Neural and behavioral probability weighting function

Agnieszka Tymula[1], **Hiroshi Yamada**[2]

Recent theoretical models challenge the existence of a probability weighting function as it was traditionally conceived in Prospect Theory in ways that are not straightforward to test using choice data. This study transcends these constraints by directly observing probability distortions in the brain, free from utility confounds. Utilizing a unique dataset comprising 64,175 decision trials and 78,067 neural measurement trials, we pinpoint neural activity (a basic biological decision processing unit) that exclusively encodes probability, independent of payoff magnitudes. Our results demonstrate that the observed neural probability weighting functions are S-shaped, which is optimal in our experimental environment, but diverge from those estimated from behavior under conventional assumptions. Incorporating biologically realistic utility functions in the estimation enhances our ability to reconstruct neural probability weighting from observed choices, offering direct biological evidence on the bases of economic decision-making.

Keywords: neuroeconomics, monkey, decision making under risk, probability

JEL codes: D91, D87, D81, D01

[1] corresponding author: School of Economics, University of Sydney, Sydney, 2006 NSW, Australia; agnieszka.tymula@sydney.edu.au
[2] Division of Biomedical Science, Institute of Medicine, University of Tsukuba
1-1-1 Tenno-dai, Tsukuba, Ibaraki, 305-8577 Japan

Acknowledgements:

Monkey FU was provided by NBRP "Japanese Monkeys" through the National Bio Resource Project of MEXT, Japan. This study was supported by JSPS KAKENHI Grant Numbers JP:15H05374 and 21H02797; the Takeda Science Foundation; Council for Addiction Behavior Studies; Narishige Neuroscience Research Foundation; Moonshot R&D JPMJMS2294 (H.Y.). We thank Tiger Mathieson for his research assistance.

1. Introduction

The probability weighting function, alongside the reference-dependent utility and discount function, is a cornerstone of the behavioral economics toolkit. In their seminal paper on Prospect Theory, Kahneman & Tversky (1979) proposed that in decisions involving small stakes and risk, people weigh probabilities differently than their objective value. They assign "decision weights" to probabilities using an inverse S-shaped probability weighting function. This explanation rapidly got adopted into both theoretical and empirical social science research. Probability weighting has attracted significant and sustained attention – a vast amount of social science research utilizes the concepts of probability weighting to inform theoretical and empirical frameworks, influencing behavior-altering policies in finance, environmental conservation, health, social sectors, and beyond. This key behavioral economics concept is taught to all students in economics, psychology, finance, and marketing.

Despite significant progress, accurately measuring the probability weighting function remains a challenge. A major obstacle is that the utility and probability weighting functions are not directly observable in standard economic data. Instead, they are typically estimated from people's choices over lotteries with varying payoffs and probabilities, under the assumption that utility and probability weighting are separable and context-independent. However, this assumption has come under increasing scrutiny. A recent meta-analysis of Prospect Theory parameters by Imai et al. (2025) found that the estimates of utility and probability weighting parameters (both elevation and likelihood sensitivity) are correlated[1] and violate procedure invariance. Similarly, Fehr-Duda et al. (2010) showed that probability weighting varies with the size of the payoff. Additionally, recent theoretical work suggests that behaviors originally attributed to probability weighting may instead be explained by context-dependent utility (Bordalo et al., 2012; Glimcher & Tymula, 2023; Kontek & Lewandowski, 2018; Schneider & Day, 2018). Supporting this view, Bruhin et al. (2022) empirically demonstrated that the frequency of Allais paradox violations can be equally well explained by salience effects as by probability weighting. Together, these findings highlight the difficulty of disentangling utility and probability weighting econometrically. They suggest that the standard Prospect Theory framework may be missspecified and raise the possibility that a distinct probability weighting function may even not exist independently of the utility of outcomes.

---

[1] Probability weighting parameters, elevation and likelihood sensitivity, are not correlated in their study.

In this paper, we integrate methods from neuroscience and economics to open the black box of probability weighting by directly measuring neural activity at the level of single neurons as participants choose between lotteries. We pursue three main goals. First, we verify whether a neural signature of probability weighting—independent of utility—exists. If no such signal were present, it would suggest that Prospect Theory is purely descriptive and misaligned with the underlying mechanism that guides decision-making. However, we find a clear neural evidence for probability coding that is separate from utility. This allows us to measure probability weighting directly, without utility confounds. Notably, the S-shaped neural probability weighting we observe closely approximates a function that would maximize earnings in our experimental environment. Second, we assess whether the neural probability weighting aligns with the probability weighting estimated from behavior using standard structural model that assume a power (or CRRA) utility function. A close match would imply that concerns about misestimation due to utility confound may be overstated. Contrary to this, we find a substantial mismatch between the neural and behaviorally estimated probability weighting. This leads us to our third aim: identifying the source of this mismatch. We find evidence suggesting that not accounting for context-dependent utility may be the reason. If, as previous work suggests, behaviors traditionally attributed to probability weighting can be explained—at least in part— by context-dependent utility, then using a structural model with a fixed, context-independent utility function is likely to introduce a bias when estimating probability weighting. Indeed, we show that once we account for salience (Bordalo et al., 2012) or value normalization (Glimcher & Tymula, 2023) in the utility function, the behaviorally estimated probability weighting functions more closely resemble those observed in neural activity.

Several features of our experimental design are particularly noteworthy. We measure probability weights directly in the brain without confounding it with the utility function, and then compare it to the probability weighting function estimated from behavior based on a large number of decisions. Our unique dataset consists of 64,175 decisions made by two monkeys and 78,067 single-lottery trials in which monkeys did not make any decisions and their brain activity was recorded while they anticipated the receipt of a lottery. Our experimental design overcomes the shortcomings of previous studies in several ways. First, in our study, brain activity is recorded using electrodes implanted directly into our subjects' brains. This means that even though the brain activity is inherently noisy, on each trial we observe neural response to our task stimuli precisely with zero measurement noise—something not achievable with fMRI or EEG.[2] Second, the lotteries are randomly drawn from a fully crossed set of 100 combinations of ten payoffs and ten probabilities, eliminating any structural correlation and ruling out alternative explanations of probability distortion, such as distributional effects in Frydman & Jin (2023), Stewart et al. (2006), and Herold & Netzer (2023). Third, by recording brain activity only during single-lottery presentations, we avoid context effects related to choice sets. Fourth, our design allows us to identify neurons that respond selectively to probability but not payoff magnitude, enabling a clean measurement of the probability weighting function. Finally, by combining neural and behavioral data, we are able to validate the predictive power of our neural measurement out-of-sample, establishing its behavioral relevance.

While the use of non-human subjects is not standard in economics, many studies provided evidence that non-human primates are a good model for human behavior (Kagel et al., 1995; Platt & Glimcher, 1999; Tymula et al., 2023; Yamada et al., 2013). Conducting studies with monkeys offers several unique advantages. Firstly, many decisions can be obtained for reliable and precise estimates of probability weighting from behavior, which is not feasible with human participants in a laboratory setting. Consequently, human studies tend to produce noisier estimates than desired. Monkeys in our study perform the task over many months, yielding a dataset of 64,175 decisions allowing us to estimate structural models with high precision. Secondly, experimental economists often question whether the relatively modest laboratory incentives motivate participants enough to elicit their true preferences as well as about the external factors unobservable to the experimenter. In our study, we have full control over the monkeys' economy and manage their hydration status (i.e., physical wealth in water content) throughout the study. The subjects are rewarded with beverage lotteries, and the liquid they earn while performing the task is their only source of hydration. They are also compensated for every decision they make. Overall, this places them in an environment with more consequential decisions spread over a much longer period compared to a traditional economic experiment. Finally, the direct neural recordings that provide a much better signal-to-noise ratio are not available for human studies except in very specific patient populations performed by a neurosurgeon as a part of medical treatment (Kandel et al., 2021). Although previous studies have estimated probability weighting functions from the behavior of non-human primates (Farashahi et al., 2018; Ferrari-Toniolo et al., 2019; Fujimoto & Minamimoto, 2019; Imaizumi et al., 2022; Stauffer et al., 2015; Tymula et al., 2023), no study to date has attempted to isolate it from neural activity. In our previous study (Imaizumi

---

[2] Electrophysiology is superior to fMRI in the quality of neural recordings due to higher temporal and spatial resolution. fMRI is also susceptible to artifacts related to motion, physiological processes, and scanner noise.

et al., 2022), we have shown that by aggregating neural activity throughout the reward circuitry, we could reconstruct from this neural activity the risk preferences estimated from monkeys' behavior, however, in that study, we did not attempt to measure probability distortions separately from utility and instead we estimated them jointly, as has typically been done in the economics literature, meaning that both the utility and probability weighting functions could be biased by both the functional assumptions and joint estimation procedure. This study makes a unique contribution that leverages the methodological benefits of non-human primate studies, to precisely isolate probability weighting function in neural activity and avoid such confounds.

By analyzing single neuron activity, which is the biological basis of the brain function, we show that the brain can encode probability and utility through distinct neural channels. Some neurons in our dataset encode probability only (and not payoff), some encode payoff only (and not probability), and some encode both probability and payoff magnitudes. Importantly, these different types of neurons are intermingled across brain regions and co-exist in relative proximity to one another. This indicates that using fMRI, another method frequently used to record human brain activity that lacks the resolution of single-neuron recordings, it would be nearly impossible to isolate the probability weighting function from the utility function and to gain a new understanding of how our neural valuation system processes probabilities.

Crucially, we use our neural data to illustrate the shape of the probability weighting function. For both subjects, the neural probability weighting is S-shaped. Although this is in contrast with Prospect Theory, we demonstrate that this shape aligns with a function that would maximize expected earnings in for a chooser with a noisy and bounded representation of probabilities. Importantly, this S-shape contrasts with the probability weighting function estimated from the same subjects' behavior. We investigate several possible explanations for this discrepancy. First, we confirm that the neural signal robustly predicts choice, ruling out the possibility that the neural measure is behaviorally irrelevant. Next, we examine whether the mismatch could arise from choice complexity (Enke & Graeber, 2023) or the restrictive assumptions on the utility function in the behavioral estimation. A growing theoretical literature (Bordalo et al., 2012; Glimcher & Tymula, 2023; Kontek & Lewandowski, 2018; Schneider & Day, 2018) shows context-dependent utility can mimic probability weighting-like behavior, even when individuals perceive probabilities objectively. We find that when using a more realistic context-dependent utility function that adapts to salience or is normalized, the gap between neural and behavioral estimates narrows substantially.

