives a reward versus when they either make an incorrect choice or have a 0 μl reward with the chosen 'Gamble' stimulus. Indeed, we found that the mean lick rate for such rewarded and unrewarded trials was significantly different in a period post-outcome (**Fig. 2.9d, $p < 0.0001$, cluster based permutation**). We next examined whether there is a difference in licking rate for different reward outcomes. To analyze this we compared the mean lick rate in the post-outcome period, revealing different lick rates for distinct reward magnitudes: 3 μl and 1.5 μl (**Fig. 2.9e, $p < 0.0001$, cluster based permutation**). Lastly, we asked whether the lick rate is different for distinct conditions in the pre-outcome time window. In order to examine this anticipatory licking, we first compared the average lick rate between a stimulus time period (-0.45 to -0.35 seconds relative to outcome) when the licking rate was low followed by the licking rate in the pre-outcome time period

(-0.1 to 0 seconds relative to outcome) which had a statistically higher licking rate than post-stimulus presentation (**Fig. 2.9h, p = 0.00000145, pairwise t-test**). Having established a relatively higher lick rate during the pre-outcome period relative to the stimulus time period, we examined if there are differences in licking rate split by conditions in the pre-outcome period. No significant difference was found in the licking rate of each of these conditions in the pre-outcome period (**Fig. 2.9g, Large vs. Small (p-value: 0.99); Large vs. Gamble (p-value: 0.98); Small vs. Gamble (p-value: 0.95), pairwise t-test**). Overall, the results indicate that lick rate is different in the post outcome delivery time scaled to the received reward magnitude. Conversely, while mice engage in pre-outcome anticipatory licking, the pre-outcome licking rate is not influenced by the expected reward associated with each stimulus.

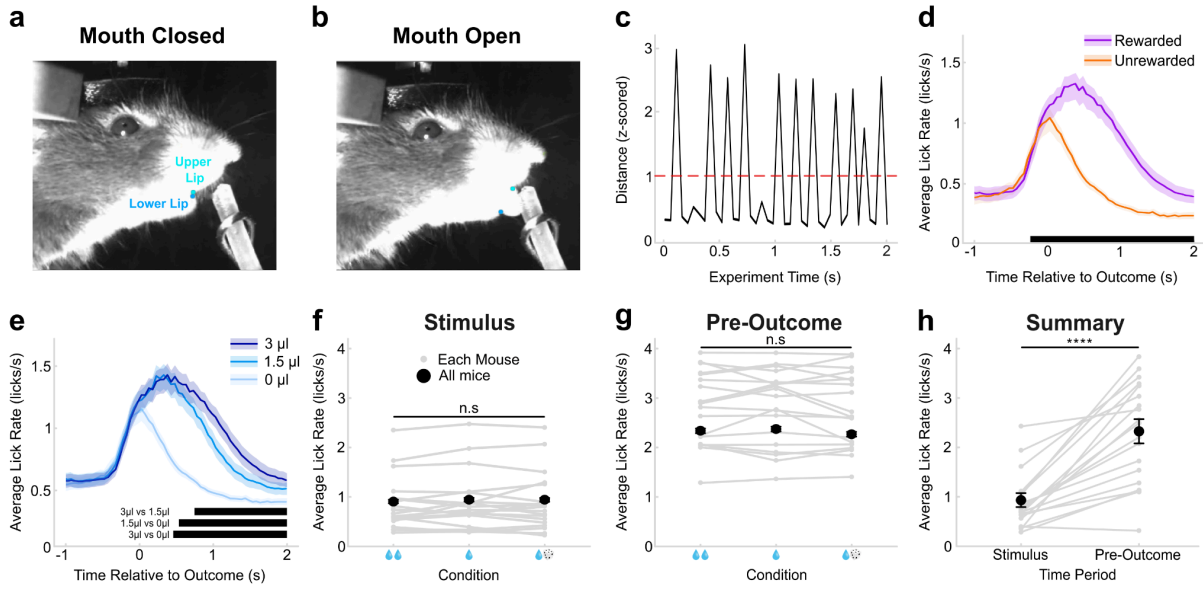


Fig. 2.9 | Facemap analysis of lick data

a, An example frame of a face video showing the absence of a lick, with the upper and lower lip keypoints labeled. **b,** An example frame of a face video showing the presence of a lick, with the upper and lower lip keypoints labeled. **c,** z-scored distance between the upper and lower lip key points plotted over 2 seconds of an example session. The threshold of 1 standard deviation, which has been used to denote a lick, is shown by a red, dashed line. **d,** Mean lick rate relative to outcome time for rewarded and unrewarded trials (n = 19 mice) calculated using a rolling mean and a bin size of 200ms. The shaded area for each line depicts the SEM. Mice exhibit a statistically significant difference in mean lick rate between rewarded and unrewarded trials ($p < 0.0001$, cluster-based permutation test). The extent of the significant cluster is shown by the black bar along the x axis. **e,** Mean lick rate for different reward magnitudes (n = 19 mice). A cluster permutation test

indicated that mice exhibit statistically significant differences ($p < 0.0001$) in mean lick rate between the various reward magnitudes. The time window of the significant clusters are shown by the dark bars along the x axis for each comparison. **f,** Mean lick rate for all three single stimulus conditions (n = 19 mice) during the stimulus time period (-0.45 to -0.35 seconds relative to outcome). Gray dots denote individual mice, black dots show group averages and error bars depict SEMs. Pairwise t-tests showed that mice overall do not exhibit statistically significant differences in mean lick rate between the three single stimulus conditions during this time period. **g,** Mean lick rate for correct trials of all three stimulus conditions during the pre-outcome time period (n = 19 mice). Gray dots denote individual mice, black dots show group averages and error bars depict SEMs. Mice did not exhibit statistically significant differences in mean lick rate between the three single stimulus conditions during this time period (pairwise t-tests, all $p > 0.05$). **h,** Mean lick rate per second (n = 19 mice) during the stimulus and pre-outcome time period. A pairwise t-test indicated that mice exhibit a statistically significant difference ($p < 0.0001$) in mean lick rate between these time periods, indicating an increase in anticipatory licking prior to reward delivery.

2.4 Discussion

Studying how neurons encode different economic variables requires a behavioral task that integrates key principles driving economic choices. Several animal studies in the past few decades have sought to create such a framework to ethologically characterize how animals make such decisions. Previous studies investigating economic decision-making have incorporated several variables crucial to making economic decisions such as reward magnitude, expected value, probability and risk (Bechara et al., 1994; Bechara et al., 1999; Stauffer et al., 2015; McCoy & Platt, 2005; Samejima et al., 2005; Bateson et al., 2002; Yamada et al., 2013; Kagel, Battalio & Green, 1995; Padoa-Schioppa & Assad, 2006; Levy et al., 2010; Peters and Büchel, 2009; Constantinople et al., 2019b; Kuwabara et al., 2020; Orsini et al., 2015). These human and animal model studies have been fruitful in growing our knowledge of choice biases, subjective value preference, risk preferences and aversion and loss aversion among many other aspects to behavioral neuroeconomics. However, previous studies in mice do not use abstract visual stimuli and instead use olfactory stimuli (Kuwabara et al., 2020) and choices rely on cue-reward associations formed with olfactory stimuli thereby making the link to work done in non-human primates incomplete. To account for these limitations in previous mouse neuroeconomics studies, we adapted a standardized mouse visual psychophysics task (Burgess et al., 2017, International Brain Laboratory, 2020) and incorporated important elements in order for our task to resemble the task developed for non-human primates (Padoa-Schioppa & Assad, 2006). This

extension to the standardized mouse head-fixed behavioral paradigm yielded several conclusions from our data, and provided a platform to use high-yield neuronal measurement methods available in mice to investigate neural bases of decision making.

We developed a visually guided economic decision-making task for head-fixed mice where animals made decisions guided by abstract visual cues, each associated with a different expected value and risk. This task was inspired by and thereby integrated elements from previous economic tasks developed for non-human primate cognitive studies. Such integrated elements include: (a) animals choosing between two stimuli presented simultaneously, and (b) learning different cue-reward magnitude and probability associations. Mice learned the first stage of this task within 1-1.5 weeks on average and the entirety of the task within 2-2.5 months. Additionally, the task yields 200-400 trials per day on average which allows for reliable measures of behavioral accuracy on different stimulus-EV associated conditions within and across sessions (**Fig. 2.2**). Head-fixing mice in this task enables future steps for electrophysiological experiments wherein one can observe how neural populations are encoding different decision-making variables.

The experiments reveal several new insights into choice behavior of mice. First, mice exhibit choice patterns that overall have high accuracy across all single-stimulus trials and differences in choice accuracy are reflective of choices being guided by economic value. Although there was only one stimulus to choose from, mice were more accurate for the stimulus with a higher EV. Mice were able to successfully make economic choices by choosing the stimulus with a higher EV between two simultaneously presented choice options. These together with other control experiments and analyses indicate that these choices were guided by the economic value and not visual properties of stimuli. Second, we show that mice exhibit a diversity of risk attitudes spanning from mice preferring the risky 'Gamble' stimulus, to preferring the risk-free 'Small' stimulus and finally to mice that are indifferent to picking either the 'Gamble' or 'Small' stimulus. Third, we show that despite mice not showing a win-stay or lose-switch strategy depending on the gamble outcome of the

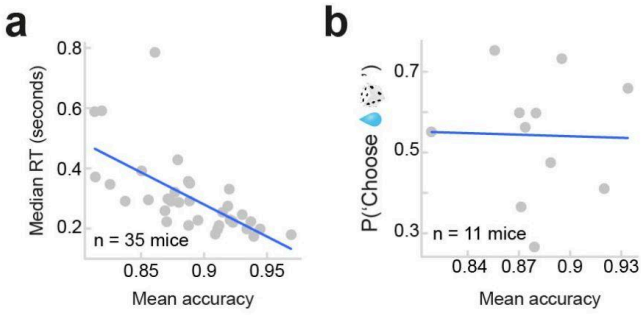
previous trial, risk attitudes shifted across sessions moving from risk-seeking towards aversion. This diverse risk attitude is comparable to recent reports in rats (Constantinople et al., 2019a). In contrast to rodent studies, most studies on head-fixed monkeys have reported mostly risk seeking behavior (Lak et al., 2014, Stauffer et al., 2014, Yamada et al., 2013). This cross-species difference requires further investigation in future studies. Mice exhibited a consistent action-based choice bias and no bias by past outcomes linked to abstract visual stimuli, suggesting that mice make economic decisions in an action space. Finally, examining the licking behavior aligned to reward outcome reveals distinct post-outcome lick rates for conditions with different reward magnitude while no difference in anticipatory licking was observed across different conditions. In particular, if we were to analyze the neural signals at outcome contrasting the post-reward lick rate between a won Gamble and won Large reward would be interesting because then you would have the same physical quantity of water but the mouse might have expected something different. The uncertainty of the Gamble might impact this post-outcome licking response so you might see differences in licking to the reward but they might have had different expectations in the two conditions. This might reflect in the neural signals - such as a reward prediction error that is different based on whether one is extremely sure that one will get a reward and or there is a 50% chance of receiving a reward. This would be a future direction in connecting the outcome-aligned licking analysis to the neural activity (**Chapter 3**). In conclusion, our results suggest that mice are able to make economic decisions and exhibit diverse risk attitudes.

2.5 Acknowledgements

This work was carried out in partnership with Caitlin Ashcroft and Lauren Strickland (2 doctoral students in the Lak Lab). When I first joined the Lak Lab in October 2020, Lauren guided me on training mice in this economic choice paradigm, and she collected data from 3 mice. Caitlin and I carried out the remaining behavioral experiments with Caitlin contributing behavioral data from 15 mice and I contributed behavioral data from 28 mice. Anita-Maria Rominto helped with training mice during the 6 months of her Masters project visit to the lab. Prof. Armin Lak supervised the

behavioral experiments throughout and also conceptualized the behavior paradigm for economic decision-making. Dr. Peter Zatka-Haas set up the MC code for running the economic decision-making behavior task with different abstract stimuli associated with different reward magnitudes and probabilities.

2.6 Supplementary Figures



Supplementary fig. 2.10 | Correlating behavioural performance with response times and risk preference

a, This analysis correlates the response times of individual mice against performance. The main question we ask here is whether a mouse that has overall better performance is faster? Median reaction times in seconds averaged across all 3 single stimulus conditions is shown on the y-axis and the mean behavioral choice accuracy averaged across all 3 single stimulus conditions is shown on the x-axis (n = 35 mice). Gray dots denote individual mice and the blue line denotes the Pearson correlation line of best fit. Mice overall exhibit higher accuracy and faster response times r(33) = -0.65, p = 2.44 x 10-5 (Pearson's correlation). **b,** This analysis correlates the performance of individual mice against their risk sensitivity (their preference towards picking the Gamble stimulus. The main question we are asking is whether the risk-sensitivity of the mouse correlates with overall task engagement (signified by overall higher accuracy across single-stimulus conditions). Proportion of choosing the gamble stimulus over the small safe stimulus is depicted on the y-axis and mean behavioral choice accuracy averaged across all 3 single stimulus conditions is depicted on the x-axis (n = 11 mice). Despite mice having overall higher task engagement, there was no correlation with the risk preference. Mice that had stronger task engagement exhibited diverse choice preferences ranging from risk-seeking (preferring the gamble stimulus) to risk-aversion (preferring the safe stimulus), but no consistent choice preference at the group level r(9) = -0.03, p = 0.94 (Pearson's correlation).

3

Representations of economic variables across the frontal cortex

Contents

3.1 Introduction . . . . . .
3.2 Methods . . . . . .
3.3 Results . . . . . .
3.4 Discussion . . . . . .
3.5 Acknowledgments . . . . . .

3.1 Introduction

A fundamental question of neuroeconomics research is how our brain makes economic decisions. Since the infancy of economic studies of decision-making in the brain, important questions regarding the structure of neuronal representations remain. Past studies have investigated how neural activity in frontal cortical areas encodes various economic variables and how these neural

signals are crucial for economic decision-making. In particular, studies from the past few decades have suggested that regions such as the ventromedial frontal lobe, and the orbitofrontal cortex, among others play a crucial role in encoding of decision parameters such as expected value (EV) and risk. While some studies propose that distinct brain areas represent different economic variables, others suggest that the neural representations of economic decision variables are distributed across several regions, with neurons in different regions encoding an overlapping array of variables such as potential payoff, the probability of success, and the cost in terms of time and effort (Kennerley et al., 2009a). However, the main limitation with this past work is that such studies have studied these representations in individual regions in isolation and used a variety of behavioral tasks, making it challenging to compare neuronal activity across different frontal regions and projections. As such, the principles that underlie such representations remain unknown. To address this, we developed a visually-guided economic decision task for head-fixed mice (detailed in **Chapter 2**) and recorded the activity of thousands of neurons across different frontal regions during this economic task. Before we delve into our neural findings, I would like to highlight the paramount lesion studies, neuroimaging and electrophysiological studies that have paved the way for studying the structure of representations of economic decision variables.

One of the first clinical studies that inspired a long history of clinical frontal lobe investigations involves the clinical history of Phineas Gage. As people followed the clinical history of Phineas Gage, one observation was of particular intrigue: in the years following his accident and recovery, his decision-making capabilities were altered with the main changes described as a "deficit in social decision-making" (Macmillan et al., 2000). Regrettably, the empirical value of an anecdotal case such as this one is limited, as the extent of the lesions and damage to the brain is not just limited to the frontal region. Following the clinical case of Phineas Gage, several studies in human patients stemming from lesion, neurodegeneration and dysfunction added to the growing characterization of behavioral changes such as suboptimal decision-making and the role of frontal cortex in different aspects of cognition such as the use of reward history to adaptively guide behavior (Culberth et al.,

2016; Bechara et al., 1994; Camille et al., 2011; Fellows et al., 2011; Hodges et al., 2001; Rahman et al., 1999; Waltz and Gold, 2016; Yu et al., 2018). Neurological damage in humans in areas such as the ventromedial prefrontal cortex, revealed impairment of decision-making within the framework of formal behavioral tasks. Such tasks often simulated real life decision-making and integrated factors such as uncertainty of outcomes, and reward and punishment (Bechara et al., 1994). In another study, patients with damage to the ventromedial frontal lobe (VMF) showed impairments on a cognitive task that requires them to maximize on acquiring rewards, known as value maximization (Camille et al., 2011). Overall, lesion studies in humans shed light on the frontal cortical regions that are important for decision-making and rational behavior.

Although the lesion studies in humans were insightful, they were rarely specific to individual brain regions, making the argument that these regions carry representations of economic variables incomplete. Targeted studies in primates furthered the insight into the particular roles that individual frontal regions play in decision-making. Impairments similar to those observed in human studies were observed in a lesion study in macaques, where different lesions in either the orbitofrontal cortex (OFC) and anterior cingulate cortex (ACA) disrupted behavior for different cognitive components on the Wisconsin Card Sorting Task analog (Buckley et al., 2009). OFC lesions specifically lent evidence to the role of the OFC in representing the rule value and updating the values of offers. In contrast, the ACA lesions in the same study highlighted the importance of this region in guiding decisions by selecting an offer and integrating recent choice-outcomes (successes or losses) into rule-based decision-making. This macaque lesion study served as a bridge between human neuropsychological and clinical lesion studies and investigations of value-based decision-making in nonhuman primates. Many other studies have pointed to the prominent role the ACA, OFC and other frontal regions plays in different types of decision-making (Sul et al., 2010; Lee et al., 2016; Rushworth et al., 2007; Rushworth and Behrens, 2008; Rangel and Hare, 2010; Wallis and Kennerley, 2010; Walton and Behrens, 2010). Overall, lesion and neurodegeneration studies point to the crucial role of different regions in the frontal cortex in making decisions. These

studies collectively suggest a critical role for various frontal regions in decision-making, but the reliance on lesion studies makes it difficult to disentangle whether the activity in these regions is related to the assessment of value and other economic variables in the provided context.

A more spatially and temporally detailed approach using single cell recordings or neuroimaging allowed for closer examination of the role of these regions in representing individual economic variables. Following the initial lesion studies, a raft of electrophysiology (Kepecs et al., 2008; Padoa-Schioppa and Assad, 2006; Padoa-Schioppa and Assad, 2008; O'Neill & Schultz, 2010), and functional neuroimaging studies (McClure et al., 2004; Kable and Glimcher, 2007; Plassman et al., 2007; Tom et al., 2007; Hare et al., 2008; Chib et al., 2009; Knutson et al., 2005; Abler et al., 2006; Yacubian et al., 2007; Tobler et al., 2006; Tobler et al., 2008; Tobler et al., 2007; Tobler et al., 2009; Rolls et al., 2008) found that activity in regions including but not limited to frontal regions, encodes different aspects in value-based decision-making. An example neuroimaging study in humans measured BOLD signals across different regions and observed that the level of activity in ventral striatum and anterior cingulate is reflective of key decision parameters such as expected value and risk (Christopoulos et al., 2009). Electrophysiological studies in primates have found the neurons in the ACA encoded the identity and the subjective value of the offer chosen by the monkey (Cai and Padoa-Schioppa, 2012). Another study in primates conducted by Platt and Glimcher (1999), sought to assess whether neurons in the LIP form a map of the economic variables of reward magnitude and probability. In this study, they found that neurons in this area encoded these economic choice variables. These results opened further investigations into what regions of the brain encode different aspects of economic decisions. Early work studying ACA suggested that this region carries a representation of value, spanning from studies examining cells found in the dorsal and ventral bank of the cingulate cortex that were differentially modulated by delivery of a reward (Niki and Watanabe, 1979; Ito et al., 2003; Matsumoto et al., 2003; Williams et al., 2004; Seo and Lee, 2007; Amiez et al., 2006; Kennerley and Wallis, 2009a; Kennerley and Wallis, 2009b).

The OFC, another subregion of the frontal cortex implicated in processing information regarding rewards and punishment, has also been closely examined in a plethora of electrophysiological studies and neuroimaging studies. However such studies often postulate that the representation of values carried in this area is different from that of the ACA. In particular, these studies point to the role of OFC in encoding subjective value and motivational salience of different cues in the environment (Padoa-Schioppa et al., 2006; Roesch and Olson, 2004, 2005; Schoenbaum et al., 1998, 1999, 2003; Tremblay and Shultz, 1999; Wallis and Miller, 2003; Montague and Berns, 2002; Tremblay and Shultz, 1999). Focusing on signatures of subjective value, OFC neurons encode different choice-related variables (e.g. value of the offers, value of the choice, and the chosen cue) independent of visuomotor properties of the choices made (Padoa-Schioppa and Assad, 2006). Distinct from the value coding observed in primate OFC studies, later studies in rats found that neurons in the lateral OFC encode reward history and outcome thereby establishing the role of the OFC in adaptive behavior (Constantinople et al., 2019b). Independent of the pure valuation signal that most of these studies carry forward as the role of the OFC, the area has also been found to be important in the encoding of spatial-motor variables that are important for carrying out goal-directed actions (Feierstein et al., 2006). In other words, neurons in the OFC do not simply just encode for the economic value of the offer but also represent the spatial locations of the goals and the actions that are necessary to obtain such goals. These studies collectively suggest multiplexed roles for the OFC in not only representing spatiomotor variables but also in representing economic value independent of spatiomotor biases. Furthermore, these studies suggest that different signals in select brain regions may transform and influence one another to permit goal-directed behavior (Dickinson and Balleine, 1994).

These early electrophysiological studies have shown that neurons in different frontal regions such as the ACA and OFC encode variables such as the identity and the subjective value of the offers and the chosen offer respectively. Importantly, neurons such as OFC neurons in monkeys encode the subjective value independent of visuomotor properties of a task (space of goods). One of the main

questions that emerges from this is how neural circuits underlie the transformation of the chosen offer into a suitable action plan (Gold and Shadlen, 2007; Cisek and Thura, 2018). In primates, the OFC - a region where goods space decisions (Padoa-Schioppa et al., 2011) may take place - is not connected to motor systems (Carmichael and Price, 1995). The portion of the frontal cortex that is well connected to the OFC is the ventral portion of the lateral prefrontal cortex (LPFCv) (Petrides and Pandya, 2006; Saleem et al., 2014). In turn the LPFCv sends outputs to the dorsal portion of the lateral prefrontal cortex (LPFCd) (Takahara et al., 2012). Finally, the LPFCv connects strongly to motor regions (Lu et al., 1994; Takada et al., 2004; Takahara et al., 2012). Given this anatomical connectivity, studies have found that the transformation of the chosen offer into a suitable action plan takes place in the LPFCv and LPFCd (Cai and Padoa-Schioppa, 2014). Essentially, the neurons in the LPFCv and LPFCd encode the chosen value in the goods space and following the monkey taking an action, neurons in both these areas encode the spatial information and the upcoming action plan for the subsequent trial. Another region of the brain where this transformation of chosen offer (in goods space) into action space occurs is the supplementary eye field (SEF) of monkeys (Chen and Stuphorn, 2015). When analyzing individual neuronal activity in the SEF, it showed that neurons in this area first encode the chosen offer, and ~100 ms later the action that allows for obtaining the rewards. This study thus supports the view that in value-based decisions, the first step is the selection of the economic good followed by the selection of the action. Overall, these studies show how different frontal regions represent economic variables in order to enable efficient economic decisions and allow for transformation of information about the chosen offer from the goods space into an action space.

As evidenced by the plethora of electrophysiological studies, detailing the contributions of individual regions across frontal regions in the encoding of economic decision variables, it appears that individual brain regions appear to play specialized roles in the decision-making process. However, these observations are based on recordings from single neurons in isolated regions. Moreover, these studies often used different behavioral tasks, making it hard to directly compare

their results. Projection and anatomical connectivity studies remind us that these regions are not operating in isolation, rather are part of an intricate nexus that controls various aspects of the decision-making process. Recording from multiple brain regions is theoretically valuable, and technical advances have only recently made this practically feasible. Therefore, in order to understand the neural circuits underlying economic decisions, it is important to record from several regions simultaneously and in the context of a standardized behavioral task.

Overall, past studies have demonstrated representation of economic decision variables in various PFC regions. Although these studies provide insights into the role of PFC in economic decision-making, it is crucial to study the representations of economic variables across scales in a standardized task (i.e., the goal of the current study). Previous studies either measured neural signals with coarse spatial and temporal resolution (e.g. using fMRI) or from small neural populations using electrophysiology in non-human primates and rodent models. Therefore, it is unclear how neural populations across PFC encode economic decisions. The issue is the disconnect between coarse, large-scale, non-invasive recordings and detailed, but small-scale single-neuron measurements. The goal is to bridge this gap with large-scale electrophysiological recordings in the mouse brain, using a complex behavioral economic decision-making task. To address this, I recorded large-scale electrophysiological signals across different regions of the frontal cortex during the behavioral task described in the previous chapter.

3.2 Methods

3.2.1 Craniotomy surgery

Mice were trained in our visual economic decision-making task as described in Chapter 2. Electrophysiological recordings were performed when mice reached expert performance on this task, defined as displaying >80% accuracy across all three single stimulus trial types. The day prior to the first recording or the morning of the recording, the mouse was anesthetized using isoflurane and

surgically prepared following which a craniotomy surgery was carried out. Mice were anesthetized with 5% isoflurane and then positioned onto a stereotaxic frame prior to intraperitoneal injection of 5 mg/kg meloxicam and 0.1 mg/kg buprenorphine. Following these injections, mice were maintained on 1.5% isoflurane for the duration of the surgical procedure and kept on a heating pad throughout the surgery. In the previous headplate implantation surgery, the locations for bregma and lambda along with locations over both hemispheres for frontal regions were marked. For the Neuropixels recording experiments, craniotomies (approximately 1 mm diameter) were made over left and right frontal cortical regions using a dental drill (AP: 2.22, ML: +/- 1.0). One craniotomy was made for the hemisphere we intended to have the first acute Neuropixels recordings from. Following the drilling of the craniotomy, we covered this with a layer of Kwikcast (World Precision Instruments) in order to prevent the craniotomy from drying up. A 3D printed cone was then implanted onto the surface of the exposed skull in order to create a watertight well to hold saline. The mouse was left to recover in a post-recovery cage placed over a heating pad and provided with soft food, water, and jelly mixed with oral Metacam. The mouse was monitored until locomotion and grooming behavior were normal and the mouse was fully recovered.

After at least three hours or overnight recovery, mice were head-fixed in the mouse holder in the experimental rig for electrophysiological recordings. With each craniotomy surgery, the Neuropixels probe was inserted for 4-5 days into each hemisphere allowing for several recording sessions in different areas of each craniotomy. Once we completed 4-5 days of recordings from one hemisphere, we performed another craniotomy surgery in the opposite hemisphere, allowing for another ~5 recording days.

3.2.2 In vivo electrophysiological recording using Neuropixels probes

The neuronal recordings were performed using SpikeGLX software. Electrophysiological activity was sampled at a rate of 30 kHz. Recordings were completed using Neuropixels (Phase 3b) electrode

array which had 960 low impedance TiN electrodes. These electrodes were distributed in a dense staggered layout along a 10 mm long shank. The tip of the shank had 384 parallel dual band (AP, LFP) low noise recording channels. The Neuropixels probe was connected to a cable via the headstage, a small interface board that permits bidirectional information transfer. The connected Neuropixels probe is mounted onto a steel rod that is attached to a micromanipulator (uMp-4 micromanipulator, Sensapex Inc.). The micromanipulator was used to lower the probe into the brain tissue at a precise and slow rate. Prior to insertion, Neuropixels the probe tip was dipped into a red-fluorescent dye (Invitrogen, Vybrant Dil Cell-Labeling Solution, V22885), a lipophilic cell membrane stain that diffuses to the entire cell membrane, in order to coat the probe tip and to allow us to localize the probe tracks during later stages when histology procedures were carried out. The probe was dipped into the dye, pulled out allowing it to dry, and then repeating the insertion into the dye two more times prior to insertion of the probe into the brain. This allowed for the probe to be thoroughly covered by the dye, enabling accurate post-mortem histology-based probe reconstructions..

Probes were soldered to an external reference to a ground and the ground was attached to an Ag/AgCl wire which was then positioned inside the printed cone. The cone was then filled with sterile saline. The saline bath covering the surface of the exposed brain tissue and skull surface was refreshed throughout the recording to prevent drying of the saline during recordings lasting longer than 30-40 minutes. Following the coverage of the skull surface with the saline bath, the probe was lowered through the saline bath and subsequently advanced through the dura and slowly lowered to the final position at a rate of 10 μm/second. Once the probe reached the most ventral position, it was retracted 100 μm in order to allow for the tissue to settle. Following the retraction of the probe, the electrode settled for 15 minutes prior to commencing the neuronal recording session.

