NBER WORKING PAPER SERIES

ON THE SOURCE AND INSTABILITY OF PROBABILITY WEIGHTING

Cary Frydman
Lawrence J. Jin

Working Paper 31573
http://www.nber.org/papers/w31573

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2023, revised April 2025

We are grateful for helpful comments from Nicholas Barberis, Benjamin Enke, Thomas Graeber, Ted O'Donoghue, Ryan Oprea, Pietro Ortoleva, Collin Raymond, Michael Woodford, George Wu, and audience members at Cornell University, Dartmouth College, New York University, the University of California Davis, the University of California Los Angeles, the University of California Santa Barbara, the University of Chicago, the University of Colorado Boulder, the University of Miami, the University of Oxford, the University of Pennsylvania, the University of Warwick, the University of Zurich, the 2025 AFA meetings, the 2024 BDRM conference, the 2024 Los Angeles Experiments Workshop at Caltech, the 2023 JDM Winter Symposium, the 2023 Behavioral Economics Annual Meeting, the 2023 NBER Behavioral Finance Spring Meeting, the 2023 Columbia University Workshop on Cognitive Noise and Economic Decisions, and the 2023 Workshop on Cognitive Economics at Ghent University. Frydman acknowledges financial support from the NSF and USC Marshall School of Business iORB fund. Jin acknowledges financial support from the Cornell SC Johnson College of Business research fund. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

On the Source and Instability of Probability Weighting
Cary Frydman and Lawrence J. Jin
NBER Working Paper No. 31573
August 2023, revised April 2025
JEL No. D03, G02, G41

ABSTRACT

We propose and experimentally test a new theory of probability distortions in risky choice. The theory is based on a core principle from neuroscience called efficient coding, which states that information is encoded more accurately for those stimuli that the agent expects to encounter more frequently. As the agent's prior beliefs vary, the model predicts that probability distortions change systematically. We provide novel experimental evidence consistent with this prediction by manipulating and measuring a subject's prior beliefs. The data reveal that lottery valuations are more sensitive to probabilities that occur more frequently under the subject's reported prior beliefs.

Cary Frydman
University of Southern California
Marshall School of Business
3670 Trousdale Parkway, BRI 204G
Los Angeles, CA 90089
cfrydman@marshall.usc.edu

Lawrence J. Jin
Johnson College of Business
Cornell University
Ithaca, NY 14850
and NBER
lawrence.jin@cornell.edu

A data appendix is available at http://www.nber.org/data-appendix/w31573

I. Introduction

Consider the prospect of winning \$25 with 5% probability and winning \$0 otherwise. Under expected utility theory, valuation of the lottery should be linear in the probability of winning the \$25 outcome. For decades, however, experimental evidence points to decision-making that is non-linear in probability. For example, subjects report a larger increase in lottery valuation as the experimenter increases the probability of winning the \$25 from 5% to 10%, compared to when she increases it from 30% to 35% (Tversky and Kahneman, 1992; Camerer and Ho, 1994; Gonzalez and Wu, 1999; Bernheim and Sprenger, 2020). The nonlinearity gives rise to probability distortions, which in turn, have been used to explain fundamental puzzles in risky choice such as the simultaneous demand for gambling and insurance. Probability distortions have also been invoked to explain a wide array of anomalies in financial markets (Barberis, 2018).

In this paper, we ask a basic question: *why* do humans distort probabilities? We explore whether the distortion can be traced to fundamental properties about how the brain processes information. Our hypothesis is that valuation is linear in *perceived* probability, which differs from objective probability due to information processing constraints. We are by no means the first to investigate the psychological underpinning of probability weighting (Gonzalez and Wu, 1999; Bordalo, Gennaioli, and Shleifer, 2012; Khaw, Li, and Woodford, 2021, 2024; Enke and Graeber, 2023).[1] But the mechanism we test—which is called efficient coding—is novel in the sense that it generates new testable hypotheses about the source and instability of probability weighting. To assess the empirical validity of our proposed mechanism, we develop a theory and present an experimental test. The data largely confirm the idea that probability weighting derives, at least in part, from the efficient allocation of cognitive resources.

The basic premise of our theory is that, when faced with a risky prospect, the decision maker (henceforth DM) bases her decision on a noisy perception of state probabilities; in general, the DM's perception does not coincide with the objective probability. This wedge between reality and perception stems from errors in encoding the state probability. That is, we assume that the DM does not have the capacity to form a precise perception of each probability that she could potentially

[1] See also important earlier work by Viscusi (1989), Tversky and Kahneman (1992), Rottenstreich and Hsee (2001), Stewart, Chater, and Brown (2006), Zhang and Maloney (2012), Steiner and Stewart (2016), and Zhang, Ren, and Maloney (2020).

be presented with. While the DM cannot form a precise perception of every probability, she can efficiently allocate cognitive resources towards accurately perceiving those probabilities that she *expects* to encounter in a given class of situations.

For example, consider someone who is often faced with binary lotteries that have probabilities that fall in the range from say, 45% to 55%. We argue that, when faced with two novel lotteries, one with 50-50 odds and another with 51-49 odds, this person will be better able to discriminate between the two lotteries, compared to someone who is often faced with skewed lotteries that have extreme odds such as 90-10 or 95-5. Our model formalizes how the DM's prior belief about the probabilities she expects to encounter will optimally distort her perception of the probabilities that she actually encounters.

Our theoretical contribution builds on our own previous work in which we demonstrated that risk taking depends on the prior distribution of monetary outcomes to which the DM is adapted (Frydman and Jin, 2022). In that model, we made the simplifying assumption that probabilities are encoded without noise. Here, we show how noisy and efficient coding can generate probability distortions. In this regard, our work is closely related to recent theory by Steiner and Stewart (2016), Zhang et al. (2020), and Khaw et al. (2021, 2024). However, an important difference is that, in our model, we show how the shape of the prior distribution over probabilities pins down the degree of probability distortion.

The model works as follows. The DM encodes a single probability p subject to a capacity constraint. We assume her objective is to choose an encoding function that maximizes the mutual information between the true probability p and the encoded signal R_p. The model predicts that the DM will encode with more precision those probabilities that she expects to encounter more often according to her prior belief. Thus, a key input to the model is the DM's prior belief about probability p.

Given the important role of the prior, our experimental test of the model involves both manipulating and directly measuring the prior. In particular, we exogenously manipulate a subject's prior belief about probabilities and then elicit certainty equivalents for lotteries of the form: $(\$25, p; \$0, 1-p)$. Efficient coding predicts that, when the DM is adapted to a distribution of probabilities that is concentrated near intermediate probabilities, she should have less difficulty discriminating between probabilities over this range compared to someone who is adapted to extreme probabilities. Crucially, because the DM does not have access to the objective probability p, she must rely only on her perceived value of p when forming a lottery valuation. These observations lead to a testable prediction: given a risky lottery, valuation depends systematically on the prior distribution to which a subject is adapted.

We find strong evidence in favor of this prediction. The slope of the measured probability weighting function over intermediate probabilities is significantly higher for subjects who are adapted to intermediate probabilities, compared to those who are adapted to extreme probabilities. Specifically, for subjects who are adapted to extreme probabilities, increasing an intermediate probability of winning \$25 by 1% causes only a \$0.18 increase in their certainty equivalent. In contrast, for subjects who are adapted to intermediate probabilities, the same 1% increase of an intermediate probability of winning \$25 leads to a \$0.30 increase in their certainty equivalent. Thus, the causal effect of a 1% increase in probability on valuation is about 66% larger for subjects who are adapted to intermediate probabilities. We find that the effect of the prior on valuation carries a similar magnitude for lotteries associated with small probabilities of winning \$25. Finally, the effect is of a smaller magnitude for lotteries associated with high probabilities of winning \$25.

A novel feature of our experimental design is that we also directly measure prior beliefs after adaptation. To our knowledge, all previous tests of efficient coding have treated the prior as unobservable and have assumed that it coincides with the empirical distribution of the relevant environment. By directly measuring the prior, we can test whether subjects have adapted to the distribution that is assumed in the theory. Furthermore, even if subjects have not perfectly adapted to the environment, we can still use their measured priors to test the model. After all, the model predicts that it is the subjective prior that distorts perception—regardless of whether that prior matches the statistical properties of a given adaptation condition.

We find that prior beliefs do not differ across conditions at the beginning of the experiment—the average subject in each condition reports a prior belief that puts more weight on intermediate probabilities compared to extreme probabilities. Yet, after only fifteen trials, we find a strong separation in beliefs across conditions, in a manner that reflects the distribution of probabilities that subjects have recently encountered. In other words, a subject's forward-looking beliefs about the next lottery's state probabilities tend to closely match the distribution of probabilities that she has recently encountered.

Importantly, we find there is also substantial heterogeneity in prior beliefs, even within an adaptation condition. According to our model, this heterogeneity in prior beliefs should lead to systematic differences in perception of probabilities and hence valuations. To test this prediction, we examine each of three ranges of probabilities: low, intermediate and high. For the low range probabilities, we compute the density that a subject attaches to this range under their elicited prior. We find that subjects who attach a higher weight to low probabilities do indeed exhibit valuations that are more sensitive to a change in probability over this range. We find similar results for the intermediate and high ranges of probabilities: subjects who state a higher likelihood of observing a given range of probabilities also report lottery valuations that are more sensitive to a change in probability over this range. These results indicate that measured priors about a given lottery are predictive of the valuation of the lottery. To be clear, even after the lottery is presented to the subject, beliefs measured 15 trials in the past are still predictive of valuations, conditional on the lottery's state probabilities.

Overall, our results suggest that probability weighting stems in part from people's imprecision in forming a cognitive representation of probability. Our key experimental finding is that valuation depends not only on objective properties of a given lottery, but also on the ex-ante beliefs about the lottery probabilities. This finding is inconsistent with a broad class of models that assume a nonlinear but stable probability weighting function (Quiggin, 1982; Yaari, 1987; Tversky and Kahneman, 1992). The fact that valuation depends systematically on prior beliefs also separates our efficient coding model from other candidate explanations for probability weighting that may also be active in the decision process. For example, the salience theory of Bordalo et al. (2012) proposes that probability weighting arises from limited attention, where distortions depend on the salience of lottery payoffs.[2] Enke and Graeber (2023) propose that agents are cognitively uncertain about valuations, which leads to a compressed linear probability weighting function. Although salience theory and cognitive uncertainty both generate context-dependent weighting functions, neither can deliver the systematic dependence on prior beliefs that we observe in our data.[3]

[2] In our baseline model, probability weighting does not depend on lottery payoffs. Khaw et al. (2024) propose a model in which the imprecision in representing a monetary amount is tied to the probability of winning that amount. In their model, fewer resources are devoted to processing payoffs that are less likely to be delivered. In Section IV.1, we sketch an extension of our model in which we investigate the implications of efficient coding over both payoffs and probabilities.

[3] Steiner and Stewart (2016) provide an alternative explanation of probability weighting that is also based on noisy perception of probability. In particular, they show that a nonlinear probability weighting function can arise as

Our model also shares similarities with the decision by sampling model proposed by Stewart et al. (2006). In that model, when the DM is presented with a probability, she draws samples of previously encountered probabilities from memory and compares the currently presented probability with these randomly drawn samples. The similarities between the two models are natural given the insight from Bhui and Gershman (2018) that efficient coding can serve as an optimizing foundation for decision by sampling. As such, we interpret our experimental results as providing novel evidence for both efficient coding and decision by sampling.

The rest of this paper is organized as follows. Section II presents a theory of efficient coding and shows theoretically that the slope of different portions of the weighting function changes as the DM's prior changes. Section III presents an experiment which uses data on lottery valuations to test whether the weighting function is malleable in the manner predicted by efficient coding. Section IV provides additional discussion, and Section V concludes.

II. The Model: Instability of Probability Weighting

In this section, we apply the efficient coding model of Heng, Woodford, and Polanía (2020) to study the DM's mental representation of probability. The model works as follows. Consider a probability p that is associated with a given lottery payoff. Before the DM is presented with the probability p, she holds a prior belief about it, denoted by $f(p)$. Then, upon presentation of p, the DM generates a noisy cognitive signal R_p; specifically, R_p is randomly drawn from a conditional distribution—or, in the language of Bayesian inference, a likelihood function—denoted by $f(R_p|p)$. The likelihood function captures the idea that the DM encodes information about probability with cognitive noise, even if the probability p is clearly presented to the DM.

Following Heng et al. (2020), we assume that the DM encodes probability p through a finite number of n "neurons," where the output state of each neuron takes the value of 0 or 1. The output states of these n neurons are assumed to be mutually independent, and each neuron takes the value 1 with probability $\theta(p)$ and 0 with the remaining probability $1 - \theta(p)$.[4] The encoded value of p is

an optimal response to perceptual noise. Khaw et al. (2021, 2024) also propose a model that derives a probability weighting function based on an optimal response to noisy coding of probabilities. One difference is that the encoding function in our model also arises from an optimization process; Section IV.3.3 shows that this is crucial in generating the model's key implications. See also McGranaghan, Nielsen, O'Donoghue, Somerville, and Sprenger (2024) for recent experimental work that motivates the development of new theories of context-dependent weighting functions.

[4] We interpret the neurons in the model as basic information processing units that emit a binary signal, rather

therefore represented by an output vector of $0s$ and $1s$, with length n. Given that the neurons are mutually independent, a sufficient statistic for the output vector is the sum across the n output values, which is defined as the noisy signal R_p. Taken together, the likelihood function is given by

$$f(R_p|p) = \binom{n}{R_p}(\theta(p))^{R_p}(1-\theta(p))^{n-R_p}, \tag{1}$$

where the noisy signal R_p can take on integer values from 0 to n.