Together, our results underscore the importance of moving beyond standard structural models when estimating probability weighting. Our findings suggest that neural data can reveal distortions in probability perception that are otherwise confounded or misrepresented in behavioral estimates. Moreover, the observed S-shape in neural probability weighting aligns with a reward-maximizing behavior in our experimental setting, offering a normative account of its origin (Vieider, 2025). By combining high-resolution neural recordings with behavioral estimation, we advance both the measurement and interpretation of probability weighting, one of the foundational constructs of behavioral economics.

In the following sections, we describe the experimental design, present the results, and then conclude.

2. Experimental design

The subjects in our study were two rhesus monkeys, named SUN and FU. They participated in two types of tasks, both involving lotteries but differentiated by the presence or absence of choice. In the "single cue no choice task," a lottery was presented without an alternative, whereas in the "choice task," the subjects chose between two lotteries presented on the screen. Since the subjects had no use for money, their rewards were given in the form of juice/water, delivered after each trial. We designed one hundred lotteries by crossing ten payoff magnitudes (ranging from 0.1 to 1.0 mL in 0.1 mL increments) with ten probability levels (from 0.1 to 1.0 in 0.1 increments). These amounts were chosen so that they are a detectable but not too large to ensure that the participants are incentivized to keep choosing for a sufficient number of trials. One mL is approximately 0.33-0.5% of our participant's daily water consumption during the experiment. The lotteries' payoff and probability magnitudes were visually communicated using pie charts: the payoff magnitude was indicated by the number of green segments, and the probability level by the number of blue segments. Figure 1 illustrates an example trial, where the subject is choosing between a lottery on the left that pays 0.5 mL of water or nothing each with a 50% chance and a lottery on the right that pays 0.2 mL of water with a 90% chance and nothing with 10% chance. After ten months of intensive training, the animals demonstrated a robust understanding of the tasks and stimuli and the probability with which they selected a lottery was increasing in its expected value and decreasing in the expected value of the other lottery (Imaizumi et al., 2022; Tymula et al., 2023). The data analyzed in this study was collected only after the completion of the training phase, ensuring the subjects' proficient comprehension of the stimuli. The monkeys performed these tasks five days a week to earn liquid rewards.

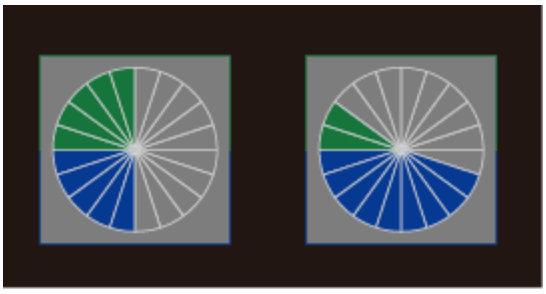

Figure 1. Example trial in choice task. The number of green pie segments represents payoff size in (each green pie = 0.1 mL) and the number of blue segments represents the probability (each blue pie = 10%).

Single cue no-choice task. The single cue no choice task was used for examining neural probability weighting function. The advantage of this task is that there is no other option in the choice set to create choice set effects in the utility function that could create confounds for the estimation of the probability weighting function. Even though the participants do not make a choice from a binary choice set in this task, they still need to stay alert and look at the lottery to receive it. To initiate each trial, monkeys had two seconds to align their gaze to the gray central fixation target. After fixation for one second, one pie chart providing information on the probability and magnitude of one lottery was presented for 2.5 seconds at the same location as the central fixation target. During these 2.5 seconds, brain activity was recorded. We calculate each neuron's activity (firing rate, FR) as the sum of the occurrences of action potentials[3] that it produced during the 2.5-second presentation of the lottery before the payoff is received (i.e., the frequency of the single neuron activity). The pie chart was then removed, and subjects received a liquid payoff, as indicated by the number of green pie chart segments, with the probability indicated by the number of blue pie chart segments. After 4–6 seconds, the next trial began. On each trial, one lottery was randomly selected, with replacement, from the set of one hundred possible lotteries. The timing of the sequence of the screens is similar to other studies of this type and achieves two goals. First, it gives participants sufficient time to respond (one second is enough for monkey to respond). Second, it provides the necessary 300-400 ms delay to observe neural activity on each trial. Overall, SUN and FU completed 38,678 and 39,389 single cue no-choice task trials respectively during which we measured the brain activity in 686 neurons, located across four brain regions associated with the computation of value: dorsal (DS) and ventral (VS) striatum, central part of orbitofrontal cortex (cOFC), and medial orbitofrontal cortex (mOFC). See Table 1 for the number

---

[3] An action potential is a brief electrical impulse that serves as the primary means of communication throughout the nervous system. It's a fundamental element of neural activity, allowing neurons (nerve cells) to transmit signals over long distances within the body.

of neurons whose activity was recorded in each region in each subject. Single cue no choice task trials were presented in blocks of 100-120 trials.

Table 1. The number of neurons which activity was recorded by brain region and subject.

	Monkey		
brain region	SUN	FU	**Total**
DS	98	96	**194**
VS	89	55	**144**
cOFC	98	92	**190**
mOFC	64	94	**158**
Total	**349**	**337**	**686**

Choice task. As in the single cue task, at the beginning of each trial, subjects had two seconds to align their gaze to the gray central fixation target. After fixation for one second, two lotteries, randomly selected from the set of one hundred, were presented to monkeys as pie charts for 2.5 seconds. After 2.5 seconds, two smaller gray circles appeared at the same location as lotteries together with a fixation target in the middle. After a 0.5-second delay, the fixation target disappeared. Subjects were then allowed 2 seconds to make their choice by shifting their gaze to the lottery that they preferred. Then, they received a reward according to the payoff and probability magnitudes they selected. After 4–6 seconds another trial began. Subjects completed approximately 30 to 60 trials of the choice task before switching to the single cue no-choice task or ending the experimental session. Our dataset includes 44,883 decisions made by monkey SUN (obtained in 884 blocks spread over 242 days) and 19,292 decisions made by monkey FU (obtained in 571 blocks spread over 127 days). During choice task neural activity was not recorded. Further details on methods, including details that would be particularly of interest to neuroscientists, are provided in (Imaizumi et al., 2022).

3. Results

3.1 Identification of decision-relevant probability and payoff neurons

We begin by questioning whether it is biologically possible to observe utility and probability signals as distinct entities in the brain. If it was not, this would raise doubts about whether traditional economic models, such as Expected Utility and Prospect Theory, accurately reflect neural processes and whether the brain truly separates the encoding of payoff magnitudes and probabilities. Regardless of the answer, resolving this issue is crucial for how we conceptualize and interpret theories of risky choice.

However, this question remains open, as previous studies on neural probability weighting have not analyzed data at the level of individual neurons to identify separate neural signatures of utility and probability weighting.

3.1.1 Identification of probability and payoff neurons

To identify neurons that respond to probability and/or payoff magnitude, we divided our one hundred lotteries into categories of high (>50%) and low (≤50%) probability of receiving the payoff, and high (>0.5ml) versus low (≤0.5ml) payoff magnitude. For each neuron, we conducted two t-tests. The first test compared its activity in trials featuring lotteries with high versus low probability, and the second test compared its activity in trials with high versus low payoff lotteries. We defined neurons as probability neurons and payoff neurons if their respective t-tests were significant at the 5% level. This method was chosen over linear regression as it does not assume a specific functional form for probability representation in the brain[4]. It only requires that the average activity between low and high-probability trials is significantly different, which should hold under any pattern of probability coding. Probability-only neurons, if they exist, are particularly interesting as their activity in response to probability is not confounded by payoff magnitude.

Our initial findings, presented in Figure 2, show that neurons sensitive only to probability and neurons sensitive only to reward magnitude exist. In our sample, 22.6% of neurons encode probability exclusively, 18.7% encode payoff only, and 14.0% encode both probability and payoff. The discovery of neurons sensitive to probability but not payoff magnitude is crucial because it allows us to observe how probability is encoded in these neurons without interference from value or expected value signals.

Another insight is that both probability and payoff neurons are distributed across all four valuation regions of the brain. None of these regions is exclusively dedicated to either probability or payoff magnitude. This suggests that if the human brain shares this property, fMRI studies, which measure aggregated activity per voxel (a unit typically consisting of about one million neurons), may not be able to isolate the neural representation of probability weighting independent of the utility function. Thus, direct recordings from single neurons provide a unique opportunity to precisely measure how a chooser processes objective probability information.

---

[4] Previous studies in neuroeconomics, often used a linear regression to identify neurons sensitive to payoff magnitude. In the appendix, we show that this type of selection makes little difference to our results. For our purposes, our selection procedure allows us to remain agnostic about the functional form.

In the main body of the paper, we present findings from analyses that include only neurons significantly responsive to probability but not to payoff magnitude. Like prior studies (Enomoto et al., 2020; Yamada et al., 2013, 2018), we observe that some neurons respond positively and others negatively to payoffs. Similarly, across all four brain regions, we identify neurons that are significantly more active, as well as those that are significantly less active, in high versus low probability trials. In the paper, we include probability only neurons that respond positively or negatively to probability. In the appendix, we further validate the robustness of our results through supplementary analyses that incorporate all neurons that positively respond to probability.

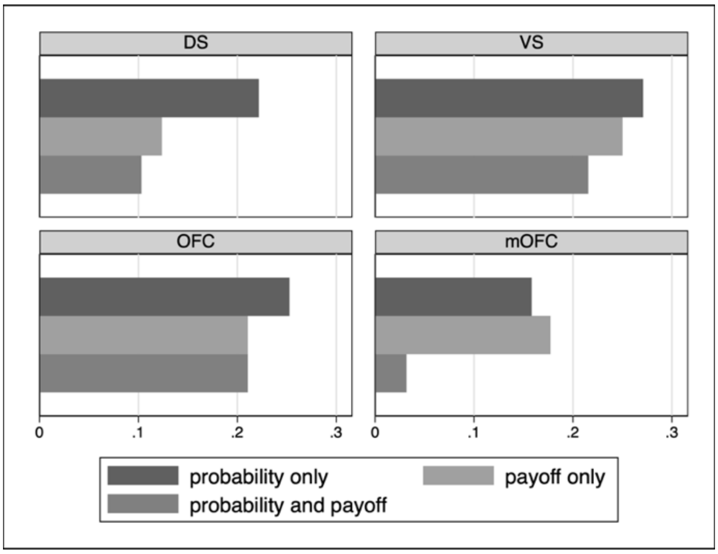

Figure 2. **Proportion of neurons in our sample that encode probability only (dark gray), payoff only (light gray), and both probability and payoff (bottom bar)** across dorsal striatum (DS), ventral striatum (VS), central orbitofrontal cortex (cOFC) and medial orbitofrontal cortex (mOFC).