3.3.3 Histological probe localization

To verify the location of the probe, histological probe localization was conducted. Mice were anesthetized and perfused with 4% paraformaldehyde followed by extraction of the brain tissue followed by fixation for 24 hours at 4°C in paraformaldehyde. The brain was embedded in 2% agarose and sliced at a thickness of 70-μm. The sections were then washed in PBS, and stained with DAPI (Merck, MBD0015). Stained slices were washed with PBS before mounting onto glass adhesion slides. Images of the DAPI-stained slices were taken at a 4x magnification using a fluorescence microscope (Leica), in two colors: blue for DAPI and red for Dil. The Dil tract allowed for localization of a probe along several slices which were analyzed using ImageJ software.

3.3.4 Probe track tracing and localization

In order to identify the locations of each probe within the framework of the Allen Common Coordinate Framework, a 'sharp-track' MATLAB user interface (https://github.com/cortex-lab/allenCCF) was used. Due to the nature of the acute electrophysiological experiments involving multiple probe insertions in each hemisphere, the 3D locations of the probes were cross referenced with probe insertions images taken the day of the experiment. Using the extracted coordinates for each probe, brain region labels were assigned to clusters in the electrophysiological data using the IBL user interface (https://github.com/int-brain-lab/iblenv).

3.3.5 Data pre-processing and analysis

The neuronal data passed through two different curation steps: the first being automatic curation followed by manual neuron selection and processing. The automatic curation of the neuronal data was performed using Kilosort2 (https://github.com/jamesjun/Kilosort2). Following the curation using Kilosort2, the neuronal data was further manually curated using Phy GUI (https://github.com/cortex-lab/phy) (**Fig. 3.1**). Additionally, the neuronal data was subject to

common-average referencing, allowing for the removal of baseline offsets in addition to removal of artifacts in the data.

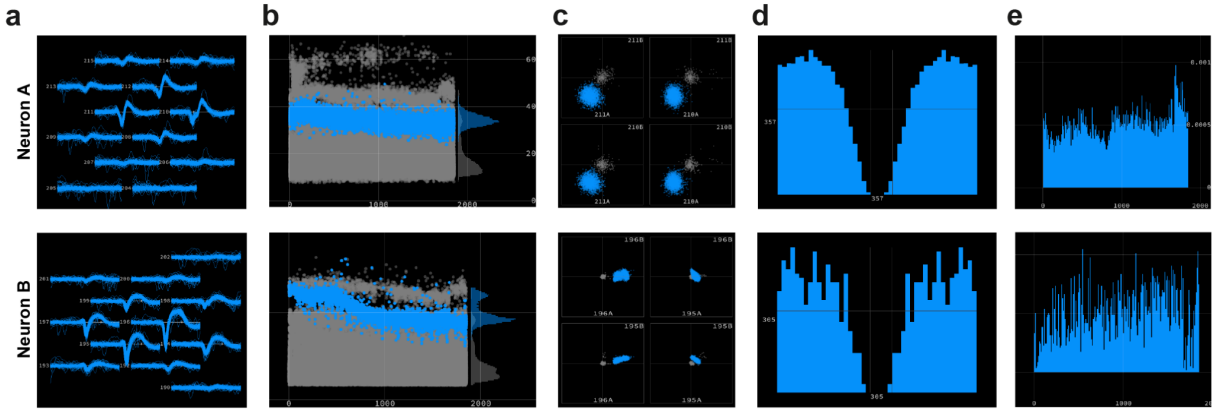


Fig. 3.1 | Example frontal neuronal clusters' parameters on Phy GUI

a, Example Neuron A and B's waveform view(top and bottom respectively) on the Phy graphical user interface. Phy is an open-source software available on Python that allows for manual curation of neuronal clusters in large electrophysiological datasets acquired from Neuropixels recordings. The 'Waveform' view shows different waveforms on channels that are close to the channel with the highest waveform amplitude. **b,** The 'Amplitude' view shows different spike amplitudes plotted as a function of time. In addition the plots show a vertical histogram on the right summing up the distribution of the amplitudes. **c,** The 'Feature' view shows the principal component features of the spikes for a particular neuron or selected clusters. In gray dots, one can find background spikes from all clusters. A good neuron will have a distinct cluster that is largely non-overlapping to the spikes from the selected neuron. **d,** In the 'Autocorrelogram' view it shows the autocorrelograms of selected clusters. In this view, the horizontal line represents the baseline firing rate. Each vertical line represents a refractory period. A good cluster will have a refractory period violation lower than 0.02. **e,** In the 'FiringRate' view, the local or instantaneous firing rate is displayed as a function of time. Normally for a 'good' cluster this firing rate view should have a sustained firing rate throughout the session.

During manual curation, each neuron was analyzed on the Phy GUI with a focus on the following parameters: waveforms, auto-correlogram and the number of spikes across time. A neuron labeled as 'good' was required to have at least one or two well shaped waveforms (**Fig. 3.1a**), and a narrow distribution of spike amplitudes, that was distinct from neighboring clusters (**Fig. 3.1b**). Furthermore, it was determined that spikes from a good neuron formed a distinct cluster. The 'Feature' view allows for determining whether the spikes from the selected neuron form a distinct cluster compared to all other background spikes from other neurons, when such spikes are projected onto their principal components (**Fig. 3.1c**). Next, the neuron's autocorrelogram which showed the refractory period of the neuron and the shape of the distribution, needed to have a

refractory period close to 0 probability (**Fig. 3.1d**). Neurons with a refractory period violation higher than 0.02 were labeled as 'multi-unit activity' as the center probability is too high to be considered a refractory period. In particular the distribution of the refractory period was crucial as a good cluster should have a clear "valley" or "rift" for bursting neurons. Neurons were labeled as 'noise' and discarded if the neuronal activity showed near-zero amplitude or waveforms that were non-physiological. A good cluster/neuron was required to have a minimum of 500 spikes over the course of the 30 to 60 minute recording. Additionally a low number of spikes usually results in waveforms that have noisy or jagged features and a higher number of overall spikes per cluster allows for a smoother waveform. Finally, neurons were excluded from the analysis if their average firing rate in the analysis window (stimulus onset to 0.4 seconds post stimulus onset) was less than 2 Hz. Additionally neuronal activity needed to be present for a majority of the recording session, with a presence ratio criterion of 90 percent being implemented (**Fig. 3.1e**). These criteria collectively ensured that subsequent analyses were focused on 'good' units.

In order to determine whether the spiking activity of a neuron was significantly modulated by different task events through the behavioral task, a series of statistical criteria were used to assess firing rate changes to different task events. The statistical tests to determine task-responsive clusters are as follows: (1) Wilcoxon sign-rank test between the firing rate of the entire trial window (average spiking rate for the entire trial window between 0 and 0.4 seconds relative to stimulus onset) and pre-trial baseline (-0.2 to 0 s relative to stimulus onset); (2) sign-rank test for the stimulus-evoked activity (period 0.05 to 0.15 s relative to the stimulus onset in each trial) and baseline rates; (3) sign-rank test for pre-choice activity (period -0.1 to 0.05 s relative to the choice onset in each trial) and and baseline rates; (4) sign-rank test for post-movement activity (period -0.05 to 0.2 s relative to the choice onset in each trial) and baseline rates; (5) sign-rank test for pre-reward window activity (period -0.15 to 0 s relative to the outcome onset in each trial) and baseline rates. A neuron was considered significantly modulated by the task if any of the statistical tests had P values that were below a Bonferroni-corrected alpha value of 0.01. In order to visualize a neuron's firing rate across

time, the activity of each individual neuron was binned at 0.01 second (counting spikes in 10 ms bins) and smoothed with a causal Gaussian filter with standard deviation of 0.025 s.

3.3.6 Kernel regression analysis

We used reduced rank kernel regression to regress out general stimulus- and action-related activity of each individual neuron in order to analyze condition-specific neural activity unrelated to general response time differences between conditions (Steinmetz et al., 2018). In the kernel regression model, the firing rate of each neuron is described as a linear sum of temporal filters that are aligned to different task events. For this study, we used stimulus and action kernels because we focused our analysis on the time period between stimulus onset and choice onset. The stimulus kernel covered a window from -0.05 to 0.35 ms relative to stimulus onset while the action kernel covered a window of -0.15 to 0.25 ms relative to choice onset. Stimulus and action kernels were agnostic to the presented condition (Large, Small, Gamble), in order to capture general stimulus- and action-related activity. Prior to kernel estimation, the firing rates of each neuron were estimated by binning the neurons' spiking activity into 10 ms bins and smoothing with a Gaussian filter (σ = 0.025s).

In order for each kernel to be fit, we constructed a Toeplitz predictor matrix. For each stimulus, the predictor matrix consists of the number of time points within the training time and the number of lags for the Action kernel. Each predictor matrix is made up of a set of diagonal stripes each of which is initiated whenever there is a specific stimulus being presented. Similar predictor matrices were constructed for the Action kernels. Once each individual matrix is constructed we horizontally concatenate the six Stimulus predictor matrices for six stimulus types (contralateral ‘Large’, contralateral ‘Small’, contralateral ‘Gamble’, ipsilateral ‘Large’, ipsilateral ‘Small’ and ipsilateral ‘Gamble’) and the two Action predictor matrices resulting in a global predictor matrix.

To estimate stimulus and action kernels, we used a reduced rank regression model. In the reduced rank model, the high dimensional kernel basis functions are projected onto a lower dimensional set of temporal basis functions that are optimized to predict the measured spiking activity of all neurons. For each neuron, we employed cross-validation to determine the optimal number of basis functions needed to capture the neurons' response profile. For this, we used elastic net regularization (using the package cvglmnet for Matlab (http://www.stanford.edu/~hastie/glmnet_matlab/), with parameters $\alpha = 0.5$ and $\lambda = 0.5$.

To estimate stimulus and action kernels, we used a reduced rank regression model. In the reduced rank model, this predictor matrix predicts the binned spikes for each individual neuron, after passing through a bottleneck which is a low-dimensional representation. Essentially the reduced rank model has learned a series of functions which are optimized to allow for prediction of the spike train. We use elastic net regularization to estimate a weight vector in order to minimize errors in fitting (using the package cvglmnet for Matlab (http://www.stanford.edu/~hastie/glmnet_matlab/), with parameters with parameters $\alpha = 0.5$ and $\lambda = 0.5$). Finally, to remove general stimulus- and action-related activity from the neural data, we computed the difference between the neural time series predicted by the cross-validated kernel regression model, and the empirically measured time series, thus yielding residuals. We then used these residuals to examine neuronal responses in a population decoding framework.

3.3.7 Linear population decoding

Following reduced rank kernel regression analysis, the residual firing rates resulting from the kernel regression analysis were used for population decoding of variables of interest. As a majority of our neurons were not recorded simultaneously, we constructed a 'pseudo session' by concatenating all neurons across sessions and counting them as part of one session for this analysis. Single stimulus trials were split into three: trials with 'Large', 'Small', and 'Gamble' stimuli. We only analyzed correct trials, in which the mouse successfully moved the stimulus to the center. In order for a

neuron to be included in this analysis, the session in which it was recorded had to have at least 30 correct trials for each condition. Furthermore, trials were balanced for left and right side stimulus presentations to avoid bias introduced by stimulus and action biases . Overall, sessions were combined by randomly subsampling trials from each session to form sets of 30 Large, Small and Gamble trials each, while also counterbalancing for left and right stimulus presentations. In order to alleviate the variability introduced by randomly subsampling trials and matching, we repeated this procedure 100 times and averaged the decoding accuracy across all 100 iterations.

3.3.7.1 Time-resolved decoding

We first determined a relevant time period relative to stimulus presentation during which single-stimulus conditions were decodable significantly above chance from the spiking activity of the entire population of recorded neurons. Time window identification was done by training a three-class linear Support Vector Machine (SVM), to classify the three condition labels (Large, Small and Gamble) based on the activity of all neurons. This classification was repeated for different time bins using binned residual spiking data (10 ms bins) ranging from -0.05 to 0.35 s after stimulus onset, using 10-fold cross-validation. A temporal generalization matrix was computed by training and testing the SVM classifier on each combination of training and testing time points (King et al., 2014). The diagonal in the matrix corresponds to the classifier being trained and tested on the same time sample. Before training and testing, we z-scored the firing rate of all neurons, separately for training and testing sets. Overall, performing the temporal generalization procedure enabled us to characterize the dynamics of the neural representations underlying the different single stimulus conditions. We next assessed the performance of the decoder and asked whether we can identify a time window during which the decoder performance is significantly above chance. In order to do so, we implemented a cluster-based permutation test (Maris et al., 2007) in which we shuffle trial labels for all three conditions and then repeat the temporal generalization decoding. In order to identify spatial clusters in the temporal generalization matrix, we noted the size of the largest spatial cluster in each of a 1,000 permutation iterations. These spatial clusters were defined as adjoining training

and testing bins whose decoding accuracy exceeded 5% above chance level decoding accuracy (33%). The size of each cluster was determined by taking the sum of the decoding accuracy across all training and testing bins in the cluster. Following this, the largest empirical cluster was then compared to the permutation distribution, and a p-value was calculated as the proportion of permutations that had a cluster that was equal to or larger than the empirical cluster.

3.3.7.2 Population and region-specific binary decoding

After identifying a relevant time window (150 ms to 350 ms post-stimulus onset) from the time-resolved decoding analysis, we then sought to perform binary population decoding and compare accuracies for conditions with different EV and risks. We asked whether we could find out which condition pairs were contributing the most to the strong decoding in that time window using a binary decoding pipeline. The goal of binary decoding of single stimulus conditions was to distinguish between three combinations of trials: different EV (Large and Small), same EV but different risk (Small and Gamble), different EV and different risk (Large and Gamble). Binary classifiers were implemented for the entire pooled data from all recorded neurons and separately for neurons recorded in individual regions. Binary classifiers were trained and tested on neural activity that was averaged across the entire time window from 150 to 350 ms post stimulus onset - the relevant time window during which three-class decoding accuracy was highest. To test whether all trial pairs contributed to the significant three-class decoding, we compared the decoding accuracy of each binary classifier against chance level (50%) by using a permutation test. The permutation test required trial labels to be randomly shuffled within each binary comparison for a total of a 1,000 permutations, each shuffled permutation followed up with a decoding analysis. At the end of running each shuffled decoding analysis, we noted the decoding accuracy from each permutation. The distribution of these permutation-derived decoding accuracies was then compared to the empirical decoding accuracy. The P-value was calculated as the proportion of permutations where the decoding accuracy was equal to or greater than the empirical accuracy, with the critical p-value being 0.05. We next asked whether the different decoding accuracy for each binary comparison are

statistically significant and different from each other (i.e. decoding accuracy for Large versus Small compared to the decoding accuracy for Small versus Gamble). We used a similar permutation test as detailed before but this time shuffling trial labels between condition pairs. For example, if we are testing the comparison of Large versus Small and Small versus Gamble - we randomly shuffled trials labels for Large and Gamble and the decoding procedure was conducted for each of a 1,000 shuffled permutations. We then compared the proportion of trials with a difference in decoding accuracy between different condition pairs that is equal to or larger than the empirical difference in decoding accuracy. This time the critical p-value used was 0.025 (two-sided test). Other binary classifications done using neural activity from single stimulus trials were decoding the (a) stimulus side (left or right side conditions) and (b) the economic value (Large, Small, Gamble) of ipsilateral and contralateral stimuli.

We further assessed whether neural decoding of risk (Small vs. Gamble classification) depended on mice's behavioral risk attitudes. For mice that performed two-stimulus trials, we calculated risk sensitivity as the absolute difference in choice preference for either Gamble or Small stimulus to the indifference point (50%). We then split the group of mice into risk-neutral and risk-sensitive subgroups by means of a median split on risk sensitivity. We then repeated risk decoding (Small vs. Gamble classification) separately for each subgroup.

We measured the decoding accuracy for neurons from individual regions, asking which brain regions are contributing most to the binary decoding of different EV conditions or different risk condition pairs. Since the number of neurons collected from each region varied, we repeated the 2-class SVM analysis using sub-sampled neural populations in steps of 50 neurons, increasing to 800 for some regions. This allowed us to compare the decoding accuracy for different regions at similar population sizes. The subsampling of neurons was repeated a 100 times in order to average out the variability introduced by randomly subsampling neurons. To test whether the performance of the

decoder for each region was above chance level, a permutation test was conducted by shuffling the trial labels in each decoding comparison and repeating to decoding.

For our main analyses, we pooled data across all recording sessions. To assess whether regional differences in decoding accuracy could be explained by behavioral variability across sessions targeting different brain regions, we conducted a control analysis. In this control analysis, we ran our decoding pipeline on simultaneously recorded pairs of brain regions, thereby precluding differences in behavior. In particular, for a given pair of simultaneously recorded brain regions, we performed two-class decoding of EV (Large v.s. Small trials) and risk (Small v.s. Gamble trials) and recorded the difference in decoding accuracy between regions. We then compared this difference in decoding accuracy to that computed using all sessions, while matching the number of neurons.

3.3.7.3 Spatial binary decoding

We next investigated the spatial structure of these neural representations of economic variables beyond brain region labels, by examining the spatial distribution of activation weights across three axes: mediolateral (ML), dorsoventral (DV) and anterior-posterior (AP) axis. For the main analysis we focused on the spatial distribution of activation weights across the ML and DV axes as we had little variance in the AP axis. We first extracted the neuronal activation weights for EV and risk from the binary classifiers that distinguished Large versus Small (different EV conditions) and Small versus Gamble (different risk) respectively. These activation weights were calculated by multiplying the SVM classifier weights by the data covariance matrix (Haufe et al., 2014). We used these derived activation weights for each neuron and mapped it using the spatial coordinates ascribed to each neuron. Then neurons were grouped into bins located in the ML and DV space with these parameters for bin size: ML bin size: 0.1 mm and DV bin size = 0.2 mm. Each bin was required to have more than 50 neurons to allow for robust observations of spatially localized neural representations. The activation weights in each bin were then averaged in order to determine the mean absolute activation weight within each bin. These mean absolute activation weights were

plotted across both axes generating an activation map. In order to test our hypotheses that these representations of economic variables displays a functional gradient across different axes, we characterized the change in activation weights across both the ML and DV axes by fitting a linear regression model to the spatial distribution of activation weights along ML and or DV axes and extracting the slope and intercept coefficients. We conducted statistical tests to test if the extracted slopes were significantly different from zero. To do so, we conducted a permutation test, shuffling spatial coordinates and repeating the linear regression in ML and DV axes. We compared the empirical slope estimates to those of the permutation data. The p-value was calculated as the proportion of permutations exceeding the empirical observed slope. The critical p-value was 0.025 (two-sided test). To compare spatial gradients between EV and risk decoders, we conducted a permutation test shuffling activation weights between EV and risk comparisons. A p-value was calculated as the proportion of permutations exhibiting a more extreme difference in EV and risk gradients than the empirically observed difference. The critical p-value was 0.025 (two-sided test).

3.3.7.4 Population and region-specific binary decoding of chosen and unchosen value

In a smaller subset of mice, we introduced choices between two simultaneously presented stimuli (two-stimulus trials). We next asked whether we could reliably decode the chosen and unchosen value in these two-stimulus trials. To do so, we used the time window of a 150 to 350 ms post stimulus onset and trained binary classifiers. For chosen value decoding, a classifier was trained to distinguish between different chosen values, while holding the unchosen value constant (chose 'Large' over 'Gamble' versus chose 'Small' over 'Gamble'). For unchosen value decoding, a classifier was trained to distinguish between different unchosen values, while holding the chosen value constant (chose 'Gamble' over 'Large' versus chose 'Gamble' over 'Small'). Due to the number of two stimuli trials within a session being fewer than single stimulus conditions, sessions were combined by randomly subsampling each session's trials to form sets of 15 trials in each condition ("chose Large over Gamble", "chose Small over Gamble" and vice versa for the unchosen

conditions). Binary classifiers distinguishing between these chosen and unchosen value conditions were implemented both for pooled neurons across all regions as well separately for neurons recorded in two regions with the largest number of recorded neurons: OFC and MOs.

3.3.8 Single neuronal analysis

We measured and visualized individual neuronal responses aligned to different task events by calculating the average residual firing rate for each neuron in 200 ms windows relative to stimulus presentation and the onset of an action. This mean residual task-event aligned firing rate was further split into different single stimuli (Large, Small and or Gamble) or two stimuli conditions (Large-Small, Large-Gamble, Small-Gamble). In order to identify neurons whose firing rate was significantly different among different pairs of conditions, we conducted a Wilcoxon rank sum test across different pairs of conditions within the windows. Following this z-tests were implemented to test whether the proportion of neurons that are modulated by different trial conditions were significantly different across neuronal populations. The spiking rate across time and aligned to important task events for individual neurons was visualized by binning the spiking activity into 10 ms bins and smoothing this with a Gaussian filter. The final results were then plotted as a peri-stimulus time histogram.

3.3 Results

3.3.1 Frontal neurons respond to the task exhibit task responsiveness and encode economic value

To examine neuronal representations during our economic decision making task, we recorded the activity of thousands of neurons (n = 8,611 neurons, 134 sessions, n = 27 mice, see Methods for details on neuronal pre-processing) across the frontal cortex using Neuropixels probes (**Fig. 3.2 a-d**). For this portion of our neuronal data analysis we exclusively focused on trials with single

stimuli, for which we observed robust behavioral manifestations of EV (**Chapter 2, Fig. 2e**) in a large cohort of mice (n = 34 mice), and we have a larger dataset of neurons with activity during these single stimulus trials. We separately analyzed neural activity in two-stimulus trials, which were introduced in a smaller cohort of mice (n = 13). The results for two-stimulus trials are presented further below. A large fraction (6,982/8,611) of neurons across different frontal brain regions responded during the task, increasing or decreasing their activity from the onset of visual stimulus until the trial outcome time (**Fig. 3.2 d-e**; correct trials only). This task-related modulation of neuronal activity was not specific to frontal cortical neurons; the activity of a large fraction of neurons recorded in olfactory regions ventral to frontal cortex were also modulated during the task.

Frontal neuron responses distinguished stimuli from each other. To examine whether the firing rate of individual neurons exhibits differential responses to different stimuli, we first inspected the activity of individual neurons associated with different single stimulus conditions. Individual neurons responded strongly at the time of stimuli, with different firing rates for different stimuli (**Fig. 3.2f**). For instance, the responses of an example neuron scaled with the EV of the stimulus, showing a larger response in trials with a stimulus with high EV, and a smaller response when encountering stimuli with low EV (**Fig. 3.2f, left**). In contrast, the responses of another example neuron were comparable for stimuli with different EVs but were scaled by the risk that the stimuli were entailing (**Fig. 3.2f, right**). Examining responses of neurons at the single cell level, showed that overall a subset of neurons displayed strong stimulus-driven responses, with 28.5% of neurons differentially modulated by different single stimulus conditions. These results show that during economic decision making neurons across the frontal cortex strongly respond to task variables.

Having observed that individual neurons exhibit distinct responses to stimuli in the task, we next asked whether the population of neurons responded differently to different visual stimuli. To do so, we first used a reduced rank kernel regression to regress out the general stimulus- and action-related activity of each individual neuron (**Fig. 3.3, 3.4**). This is necessary because in our task animals had

different response times (RTs) for different stimuli, and as such neuronal responses might reflect these RTs, rather than the reward EV or risk. To this end, we fit each neuron's activity with general stimulus and action kernels, independent of the trial condition (Large, Small, Gamble). In the regression analysis for the example neuron, we show a stimulus and action kernel used to predict the neuron's activity (**Fig. 3.3a**). The mean firing rate of the example neuron overlaid with a cross-validated prediction of the regression model using both stimulus and action kernels, shows a good fit (**Fig. 3.3b**). To regress out the general stimulus- and action-related activity, we then subtracted the predictions of the kernel regression model from the raw empirical data. We then used the residuals of this regression to examine whether neuronal responses represent the economic value of stimuli. The residuals represent the neuron's responses that are not explained by the general stimulus and action kernels (**Fig. 3.3c**).

Another line of evidence to ensure that the neuronal responses are reflective of economic value rather than RTs, involves decoding the stimulus condition using the RTs (**Fig. 3.4**). To do so, we used a logistic regression model to predict stimulus conditions based on RTs, separately for each session, and computed RT-based decoding accuracies using 10-fold cross validation. We find that the RT decoding yielded an accuracy that is slightly above chance level, however, it is far below the accuracy derived from the neural decoding (**Fig. 3.4a**). Next, we asked whether removing the action kernel from the neural data impacts the decoding accuracy. To do so, we repeat the 2-class SVM to classify two stimuli from each other and input either the raw neural data and or the residual data. We found that removing the Action kernel leads to a slight decrease in decoding accuracy across all three sets of condition pairs (**Fig. 3.4b**). Having found that decoding using RTs leads to much lower accuracy compared to neural decoding and that removing the Action kernel only leads to a slight decrease in decoding accuracy, we next asked whether the kernel regression is removing RT signals from the neural activity. To do so, we compared the relationship between the session-by-session RT decoding accuracy and SVM weights derived from classifiers trained on the raw or residualized neural data. If RT-related signals contributed to the neural decoding, neurons

recorded in sessions with higher RT-based decoding accuracy should exhibit higher SVM weights during neural decoding. Further, if removing the Action kernel from the neural data removed RT-based signals, such a correlation between RT decoding accuracy and SVM weights should be weakened. Although RT-based decoding accuracy was only weakly correlated with neural SVM weights, removing the action kernels from the neural data indeed further decreased this correlation between the classifier weights and the RT-based decoding accuracy (**Fig. 3.4c**). This suggests that the kernel regression is partially successful in removing RT-related artifacts or movement related signals from the neural data.

We subsequently employed a linear population decoding technique, Support Vector Machines (SVM), to test how well we can distinguish the three different stimulus conditions from the neural activity of all recorded neurons (3-class decoding corresponding to three stimuli, see Methods). We determined a specific time window of the trial during which single stimulus conditions were decodable from the entire population of recorded neurons, using a temporal generalization approach. This temporal generalization approach revealed a crucial time window, 150 to 350 ms post stimulus onset, during which the stimuli were decodable above chance (**Fig. 3.2g, 1,000 permutations of shuffled trial labels, $P < 0.001$**).

To examine whether this decoding is reflective of the economic value (rather than stimulus identity), we asked whether the decoding accuracy is sensitive to the value of stimuli. We reasoned that if neural responses encoded the value signaled by visual stimuli, neural activity patterns should be more distinguishable, and decoding accuracy higher, when the expected value differed rather than when it was equal between two sets of trials. To examine this hypothesis, we used a 2-class SVM to classify two stimuli from each other using the previously identified time window of neural activity. Indeed, the SVM's decoding accuracy was sensitive to economic value; it was higher when classifying stimuli with different EVs rather than stimuli with the same EV (**Fig. 3.2h**). Moreover, the accuracy of the classifier among stimuli with different risk but the same EV depended on the

animal's risk attitude; it was higher in risk-sensitive mice (i.e. both risk-seeking and risk-averse), compared to risk-neutral mice (**Fig. 3.2h**), providing further support for frontal population neuronal responses being reflective of economic value. Thus, consistent with several past studies, frontal population neuronal responses reflected animals' preference between stimuli.