A signature feature of efficient coding is that the DM endogenously chooses the likelihood function $f(R_p|p)$ as a function of the prior $f(p)$. We assume that the DM chooses $\theta(p)$ in (1) to maximize the mutual information between probability p and its noisy signal R_p

$$\max_{\theta(p)} I(p; R_p), \tag{2}$$

where the mutual information $I(p; R_p)$ is defined as the difference between the marginal entropy of R_p and the entropy of R_p conditional on p.[5] Intuitively, the objective function in (2) leads the DM to better discriminate between probability values that she expects to encounter more frequently given her prior belief. A large literature in sensory perception documents strong support for this objective function (Laughlin, 1981; Girshick, Landy, and Simoncelli, 2011; Wei and Stocker, 2015). Heng et al. (2020) show that the optimal coding rule that maximizes $I(p; R_p)$ is given by

$$\theta(p) = \left(\sin\left(\frac{\pi}{2}F(p)\right)\right)^2, \tag{3}$$

where $F(p)$ is the cumulative distribution function of the prior belief $f(p)$.

Given the prior belief and the noisy signal, the DM follows Bayes' rule to generate a posterior belief about p

$$f(p|R_p) = \frac{f(R_p|p)f(p)}{\int_0^1 f(R_p|p)f(p)dp}. \tag{4}$$

For a given objective probability value p, the DM draws the noisy signal R_p according to the likelihood function $f(R_p|p)$ from equation (1). Then, for each noisy signal R_p, the DM forms a

than literally as cells in the brain that have more complex properties.

[5] We discuss an alternative performance objective, namely maximizing expected payoff, in Section IV.3.1.

posterior distribution about p given by equation (4). Together, these two steps imply that the average subjective valuation of p is

$$v(p) = \sum\nolimits_{R_p=0}^{n} f(R_p|p) \cdot \mathbb{E}[\tilde{p}|R_p], \tag{5}$$

where

$$\mathbb{E}[\tilde{p}|R_p] \equiv \int_0^1 f(p|R_p)pdp = \frac{\int_0^1 f(R_p|p)f(p)pdp}{\int_0^1 f(R_p|p)f(p)dp} \tag{6}$$

is the posterior mean of p conditional on R_p. Note that $v(p)$ in (5) represents the DM's subjective valuation of p averaged across different values of R_p, and in general, $v(p)$ is a nonlinear function that maps objective probabilities into distorted, decision-relevant probabilities.[6]

It is important to note that, in an efficient coding model, the DM's subjective valuation of probability $v(p)$ depends on her prior belief $f(p)$ through *two* channels. First, the posterior belief $f(p|R_p)$ is directly affected by the prior belief $f(p)$, as shown in equation (4). Second and more importantly, $f(p)$ indirectly affects the posterior belief $f(p|R_p)$ through the coding rule $\theta(p)$. In particular, equation (3) shows that the coding rule takes the prior as an input, so that the same value of p is encoded differently depending on the prior. Thus, the prior continues to play its traditional and direct role in Bayesian inference, but it also affects information processing through the endogenously determined likelihood function given by equations (1) and (3). Taken together, the efficient coding model described by equations (1) to (6) shows that the DM's prior belief $f(p)$ affects her perception of probability and hence her subjective valuation $v(p)$. In other words, if the DM's prior belief $f(p)$ changes, $v(p)$ will change accordingly.

Below, we formally show how efficient coding generates a probability weighting function that is *malleable*. We begin by assuming that the DM's prior belief takes the following form

$$f(p) = \xi \cdot \underbrace{f_1(p)}_{\text{a stable component}} + (1-\xi) \cdot \underbrace{f_2(p)}_{\text{a fast-moving component}}. \tag{7}$$

We now refer to $f(p)$ as a "mixed" prior, which captures the principle that adaptation can take

[6]For the rest of the paper, we focus on the implications of our model for the case when $0 < p < 1$. We remain agnostic about whether the same noisy coding machinery operates for the cases of impossibility ($p = 0$) and certainty ($p = 1$).

place at multiple timescales (Wark, Fairhall, and Rieke, 2009; Weber, Krishnamurthy, and Fairhall, 2019). The mixed prior contains two components. The first component is stable and denoted as $f_1(p)$. This component represents the prior that the DM has formed based on encountering different probability values with different frequencies over a long timescale. The second, fast-moving component is denoted as $f_2(p)$; it captures transient changes in statistics of the local environment. For example, $f_2(p)$ can correspond to the empirical distribution of probabilities that we present in the lab experiment in Section III.[7] The parameter ξ in equation (7) represents the weight the DM puts on the stable component; this weight may depend on the rate at which the statistics of the environment change.

For the stable component $f_1(p)$, we assume, for simplicity, that it takes the form of a uniform distribution between 0 and 1. The model's main implications for the malleability of the probability weighting function are robust to alternative specifications of $f_1(p)$. For the fast-moving component $f_2(p)$, our main goal is to demonstrate how shocks to this component drive differences in valuation. As such, we consider two specifications of the fast-moving component that have very different shapes. We investigate an "intermediate" prior component, which has density only in the center of the unit interval. Formally, for this "intermediate" prior component, we assume

$$f_2(p) = \begin{cases} \dfrac{1}{p_h - p_l} & p_l \leq p \leq p_h \\ 0 & \text{otherwise} \end{cases}, \tag{8}$$

where p_l and p_h take intermediate values and $0 < p_l < p_h < 1$.

We also investigate an "extreme" prior component, which has density only near the endpoints of the unit interval. For this "extreme" prior component, we assume

$$f_2(p) = \begin{cases} \dfrac{1}{(p_{h,2} - p_{h,1}) + (p_{l,2} - p_{l,1})} & p_{l,1} \leq p \leq p_{l,2} \\ \dfrac{1}{(p_{h,2} - p_{h,1}) + (p_{l,2} - p_{l,1})} & p_{h,1} \leq p \leq p_{h,2} \\ 0 & \text{otherwise} \end{cases}, \tag{9}$$

[7] A growing literature has argued that people adapt, at least in part, towards the local context that changes over the course of an experimental session (Burke, Baddeley, Tobler, and Schultz, 2016; Zimmermann, Glimcher, and Louie, 2018; Conen and Padoa-Schioppa, 2019; Payzan-LeNestour and Woodford, 2022; Frydman and Jin, 2022).

where $p_{l,1}$ and $p_{l,2}$ take low values, $p_{h,1}$ and $p_{h,2}$ take high values, and $0 \leq p_{l,1} < p_{l,2} < 0.5 < p_{h,1} < p_{h,2} \leq 1$.

The two different specifications of the fast-moving component give rise to two different mixed priors, which we refer to as the intermediate prior and the extreme prior. We specify the parameter values for these two priors as follows. As stated before, the stable component $f_1(p)$ of each mixed prior is assumed to be a uniform distribution between 0 and 1. For the fast-moving intermediate prior component $f_2(p)$, we set $p_l = 0.38$ and $p_h = 0.62$. For the fast-moving extreme prior component, we set $p_{l,1} = 0.1$, $p_{l,2} = 0.21$, $p_{h,1} = 0.79$, and $p_{h,2} = 0.9$. Finally, we set ξ, the weight on the stable component in the mixed prior, to 0.5.

Armed with two specific mixed priors, we now use them as inputs to our efficient coding model and assess the differences in perception of probability. For each mixed prior $f(p)$, we substitute it into the coding rule in equation (3). The two different coding rules—one for each specification of the mixed prior—give rise to a distinct set of efficient likelihood functions. Then, each prior, together with its corresponding likelihood functions, generates the probability weighting function $v(p)$. Figure I summarizes these results: it plots, for both the intermediate prior and the extreme prior, the prior distribution $f(p)$, its coding rule $\theta(p)$, and the implied weighting function $v(p)$.

Figure I demonstrates the *malleability* of probability weighting: as the DM's prior beliefs change (the upper graph), the coding function $\theta(p)$ (the middle graph) and the implied probability weighting function $v(p)$ (the lower graph) change significantly. Importantly, the way in which the weighting function changes is governed by efficient coding. For intermediate probabilities, the slope of the weighting function is steeper when the fast-moving prior component increases the density for intermediate probabilities; however, for extreme probabilities, the slope of the weighting function is steeper when the fast-moving prior component puts extra mass on the extreme portions of the unit interval, near 0 and 1.

To understand these results, we first note that, by construction, intermediate probabilities occur much more frequently under the intermediate prior, compared to the extreme prior. Thus, under the intermediate prior, the coding rule $\theta(p)$ has a much steeper slope for intermediate probabilities, causing the likelihood function $f(R_p|p)$ to shift substantially as p varies over the intermediate range. This greater separation of likelihood functions for nearby probabilities gives rise to the higher slope of the probability weighting function under the intermediate prior.

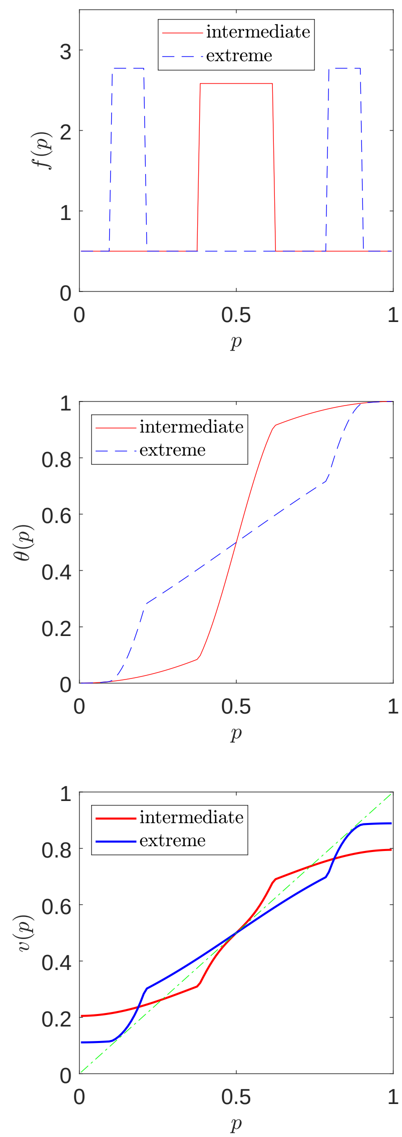

Figure I
Prior distribution, coding function, and value function: Efficient coding and mixed prior

The upper graph plots two mixed prior distributions in the form of (7). For each mixed prior, the first, stable component takes the form of a uniform distribution between 0 and 1. The second, fast-moving component takes the form of (8) for the intermediate prior and the form of (9) for the extreme prior; the parameter values are: $p_l = 0.38$, $p_h = 0.62$, $p_{l,1} = 0.1$, $p_{l,2} = 0.21$, $p_{h,1} = 0.79$, and $p_{h,2} = 0.9$. The weight ξ the DM assigns to the stable component is 0.5. The middle graph plots, for both the intermediate prior and the extreme prior, the corresponding coding rule $\theta(p)$ in the form of (3). The lower graph plots, for both priors, the subjective valuation implied by efficient coding, $v(p)$. We set the parameter n to 10. The green dash-dot line is the forty-five degree line.

The same intuition holds when examining perception of probabilities closer towards the boundaries of the unit interval. By construction, low or high probabilities occur more frequently under the extreme prior, compared to the intermediate prior. Therefore, as p varies over the low or high range of probability, the likelihood function shifts to a greater extent under the extreme prior, generating a steeper slope of the probability weighting function.[8]

In the next section, we design an experiment that exogenously manipulates the fast-moving component of a subject's prior. The particular specifications of the fast-moving component follow closely from the intermediate and extreme priors we have theoretically considered above. Using data on lottery valuations and prior beliefs, the experiment will allow us to test whether the slope of different portions of the probability weighting function is malleable in the manner predicted by our theory of efficient coding.

III. Experimental Manipulation of Probability Distortions

III.1. Design

III.1.1. Manipulation of prior

Our experimental design is intended to test whether subjects' prior beliefs about probability causally affect their valuations of risky lotteries. For simplicity, here we assume that subjects in our experiment only have noise when encoding probability; their encoding of lottery payoffs is noiseless. To justify this assumption, we consider only binary lotteries with two possible payoffs, and crucially, the two possible payoffs are kept constant across all trials. As such, any noise arising during the encoding of lottery payoffs should be minimal. In Section IV.1, we discuss the more general case in which subjects encode both probability and lottery payoffs with noise.

We now turn to the details of our design. On each trial, a subject is presented with a risky lottery of the following form

$$(\$25, p; \$0, 1 - p). \tag{10}$$

The subject is then asked to provide her certainty equivalent for the lottery by using a slider

[8] The malleability of probability weighting presented in Figure I uses the coding rule of (3) and sets the capacity constraint parameter n to 10. We note that Heng et al. (2020) derive the coding rule of (3) when n is asymptotically large. However, the consequence of this apparent inconsistency is small: Appendix 7 of Heng et al. (2020) shows that, with any finite n that is greater than or equal to 5, the coding rule of (3) remains approximately optimal.

bar. Importantly, information about the probability p is displayed to the subject through both a numerical and a graphical representation; see Online Appendix D Figure D.1 for a screenshot of an example trial. We choose to use the slider method rather than the commonly used multiple price list, as it allows us to elicit an exact valuation instead of a switching interval. Khaw et al. (2024) also adopt the slider method for similar reasons.

Each subject in the experiment completes a total of thirty-seven trials, and thus submits a total of thirty-seven certainty equivalents for various values of probability p. Subjects are incentivized to report their certainty equivalents using a Becker-Degroot-Marschak mechanism (Becker, DeGroot, and Marschak, 1964). We define the first thirty trials as "adaptation trials," and the subsequent six trials are "test trials."[9] The core of our design is to exogenously manipulate the distribution of p on the adaptation trials and test the impact on valuation in the test trials. The distribution of p on adaptation trials plays the role of the fast-moving component of the subject's mixed prior, as we previously described in Section II. Importantly, for different sets of adaptation trials, we hold constant the set of test trials. In this way, the only feature of the experiment that varies is the distribution of p that subjects experience before they provide their valuations on the six test trials.

When constructing our design, we face a tradeoff in choosing the number of test trials per subject. On the one hand, more test trials provide us with a greater number of data points on which to estimate the treatment effect of the prior. On the other hand, the more test trials we use, the weaker the treatment effect will be for the final few test trials. This is because early test trials will begin to contaminate the prior that we attempt to induce with the adaptation trials. Thus, while we would ideally like to construct a large set of test trials that sample p across the entire unit interval, we are constrained by the number of test trials that we can use *per subject*.