3.1.2 Neural prediction of lottery choice

The activity of our probability and payoff neurons in the single-cue no choice task is interesting to economists only if they predict choice out of the sample. We find that their activity in the no choice task correctly predicts about 87% of decisions in the choice task. To calculate this percentage, for each of the 100 lotteries we calculated its associated average neural activity, which we will now refer to as neural subjective value as a shortcut. Since neural activity is observable, there is no need for estimation. Simply, for each of the 100 lotteries, and for each subject separately, we calculate the average neural activity across all instances when this lottery was presented. We base this measure on the activity of all neurons that either coded probability or payoff magnitude (not both). The activity of neurons that have a negative relationship with probability and/or payoff magnitudes is subtracted. To make sure that all neurons whose activity was measured in the study have the same impact on the neural subjective value, we first calculate the average activity for each lottery within a neuron and then average it across neurons. Using the calculated neural subjective value, we predicted that subjects would choose the lottery with the higher neural subjective value in each trial. The prediction accuracy was remarkably high, with subjective values predicting 87.22% of choices correctly for SUN and 87.74% for FU.

Another way to check whether neural data predicts choice is to run a logistic regression with choice as a dependent variable and average neural subjective value of each lottery as independent variables and compare this to a traditional logit with lottery information (probabilities and magnitudes) as independent variables. We present this analysis in Table 2. The dependent variable is equal to 1 if participant selected lottery 1 and 0 if participant selected lottery 2. In “neurons” model, we include only the neural subjective value variables (calculated according to the description in the paragraph above). In “lottery info” model, we include only the information about the lottery probability and magnitude that participants saw on the screen. In “both” model, we include both. Comparing the performance of “neurons” and “lottery info” models, it is evident that our out-of-sample neural measurements of subjective value are as good as the probability and magnitude information in predicting choice. Moreover, comparing the loglikelihood of “both” and “lottery info” models, we see that adding neural measurements to a traditional model with just lottery parameters (probability and magnitude) as explanatory variables, substantially increases loglikelihood, demonstrating the additional predictive power from neural measurement.

Table 2. **Neural and behavioral predictors of choice**. Logistic regression with dependent variable = 1 (=0) if participant selected lottery 1 (2). FR_n is the neurally measured subjective value of lottery *n*. p_n and m_n are the probability and magnitude of lottery *n*.

		SUN			FU	
	neurons	lottery info	both	neurons	lottery info	both
FR_1	0.8349***		0.5579***	0.8528***		0.4514***
	(0.0085)		(0.0156)	(0.0136)		(0.0197)
FR_2	-0.7701***		-0.4984***	-0.8590***		-0.6038***
	(0.0082)		(0.0156)	(0.0137)		(0.0208)
p_1		6.4320***	2.3598***		7.4183***	4.8205***
		(0.0761)	(0.1362)		(0.1284)	(0.1916)
m_1		7.1348***	3.4433***		6.6290***	4.7988***
		(0.0799)	(0.1301)		(0.1212)	(0.1622)
p_2		-5.7977***	-2.2544***		-7.2448***	-3.3076***
		(0.0725)	(0.1355)		(0.1269)	(0.1869)

m_2		-6.5482***	-3.3600***		-6.4980***	-3.6982***
		(0.0767)	(0.1301)		(0.1204)	(0.1589)
constant	0.1078***	-0.6546***	-0.0064	0.1740***	-0.0052	-1.0355***
	(0.0168)	(0.0572)	(0.2169)	(0.0274)	(0.0893)	(0.2010)
N	44,883	44,883	44,883	19,292	19,292	19,292
pseudo R2	0.5796	0.5679	0.5679	0.5904	0.593	0.6474
log likelihood	-13080	-13443	-12302	-5475	-5440	-4713

The accuracy with which our neural measurements predict choice is notable for several reasons. Firstly, our predictions are done out-of-sample. The neural recordings are collected in the single cue no-choice tasks and are then used to predict decisions in a separate task when participants decide which lottery they prefer. This creates a time separation between the measurements. Additionally, the context of presenting two lotteries at a time in the choice task likely influences the subjective value in ways not detectable in neural data from single-lottery presentations (Louie et al., 2013). Finally, we recorded only from a subset of neurons in the brain's value regions. The ability of this subset to predict out-of-sample behaviors so accurately is impressive.

In conclusion, we find evidence that the brain is capable of separately encoding probability and payoff magnitude, as commonly assumed in popular economic models of choice. For the first time, we demonstrate that the brain can be conceptualized as dedicating distinct neurons to encoding payoff probability and others to encoding payoff magnitude. Leveraging this discovery, we utilize the neural probability-payoff separability to describe the shape of the neural probability weighting function in the following section.

3.2 Neural probability weighting function

We have established that our neural measurements can predict choices and that there are neurons specifically sensitive to probability which activity is not confounded by value signals. In Figure 3, we present the neural probability weighting function. This figure is based on the activity of neurons that solely encode probability and do not encode payoff, providing the cleanest measure of brain’s response to probability. This dataset comprises 17,958 neural measurements (9,304 for SUN and 8,654 for FU) collected from 155 neurons (82 for SUN and 73 for FU). To construct Figure 3, we first calculate the average activity of each neuron at each probability level. We then compute the overall average activity at each probability level across all neurons, incorporating neurons that negatively respond to probability with a negative sign. This method ensures that all neurons contribute equally, regardless of the number of measurements recorded from each. To facilitate a direct comparison between probability weighting functions measured behaviorally and in the brain function, we transform this neural measure to a scale of [0,1] using the formula for the neurobiological probability response ratio (NPRR) previously used by Berns (2008):

$$FR_{norm} = \frac{FR - FR_{min}}{FR_{max} - FR_{min}}$$

where FR is the actual average firing rate at a given probability level, $FR_{min}$is the lowest and FR_{max} is the highest average firing rate across all probability levels. As the NPRR is an affine transformation, it does not distort the curvature of the neural probability weighting function. We find that the neural probability weighting in each monkey is slightly S-shaped. We confirm this by fitting FR_{norm} with a two-parameter probability weighting function[5], $w(p) = \frac{\delta p^{\gamma}}{\delta p^{\gamma} + (1-p)^{\gamma}}$, using nonlinear least-squares estimation via nl command in Stata (see Table 3 "neural" for the parameter estimates). The γ parameter, which gives the function S-shape does not statistically differ between the subjects. The δ parameter is significantly higher for FU ($p < 0.001$) which captures that the fixed point, $w(p) = p$, occurs for lower probability for FU.

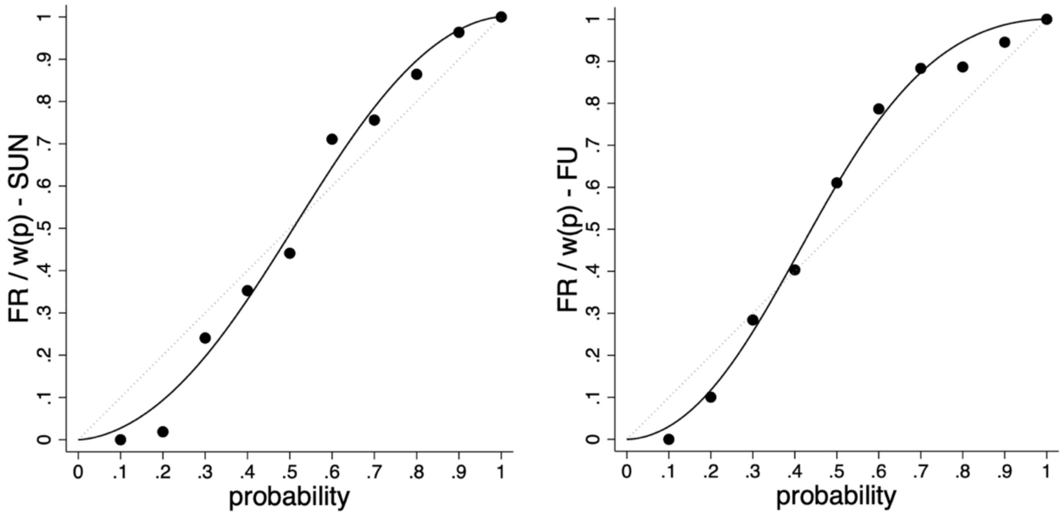

Figure 3. **Neural probability weighting function.** Includes data from all neurons that significantly respond to probability but do not respond to payoff magnitude. The dots represent FR_{norm} and the solid curves are Goldstein and Einhorn (1987) two-parameter probability weighting functions that best fit the neural data.

---

[5] In Tymula et al. (2023), we conducted a horserace between the expected value, expected utility, and prospect theory models with four different probability weighting functions using the same behavioral dataset (but we did not analyze the neural data). We used power utility function and various probability weighting functions (Goldstein & Einhorn, 1987; Kahneman & Tversky, 1979; Prelec, 1998). We found that prospect theory models with a two-parameter probability weighting function provided the best fit to behavioral data according to the Bayesian Information Criterion (BIC) (see Fig 2 in Tymula et al. (2023)). Following this finding, here we use the two-parameter probability weighting function (Goldstein & Einhorn, 1987).

In the Appendix, we conduct several robustness checks to validate our findings through three alternative approaches. First, we analyze the data using only those probability-only neurons that show a positive response to increasing probability, excluding any neurons whose activity decreases with probability magnitude (Figure S1). Second, we broaden our dataset to include all neurons that react to probability magnitude, encompassing those that also respond to both probability and payoff magnitudes (Figure S2). Third, we employ a different classification method for neurons; here, each neuron is classified as coding for probability and/or payoff magnitude based on whether its activity significantly ($p < 0.05$) correlates with either lottery probability or payoff magnitude. This classification is determined by a linear regression model ($FR = b_0 + b_1 * payoff + b_2 * probability$) run separately for each neuron (Figure S3). The results from these varied approaches are qualitatively similar, i.e. that the neural probability weighting function is slightly S-shaped, reinforcing the consistency of our original findings (Figures S1 - S3). These additional analyses enhance our confidence in the robustness of our results and that our original observations are not an artifact of the specific subset of neurons, or the method of analysis used in the main study.