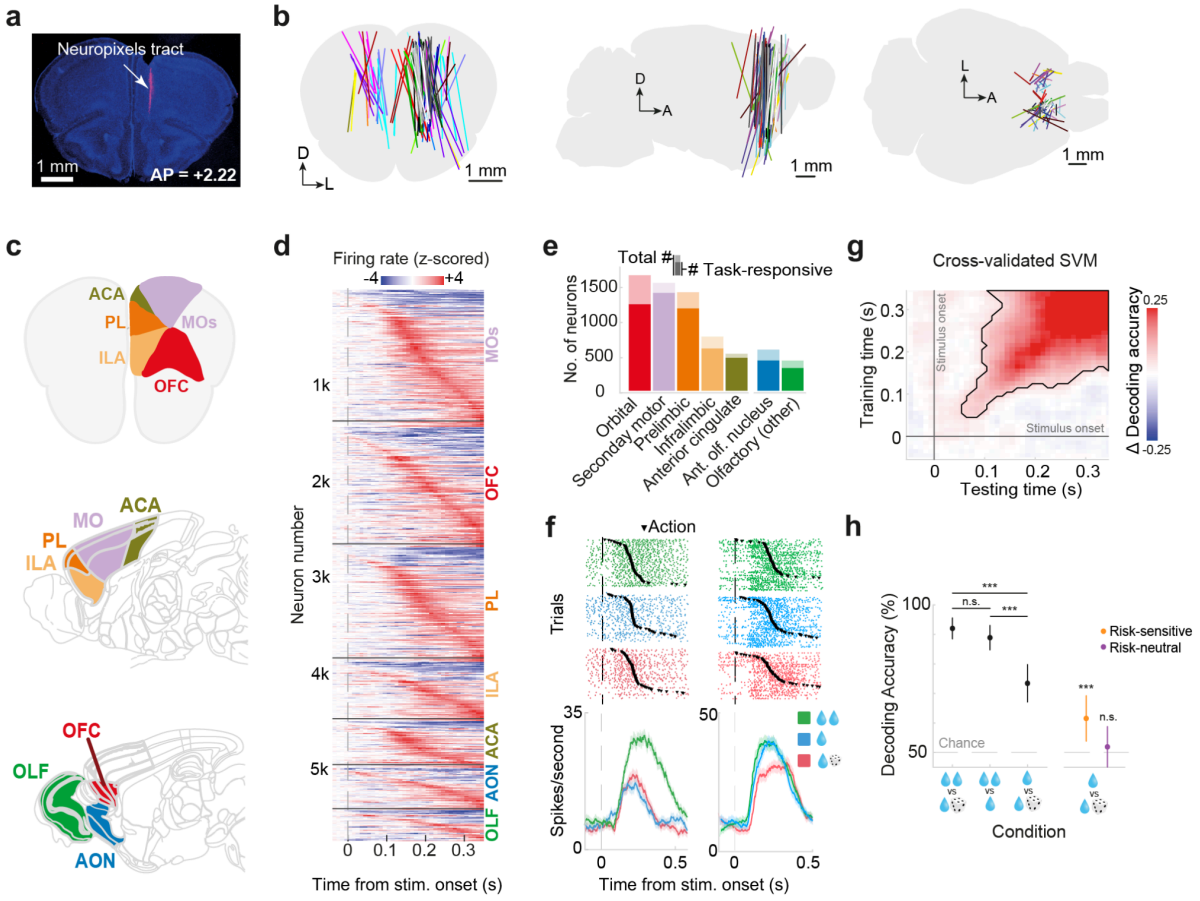


Fig. 3.2 | Frontal neuronal responses reflect the economic value of stimuli

a, Image of a coronal frontal section of a mouse brain with a probe tract shown in red (Dil). **b,** Reconstructed probe tract locations aligned to the Allen Common Coordinate Framework 3D space shown in a coronal section (left), sagittal section (middle) and top down view (right). Each group of colored lines corresponds to various probe insertions across different frontal areas throughout several acute recordings. **c,** Brain regions from which electrophysiology data was primarily collected from, shown in a coronal section (top) and two sagittal sections (bottom two). **d,** Average normalized (z-scored) responses of task responsive neurons from five major frontal areas and olfactory areas, aligned to stimulus onset. Each row represents one neuron. **e,** Overview of recorded neurons and proportion of neurons identified as task responsive. **f,** PSTH and raster plots for two example neurons. Each row in the raster plots show spiking activity in one trial. Black inverted triangles denote the time of choice onset (i.e. wheel movement to indicate a choice). The PSTH show firing rate split by stimulus conditions. Shaded regions show the s.e.m. across trials. **g,** Temporal

generalization matrix representing the decoding performance of a SVM classifier at different time points. The matrix shows training time on the y-axis and testing time on the x-axis with zero denoting the time of stimulus onset. Decoding accuracy is significantly exceeding chance level in a time period 150 to 350 ms post stimulus onset (black line outlines significant cluster). **h,** Binary decoding of stimulus conditions. Decoding accuracy was significantly reduced when classifying trials with the same compared to different expected values, indicating that neural representations were sensitive to value. Moreover, the accuracy of the decoding between two stimuli with the same EV but different risks depended on the animals' risk attitude, being higher in risk-sensitive animals compared to risk-neutral animals.

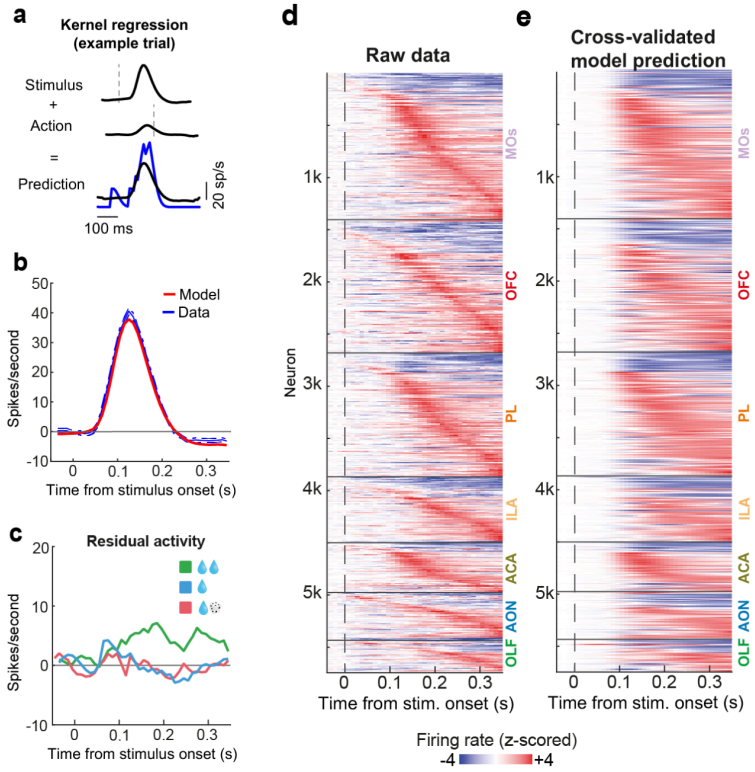


Fig. 3.3 | Residuals from the reduced rank kernel regression method show stimulus-locked responses

a, An example stimulus and action kernel being used to predict the example neuron's spiking activity. Bottom panel shows an example fit of spiking data using these two different kernels (stimulus and action). Blue trace shows the neuron's raw spiking data smoothed with a causal filter and the black line denotes the prediction of the model. **b,** The average firing rate of the example neuron with an overlay of the model's prediction. **c,** Subtracting the predictions of the kernel regression model from the raw empirical data removed general stimulus- and action-related activity, while maintaining differences in neural activity across stimulus conditions. **d,** All recorded neurons from the dataset with z-scored firing rate and their responses aligned to stimulus onset. Neurons are grouped according to brain region. All neurons exhibit robust stimulus-aligned responses with variable delays. **e,** All recorded neurons from the dataset with the cross-validated model prediction residuals plotted aligned to stimulus onset. Neurons are grouped according to brain region.

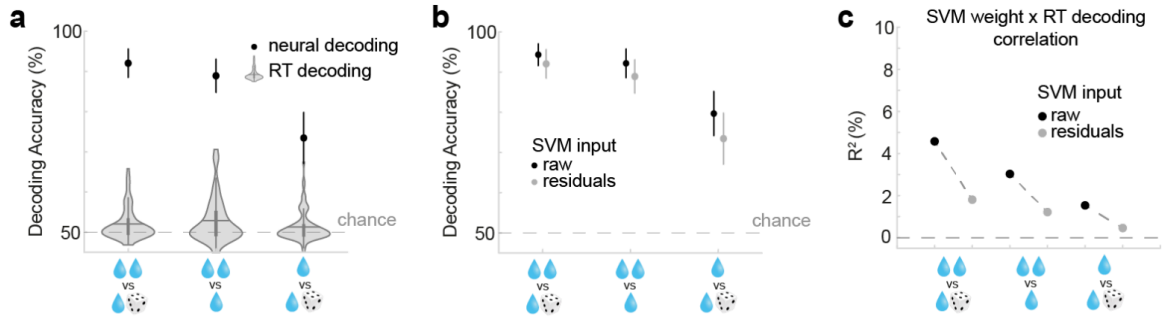


Fig. 3.4 | Kernel regression partially removes reaction-time related signals

a, Binary decoding of single stimulus conditions from response times (RTs). Decoding accuracy was slightly above chance level for classifying trials with same and different expected values but considerably lower compared to decoding accuracy provided by the neural activity. RT-based decoding accuracy was computed for each session, and violin plots show the distribution across sessions. **b,** Binary decoding of single stimulus conditions from raw data versus residuals derived from removing action kernels from the raw neural activity. Decoding accuracy was slightly reduced for SVMs based on the residuals (action kernel removed from the raw neural data). **c,** Comparison of the coefficient of determination (R^2) from a linear regression examining the relationship between either the SVM weight from the raw neural data and RT decoding accuracy or the relationship between the residuals SVM weight and the RT decoding accuracy. The R^2 was higher for the raw SVM weights versus the RT decoding accuracy compared to the residual SVM weights versus the RT decoding accuracy. This suggests that the kernel regression analysis was successful for partially removing RT-related signals.

3.3.2 Graded representation of economic value across regions of frontal cortex

Our previous analyses utilized neurons across all recorded brain regions. Focusing on the decodable time period (0.15 to 0.35s post-stimulus), we then investigated which brain areas were contributing to the decoding by classifying the three single-stimulus trials ('Large', 'Small' and 'Gamble') based on the spiking activity of individual areas (**Fig. 3.5a**). Cross-validated decoding accuracy was particularly high in medial prefrontal regions (ACA, PL), followed by secondary motor cortex and orbitofrontal regions (**Fig. 3.5a**). In contrast, decoding from olfactory regions was much lower compared to decoding from other frontal regions. One main question we ask using this analysis is whether one can distinguish two stimuli better than chance level (33%). We answer this using a permutation test where we shuffle the two labels across the conditions and then repeat the decoding.

The null hypothesis is that when one shuffles the trial labels, this should not make a difference in the decoding accuracy. We compare the empirical difference to the shuffled distribution and find that it is outside of the shuffled distribution allowing for statistical evidence that different regions are significantly decoding the three single-stimulus trials above chance level. This being said, it is unclear how a lower decoding accuracy translates to behavioral preference and it can be said that while there is information to decode, the information may appear to be too weak to influence the behavior of the mouse. While the decoding accuracy was statistically significant from chance across several brain areas, a smaller decoding accuracy metric is probably due to the small amount of data that I have and it is unclear how this relates to behavioural performance.

Since we constructed a 'pseudo-mega session' by concatenating all neurons across sessions and counting them as part of one session for this analysis, we next asked whether the strength of representations of economic value across regions is maintained in sessions where neurons are simultaneously recorded from. Indeed, the difference in decoding accuracy observed in the 'pseudo-session' was maintained in recording sessions where pairs of regions are recorded simultaneously (**Fig. 3.5 b-d**). We look at the difference in decoding between areas to see whether there is worse decoding in one area compared to another brain area. Since this is a difference in decoding accuracy between regions, there is no reason to assume that with a lower number of neurons that the difference in decoding accuracy should sway in one way or another. The variance could increase but not the difference in decoding accuracy. Therefore, our control analyses revealed that the difference in the strength of representations between areas is largely maintained.

We next used SVM classifiers to investigate which brain areas were contributing to the decoding of EV. This brain-region specific analysis was completed by classifying single-stimulus trials with different reward magnitude but same probability ('Large' and 'Small' single stimulus trials) based on the spiking activity of individual areas. Cross-validated decoding accuracy was particularly high in medial prefrontal areas (anterior cingulate cortex, ACA) followed by secondary nmotor cortex

(MOs), prelimbic cortex (PL) and orbital regions (OLF; **Fig. 3.6a**). Unlike prefrontal regions where decoding accuracies for EV coding were high, decoding from olfactory regions was poor, suggesting that value was not strongly encoded in these olfactory regions. Overall, decoding accuracy was particularly high in medial prefrontal regions such as anterior cingulate and prelimbic cortex, compared with more lateral regions such as orbitofrontal cortex (**Fig. 3.6a**).

Having established that population frontal neuronal responses reflect the value of stimuli, we next asked how these representations of different economic variables (i.e. EV and risk) are structured across the frontal cortex. To do so, we repeated our two-class SVM analysis on neurons recorded in different frontal regions. Given that we recorded different numbers of neurons in different frontal regions, we performed this analysis on a matched number of neurons across regions. The decoding accuracy for stimuli with different EV but the same risk was significantly higher than chance in many regions of the frontal cortex (**Fig. 3.6a**). Similarly, the decoding accuracy for stimuli with different risk but the same EV was significantly higher than chance in many regions (**Fig. 3.6b**). For this analysis, we sought to classify single-stimulus trials with a difference in risk ('Small' and 'Gamble' trials). Risk decoding accuracy was particularly high in medial frontal regions (ACA, PL and ILA) and was notably decreased in MOs (**Fig. 3.6b**). The low risk decoding accuracy in dorsally-located MOs stood in stark contrast to EV decoding, in which MOs was among the regions with highest decoding accuracy. Conversely, the more ventrally located area ILA provided comparatively good risk decoding, but notably low decoding accuracy for EV. These neuronal representations of EV and risk showed both striking similarities and differences. For instance, ACA showed similarly high accuracy for both variables whereas MOs exhibited high accuracy for EV but not risk (**Fig. 3.6a,b**). The results thus show different representations of economic variables across different regions of the frontal cortex. These observations prompted us to investigate the spatial structure of EV and risk encoding beyond individual brain regions.

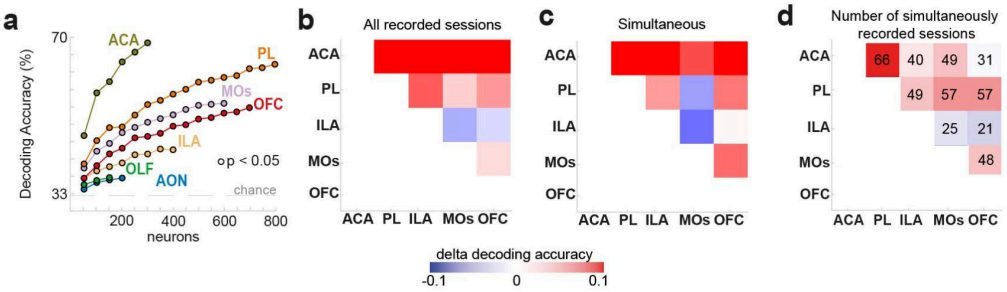


Fig. 3.5 | Strength of representations are maintained in simultaneously recorded sessions

a, Decoding of three single stimulus conditions (Large, Small and Gamble) from individual prefrontal cortical regions and olfactory regions. While decoding accuracy is significantly above chance for all regions, decoding accuracy is much higher for frontal cortical regions such as ACA and PL compared to olfactory regions. **b,** A matrix showing the difference in decoding accuracy between two regions for the 3-class SVM (decoding single stimulus trial condition). Region pairs were matched with the minimum number of neurons among each pair. This matrix is from all recorded sessions where all neurons are combined into one pseudo session across individual sessions. **c,** A matrix with paired regions and the difference in decoding accuracy from region pairs during recording sessions in which two regions are simultaneously recorded from. Population size was matched between the two regions followed by taking the difference in the decoding difference at the highest population size. **d,** A matrix with paired regions on the axes indicating regions recorded from simultaneously. The number indicates the number of sessions in which region pairs were recorded from.

3.3.3 Neuronal responses exhibit distinct spatial gradients for expected value and risk

We examined the spatial structure of economic variable representations representations across the frontal cortex, asking whether there is a different spatial structure or gradient for the encoding of difference in EV ('Large' and 'Small') and the difference in risk ('Small' and 'Gamble'). In order to examine the spatial structure we grouped neurons into spatial bins across the three axes of the frontal cortex - medial-lateral (ML), and dorsal-ventral (DV) (**Fig. 3.6 c,d**). We generated 2D spatial maps of the absolute difference in activation patterns (Haufe et al., 2014) between stimuli associated with high and low EV or high and low risk along different axes of the frontal cortex (see Materials and Methods).

EV coding was most robust in the dorsal- but not ventral - regions of the frontal cortex therefore exhibiting a dorsal-to-ventral gradient (**Fig. 3.6c**). Alongside the gradient of EV coding, we examined the representation of risk by contrasting the activation patterns to stimuli with the same EV but different risk. In contrast to the representation of EV, the representation of risk exhibited a medial-to-lateral gradient, with risk coding most strongly observed in the medial - but not lateral - regions (**Fig. 3.6d**). Averaging the activation weights across the DV axis revealed that the EV coding was significantly reduced in the dorsal-to-ventral direction (**Fig. 3.6e; permutation test of linear slope against zero, $p < 0.001$**), while risk did not (**Fig. 3.6e; permutation test of linear slope against zero, $p = 0.086$**). In contrast, averaging the activation weights across the ML axis, revealed that risk coding was significantly reduced in the medial-to-lateral direction (**Fig. 3.6f; permutation test of linear slope against zero, $p = 0.012$**), while EV did not show a change in medial-to-lateral direction (**Fig. 3.6f; permutation test of linear slope against zero, $p = 0.69$**). These activation gradients were significantly different between EV and risk in both axes (**Fig. 3.6 e-f; DV difference - permutation test, $p < 0.00001$; ML difference - permutation test, $p < 0.016$**). Moreover, these spatial gradients were similarly expressed along the posterior-to-anterior axes that we examined (**Fig. 3.6 c,d**). Overall, we demonstrate that representations of different economic variables exhibit distinct spatial maps across the frontal cortex.

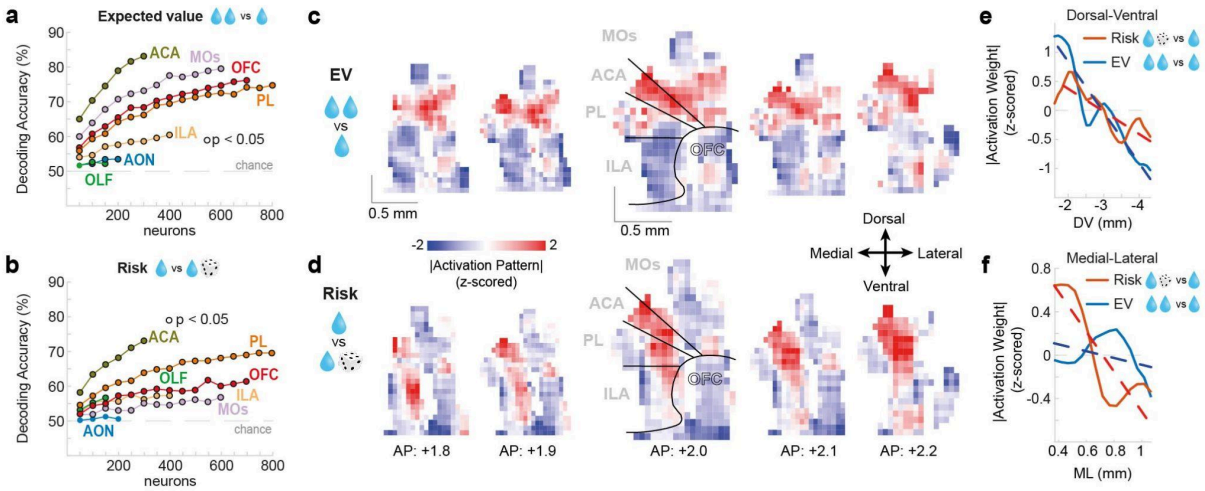


Fig. 3.6 | Representations of economic variables across frontal cortex exhibit distinct spatial gradients

a, Decoding of single stimulus conditions with different EV but same risk from individual prefrontal cortical regions. **b,** Decoding of single stimulus conditions with the same EV but different risk from individual prefrontal cortical regions. **c,** Spatial decoding plot showing absolute activation weights averaged in spatial bins with an atlas image overlay. Distinct dorsal-to-ventral gradient for expected value representations was observed and maintained across the anterior-to-posterior axes. **d.** Spatial decoding plots for classifying conditions with the same EV but different risk. Higher activation weights were observed in medial compared to lateral regions and maintained across the anterior-to-posterior axes. **e,** Activation weights (z-scored) across the dorsal-ventral axes split by risk and EV coding. The dotted line represents a linear regression model capturing the change in the activation weights along the DV axis. **f,** Same as in e for the ML axis.

3.3.4 Frontal neuronal responses display stronger contralateral representation of economic variables

Having observed that frontal neurons encode economic variables - we next asked whether this encoding is occurring in a lateralized manner. We further characterized the representations of EV and risk across the frontal cortex, asking whether there is a stronger response to contralateral visual stimuli compared to responses to ipsilateral stimuli. To do so, we first examined individual neuron stimulus-aligned firing rates split by contralateral and ipsilateral trials (relative to recording hemisphere). Interestingly, we found that individual neurons exhibited strong stimulus-aligned responses to contralateral visual stimuli. For example, the neuron displayed in **Fig. 3.7a** shows higher firing rates for all three visual stimuli when presented in the hemifield contralateral to the

recording site, and a smaller response when encountering visual stimuli in the same hemisphere as the recording site (**Fig. 3.7a**). This example demonstrates that during economic decision making, individual neurons across the frontal cortex respond strongly to economic variables but in a lateralized manner.

Having observed distinct lateralized responses to stimuli in the task in the response profiles of example individual neurons, we next asked whether the population of neurons responded differently to different visual stimuli. To investigate whether lateralized stimulus responses were prominent across PFC, we turned to population level analyses. To do so, we first used a 2-class SVM to classify the spatial location of stimuli (ipsilateral or contralateral to the recorded hemisphere) independent of their associated economic value. Specifically, we investigated which brain areas were contributing to the decoding of spatial location by classifying ipsilateral or contralateral trials utilizing the spiking activity across different brain regions in a 2-class SVM analysis (**Fig. 3.7b**). Cross-validated decoding accuracy was particularly high in dorsal prefrontal regions (ACA, MOs) followed by ventral prefrontal regions (PL, OFC, ILA). Decoding from olfactory regions was slightly lower compared to decoding from other frontal regions.

Having established that neurons strongly differentiated the spatial location of the stimulus, we next asked whether the representation of economic value was sensitive to spatial location. First, when examining decoding accuracies for contralateral conditions, we find that all three pairs of conditions have accuracies that are overall high and significantly higher than chance level (**Fig. 3.7c; binary decoding accuracy for contralateral trials ['Large' and 'Small']: 0.894 ± 0.0627, $p < 0.0001$; ['Large' and 'Gamble': 0.945 ± 0.05, $p < 0.0001$; ['Small' and 'Gamble': 0.734 ± 0.079, $p = 0.0060$; mean $\pm$ SEM, two-sided paired t-test**). Next, we find that the SVM accuracy for two of three ipsilateral condition pairs is significantly above chance, notably, those two ipsilateral condition pairs that have a difference in EV but same probability (**Fig. 3.7d; binary decoding accuracy for ipsilateral trials ['Large' and 'Small']: 0.696 ± 0.098, $p = 0.029$; ['Large' and**

'Gamble']: 0.716 ± 0.092, p = 0.006; ['Small' and 'Gamble']: 0.477 ± 0.101, p = 0.1350; mean ± SEM, two-sided paired t-test). Overall the SVM's decoding accuracy was much higher for classifying contralateral conditions compared to ipsilateral conditions, providing support for frontal neuronal responses being reflective of economic value but in a lateralized fashion. Thus, frontal population neuronal responses reflected economic value in a lateralized manner. These observations prompted us to investigate the spatial structure of EV and risk encoding for contralateral and ipsilateral stimuli beyond individual brain regions.

Stronger spatial representations of EV and risk were observed in contralateral trials compared to ipsilateral. We analyzed the spatial structure of the representations of these economic variables across frontal cortical regions separately for ipsilateral and contralateral trials, asking whether these gradients were maintained in both contra- and ipsilateral trials. We first analyzed the different spatial structure or gradient for the encoding of difference in EV ('Large' and 'Small') and the difference in risk ('Small' and 'Gamble') in contralateral trials (**Fig. 3.7e**). Similar to previous spatial structure analysis, neurons were grouped into spatial bins across the three axes - medial-lateral (ML), dorsal-ventral (DV) and anterior-posterior (AP) - of the frontal cortex (**Fig. 3.7e,f**). Finally, 2D spatial maps of the absolute difference in activation patterns between stimuli associated with high and low EV were generated. Among contralateral trials, the EV coding was robust in dorsal- but not ventral- space in the frontal cortex, leading to a dorso-to-ventral gradient (**Fig. 3.7e, left**). Alongside the gradient of EV coding, we analyzed the representation of risk in contralateral trials by contrasting the activation patterns to stimuli with different risks and the same EV. We observed that representations of risk were robust in the medial areas compared to more lateral areas. Therefore, risk coding in contralateral trials displayed a distinct medial-to-lateral gradient (**Fig. 3.7e, right**).

Having discovered distinct spatial gradients for contralateral trials for different economic variables, we next investigated whether ipsilateral trials exhibited similar or distinct spatial gradients. Previously having noticed that the neural population more strongly encodes for value on the

contralateral side compared to ipsilateral, we reasoned that spatial gradients may exist in the ipsilateral trials although the robustness of such gradients may not be as strong as those found in mappings of activation weights from contralateral trials. In order to examine this hypothesis, we once again constructed 2D spatial maps with the absolute difference of averaged activation weights between stimuli with different EV taken (**Fig. 3.7f, left**). In ipsilateral trials, there was a stronger representation of EV in the dorsal and not ventral space, mirroring what was observed in contralateral EV spatial representations, albeit with less strength in representations. In contrast, the spatial gradient formed by representations of risk in the 2D space was found mostly in dorsomedial areas (**Fig. 3.7f, right**). Comparatively the spatial gradients observed in contralateral trials were stronger than those observed in ipsilateral trials. Therefore, we find that values are encoded more strongly in the contralateral side. Overall, we find that frontal cortical neurons encode economic variables in a lateralized manner with stronger representations maintained in the hemisphere contralateral to the recording site.

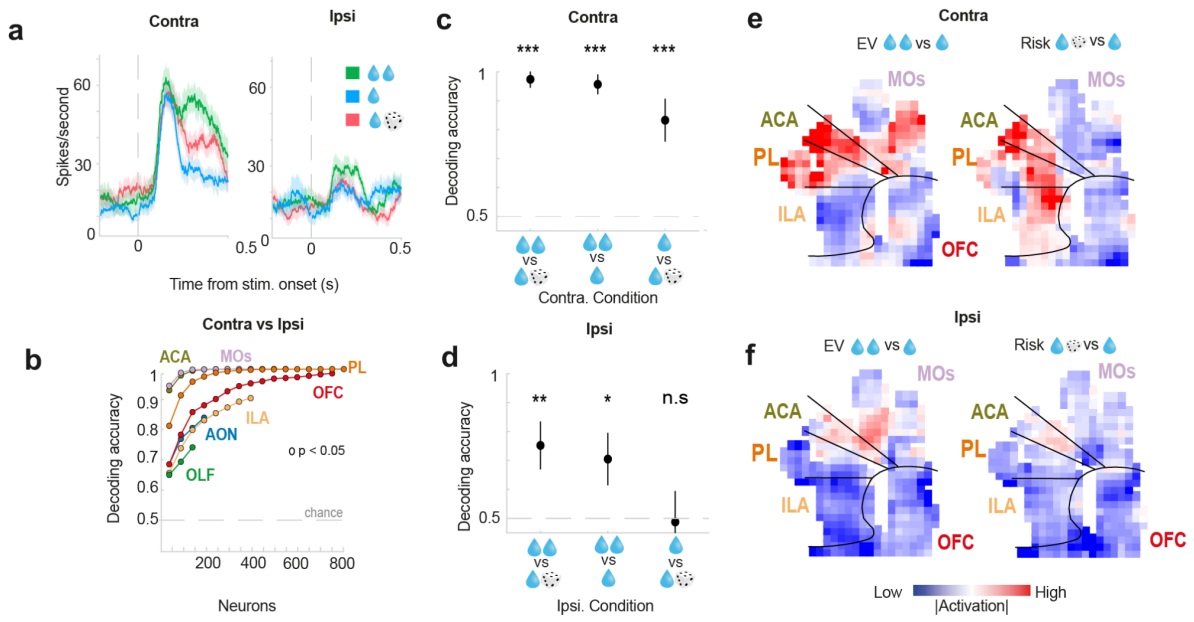


Fig. 3.7 | Stronger contralateral value representations observed in frontal cortical neurons

a, PSTH plots for an example neuron encoding. The PSTH shows firing rate split by three single stimulus conditions. The PSTH on the left depicts the firing rates of the neuron for contralateral trials (relative to the recording hemisphere). The PSTH on the right depicts the firing rate for the same neuron but for ipsilateral trials. **b,** Decoding of ipsilateral or contralateral trial conditions from individual prefrontal cortical regions and olfactory regions. While decoding accuracy is significantly above chance for all grouped regions, decoding accuracy is much higher for frontal cortical regions such as ACA and PL compared to olfactory regions. **c,** Binary decoding of single stimulus contralateral conditions. Decoding accuracy was significantly higher compared to chance when classifying all trial pairs. However the accuracy was notably higher for condition pairs with a difference in EV compared to the same EV pair, indicating that neural representations were sensitive to value even when considering contralateral visual stimuli. **d,** Similar binary decoding for different single stimulus condition pairs but within ipsilateral conditions. Decoding accuracy was significantly above chance for condition pairs with a difference in EV but not for the same EV pair. Notably the lower decoding accuracy on all condition pairs compared to the contralateral trials indicates that neural representations are sensitive to value in a lateralized manner. **e (left),** Spatial decoding plot showing absolute activation weights averaged in spatial bins that are overlaid with an image of a brain atlas. In this spatial decoding plot, we classify for contralateral conditions with different EV but same risk. Higher activation weights were observed in the dorsal compared to the ventral regions. **e (right),** Spatial decoding plot showing results of a classification of contralateral conditions with the same EV but different risk. Higher activation weights were observed in the medial compared to lateral space in the atlas view. **f (left),** Spatial decoding plot showing activation weights for classifying ipsilateral conditions with different EV and same risk. Activation weights are higher in the dorsal compared to ventral space. **f (right),** Spatial decoding plot displaying activation weights for classifying ipsilateral conditions with same EV and different risks. In this plot, the highest activation weights are observed in the dorso-medial space.