Because this design constraint operates only at the subject level, we choose to construct three separate test trial ranges, and randomly assign subjects to one of the three ranges. In this manner, we can span a large range of the unit interval while minimizing the number of test trials needed per subject. Specifically, we define a *low* range for which $p \in \{0.11, 0.15, 0.19, 0.23, 0.27, 0.31\}$, an *intermediate* range for which $p \in \{0.38, 0.42, 0.47, 0.53, 0.58, 0.62\}$, and a *high* range for which $p \in \{0.69, 0.73, 0.77, 0.81, 0.85, 0.89\}$.

[9]The final and thirty-seventh trial serves as an attention check. Footnote 12 in Section III.3 describes it in more detail.

After subjects are randomized into one of the three test trial ranges, we further randomize them into an adaptation condition. Each test trial range is associated with one of two possible adaptation conditions. Table D.1 in Online Appendix D provides the exact values of p that we use in each adaptation condition.[10] For the *low* test trial range, we adapt subjects to either a *low* or *intermediate* range of adaptation trials. The intuition is that, for a subject who is adapted to low values of p, she should be able to easily discriminate between low values of probability once the low range test trials arrive. In contrast, for a subject who is adapted to intermediate values of p, she should be less able to discriminate between low values of probability. Thus, conditional on being randomized into the *low* test trial range, our model predicts that valuation will respond more strongly to changes in p among subjects in the *low* adaptation condition, compared to those in the *intermediate* adaptation condition.

A similar intuition applies for the other two test trial ranges. For those subjects randomized into the *high* test trial range, we adapt them to either a *high* or *intermediate* range of adaptation trials. For a subject who is adapted to high values of p, she should be better able to discriminate between probabilities in this high range once the test trials arrive, compared to a subject who is adapted to intermediate values of p.

Finally, for those subjects randomized into the *intermediate* test trial range, we adapt them to either an *extreme* or *intermediate* range of adaptation trials. A subject who is adapted to extreme values of p—closer to either of the two endpoints of the unit interval—should be less able to discriminate between probabilities in the intermediate range once the test trials arrive, compared to a subject who is adapted to intermediate values of p.

III.1.2. Directly measuring prior beliefs

A novel feature of our design, relative to previous tests of efficient coding, is that we directly measure prior beliefs. After all, the key input to the efficient coding model described in Section II is the DM's prior beliefs, which pin down the likelihood functions and shape of the probability weighting function. Measuring prior beliefs enables us to test whether endogenously reported beliefs

[10] There are 24 distinct values of p in each adaptation condition. For each subject, we draw 30 values with replacement from this set of 24, and we randomize the ordering. Thus, even within an adaptation condition, there exists exogenous variation across subjects in the sequence of lotteries experienced, which could in turn, provide a source of variation in prior beliefs at the test trial stage.

about an upcoming choice set have predictive power for valuations. Moreover, the reported beliefs can validate whether subjects have adapted to the particular distribution that we assume in our key empirical analyses.

You are approximately halfway through the survey. We would like to get your sense of which gamble you think you'll see on the next round. As in all previous questions, in the next question you will see a gamble which pays $25 with some chance, which we will denote as **C**. The value of **C** will be somewhere between 1% and 99%. We are interested in what you think the probability of observing **C** is, for 10 different ranges.

The first row below describes the range of **C** between 1% - 10%. Please enter the probability that you think **C** will be drawn from that range. For the second row, please enter the probability that you think the value of **C** will be in the range 11% - 20%. The same goes for the remaining rows, and the total must equal 100%.

Please think carefully about each of the 10 rows and then begin to enter your responses.

1% - 10%	0	%
11% - 20%	0	%
21% - 30%	0	%
31% - 40%	0	%
41% - 50%	0	%
51% - 60%	0	%
61% - 70%	0	%
71% - 80%	0	%
81% - 90%	0	%
91% - 99%	0	%
Total	0	%

Figure II
Procedure for eliciting prior beliefs

The figure presents a screenshot of the belief elicitation procedure that we implement after trial fifteen. For each of the ten bins, subjects are asked to report their belief that the upside probability on the next trial will fall in that particular bin. We require that the percentages sum up to 100% before the subject can submit their response.

We elicit prior beliefs twice during the experiment: immediately before the first adaptation trial and immediately after trial fifteen. Specifically, we measure beliefs by asking subjects for the probability they assign to each of ten equally spaced bins over the unit interval. The probability in each bin is therefore the perceived chance that the probability associated with the upside payoff of the risky lottery on the next trial will fall in that particular bin. Figure II shows a screenshot of belief elicitation. We randomize across subjects whether the order of bins is increasing or decreasing.

Asking subjects for a probability distribution over probabilities is a conceptually challenging task. Thus, belief elicitation before the first adaptation trial is primarily to familiarize subjects with the elicitation technique that they will engage with again immediately after trial fifteen. Moreover, the belief measurement is unincentivized. We chose not to incentivize beliefs because, as assumed in our model, reported beliefs likely contain a component of long-term prior beliefs that are difficult to incentivize.

III.2. Procedures

We recruit 750 subjects from Prolific. Subjects are paid \$3.00 for completing the experiment. In addition, 10% of the subjects are randomly selected to receive a bonus whose amount is based on their decision in one randomly selected trial. Specifically, the bonus amount is the outcome of a Becker-Degroot-Marschak mechanism: on the randomly selected trial, we draw a monetary amount randomly from a uniform distribution between $[0, 25]$. If the subject's certainty equivalent is below or equal to this monetary amount, the subject receives it. If the subject's certainty equivalent is above the monetary amount, the computer plays the lottery and pays the subject either \$25 or \$0 with the trial-specific probabilities.

At the beginning of the experiment, subjects are presented with instructions and then need to pass a two-question comprehension check; the experimental instructions and the comprehension check are provided in Online Appendix E. On average, subjects completed the experiment in about 15 minutes. Conditional on being chosen to receive a bonus, they received an average total earning of \$17.33. The sample size, main analyses, and data exclusion criteria are pre-registered on Aspredicted.org.[11]

[11] The pre-registration document for this experiment is at: https://aspredicted.org/5rwf-bd85.pdf.

III.3. Results

Following the pre-registration, we apply three exclusion criteria before analyzing the data. First, we exclude those subjects who violated a basic monotonicity property in their responses. For each subject, we estimate a linear regression of the certainty equivalent on p based on the first thirty adaptation trials. We then record the regression coefficient for each subject, and drop all subjects for whom the coefficient is negative; this restriction excludes 83 of the 750 subjects. Note that we estimate the regression for each subject only on the first thirty trials, so that we do not select on certainty equivalents on test trials, the dependent variable of interest. Next, of the remaining 667 subjects, we exclude 135 additional subjects based on an attention check administered after the final test trial.[12] Finally, we drop all test trial observations for which the subject reported their certainty equivalent in less than one second; this restriction further excludes 4 observations at the subject-trial level. Taken together, we are left with a data set that contains 532 subjects and 19,670 trials, of which 3,188 are test trials.

III.3.1. Manipulation of probability weighting function

We now turn to testing our main hypothesis for each of the three different test trial ranges. In order to arrive at an estimate of the subject's perceived value of p, here we assume that the subject has a linear utility function. This implies that the subject's perception of p can be estimated simply as $\frac{CE}{25}$, where CE represents the certainty equivalent that the subject provides.[13]

Figure III shows a graphical depiction of the causal effect of the prior on perception of probability. The figure has three panels; each panel plots the perceived probability against the objective

[12] After the last test trial, namely trial thirty-six, we inserted the final and thirty-seventh trial that acts as an attention check. The lottery on this final trial was designed in a way that it should appear as an "outlier" relative to the first thirty-six trials they've encountered. As such, subjects should be surprised and hence slow to respond. Specifically, the lottery on the final trial has an upside probability that is far from any upside probability in the first thirty-six trials: subjects in the *low* adaptation condition see a lottery that has an upside probability of 94%; subjects in the *high* adaptation condition see a lottery that has an upside probability of 6%; and subjects in the *intermediate* adaptation condition see a lottery that has an upside probability of either 6% or 94%—each with equal chance. In each of the three cases, the upside probability on the final trial is extreme and surprising relative to the upside probabilities experienced by subjects in the first thirty-six trials. If subjects were paying sufficient attention, we reasoned that the response time on this thirty-seventh trial should be longer than that on the previous, thirty-sixth trial. Thus, our pre-registration excludes subjects if their response time on trial thirty-seven is shorter than that on trial thirty-six. Our main experimental tests are robust to this exclusion criterion.

[13] Later in this section, we consider a more general case that allows the subject to have intrinsic risk aversion over the lottery payoff; for a payoff of X, the utility is $u(X) = X^{\alpha}$, where $\alpha \leq 1$. We find that, under nonlinear utility specifications, our experimental data continue to support the theoretical predictions from Section II and Figure I.

probability for two sets of subjects. The left panel involves two sets of subjects who are both randomized into the low test trials; one set is adapted to the *low* adaptation trials, while the other set is adapted to the *intermediate* adaptation trials. The right panel involves two sets of subjects who are both randomized into the high test trials; one set is adapted to the *high* adaptation trials, while the other set is adapted to the *intermediate* adaptation trials. Finally, the middle panel involves two sets of subjects who are both randomized into the intermediate test trials; one set is adapted to the *extreme* adaptation trials, while the other set is adapted to the *intermediate* adaptation trials.

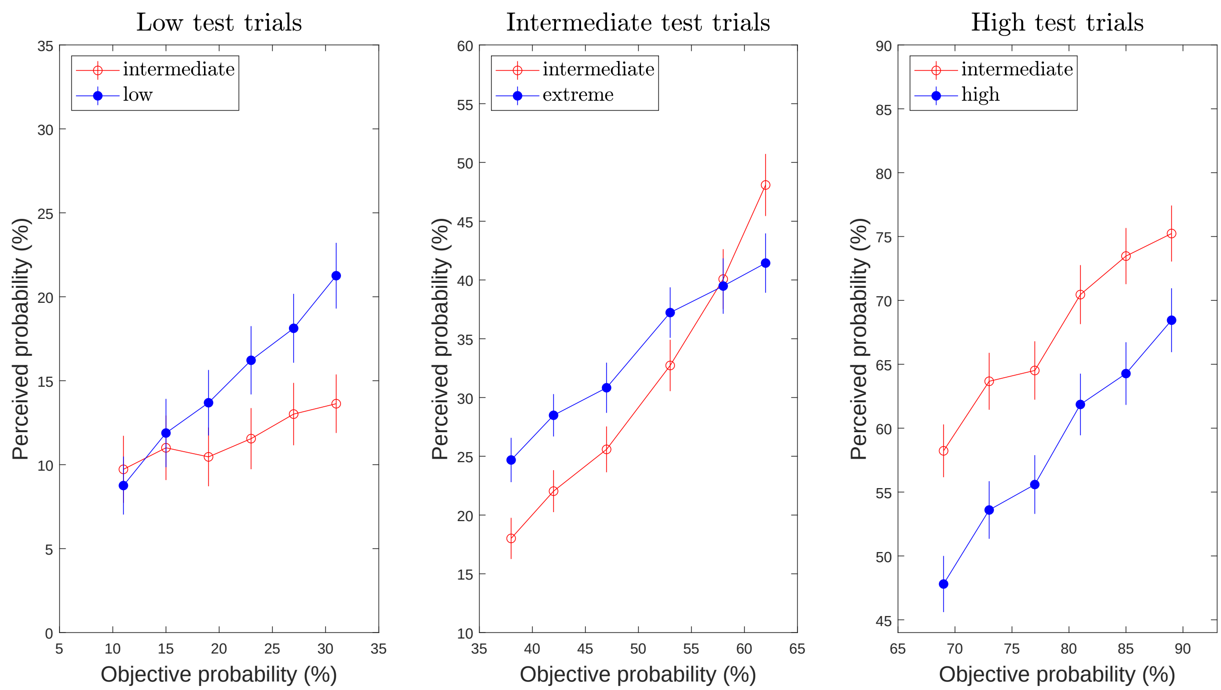

Figure III
Causal effect of prior experience on probability distortions

Each panel plots perceived probabilities against objective probabilities from one of the three test trial ranges. The legend in the upper left of each panel denotes the adaptation condition. Vertical bars denote two standard errors of the mean.

Recall from Section II and Figure I our theoretical prediction: subjects who are adapted to intermediate probabilities should have greater difficulty discriminating between low probabilities, compared to those who are adapted to low probabilities. In other words, when the perceived probability is plotted against the objective probability, the slope should be lower for subjects who are adapted to intermediate probabilities; the left panel of Figure III confirms this prediction. Similarly, subjects who are adapted to intermediate probabilities should also have greater difficulty discriminating between high probabilities, compared to those who are adapted to high probabilities. The right panel of Figure III confirms this prediction, although the slope difference is weaker compared to that estimated among low probabilities in the left panel. Finally, subjects who are adapted to intermediate probabilities should have *less* difficulty discriminating between intermediate probabilities, compared to those who are adapted to extreme probabilities; the middle panel of Figure III confirms this prediction. In sum, the three panels in Figure III provide experimental evidence for the theoretical prediction that the slope of different portions of the probability weighting function is malleable in the manner predicted by the theory of efficient coding.

Table I
Malleability of probability weighting

	(1)	(2)	(3)
Dependent variable: "Perceived probability"	Low test trials sample	Intermediate test trials sample	High test trials sample
p	0.585***	0.713***	1.012***
	(0.056)	(0.073)	(0.085)
intermediate	4.741	–27.660***	20.987**
	(3.033)	(6.103)	(8.526)
$p \times$*intermediate*	–0.395***	0.501***	–0.152
	(0.072)	(0.129)	(0.107)
Constant	2.825	–1.922	–21.318***
	(2.054)	(3.455)	(6.668)
Observations	1,120	899	1,169

Notes. The table reports results from mixed effects linear regressions, in which the dependent variable is the perceived probability, estimated from the certainty equivalent on each test trial, and the independent variables include p, *intermediate*, and the interaction between the two. The variable p takes the objective value of the probability associated with the risky lottery's upside payoff. The dummy variable *intermediate* takes the value of one if the trial belongs to the intermediate adaptation condition, and zero otherwise. The dependent variable and the independent variable p are both multiplied by 100 (in percentage). Only data from test trials are included. There are random effects on the independent variable p and the intercept. Standard errors of the fixed effect estimates are clustered at the subject level and reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

To formally test the efficient coding hypothesis, Table I presents results from three mixed effects linear regressions, one for each test trial range. For these regressions, the dependent variable is our estimate of the subject's perceived valuation of p from each test trial; the independent variables include the objective probability p, a dummy variable labeled *intermediate* which takes the value of one if the subject is randomized into the intermediate adaptation condition (and zero otherwise), and the interaction between p and *intermediate*. We include only test trials in the regressions, which allow us to estimate the effect of adaptation trials on valuation.