3.3 Neural probability weighting – why is it curved?

Our neural probability weighting function clearly demonstrates that neurons distort objective probabilities, but the distortion pattern contradicts the commonly accepted inverse-S shape found in the literature. One possible explanation for this discrepancy is that the probability weighting function is not driven by fixed probability weighting parameters but is instead adaptive, shaped by the specific environment we created in the laboratory. Several contemporary theories suggest that probability weighting functions adapt to the complexity and statistical properties of the environment. Therefore, next we examine whether the observed neural probability weighting functions align with these theories given the experimental conditions we established for our participants.[6]

3.3.1 Cognitive complexity

Enke & Graeber (2023) argue that as cognitive uncertainty increases, probability weighting functions become attenuated and regress toward the intermediate option. In other words, in decisions characterized by higher cognitive uncertainty, the probability weighting functions should be shallower.

---

[6] We postpone the discussion of theories that suggested that probability weighting depends on other alternatives in the choice set (Bordalo et al., 2012; Glimcher & Tymula, 2023; Kontek & Lewandowski, 2018) to behavioral data analysis in a later section because neural recordings were made when only one lottery was shown and available.

They further propose that cognitive uncertainty is inherently linked to the complexity of the decision problem. Our study was not specifically designed to examine the effects of choice complexity. We did not purposefully manipulate complexity in our experimental design, and the single-cue no-choice trials do not involve meaningful complexity. However, since blocks of the single-cue no-choice task are interleaved with blocks of the choice task, it is possible that naturally occurring complexity or cognitive uncertainty from the choice task blocks could have influenced the neural probability weighting function measured in different blocks. Given the novelty of our data and the growing interest in complexity and cognitive uncertainty, we explore this aspect within our experimental framework.

First, we investigated whether certain lottery probabilities are systematically associated with more difficult choices on average. We define trials as more cognitively uncertain or complex when the expected values of the two lotteries being compared are harder to distinguish. For each trial within the choice task, we calculate the absolute value of the difference between the expected values of the two lotteries. We then average these differences across all trials and participants. In Figure S4, we plot this difference as a function of lottery probability.

Our analysis reveals that the most challenging decisions occur most frequently for probabilities around 0.3–0.4. According to Enke and Graeber's (2023) argument, we would expect the probability weighting function to be the shallowest in this range and steeper for both lower and higher probabilities. However, this is generally inconsistent with the shape of our observed neural probability weighting function, likely because the potential effects of complexity may be too subtle to detect in our experiment.

3.3.2 Efficient adaptation to the distribution of probabilities

A series of theoretical models, choice data, and neural data show that payoff and/or probability encoding functions are flatter (i.e. leading to discrimination difficulties) for values that are experienced less frequently or lead to lower financial benefits in given choice environment (Frydman & Jin, 2023; Glimcher, 2022; Glimcher & Tymula, 2023; Khaw et al., 2021; Netzer et al., 2025; Robson et al., 2023; Stewart et al., 2006; Vieider, 2025; Woodford, 2012). The biological explanation for this phenomenon is based on the fact that brain is noisy and has a limited precision with which it encodes information for decision making. As a result, this body of work argues, choices are better (i.e. lead to higher expected payoffs or fewer errors) if they are guided by utility and probability weighting functions that adapt to the statistics of the environment. Empirical evidence supports this research with, for example, Zhang et al. (2020) showing that probability weighting functions adapt to the range of probabilities used in the task.

The exact methodological details differ among the papers that theoretically establish the optimal shape of the probability weighting function (Vieider, 2025). However, what they all have in common is their focus on justifying the frequently observed inverse S-shaped functions. For example, Stewart et al. (2006) observe that small and large probabilities are more frequently experienced than intermediate probabilities in language, in real life, and in the experiments that elicit probability weighting. Therefore, the probability weighting function is inverse-S shaped (thus steeper for low and high probabilities) to increase discriminability in these more frequently encountered probabilities. Frydman & Jin (2023) and Khaw et al. (2021) additionally argue that when cognitively uncertain choosers resort to the intermediate prior of 50% chance which leads to an inverse-S shape. So why do we instead observe S-shaped neural probability weighting function in our study?

We suggest that the reason is because our task is unique relative to other experiments in that we uniformly sample the full space of probabilities (in 0.1 increments). This means that none of the probabilities is more frequently experienced than other probabilities while participants perform the study. Therefore, small and large probabilities should not be prioritized according to discriminability because they are not experienced more frequently. Another reason why our study is different to prior research, is because unlike in the experiments with human participants, we have full control of the prize distribution in the monkeys' environment. They have spent many months learning the stimuli and then performing the task, all the time experiencing uniformly distributed probabilities. Over this period, on weekdays, all of their liquid consumption came from the decisions they made in the study. This gives us control over participants' payoff environment that is unmatched by any study with human participants. As we demonstrate below, S-shaped probability weights should optimally emerge for choosers who want to maximize their payoffs in our experimental environment. This provides further support that the statistical probabilities of the choice environment influence observed probability weighting functions.

To establish the optimal probability weighting function, we follow the approach in Glimcher et al. (2025). We assume that participants' goal is to maximize their earnings. Since there is an infinite number of the candidate probability weighting functions to consider, to make our analytical approach feasible, we make the following simplifications: (1) we assume that the probability weighting function is increasing in probability, (2) we follow the convention that $w(0.1) = 0.1$[7] and $w(1) = 1$, (3) we calculate the optimal probability weights for the remaining eight probability levels used in our experiment (from 0.2 to 0.9 in 0.1 increments), (4) we allow $w(p)$ to take values from 0.1 to 1 in 0.025 increments, (5) without loss of generality, we assume that the payoff is always equal to 1, and finally (6) we assume that choosers encode utilities with and additive noise, $\varepsilon \sim N(0, \sigma^2)$[8] and vary the standard deviation of noise from 0.05 to 0.9. While some of our assumptions may seem restrictive, they nevertheless leave us with 177,232,627 candidate probability weighting functions to evaluate.

We proceed as follows: First, we generate all possible candidate probability weighting functions according to assumptions 1-4. Second, we generate all possible 45 binary choice sets that the participant could have faced in our study (we disregard choice sets where both options are equal). Since in our environment all probabilities are equally likely to occur, all the possible binary choice sets are also equally likely to be encountered. Next, for each possible binary choice set and for each candidate probability weighting function, we calculate expected earnings at a given level of noise. The expected earnings when choosing between getting the prize with probability p_i or p_j are given by a standard formula: $EV\{p_i, p_j\} = \Phi\left(\frac{w(p_j)-w(p_i)}{\sqrt{2\sigma^2}}\right) p_j u(1) + (1 - \Phi\left(\frac{w(p_j)-w(p_i)}{\sqrt{2\sigma^2}}\right)) p_i u(1)$, where $\Phi\left(\frac{w(p_j)-w(p_i)}{\sqrt{2\sigma^2}}\right)$ is the probability that the chooser picks p_j. We calculate the expected earnings under a given probability weighting function at a given level of noise by adding the expected earnings for all 45 possible choice sets $\{p_i, p_j\}$. The program, written in R, is available in the online appendix.

Figure 4 illustrates the probability weighting functions that maximize earnings in our experimental environment for different levels of noise. As one would expect when the level of noise is low, a linear probability weighting function (lightest gray) is optimal. When the noise level is high (dark line), the probability weighting function is a step-function—the noisiest chooser distinguishes between low and high probabilities but not within each of these categories. For the intermediate levels of noise, the optimal probability weighting function is S-shaped, just as observed in our neural data.

---

[7] The lowest probability stimulus in the study is 0.1 which is why we do not start at 0.

[8] The cumulative distribution function (CDF) of this distribution is denoted by $\Phi(.)$ and the probability distribution function (PDF) as $\phi(.)$.

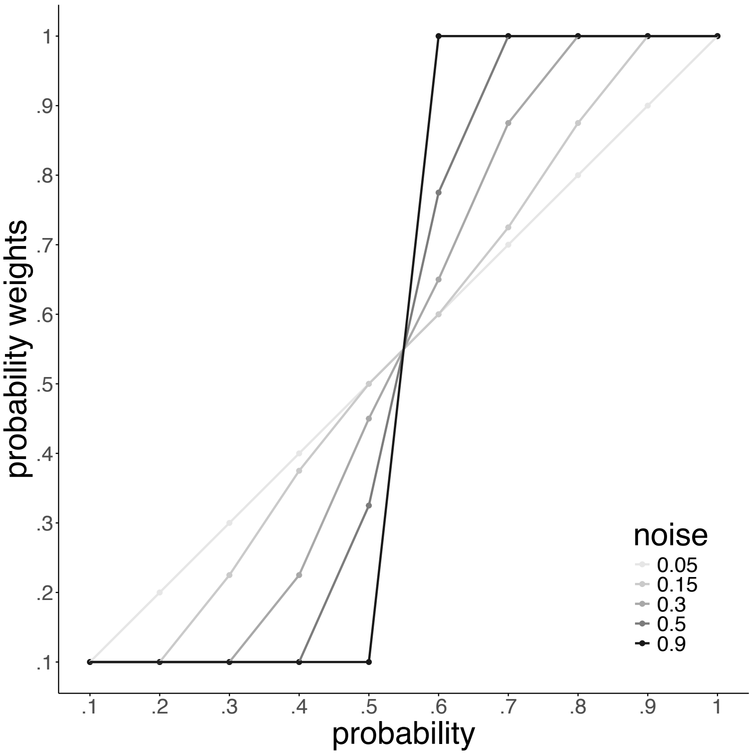

Figure 4. **Optimal probability weighting functions for different levels of noise**.

3.4 Traditional behaviorally estimated probability weighting versus brain activity

A series of theoretical papers (Bordalo et al., 2012; Glimcher & Tymula, 2023; Kontek & Lewandowski, 2018; Schneider & Day, 2018) have argued that neglecting the influence that all options in the choice set have on utility can lead to distorted estimates of the probability weighting function. Therefore, we next investigate whether the probability weighting functions observed in neural data in response to the presentation of a single lottery resemble those estimated from binary choices using standard behavioral economics methods which typically employ utility functions that do not account for the impact of the other alternatives on utility. If we see differences between the neural measured and behaviorally estimated probability weighting functions this may warrant an investigation into the unaccounted-for salience or normalization of value.