3.3.5 Frontal neuronal responses reflect the chosen and unchosen value

Thus far, we have focused on trials where mice were shown a single stimulus with different values and risk and have revealed how these variables are encoded in the frontal cortex. We next focused on a situation in which mice were able to choose from two simultaneously presented stimuli with different economic value. This allowed us to probe the neural representation of chosen and unchosen value, when choosing between two alternative offers. For our neuronal data analysis in this section, we exclusively focused on trials with two stimuli trials wherein these variables could be estimated, and for which we successfully trained a subset of mice (n = 13) to perform these trials (**Chapter 2, Fig. 2.3**). Moreover this analysis is conducted using a subset of mice (n = 13) trained with two stimuli trials and exhibiting high overall single stimulus performance to show overall task engagement. We have a smaller set of neurons recorded in this analysis and neurons were primarily recorded from two regions: secondary motor cortical areas (MOs) and orbitofrontal cortex (OFC). Individual neurons show strong stimulus-aligned responses, with different firing rates for different chosen stimulus conditions. For example, the OFC neuron in **Fig. 3.8a** shows a larger response in trials when the mouse is choosing the 'Large' over the 'Gamble' stimulus in the 'Large versus Gamble' two stimuli condition compared to trials when the mouse is picking the 'Small' over the 'Gamble stimulus in the 'Small versus Gamble' two stimuli condition (**Fig. 3.8a, left**). Similarly the same OFC neuron shows a strong response for trials when the 'Gamble' stimulus is chosen over the 'Large' stimulus compared to trials when choosing the 'Gamble' over the 'Small' (**Fig. 3.8a, right**). This neuron therefore appears to encode the expected value of stimuli regardless of whether they are ultimately chosen. In contrast, the MOs neuron demonstrates larger responses for the conditions for chosen value but similar response profiles for unchosen value trials, suggesting modulation by only chosen value conditions in the post-stimulus period (**Fig. 3.8b**). This preliminary analysis of individual neurons suggests that during the economic decision-making task, in the time period post-stimulus presentation, neurons are modulated by the value of the chosen and unchosen offer.

We next sought to investigate the representations of chosen and unchosen values at the population level.

In order to examine these representations on a population level, we applied the linear population decoding technique, Support Vector Machines (SVM) to see how well we can decode the chosen and unchosen variables from the spiking activity of all recorded cells. To do so, we repeated the two-class SVM analysis using two stimuli trials and classifying for chosen and unchosen value conditions, using neurons across the frontal cortex. For chosen value decoding, we trained a classifier to classify different chosen values, while holding the unchosen value constant (chose 'Large' over 'Gamble' versus chose 'Small' over 'Gamble'). For unchosen value decoding, a classifier was trained to distinguish between different unchosen values, while holding the chosen value constant (chose 'Gamble' over 'Large' versus chose 'Gamble' over 'Small'). We used the time window of 150 to 350 ms post-stimulus, during which single stimulus conditions were decodable above chance. Focusing on this time window, we find that one can decode chosen and unchosen values using the entire population of neurons in the dataset above chance (**Fig. 3.8c; mean binary decoding accuracy for chosen value: 0.6877 $\pm$ 0.0929, $p < 0.001$, two-sided permutation test; mean binary decoding accuracy for unchosen value: 0.6277 $\pm$ 0.0844, mean $\pm$ SEM, $p < 0.001$, two-sided paired t-test**). Finally, we found a stronger representation of chosen compared to unchosen value (**Fig. 3.8c; Δ Mean decoding accuracy [chosen value] - [unchosen value] = 0.0690, $p < 0.001$, two-sided paired t-tests**). Thus, population decoding in the frontal cortex revealed representations of chosen and unchosen value.

Having established that the general population of frontal neuronal responses reflect the value of the chosen or unchosen stimuli, we next asked how these representations of different economic variables are structured across the frontal cortex. To do so, we repeated our two-class SVM analysis on neurons recorded in different frontal regions. Since we recorded these neurons in a subset of mice in the dataset, we used neurons from two regions with enough neurons to complete the

decoding analysis (i.e. OFC and MOs) and since we acquired a different number of neurons from each region, we performed this analysis on a matched number of neurons across these two regions. The decoding accuracy for the chosen value of the stimuli was significantly higher than chance in both regions of the frontal cortex (**Fig. 3.8d**). Meanwhile, the decoding accuracy for the unchosen value of the stimuli was relatively maintained in OFC neurons but dropped close to chance level in MOs, suggesting a dominant role of OFC in encoding the unchosen value compared to MOs (**Fig. 3.8e**). In single-stimulus EV decoding, MOs provides higher decoding accuracy compared to OFC. However, in two-stimulus EV decoding of the chosen option OFC is higher than MO. This may suggest that the neural codes underlying the choice between two stimuli is at least partially different from the neural representation of the value of a single stimulus.

Since OFC neurons significantly encode chosen and unchosen values above chance, we next asked whether the chosen and unchosen values were encoded in a similar or different format in OFC neurons. In order to answer this question, we performed cross-decoding analysis, i.e. training the SVM with chosen value data to predict unchosen value conditions, and vice versa. Examining the cross decoding results reveals that the neural code of chosen and unchosen value is largely non-overlapping (**Fig. 3.8f**). The cross decoding results are only slightly above chance level compared to the chosen and unchosen value decoding accuracies (**Fig. 3.8f; OFC decoding accuracy for chosen value: 0.696 ± 0.0929, OFC decoding accuracy for unchosen value: 0.628 + 0.0965, OFC cross-decoding accuracy: 0.527 + 0.0722, p = 0.002, mean ± SEM**). Cross-decoding by training from chosen value and testing on unchosen value, and training of unchosen value and testing on chosen value, yielded similar decoding accuracies that are slightly significant above chance (**Fig. 3.8f; mean cross-decoding [chosen (train) -> unchosen (test)]: 0.532 ± 0.0686, P = 0.0260 , mean cross-decoding [unchosen (train) -> chosen (test)]: 0.511 + 0.0783, P < 0.001; mean ± SEM, two-sided paired t-tests**). Therefore, we found that frontal cortical neurons encode chosen and unchosen offer value with stronger representations for

both variables in orbital regions compared to MOs. Furthermore, this strong representation of chosen and unchosen offer value in the OFC population appear to be non-overlapping.

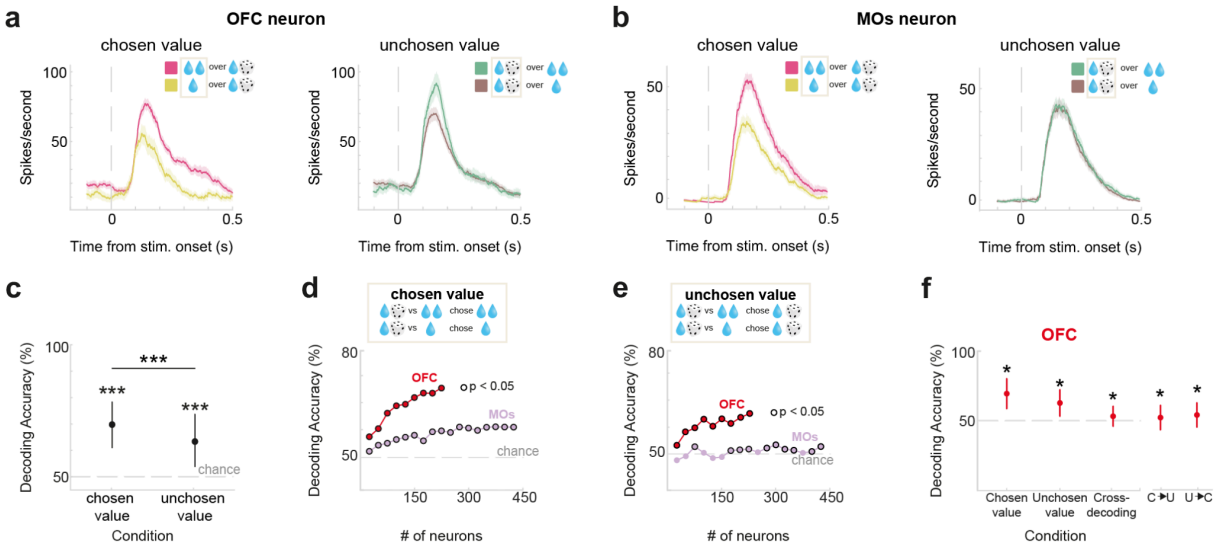


Fig. 3.8 | Distinct representations of chosen and unchosen value in the frontal cortex

a (left), PSTH plot for an example OFC neuron encoding chosen value. The PSTH shows firing rate split by chosen value conditions. Shaded regions show the s.e.m. across trials. **a (right),** PSTH plot for the same example OFC neuron in (a, left) encoding unchosen value. Overall, this individual OFC neuron shows strong stimulus-aligned responses, with different firing rates for different chosen/unchosen stimulus conditions. **b (left),** PSTH plot for an example MOs neuron encoding chosen value. The PSTH shows firing rate split by chosen value conditions. Shaded regions show the s.e.m. across trials. **b (left),** PSTH plot for the same example MOs split by conditions in the unchosen value categories. Unlike the OFC neuron showing encoding of chosen and unchosen value in the post-stimulus period, the MOs only show encoding of chosen value but not unchosen value. **c,** Binary decoding of chosen/unchosen values in two stimulus conditions. Decoding accuracy was significantly reduced when classifying trials with the unchosen value compared to chosen values, indicating that neural representations were sensitive to the chosen offer. However, one can still decode chosen and unchosen values using the entire population of neurons in the dataset significantly above chance; with a significantly higher decoding accuracy for chosen value. **d,** Decoding of two stimuli trails with chosen value conditions from two prefrontal cortical regions (OFC and MOs). OFC and MOs neurons have high decoding accuracies for chosen value conditions but the OFC population displays much higher decoding accuracies than MOs. **e,** Decoding of two stimuli trials with unchosen value conditions from the same individual prefrontal cortical regions. OFC neurons display high decoding for unchosen value conditions but decoding for unchosen value from MOs is far below the accuracy provided by OFC neurons. **f,** Comparison of decoding of chosen value, unchosen value, and cross-decoding of chosen value using unchosen value trails for training, and vice versa. Examining the population of OFC neurons that strongly encode chosen and the population that encodes unchosen value reveals that the population is largely non-overlapping. The cross decoding results are only slightly above chance level compared to the chosen/unchosen value decoding accuracies for chosen and unchosen value using OFC neurons.

3.4 Discussion

The aim of this study was to determine the structure of representations of economic variables across frontal cortical regions during a head-fixed visually guided economic decision-making task. I found that individual frontal neurons encode different economic variables such as the expected value of the offer as well as the risk associated with a stimulus. These neurons encoding such task variables were found to be distributed across frontal cortical regions, spanning across regions such as the OFC and the ACA. These findings are consistent with Padoa-Schioppa and Assad et al. (2006), a study which used electrophysiological measures of neural activity in primate OFC to study how values are represented neurally in a particular frontal region. The authors found that during economic choice, individual neurons in area 13 of the primate OFC not only displayed responsiveness to the task variables, but also encoded the value of the juice offered to the monkey. Here, I have further demonstrated that the activity of neurons located in multiple regions across the frontal cortex are reflective of similar economic variables.

We find that on a population level, frontal neuronal responses encoded economic value, reflecting animals' preference among stimuli. Moreover, we show that these representations of economic value assume a graded representation across frontal regions. Additionally, the neuronal responses assume distinct graded spatial gradients for expected value and risk, suggesting that the spatial gradient is a crucial part of understanding the structure of representations in frontal cortical areas. In addition to encoding expected value and risk, frontal neurons also encode the chosen and the unchosen value of the offers, pointing towards the encoding of multiple economic variables by frontal populations. Lastly, the expected value and risk coding assume stronger representations in the hemisphere contralateral to the recording site, pointing towards a hemisphere-specific representation of these economic variables.

3.4.1 Distributed representations of expected value and risk across frontal cortex

Individual neurons exhibit responsiveness to different task events and responses are reflective of the mouse's choices, i.e. responses are scaled with either the value or the probability associated with the stimulus-reward pairing in our task. Our findings that neurons in the frontal cortex show task-modulated responses to different economic variables such as expected value and risk are in line with past research indicating that frontal neurons encode different decision-variables (Padoa-Schioppa and Assad, 2006; O'Neill and Shultz, 2010; Kepecs et al., 2008; McClure et al., 2004; Tobler et al., 2006; Kable and Glimcher, 2007; Walton et al., 2002; Wallis et al., 2001). These studies highlight the roles of neurons in various regions in encoding key decision variables such as the values of the offers, the value of the chosen offer, the identity of the chosen cue, risk, reward history, the cost to obtain a larger reward, abstract rule coding. However, these studies are limited in their role in explaining how frontal regions contribute to decision-making as often these studies comprise recordings from a small number of neurons in particular regions in isolation and or use multiple behavioral tasks making it challenging to compare neural activity aligned to different task events. To account for this, we developed a few extensions to previous studies: we trained mice to make economic choices in our visually guided economic decision-making task and recorded from several regions simultaneously while mice made choices.

We observed that population activity aligned to the onset of stimulus presentation is reflective of the value or the reward probability associated with each stimulus condition. Our study highlights the distribution of the neural representations for such variables across multiple regions and not just confined to one region - thereby emphasizing the importance of analyzing the neural code and its distribution across brain areas. Furthermore, we show quantitative differences in economic variable encoding among these areas, suggesting that some areas may play a more prominent role in encoding of EV and risk. For example, we observe that ACA and MOs - two regions located in the

dorsal portion of the mouse frontal cortex - exhibit higher decoding for EV compared to OFC - a region largely dominant in the encoding of EV variables in the literature. Risk decoding shows the highest decoding accuracy from ACA and PL - two regions located in the medial portion of the mouse prefrontal cortex- compared to OFC. However, it is important to note that these distributions of key decision variables across frontal cortical regions are also present in non-frontal areas.

We have presented several pieces of evidence indicating that neuronal responses encode economic variables, rather than pure visual properties. First, neural responses in conditions with different EV were more distinguishable than conditions with the same EV, despite visual stimuli being randomly assigned to conditions across animals. If neural activity would purely reflect visual features, all conditions should be equally distinguishable. Second, we observed distinct, orthogonal spatial gradients for the coding of EV and Risk. This is difficult to reconcile with a neural code that purely reflects the visual features of our stimuli, especially given that different visual features signalled different economic variables across animals. Having said this, we would like to propose future experiments that will help disentangle that neuronal responses indeed scale with stimulus value, rather than encoding stimulus visual properties. There are ongoing experiments using two-photon imaging of dorsomedial frontal cortex (dmFC) throughout a classical conditioning task which may further disentangle the visual properties from the neural responses related to economic variables. In the future, we will perform experiments in which we image dmFC neuronal signals in well-trained mice in a task in which two visual stimuli have the same value, while the third one has a lower value. We would expect the neuronal responses to the two stimuli with the same value to be largely similar, and differ from the 3rd stimulus, thereby allowing us to further disentangle the neuronal responses for economic variables.

Despite our use of a kernel regression model, to remove the effects of stimulus and action from EV, we are left with a number of unmeasured variables. In particular, we are using a linear model and it

does not remove non-linear movement related signals and does not remove signatures that are not related to the RT (e.g. licking and other orofacial movements). We removed the neural activity that is related to the response time at the onset of the wheel movement. There is a caveat using this approach which is that there is a potential for several confounds that might not be removed entirely. Specifically, there are potential movement related signals (e.g. licking), mental state (e.g. motivation, arousal), neural noise and several other unmeasured variables which may impact the neural activity. If for example, neurons respond strongly to faster choices and much less so to slower choices, then that is something that we do not remove. We assume that the neural response to an action is always the same and just differs in the timing depending on when the mouse decides to move the wheel. Therefore we do not remove non-linear movement related signals and or signals related to movement that doesn't manifest in the wheel movement or response times.

Overall, our analysis addresses the issue of movement confounds using state-of-the-art methods, such as kernel regression. However, there is a potential caveat that we have not controlled for all movement related signals beyond RTs. In particular in the future we plan to probe this in new experiments in which we measure dmFC neuronal signals using two-photon imaging, in a classical conditioning experiment in which visual stimuli are associated with different reward probabilities. These experiments will help to further disentangle whether dmFC neuronal responses indeed reflect the economic value of stimuli. Despite this, we have not measured all movements of the mouse such as arm position and this will be crucial in the future for us to further control for movement-related signals.

Even though we observe robust decoding accuracies for EV and risk across these frontal populations - it would be crucial to observe the timescales during which these representational activations are most pronounced. Even though EV and risk conditions are decoded using a 150 ms to 350 ms post stimulus time window - there may be a lag in the activations across particular brain regions or across all of them collectively. A paper that investigated activations in dopaminoceptive structures for EV

and risk by Preuscoff et al., 2006 found that on a temporal scale the activation related to EV came first followed by a delayed activation for risk. This study suggests a potential avenue for us to investigate these representations and identify not only spatially the robustness of these representations but also on a temporal scale. If we identify a lag among these activations it could suggest that either different cell types or neuromodulators may be recruited in mediating each of these representations or even an activation from one region provides information to other regions. Even though preliminary analysis of time-resolved decoding across frontal areas did not reveal stark differences in the emergence of economic variable representations, expanded analyses on the timescales of multiple representations may show interesting patterns emerging.

Past studies have been crucial in characterizing the role of different frontal regions, with some studies suggesting distinct roles for distinct regions in encoding of economic variables, while others suggest that different regions play multiple roles in encoding of different variables in economic decision-making. Neurons in the frontal cortex have been shown to exhibit diverse roles in encoding of different cognitive variables during decision-making (Yang et al., 2022). This past study suggests that individual neurons that encode such variables in the mouse anterior lateral motor cortex exhibit distinct and robust coding of different decision variables rather than displaying randomly mixed selectively. This organized and distinctly selective code displayed by neurons in a particular region during a decision-making task involving tactile stimuli, implies that the structure for such representations in a region assumes a distinct neural code. Having observed distinct representations in the past study, prompts the question of whether there may be either singular or multiplexed roles of each region in representing decision variables in the frontal cortex. In the frontal cortex, encoding of decision variables often occurs with neurons in certain regions showing responses tied to multiple decision variables, supporting the multiplexed role of these regions. A paper that demonstrates this is by Kennerley and Wallis, (2009a) which uses electrophysiology in macaques to study how frontal neurons encode multiple decision variables. The authors found that neurons in the ACA, OFC and lateral PFC encode the potential payoff of the stimulus, the reward probability,

as well the cost of making a choice. However, most interestingly, the signals for these three representations were most robust in the ACA and individual ACA neurons showed multiplexed representations of these three decision variables. This study suggests that particular regions in the frontal cortex play multiple roles in decision-making. Although we do not directly address the potential multiplexed roles of the frontal neurons in our dataset - we hypothesize that the structure underlying such robust representations are intertwined and dependent on one another. One step that we can take is to see whether we can cross decode using a training population that robustly codes for EV single stimulus conditions and test on two stimulus trials with chosen value. In our analyses we did not conduct this cross decoding between single and two stimulus populations due to us having fewer neurons recorded from during the two stimulus conditions.

3.4.2 Distinct spatial gradients for representations of economic variables

EV and risk decoding display distinct spatial gradients that assume orthogonal conformations to each other in our study. EV coding assumes a dorsal to ventral gradient with the highest decoding accuracy maintained across the AP axes in more dorsal regions such as the ACA and the MOs, while coding for risk assumes a strong medial-to-lateral gradient with the strongest coding present in medial structures such as ACA, PL and IL. Interestingly both EV and risk spatial activation gradients share one commonality - they overlap with the strongest coding for both decision variables in the dorso-medial portion of the spatial map with the ACA being present in that portion.

These observations of stronger EV coding in the dorsal portion of the spatial map may be due to the way that learned value information is relayed through the frontal cortical spatial map. One reason for this stronger spatial representation of EV in the dorsal portion could be due to the fact that these mice are trained in this task and have learned value information across many training sessions. Since they have been strongly trained - these value representations are maintained throughout their frontal circuits with a possibility that on a spatial scale - the ventral portion of frontal cortical areas

is important for maintenance of EV representations but the dorsal portions representations of EV are more strongly coupled to the transformation of the value information into an action. In fact a study by Tang et al., 2022 demonstrates these points during their examination of corticostriatal circuits. They discover that learned value and stimulus information is strongly represented in the ventral corticostriatal circuit in areas such as the ventral striatum, OFC and amygdala. Meanwhile, this stored value information is converted into an actionable goal in the dorsal corticostriatal circuit in the LPFC. Additionally a plausible explanation for the stronger EV activation in dorsal frontal cortex is the connectivity between dorsal and ventral portions of frontal cortical circuits. In fact, ventral regions such as the OFC send outputs to dorsal areas such as the ACA, a region from which information is directly routed to motor areas where action plans to obtain rewards can be executed. These anatomical connections support the view that information about value is represented across multiple areas but assumes a functional gradient shaped by the functional connectivity and inputs and outputs to areas.

Furthermore, our finding of distinct dorsal-to-ventral coding of risk gradient in the frontal cortex is not the first report of such a phenomenon. In a monkey study by Maisson et al., 2021, they train monkeys in a two-option risky decision-making task and examine neural activity across different regions of the frontal cortex. In this paper they identify a functional gradient that exists along the ventrodorsal axis of the medial frontal cortex with the strongest decodability for two economic variables - expected value and chosen offer - appearing in dorsal regions and lowering towards the ventral regions. They find this functional gradient exists for different task variables which includes not only economic variables but also encompasses variables such as the spatial information regarding the offers at hand. Studies such as this one highlight the relevance of a functional gradient and the role it serves in how information is transmitted about encoding of different decision variables.

3.4.3 Neural correlates of post-decision variables

Our everyday decisions often require us to make choices between two or more offered options and each option is associated with a different action. There are two schools of thought surrounding the way the brain enables us to make such decisions. The first view suggests that the brain makes decisions by computing and then comparing the values of the offers in a distinct 'goods space' before making an action. In other words, the first view suggests that one compares the values of the offers and makes a choice among them even before a motor response association is made. The second view posits that the brain first assigns values to the associated actions and then makes a choice between the value-actions.

We examine how decisions are formed across frontal regions in a portion of our task design where mice encounter two simultaneously presented stimuli on the screen and have to make a choice between them. We find that OFC and MOs neurons encode chosen values during the post-stimulus period. This finding of chosen value encoding by OFC neurons is consistent with previous work by Padoa-Schioppa and Assad, (2006) - a study that trains monkeys to make choices between two juice offers and simultaneously take electrophysiological measures of neural activity. The authors find that during economic choice, individual neurons in the primate OFC encode the value of the chosen good. Interestingly in our study we can strongly decode the value of the chosen stimulus in the period post stimulus presentation. This finding of a post-decision variable signal being present during the post-stimulus presentation time period, is consistent with previous work by Wunderlich et al., 2010, a study which uses fMRI in humans to characterize the nature of decision-making in the human brain. The authors in this study find that neural signals in the ventromedial prefrontal cortex indicate the value of the chosen offer in the time period prior to the stimulus-action pairing. This previous study provides support for the view that values are computed in a goods space. However, we cannot yet conclude in our study whether the mice in our task compute value in an abstract goods space and not in a spatially- or action-specific format. Furthermore, our neural representations exhibiting some similarities and differences to past studies in monkeys and humans

makes sense given the discrepancies and homologies in the anatomy of frontal regions across these species as well as the differences in behavioral task structures between monkeys and rodents.

Additionally, it is interesting to compare decoding accuracy for chosen value among the two regions: MOs and OFC. First, OFC is the region that strongly encodes both chosen and the unchosen value of the offers. In contrast, neurons in MOs mainly encoded one of the two post-decision variables: chosen value. This suggests a stronger role that OFC plays in encoding of economic decisions compared to motor regions where more action-related information related to these decisions may be present. Since the neural code for chosen value is stronger in OFC compared to MOs, it can be hypothesized that the neural signal in OFC is tied to the goods space while the MOs neural code might be intertwined with the action space. One way to examine this is to establish the time window important for decoding chosen values in OFC neurons compared to MOs. However, in the scope of our study, we did not implement this time-resolved decoding for chosen value mainly because the sample size was too small.

3.4.5 Summary

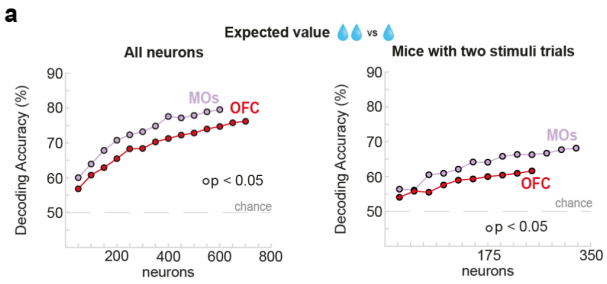
This chapter shows that individual neurons across frontal cortical areas exhibit stimulus-aligned responses to different task variables such as EV and risk. Furthermore, population activity can discriminate with high decoding accuracy, conditions with different EV, suggesting that population activity across multiple areas carries information about value. Furthermore, examining population activity and splitting by mice that exhibit risk sensitivity and neutrality reveals that regions are also carrying information about risk and individual risk-preferences in our task. These representations of two key variables assume distinct spatial gradients that are (A) orthogonal to one another and (B) strongest for contralateral stimuli. Furthermore, we identify signatures of two more variables in addition to EV and risk from our neural activity: value of the chosen offer and the value of the unchosen offer.

In order to fully understand the relationship between encoding of these four economic variables, investigations of: (a) the spatial structure of chosen and unchosen value, (b) the degree to which these neuronal population carry information about multiple variables or solely represent one variable and (c) the temporal nature of the timescales in which these representations appear in the brain, will need to be investigated. Overall the distinct spatial gradients observed for encoding of expected value can be explored in other contexts in our study. Since the representation of expected value is higher in dorsal compared to ventral areas it can be asked: does this gradient expand to encoding of other decision making variables such as chosen value? In our study we do not examine this gradient due to the smaller population of neurons in our two stimuli dataset. However, this would be a crucial avenue to explore. Additionally we anticipate that motivation plays an important role in the encoding of economic variables. Motivation and salience of different stimuli as well as general task engagement would not only explain behavioral fluctuations but may shed light on shifts in the representations of variables that we note in our study.