Column (1) reports the regression results for the low test trial range. The key variable of interest, $p \times intermediate$, has a significantly negative coefficient. This indicates that subjects who are adapted to intermediate probabilities form perception of low probability that is less sensitive to changes in the objective probability, compared to subjects who are adapted to low probabilities. Column (3) reports the regression results for the high test trial range. Again, the interaction term, $p \times intermediate$, has a negative coefficient, although it is not significant at conventional levels (p-value $= 0.16$ for a two-tailed test against the null of zero). Finally, Column (2) reports the regression results for the intermediate test trial range. Here, we see that the interaction term has a significantly *positive* coefficient. Consistent with the prediction of efficient coding, this result indicates that subjects who are adapted to intermediate probabilities form perception of intermediate probability that is *more* sensitive to changes in the objective probability, compared to subjects who are adapted to extreme probabilities.[14] Online Appendix A re-estimates the three regressions of Columns (1) to (3) using an alternative exclusion criterion, in which the filter is based on subjects' behavior on test trials rather than on adaptation trials. We find that the treatment effects become stronger: the coefficient on the interaction term in each of the three regressions is of the hypothesized sign; and it is statistically significant at the 5% level in one regression and at the 1% level in two other regressions.

In the above analyses, our key dependent variable is the subject's perceived value of p. However, because this perceived value is unobservable, we estimate it using the certainty equivalent data—under the assumption that the subject has a linear utility function. If, instead, we relax this linear utility assumption and allow the subject to have intrinsic risk aversion over the lottery payoff X—

[14]Section III.3.3 presents a replication of our main treatment effects using a separate sample of 600 subjects and a slightly modified experimental design.

that is, to assume a utility function $u(X) = X^\alpha$—where $X = 25$ is associated with probability p and $X = 0$ is associated with probability $1 - p$, then, on each trial, the certainty equivalent that the subject provides is

$$CE = [\mathbb{E}(\tilde{p}|R_p) \cdot (25)^\alpha]^{1/\alpha} = [\mathbb{E}(\tilde{p}|R_p)]^{1/\alpha} \cdot 25, \tag{11}$$

where α is the risk aversion parameter and $\alpha \leq 1$. Equation (11) implies that $\mathbb{E}(\tilde{p}|R_p) = (\frac{CE}{25})^\alpha$.

We re-estimate our main regressions in Table I, but now using $(\frac{CE}{25})^\alpha$ as the dependent variable. We have examined a variety of values for α, including $\alpha = 0.88$, the value reported in Tversky and Kahneman (1992), and $\alpha = 0.7$, the value used by Barberis, Jin, and Wang (2021) to match the behavior of real-world investors. We find that our regression results are robust to these alternative nonlinear utility specifications: the coefficient on $p \times \textit{intermediate}$ remains significantly negative for the low test trial range and it remains significantly positive for the intermediate test trial range. For the high test trial range, the result becomes stronger and attains significance at the 10% level when assuming $\alpha = 0.7$. We also note that higher risk aversion (i.e., a lower α) leads to an increase in our estimate of $v(p)$. For example, when $\alpha = 0.7$, we estimate that small probabilities are distorted upwards: $v(p) > p$ when p is small.

III.3.2. Empirical link between beliefs and valuations

In our efficient coding model, heterogeneity in prior beliefs is the main source of cross-subject variation in the probability weighting function. However, the regression analysis in Table I has treated prior beliefs as unobservable. In this section, we leverage our directly measured prior beliefs to provide additional tests of the efficient coding mechanism. We first describe the data on elicited beliefs. We then proceed to investigate how the beliefs data relate to observed valuations.

Recall that, for each of the three test trial ranges, subjects are randomized into one of two adaptation conditions. The "intermediate" condition involves probabilities drawn from around the center of the unit interval. The other condition draws probabilities from more extreme values of the unit interval. Figure IV shows the prior beliefs averaged across subjects within a condition, measured immediately before trial one (left panel) and immediately after trial fifteen (right panel). In each panel, we pool all subjects from the intermediate condition into one group and label it “intermediate,” and we pool all subjects from the non-intermediate conditions (i.e., low, high, or extreme) into the other group and label it “extreme.”

Figure IV generates two observations. First, the left panel shows no difference in measured beliefs across the groups, which is expected given the experimental randomization. While the average priors in the left panel are consistent with subjects putting greater weight on intermediate probabilities, there exists extensive variation across subjects in their reported priors. For example, while the average prior is hump-shaped, approximately 23% of subjects report a uniform prior before trial one.[15]

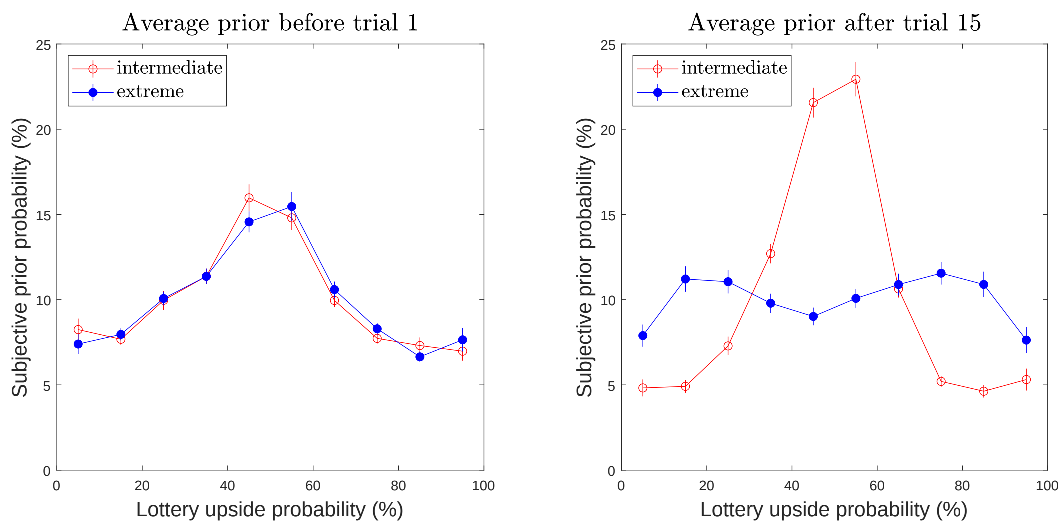

Figure IV
Measured prior beliefs by experimental condition

Prior beliefs averaged across subjects and disaggregated by adaptation condition. Each data point represents one of the ten probability bins (placed on the x-axis at the midpoint of the bin). The left panel reports average beliefs elicited at the beginning of the experiment (before trial one). The right panel reports average beliefs elicited after 15 adaptation trials. Vertical bars denote two standard errors of the mean.

The second and more important observation is that, after 15 adaptation trials, there is a strong separation in prior beliefs across the intermediate and extreme groups. Subjects who faced lotteries with the probability of the upside payoff drawn from intermediate values of p state a high likelihood of observing an intermediate probability on the next trial. Conversely, subjects who faced lotteries with the probability of the upside payoff drawn from extreme values of p state a *low* likelihood of

[15]By contrast, only 12% of subjects report a uniform prior after trial fifteen.

observing an intermediate probability on the next trial. Notice also that for the extreme group in the right panel, subjects attach the lowest likelihood to the probability bins (1-10%) and (91-99%). This result is consistent with the empirical distribution of probabilities on adaptation trials: subjects never observe a probability below 10% or above 90% in the low, high, or extreme condition.

Next, we use the prior beliefs elicited after trial fifteen to test the efficient coding mechanism. We begin with a reduced form test for each of the three test trial ranges. For each subject who is randomized into the *low* test trials with $p \in [0.11, 0.31]$, we measure how much weight they attach to this range in their prior beliefs. The hypothesis is that a subject who attaches a larger weight to this range should endogenously be more sensitive to changes in p within the range. It follows that lottery valuations from such a subject should be more sensitive to changes in p, compared to those from another subject who attaches a smaller weight to observing a value of $p \in [0.11, 0.31]$.

To measure how much weight a subject attaches to these low probabilities, we sum the density she attaches to the 11-20% bin and the 21-30% bin in her elicited beliefs. We label this sum $localDensity$. We then run a mixed effects linear regression, in which the dependent variable is our estimate of the subject's perceived valuation of p from each low test trial, and the independent variables include the objective probability p, the variable $localDensity$, and the interaction between p and $localDensity$. A positive coefficient on the interaction term would suggest that a one unit increase in p generates a larger increase in valuation for subjects who attach a higher likelihood to facing an upside probability in this range. Importantly, this analysis does not condition on the adaptation condition, which means that such a test can, in principle, be conducted in more general settings that do not require exogenous variation in prior beliefs.

Column (1) of Table II provides the regression results. We observe that the interaction term, $p \times localDensity$, has a positive coefficient that is significant at the 1% level. This confirms the hypothesis that measured beliefs about an upcoming choice set can explain variation in valuations. To our knowledge, this is the first result in the cognitive economics literature that establishes a link between measured prior beliefs about a choice set and behavior when the subject ultimately faces that choice set. We also emphasize that a sizable gap exists between the time that subjects report beliefs and the time that they report lottery valuations on test trials. Specfically, the results in Column (1) indicate that beliefs elicited after trial fifteen correlate with subjects' lottery valuations measured 15 trials later.[16]

Table II
Predicting lottery valuations using prior beliefs

	(1)	(2)	(3)
Dependent variable: "Perceived probability"	Low test trials sample	Intermediate test trials sample	High test trials sample
p	0.285***	0.840***	0.842***
	(0.045)	(0.111)	(0.076)
localDensity	–0.141***	–0.245***	–0.338
	(0.041)	(0.080)	(0.219)
$p \times$*localDensity*	0.004***	0.003*	0.005*
	(0.002)	(0.002)	(0.003)
Constant	8.460***	–2.654	–4.340
	(2.093)	(5.518)	(6.147)
Observations	1,120	899	1,169

Notes. The table reports results from mixed effects linear regressions, in which the dependent variable is the perceived probability, estimated from the certainty equivalent on each test trial, and the independent variables include p, *localDensity*, and the interaction between the two. The variable p takes the objective value of the probability associated with the risky lottery's upside payoff. The variable *localDensity* is the sum of the densities that a subject attaches to the range of upside probabilities that she faces during subsequent test trials; an upside probability is the probability associated with the upside payoff of the risky lottery. Specifically, for the low test trials, *localDensity* is the sum of the densities that a subject attaches to the 11-20% and 21-30% bins. For the intermediate test trials, *localDensity* is the sum of the densities that a subject attaches to the 31-40%, 41-50%, 51-60% and 61-70% bins. For the high test trials, *localDensity* is the sum of the densities that a subject attaches to the 71-80% and 81-90% bins. The dependent variable and the independent variables p and *localDensity* are all multiplied by 100 (in percentage). Only data from test trials are included. There are random effects on the independent variable p and the intercept. Standard errors of the fixed effect estimates are clustered at the subject level and reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

For the intermediate test trials, subjects face lotteries where the upside probability is drawn from the range [0.38, 0.62]. In this case, *localDensity* is the sum of the densities that a subject attaches to the 31-40%, 41-50%, 51-60%, and 61-70% bins. Column (2) of Table II shows that the

[16]Two reasons lead our experimental design to have a lag between belief measurement and test trials. First, we did not want subjects' attention to be fully directed to the belief measurement task right before the crucial set of test trials: this may interrupt the adaptation process. Second, a recent study provides intriguing evidence that the encoding function may adapt more slowly than the prior (Hahn and Wei, 2024).

interaction term again has a positive coefficient, which is significant at the 10% level. Finally, for the high test trials, subjects face lotteries where the upside probability is drawn from the range [0.69, 0.89]. In this case, *localDensity* is the sum of the densities that a subject attaches to the 71-80% and 81-90% bins. Column (3) of Table II shows that the coefficient on the interaction term is again significantly positive at the 10% level.

Such reduced form tests are useful as they are free of any parametric assumption; for example, they do not require specifying the capacity parameter n from equation (1). At the same time, the simplicity comes at the cost of discarding valuable information about subjects' prior beliefs: the regression analysis in Table II only uses information about subjects' prior beliefs in the "local area" of the unit interval that heavily overlaps with the range of upside probabilities presented during test trials. The efficient coding model, instead, takes into account the entire prior distribution when making predictions about perceived probabilities.