Here, we use data from the choice task that contains the same subjects' decisions made in different blocks of trials in the same experimental sessions. Using this data, we estimate probability weighting functions from behavior within a random utility framework. A lottery, denoted as $L(m, p)$ represents a gamble that pays m with probability p, and 0 otherwise. We adopt a popular two-parameter probability weighting function (Goldstein & Einhorn, 1987)[9]:

$$w(p) = \frac{\delta p^{\gamma}}{\delta p^{\gamma} + (1-p)^{\gamma}}$$

For the utility function, we use a power utility function, commonly utilized in economic literature:

$$u(m) = m^{r}$$

The expected utility of the lottery is then calculated as $U(L) = w(p)u(m)$. The probability that participants choose the lottery on the right side (L_R) over the one on the left (L_L) is estimated using a logistic choice function, $P(L_R) = \frac{1}{1+e^{-Z}}$ where $Z = \frac{U(L_R)-U(L_L)}{\varepsilon}$ and free parameter ε captures the degree of stochasticity observed in choice. We fit the data by maximizing the log-likelihood of the observed choices. All estimations are done in Stata 18.

In Figure 5, just as in Figure 3, we replot the observed neural probability weighting with black dots, and the two-parameter probability weighting function that best fits the neural data is depicted as a black solid line. The black longdash line represents the probability weighting function estimated from behavior using the above model with power utility function. Visually, Figure 5 clearly suggests that the behavioral probability weighting function (longdash) substantially differs from both the observed neural activity (dot) and the estimated (solid line) neural probability weighting functions. One can further assess the statistical significance of these differences by comparing each subject's probability weighting parameters estimated from the neural data against those derived from behavioral measurements—see Table 3 under the headings 'neural' versus 'power'. For SUN, δ and γ parameters estimated from neural activity and behavior are significantly different at 0.001 and 0.1 significance levels respectively. For FU, δ and γ parameters estimated from neural activity and behavior are significantly different at 0.01 and 0.001 significance levels respectively.

There is a stark mismatch between neurally measured and behaviorally estimated probability weighting and the behaviorally estimated probability weighting is markedly different form probability weighting that would be an efficient adaptation to the statistical properties of our experimental

---

[9] In Tymula et al. (2023), we conducted a horserace between the expected value, expected utility, and prospect theory models with four different probability weighting functions using the same behavioral dataset (but we did not analyze the neural data). We used power utility function and various probability weighting functions (Goldstein & Einhorn, 1987; Kahneman & Tversky, 1979; Prelec, 1998). We found that prospect theory models with a two-parameter probability weighting function provided the best fit to behavioral data according to the Bayesian Information Criterion (BIC) (see Fig 2 in Tymula et al. (2023)). Following this finding, here we use the two-parameter probability weighting function (Goldstein & Einhorn, 1987).

environment (Figure 4). A potential explanation is that the neural measurements were done when only a single lottery was presented, and the behavioral measurements are estimated based on a choice task where *two* lotteries were presented at a time. It is possible that the other option affects utility and when this effect is unaccounted for, we get distorted probability weights estimated from behavior. Therefore, next we explore whether using more realistic utility functions that account for context created by other alternatives in the choice set will result in behaviorally estimated probability weighting functions that better match brain activity.

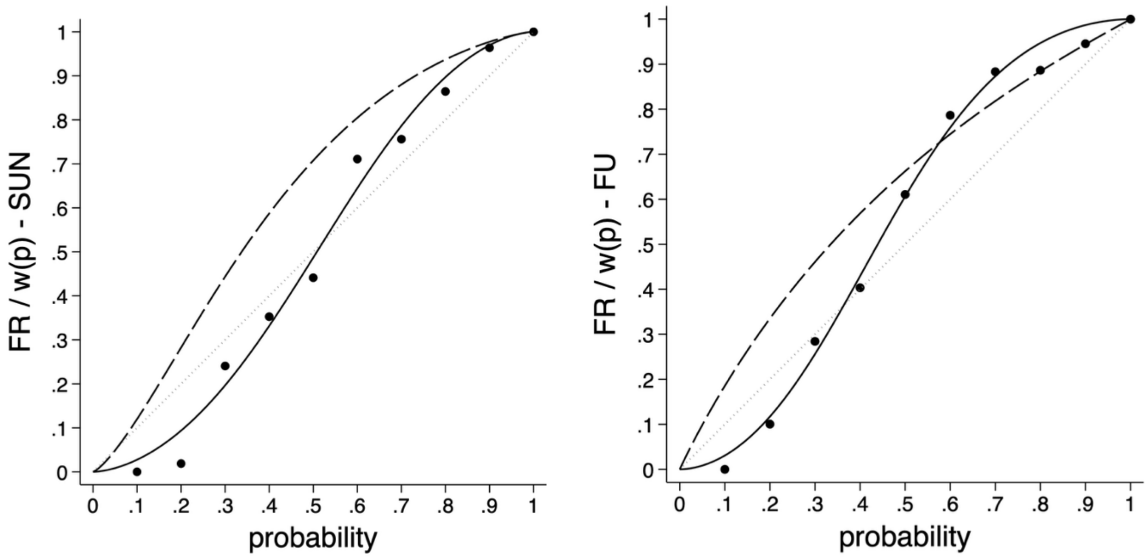

Figure 5. **Neural versus behavioral probability weighting function.** Includes data from neurons that respond to probability but not to payoff magnitude. The dots are FR_{norm} and the solid curve is Goldstein and Einhorn (1987) probability weighting function fit to neural data. The long dash and curve is Goldstein and Einhorn (1987) probability weighting function estimated using behavioral data (with power utility function).

Table 3. Neural and behavioral estimates of two-parameter Goldstein and Einhorn (1987) probability weighting functions. For the probability weighting parameters, α and r stars indicate significant differences from 1. For the rest of the parameters from 0. Neural is the best fit for the neural data, power are the behavioral estimates using the power utility function, ESVT are behavioral estimates using the ESVT utility function, salience are the behavioral estimates using salience model.

	SUN				FU			
	neural	power	DN	salience	neural	power	DN	salience
				probability weighting parameters				
δ	0.949	2.419***	1.124+	1.598***	1.539***	1.963***	0.915	1.158
	(0.091)	(0.055)	(0.065)	(0.035)	(0.118)	(0.075)	(0.085)	(0.038)
γ	1.599**	1.314***	1.715***	1.354***	1.768***	0.982	1.219***	1.034
	(0.152)	(0.022)	(0.045)	(0.022)	(0.125)	(0.023)	(0.043)	(0.022)
				utility parameters				
α or r		0.879***	1.657***	0.614***		0.578***	1.046	0.335***
		(0.016)	(0.071)	(0.016)		(0.019)	(0.069)	(0.018)
ω			0.093***				0.070***	
			(0.007)				(0.009)	
ε		0.065***	0.015***	0.051***		0.064***	0.022***	0.053***
		(0.001)	(0.002)	(0.001)		(0.001)	(0.003)	(0.001)
N	10	44883	44883	44883	10	19292	19292	19292
BIC		19678	19494	19633		8681	8621	8949

+ $p<0.1$, * $p<0.05$, ** $p<0.01$, *** $p<0.001$

3.4.1 Hidden influence of context-effects on behavioral probability weighting

While it may be puzzling that neural and behaviorally estimated probability weighting functions do not align, Bruhin et al. (2022) have recently provided evidence from incentivized Allais Paradox experiments that choices that depart from Expected Utility are equally often driven by probability weighting and choice set context dependency, conceptualized in their case as salience (Bordalo et al., 2012). This suggest that the context created by the other alternatives—which is usually ignored in the estimation of Prospect Theory parameters—can often explain observed "anomalies" as well as, or better than, probability weights. In addition to the salience theory (Bordalo et al., 2012), divisive normalization (Glimcher & Tymula, 2023), range normalization (Kontek & Lewandowski, 2018), and target adjusted utility (Schneider & Day, 2018) are other recent examples of models that theoretically and through simulated examples illustrated how unaccounted-for choice set context dependency in the utility function can bias the estimates of the probability weighting function. Next, we investigate whether using a more neurobiologically realistic utility function removes the mismatch between the neural and behaviorally-estimated probability weighting. As examples, we consider salience model for its popularity in economics (Bordalo et al., 2012) and divisive normalization model for its popularity in neuroscience and neuroeconomics (Carandini & Heeger, 2012; Louie et al., 2013).[10] We follow the same structural estimation procedures as in the beginning of this section, with some modifications for each model, as described below.

To assess how incorporating salience influences the estimated probability weighting function, we follow the original salience model by (Bordalo et al., 2012). In the salience model, the expected utility of lottery L_i is then given by:

$$U^s(L_i) = w(p_i)s(m_i)u(m_i)$$

where $s(m_i)$ is the relative salience of the states with payoff m_i and the probability weighting and the utility function are defined as before. Let $\sigma(x_1, x_2) = \frac{|x_1 - x_2|}{x_1 + x_2}$ be the salience function of the state where L_1 pays x_1 and L_2 pays x_2 and $\sigma(0,0) = 0$. Then $s(m_1) = \frac{\sigma(m_1,m_2)+\sigma(m_1,0)}{\sigma(m_1,m_2)+\sigma(m_1,0)+\sigma(0,m_2)+\sigma(0,0)}$ and $s(m_2) = \frac{\sigma(m_1,m_2)+\sigma(0,m_2)}{\sigma(m_1,m_2)+\sigma(m_1,0)+\sigma(0,m_2)+\sigma(0,0)}$.

In the divisive normalization model, the expected utility of lottery L_1 when choosing between L_1 and L_2 is:

$$U^{DN}(L_1, L_2) = \frac{w(p_1)m_1^\alpha}{\varepsilon + \omega(w(p_1)m_1^\alpha + w(p_2)m_2^\alpha)}$$

where α is a parameter called predisposition. We follow Webb et al. (2021) and use parameter ω to nest the power utility model. If we estimate that $\omega = 0$, we are back in the prospect theory framework with a power utility function where ε becomes the noise level. If $\omega \neq 0$, then we conclude that the utility of the lottery L_1 is normalized by itself and the other elements of the choice set. The remaining details of the structural estimation are the same as before.