3.5 Acknowledgments

This work was carried out in partnership with Caitlin Ashcroft (a doctoral student in the Lak Lab) and Dr. Matthias Fritsche (postdoctoral fellow in the Lak Lab). Caitlin Ashcroft and I carried out electrophysiological experiments with Caitlin contributing electrophysiological data from 14/27 mice in the entire dataset. Toluwani Duckworth-Esselfie helped with histology and imaging along with tracing of probe tracks using the SHARP track pipeline. Anita-Maria Rominto and Jake Swann helped with manual curation of neuronal clusters using the Phy GUI along with probe tracing and alignment to the Allen Brain Atlas. Prof. Armin Lak supervised the electrophysiology experiments and Dr. Peter Zatka-Haas wrote the neuronal pre-processing code as well as the common reference averaging code. Dr. Matthias Fritsche designed and implemented the kernel regression and population decoding analysis pipeline. Caitlin Ashcroft created the analysis pipeline for population decoding of value from ipsilateral/contralateral trials and stimulus side decoding. A manuscript of this chapter is in preparation for publication.

3.6 Supplementary Figures



Supplementary fig. 3.9 | Decoding accuracies for mice with two stimuli conditions mirrors those in the general population

a, This analysis supports the observations for two stimuli condition decoding of chosen and unchosen values. In our experiment we had two groups of mice: one cohort only experienced the single-stimulus trials, while the second cohort also experienced the two stimuli trials. One concern is that the two stimuli condition decoding accuracies are driven by the variability in behavior among these two cohorts. To address this potential confound, we compared the decoding of single stimulus conditions with different EV from all mice with mice that made choices among two stimuli. We reasoned that if the two stimuli decoding accuracy is reflective of behavioral variability of the subset of mice with two stimuli trials then we should observe higher decoding accuracy for OFC compared to MOs in the single stimulus different EV condition decoding using the subset of mice with the two stimuli trials. We concluded that the higher decoding accuracy for the OFC compared to MOs is not due to the behavioral variability from the subset of mice that made choices among two stimuli. Decoding of single stimulus conditions with different EV but same probability from two individual prefrontal cortical regions (MOs in lilac and OFC in red) with data from all neurons across all recordings. The decoding accuracy for MOs is higher than for OFC in the general population. We next ask whether this improved decoding accuracy for MOs compared to OFC is due to behavioral variability between all mice and mice that had two stimuli conditions. **b,** Decoding of single stimulus conditions with the same EV but different risk from the same two prefrontal regions but this time using neurons from recordings with mice that had two stimuli conditions. Decoding accuracies are slightly lower for the population of neurons from mice that had two stimuli conditions compared to the general population suggesting that behavioral accuracy does play a role in the decoding accuracy. However, the decoding accuracy of MOs is higher than for OFC even in the neuronal population from mice that had two stimuli conditions. Since the order of accuracy is maintained across both - it suggests that behavioral variability among groups of mice did not impede on the observations of a stronger representation of EV from the MOs population.

4

Optogenetic inactivation of frontal regions

Contents

4.1 Introduction . . . . .
4.2 Methods . . . . .
4.3 Results . . . . .
4.4 Discussion . . . . .
4.5 Acknowledgements . . . . .

4.1 Introduction

In the previous chapter, we demonstrated that the frontal cortical representations of economic variables (i.e. expected value (EV) and risk) assume orthogonal spatial gradients. Specifically, we have demonstrated a medial-to-lateral gradient for representation of risk and a dorsal-to-ventral gradient for representation of EV, hence establishing the strongest level of representation of such variables in the neural activity of the dorsomedial frontal cortex. Having shown that spatial gradients in the encoding of economic variables across prefrontal cortex (PFC), we next asked whether this spatial structure characterizes the causal contribution of these regions to economic decision-making. To investigate this, we used optogenetics tools available in mice to investigate whether prefrontal cortical regions causally contribute to decision-making in this task and ask whether we can replicate the spatial gradient observed in the neural activity. We hypothesize that by inactivating frontal

regions during the economic decision-making task, we will disrupt choice behavior that rely on the computation of economic value. We focused on conditions with a different EV (Large and Small single stimulus trials) to run our pilot experiments with the hypothesis that inactivation of frontal regions would in particular disrupt choice accuracy on different EV conditions.

In our previous chapter we noted similar decodability of EV from neural activity in prelimbic cortex (PL) and orbitofrontal cortex (OFC), thereby prompting us to choose these areas as candidate areas in our inactivation experiments. We hypothesize that given the similarity in decodability of EV from these two areas, that inactivation of these areas might show similar disruption of economic choices. Overall, assessing whether neural representations in frontal regions play a causal role in economic choice in our head-fixed visually guided economic task is important because we demonstrate a distributed spatial network of these representations across the frontal cortex. More importantly, we identify a critical time window post stimulus presentation of 150 to 350 ms during which we can decode the expected value significantly above chance from the spiking data across regions. It is therefore unclear whether manipulating neural activity in this specific time window will adequately disrupt computation of economic value. It is also unclear whether the causal role of these representations in the dorsomedial spatial area, an area with the strongest representations of EV and risk, will have the strongest impact on decision-making behavior or whether we need to adopt the idea of a distributed code of value computations in the brain. In our study we unilaterally inactivate one target region at a time, leading to a possible caveat: inactivating one area may lead to compensatory roles from other regions to allow for effective decision-making. However, addressing the causal role of frontal representations in a mouse model has its advantages. First we have the availability of genetic mouse lines that allow us to activate specific cell types. Second, using targeted injections and fiber implantations above target areas is useful in investigating the causal nature of these regions implicated in the decision making circuit. Before I detail our causal experiments, I would like to briefly describe previous studies that have investigated the causal role.

Despite the many observations of frontal regions encoding variables such as value, it remains unclear whether such neural representations are causally related to choices. Several studies in the past have provided correlational evidence connecting economic decisions to representations of value in the brain. In studies in non-human primates and rodents, there is considerable evidence pointing to the crucial role that different brain regions play in carrying representations of economic value. Since then, studies have sought to show that neuronal signals in such frontal regions are causally related to economic choices. In order to investigate the causal role of these areas, different methods to activate or inactivate brain regions were employed.

One of the earliest methods employed is electrical stimulation. Electrical stimulation of frontal areas shows how neural signals in certain areas are causal to economic choices. A paper that demonstrates this is Ballesta et al., (2020) which used two levels of electrical stimulation of OFC neurons in rhesus monkeys performing an economic decision making task. The task requires monkeys to choose between two juices which were offered in different amounts. These two offers were presented to the monkey sequentially in the center of a computer screen with the sequential order of the two juices varying pseudorandomly. In this study, electrically stimulating OFC leads to monkeys showing biased choice and increased subjective value for specific offers (Ballesta et al., 2020). Specifically, they found that a low current stimulation increased the subjective values associated with the offers and biased the monkey's choice behavior. Meanwhile, a higher level of stimulation led to the disruption and computation of subjective values leading to disruption of choice behavior. This study quite interestingly shows how different degrees of stimulation can progressively disrupt choice behavior and disrupt a chain of processes that are integral to proper computation and comparison of economic value. The lower degree of stimulation pointed out the causal role of OFC. A caveat in such a study is the organization of OFC and the presence of mixed populations of neurons in these areas encoding different variables might not necessarily target OFC neurons that solely encode value. Other primate OFC stimulation studies point out the casual role of OFC in value computations and evaluations. In another study from the same lab using a similar protocol of

electrical stimulation in OFC and the same behavioral task as Ballesta et al., (2020), they found that electrical stimulation of OFC neurons disrupts the value comparison process and does not impact the value of each individual offer - pointing to the important role of the OFC in value comparisons (Ballesta et al., 2022). Both Ballesta et al., (2020) and Ballesta et al., (2022) point out the importance of activation studies for studying the causal role of neural representations in decision-making. Additionally, they point out the causal role of neural representations in OFC in economic choices.

Manipulation experiments have since expanded into rodent models where targeted manipulations are possible given genetic tools. Rodent studies, using rats, modeled an economic task after many of the central principles used in previous primate studies in order to isolate the causal role of the OFC in making economic choices. In a study by Gardner et al., 2017, rats were trained in an economic choice task in which they had to choose between different numbers of flavored food pellets and this allowed for measurement of the relative preference between different goods. In expert rats that learned the task, optogenetic inactivation of lateral OFC revealed no impact on animal's choice behavior (Gardner et al., 2017). However, the same study showed that lateral OFC is critical for reinforcer devaluation on Pavlovian conditioned responses (Gardner et al., 2017). The Pavlovian reinforcer devaluation task is a behavioral task that assesses the animals' ability to alter their choices when the value of the offer has changed by pairing the stimulus with an aversive stimulus, in this case by conditioned taste aversion. An important consideration is that the region of OFC targeted is analogous to the primate area 13 - where the neural correlates of economic choice are normally observed (Padoa-Schioppa et al., 2009; Padoa-Schioppa and Assad, 2006, 2008; Rich and Wallis, 2016; Tremblay and Shultz, 1999; Xie and Padoa-Schioppa, 2016) as evidenced by anatomical studies (Heilbronner et al., 2016; Preuss et al., 1995; Stalnaker et al., 2015). Another causal manipulation study, in rats, inactivated neurons in the lateral OFC during an economic task, which required rats to choose between different visual cues and food associations (Gardner et al., 2019). When they inactivated the lateral OFC during this task, they found that lateral OFC is not necessary for economic choice unless the choice involves integrating new information about the outcome

(Gardner et al., 2019). These studies provide conflicting evidence to the causal role of OFC in economic choice behavior. A likely explanation for such conflicting evidence is that the lateral OFC is part of a larger network of distributed representations of economic value and that it may not play as dominant a role as previously proposed in many studies. Even if it is part of a network with a distribution of representations, the OFC may still play an integral role under some circumstances.

Weak effects of inactivation or stimulation of OFC, commonly observed in studies above, may point out even more that calculation of economic value is not solely performed by a single brain region but also observed in other regions (Hunt & Hayden, 2017). Indeed other studies have investigated the causal role of representations in other regions in the context of economic choice. A study in rats, inactivated the frontal orienting field (FOF) and parietal cortex (PPC) during a behavioral task with trials where the rat had to choose between a small safe sure bet and a gamble (Bao et al., 2023). By inactivating FOF, they noticed a change in the curvature of the utility function, a common framework to explain risk sensitivity, for each rat - implying that FOF neurons encode the value of the gamble during this task. Specifically they noticed that silencing of FOF made rats risk-averse, caused by an increase in the concavity of the utility function. This experiment demonstrates how value signals in the brain are causally related to economic choices with a focus on the decision variable of risk. It is crucial to note that there is a previous study that has studied frontal inactivation during economic choice tasks in mice (Kuwabara et al., 2020). This study had specifically sought to investigate the causal role of lateral orbital area in the economic decision-making process. In their economic choice task, mice were trained to choose between different juices offered in varying amounts and each juice was associated with a different odor cue, with the odor concentration denoting the magnitude of reward. Furthermore, they demonstrated that silencing of lateral orbital area in mice using optogenetic inactivation methods led to disruption of choice behavior. Given the dominant role of OFC in economic decision making literature, it makes sense that some of the first causal experiments were designed to explore OFC in particular, however, as we demonstrate in Chapter 3, there are neural representations of EV and risk across

several frontal regions with some areas showing stronger representations than OFC for coding of economic variables. Therefore, this prompts investigation of the causal role of representations in regions beyond OFC in economic choices.

The timescales of representations are also integral to the decision making process and can differ across distinct regions. In order to target the time window of these representations and see whether this disrupts choices we need tools to silence neural activity with temporal specificity. Optogenetic inactivation of neural activity which uses the millisecond precision of temporally activating or inactivating neural activity is important for experiments where manipulation is targeted at specific time points. Overall, to study how neural representations are causally related to economic choices, we not only need the right silencing tools, but also the genetic technology to enable both spatial and temporal specificity during the interrogation of circuits. More work needs to be done to characterize the causal role of multiple frontal regions in economic decision making and to further understand the structure of the neural circuit and interactions between these regions. It is crucial to understand how representations in each region inform connected regions and regulate the decision-making process.

Here we present a mouse model for optogenetic inactivation of frontal regions during an economic decision-making task. We used optogenetics to study the effect of targeted inactivation of OFC and regions in the medial prefrontal cortices on the choice behavior in the head-fixed economic decision-making task. First, we adopt a modified version of our economic decision-making to probe how inactivation disrupts EV computations. Mice learned this modified version of the task rapidly and overall exhibited consistent choice behavior. Second, we used optogenetics in PV interneurons to examine the effects of inactivating OFC and mPFC areas on choices. Our target delivery of a Cre-dependent channelrhodopsin virus ensured a targeted spatial approach. Inactivation was induced by optogenetically activating PV interneurons which reliably induce a temporally precise

shutdown of neural activity in these areas. Overall, this approach allows for a careful evaluation of the causal role of neural representations of frontal regions in economic choices.

4.2 Methods

4.2.1 Animals

The data presented here is from 6 mice (all male) aged between 10-24 weeks. Pvalb-Cre mice backcrossed with C57/BL6 were used in this study. Mice were housed in a facility that maintained a 12 hour light/dark cycle, a room ambient temperature of 20-24℃ and a controlled humidity level of around 40%. All experiments were conducted under the UK Animals Scientific Procedures Act (1986) under specific personal and project licenses for this project.

4.2.2 Surgeries

All mice were implanted with custom-made headplates to allow for head fixation during behavioral experiments. A headplate implantation surgery was conducted in order to implant the custom-made steel headplate onto the skull. At the time of the headplate surgery, mice were 8 weeks of age. Mice were anesthetized with 5% isoflurane and then positioned onto a stereotaxic frame prior to intraperitoneal injection of 5 mg/kg meloxicam and 0.1 mg/kg buprenorphine. Following these injections, mice were maintained on 1.5% isoflurane (1-2% oxygen) for the duration of the surgical procedure and kept on a heating pad throughout the surgery. The dorsal surface of the skull was exposed by removing the overlying skin and periosteum. Following this, removal of any remaining tissue was done by using a brief application of an acid primer (Super-Bond C&B; Sun Medical) followed by washing the skull surface with sterile cortex buffer. We marked the location of bregma and lambda and then leveled the skull based on the DV coordinates at both of these marked locations.

For the injection and fiber implantation part of the surgery, mice were injected with 300 nl of an adeno-associated viral vector expressing hChR2, a channelrhodopsin virus, and a EYFP fluorescent tag (pAAV-EF1a-double floxed-hChR2(H134R)-EYFP-WPRE-HGHpA, titer = 1×10^{13} vg/mL) in medial PFC in one hemisphere and in OFC in the opposite hemisphere. The injection in the mPFC consisted of three injections of 100 nl at coordinates: AP: +1.7 mm, ML: $\pm$ 0.5 mm, DV: [-1.8 mm, -2.2, -2.6 mm] (**Fig. 4.1a**). The OFC injection consisted of one injection of 300 nl at coordinates: AP: + 2.7 mm, ML: $\pm$ 1.4 mm, DV: -1.8 mm (**Fig. 4.1b**). Prior to the injection, we made craniotomies over the two target areas (OFC and mPFC) using a dental drill. Following the injection, an optic fiber (200 μm diameter, 0.37 NA, Neurophotometrics) was chronically implanted in either of the two target areas. The coordinates for fiber implantation above mPFC were: AP: +1.7 mm, ML: $\pm$ 0.5 mm, DV: -1.3 mm. The coordinates for implantation of optic fiber over OFC were: AP: + 2.7 mm, ML: $\pm$ 1.4 mm, DV: -1.3 mm. The optic fibers were cut at a length of 2 mm to allow for appropriate targeting of frontal regions. Once inserted, we applied opaque dental cement surrounding the fiber, leaving room in the back portion of the exposed skull for implantation of the custom-made headplate. Following completion of the virus injection and fiber implantation in both hemispheres, the custom-made headplate was then attached to the dorsal surface of the skull using dental cement (Super Bond C&B; Sun Medical). The mouse was left to recover in a post-recovery cage placed over a heating pad. The mouse was provided with soft food, water, and jelly mixed with oral Metacam. The mouse was monitored until locomotion and grooming behavior were normal and the mouse was fully recovered. Following surgery, mice recovered for 5 days prior to initiating water restriction. Once water restriction was started, mice were handled regularly followed by acclimatization to the behavioral boxes by head-fixing. From the first day of head-fixation, mice were started with the initial shaping stage.

4.2.3 Behavioral task

In this task, mice (n = 6) encountered single stimulus trials, and the task included two visual stimuli. One stimulus predicted a large reward of 3 μL delivered with a probability of 100% referred to as a

'Large' stimulus with EV = 3 μL. The second stimulus used in this portion of the task is associated with a smaller reward of 1.5 μL with a 100% probability and an EV of 1.5 μL. We refer to this stimulus as a 'Small' stimulus. As before in our main behavior paradigm described in Chapter 2, we counterbalanced the mapping between stimuli and the associated reward across mice. Once mice were trained to expert level on the single stimulus portion of the task we initiated manipulation experiments. In each manipulation session, the inactivation ('laser on') and control ('laser off') trials were randomly interleaved (**Fig. 4.1c**). Additionally, inactivation trials consisted of 30% of all trials within a session on average. We alternated between unilaterally inactivating mPFC and ORF across days.

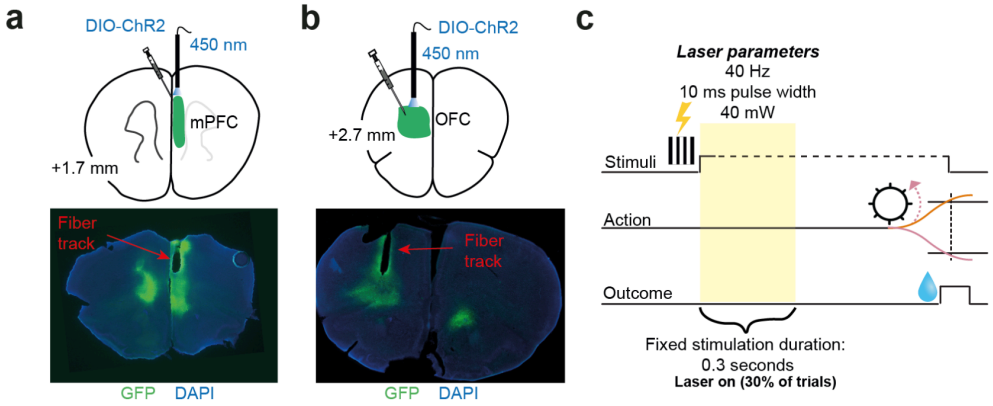


Fig. 4.1 | Histological confirmation of virus expression and fiber placement

a (top), Schematic of a local injection in mPFC with a channelrhodopsin virus. **a (bottom)**, Image of a frontal coronal section (AP: +1.7 mm) stained with DAPI (blue) and with expression of the pAAV-EF1a-double floxed-hChR2 (green). **c,** Example trial timeline showing different events such as stimulus onset, action onset and outcome time. In the yellow box, we show the time period when the laser is turned on and there are specific parameters used for this laser on for a fixed duration of 0.3 seconds post stimulus onset on 'laser on' trials. In this experiment we used single stimulus trials with two conditions ('Large' and 'Small') and overall 30% of trials had optogenetic stimulation with the laser being on for those trials while 70% of trials were control trials and had no laser. Stimulation and non-stimulation trials were randomly interleaved during each optogenetic experimental session.

4.2.4 Optogenetic manipulation experiments

In order to inactivate mPFC and OFC responses, we injected 300 nl of diluted virus containing ChR2 (pAAV-EF1a-double floxed-hChR2(H134R)-EYFP-WPRE-HGHpA) unilaterally in the

mPFC (in one hemisphere) and the OFC (in the opposing hemisphere) of Pvalb-Cre mice backcrossed with wild type, C57/BL6 (**Fig. 4.1a,b**). We implanted an optic fiber (200 mm, Doric Lenses Inc) in these target areas, keeping the fiber implanted roughly 0.4 mm from the injection site. Following injection and fiber implantation, the virus fully expressed over the course of 2 weeks. In the meantime, mice started behavioral training. After the mice attain overall high accuracy (>= 70% on average of 'Large' and 'Small' single stimulus trials) on the single stimulus economic decision-making task, we began optogenetic manipulation experiments.

For the optogenetic manipulation protocol, we used laser pulses with the following parameters: 450 nm (Neurophotometrics, FP3002), number of pulses: 10, each pulse lasting 10 ms and separated by 25 ms, laser power ~40 mW (measured at the tip of the optic fiber) (**Fig. 4.1c**). These laser pulses occurred during a behavior session on 30% of randomly chosen trials within a session. Furthermore, these laser pulses were applied from the time of stimulus presentation, for a fixed duration of 0.3 seconds. The optogenetic experiments involved unilateral inactivation of target regions therefore, only one hemisphere was inactivated on a specific experiment day, and we switched which hemisphere we targeted each day.

4.2.5 Histological verification of anatomical localization of fiber placement and viral expression

In order to verify the expression of the injected ChR2 virus in both the mPFC and the OFC as well as the location of the implanted fiber, we performed post-mortem histology (**Fig. 4.1a,b**). Mice were terminally anesthetized and perfused with 4% paraformaldehyde in 0.1M PBS followed by overnight fixation of the head to allow for more careful removal of the optical fibers. The following day, we extracted the brain tissue and carefully removed the implanted probes. The brain tissue was then fixed in paraformaldehyde for 24 hours and kept at 4°C during this overnight fixation. Following overnight fixation, the brain was embedded in 2% agarose and sliced at a thickness of 70-μm using a vibratome (Leica Microsystems). In the next stage the brain was washed in PBS and

subsequently stained for DAPI (Merck, MBD0015). The DAPI-stained slices were then washed with PBS before mounting onto glass adhesion slides and cover slipped with the mounting medium. The mounted slides were then left to dry overnight. The following day, we carried out fluorescence imaging using a fluorescence microscope (Leica) and images were taken using a 1.6 and 10X magnification to image the DAPI stain, the fluorescence of the fluorophore attached to the channelrhodopsin virus and the optic fiber tracts.

4.3 Results

4.3.1 Effects of optogenetic manipulation on economic choice

In order to understand the role of frontal areas such as mPFC and OFC in economic choices, we observed how manipulation of these areas affects choices. Manipulation experiments were carried out using mice expressing Channelrhodopsin-2 (Chr2) in Pvalb-expressing inhibitory neurons. This approach allowed us to activate parvalbumin fast-spiking (PV) interneurons that express ChR2, thereby inactivating targeted brain areas (Guo et al., 2014). We chose this protocol because activating GABAergic interneurons, specifically PV interneurons, induces a reliable shutdown of neural activity. Examining the performance of animals in either control ('Laser off') and inactivation ('Laser on') trials confirmed that mice were able to perform the action to obtain reward on single stimulus trials by moving the stimulus to the center of the screen, as signified by high accuracy levels on both trial types (**Fig. 4.2a;** 0.90 $\pm$ 0.02, and 0.87 $\pm$ 0.03, mean $\pm$ SEM for 'Laser off' and 'Laser on' trials, respectively). Crucially, there was an effect of laser stimulation: mice consistently displayed a decrease in accuracy on trials where the laser was turned on (**Fig. 4.2a,** accuracy difference - ΔLaser off - Laser on: $t(5) = 5.43$, $p = 2.87 \times 10\text{-}3$). These results confirmed that mice exhibited an effect on their choice accuracy on single stimulus trials regardless of which region we stimulated.

Having observed an effect of laser stimulation on performance, we next asked whether this decline in performance is stronger in one targeted region than in the other. Examining the single stimulus performance split by targeted region, revealed that inactivating mPFC did not significantly influence choice behavior (**Fig. 4.2b;** 0.90 ± 0.02, and 0.87 ± 0.03, mean ± SEM for 'Laser off' and 'Laser on' trials, respectively, accuracy difference - ΔLaser off - Laser on: $t(4) = 2.41$, $p = 0.07$). Even though the results for the mPFC manipulation did not reach significance, there is a noticeable decrease in choice accuracy when comparing control versus inactivation trials. Inactivating orbitofrontal cortical neurons had a similar effect on performance but reached statistical significance (**Fig. 4.2c;** 0.90 ± 0.03, and 0.88 ± 0.03, mean ± SEM for 'Laser off' and 'Laser on' trials, respectively, accuracy difference - ΔLaser off - Laser on: $t(5) = 4.8$, $p = 4.89 \times 10\text{-}3$). When we test the difference between regions, we find that the interaction is not significant (Δ accuracy difference - ΔmPFC - OFC: $t(4) = 0.43$, $p = 0.69$). Overall, we find that laser stimulation reduced performance largely independent of which region was targeted.

Having observed that laser stimulation decreases performance, we next asked whether the mouse's understanding or computation relating to economic value is influenced or changed by laser stimulation. We hypothesize that during laser stimulation-induced inactivation of regions instrumental to encoding of economic value, the difference in accuracy between Large and Small conditions will diminish suggesting that there is disruption of the comprehension of economic value. To test this hypothesis, we examine the difference in performance on 'Large' and 'Small' single stimulus trials dependent on whether the laser was on or off during a trial. Examining the difference in performance on these different EV conditions split by laser condition revealed that mPFC inactivation numerically reduced, but not significantly diminished the accuracy difference between 'Large' and 'Small' single stimulus conditions (**Fig. 4.2d;** 0.06 ± 0.02, and 0.04 ± 0.03, mean ± SEM for 'Large - Small: Laser off' and 'Large - Small: Laser on' trials, respectively; Δ accuracy difference - Δlaser off - laser on: $t(4) = 0.98$, $p = 0.38$). Inactivating OFC numerically increased the accuracy difference between 'Large' and 'Small' single stimulus conditions, but this effect likewise

did not reach significance (**Fig. 4.2e;** 0.05 ± 0.02, and 0.07 ± 0.01, mean ± SEM for 'Large - Small: Laser off' and 'Large - Small: Laser on' trials, respectively; Δ accuracy difference - Δlaser off - laser on: $t(4) = -1.96$, $p = 0.11$). Interestingly, when we examine the interaction effect between the regions we stimulate and the difference in laser off and on accuracy for 'Large - Small' trials, we notice that there is a tendency towards opposing effects of turning on the laser in particular regions but this was not statistically significant. Specifically, we note that turning the laser on in the mPFC regions results in a smaller accuracy difference between different EV conditions. Meanwhile, inactivating OFC results in a larger accuracy for different EV conditions(**Fig. 4.2f;** -0.01 ± 0.01, and 0.03 ± 0.01, mean ± SEM for mPFC and OFC, respectively; Δ accuracy difference - ΔMPFC - OFC: $t(4) = -2.07$, $p = 0.11$); ($F(1, 4) = [4.287]$, $p = 0.107$)). While the region x laser interaction was not statistically significant, 4 out of 5 mice exhibit an increased difference between Large and Small choice accuracy when stimulating OFC compared to mPFC. This suggests that there could be spatially specific modulations of economic choice behavior, which needs to be confirmed with larger sample sizes.