We now turn to a model-based test of the efficient coding mechanism, which takes as input the entire distribution of prior beliefs that subjects provide after 15 adaptation trials. Specifically, for each subject, we take the ten densities that she attaches to the ten probability bins and construct a prior distribution. We assume that, within each bin, the prior distribution is uniform. We then plug this prior distribution into the efficient coding model described in Section II and compute, for each test trial, the model-implied valuation of p according to equation (5); we denote this model-implied valuation $v(p|prior)$. Here, p takes the objective value of the upside probability in the test trial, and we set n, the capacity parameter in equation (1), to 10 when computing $v(p|prior)$.[17]

In Table III, we regress the perceived probability, estimated from the certainty equivalent on each test trial, on p and $v(p|prior)$. The latter variable is the model-implied valuation given the prior distribution elicited after 15 adaptation trials. Column (1) of Table III provides a benchmark of the sensitivity of the perceived probability to the objective probability p. It shows that a one unit increase in p leads to a 0.76 units increase in the perceived probability. Column (2) shows that a one unit increase in $v(p|prior)$ leads to a 0.91 units increase in the perceived probability. The crucial test is shown in Column (3): after controlling for the objective probability p, the model-implied valuation $v(p|prior)$ continues to explain variation in the lottery valuation and hence the

[17] The results in Table III are robust to different values of n.

perceived probability.[18]

Table III
Predicting valuations using model-implied distorted probabilities

	(1)	(2)	(3)
Dependent variable: "Perceived probability"			
p	0.764***		0.667***
	(0.029)		(0.045)
$v(p\|prior)$		0.907***	0.142**
		(0.037)	(0.059)
Constant	–3.456**	–8.719***	–5.635***
	(1.501)	(1.946)	(1.883)
Observations	3,188	3,188	3,188

Notes. The table reports results from mixed effects linear regressions, in which the dependent variable is the perceived probability, estimated from the certainty equivalent on each test trial, and the independent variables include p and $v(p|prior)$. The variable p takes the objective value of the probability associated with the risky lottery's upside payoff. The variable $v(p|prior)$ is the model-implied valuation of p. The dependent variable and the independent variables p and $v(p|prior)$ are all multiplied by 100 (in percentage). Only data from test trials are included. There are random effects on the independent variable p and the intercept. Standard errors of the fixed effect estimates are clustered at the subject level and reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

It is worth noting that the link between beliefs and lottery valuations is also at the core of alternative theories of risky choice, such as the Kőszegi and Rabin (2007) model of expectation-based reference points. However, there exists a crucial distinction between our model and the Kőszegi and Rabin (2007) model. In their model, the DM perfectly observes the choice set and state probabilities, and the valuation of a risky lottery is driven by the DM's beliefs about the *future* outcomes of the lottery. By contrast, in our model, the prior beliefs about a lottery's state probabilities drive lottery valuations during subsequent trials—even *after* the choice set and state probabilities on these trials have been presented to the DM. Put differently, for any lottery we display on a test trial, the Kőszegi and Rabin (2007) model predicts that valuation is independent of the prior

[18] Online Appendix A.2 reports an alternative pre-registered, but less powerful, test of the the empirical link between beliefs and valuations.

belief we measure after trial fifteen, whereas our model predicts dependence of valuation on the prior belief. We provide further discussion of alternative theories in Section IV.3.

III.3.3. Replication of main treatment effects

Table IV
Conceptual replication of treatment effects from earlier experiment in Frydman and Jin (2023)

	(1)	(2)	(3)
Dependent variable: "Perceived probability"	Low test trials sample	Intermediate test trials sample	High test trials sample
p	0.620***	0.664***	0.983***
	(0.060)	(0.059)	(0.061)
intermediate	1.710	–18.531***	14.857**
	(1.901)	(4.092)	(6.745)
p×*intermediate*	–0.312***	0.367***	–0.166**
	(0.070)	(0.094)	(0.084)
Constant	0.352	–5.279*	–21.057***
	(1.217)	(2.722)	(4.969)
Observations	1,127	1,096	1,133

Notes. The table reports results from mixed effects linear regressions using data from a separate experiment discussed in more detail in Frydman and Jin (2023). In these regressions, the dependent variable is the perceived probability, estimated from the certainty equivalent on each test trial, and the independent variables include p, *intermediate*, and the interaction between the two. The variable p takes the objective value of the probability associated with the risky lottery's upside payoff. The dummy variable *intermediate* takes the value of one if the trial belongs to the intermediate adaptation condition, and zero otherwise. The dependent variable and the independent variable p are both multiplied by 100 (in percentage). Only data from test trials are included. There are random effects on the independent variable p and the intercept. Standard errors of the fixed effect estimates are clustered at the subject level and reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

In an earlier working paper (Frydman and Jin, 2023), we conducted a similar experiment to the one reported above. The main difference is that the earlier experiment did not elicit prior beliefs.[19] We found very similar results in that experiment, whereby the sensitivity of the perceived probability to the objective probability was significantly modulated by adaptation trials in the

[19] Another difference is that the earlier experiment included 24 adaptation trials, whereas the current experiment includes 30.

manner predicted by the theory of efficient coding. Table IV shows that the results from the earlier experiment are analogous to those presented above in Table I; if anything, the earlier results are even stronger. We view this earlier experiment, which we conducted on a separate sample of 600 subjects from Prolific in December 2022, as a conceptual replication of our main treatment effects.[20]

IV. Discussion

IV.1. Noisy coding of payoffs and probabilities

The model described in Section II assumes that noise enters the decision process only when the DM encodes probabilities; that is, it makes the simplifying assumption that there is no noise in encoding lottery payoffs. To justify this assumption, we use the same lottery upside of \$25 on each trial of our experiment, and thus noisy coding of this fixed amount of \$25 should be minimal. However, previous studies have shown that, when payoffs vary across decision problems, noisy encoding of payoffs has important consequences for behavior (Khaw et al., 2021; Frydman and Jin, 2022). In this section, we study a more general case in which the DM encodes *both* payoffs and probabilities with noise.

Consider a risky lottery of the form $(\$X, p; \$0, 1-p)$; it pays $X > 0$ with probability p and zero dollars with probability $1 - p$. Suppose that the DM holds prior beliefs about X and p, denoted by $f(X, p)$. When the lottery is revealed to the DM, she draws two noisy signals: R_x, a noisy signal of X, and R_p, a noisy signal of p. The two signals are drawn from their respective likelihood functions, $f(R_x|X, p)$ and $f(R_p|X, p)$. Further suppose that the DM encodes X and p through a total number of n "neurons"—following Heng et al. (2020)—and that she chooses $f(R_x|X, p)$ and $f(R_p|X, p)$ to maximize the mutual information between the payoff-probability pair, (X, p), and its noisy representation, (R_x, R_p). Online Appendix B.1 formalizes this maximization problem.

We analyze a special case of the above model in which X and p are statistically independent; in this case, $f(X, p) = f(X) \cdot f(p)$. This special case is easier to solve, yet it remains a generalized version of our baseline model as it allows for variation in X and noisy coding of X. We solve it in Online Appendix B.2. We find that, under this model, our main predictions for the instability of

[20]See Frydman and Jin (2023) for more details on the design of the earlier experiment and the associated pre-registration document.

probability weighting continue to hold.

To illustrate our finding, we provide a numerical example. Suppose the DM holds a prior that X is drawn uniformly from $[23, 27]$ and that X and p are drawn independently. Also suppose that the DM has a total of $n = 15$ neurons for the encoding of X and p. Then, under the intermediate prior about p which is a mixture of the stable component that takes the form of a uniform distribution and the fast-moving component that takes the form in equation (8), the DM optimally allocates $n_x = 8$ neurons to encode X and $n_p = 7$ neurons to encode p. Under the extreme prior about p which is a mixture of the stable component that takes the form of a uniform distribution and the fast-moving component that takes the form in equation (9), the DM again optimally allocates $n_x = 8$ neurons to encode X and $n_p = 7$ neurons to encode p. Figure B.1 in Online Appendix B.2 shows that the weighting functions implied by the intermediate and the extreme priors are quantitatively similar to those in the lower graph of Figure I. As such, the theoretical predictions discussed in Section II regarding the malleability of probability weighting remain robust to allowing for noisy coding of X.[21]

We emphasize that the above numerical analysis assumes that X and p are statistically independent. In the more general model formalized in Online Appendix B.1, which allows for any correlation between X and p, the optimal encoding rules will change. Solving this general model and deriving its predictions are beyond the scope of the paper.[22] One potentially interesting direction for future work is to experimentally manipulate the correlation between X and p in order to assess the extent to which information about p affects the encoding precision of X. This, in turn, could guide future theorizing about how the DM jointly encodes payoffs and probabilities.

IV.2. Belief distortion as a source of probability weighting

The model that we present, which builds on earlier work by Khaw et al. (2021) and Frydman and Jin (2022), has only one free parameter n. This parameter measures the number of "neurons" that the DM uses to encode probability, so it can be interpreted as a capacity constraint. Our

[21] We note that the results from Figure B.1 are largely *invariant* with respect to the range over which X is drawn. Suppose the DM holds a prior that X is drawn uniformly from $[25 - \Delta, 25 + \Delta]$. We find that, for any Δ between 0 and 25 and for both the intermediate and the extreme priors about p, the DM always sets $n_x = 8$ and $n_p = 7$.

[22] For related theoretical work, see Khaw et al. (2024) who allow the encoding precision of X to depend on the noisy signal for p. Note that their model exogenously assumes the encoding functions; in contrast, our model endogenously derives the encoding functions from the DM's prior.

assumption that n is finite—and thus information processing is constrained—is responsible for the prediction that the average perception of p does not always equal p. As such, our model of probability weighting derives exclusively from a friction in information processing and thus points to a belief-based channel as an important driver of probability distortions.

Such a belief-based channel is worth emphasizing, given that the literature beginning with Kahneman and Tversky (1979) has largely interpreted probability weighting as an expression of preferences, rather than a misperception of beliefs. Specifically, our results suggest that the degree of probability weighting should be modulated by the agent's prior expectations, and thus information provision and past experience may play an important role in characterizing the shape of the weighting function. This interpretation is consistent with recent work by Enke and Graeber (2023), who show that the weighting function can be modulated by the complexity of the lottery under consideration. At the same time, our data do not rule out preferences as a source of probability weighting. That is, even if we shut down our noisy encoding channel, it may well be that valuation is still passed through an exogenous nonlinear probability weighting function.

IV.3. Comparison with alternative theories

IV.3.1. Efficient coding with an alternative performance objective

The efficient coding model in Section II assumes that the DM chooses the coding rule $\theta(p)$ that maximizes the mutual information between probability p and its noisy signal R_p. However, there are other plausible performance objectives that the DM may use instead (Ma and Woodford, 2020). For example, the DM may choose the coding rule so as to maximize the expected payoff from the task at hand (Heng et al., 2020). In this section, we examine our model's implications when the DM chooses the coding rule $\theta(p)$ that maximizes the expected payoff from the experiment.

Specifically, given the Becker-Degroot-Marschak (Becker et al., 1964) incentive scheme implemented in the experiment, the DM chooses $\theta(p)$ to maximize the expected payoff

$$\mathbb{E}[\text{payoff}] = \frac{25}{2}\int_0^1 (1+p^2) f(p)dp - \frac{25}{2}\int_0^1 \left(\sum\nolimits_{R_p=0}^{n} (\mathbb{E}[\tilde{p}|R_p] - p)^2 f(R_p|p)\right) f(p)dp, \tag{12}$$

where $f(p)$ is the DM's prior belief, and $f(R_p|p)$ and $\mathbb{E}[\tilde{p}|R_p]$ are given by equations (1) and (6); we derive equation (12) in Online Appendix C. Because the coding rule only affects the second term of equation (12), the performance objective can be rewritten so as to minimize the following expected error

$$\int_0^1 \left(\sum\nolimits_{R_p=0}^{n} (\mathbb{E}[\tilde{p}|R_p] - p)^2 f(R_p|p) \right) f(p)dp. \tag{13}$$

We now examine whether our efficient coding model still gives rise to malleability of probability weighting under the alternative assumption that the DM chooses $\theta(p)$ to maximize the expected payoff in equation (12). We again consider the two mixed priors specified in Section II and Figure I. One mixed prior applies equal weights to a stable component that is uniform and a fast-moving component that takes the form of (8) with $p_l = 0.38$ and $p_h = 0.62$. The other mixed prior we examine applies equal weights to the same uniform stable component and a fast-moving component that takes the form of (9) with $p_{l,1} = 0.1$, $p_{l,2} = 0.21$, $p_{h,1} = 0.79$, and $p_{h,2} = 0.9$. Figure C.1 in Online Appendix C plots these two mixed priors, the optimal coding rules $\hat{\theta}(p)$, as well as the subjective valuations $\hat{v}(p)$ implied by $\hat{\theta}(p)$.[23] The figure shows that two very different priors both lead to similar probability weighting functions. In other words, when the DM chooses the optimal code to maximize the expected payoff, our efficient coding model does not generate malleability of probability weighting. This finding, in conjunction with our experimental results from Section III, implies that subjects in our experiment are likely maximizing mutual information rather than the expected payoff.[24]

IV.3.2. Expectations-based reference points: Kőszegi and Rabin (2007)

Consider a risky lottery $\mathcal{L}$ that has N potential outcomes, x_1, x_2,..., x_N; outcome x_n is associated with probability p_n. Kőszegi and Rabin (2007) propose a model of the reference point, whereby the DM equates the lottery's payoff distribution with the reference point distribution.

[23]We apply a projection method with Chebyshev polynomials to numerically solve for the $\hat{\theta}(p)$ that minimizes (13). See Mason and Handscomb (2003) for a detailed discussion of the properties of Chebyshev polynomials.

[24]When comparing different performance objectives that can be implemented in an efficient coding model, one caveat is that there are also multiple possible information processing constraints for each performance objective. The particular constraint we impose in our model is inherited from Heng et al. (2020), but other possible constraints—for example, the one imposed in Wei and Stocker (2015)—have also been examined. For further discussion on the comparison between minimizing error and maximizing discriminability in the context of probabilities, see Zhang et al. (2020).

The DM's lottery valuation is, therefore, given by

$$v(\mathcal{L}) = \sum_{i=1}^{N}\sum_{j=1}^{N} p_i p_j (x_i + \mu(x_i - x_j)), \tag{14}$$

where

$$\mu(y) = \begin{cases} \eta \cdot y & \text{if } y \geq 0 \\ (\eta\lambda) \cdot y & \text{if } y < 0 \end{cases}, \tag{15}$$

η measures the relative importance of the gain-loss utility, and λ measures the degree of loss aversion.

Under this specification, the lottery valuation $v(\mathcal{L})$ depends only on the payoff and probability distributions $\{x_n, p_n\}_{n=1}^{N}$, which, in our experiment, is held constant across different adaptation conditions for a given test trial. Yet, our experiment shows that, after subjects are presented with the payoff and probability distributions $\{x_n, p_n\}_{n=1}^{N}$, lottery valuations are still affected by prior beliefs about the probabilities. As such, the model of Kőszegi and Rabin (2007) cannot explain our experimental result of malleability with respect to the prior.