Table 3 summarizes all the estimates. Regarding the divisive normalization model, for both SUN and FU, we find that the parameter ω is significantly different than zero. Moreover, the Bayesian

---

[10] Divisive normalization is a well-established neural computation found throughout the neural pathway from the retina to the cortex (Carandini & Heeger, 2012; Louie et al., 2011; Yamada et al., 2018). It has been theoretically demonstrated to be an efficient computation for encoding value (Steverson et al., 2019), especially with Pareto-distributed payoffs (Bucher & Brandenburger, 2022). Divisive normalization is a relevant model because it can account for many behaviors associated with Prospect Theory but without the probability weighting function or the discontinuity created by the loss aversion parameter (Glimcher & Tymula, 2023). It also effectively captures various choice set effects (Guo & Tymula, 2021; Khaw et al., 2017; Landry & Webb, 2021; Louie et al., 2013; Webb, 2020; Webb et al., 2021), and recent research suggests that our brains may universally use this encoding mechanism (Kurtz David et al., 2023). A comprehensive summary of this model and its historical context is provided elsewhere (Carandini & Heeger, 2012; Glimcher, 2022; Glimcher & Tymula, 2023; Louie et al., 2015). We note that in our particular design where the other lottery payoff is always zero, other normalization models would not explain probability weighting.

Information Criterion is lower than in the traditional power utility function model, indicating a better fit. Based on these results—significant parameter values of ω, slightly lower BIC—we conclude that behavioral data is more accurately estimated using the neurobiologically realistic utility function as specified by the divisive normalization model than using a simple power utility function. Regarding the salience model, we find that it offers a better behavioral fit for SUN, but not for FU.

Next, we compare all models based on how well their estimates approximate the probability weighting function estimated in the brain. In Figure 6, we plot the probability weighting function in addition to the neural probability weighting (black dots and black solid best fit curve) and the behaviorally estimated probability weighting function with power utility function (dashed black line), we plot the estimated probability weighting function in the salience (red) and divisive normalization (blue) models. Visually both salience and divisive normalization models improve the match. In particular, for SUN, the blue curve estimated in the divisive normalization model closely aligns with the average neural probability weighting signal represented by black dots, indicating a near-perfect match.

To formally assess how accurately the probability weighting functions estimated with salience, divisive normalization, and traditional power utility models approximate the neural measurements, we calculate residuals for each model. These residuals are determined by comparing each model's estimate of the probability weight to the averaged and normalized neural measurement at each probability level. We then compute the sum of square errors for these residuals.

The results show that for both participants, the sum of square errors is lower for salience model than for the power utility model (0.007 vs. 0.028, p=0.020 for SUN and 0.008 vs. 0.016, p=0.286 for FU). Similarly, it is lower for the divisive normalization model than for the power utility model (0.002 vs. 0.028, p=0.015 for SUN and 0.009 vs 0.016, p=0.468 for FU). This suggests that estimating behavior with the salience and divisive normalization models results in probability weighting functions that are a better description of the underlying neural activity. Comparing between the divisive normalization and salience models, we find that the prediction is significantly better for divisive normalization for SUN (p=0.043) but there is no difference (p=0.678) for FU. These findings support the superior ability of the recent models that account for the contextual effects created by the choice set to estimate the neural probability weighting from behavior, particularly evident in the SUN dataset.

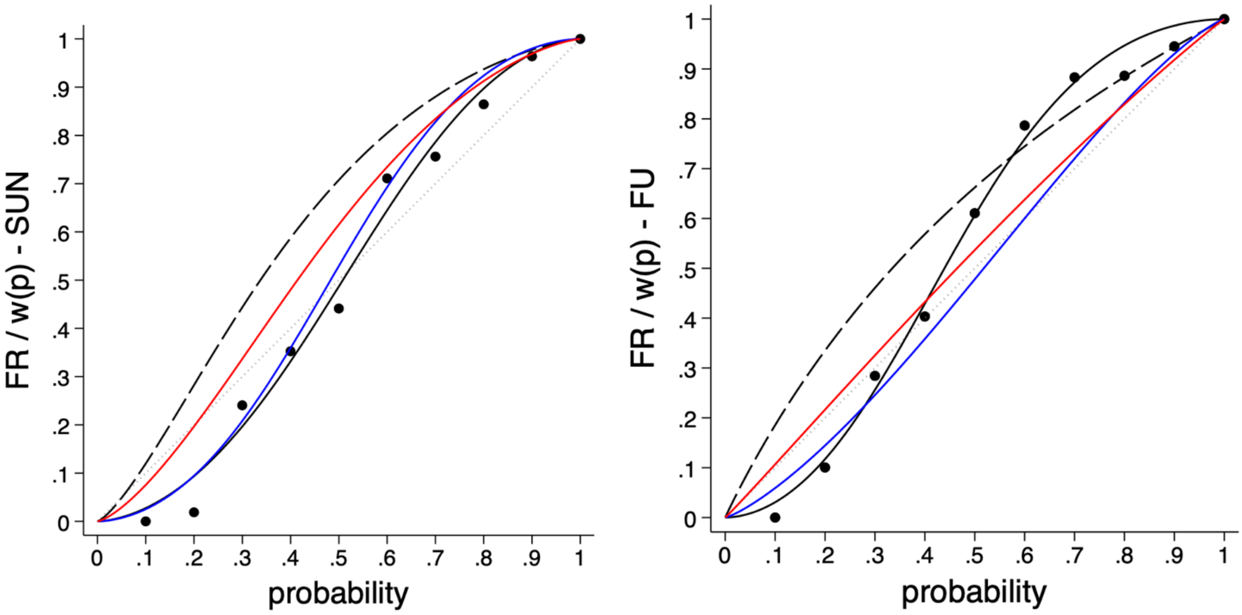

Figure 6. **Neural versus behavioral probability weighting functions.** Includes data from neurons that respond to probability but not to payoff magnitude. The dots are FR_{norm} and the solid curve is Goldstein and Einhorn (1987) probability weighting function fit to neural data. The long dash and curve is Goldstein and Einhorn (1987) probability weighting function estimated using behavioral data (with power utility function). The red and blue curves are Goldstein and Einhorn (1987) probability weighting function estimated using behavioral data using salience model (in red) and divisive normalization model (in blue).

4. Discussion

We demonstrate that probability weighting exists in the brain separately to the utility of payoffs. The observed probability weighting function is not linear. Instead, it is S-shaped. The normative literature on probability weighting has so far mainly focused on demonstrating under what circumstances the traditional inverse-S shape is efficient. For instance, Stewart et al. (2006) and Frydman & Jin (2023) argue that individuals are more likely to encounter small or large probabilities, leading to an argument for an inverse S-shaped probability weighting function that helps to distinguish between more frequently occurring probabilities. In our experimental design, probabilities are uniformly drawn from a range of 10% to 100%, in 10% increments, thus eliminating such advantage. We show that in such an environment S-shaped probability weighting would maximize earnings and this is the shape that we observe in brain activity. It is perhaps not surprising given that our subjects had an extensive experience with the stimuli which included 10-months-long training and then performed the task daily for another couple of months. The S-shaped distortions that we document offer a new outlook on probability weighting which is particularly relevant given that the existing theoretical explanations of probability weighting largely focused on figuring out under what conditions probability weighting should be inverse S-shaped (Blavatskyy, 2007; Enke & Graeber, 2023; Herold & Netzer, 2023; Steiner & Stewart, 2016; Zhang et al., 2020) highlighting it may be additionally worthwhile to consider the flexibility in the shapes of probability weighting as well (Bordalo et al., 2012; Frydman & Jin, 2023; Glimcher & Tymula, 2023).

Our paper represents a unique contribution. The measurement of probability weighting has been seen as a challenge for a long time. Previous researchers provided solutions by for example designing new nonparametric ways to estimate the utility and probability weighting functions without functional assumptions (Gonzalez & Wu, 1999), designing new experimental methods to estimate Prospect Theory parameters(Abdellaoui, 2000), or by asking participants for frequency judgements (Zhang et al., 2020). Ours is the first study to measure the brain's response to probability through single-neuron activity dedicated solely to probability, and not payoff magnitude. This is the cleanest and most direct measurement of how the brain encodes probability that is possible using current technology. Consequently, the probability weighting functions we document are directly observed, rather than estimated from behavior, and are not confounded by or entangled with utility.

Demonstrating that such signature of probability weighting can be measured in the brain opens the door to answering more complex questions in the future such as exploring how efficient coding constraints (Frydman & Jin, 2023; Glimcher, 2022; Glimcher & Tymula, 2023; Louie & Glimcher, 2012; Polanía et al., 2019), perceptual factors (Oprea, 2024), and salience (Bordalo et al., 2012) relate to probability encoding in the brain. In our paper, we demonstrate the importance of using a realistic context-dependent utility function when recovering probability weighting from behavior. Notably, for both subjects, the probability weighting estimated from behavior closely aligns with (and for one subject, essentially mirrors) the neural measurement when we model utility using the canonical, divisive normalization utility model from neuroscience and neuroeconomics or incorporate salience. There is a big mismatch in the neural and behaviorally-estimated probability weighting if we instead use the usual power or CRRA utility function. This suggests that as our understanding of how the brain encodes value improves, we will become increasingly better at recovering probability weighting from behavior.

Obtaining such high quality neural (and behavioral) measurements has been possible because the subjects in our study are macaques. While rhesus macaques are not human, they are primates, and we share roughly 93% of our DNA sequences (Wang et al., 2014). Many studies demonstrated a close parallel between human and monkey behaviors and brain function. Economists many years before argued how economics can learn from animal studies (Kagel et al., 1995). This study is a perfect example that illustrates these advantages – we directly observe a noiseless probability weighting signal in the brain that is not accessible with choice data. We also observe a lot of decisions, ensuring that our estimated probability weighting from behavior is more reliable. Previous monkey studies estimated probability weighting from behavior (Farashahi et al., 2018; Ferrari-Toniolo et al., 2019; Fujimoto & Minamimoto, 2019; Imaizumi et al., 2022; Stauffer et al., 2015; Tymula et al., 2023) but none of them measured or reported probability weighting in single neuron activity. Additionally, our lotteries were constructed with the highest number of different probability levels in any monkey study so far, which enhances the precision of our measurement in both brain activity and behavior. The existing studies with monkey subjects were not conclusive about the shape of the probability weighting estimated from behavior but demonstrated that some features of how choices are presented affect the estimates, reinforcing the need for studies that establish how the perception of probabilities changes based on environment or experience (Ferrari-Toniolo et al., 2019; Tymula et al., 2023).