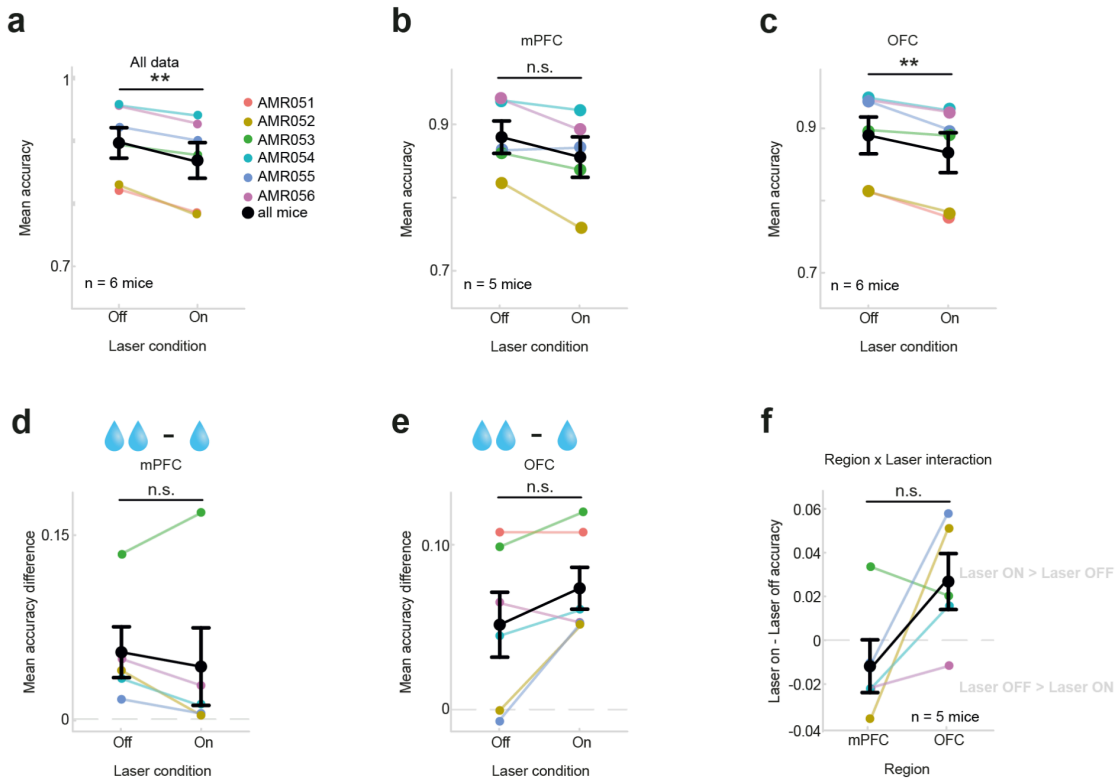


Fig. 4.2 | Effects of unilateral inactivation of OFC and medial PFC areas

a, Behavioral choice accuracy for pooled data from both single stimulus conditions split by control (laser off) and inactivation (laser on) trials (n = 6 mice). Colored dots denote individual mice, and black dots show group averages. Error bars in all panels depict SEMs. Mice overall exhibit high accuracy for both control and inactivation trials, suggesting that they are able to successfully move a single stimulus to the screen center. Despite this overall high level of accuracy in both laser conditions, mice are significantly more accurate in the control condition (laser off) compared to the inactivation condition (laser on). **b,** Behavioral choice accuracy for mPFC data for both single stimulus conditions split by control (laser off) and inactivation (laser on) trials (n = 5 mice). Inactivation trials did not show a significant impact on choice accuracy. **c,** Behavioral choice accuracy for OFC data for both single stimulus conditions split by control (laser off) and inactivation (laser on) trials (n = 6 mice). OFC inactivation significantly disrupted choice accuracy. **d,** Difference in mPFC data choice accuracy between two single stimulus conditions ('Large' and 'Small') split by control and inactivation trials. Mice do not exhibit a larger difference in accuracy between these different EV conditions, compared to trials where mPFC is inactivated. However there is a downward trend exhibited by 4 out of 5 mice with a smaller mean accuracy difference for laser on trials. **e,** Difference in OFC data choice accuracy between two single stimulus conditions ('Large' and 'Small') split by control and inactivation trials. Mice do not exhibit a larger difference in accuracy between these different EV conditions, compared to trials where OFC is inactivated. Among 4/6 mice there is an increase in mean difference in accuracy for inactivation trials. Having noticed a smaller mean accuracy difference for different EV conditions when inactivating mPFC and a larger difference when inactivating OFC, we asked whether there is an interaction effect between the region we inactivated and the laser setting. **f,** Difference in the accuracy between two laser conditions (laser on trials

accuracy - laser off trials accuracy) split by inactivated region. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, **** $p < 0.0001$, one-sample or paired t-tests.

4.3.2 Effects of optogenetic manipulation on response times

Examining the response times - computed by taking the difference between the choice complete time and the stimulus presentation time - for each trial, across both single stimulus conditions revealed overall fast reaction times. We asked whether inactivation of mPFC regions has an effect on these response times by comparing the mean response times between the two laser conditions. There was no difference between the mean response times for laser stimulation and control trials (**Fig. 4.3a**; 0.25 ± 0.01, and 0.26 ± 0.01, mean ± SEM for 'Laser off' and 'Laser on' trials, respectively, response time difference - ΔLaser off - Laser on: $t(4) = -0.33$, $p = 0.76$). Similarly, for OFC stimulation, there was no significant effect observed on response times during laser stimulation compared to control trials (**Fig. 4.3b**; 0.26 ± 0.02, and 0.25 ± 0.01, mean ± SEM for 'Laser off' and 'Laser on' trials, respectively, response time difference - ΔLaser off - Laser on: $t(5) = 1.49$, $p = 0.20$). Overall, examining response times across all single stimulus conditions revealed that there is no effect of optogenetic stimulation.

Having observed that laser stimulation does not impact response times regardless of which region is stimulated, we next asked whether the response times related to economic value are impacted by laser stimulation. Previously in our main behavioral results chapter, we showed faster response times for higher EV stimuli compared to lower EV stimuli(**Chapter 2, see Results section 2.3.1**). We reasoned a priori that laser stimulation should diminish the difference in response times between the 'Large' and 'Small' condition. To test this hypothesis, we examine the difference in response times on 'Large' and 'Small' single stimulus trials dependent on whether the laser was on or off. Examining the difference in response times on these different EV conditions split by laser condition revealed that mPFC stimulation did not significantly diminish the response time difference between 'Large' and 'Small' single stimulus conditions (**Fig. 4.3c;** -0.02 ± 0.01, and -0.02 ± 0.01, mean ± SEM for 'Large - Small: Laser off' and 'Large - Small: Laser on' trials, respectively; response time

difference: ΔLaser off (Large - Small) - Laser on (Large - Small): t(4) = -0.021, p = 0.85). Similarly, we observed no significant difference between the response times of the different EV conditions when optogenetically inactivating OFC (**Fig. 4.3d;** -0.05 ± 0.03, and -0.04 ± 0.02, mean ± SEM for 'Large - Small: Laser off' and 'Large - Small: Laser on' trials, respectively; response time difference: ΔLaser off (Large - Small) - Laser on (Large - Small): t(5) = -1.50, p = 0.20). However, since we did not observe a significant difference between Large and Small laser off trials (**Fig. 4.3e,f;** mPFC: (F(3, 12) = [2.291], p = 0.13); OFC: (F(1.04, 5.19) = [2.547], p = 0.17)), we were not able to replicate our previous findings in Chapter 2 of faster reaction times for the Large compared to the Small trials. Therefore, overall we did not observe any effect of laser inactivation on reaction times and additional experiments are required to further investigate this. Therefore, the current data provides no evidence that inactivating frontal regions disrupts economic evaluation.

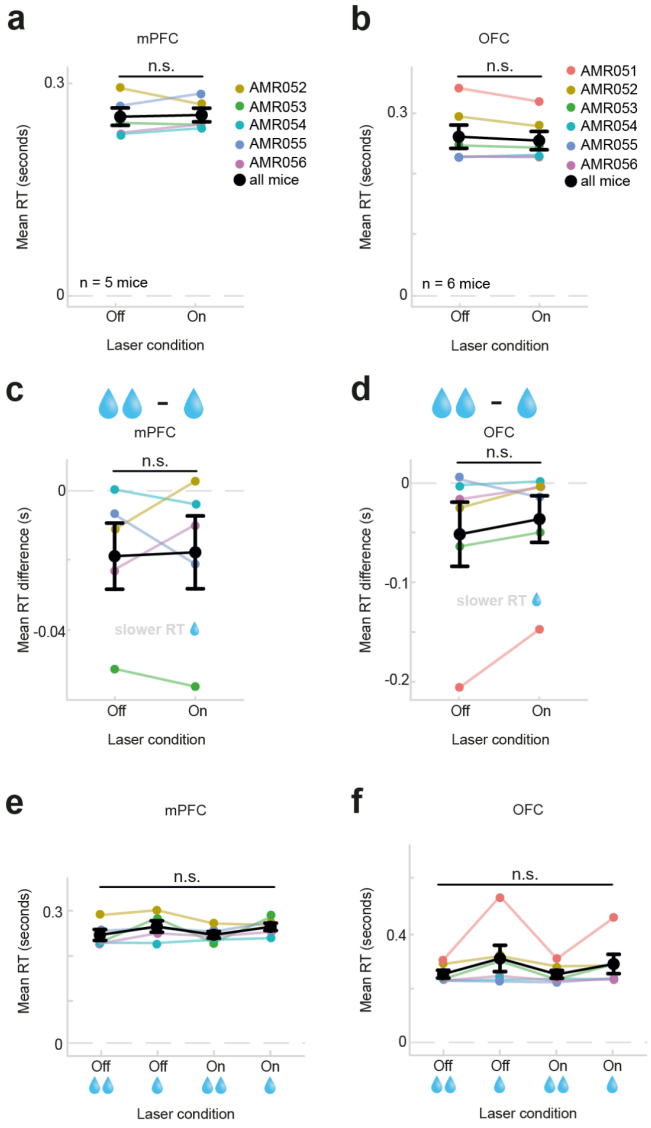


Fig. 4.3 | Effects of unilateral inactivation of OFC and medial PFC areas on response times

a, Mean response times for mPFC data from both single stimulus conditions split by control (laser off) and inactivation (laser on) trials (n = 5 mice). Colored dots denote individual mice, and black dots show group averages. Error bars in all panels depict SEMs. Mice overall exhibit comparable response times for both the control and the inactivation trials, suggesting that inactivation of mPFC does not impact the response times in making choices overall in the single stimulus portion of the task. **b,** Mean response times for OFC data for both single stimulus conditions split by control (laser off) and inactivation (laser on) trials (n = 6 mice). Similar to the mPFC inactivation results, OFC inactivation did not result in a significant change in response times. **c,** Difference in mPFC response times between two single stimulus conditions ('Large' and 'Small') split by control and inactivation trials. Mice do not exhibit a significant difference in response times between different EV conditions when comparing laser on and off trials. **d,** Difference in OFC response times between two single stimulus conditions ('Large' and 'Small') split by control and inactivation trials. Mice do not

exhibit a significant difference in response times between different EV conditions when comparing laser on and off trials. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, **** $p < 0.0001$, one-sample or paired t-tests. **e**, Mean reaction time for mPFC data from all four laser/ single stimulus conditions (n = 5 mice). Colored dots and lines denote individual mice, and black dots and lines denote the group average. Error bars in all panels depict SEMs. Overall, mice display comparable reaction times for both control and inactivation trials for both Large and Small single stimulus conditions. **f,** Same as (e) but with data from OFC inactivation experiments (n = 6 mice). Notably, there was no difference in reaction times between the Large and Small control trials, and similarly, there was no significant decrease in the delta reaction time between 'Large - laser on' and 'Small - laser on'.

4.4 Discussion

Extensive research has shown that frontal cortical regions play an important role in guiding economic decisions. In the previous chapter, we demonstrate that the neuronal representation of economic variables such as EV and risk are present across all frontal regions and assume different functional gradients. Since we show that the representations of economic value exist across different frontal regions, we asked whether the representations in these regions play a casual role in encoding decision variables. We hypothesize that inactivating frontal regions during this economic decision-making task would affect choice accuracy and in particular disrupt choices that rely on computation of economic value. Specifically, inactivation would disrupt the difference in choice accuracy on different EV conditions relative to control trials. In addition, we noted similar decodability of EV from neural activity in PL and OFC (**Chapter 3, Fig. 3.6a**) - two regions we directly inactivate in our optogenetic inactivation experiment. Therefore, we assume that inactivation of these areas may show similar disruption of economic decisions.

In order to test our hypotheses, we designed a modified version of the economic decision-making task detailed in **Chapter 2** and trained mice to make choices with two single stimulus conditions which have a difference in EV. Once mice are fully trained in this task, we unilaterally inactivate either one of two target regions: mPFC and OFC. We present three main results. First, we show that laser stimulation affects choice accuracy irrespective of the region inactivated. Specifically, optogenetic inactivation resulted in a significant decrease in choice accuracy across 6 mice in the

dataset. Second, we found inconclusive results that inactivating specific frontal regions disrupts economic evaluation. Optogenetic inactivation of OFC consistently decreased choice accuracy. In contrast, inactivation of medial PFC areas did not significantly decrease choice accuracy. Third, we show that optogenetic inactivation of OFC and mPFC show a trend towards opposing effects on evaluation of expected value which is not statistically significant. Specifically, inactivation of mPFC induced a decrease in delta accuracy between Large and Small single stimulus conditions, while inactivation of OFC resulted in a larger difference between these different EV conditions. Taken together, these results suggest that inactivation of frontal regions during a time window post stimulus presentation may cause a disruption of economic choice behavior, but overall our observations are not statistically significant and require further experimentation.

It is particularly important to note that previous studies have sought to examine the causal roles of representations across frontal cortical areas in economic decisions. For example, Kuwabara et al. (2020) trained mice in an economic decision-making task and showed that optogenetic inactivation of OFC disrupted economic choice. In particular, they noted that OFC manipulation increased choice variability. In addition, lesion studies and inactivation studies from the past decades indicate the crucial role that frontal regions such as OFC play in decision-making. These studies have shown that inactivation of OFC disrupts performance on decision making tasks (O'Doherty et al., 2014; Padoa-Schioppa and Schoenbaum, 2015). However, there are other studies that have conflicting findings, showing that inactivation of OFC does not disrupt economic choice (Gardner et al., 2017).

There are several limitations to our task design and our optogenetic study. In the following paragraphs I will briefly outline these limitations alongside proposed experimental modifications for refining our experimental design for future causal manipulation experiments.

In Chapter 3, we found that stimulus conditions were decodable above chance level shortly after stimulus onset, but were most robustly and stably decodable in a time window from 150 to 350 ms after stimulus onset. Given that we established a time window during which frontal neuronal responses differentiated between visual stimuli, we reasoned that by optogenetically manipulating neural activity within this time window, we may disrupt economic evaluation of stimuli. Thereby, we optogenetically stimulated PV interneurons in a fixed time window for 300 ms post stimulus presentation. Our study stands in contrast to other inactivation studies since we use brief optogenetic pulses compared to other studies which use a larger time window for inactivation (Bao et al., 2023; Kuwabara et al., 2020; Constantinople et al., 2019b; Gore et al., 2023; Ballesta et al., 2020; Gardener et al., 2017; Gardner et al., 2019). Since we tried to stay close to the time window of neural representations in our study, we might not have manipulated the entire time window during which these economic valuation computations are occurring and perhaps in future studies we can modify and integrate a larger time window of optogenetic inhibition. In addition, in our optogenetic inactivation protocol we only used unilateral inhibition. This deviates from other inactivation studies where both bilateral (van Holstein et al., 2020; Pinto et al., 2019; Verharen et al., 2020; Cole et al., 2024) and unilateral inactivation trials are randomly interleaved within a session. For future versions of our optogenetic inactivation, we will integrate bilateral inactivation with each session.

In Chapter 3 of this thesis, we find a medial-to-lateral gradient for representation of risk and a dorsal-to-ventral gradient for representation of EV, hence demonstrating the strongest representation of these variables in the neuronal activity of dorsomedial frontal cortex (dmFC). Therefore, the next step in examining the causal role of these neuronal representations would involve studying how bilateral inactivation of the dmFC impacts choice behavior. In the scope of our pilot optogenetic inactivation study we did not inactivate the dmFC which had the strongest encoding. However in the near future, I will aim to inactivate ACC both unilaterally and bilaterally in order to ascertain whether inactivating the areas with the highest degree of representations

impedes economic decision-making. Furthermore, one of the frontal areas that we target in our study is the PL, located in the medial portion of the frontal cortex. Previously we established the strongest representation of risk in the medial portion of the frontal cortex compared to the lateral areas. In our optogenetic study, we have not included a risky Gamble stimulus, thus we are not investigating how optogenetic inhibition of the areas with the highest representations of risk may influence the risk attitude of each animal. This could be a potential avenue for future studies. In fact, a study by Passecker et al., (2019), showed that optogenetic silencing of the prelimbic cortex exclusively timed to the encounters of no reward, provoked animals to excessive gambling for large rewards. Thus, I will plan to reintegrate the single stimulus condition with probabilistic reward outcomes (Gamble) in order to understand the causal role of frontal areas in risky decision-making.

Another plausible explanation for our study not being able to establish the causal roles of different frontal neural representations in economic decisions may be due to the use of over-trained mice during the optogenetic experiments. In particular optogenetic inactivation experiments were performed when the mice reached expert performance on this task,defined as displaying ≥80% across all trial types. It may be plausible that optogenetic inactivation would have a greater effect on such neural representations and subsequently on choice behavior during learning of such a task. In fact a previous study by Lak et al., (2020) highlights the role of mPFC in predicting value signals during a task that requires continuous learning. While the structure of our task and the use of over trained mice may offer a plausible explanation for the lack of effect of optogenetic inactivation during a portion of the task where there may be less learning of abstract rules, our results nonetheless point towards other avenues of refining our task design in order to study the causal role of frontal regions in economic choices.

Another concern is that our viral injections are variable, meaning that they could extend outside of the areas that we are targeting, leading to the concern about whether the inactivation effect is solely driven by either mPFC or OFC regions. Lastly, another concern in our study is the scarcity of PV+

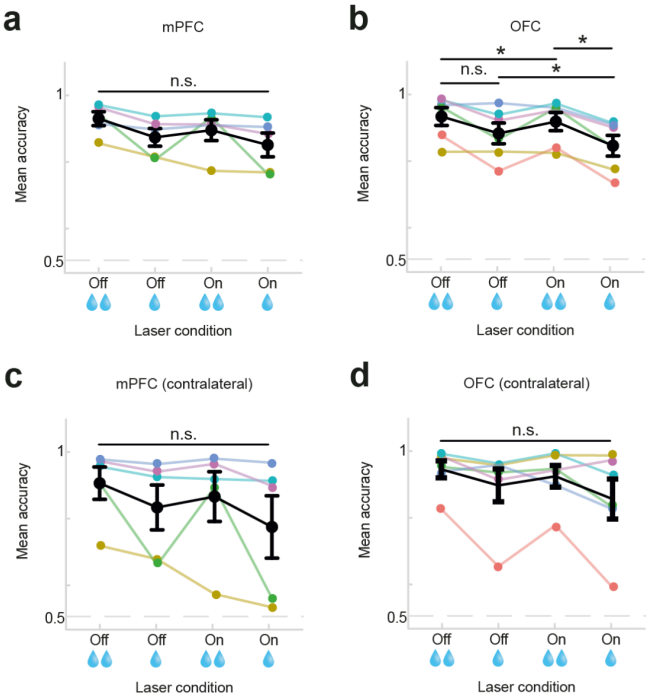
interneurons in medial frontal regions such as PL compared to other regions where there is a more prominent presence of this subpopulation of interneurons (Kamalova et al., 2024; Kim et al., 2017). Since our approach is reliant on effective silencing of neural activity in these areas by using PV interneurons, it would be crucial to establish how this impacts the mechanisms of inhibition in an area.

The timescales of representations are also integral to the decision making process and can differ across distinct regions. In order to target the time window of these representations and see whether this disrupts choices we need tools to silence neural activity with temporal specificity. Optogenetic inactivation of neural activity which uses the millisecond precision of temporally activating or inactivating neural activity is important for experiments where manipulation is targeted at specific time points. Overall, to study how neural representations are causally related to economic choices, we not only need the right silencing tools, but also the genetic technology to enable both spatial and temporal specificity during the interrogation of circuits. More work needs to be done to characterize the causal role of multiple frontal regions in economic decision making and to further understand the structure of the neural circuit and interactions between these regions. It is crucial to understand how representations in each region inform connected regions and regulate the decision-making process.

4.5 Acknowledgements

This work was carried out with assistance from Rathi Puliyadi (lab manager in the Lak Lab) and Olena Didenko (research assistant in the Lak Lab). I conducted the injections and fiber implantation surgeries. Rathi and Olena trained mice in the initial shaping stages of the economic decision-making task. Prof. Armin Lak supervised the optogenetic experiments. Dr. Matthias Fritsche helped set up the MC code for running the economic decision-making behavior task with a modification for stimulation parameters.

4.6 Supplementary Figures



Supplementary fig. 4.4 | Effects of unilateral inactivation of OFC and medial PFC on individual single stimulus conditions

a, Mean choice accuracy for mPFC data from all four laser/ single stimulus conditions (n = 5 mice). Colored dots and lines denote individual mice, and black dots and lines denote the group average. Error bars in all panels depict SEMs. Overall, mice display comparable choice accuracy for both control and inactivation trials for both Large and Small single stimulus conditions. **b,** Same as (a) but with data from OFC inactivation experiments (n = 6 mice). Mice show a significant decrease in choice accuracy on 'Large - laser on' trials compared to 'Large - laser off'. Additionally, mice show a similar and significant decrease in choice accuracy for 'Small - laser on' trials compared to 'Small - laser off' trials. Notably, there was no difference in choice accuracy between the Large and Small control trials, but comparatively, there was a significant decrease in choice accuracy between 'Large - laser on' and 'Small - laser on'. **c,** Having observed that value representations assume a lateralized manner and the representations are stronger in contralateral trials - we next asked whether inactivation of frontal regions has a stronger impact on disrupting choice accuracy on contralateral trials. Mean choice accuracy for mPFC contralateral trial data from all four laser/ single stimulus condition combinations. Overall, stimulation of mPFC does not result in a significant change in choice accuracy relative to control trial accuracy. **d,** Same as (c) but with OFC contralateral data. OFC stimulation does not impact the mean choice accuracy relative to control trials.

5

General Discussion

The goal of this thesis was to elucidate the structure of neuronal representations of decisions guided by economic variables. By using a visually guided economic decision making task, and combining it with high density electrophysiology (Neuropixels) recordings, I investigated how mice make economic choices and how neural activity across frontal areas may reflect different economic variables. Additionally, by optogenetically targeting select brain areas, my work has started to explore how disrupting the neural activity in individual frontal areas may point to a causal role of representations in these regions in making economic choices.

In **Chapter 2**, I investigated how mice make economic choices. I found that mice can make economic decisions between abstract visual stimuli. Specifically, I examined whether the animals' performance scaled with EV and was independent of the visuospatial properties of the task. Animal's choice accuracy scaled with EV and was not impacted by differences in visual stimulus features or the spatial location of the stimuli. Furthermore, lick rates were higher for conditions with higher EV during the outcome period. Collectively, these results suggest that mice learned and discriminated the value of the visual stimuli. This observation of mice showing signatures of responses scaling with EV on single stimulus trials was compounded with their performance on two stimuli trials in which mice were successfully able to choose the stimulus with the higher EV on a majority of trials. In addition, their choices satisfied choice transitivity, suggesting that not only do mice comprehend the value of the stimuli but they also make consistently rational choices. These two stimuli trials also showed that mice exhibit diverse risk attitudes that evolve over time. To

further interrogate how mice make economic choices, we asked whether two stimuli choices are impacted by a previous rewarded action or by past outcomes linked to abstract visual stimuli, independent of action. Mice exhibited a consistent action-based choice bias, repeating previously rewarded actions, in line with continual updating of action value (Lak et al., 2020; Fritsche et al., 2024). Conversely, animals showed no bias by past outcomes linked to abstract visual stimuli, suggesting that stimulus-specific value representations were more static over short timescales.

In **Chapter 3**, I examined the structure of neural representations of different economic decision variables. Specifically, in order to examine neuronal representations of economic variables we recorded the activity of thousands of neurons across frontal regions using Neuropixels probes. We found that individual neurons exhibited strong stimulus-driven responses that were modulated by trial conditions. In order to ascertain how well trial conditions could be classified from the activity of all recorded neurons, we employed a linear population decoding approach (SVMs) and found that trial conditions could be significantly decoded shortly after stimulus onset. Focusing on neural activity in this time window, we found that population activity carries information about the economic value rather than the identity of the visual stimulus and also reflects individual risk preferences. Having shown coding of economic value and risk in frontal cortical neurons we investigated how these representations are structured across the cortex. We found graded representations of EV and risk in our region-specific decoding analysis - the representations of economic variables are differentially distributed across the frontal cortex. Moving past these region specific representations, we find that the frontal cortex exhibits orthogonal spatial gradients for the representation of EV and risk. Moreover, these spatial gradients were mainly driven by neuronal responses towards stimuli presented in the visual field contralateral to the recording hemisphere. So far all of these conclusions came from neural activity from single stimulus trials with different values and risk. We next focused on the neural activity during two stimuli trials allowing us to probe whether the neural activity is reflective of post-decision variables such as chosen and unchosen value, when choosing between two alternative offers. We found that neural activity in the

orbitofrontal cortex (OFC) is modulated both by the value of the chosen and unchosen offer. Together, these findings suggest that encoding of economic variables spans across multiple frontal areas and that distinct spatial gradients underlie the representations of these variables.

Finally in **Chapter 4**, I examined the causal role of these frontal regional representations in economic choices. In **Chapter 3**, we show that different frontal regions encode economic variables in a distributed manner, so we asked whether these distributed representations are causally related to choices. To examine this, we used optogenetic approaches and inactivated frontal regions during the economic decision task. Laser-induced inactivation reduced performance largely independent of which regions were targeted. We found inconclusive results that inactivating specific frontal regions disrupts economic evaluation. However, despite the region-specific inactivation not showing any significant disruption of economic choices, a majority of mice show an increased difference in choice accuracy between different EV conditions when stimulating OFC compared to mPFC, suggesting that there might be a spatially specific modulation of economic choice behavior, which needs to be confirmed with a larger sample size.

Can mice learn the economic value of abstract visual stimuli?

This thesis has shown that mice are able to learn the task structure within the framework of an economic task which integrates and manipulates two key variables: EV and risk (derived from the probability of reward). Indeed, economic decision-making involves the computation and comparison of options each of which may have differences in their reward value or reward probability, or both. In our task, we adapt and integrate elements from different previous paradigms to create a head-fixed economic task for mice. We first adapted a task for head-fixed mice designed to probe detection of visual stimuli (Burgess et al., 2017) using behavioral task elements from economic decision-making designed for non-human primates (O' Neill & Shultz, 2011; Padoa-Schioppa & Assad, 2006). Although several economic decision tasks for rodents have been developed in the past decade (Kuwabara et al., 2020; Constantinople et al., 2019a; Bao et al., 2023;

Verharen et al., 2020; Gore et al., 2023; Gardner et al., 2017; Choi et al., 2023), our task is different because we required mice to learn about economic value using abstract visual stimuli while previous tasks used sensory cues whose intensity directly correlated with the magnitude of the rewards they are associated with. Instead our usage of abstract visual stimuli mirror the cues used in previous monkey economic tasks (O'Neill & Shultz., 2010). In this study, they present monkeys with abstract fractal images for predicting the outcome to control for the visual components. Our task is one of the first rodent economic tasks that takes direct inspiration from such non-human primate studies and ensures that mice are making choices based on economic values of stimuli, rather than their low-level sensory features. This control of visual features helps in our interpretation of neural activity as we show that neurons are encoding economic features of our task and such economic variable representations are less contaminated by sensory features. Despite the implementation of abstract visual stimuli in our task, it would be interesting to expand the number of stimuli and introduce even more ranges of reward magnitudes and probabilities. This will allow us to directly compare behavioral signatures of risk sensitivity to those observed in other rodent and non-human primate studies where they have an expanded number of stimuli. The reason we have not implemented this is due to our use of abstract stimuli and the amount of time it takes for mice to learn our task. In particular, we notice that the learning trajectories associated with learning economic values of 3 abstract stimuli already pose challenges and expanding to a larger number of conditions may considerably increase training time. However, future work may implement changes in task designs to allow for more in depth observations of risk behavior.

In **Chapter 2**, we assessed individual risk sensitivity by quantifying the preference for the risky stimulus over the safe stimulus. We found diverse risk attitudes that span from risk-seeking to risk-aversion. This diversity in risk attitudes has been explored in previous rodent studies (Zalocusky et al., 2016) - where rats made choices between risky and safe offers - with conflicting results. A study by Constantinople et al., (2019a) shows that rats display a similar diversity in risk attitudes. Rats in the Constantinople study demonstrated a reference dependence showing that following an

unrewarded gamble trial, the probability that they will switch the port they go to, increases. However, our study shows that individual risk attitude did not show a reference dependence - the risk attitude on a short-term scale was not impacted by a loss experienced in a previous trial gamble outcome. Our finding that the outcome of the previous trial did not impact the current trial probability of choosing the gamble is mirrored by another study in rats (Bao et al., 2023). In addition to the range of risk attitudes we show in our task, it is imperative to remember that there are ethological considerations that would help us understand why foraging animals are prone to risk sensitivity when considering food offerings in their environments. These behavioral phenomena have been studied and combined with an ethological/evolutionary framework - finding that foraging decisions of animals when picking between different offers in their environment is contingent on a combination of psychological and fitness consequences of making alternate decisions (Bateson et al., 2002). This motivates future studies to assess risk sensitivity as a function of these variables.