IV.3.3. Endogenizing the encoding function

There have been many earlier and ongoing attempts to microfound the probability weighting function. Our model is closest in spirit to recent models of noisy and efficient coding of probability. In particular, Khaw et al. (2021, 2024) assume that the DM encodes probabilities using a log odds encoding function. Specifically, in their models, a noisy cognitive signal R_p is drawn as follows:

$$f(R_p|p) = \frac{1}{\sqrt{2\pi v_z^2}} \exp\left(-\frac{1}{2v_z^2}\left[R_p - \log\left(\frac{p}{1-p}\right)\right]^2\right). \tag{16}$$

That is, the noisy signal R_p is drawn from a Normal distribution with mean of $\log(\frac{p}{1-p})$ and standard deviation of v_z. Figure V presents the two mixed prior distributions specified in Figure I and the probability weighting functions implied by these priors and the likelihood function of (16).

Figure V leads to two observations. First, a log odds encoding function tends to always produce an inverse S-shaped weighting function, hence leaving little room for malleability of the general *shape* of the weighting function. Second, when the likelihood function is exogenously assumed as

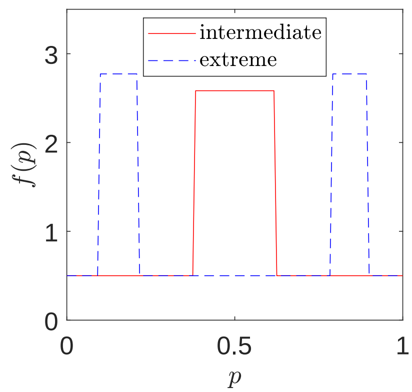

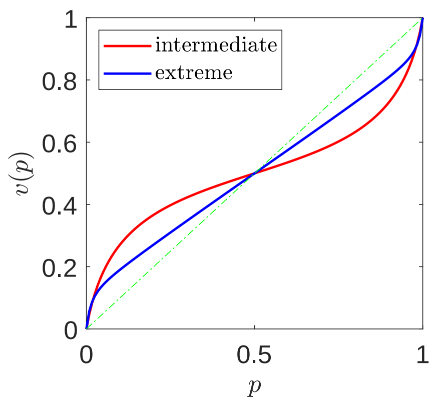

Figure V
Prior distribution and value function implied by the log odds encoding function from Khaw et al. (2021, 2024)

The upper graph plots two mixed prior distributions $f(p)$ in the form of (7). For each mixed prior, the first, stable component is assumed to be a uniform distribution between 0 and 1. The second, fast-moving component takes the form of (8) for the intermediate prior and the form of (9) for the extreme prior; the parameter values are: $p_l = 0.38$, $p_h = 0.62$, $p_{l,1} = 0.1$, $p_{l,2} = 0.21$, $p_{h,1} = 0.79$, and $p_{h,2} = 0.9$. The weight ξ the DM assigns to the stable component is 0.5. The lower graph plots, for both the intermediate prior and the extreme prior, the probability weighting function implied by the mixed prior and the log odds encoding function of (16); here we set $v_z = 1$. The green dash-dot line is the forty-five degree line.

a log odds function, the intermediate prior gives rise to a lower slope of the weighting function for intermediate probabilities; this implication is inconsistent with the experimental finding reported in Table I that, in the data, the intermediate prior actually gives rise to a *higher* slope of the weighting function for intermediate probabilities.

The recent work of Enke and Graeber (2023) also provides a microfoundation of the probability weighting function. In their model, the noisy signal R_p is drawn according to:

$$f(R_p|p) = \binom{n}{R_p}(p)^{R_p}(1-p)^{n-R_p}. \tag{17}$$

This likelihood function is exogenously specified and not tied to the prior distribution. Figure VI plots the weighting function $v(p)$ implied by this likelihood function and either of the two mixed priors specified in Figure I.

Figure VI shows that, similar to the result in Figure V, the model of Enke and Graeber (2023) generates the same counterfactual implication: the intermediate prior implies a lower slope of the weighting function for intermediate probabilities, while in the data, the intermediate prior implies a higher slope. Taken together, Figures V and VI highlight the importance of allowing the likelihood function to be endogenously tied to the prior distribution—this is the approach our paper takes.

IV.4. Broader implications of imprecision of ratios

Throughout the paper, we have mainly been concerned with deriving and testing implications for risky choice when the DM has an imprecise representation of *probability*. To this end, our experiment uses explicit representations of state probabilities, both numerically (displaying percentages) and visually (displaying the ratio of two colored bars).

However, our theory of efficient coding can be applied more broadly to real numbers in the unit interval that do not necessarily have to represent probabilities. For example, our framework of noisy and efficient coding will also make predictions about perception of proportions when there is no intrinsic risk to consider. To make ideas concrete, consider the task in Oprea (2024) where the DM must report a valuation of a riskless—but disaggregated—set of monetary payments. In particular, subjects are asked to provide a valuation for an asset that is comprised of 100 disaggregated payments. A proportion p of these payments are each worth \$2.50 and the remaining proportion $1-p$ of the payments are worth \$0. Thus, the subject's task is simply to aggregate the 100 payments into a single valuation.

If the subject encodes the proportion p with noise, this can drive a wedge between the subject's reported valuation and the "true" valuation—the latter is simply the sum of the 100 disaggregated payments. This is consistent with the finding of Oprea (2024) that subjects report valuations—what

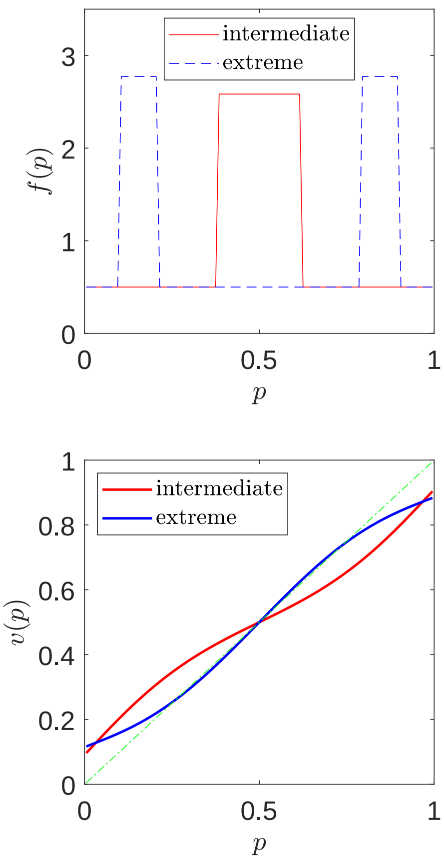

Figure VI
Prior distribution and value function implied by the encoding function from Enke and Graeber (2023)

The upper graph plots two mixed prior distributions $f(p)$ in the form of (7). For each mixed prior, the first, stable component is assumed to be a uniform distribution between 0 and 1. The second, fast-moving component takes the form of (8) for the intermediate prior and the form of (9) for the extreme prior; the parameter values are: $p_l = 0.38$, $p_h = 0.62$, $p_{l,1} = 0.1$, $p_{l,2} = 0.21$, $p_{h,1} = 0.79$, and $p_{h,2} = 0.9$. The weight ξ the DM assigns to the stable component is 0.5. The lower graph plots, for both the intermediate prior and the extreme prior, the probability weighting function implied by the mixed prior and the likelihood function of (17); here we set the parameter n to 10. The green dash-dot line is the forty-five degree line.

he calls "simplicity equivalents"—that seem to accord with the four-fold pattern of risk-taking, even in the absence of any intrinsic risk.[25] Moreover, our model makes an additional untested prediction, which is that the pattern of reported simplicity equivalents can be modulated by the DM's prior

[25] See also Vieider (2024) who replicates the Oprea (2024) findings using binary choice rather than the multiple price list approach.

belief about proportions.

IV.5. Applications to financial economics

While probability weighting has been applied in a variety of areas of economics, including insurance and betting markets, it has been particularly fruitful in explaining facts in financial economics (Barberis, 2013; O'Donoghue and Somerville, 2018). The core component of probability weighting that most researchers in financial economics have invoked is the overweighting of small probabilities. Models with this basic assumption can generate the high equity premium (De Giorgi and Legg, 2012), abnormal returns of individual stocks (Barberis and Huang, 2008; Barberis, Mukherjee, and Wang, 2016; Barberis et al., 2021), overpricing of out-of-the-money stock options (Baele, Driessen, Ebert, Londono, and Spalt, 2019), and time-inconsistent risk-taking (Barberis, 2012; Ebert and Strack, 2015; Heimer, Iliewa, Imas, and Weber, 2025). Importantly, in all these studies, researchers assume that probability distortions are "hard-wired" and do not shift over time.

Our results suggest that the degree of probability distortions depends critically on the DM's prior belief about probability, which can vary over time. It follows that the financial phenomena described above should also vary over time—to the extent that they are generated by probability weighting. This insight leads to two concrete paths forward for future empirical work. First, our model predicts *when* probability distortions should arise, and thus when we should observe anomalous asset prices and risk-taking behavior. For example, when combined with the theory in Barberis et al. (2021), our model predicts conditional moments of asset prices by identifying periods when overweighting of small probabilities should be more pronounced. During these periods—when perception of probability is heavily distorted—we should observe average returns of individual stocks that are consistent with those returns documented in anomalies that can be explained by probability weighting, for example, the idiosyncratic volatility and failure probability anomalies. But during other periods, when investors' prior beliefs induce less perceptual distortions, the asset pricing anomalies should be weaker.

Second, by connecting probability distortions to investors' prior beliefs, our model also makes predictions about individual heterogeneity (Bruhin, Fehr-Duda, and Epper, 2010). In particular, those investors who hold prior beliefs that differ significantly from a uniform distribution will distort probabilities more than investors who believe that probabilities are uniformly distributed.

One challenge, of course, in testing these predictions is to obtain accurate measures of real-world investors' prior beliefs about probabilities.

V. Conclusion

We have provided a new explanation for probability weighting based on a core principle from neuroscience called efficient coding. We first develop a model to show that efficient coding readily generates nonlinear and malleable probability weighting. In particular, as the DM's prior beliefs change, the slope of different portions of the probability weighting function changes in the manner predicted by efficient coding.

Our main experimental contribution is to unveil the close connection between prior beliefs and lottery valuations. To demonstrate this connection, we exogenously manipulate the distribution of lottery probabilities faced by subjects and then directly measure prior beliefs. We find strong evidence that valuations, and hence the shape of the probability weighting function, are systematically unstable. In particular, as subjects report a higher likelihood of observing a given range of probabilities for an upcoming lottery, their lottery valuations become more sensitive to a change in probability over this range. This empirical correlation between *ex-ante* beliefs and subsequent valuations validates a critical but previously untested prediction from efficient coding models.

In closing, it is useful to highlight the connection between the experimental results in our paper and the patterns documented in recent experiments by Payzan-LeNestour and Woodford (2022) and Frydman and Jin (2022). Each of the two latter papers provides experimental support in favor of efficient coding as a driver of diminishing sensitivity, which is another core ingredient of prospect theory. Thus, when viewed against this broader perspective, our data add credence to the hypothesis that efficient coding serves as a common mechanism underlying multiple features of prospect theory (Woodford, 2012a,b).

REFERENCES

Baele, Lieven, Joost Driessen, Sebastian Ebert, Juan M Londono, and Oliver G Spalt, 2019, Cumulative prospect theory, option returns, and the variance premium, *Review of Financial Studies* 32, 3667–3723.

Barberis, Nicholas, 2013, Thirty years of prospect theory in economics: A review and assessment, *Journal of Economic Perspectives* 27, 173–196.

Barberis, Nicholas, 2018, Psychology-based models of asset prices and trading volume, in Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds., *Handbook of Behavioral Economics* (North Holland, Amsterdam).

Barberis, Nicholas C., 2012, A model of casino gambling, *Management Science* 58, 35–51.

Barberis, Nicholas C., and Ming Huang, 2008, Stocks as lotteries: The implications of probability weighting for security prices, *American Economic Review* 98, 2066–2100.

Barberis, Nicholas C., Lawrence J. Jin, and Baolian Wang, 2021, Prospect theory and stock market anomalies, *Journal of Finance* 76, 2639–2687.

Barberis, Nicholas C., Abhiroop Mukherjee, and Baolian Wang, 2016, Prospect theory and stock returns: An empirical test, *Review of Financial Studies* 29, 3068–3107.

Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak, 1964, Measuring utility by a single-response sequential method, *Behavioral Science* 9, 226–232.

Bernheim, Douglas, and Charles Sprenger, 2020, On the empirical validity of cumulative prospect theory: Experimental evidence of rank-independent probability weighting, *Econometrica* 88, 1363–1409.

Bhui, Rahul, and Samuel Gershman, 2018, Decision by sampling implements efficient coding of psychoeconomic functions, *Psychological Review* 125, 3789–3811.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2012, Salience theory of choice under risk, *Quarterly Journal of Economics* 127, 1243–1285.

Bruhin, Adrian, Helga Fehr-Duda, and Thomas Epper, 2010, Risk and rationality: Uncovering heterogeneity in probability distortion, *Econometrica* 78, 1375–1412.

Burke, Christopher J., Michelle Baddeley, Philippe N. Tobler, and Wolfram Schultz, 2016, Partial adaptation of obtained and observed value signals preserves information about gains and losses, *Journal of Neuroscience* 36, 10016–10025.

Camerer, Colin F., and Teck-Hua Ho, 1994, Violations of the betweenness axiom and nonlinearity in probability, *Journal of Risk and Uncertainty* 8, 167–196.

Conen, Katherine E., and Camillo Padoa-Schioppa, 2019, Partial adaptation to the value range in the macaque orbitofrontal cortex, *Journal of Neuroscience* 39, 3498–3513.

De Giorgi, Enrico G., and Shane Legg, 2012, Dynamic portfolio choice and asset pricing with narrow framing and probability weighting, *Journal of Economic Dynamics and Control* 36, 951–972.

Ebert, Sebastian, and Philipp Strack, 2015, Until the bitter end: On prospect theory in a dynamic context, *American Economic Review* 105, 1618–1633.

Enke, Benjamin, and Thomas Graeber, 2023, Cognitive uncertainty, *Quarterly Journal of Economics* 138, 2021–2067.

Frydman, Cary, and Lawrence J. Jin, 2022, Efficient coding and risky choice, *Quarterly Journal of Economics* 137, 161–213.

Frydman, Cary, and Lawrence J. Jin, 2023, On the source and instability of probability weighting, SSRN working paper No. 4493313.