Kahneman and Tversky, when they conceived prospect theory in the 1970s, could only rely on observed choices to guide their assumptions about how probability and payoff value are encoded. Since then, science has made remarkable progress. We can now directly observe neuronal activity that guides choice using a range of tools. A series of studies in the 2000s set out to measure and describe the functional properties of the probability weighting function using non-invasive functional magnetic resonance imaging (fMRI) techniques that measure brain activity by tracking blood oxygenation levels throughout the brain (Abler et al., 2006; Berns et al., 2008; Hsu et al., 2009; Preuschoff et al., 2006; Tobler et al., 2008). Remarkably, these neuroeconomic studies generally agree that value-coding areas of the brain are capable of encoding probability, both with and without distortions. Consistent with our finding that probability is encoded by neurons across different value regions in the brain, none of these earlier fMRI studies pointed to a specific brain region that encodes probability alone. Given that a typical unit of brain activity in an fMRI study (voxel) aggregates the activity of approximately half million neurons, it is unlikely that an fMRI study would be able to pick up probability signal independent of utility in a lottery choice task. We overcome this shortcoming by using a direct measurement of brain activity, employing the largest number of probability levels used to date with monkey subjects. While it is possible that our probability-only neurons, have other functions than just encoding probability, what is most important for our research is that they do not confound the probability and payoff magnitude. As such, they allow us to directly observe the probability weighting function, with as little measurement error as possible using the current technology.

Bibliography

Abdellaoui, M. (2000). Parameter-Free Elicitation of Utility and Probability Weighting Functions. *Management Science*, *46*(11), 1497–1512.

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. https://doi.org/10.1016/j.neuroimage.2006.01.001

Berns, G. S., Capra, C. M., Chappelow, J., Moore, S., & Noussair, C. (2008). Nonlinear neurobiological probability weighting functions for aversive outcomes. *NeuroImage*, *39*(4), 2047–2057. http://linkinghub.elsevier.com/retrieve/pii/S1053811907009615

Blavatskyy, P. R. (2007). Stochastic expected utility theory. *Journal of Risk and Uncertainty*, *34*(3). https://doi.org/10.1007/s11166-007-9009-6

Bordalo, P., Gennaioli, N., & Shleifer, A. (2012). Salience Theory of Choice Under Risk. *The Quarterly Journal of Economics*, *127*(3), 1243–1285. https://doi.org/10.1093/qje/qjs018

Bruhin, A., Manai, M., & Santos-Pinto, L. (2022). Risk and rationality: The relative importance of probability weighting and choice set dependence. *Journal of Risk and Uncertainty*, *65*(2). https://doi.org/10.1007/s11166-022-09392-x

Bucher, S., & Brandenburger, A. (2022). Divisive Normalization is an Efficient Code for Multivariate Pareto-Distributed Environments. *Proceedings of the National Academy of Sciences*, *119*(40).

Carandini, M., & Heeger, D. J. (2012). Normalization as a canonical neural computation. *Nature Reviews. Neuroscience*, *13*(1), 51–62. https://doi.org/10.1038/nrn3136

Enke, B., & Graeber, T. (2023). Cognitive Uncertainty. *Quarterly Journal of Economics*, *138*(4). https://doi.org/10.1093/qje/qjad025

Enomoto, K., Matsumoto, N., Inokawa, H., Kimura, M., & Yamada, H. (2020). Topographic distinction in long-term value signals between presumed dopamine neurons and presumed striatal projection neurons in behaving monkeys. *Scientific Reports*, *10*(1). https://doi.org/10.1038/S41598-020-65914-0

Farashahi, S., Azab, H., Hayden, B., & Soltani, A. (2018). On the Flexibility of Basic Risk Attitudes in Monkeys. *Journal of Neuroscience*, *38*(18), 4383–4398. https://doi.org/10.1523/JNEUROSCI.2260-17.2018

Fehr-Duda, H., Bruhin, A., Epper, T., & Schubert, R. (2010). Rationality on the rise: Why relative risk aversion increases with stake size. *Journal of Risk and Uncertainty*, *40*(2). https://doi.org/10.1007/s11166-010-9090-0

Ferrari-Toniolo, S., Bujold, P. M., & Schultz, W. (2019). Probability Distortion Depends on Choice Sequence in Rhesus Monkeys. *Journal of Neuroscience*, *39*(15), 2915–2929. https://doi.org/10.1523/JNEUROSCI.1454-18.2018

Frydman, C., & Jin, L. J. (2023). On the Source and Instability of Probability Weighting. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4546440

Fujimoto, A., & Minamimoto, T. (2019). Trait and state-dependent risk attitude of monkeys measured in a single-option response task. *Frontiers in Neuroscience*, *13*(JUL), 467953. https://doi.org/10.3389/FNINS.2019.00816/BIBTEX

Glimcher, P. (2022). Efficiently irrational: illuminating the riddle of human choice. *Trends in Cognitive Sciences*, *26*(8), 669–687. https://doi.org/10.31234/OSF.IO/ZHYMG

Glimcher, P., Sinha, S., & Tymula, A. (2025). *Optimal Utility: Endogenizing the Cardinal Representation of Riskless Subjective Value in Cognitively Constrained Chooser*.

Glimcher, P., & Tymula, A. (2023). Expected Subjective Value Theory (ESVT): a representation of decision under risk and certainty. *Journal of Economic Behavior and Organisation*, *207*, 110–128.

Goldstein, W. M., & Einhorn, H. J. (1987). Expression Theory and the Preference Reversal Phenomena. *Psychological Review*. https://doi.org/10.1037/0033-295X.94.2.236

Gonzalez, R., & Wu, G. (1999). On the shape of the probability weighting function. *Cognitive Psychology*, *38*(1), 129–166. https://doi.org/10.1006/cogp.1998.0710

Guo, J., & Tymula, A. (2021). Waterfall illusion in risky choice – exposure to outcome-irrelevant gambles affects subsequent valuation of risky gambles. *European Economic Review*, *139*, 103889. https://doi.org/10.1016/J.EUROECOREV.2021.103889

Herold, F., & Netzer, N. (2023). Second-best probability weighting. *Games and Economic Behavior*, *138*. https://doi.org/10.1016/j.geb.2022.12.005

Hsu, M., Krajbich, I., Zhao, C., & Camerer, C. F. (2009). Neural response to reward anticipation under risk is nonlinear in probabilities. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *29*(7), 2231–2237.

Imai, T., Nunnari, S., Wu, J., & Vieider, F. (2025). *Meta-Analysis of Risk-Taking Propensities*.

Imaizumi, Y., Tymula, A., Tsubo, Y., Matsumoto, M., & Yamada, H. (2022). A neuronal prospect theory model in the brain reward circuitry. *Nature Communications*, *13*(1). https://doi.org/10.1038/s41467-022-33579-0

Kagel, J. H., Battalio, R. C., & Green, L. (1995). *Economic choice theory: An experimental analysis of animal behavior*. Cambridge University Press.

Kahneman, D., & Tversky, A. (1979). Prospect Theory - analysis of decision under risk. *Econometrica*, *47*(2), 263–292.

Kandel, E. R., Koester, J. D., Mack, S. H., & Siegelbaum, S. A. (2021). *Principles of Neural Science* (6th ed.). McGraw Hill / Medical.

Khaw, M. W., Glimcher, P. W., & Louie, K. (2017). Normalized value coding explains dynamic adaptation in the human valuation process. *Proceedings of the National Academy of Sciences*, *114*(48), 12696–12701. https://doi.org/10.1073/pnas.1715293114

Khaw, M. W., Li, Z., & Woodford, M. (2021). Cognitive Imprecision and Small-Stakes Risk Aversion. *Review of Economic Studies*, *88*(4). https://doi.org/10.1093/restud/rdaa044

Kontek, K., & Lewandowski, M. (2018). Range-Dependent Utility. *Management Science*, *64*(6), 2473–2972. https://doi.org/10.2139/ssrn.2307858

Kurtz David, V., Alladi, V., Sinha, S., Bucher, S., Brandenburger, A., Louie, K., Dewan, A., Glimcher, P., & Tymula, A. (2023). *Inefficient Divisive Normalization: Human Choosers Employ Divisive Normalization Even When They Should Not*. 544.

Landry, P., & Webb, R. (2021). Pairwise Normalization: A Theory of Multi-Attribute Choice. *Journal of Economic Theory*, *193*.

Louie, K., & Glimcher, P. W. (2012). Efficient coding and the neural representation of value. *Annals of the New York Academy of Sciences*, *1251*, 13–32. https://doi.org/10.1111/j.1749-6632.2012.06496.x

Louie, K., Glimcher, P. W., & Webb, R. (2015). Adaptive neural coding: From biological to behavioral decision-making. *Current Opinion in Behavioral Sciences*, *5*, 91–99. https://doi.org/10.1016/j.cobeha.2015.08.008

Louie, K., Grattan, L. E., & Glimcher, P. W. (2011). Reward value-based gain control: Divisive normalization in parietal cortex. *The Journal of Neuroscience*, *31*(29), 10627–10639. https://doi.org/10.1523/JNEUROSCI.1237-11.2011

Louie, K., Khaw, M. W., & Glimcher, P. W. (2013). Normalization is a general neural mechanism for context-dependent decision making. *Proceedings of the National Academy of Sciences of the United States of America*, *110*(15), 6139–6144. https://doi.org/10.1073/pnas.1217854110

Netzer, N., Robson, A., Steiner, J., & Kocourek, P. (2025). Risk perception: measurement and aggregation. *Journal of the European Economic Association*.

Oprea, R. (2024). Decisions Under Risk are Decisions Under Complexity. *American Economic Review*, *12*(114), 3789–3811.

Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, *400*(6741), 233–238. https://doi.org/10.1038/22268

Polanía, R., Woodford, M., & Ruff, C. C. (2019). Efficient coding of subjective value. *Nature Neuroscience*, *22*(1). https://doi.org/10.1038/s41593-018-0292-0

Prelec, D. (1998). The probability weighting function. *Econometrica*, 497–527. http://www.jstor.org/stable/10.2307/2998573

Preuschoff, K., Bossaerts, P., & Quartz, S. R. (2006). Neural Differentiation of Expected Reward and Risk in Human Subcortical Structures. *Neuron*, *51*(3), 381–390. https://doi.org/10.1016/j.neuron.2006.06.024

Robson, A. J., Whitehead, L. A., & Robalino, N. (2023). Adaptive utility. *Journal of Economic Behavior and Organization*, *211*. https://doi.org/10.1016/j.jebo.2023.04.023

Schneider, M., & Day, R. (2018). Target-adjusted utility functions and expected-utility paradoxes. *Management Science*, *64*(1). https://doi.org/10.1287/mnsc.2016.2588

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. https://doi.org/10.1523/JNEUROSCI.3653-14.2015

Steiner, J., & Stewart, C. (2016). Perceiving prospects properly. *American Economic Review*, *106*(7). https://doi.org/10.1257/aer.20141141

Steverson, K., Brandenburger, A., & Glimcher, P. (2019). Choice-theoretic foundations of the divisive normalization model. *Journal of Economic Behavior & Organization*, *164*, 148–165.