In **Chapter 2**, I found that mice make economic decisions in an action space, meaning that they exhibit a propensity to repeat a previously rewarded action (rather than stimulus) for their current choice when making a choice between two simultaneously presented stimuli. This finding was previously found using an olfactory-guided economic decision task for mice, showing that choice behavior can be influenced by factors other than the identity and magnitude of the cue (Kuwabara et al., 2020). It was shown that mice display a consistent direction hysteresis, despite it being smaller in its influence when compared to variables such as the magnitude of the stimulus. In a similar way, our study demonstrates that mice display a small but consistent direction hysteresis but virtually no stimulus-based hysteresis, suggesting that mice may assign value to recently rewarded actions, alongside a more stable value representation associated with visual stimuli. An intriguing hypothesis is whether the direction hysteresis observed predominates during a specific portion of each behavioral session. It would be interesting to study whether the action-based updating occurs in the beginning of a task, a time when they often exhibit more impulsive behavior. This distinct action-based strategy might be prominent during a ‘state’ of disengagement with the task (Ashwood

et al., 2022) wherein the mice rely less on economic strategies and instead bias their choices. In addition, future work may address whether this action-updating strategy alternates with a strategy more reliant on value computations.

What is the structure of neural representations of economic variables?

Our neural recordings revealed that different frontal regions encoded two economic variables that are integral to the decision-making process. We found in **Chapter 3** that population activity is reflective of the decision input (the value of the offer - expected value) and the probability associated with each offer (risk). In addition to these two variables, neurons encoded other variables such as the position of the stimulus (on either the right or left side of the screen), the decision input contralateral to the recorded neuron's hemisphere (offer value contra) and the decision output (the value of the chosen and unchosen offer). These decision variables for which we observe frontal cortical representations are analogous to those identified in economic task studies based in non-human primates (Padoa-Schioppa and Assad, 2006).

Large-scale electrophysiological recordings revealed that individual neurons in the frontal cortex show responses that are modulated by economic variables, and the population activity is reflective of such modulations. Importantly, we show the distribution of coding of economic variables takes place across multiple regions, with substantial quantitative differences in the level of coding of these variables between regions. For example, for EV coding we show a higher degree of decodability from anterior cingulate (ACA) neurons compared to orbitofrontal (OFC) neurons. Our finding echoes previous findings of coding of multiple decision variables across various regions of frontal cortical areas as well as quantitative differences across regions (Kennerley et al., 2009a; Kennerley et al., 2009b). Interestingly, these quantitative differences often highlight the OFC as a region that plays a dominant role in carrying neural signals related to value-based decision-making (Sul et al., 2010). However, having an OFC-centric view of value representations does not paint the whole picture of how an animal makes economic choices, as this region is connected to other frontal areas permitting

interactions that carry value representations across several areas. Even though we do not characterize interactions between the frontal areas in our study, we note the strongest representations in a dorsomedial frontal region: ACA. This finding resembles the findings of other studies that highlight the role of the ACA in encoding decision variables (Cai & Padoa-Schioppa, 2012; Balewski et al., 2023; Christopoulos et al., 2009; Kennerley et al., 2011; Luk & Wallis, 2012). Despite our observation of distributed coding of economic variables across frontal regions, we also observe some neurons in olfactory areas whose activity is reflective of EV and risk. This observation is consistent with studies that find encoding of decision variables in non-frontal areas (Bao et al., 2023; Tobler et al., 2007). Overall our study deviates from previous studies that typically found such representations predominantly in OFC, while we show a distribution of such representations across distinct regions of the frontal cortex, suggesting that other structures in the frontal cortex play an important role in economic decision-making.

Since our region-specific analyses yielded interesting findings about a distributed and quantitatively graded representation of economic variables across the frontal cortex, we abandoned the region labels in order to explore the spatial structure of these representations overlaid on a spatial map of the frontal cortex spanning across the ML, DV and AP axes. We find that the frontal cortex exhibits orthogonal spatial gradients for the representation of EV and risk with EV decoding displaying a dorsal-to-ventral gradient and risk a medial-to-lateral gradient. These findings of distinct functional gradients are consistent with a study by Maisson et al., (2021) where they describe a functional gradient along the ventromedial axis of medial frontal cortex in macaques performing an economic decision-making task with the decodability of task variables increasing in the ventral to dorsal direction, suggesting strong decoding in dorsomedial regions. Both our study and the macaque study endorse a gradient-based organization of economic variables in the frontal cortex and argue against a view where regions have distinct roles in encoding of economic variables. Our finding of strong activation of risk representations in the medial portion of the frontal cortex leaves an open question about whether this gradient differs across mice that exhibit risk-sensitive behavior and

those that exhibit risk-neutral behavior. One reason why this question may be interesting to pursue is due to a study by Tobler et al., (2006) finding higher risk-related activations in lateral OFC in risk averse individuals while risk-seeking individuals showed higher activations in the medial wall. Our study does not compare functional gradients among risk sensitive and neutral mice, as our sample size needs to be expanded for examination of such a question.

Regarding the lateralized representations, we note that values are strongly coded for visual stimuli presented in the contralateral visual field. An interesting question that emerges is: what does this lateralized distribution mean for representations in the brain? One possibility is that lateralized representations may suggest that there is competition between the two hemispheres. Importantly, there are a few studies that observe this competition between hemispheres for representation of information (van der Knaap & van der Ham, 2011). Lateralization is thought to be an evolutionary feature and ensures each hemisphere is able to process information without interference from the other. Another possibility for the function of lateralized representations is that it can allow coordinated effort between hemispheres to allow for proper integration of information and subsequently permit an appropriate response to stimuli. The lateralization of responses we observe is also a reflection of how visual information is relayed from the visual cortex to the frontal lobes (Harris et al., 2019). In our examination of lateralized representations, we focus solely on such representations in single stimulus trials, wherein the mouse has to pay attention to where the presented stimulus appears (either the right or left side of the screen). For the two stimuli trials, where two stimuli are presented simultaneously, lateralization of representations may be important. This is because in order for the two stimuli trial choice behavior to be optimal there needs to be interhemispheric cooperation and information transfer. In particular, given that we observe stronger value representations to stimuli presented contralaterally in single stimulus trials, it would be interested to learn if the value of stimuli presented contralaterally in two stimuli cases assume a similar stronger representation relative to stimuli presented ipsilaterally relative to the recording hemisphere. Previously fMRI studies have shown greater bilateral activation in complex tasks

(equivalent to our two stimuli trials) versus simpler tasks (equivalent to our single stimulus trials) (Braver et al., 1997; Banich et al., 1998). Therefore in the future, more work needs to be done to assess how these lateralized responses emerge in a choice task that integrates more than one choice option. Additionally, to address the interhemispheric communication or crosstalk between hemispheres, we can address this by bilaterally recording from similar regions and understanding how value information is distributed among the two hemispheres to allow for efficient and rational choices to be made.

Do frontal cortical representations of economic variables have a causal role in economic choices?

We used optogenetic inactivation methods (**Chapter 4**) to inactivate frontal areas in order to interrogate whether the previously observed representations of EV in frontal areas (**Chapter 3**) play a causal role in economic decision making. Inactivation overall impacted choice accuracy but there was no conclusive evidence for a difference between inactivating OFC and mPFC. While our study did not show any disruption of economic choices induced by optogenetic inactivation, previous studies demonstrate the causal role of regions in making economic decisions. A previous study by Bao et al., (2023) demonstrates that inactivation of the frontal orienting field (FOF) in rats during an economic task reduced the probability of rats picking the riskier option, suggesting that FOF is crucial for making choices under risk. Importantly, there are multiple explanations for why they found a positive effect compared to our null effect. In this experiment, they used one of two methods: pharmacological methods using muscimol infusions and optogenetic silencing of the FOF using halorhodopsin. First, the muscimol inactivation is a drastic pharmacological interference because it lasts for the entirety of the session rather than for specific periods within a trial, thus it is probably a stronger manipulation than what I have used and could explain why they have found an effect. However, this method is not a temporally targeted intervention for investigating critical time periods during decision formation. For investigating the causal role of the representations we observed in a specific time window post-stimulus onset as described in **Chapter 3**, we need

temporally precise methods such as optogenetics. Second, notably although inactivation using halorhodopsin is less effective than inactivation through channelrhodopsin-induced activation of PV interneurons (Raimondo et al., 2012; Witten et al., 2017), an approach which we employed in our task design, they do find an effect. Despite the limitations of these approaches in the rat risk decision task study, another study in mice used channelrhodopsin stimulation of PV interneurons in OFC and showed significant changes in relative value and disruption of choice accuracy (Kuwabara et al., 2020). Of note, in this study in mice they introduced the stimulation of PV interneurons starting from the time of presentation of the odor cue and through the entire offer period plus 1 second extra. In contrast, our inactivation protocol employs a fixed laser stimulation of 300 milliseconds throughout the offer period, as we found a time period during which stimulus conditions were decodable above chance from 150 to 350 milliseconds post stimulus presentation (shown in **Chapter 3** results). Taking a step back and comparing our manipulation to that in Kuwabara's study points out the obvious that the latter study uses a longer stimulation period compared to the time period used in our optogenetic manipulation experiment. It is possible that if we were to implement a similarly long time period for inactivation, then we might find stronger effects on economic decisions.

Furthermore, we reason that by disrupting neural activity in a specific time window post stimulus presentation, that we will impact evaluation of economic value. Our line of reasoning may have several shortcomings. **First**, we unilaterally inactivate regions, whereas bilaterally inactivating regions may prove to have a larger effect on behavior and prevent any compensatory effect from the non-stimulated hemisphere to take place. A question that unilateral inaction may be able to address is: whether choice behavior for stimuli presented contralateral to the inactivation hemisphere impaired to a larger relative to choice behavior for stimuli presented ipsilaterally to the inaction hemisphere. We hypothesize that unilateral disruption might have more severe consequences for contralateral stimuli as we previously observed the representations of EV are strongest for contralateral trials (see Chapter 3 results), however in the scope of our task we did not acquire

enough data to analyze this. However, overall it would be useful to compare how unilateral inactivation compares to bilateral inactivation with a focus on the question of whether one has a higher impact on disrupting economic choices. **Second**, these representations, although robust in this offer period, are not just confined to this time period as these representations need to be transformed into an action plan. **Third**, we are inactivating using PV interneurons, and this population of inhibitory neurons may not have the same prevalence in the areas that we are inactivating in this study, particularly in mPFC (Kamalova et al., 2024). **Fourth**, we have not considered whether there is any behavioral impact of rebound excitation in our study as we stimulate for a fixed period but the post-stimulation period of neural activity might play a role in these economic decisions. Therefore in a future cohort it would be useful to account for any variability in injection sites across mice and its potential impact on disruptions to choice behavior. Overall, our causal experiments are an important next step but require more refinement. As we observe the highest strength of representations in the dorsomedial areas in the frontal cortex, our next experiments should bilaterally and unilaterally target areas such as ACA and observe whether this disrupts economic choices.

5.1 Future directions

Is there a switch from an action bias to an unbiased economic strategy within a session?

In our preliminary analyses, we identify a bias to repeat previous actions but no stimulus-based bias which suggests that the mouse's understanding of value is relatively stable over the course of a short timescale. One question that emerges is whether this action bias is reduced during certain phases of the task so that the mouse becomes unbiased and mainly makes choices based on the value associated with the stimulus. Overall, we find that mice can make economic choices using these abstract visual stimuli but in the scope of our study we do not characterize fluctuations in behavior during our task. We intend to understand if there are any behavioral fluctuations in our task and if so, quantify this phenomena in a descriptive manner within and across mice. We hypothesize that fluctuations in choice accuracy on single and two stimuli trials might point towards overall

engagement with the task and point towards different strategies that mice rely on to make choices at different stages of the task. For example, one question is whether, in the beginning of the session when the mice are exhibiting more random choices and more impulsive behavior, we can characterize whether the previously described action bias dominates at that time? Perhaps mice reach a threshold at which this action-updating switches to a state where mice mainly care about the economic values of the stimuli with little to no contamination of choice biases related to abstract stimuli when they are highly engaged in the task.

Does pre-stimulus neural activity predict a mouse's choice behavior?

One interesting question for future research is how neural activity predicts temporally local performance on this economic decision-making task. We want to ask whether we can predict performance on single and two stimuli trials using activity from frontal areas. Past studies have often found that different frontal regions integrate information about past rewards and also help with updating choices (Constantinople et al., 2019b; Wallis and Kennerley, 2010; Shohamy and Daw, 2015; Lee et al., 2007). It would be useful to examine how neural activity during the inter-trial interval may signal different variables such as future choices or the accumulation of reward.

How do neural representations of economic value representations get established during learning?

In addition to examining how neural activity may predict behavior on a short-term scale, we can extend into a long-term learning study. In particular, one can ask how these spatial representations of EV and risk are shaped across learning. In our task, we only look at these representations in an expert mouse - one that has learned the rules associated with our economic task. In our study, we identify different degrees of representations of economic value and risk across various regions in an expert animal. However, by utilizing chronic electrophysiology and imaging methods, we can study how these representations may change across learning (Reinert et al., 2021; Miller et al., 2022). For example, perhaps the OFC may be more strongly implicated in the learning phases of the task and

later shows a smaller activation for different representations compared to medial and dorsal frontal areas. Perhaps the mapping of the activations and the spatial gradients may shift across learning suggesting that information flow changes as mice learn different abstract task rules and learn to associate different stimuli with different reward magnitudes and probabilities.

Is the neural population encoding expected value overlapping with neurons that encode chosen value?

It would be interesting to see whether there are separate populations guiding single and two stimuli performance. Specifically, we want to ask whether there is generalization between single and two stimuli trials - can you decode on EV using single stimulus trials and use the readout weights to decode the chosen value on the two stimuli trials. Our preliminary results of the two stimuli decoding analysis, where we observe that OFC can encode information about the two stimulus values more strongly in relation to MOs compared to single stimulus trials, might already suggest that these neural populations might be partially non-overlapping. However understanding whether the degree to which these populations are overlapping would be an interesting avenue of future research.

What are the neural representations underlying outcome processing?

In my thesis we have this far focused on the neural signals before the outcome was delivered in order to decode pre-outcome variables. However there may be a distinction between these pre-outcome variables which reflect the internal representation of value or reward expectation versus post-outcome variables which may reflect reward prediction errors. Thus this prompts the question: are the same neurons that encode pre-outcome variables (i.e. EV, risk, chosen value, stimulus location) also involved with processing the outcome, for instance encoding the difference between actual and expected outcome? To test this, we can develop a behavioral paradigm wherein there are random violations of reward expectation. An example of this would be if the mouse is expecting a large, guaranteed reward of 3 μl, but then this expectation is randomly violated and set to 1.5 μl

instead. In this case, one can study how the neural activity of a neuron whose responses are modulated by the size of the reward are changed when violating the expected reward. Another way of testing this would be to analyze this in the context of the risky stimulus condition (Gamble) which delivers 3 μL or no reward with equal probability. Given this switch between a certain and uncertain reward, we can ask whether individual neurons' outcome responses are different between these two scenarios. Therefore, a future direction would be to analyze the outcome aligned responses in this task which we have not yet analyzed.

How do prefrontal areas exchange information?

We are interested in characterizing the dynamics of information flow across prefrontal cortical areas. In **Chapter 3**, we show distinct spatial gradients for different economic variables. We observe a distinct medial-to-lateral gradient in risk with the highest decoding weights observed in the medial portion of the frontal cortex and for EV we observed a distinct dorsal-to-ventral gradient with the highest activation weights observed in the dorsal portion of the frontal cortex. The next question we can ask using these gradients is whether information flow is larger for risk information between medial areas compared to lateral areas and likewise for information exchange about EV being more pronounced in the dorsal areas relative to ventral areas. It would be interesting to observe whether the information flow also has a certain directionality within the part of the gradient with the strongest representations of such economic variable information. We can hypothesize that perhaps risk information flows more from areas such as the ILA towards areas such as the ACA, a region that has been previously implicated in relaying task relevant information for formation of choices. By studying how information flows between regions may allow for predictions about which areas are more strongly interconnected. One main caveat is that these orthogonal EV and risk spatial gradients are computed across a large time window spanning from 150 to 350 milliseconds post stimulus presentation. Therefore, it is currently unclear how information is routed through the network of frontal cortical areas. In our current study we were not able to address this issue due to the method we implemented for data collection. We utilized a single shank Neuropixels probe,

therefore, we have a limited set of regions that we were able to record from simultaneously. There are recent developments such as the 4-shank Neuropixels probe (Steinmetz et al., 2021) which allow for greater coverage of frontal cortical areas and allow for recording neural activity across even more regions simultaneously. Another issue related to why we have not yet looked into information flow across these frontal areas in our study is the number of trials in our task. Typically in functional connectivity studies, one needs a large number of trials because it allows for reliable observations of co-fluctuations of neuronal activity Therefore, one would need a high number of trials per condition to see how these neurons co-fluctuate across time and to garner whether there is information exchange occurring in a specific way. The main bottleneck in our task is that mice typically perform on average 300 to 400 trials per session, yielding relatively low numbers of trials per task conditions. To address this limitation, in a future experiment we can run a reduced task such as having only high EV and low EV single stimulus conditions (Small and Large) or a task with only different EV two stimulus conditions (Large versus Small). Another task modification that can be implemented may involve reducing the volume of reward that the mice receive. Currently the mice receive either 1.5 or 3 μl of water but there is a possibility of changing it to 1 and 2 μl instead. This modification may motivate mice to perform more trials to meet their daily water requirement. Both these proposed modifications may allow for analyses on how information is transferred across prefrontal areas in our task. In general, it would be great to integrate chronic recordings into our task design to study the question about information flow across areas, mainly due to the nature of such recordings: we can record from and track the same neuron across several days, which would allow for more trial repetitions. In fact there are studies that have implemented chronically implanting Neuropixels probes (Song et al., 2024; Juavinett et al., 2019) and such studies are able to track a large number of isolated units across recording days. Another possible method that can be used is 2-photon imaging of cells located in the dorsomedial regions of the frontal cortex or implanting GRIN lenses (Hjort et al., 2024) in order to image deeper regions across multiple sessions. Overall, chronic recordings or imaging techniques would enable longitudinal tracking of

neurons across days, yielding a higher number of trial repetitions and also enable functional connectivity analysis.

Are there other principles governing the structure of economic value representations?

In addition to the representations of economic variables being present across regions of the frontal cortex, other possibilities for the structures of these representations may include a projection-level and cell-type specific level of representations.

During economic decision making, several regions are shown to carry representations of economic variables. However, it is unclear how such representations are connected across the brain. During decision-making, it is interesting to study how different regions are interacting with one another and carry forward representations of variables involved in decision-making paradigms such as value. Previous studies characterize the interaction between frontal regions ACC and OFC during value-based decision-making and find that during behaviorally relevant time scales, the ramping activity of ACC is impacted by the representation of value in the OFC (Balewski et al., 2023). Another level of representation of economic variables in value-based decision-making involves studying the projections from frontal regions to regions where values can be transformed into action outputs, such as the fronto-striatal circuits. Prelimbic cortex (PL) is a frontal region that has robust projections to the striatum (particularly the dorsomedial striatum, DMS, and the dorsolateral striatum, DLS). This connection has been implicated in distributed coding of value-based signals - in particular, the PL-DMS circuit carries representations of key decision-variables and are modulated during choice epochs (Choi et al., 2023). Meanwhile, PL neurons targeting the DLS encode internal representations of value. In yet another study this time focusing on the projections from the ventrolateral OFC to the dorsomedial striatum (DMS) - these connected regions were found to be required for economic decision-making (Gore et al., 2023) with neurons in these areas strongly encoding the spatial features and choice direction. Interestingly, the neural encoding of the choice direction in the OFC precedes that in the DMS - thus predicting

choice accuracy. Therefore, projections from one region to another as well as studying interactions of regions implicated in the value-based decision-making process reveals a distributed code of economic variables as well as a sequential order following anatomical connectivity in how information is transformed.

Overall, several studies have suggested that projection-defined populations of frontal cortical cells are differentially engaged during economic decision-making behavior - suggesting that neural representations are structured according to downstream computations (Hirokawa et al., 2019). The main limitation is that these projection studies focus on connections between specific upstream and downstream brain regions, making it challenging to disentangle the relationship between projection- and region-specific signals. A study in the lab is looking at whether the regional representations depend on the projection target (Ashcroft et al., SFN 2023). Other factors such as cell type specificity may also shape frontal neuronal representations of economic variables. Previous studies have shown that in other decision-making or simpler tasks, the encoding of task-relevant information is cell-type specific (Kvitsiani et al., 2013). It would be interesting to test whether there are principles of organization across cell types in our economic task with respect to these economic variables. All of these future directions would enrich the interpretations of our neural dataset and advance our understanding of the structure and temporal dynamics of representations of core economic variables.

Summary:

In conclusion, we combine behavioral training, electrophysiology and optogenetic manipulations to understand the structure of neural representations of economic value in the regions of the frontal cortex and test their causal roles. Our results show that mice can make economic choices by learning the value associated with different abstract stimuli in our task. These economic choices are guided not only by the expected value of different stimuli but also by risk, and we show diverse individual risk attitudes. We observe distinct neural representations of both EV and risk across several frontal

regions. These representations ultimately form orthogonal functional gradients for each economic variable. Lastly, we try to understand the causal role of such representations with no conclusive evidence that inactivating frontal regions during our task disrupts economic choices. In conclusion, we develop an useful animal model of economic choice behavior which allows us to understand how the brain structures representations of economic variables. Future work will address not only the causal role of these representations, but also investigate how the structure of representations of decision variables emerges during learning. These findings have implications for understanding the neural processes underlying impairment of decision-making in patients.

References

1. Abler, B., Walter, H., Erk, S., Kammerer, H., & Spitzer, M. (2006). Prediction error as a linear function of reward probability is coded in human nucleus accumbens. Neuroimage, 31(2), 790-795.
2. Ainslie, G. (1992). Picoeconomics: The strategic interaction of successive motivational states within the person. Cambridge University Press.
3. Amiez, C., Joseph, J. P., & Procyk, E. (2006). Reward encoding in the monkey anterior cingulate cortex. Cerebral cortex, 16(7), 1040-1055.
4. Ashwood, Z. C., Roy, N. A., Stone, I. R., International Brain Laboratory, Urai, A. E., Churchland, A. K., ... & Pillow, J. W. (2022). Mice alternate between discrete strategies during perceptual decision-making. Nature Neuroscience, 25(2), 201-212.
5. Balewski, Z. Z., Elston, T. W., Knudsen, E. B., & Wallis, J. D. (2023). Value dynamics affect choice preparation during decision-making. Nature neuroscience, 26(9), 1575-1583.
6. Ballesta, S., Shi, W., Conen, K. E., & Padoa-Schioppa, C. (2020). Values encoded in orbitofrontal cortex are causally related to economic choices. Nature, 588(7838), 450-453.
7. Ballesta, S., Shi, W., & Padoa-Schioppa, C. (2022). Orbitofrontal cortex contributes to the comparison of values underlying economic choices. Nature communications, 13(1), 4405.
8. Banich, M. T. (1998). The missing link: the role of interhemispheric interaction in attentional processing. Brain and cognition, 36(2), 128-157.
9. Bateson, M. (2002). Recent advances in our understanding of risk-sensitive foraging preferences. Proceedings of the Nutrition Society, 61(4), 509-516.
10. Bao, C., Zhu, X., Mōller-Mara, J., Li, J., Dubroqua, S., & Erlich, J. C. (2023). The rat frontal orienting field dynamically encodes value for economic decisions under risk. Nature Neuroscience, 26(11), 1942-1952.
11. Bechara, A., Damasio, A. R., Damasio, H., & Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. Cognition, 50(1-3), 7-15.

12. Bechara, A., Damasio, H., Damasio, A. R., & Lee, G. P. (1999). Different contributions of the human amygdala and ventromedial prefrontal cortex to decision-making. Journal of neuroscience, 19(13), 5473-5481.
13. Bhagat, J., Wells, M. J., Harris, K. D., Carandini, M., & Burgess, C. P. (2020). Rigbox: an Open-Source toolbox for probing neurons and behavior. Eneuro, 7(4).
14. Braver, T. S., Cohen, J. D., Nystrom, L. E., Jonides, J., Smith, E. E., & Noll, D. C. (1997). A parametric study of prefrontal cortex involvement in human working memory. Neuroimage, 5(1), 49-62.
15. Brito e Abreu, F., & Kacelnik, A. (1999). Energy budgets and risk-sensitive foraging in starlings. Behavioral Ecology, 10(3), 338-345.
16. Buckley, M. J., Mansouri, F. A., Hoda, H., Mahboubi, M., Browning, P. G., Kwok, S. C., ... & Tanaka, K. (2009). Dissociable components of rule-guided behavior depend on distinct medial and prefrontal regions. Science, 325(5936), 52-58.
17. Burgess, C. P., Lak, A., Steinmetz, N. A., Zatka-Haas, P., Reddy, C. B., Jacobs, E. A., ... & Carandini, M. (2017). High-yield methods for accurate two-alternative visual psychophysics in head-fixed mice. Cell reports, 20(10), 2513-2524.
18. Bussey, T. J., Wise, S. P., & Murray, E. A. (2001). The role of ventral and orbital prefrontal cortex in conditional visuomotor learning and strategy use in rhesus monkeys (Macaca mulatta). Behavioral neuroscience, 115(5), 971.
19. Cai, X., & Padoa-Schioppa, C. (2012). Neuronal encoding of subjective value in dorsal and ventral anterior cingulate cortex. Journal of Neuroscience, 32(11), 3791-3808.
20. Cai, X., & Padoa-Schioppa, C. (2014). Contributions of orbitofrontal and lateral prefrontal cortices to economic choice and the good-to-action transformation. Neuron, 81(5), 1140-1151.
21. Camerer, C. F. (2003). Behavioral studies of strategic thinking in games. Trends in cognitive sciences, 7(5), 225-231.

22. Camille, N., Griffiths, C. A., Vo, K., Fellows, L. K., & Kable, J. W. (2011). Ventromedial frontal lobe damage disrupts value maximization in humans. Journal of Neuroscience, 31(20), 7527-7532.
23. Carmichael, S. T., & Price, J. L. (1995). Limbic connections of the orbital and medial prefrontal cortex in macaque monkeys. Journal of comparative neurology, 363(4), 615-641.
24. Carmichael, S. T., & Price, J. L. (1996). Connectional networks within the orbital and medial prefrontal cortex of macaque monkeys. Journal of Comparative Neurology, 371(2), 179-207.
25. Chen, X., & Stuphorn, V. (2015). Sequential selection of economic good and action in medial frontal cortex of macaques during value-based decisions. Elife, 4, e09418.
26. 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. Journal of Neuroscience, 29(39), 12315-12320.
27. Choi, K., Piasini, E., Díaz-Hernández, E., Cifuentes, L. V., Henderson, N. T., Holly, E. N., ... & Fuccillo, M. V. (2023). Distributed processing for value-based choice by prelimbic circuits targeting anterior-posterior dorsal striatal subregions in male mice. Nature Communications, 14(1), 1920.
28. 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.
29. Cisek, P., & Thura, D. (2018). Neural circuits for action selection. In Reach-to-Grasp Behavior (pp. 91-118). Routledge.
30. Cole, N., Harvey, M., Myers-Joseph, D., Gilra, A., & Khan, A. G. (2024). Prediction-error signals in anterior cingulate cortex drive task-switching. Nature Communications, 15(1), 7088.
31. Constantinople, C. M., Piet, A. T., & Brody, C. D. (2019a). An analysis of decision under risk in rats. Current Biology, 29(12), 2066-2074.