Girshick, Ahna R., Michael S. Landy, and Eero P. Simoncelli, 2011, Cardinal rules: Visual orientation perception reflects knowledge of environmental statistics, *Nature Neuroscience* 14, 926–932.

Gonzalez, Richard, and George Wu, 1999, On the shape of the probability weighting function, *Cognitive Psychology* 38, 129–166.

Hahn, Michael, and Xue-Xin Wei, 2024, A unifying theory explains seemingly contradictory biases in perceptual estimation, *Nature Neuroscience* 27, 793–804.

Heimer, Rawley, Zwetelina Iliewa, Alex Imas, and Martin Weber, 2025, Dynamic inconsistency in risky choice: Evidence from the lab and field, *American Economic Review* 115, 330–363.

Heng, Joseph, Michael Woodford, and Rafael Polanía, 2020, Efficient sampling and noisy decisions, *eLife* 1–49.

Kahneman, Daniel, and Amos Tversky, 1979, Prospect theory: An analysis of decision under risk, *Econometrica* 47, 263–291.

Khaw, Mel Win, Ziang Li, and Michael Woodford, 2021, Cognitive imprecision and small-stakes risk aversion, *Review of Economic Studies* 88, 1979–2013.

Khaw, Mel Win, Ziang Li, and Michael Woodford, 2024, Cognitive imprecision and stake-dependent risk attitudes, NBER working paper No. 30417.

Kőszegi, Botond, and Matthew Rabin, 2007, Reference-dependent risk attitudes, *American Economic Review* 97, 1047–1073.

Laughlin, Simon, 1981, A simple coding procedure enhances a neuron's information capacity, *Zeitschrift fur Naturforschung* 36, 9–10.

Ma, Wei Ji, and Michael Woodford, 2020, Multiple conceptions of resource rationality, *Behavioral and Brain Sciences* 43, 1–60.

Mason, John C., and David C. Handscomb, 2003, *Chebyshev Polynomials* (Chapman & Hall/CRC, New York).

McGranaghan, Christina, Kirby Nielsen, Ted O'Donoghue, Jason Somerville, and Charles Sprenger, 2024, Distinguishing common ratio preferences from common ratio effects using paired valuation tasks, *American Economic Review* 114, 307–347.

O'Donoghue, Ted, and Jason Somerville, 2018, Modeling risk aversion in economics, *Journal of Economic Perspectives* 32, 91–114.

Oprea, Ryan, 2024, Decisions under risk are decisions under complexity, *American Economic Review* 114, 3789–3811.

Payzan-LeNestour, Elise, and Michael Woodford, 2022, Outlier blindness: A neurobiological foundation for neglect of financial risk, *Journal of Financial Economics* 143, 1316–1343.

Quiggin, John, 1982, A theory of anticipated utility, *Journal of Economic Behavior & Organization* 3, 323–343.

Rottenstreich, Yuval, and Christopher K. Hsee, 2001, Money, kisses, and electric shocks: On the affective psychology of risk, *Psychological Science* 12, 185–268.

Steiner, Jakub, and Colin Stewart, 2016, Perceiving prospects properly, *American Economic Review* 106, 1601–1631.

Stewart, Neil, Nick Chater, and Gordon Brown, 2006, Decision by sampling, *Cognitive Psychology* 53, 1–26.

Tversky, Amos, and Daniel Kahneman, 1992, Advances in prospect theory: Cumulative representation of uncertainty, *Journal of Risk and Uncertainty* 5, 297–323.

Vieider, Ferdinand M., 2024, Decisions under uncertainty as bayesian inference on choice options, *Management Science* 70, 9014–9030.

Viscusi, W Kip, 1989, Prospective reference theory: Toward an explanation of the paradoxes, *Journal of Risk and Uncertainty* 2, 235–263.

Wark, Barry, Adrienne Fairhall, and Fred Rieke, 2009, Timescales of inference in visual adaptation, *Neuron* 61, 750–761.

Weber, Alison I, Kamesh Krishnamurthy, and Adrienne L Fairhall, 2019, Coding principles in adaptation, *Annual Review of Vision Science* 5, 427–449.

Wei, Xue-Xin, and Alan A. Stocker, 2015, A bayesian observer model constrained by efficient coding can explain 'anti-bayesian' percepts, *Nature Neuroscience* 18, 1509–1517.

Woodford, Michael, 2012a, Inattentive valuation and reference-dependent choice, Working paper.

Woodford, Michael, 2012b, Prospect theory as efficient perceptual distortion, *American Economic Review Papers and Proceedings* 102, 41–46.

Yaari, Menahem E, 1987, The dual theory of choice under risk, *Econometrica* 95–115.

Zhang, Hang, and Laurence T. Maloney, 2012, Ubiquitous log odds: a common representation of probability and frequency distortion in perception, action, and cognition, *Frontier in Neuroscience* 6, 1–14.

Zhang, Hang, Xiangjun Ren, and Laurence T. Maloney, 2020, The bounded rationality of probability distortion, *Proceedings of the National Academy of Sciences* 117, 22024–22034.

Zimmermann, Jan, Paul W. Glimcher, and Kenway Louie, 2018, Multiple timescales of normalized value coding underlie adaptive choice behavior, *Nature Communications* 9, 3206.

Online Appendix

A. Additional Empirical Analyses of Experimental Data

A.1. Re-estimating main regressions with an alternative exclusion criterion

Our main tests of the malleability of probability weighting are summarized in Table I of the main text, whereby we followed the pre-registered exclusion criterion by taking out subjects who violated a basic monotonicity property in their valuations. Specifically, for each subject, we regressed her certainty equivalents, elicited from the adaptation trials (trials 1-30), on the lottery's upside probabilities. We then excluded any subject with an estimated negative slope. In this section, we re-estimate the main regressions by applying a filter that takes out subjects based on their behavior on test trials rather than on adaptation trials; this alternative exclusion criterion alleviates the concern that, across different adaptation conditions, our filter is applied to different lotteries.

Table A.1
Malleability of probability weighting: An alternative exclusion criterion

	(1)	(2)	(3)
Dependent variable: "Perceived probability"	Low test trials sample	Intermediate test trials sample	High test trials sample
p	0.651***	0.824***	1.113***
	(0.056)	(0.066)	(0.076)
intermediate	1.934	–20.405***	24.262**
	(2.079)	(5.939)	(7.557)
$p\times$*intermediate*	–0.336***	0.356***	–0.187**
	(0.065)	(0.120)	(0.095)
Constant	0.666	–5.741*	–31.084***
	(1.610)	(3.280)	(5.904)
Observations	990	894	1,115

Notes. The table reports results from mixed effects linear regressions, in which the dependent variable is the perceived probability, estimated from the certainty equivalent on each test trial, and the independent variables include p, *intermediate*, and the interaction between the two. The variable p takes the objective value of the probability associated with the risky lottery's upside payoff. The dummy variable *intermediate* takes the value of one if the trial belongs to the intermediate adaptation condition, and zero otherwise. The dependent variable and the independent variable p are both multiplied by 100 (in percentage). Only data from test trials are included, and we apply the alternative exclusion criterion described in Online Appendix A.1. There are random effects on the independent variable p and the intercept. Standard errors of the fixed effect estimates are clustered at the subject level and reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

For each subject, we regress her certainty equivalents, elicited from the test trials (trials 31-36), on the lottery's upside probabilities. The regression coefficient is negative for 117 out of 750 subjects in our sample. We exclude these 117 subjects. For the remaining 633 subjects, we re-run the regressions from Table I and report the results in Table A.1. We find that the treatment effects are robust to the alternative exclusion criterion. If anything, the results become stronger: the p-values on the treatment effects for low, intermediate, and high test trials are 0.001, 0.003, and 0.049, respectively.

A.2. Pre-registered test linking beliefs and valuations

In this section, we report results from a pre-registered test linking beliefs and valuations. Specifically, for each subject, we regress $v(p|prior)$, which is the model-implied valuation of p, on the objective probability p, using test trials only. We then record the subject-specific coefficient from this regression; the coefficient for subject s is denoted as β_s. For each test trial condition, we conduct a median split on β_s. Those subjects who are above the median are classified as "sensitive." Those who are at or below the median are classified as "insensitive." Hence, the classification of "sensitive" is done without any data on certainty equivalents as it relies on beliefs data only.

We then test, for each test trial condition, whether "sensitive" subjects display a stronger correlation between their certainty equivalents and the lottery's upside probabilities, compared to insensitive subjects. To implement this test, we re-run the regressions in Table I, replacing the "intermediate" dummy with the "sensitive" dummy. The results are shown in Table A.2. Under this specification, we do not find that beliefs can be used to significantly predict valuations. One potential reason for this null effect is that the test here is less powerful than the one we report in the main text. In particular, the regression we report here relies crucially on the "sensitive" dummy variable, which contains a nontrivial amount of estimation error. The source of this estimation error can be in part traced to the fact that the "sensitive" dummy variable is estimated from a regression for each subject that contains only six data points.

Table A.2
Pre-registered regressions linking beliefs and valuations

	(1)	(2)	(3)
Dependent variable: "Perceived probability"	Low test trials sample	Intermediate test trials sample	High test trials sample
p	0.334***	1.035***	0.948***
	(0.044)	(0.104)	(0.072)
sensitive	–0.093	5.648	5.178
	(3.075)	(6.565)	(8.740)
$p\times$*sensitive*	0.098	–0.123	–0.019
	(0.077)	(0.136)	(0.109)
Constant	5.314***	–19.125***	–13.750**
	(1.927)	(4.961)	(5.735)
Observations	1,120	899	1,169

Notes. The table reports results from mixed effects linear regressions, in which the dependent variable is the perceived probability, estimated from the certainty equivalent on each test trial, and the independent variables include p, *sensitive*, and the interaction between the two. The variable p takes the objective value of the probability associated with the risky lottery's upside payoff. The dummy variable *sensitive* takes the value of one if a subject is classified as "sensitive" according to the definition described in Online Appendix A.2, and zero otherwise. The dependent variable and the independent variable p are both multiplied by 100 (in percentage). Only data from test trials are included. There are random effects on the independent variable p and the intercept. Standard errors of the fixed effect estimates are clustered at the subject level and reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

B. Multi-Dimensional Efficient Coding

B.1. A general setup

Consider a binary risky lottery: ($\$X, p; \$0, 1-p$). In this section, we analyze an efficient coding model with two dimensions: the first dimension is X, the lottery upside; the second dimension is p, the probability that the risky lottery delivers X. Suppose that the DM holds prior beliefs about X and p, denoted by $f(X,p)$. When the lottery is revealed to the DM, she draws a noisy signal, R_x, of X, from the likelihood function $f(R_x|X,p)$ and a noisy signal, R_p, of p, from the likelihood function $f(R_p|X,p)$. We then assume that the DM encodes X and p through a total number of n "neurons" and that she maximizes

$$I((X,p);(R_x,R_p)), \tag{B.1}$$

which is the mutual information between the payoff-probability pair, (X,p), and its noisy representation (R_x, R_p).

Specifically, the DM chooses n_x, the number of neurons used to encode X, and $n_p = n - n_x$, the number of neurons used to encode p. Given n_x, n_p, and the coding constraint from Heng et al. (2020), the DM then chooses the two likelihood functions

$$\begin{aligned} f(R_x|X,p) &= \binom{n_x}{R_x}(\theta_x(X,p))^{R_x}(1-\theta_x(X,p))^{n_x-R_x} \\ f(R_p|X,p) &= \binom{n_p}{R_p}(\theta_p(X,p))^{R_p}(1-\theta_p(X,p))^{n_p-R_p} \end{aligned} \tag{B.2}$$

by choosing the two coding functions $\theta_x(X,p)$ and $\theta_p(X,p)$, each bounded between 0 and 1. In other words, the DM chooses n_x, an integer between 0 and n, and two coding functions $\theta_x(X,p)$ and $\theta_p(X,p)$ to maximize the mutual information in (B.1).

B.2. A special case with independent priors

We now analyze a special case of the above model in which the DM has a prior that X and p are drawn independently; that is, $f(X,p) = f(X)f(p)$. In this case, p contains no information about X and therefore $f(R_x|X,p) = f(R_x|X)$; similarly, X contains no information about p and therefore $f(R_p|X,p) = f(R_p|p)$. Moreover, it is easy to show that, when X and p are drawn independently,

$$I((X,p);(R_x,R_p)) = I(X;R_x) + I(p;R_p). \tag{B.3}$$

That is, the overall mutual information in (B.1) is equal to the sum of $I(X;R_x)$, the mutual information between X and R_x, and $I(p;R_p)$, the mutual information between p and R_p.

The two likelihood functions in (B.2) are reduced to

$$\begin{aligned} f(R_x|X) &= \binom{n_x}{R_x}(\theta(X))^{R_x}(1-\theta(X))^{n_x-R_x}, \\ f(R_p|p) &= \binom{n_p}{R_p}(\theta(p))^{R_p}(1-\theta(p))^{n_p-R_p}, \end{aligned} \tag{B.4}$$

where

$$\theta(X) = \left(\sin\left(\frac{\pi}{2}F(X)\right)\right)^2 \quad \text{and} \quad \theta(p) = \left(\sin\left(\frac{\pi}{2}F(p)\right)\right)^2 \tag{B.5}$$

are derived in Heng et al. (2020). As such, the constrained optimization problem faced by the *DM* reduces to

$$\max_{n_x>0, n_p>0}\{I(X;R_x) + I(p;R_p)\}, \tag{B.6}$$

subject to equations (B.4) and (B.5) and $n_x + n_p = n$. We solve this constrained optimization problem numerically.

Figure B.1 provides a numerical example of the model's solution. Suppose the *DM* holds a prior that X is drawn uniformly from $[23, 27]$ and that X and p are independent. Further suppose that the *DM* has a total of $n = 15$ neurons to encode X and p. We then consider the two mixed priors, which take the form of (7) described in Section II. For both mixed priors, the stable component is assumed to be a uniform distribution between 0 and 1. For the intermediate prior, the fast-moving component takes the form of (8). For the extreme prior, the fast-moving component takes the form of (9). The parameter values are: $p_l = 0.38$, $p_h = 0.62$, $p_{l,1} = 0.1$, $p_{l,2} = 0.21$, $p_{h,1} = 0.79$, and $p_{h,2} = 0.9$. The weight ξ the *DM* assigns to the stable component is 0.5.