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, *53*(1). https://doi.org/10.1016/j.cogpsych.2005.10.003

Tobler, P. N., Christopoulos, G. I., O'Doherty, J. P., Dolan, R. J., & Schultz, W. (2008). Neuronal distortions of reward probability without choice. *The Journal of Neuroscience*, *28*(45), 11703–11711. http://www.jneurosci.org/content/28/45/11703.full

Tymula, A., Wang, X., Imaizumi, Y., Kawai, T., Kunimatsu, J., Matsumoto, M., & Yamada, H. (2023). Dynamic prospect theory: Two core decision theories coexist in the gambling behavior of monkeys and humans. *Science Advances*, *9*(20). https://doi.org/10.1126/sciadv.ade7972

Vieider, F. (2025). *Neuro-Biological Origins of PT functionals*.

Wang, J., Xia, Y., Wang, G., Zhou, T., Guo, Y., Zhang, C., An, X., Sun, Y., Guo, X., Zhou, Z., & Sha, J. (2014). In-depth proteomic analysis of whole testis tissue from the adult rhesus macaque. *Proteomics*, *14*(11). https://doi.org/10.1002/pmic.201300149

Webb, R. (2020). The (neural) dynamics of stochastic choice. *Management Science*, *65*(1). https://doi.org/10.1287/mnsc.2017.2931

Webb, R., Glimcher, P. W., & Louie, K. (2021). The Normalization of Consumer Valuations: Context-Dependent Preferences from Neurobiological Constraints. *Management Science*, *67*(1). https://doi.org/10.2139/ssrn.2462895

Woodford, M. (2012). Prospect Theory as efficient perceptual distortion. *American Economic Review*, *102*(3), 41–46. https://doi.org/10.1257/aer.102.3.41

Yamada, H., Louie, K., Tymula, A., & Glimcher, P. W. (2018). Free choice shapes normalized value signals in medial orbitofrontal cortex. *Nature Communications*, *9*(1). https://doi.org/10.1038/s41467-017-02614-w

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 of the United States of America*, *110*(39). https://doi.org/10.1073/pnas.1308718110

Zhang, H., Ren, X., & Maloney, L. T. (2020). The bounded rationality of probability distortion. *Proceedings of the National Academy of Sciences of the United States of America*, *117*(36). https://doi.org/10.1073/pnas.1922401117

Appendix

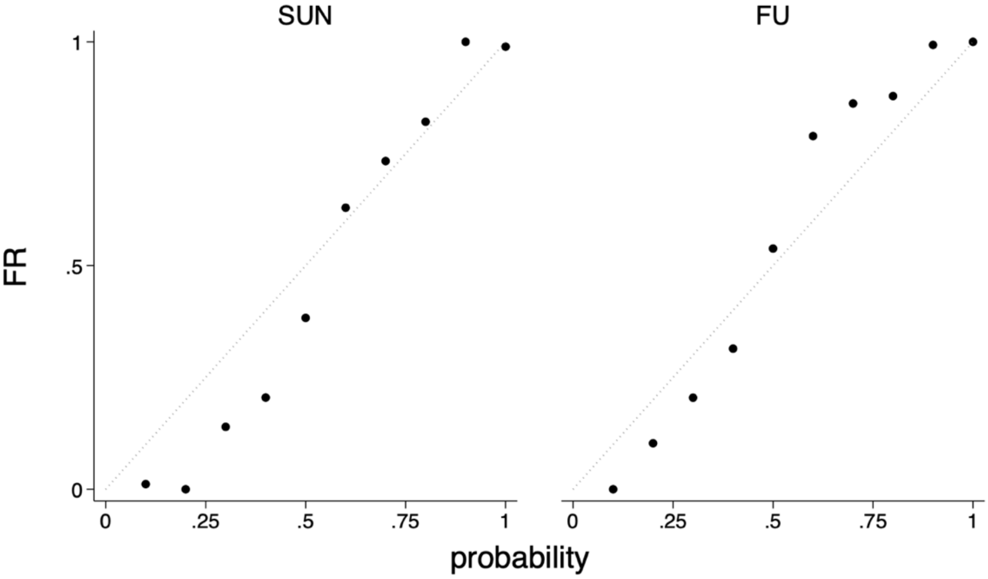

Figure S1. Neural probability weighting function. Includes data from all neurons that significantly and positively (i.e. higher probability higher activity) respond to probability and do not respond to payoff magnitude. Best fit probability weighting parameters are $\delta_{SUN} = 0.634$ with SE=0.059, $\delta_{FU} = 1.219$ with SE=0.106, $\gamma_{SUN} = 1.812$ with SE=0.153, and $\gamma_{FU} = 1.947$ with SE=0.160.

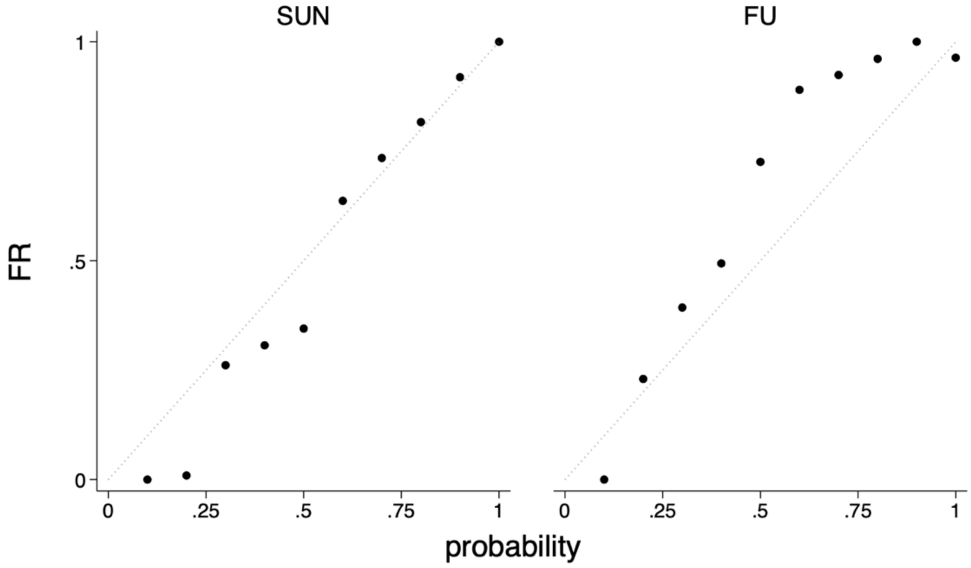

Figure S2. Neural probability weighting function. Includes data from all neurons that significantly respond to probability (positively and negatively) including those that respond to payoff magnitude. Best fit probability weighting parameters are $\delta_{SUN} = 0.747$ with SE=0.089, $\delta_{FU} = 2.627$ with SE=0.306, $\gamma_{SUN} = 1.452$ with SE=0.172, and $\gamma_{FU} = 1.775$ with SE=0.157.

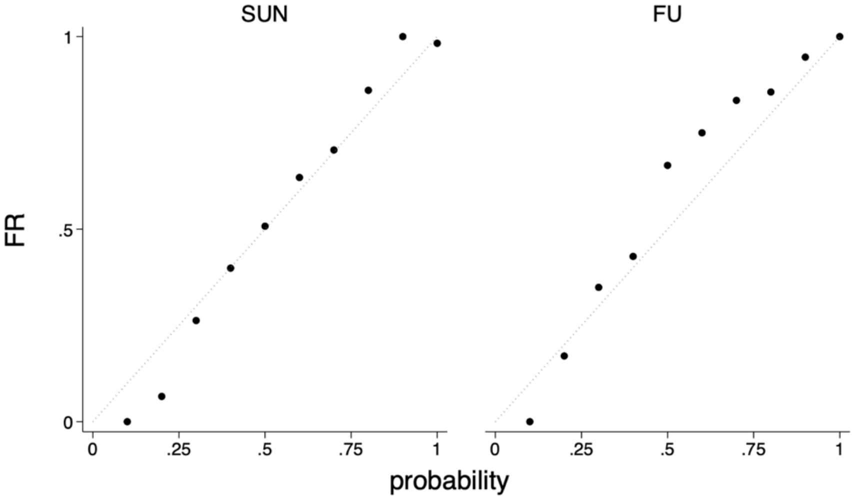

Figure S3. Neural probability weighting function. Includes data from all neurons that significantly respond to probability and do not respond to payoff magnitude. Neurons are included if they positively respond to probability as determined by a linear regression. Best fit probability weighting parameters are $\delta_{SUN} = 0.983$ with SE=0.080, $\delta_{FU} = 1.582$ with SE=0.143, $\gamma_{SUN} = 1.384$ with SE=0.115, and $\gamma_{FU} = 1.437$ with SE=0.124.

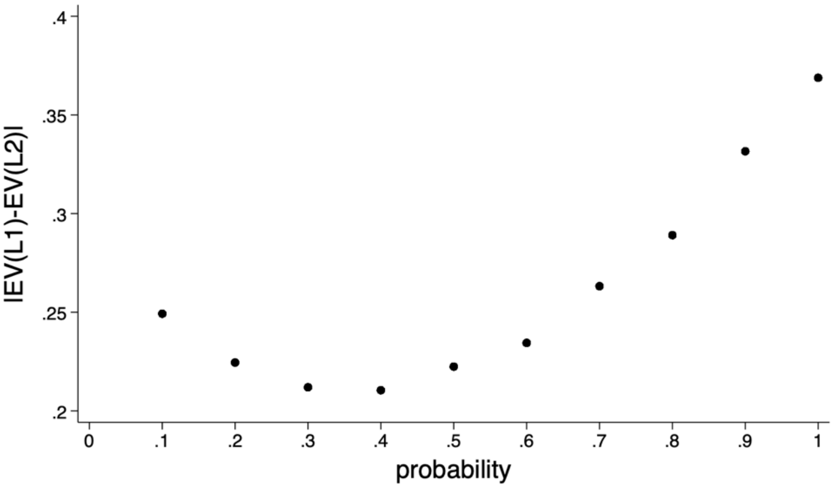

Figure S4. **Complexity in the choice task by probability level.** Dots illustrate the absolute value of the difference in expected values of the lotteries on offer, averaged across all trials for both participants.