32. Constantinople, C. M., Piet, A. T., Bibawi, P., Akrami, A., Kopec, C., & Brody, C. D. (2019b). Lateral orbitofrontal cortex promotes trial-by-trial learning of risky, but not spatial, biases. Elife, 8, e49744.
33. Culbreth, A. J., Westbrook, A., Daw, N. D., Botvinick, M., & Barch, D. M. (2016). Reduced model-based decision-making in schizophrenia. Journal of abnormal psychology, 125(6), 777.
34. Damasio, H., Grabowski, T., Frank, R., Galaburda, A. M., & Damasio, A. R. (1994). The return of Phineas Gage: clues about the brain from the skull of a famous patient. Science, 264(5162), 1102-1105.
35. Dickinson, A., & Balleine, B. (1994). Motivational control of goal-directed action. Animal learning & behavior, 22(1), 1-18.
36. Feierstein, C. E., Quirk, M. C., Uchida, N., Sosulski, D. L., & Mainen, Z. F. (2006). Representation of spatial goals in rat orbitofrontal cortex. Neuron, 51(4), 495-507.
37. Fellows, L. K. (2011). Orbitofrontal contributions to value-based decision making: evidence from humans with frontal lobe damage. Annals of the new York Academy of Sciences, 1239(1), 51-58.
38. Fritsche, M., Majumdar, A., Strickland, L., Liebana Garcia, S., Bogacz, R., & Lak, A. (2024). Temporal regularities shape perceptual decisions and striatal dopamine signals. Nature Communications, 15(1), 7093.
39. Gaffan, D., & Harrison, S. (1989). A comparison of the effects of fornix transection and sulcus principalis ablation upon spatial learning by monkeys. Behavioural brain research, 31(3), 207-220.
40. Gardner, M. P., Conroy, J. S., Shaham, M. H., Styer, C. V., & Schoenbaum, G. (2017). Lateral orbitofrontal inactivation dissociates devaluation-sensitive behavior and economic choice. Neuron, 96(5), 1192-1203.
41. Gardner, M. P., Conroy, J. C., Styer, C. V., Huynh, T., Whitaker, L. R., & Schoenbaum, G. (2018). Medial orbitofrontal inactivation does not affect economic choice. Elife, 7, e38963.

42. Gardner, M. P., Conroy, J. C., Sanchez, D. C., Zhou, J., & Schoenbaum, G. (2019). Real-time value integration during economic choice is regulated by orbitofrontal cortex. Current Biology, 29(24), 4315-4322.
43. Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. Annu. Rev. Neurosci., 30(1), 535-574.
44. Gore, F., Hernandez, M., Ramakrishnan, C., Crow, A. K., Malenka, R. C., & Deisseroth, K. (2023). Orbitofrontal cortex control of striatum leads economic decision-making. Nature Neuroscience, 26(9), 1566-1574.
45. Guo, Z. V., Li, N., Huber, D., Ophir, E., Gutnisky, D., Ting, J. T., ... & Svoboda, K. (2014). Flow of cortical activity underlying a tactile decision in mice. Neuron, 81(1), 179-194.
46. Harbaugh, W. T., Krause, K., & Berry, T. R. (2001). GARP for kids: On the development of rational choice behavior. American Economic Review, 91(5), 1539-1545.
47. 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.
48. Harlow, J. M. (1974). Recovery after severe injury to the head. Van der Kloot, WG, Walcott, C., and Dane, B.. Readings in Behavior. New York: Holt, Rinehart and Winston, 291-309.
49. Harris, J. A., Mihalas, S., Hirokawa, K. E., Whitesell, J. D., Choi, H., Bernard, A., ... & Zeng, H. (2019). Hierarchical organization of cortical and thalamic connectivity. Nature, 575(7781), 195-202.
50. Haufe, S., Meinecke, F., Görgen, K., Dähne, S., Haynes, J. D., Blankertz, B., & Bießmann, F. (2014). On the interpretation of weight vectors of linear models in multivariate neuroimaging. Neuroimage, 87, 96-110.
51. Hayden, B. Y., & Platt, M. L. (2007). Temporal discounting predicts risk sensitivity in rhesus macaques. Current Biology, 17(1), 49-53.

52. Heilbronner, S. R., Rodriguez-Romaguera, J., Quirk, G. J., Groenewegen, H. J., & Haber, S. N. (2016). Circuit-based corticostriatal homologies between rat and primate. Biological psychiatry, 80(7), 509-521.
53. Hirokawa, J., Vaughan, A., Masset, P., Ott, T., & Kepecs, A. (2019). Frontal cortex neuron types categorically encode single decision variables. Nature, 576(7787), 446-451.
54. Hjort, M. M., Gowrishankar, R., Tian, L., Gordon-Fennell, A., Namboodiri, V. M., Bruchas, M. R., & Stuber, G. D. (2024). Microprisms enable enhanced throughput and resolution for longitudinal tracking of neuronal ensembles in deep brain structures. Neurophotonics, 11(3), 033407-033407.
55. Hodges, J. R., & Miller, B. (2001). The neuropsychology of frontal variant frontotemporal dementia and semantic dementia. Introduction to the special topic papers: Part II. Neurocase, 7(2), 113-121.
56. Hunt, L. T., & Hayden, B. Y. (2017). A distributed, hierarchical and recurrent framework for reward-based choice. Nature Reviews Neuroscience, 18(3), 172-182.
57. Hurly, T. A., & Oseen, M. D. (1999). Context-dependent, risk-sensitive foraging preferences in wild rufous hummingbirds. Animal Behaviour, 58(1), 59-66.
58. International Brain Laboratory, Aguillon-Rodriguez, V., Angelaki, D., Bayer, H., Bonacchi, N., Carandini, M., ... & Zador, A. M. (2021). Standardized and reproducible measurement of decision-making in mice. Elife, 10, e63711.
59. Ito, S., Stuphorn, V., Brown, J. W., & Schall, J. D. (2003). Performance monitoring by the anterior cingulate cortex during saccade countermanding. Science, 302(5642), 120-122.
60. Izquierdo, A., & Murray, E. A. (2004). Combined unilateral lesions of the amygdala and orbital prefrontal cortex impair affective processing in rhesus monkeys. Journal of neurophysiology, 91(5), 2023-2039.
61. Juavinett, A. L., Bekheet, G., & Churchland, A. K. (2019). Chronically implanted Neuropixels probes enable high-yield recordings in freely moving mice. Elife, 8, e47188.

62. Jun, J. J., Steinmetz, N. A., Siegle, J. H., Denman, D. J., Bauza, M., Barbarits, B., ... & Harris, T. D. (2017). Fully integrated silicon probes for high-density recording of neural activity. Nature, 551(7679), 232-236.
63. Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. Nature neuroscience, 10(12), 1625-1633.
64. Kacelnik, A., & Bateson, M. (1996). Risky theories—the effects of variance on foraging decisions. American zoologist, 36(4), 402-434.
65. Kagel, J. H., Battalio, R. C., & Green, L. (1995). Economic choice theory: An experimental analysis of animal behavior. Cambridge University Press.
66. Kahneman, D., & Tversky, A. (2013). Prospect theory: An analysis of decision under risk. In Handbook of the fundamentals of financial decision making: Part I (pp. 99-127).
67. Kamalova, A., Manoocheri, K., Liu, X., Casello, S. M., Huang, M., Baimel, C., ... & Carter, A. G. (2024). CCK+ Interneurons Contribute to Thalamus-Evoked Feed-Forward Inhibition in the Prelimbic Prefrontal Cortex. Journal of Neuroscience, 44(23).
68. Kennerley, S. W., Dahmubed, A. F., Lara, A. H., & Wallis, J. D. (2009a). Neurons in the frontal lobe encode the value of multiple decision variables. Journal of cognitive neuroscience, 21(6), 1162-1178.
69. Kennerley, S. W., & Wallis, J. D. (2009b). Evaluating choices by single neurons in the frontal lobe: outcome value encoded across multiple decision variables. European Journal of Neuroscience, 29(10), 2061-2073.
70. Kennerley, S. W., Behrens, T. E., & Wallis, J. D. (2011). Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. Nature neuroscience, 14(12), 1581-1589.
71. Kepecs, A., Uchida, N., Zariwala, H. A., & Mainen, Z. F. (2008). Neural correlates, computation and behavioural impact of decision confidence. Nature, 455(7210), 227-231.

72. Kim, Y., Yang, G. R., Pradhan, K., Venkataraju, K. U., Bota, M., Del Molino, L. C. G., ... & Osten, P. (2017). Brain-wide maps reveal stereotyped cell-type-based cortical architecture and subcortical sexual dimorphism. Cell, 171(2), 456-469.
73. King, J. R., & Dehaene, S. (2014). Characterizing the dynamics of mental representations: the temporal generalization method. Trends in cognitive sciences, 18(4), 203-210.
74. Kvitsiani, D., Ranade, S., Hangya, B., Taniguchi, H., Huang, J. Z., & Kepecs, A. (2013). Distinct behavioural and network correlates of two interneuron types in prefrontal cortex. Nature, 498(7454), 363-366.
75. Knutson, B., & Cooper, J. C. (2005). Functional magnetic resonance imaging of reward prediction. Current opinion in neurology, 18(4), 411-417.
76. Kuwabara, M., Kang, N., Holy, T. E., & Padoa-Schioppa, C. (2020). Neural mechanisms of economic choices in mice. Elife, 9, e49669.
77. Lak, A., Costa, G. M., Romberg, E., Koulakov, A. A., Mainen, Z. F., & Kepecs, A. (2014). Orbitofrontal cortex is required for optimal waiting based on decision confidence. Neuron, 84(1), 190-201.
78. Lak, A., Okun, M., Moss, M. M., Gurnani, H., Farrell, K., Wells, M. J., ... & Carandini, M. (2020). Dopaminergic and prefrontal basis of learning from sensory confidence and reward value. Neuron, 105(4), 700-711.
79. Lau, B., & Glimcher, P. W. (2007). Action and outcome encoding in the primate caudate nucleus. Journal of Neuroscience, 27(52), 14502-14514.
80. Lee, D., Rushworth, M. F., Walton, M. E., Watanabe, M., & Sakagami, M. (2007). Functional specialization of the primate frontal cortex during decision making. Journal of Neuroscience, 27(31), 8170-8173.
81. Lee, D., & Seo, H. (2016). Neural basis of strategic decision making. Trends in neurosciences, 39(1), 40-48.
82. Le Merre, P., Ährlund-Richter, S., & Carlén, M. (2021). The mouse prefrontal cortex: Unity in diversity. Neuron, 109(12), 1925-1944.

83. Levy, I., Snell, J., Nelson, A. J., Rustichini, A., & Glimcher, P. W. (2010). Neural representation of subjective value under risk and ambiguity. Journal of neurophysiology, 103(2), 1036-1047.
84. Lu, M. T., Preston, J. B., & Strick, P. L. (1994). Interconnections between the prefrontal cortex and the premotor areas in the frontal lobe. Journal of Comparative Neurology, 341(3), 375-392.
85. Luce, R. D. (1995). Four tensions concerning mathematical modeling in psychology. Annual Review of Psychology, 46, 1.
86. Luce, R. D. (1997). Several unresolved conceptual problems of mathematical psychology. Journal of mathematical psychology, 41(1), 79-87.
87. Luk, C. H., & Wallis, J. D. (2013). Choice coding in frontal cortex during stimulus-guided or action-guided decision-making. Journal of Neuroscience, 33(5), 1864-1871.
88. Macmillan, M. (2000). Restoring phineas gage: A 150th retrospective. Journal of the History of the Neurosciences, 9(1), 46-66.
89. Maisson, D. J. N., Cash-Padgett, T. V., Wang, M. Z., Hayden, B. Y., Heilbronner, S. R., & Zimmermann, J. (2021). Choice-relevant information transformation along a ventrodorsal axis in the medial prefrontal cortex. Nature communications, 12(1), 4830.
90. Maris, E., Schoffelen, J. M., & Fries, P. (2007). Nonparametric statistical testing of coherence differences. Journal of neuroscience methods, 163(1), 161-175.
91. Matsumoto, K., Suzuki, W., & Tanaka, K. (2003). Neuronal correlates of goal-based motor selection in the prefrontal cortex. Science, 301(5630), 229-232.
92. McClure, S. M., Li, J., Tomlin, D., Cypert, K. S., Montague, L. M., & Montague, P. R. (2004). Neural correlates of behavioral preference for culturally familiar drinks. Neuron, 44(2), 379-387.
93. McCoy, A. N., & Platt, M. L. (2005). Risk-sensitive neurons in macaque posterior cingulate cortex. Nature neuroscience, 8(9), 1220-1227.

94. Miller, K. J., Botvinick, M. M., & Brody, C. D. (2022). Value representations in the rodent orbitofrontal cortex drive learning, not choice. Elife, 11, e64575.

95. Milner, B. (1963). Effects of different brain lesions on card sorting: The role of the frontal lobes. Archives of neurology, 9(1), 90-100.

96. Montague, P. R., & Berns, G. S. (2002). Neural economics and the biological substrates of valuation. Neuron, 36(2), 265-284.

97. Murray, E. A., Bussey, T. J., & Wise, S. P. (2000). Role of prefrontal cortex in a network for arbitrary visuomotor mapping. Executive Control and the Frontal Lobe: Current Issues, 114-129.

98. Neylan, T. C. (1999). Frontal lobe function: Mr. Phineas Gage's famous injury. The Journal of neuropsychiatry and clinical neurosciences, 11(2), 280-281.

99. Niki, H., & Watanabe, M. (1979). Prefrontal and cingulate unit activity during timing behavior in the monkey. Brain research, 171(2), 213-224.

100. O'Doherty, J. P. (2014). The problem with value. Neuroscience & Biobehavioral Reviews, 43, 259-268.

101. O'Neill, M., & Schultz, W. (2010). Coding of reward risk by orbitofrontal neurons is mostly distinct from coding of reward value. Neuron, 68(4), 789-800.

102. Orsini, C. A., Moorman, D. E., Young, J. W., Setlow, B., & Floresco, S. B. (2015). Neural mechanisms regulating different forms of risk-related decision-making: Insights from animal models. Neuroscience & Biobehavioral Reviews, 58, 147-167.

103. Padoa-Schioppa, C., Jandolo, L., & Visalberghi, E. (2006). Multi-stage mental process for economic choice in capuchins. Cognition, 99(1), B1-B13.

104. Padoa-Schioppa, C., & Assad, J. A. (2006). Neurons in the orbitofrontal cortex encode economic value. Nature, 441(7090), 223-226.

105. Padoa-Schioppa, C. (2009). Range-adapting representation of economic value in the orbitofrontal cortex. Journal of Neuroscience, 29(44), 14004-14014.

106. Padoa-Schioppa, C., & Assad, J. A. (2008). The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. Nature neuroscience, 11(1), 95-102.

107. Padoa-Schioppa, C. (2011). Neurobiology of economic choice: a good-based model. Annual review of neuroscience, 34(1), 333-359.

108. Padoa-Schioppa, C., & Schoenbaum, G. (2015). Dialogue on economic choice, learning theory, and neuronal representations. Current opinion in behavioral sciences, 5, 16-23.

109. Passecker, J., Mikus, N., Malagon-Vina, H., Anner, P., Dimidschstein, J., Fishell, G., ... & Klausberger, T. (2019). Activity of prefrontal neurons predict future choices during gambling. Neuron, 101(1), 152-164.

110. Peters, J., & Büchel, C. (2009). Overlapping and distinct neural systems code for subjective value during intertemporal and risky decision making. Journal of Neuroscience, 29(50), 15727-15734.

111. Petrides, M., & Milner, B. (1982). Deficits on subject-ordered tasks after frontal-and temporal-lobe lesions in man. Neuropsychologia, 20(3), 249-262.

112. Petrides, M., & Pandya, D. N. (2006). Efferent association pathways originating in the caudal prefrontal cortex in the macaque monkey. Journal of Comparative Neurology, 498(2), 227-251.

113. Pinto, L., Rajan, K., DePasquale, B., Thiberge, S. Y., Tank, D. W., & Brody, C. D. (2019). Task-dependent changes in the large-scale dynamics and necessity of cortical regions. Neuron, 104(4), 810-824.

114. 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.

115. Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. Nature, 400(6741), 233-238.

116. Polanía, R., Woodford, M., & Ruff, C. C. (2019). Efficient coding of subjective value. Nature neuroscience, 22(1), 134-142.

117. Preuss, T. M. (1995). Do rats have prefrontal cortex? The Rose-Woolsey-Akert program reconsidered. Journal of cognitive neuroscience, 7(1), 1-24.

118. Preuschoff, K., Bossaerts, P., & Quartz, S. R. (2006). Neural differentiation of expected reward and risk in human subcortical structures. Neuron, 51(3), 381-390.

119. Rahman, S., Sahakian, B. J., Hodges, J. R., Rogers, R. D., & Robbins, T. W. (1999). Specific cognitive deficits in mild frontal variant frontotemporal dementia. Brain, 122(8), 1469-1493.

120. Raimondo, J. V., Kay, L., Ellender, T. J., & Akerman, C. J. (2012). Optogenetic silencing strategies differ in their effects on inhibitory synaptic transmission. Nature neuroscience, 15(8), 1102-1104.

121. Rangel, A., & Hare, T. (2010). Neural computations associated with goal-directed choice. Current opinion in neurobiology, 20(2), 262-270.

122. Regenwetter, M., Dana, J., & Davis-Stober, C. P. (2010). Testing transitivity of preferences on two-alternative forced choice data. Frontiers in Psychology., 1, 148.

123. Reinert, S., Hübener, M., Bonhoeffer, T., & Goltstein, P. M. (2021). Mouse prefrontal cortex represents learned rules for categorization. Nature, 593(7859), 411-417.

124. Rich, E. L., & Wallis, J. D. (2016). Decoding subjective decisions from orbitofrontal cortex. Nature neuroscience, 19(7), 973-980.

125. Roesch, M. R., & Olson, C. R. (2004). Neuronal activity related to reward value and motivation in primate frontal cortex. Science, 304(5668), 307-310.

126. Roesch, M. R., & Olson, C. R. (2005). Neuronal activity in primate orbitofrontal cortex reflects the value of time. Journal of neurophysiology, 94(4), 2457-2471.

127. Rolls, E. (2008). Functions of the orbitofrontal and pregenual cingulate cortex in taste, olfaction, appetite and emotion. Acta Physiologica Hungarica, 95(2), 131-164.

128. Rose, J. E., & Woolsey, C. N. (1949). Organization of the mammalian thalamus and its relationships to the cerebral cortex. Electroencephalography and clinical neurophysiology, 1(1-4), 391-404.

129. Rothschild, M., & Stiglitz, J. E. (1978). Increasing risk: I. A definition. In Uncertainty in economics (pp. 99-121). Academic Press.

130. Rudebeck, P. H., & Murray, E. A. (2011). Dissociable effects of subtotal lesions within the macaque orbital prefrontal cortex on reward-guided behavior. Journal of Neuroscience, 31(29), 10569-10578.

131. Rudebeck, P. H., Walton, M. E., Smyth, A. N., Bannerman, D. M., & Rushworth, M. F. (2006). Separate neural pathways process different decision costs. Nature neuroscience, 9(9), 1161-1168.

132. Rushworth, M. F., Behrens, T. E. J., Rudebeck, P. H., & Walton, M. E. (2007). Contrasting roles for cingulate and orbitofrontal cortex in decisions and social behaviour. Trends in cognitive sciences, 11(4), 168-176.

133. Rushworth, M. F., & Behrens, T. E. (2008). Choice, uncertainty and value in prefrontal and cingulate cortex. Nature neuroscience, 11(4), 389-397.

134. Saleem, K. S., Miller, B., & Price, J. L. (2014). Subdivisions and connectional networks of the lateral prefrontal cortex in the macaque monkey. Journal of Comparative Neurology, 522(7), 1641-1690.

135. Samejima, K., Ueda, Y., Doya, K., & Kimura, M. (2005). Representation of action-specific reward values in the striatum. Science, 310(5752), 1337-1340.

136. Samuelson, P. A. (1938). A note on the pure theory of consumer's behavior. Economica, 5(17), 61-71.

137. Seo, H., & Lee, D. (2007). Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. Journal of neuroscience, 27(31), 8366-8377.

138. Schoenbaum, G., Chiba, A. A., & Gallagher, M. (1998). Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. Nature neuroscience, 1(2), 155-159.

139. Schoenbaum, G., Chiba, A. A., & Gallagher, M. (1999). Neural encoding in orbitofrontal cortex and basolateral amygdala during olfactory discrimination learning. Journal of Neuroscience, 19(5), 1876-1884.

140. Schoenbaum, G., Setlow, B., Saddoris, M. P., & Gallagher, M. (2003). Encoding predicted outcome and acquired value in orbitofrontal cortex during cue sampling depends upon input from basolateral amygdala. Neuron, 39(5), 855-867.

141. Shohamy, D., & Daw, N. D. (2015). Integrating memories to guide decisions. Current Opinion in Behavioral Sciences, 5, 85-90.

142. Skinner, B. F. (1957). The experimental analysis of behavior. American scientist, 45(4), 343-371.

143. Song, Z., Alpers, A., Warner, K., Iacobucci, F., Hoskins, E., Disterhoft, J. F., ... & Widge, A. S. (2024). Chronic, Reusable, Multiday Neuropixels Recordings during Free-Moving Operant Behavior. Eneuro, 11(1).

144. Stalnaker, T. A., Cooch, N. K., & Schoenbaum, G. (2015). What the orbitofrontal cortex does not do. Nature neuroscience, 18(5), 620-627.

145. Stauffer, W. R., Lak, A., & Schultz, W. (2014). Dopamine reward prediction error responses reflect marginal utility. Current biology, 24(21), 2491-2500.

146. Stauffer, W. R., Lak, A., Bossaerts, P., & Schultz, W. (2015). Economic choices reveal probability distortion in macaque monkeys. Journal of Neuroscience, 35(7), 3146-3154.

147. Steinmetz, N. A., Aydin, C., Lebedeva, A., Okun, M., Pachitariu, M., Bauza, M., ... & Harris, T. D. (2021). Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. Science, 372(6539), eabf4588.

148. Steinmetz, N. A., Zatka-Haas, P., Carandini, M., & Harris, K. D. (2018). Distributed correlates of visually-guided behavior across the mouse brain. BioRxiv, 474437.

149. Sul, J. H., Kim, H., Huh, N., Lee, D., & Jung, M. W. (2010). Distinct roles of rodent orbitofrontal and medial prefrontal cortex in decision making. Neuron, 66(3), 449-460.

150. Syeda, A., Zhong, L., Tung, R., Long, W., Pachitariu, M., & Stringer, C. (2024). Facemap: a framework for modeling neural activity based on orofacial tracking. Nature Neuroscience, 27(1), 187-195.

151. Takada, M., Nambu, A., Hatanaka, N., Tachibana, Y., Miyachi, S., Taira, M., & Inase, M. (2004). Organization of prefrontal outflow toward frontal motor-related areas in macaque monkeys. European Journal of Neuroscience, 19(12), 3328-3342.

152. Takahara, D., Inoue, K. I., Hirata, Y., Miyachi, S., Nambu, A., Takada, M., & Hoshi, E. (2012). Multisynaptic projections from the ventrolateral prefrontal cortex to the dorsal premotor cortex in macaques–anatomical substrate for conditional visuomotor behavior. European Journal of Neuroscience, 36(10), 3365-3375.

153. Tang, H., Costa, V. D., Bartolo, R., & Averbeck, B. B. (2022). Differential coding of goals and actions in ventral and dorsal corticostriatal circuits during goal-directed behavior. Cell reports, 38(1).

154. Tobler, P. N., O'Doherty, J. P., Dolan, R. J., & Schultz, W. (2006). Human neural learning depends on reward prediction errors in the blocking paradigm. Journal of Neurophysiology, 95(1), 301-310.

155. Tobler, P. N., Fletcher, P. C., Bullmore, E. T., & Schultz, W. (2007). Learning-related human brain activations reflecting individual finances. Neuron, 54(1), 167-175.

156. Tobler, P. N., Christopoulos, G. I., O'Doherty, J. P., Dolan, R. J., & Schultz, W. (2008). Neuronal distortions of reward probability without choice. Journal of Neuroscience, 28(45), 11703-11711.

157. Tobler, P. N., Christopoulos, G. I., O'Doherty, J. P., Dolan, R. J., & Schultz, W. (2009). Risk-dependent reward value signal in human prefrontal cortex. Proceedings of the National Academy of Sciences, 106(17), 7185-7190.

158. 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.

159. Tremblay, L., & Schultz, W. (1999). Relative reward preference in primate orbitofrontal cortex. Nature, 398(6729), 704-708.

160. Tversky, A., & Kahneman, D. (Eds.). (2000). Choices, values, and frames (pp. 375-406). New York: Cambridge University Press.

161. Wallis, J. D., Anderson, K. C., & Miller, E. K. (2001). Single neurons in prefrontal cortex encode abstract rules. Nature, 411(6840), 953-956.

162. Wallis, J. D., & Kennerley, S. W. (2010). Heterogeneous reward signals in prefrontal cortex. Current opinion in neurobiology, 20(2), 191-198.

163. Wallis, J. D., & Miller, E. K. (2003). Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. European Journal of Neuroscience, 18(7), 2069-2081.

164. Walton, M. E., Bannerman, D. M., & Rushworth, M. F. (2002). The role of rat medial frontal cortex in effort-based decision making. Journal of Neuroscience, 22(24), 10996-11003.

165. Walton, M. E., Behrens, T. E., Buckley, M. J., Rudebeck, P. H., & Rushworth, M. F. (2010). Separable learning systems in the macaque brain and the role of orbitofrontal cortex in contingent learning. Neuron, 65(6), 927-939.

166. Waltz, J. A., & Gold, J. M. (2016). Motivational deficits in schizophrenia and the representation of expected value. Behavioral neuroscience of motivation, 375-410.

167. Weber, E. U., Shafir, S., & Blais, A. R. (2004). Predicting risk sensitivity in humans and lower animals: risk as variance or coefficient of variation. Psychological review, 111(2), 430.

168. Williams, Z. M., Bush, G., Rauch, S. L., Cosgrove, G. R., & Eskandar, E. N. (2004). Human anterior cingulate neurons and the integration of monetary reward with motor responses. Nature neuroscience, 7(12), 1370-1375.

169. Witten, I. B., Steinberg, E. E., Lee, S. Y., Davidson, T. J., Zalocusky, K. A., Brodsky, M., ... & Deisseroth, K. (2011). Recombinase-driver rat lines: tools, techniques, and optogenetic application to dopamine-mediated reinforcement. Neuron, 72(5), 721-733.

170. Wunderlich, K., Rangel, A., & O'Doherty, J. P. (2010). Economic choices can be made using only stimulus values. Proceedings of the National Academy of Sciences, 107(34), 15005-15010.

171. Van De Werd, H. J. J. M., Rajkowska, G., Evers, P., & Uylings, H. B. (2010). Cytoarchitectonic and chemoarchitectonic characterization of the prefrontal cortical areas in the mouse. Brain Structure and Function, 214, 339-353.

172. van der Knaap, L. J., & van der Ham, I. J. (2011). How does the corpus callosum mediate interhemispheric transfer? A review. Behavioural brain research, 223(1), 211-221.

173. van Holstein, M., & Floresco, S. B. (2020). Dissociable roles for the ventral and dorsal medial prefrontal cortex in cue-guided risk/reward decision making. Neuropsychopharmacology, 45(4), 683-693.

174. Verharen, J. P., den Ouden, H. E., Adan, R. A., & Vanderschuren, L. J. (2020). Modulation of value-based decision making behavior by subregions of the rat prefrontal cortex. Psychopharmacology, 237, 1267-1280.

175. Von Neumann, J., & Morgenstern, O. (1947). Theory of games and economic behavior, 2nd rev.

176. Xie, J., & Padoa-Schioppa, C. (2016). Neuronal remapping and circuit persistence in economic decisions. Nature neuroscience, 19(6), 855-861.

177. Yamada, H., Tymula, A., Louie, K., & Glimcher, P. W. (2013). Thirst-dependent risk preferences in monkeys identify a primitive form of wealth. Proceedings of the National Academy of Sciences, 110(39), 15788-15793.

178. Yang, W., Tipparaju, S. L., Chen, G., & Li, N. (2022). Thalamus-driven functional populations in frontal cortex support decision-making. Nature neuroscience, 25(10), 1339-1352.

179. Yu, L. Q., Dana, J., & Kable, J. W. (2018). Individuals with ventromedial frontal damage have more unstable but still fundamentally transitive preferences. bioRxiv, 384024.

180. Yacubian, J., Sommer, T., Schroeder, K., Gläscher, J., Braus, D. F., & Büchel, C. (2007). Subregions of the ventral striatum show preferential coding of reward magnitude and probability. Neuroimage, 38(3), 557-563.

181. Zalocusky, K. A., Ramakrishnan, C., Lerner, T. N., Davidson, T. J., Knutson, B., & Deisseroth, K. (2016). Nucleus accumbens D2R cells signal prior outcomes and control risky decision-making. Nature, 531(7596), 642-646.