Following the constrained optimization described by (B.6), the *DM* optimally allocates $n_x = 8$ neurons to encode X and $n_p = 7$ neurons to encode p under the intermediate prior; she optimally allocates $n_x = 8$ neurons to encode X and $n_p = 7$ neurons to encode p under the extreme prior. Figure B.1 plots the probability weighting functions implied by the intermediate and the extreme priors. Comparing Figure B.1 with Figure I, we observe that the theoretical predictions discussed in Section II regarding the malleability of probability weighting remain robust to allowing for noisy coding of X.

Interestingly, the results presented in Figure B.1 are *invariant* with respect to the range over which X is drawn. Specifically, suppose the *DM* holds a prior that X is drawn uniformly from $[25-\Delta, 25+\Delta]$ and that X and p are independent. We find that, for any Δ between 0 and 25 and for both the intermediate prior and the extreme prior, the *DM* always chooses to allocate $n_x = 8$ neurons to encode X and $n_p = 7$ neurons to encode p. In this sense, the predicted malleability of probability weighting holds for a large class of uniform priors over X. This result is notable because it implies that, as the support of the prior over X expands, the *DM* does not choose to reallocate any coding resources away from the p dimension.

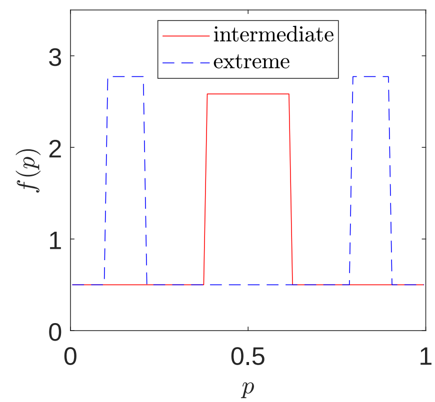

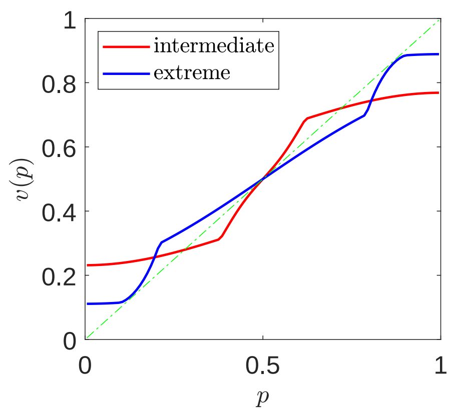

Figure B.1

Prior distribution and value function: Multi-dimensional efficient coding

The upper graph plots two mixed prior distributions about p in the form of (7) described in the main text. For each mixed prior, the first, stable component is assumed to be a uniform distribution between 0 and 1. The second, fast-moving component takes the form of (8) for the intermediate prior and the form of (9) for the extreme prior; the parameter values are: $p_l = 0.38$, $p_h = 0.62$, $p_{l,1} = 0.1$, $p_{l,2} = 0.21$, $p_{h,1} = 0.79$, and $p_{h,2} = 0.9$. The weight ξ the DM assigns to the stable component is 0.5. The lower graph plots, for both the intermediate prior and the extreme prior, the probability weighting function implied by the multi-dimensional efficient coding model described in Online Appendix B; here we set $n = 15$. For both the intermediate and the extreme priors about p, X is drawn uniformly from $[23, 27]$. In both cases, the DM optimally chooses $n_x = 8$ and $n_p = 7$. The green dash-dot line is the forty-five degree line.

C. Expected Payoff from Experiment

In this section, we derive the expected payoff a subject receives from the experiment. As described in Section III, each subject receives a fixed amount of \$3.00 for participating in the experiment. As such, when computing the subject's expected payoff, we focus only on the bonus component.

Suppose, for a given trial, the objective probability is p. Given p, the subject's perceptual system generates a noisy signal R_p from $f(R_p|p)$. Then, given R_p, the subject reports a certainty equivalent of $25 \cdot \mathbb{E}[\tilde{p}|R_p]$. With the Becker-Degroot-Marschak (Becker et al., 1964) incentive scheme, the expected bonus payoff the subject receives, conditional on p and R_p, is

$$\begin{aligned}
& \frac{1}{25}\int_{25\cdot\mathbb{E}[\tilde{p}|R_p]}^{25} q dq + \frac{1}{25}\int_{0}^{25\cdot\mathbb{E}[\tilde{p}|R_p]} (25\times p) dq \\
= \quad & \frac{1}{25}\int_{25p}^{25} q dq + \frac{1}{25}\int_{0}^{25p} (25\times p) dq - \frac{1}{25}\int_{25p}^{25\cdot\mathbb{E}[\tilde{p}|R_p]} (q - 25p) dq \\
= \quad & \frac{25}{2}(1+p^2) - \frac{25}{2}(\mathbb{E}[\tilde{p}|R_p] - p)^2. \qquad \text{(C.1)}
\end{aligned}$$

Note that equation (C.1) corresponds to the limiting case of the 2,500-row payoff scheme we implemented in the experiment; the details of this 2,500-row payoff scheme are provided in Online Appendix E. As the number of rows increases from 2,500 to infinity, the expected bonus payoff approaches the expression given by equation (C.1).

Averaging across different values of R_p for a given p, and further averaging across different values of p drawn from the subject's prior belief $f(p)$, the expected bonus payment is

$$\mathbb{E}[\text{payoff}] = \frac{25}{2}\int_0^1 (1+p^2) f(p) dp - \frac{25}{2}\int_0^1 \left(\sum\nolimits_{R_p=0}^{n} (\mathbb{E}[\tilde{p}|R_p] - p)^2 f(R_p|p)\right) f(p) dp, \qquad \text{(C.2)}$$

which is equation (12) in the main text.

We consider the two mixed prior distributions—one intermediate and one extreme—presented in Figure I of the main text. Given the DM's objective of maximizing the expected payoff in equation (C.2), Figure C.1 plots the two mixed priors, the optimal coding rules $\hat{\theta}(p)$, which we solve numerically, as well as the subjective valuations $\hat{v}(p)$ implied by the coding rules.

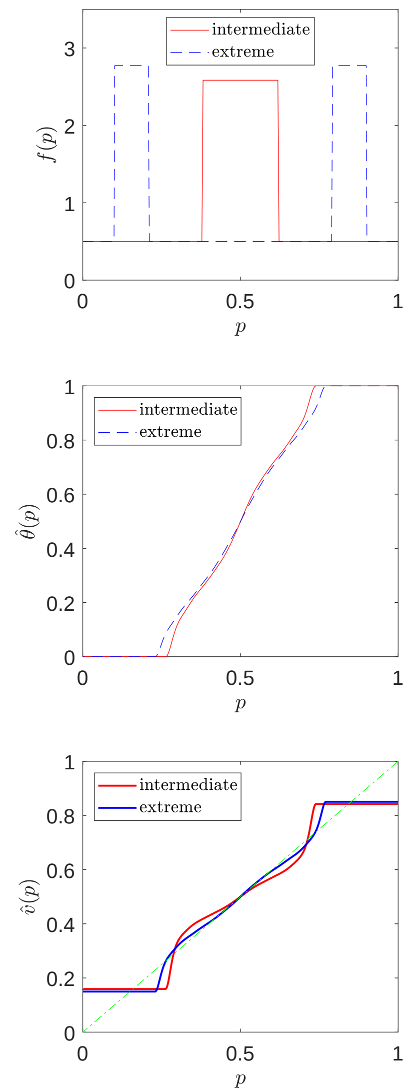

Figure C.1
Prior distribution, coding function, and value function: Efficient coding with an alternative performance objective

The upper graph plots two mixed prior distributions $f(p)$ in the form of (7) described in the main text. For each mixed prior, the first, stable component takes the form of a uniform distribution between 0 and 1. The second, fast-moving component takes the form of (8) for the intermediate prior and the form of (9) in the extreme prior; the parameter values are: $p_l = 0.38$, $p_h = 0.62$, $p_{l,1} = 0.1$, $p_{l,2} = 0.21$, $p_{h,1} = 0.79$, and $p_{h,2} = 0.9$. The weight ξ the DM assigns to the stable component is 0.5. The middle graph plots, for both the intermediate prior and the extreme prior, the coding rule $\hat{\theta}(p)$, numerically solved for maximizing the expected payoff given by equation (C.2). The lower graph plots, for both priors, the subjective valuation $\hat{v}(p)$ implied by $\hat{\theta}(p)$; when computing $\hat{v}(p)$, we set the parameter n to 10. The green dash-dot line is the forty-five degree line.

D. Experimental Designs

Figure D.1
Screenshot of trial

In this example, the risky lottery pays $25 with 53% probability and $0 with 47% probability. The slider position indicates that the subject's reported certainty equivalent on this trial is $10.12.

Table D.1
Design parameters

Adaptation conditions				Test conditions		
Extreme	Low	Intermediate	High	Low	Intermediate	High
10	10	38	67	11	38	69
11	11	39	68	15	42	73
12	12	40	69	19	47	77
13	13	41	70	23	53	81
14	14	42	71	27	58	85
15	15	43	72	31	62	89
16	16	44	73			
17	17	45	74			
18	18	46	75			
19	19	47	76			
20	20	48	77			
21	21	49	78			
79	22	51	79			
80	23	52	80			
81	24	53	81			
82	25	54	82			
83	26	55	83			
84	27	56	84			
85	28	57	85			
86	29	58	86			
87	30	59	87			
88	31	60	88			
89	32	61	89			
90	33	62	90			

Notes. The table provides the specific values of probability (in percentage) used in each of the adaptation and test conditions. Each of the four columns on the left corresponds to an adaptation condition, and it contains 24 distinct values. For each subject, we sample 30 values with replacement from the corresponding adaptation condition (a different sample is drawn for each subject). Each of the three columns on the right corresponds to a test condition, and it contains 6 trials. Each entry in the table denotes a probability p associated with the upside payoff of \$25. Subjects who are randomized into the low test condition are further randomized into either the low or intermediate adaptation condition; subjects who are randomized into the high test condition are further randomized into either the high or intermediate adaptation condition; subjects who are randomized into the intermediate test condition are further randomized into either the extreme or intermediate adaptation condition.

E. Experimental Instructions

Instructions (Page 1 of 3)

Please read these instructions carefully. There will be a short comprehension quiz at the end of the instructions, which you need to pass before continuing to the study.

The study consists of approximately 35 rounds. In each round, you will see a gamble which pays either $25 or $0, each with a different probability. For example, if you see the following gamble:

this means that the gamble pays $25 with 53% chance, and $0 with 47% chance. Importantly, there are no right or wrong answers in this study, please answer all questions depending on your personal opinions.

→

Instructions (page 2 of 3)

On each round, you will be asked to tell us what amount of money -- in your opinion -- the gamble is worth to you. In order to help you form your opinions, think about the amount of money as the maximum price you'd be willing to pay to play the gamble. We will ask you to enter your decision on each round using a slider that goes from $0 (on the left) to $25 (on the right), like in the example below.

Selected Value: **$10.12**

In the above example, the person indicated that a 53% chance of winning $25 is worth $10.12 to them. In other words, the maximum amount they'd be willing to pay to play the gamble is $10.12. We emphasize that different people will give different answers to the same question, since answers depend on your own opinions. There are no right or wrong answers in this survey.

→

Instructions (Page 3 of 3)

10% of all participants will be chosen to receive a bonus, in addition to the $3 for completing the survey.

We have developed the bonus payment method in a way that makes it optimal for you to report your true valuation of the gamble. If you don't care about the exact mechanism that we use to compute the bonus payment, you can feel free to skip the rest of the details on this screen (the comprehension quiz will not ask about how bonuses are computed). If you would like to know exactly how we determine the bonus, then please read below.

Details on how bonus is computed:

As you know, you will see a new gamble on each round and you will be asked to give a certain payment amount that is worth as much as the gamble to you. We will use your response to fill in the table below which actually contains 2,500 questions (for example, see the table below which corresponds to a round in which the gamble offers a 53% chance of winning $25 and 47% chance of winning $0). Question #1 asks if you'd rather have the gamble or $0.01 with certainty. Question #2 asks if you'd rather have the gamble or $0.02 for with certainty. Question #2,500 asks if you'd rather have the gamble or $25.00 with certainty.

Question #		**Option A**		**Option B**
1	Would you rather have:	With probability **53%: Get $ 25** With probability **47%: Get $ 0**	or	$0.01 with certainty
2	Would you rather have:		or	$0.02 with certainty
3	Would you rather have:		or	$0.03 with certainty
4	Would you rather have:		or	$0.04 with certainty
…	…		…	…
2,499	Would you rather have:		or	$24.99 with certainty
2,500	Would you rather have:		or	$25.00 with certainty

Once we get your response on the slider, we will fill in the answer to all 2,500 questions as follows. Suppose your response on the slider is denoted by "x". Then, we will assume you would have chosen Option A for all questions for which the certain amount in Option B is less than or equal to x. We will also assume you would have Chosen Option B for all questions for which the certain amount in Option B is greater than x.

If you are chosen to receive a bonus, then we will then randomly pick one of the 2,500 questions (from one randomly chosen round) and pay you according to what you chose on that one question. Each question and each round is equally likely to be chosen for payment. If you chose Option B on the randomly selected question, you would be paid the certain amount. If you chose Option A on the randomly selected question, then the gamble would be played out by the computer. Using the above table as an example, if you chose Option A then you would receive $25 with 53% chance and $0 with 47% chance.

→

Comprehension Quiz

The following two questions test your understanding of the instructions. Please answer them correctly in order to continue to the survey. Suppose that someone is presented with the gamble below and answers with the following slider location.

Selected Value: **$7.43**

1. Which one of the following statements is correct?

- The person values the gamble at exactly $7.43
- The person values the gamble at more than $7.43
- The person values the gamble at less than $7.43

2. Which one of of the following statements is correct?

- The person would rather have $4.24 with certainty than play the gamble
- The person would rather have $8.43 with certainty than play the gamble
- The person would rather have $8.34 with certainty than play the gamble

→

You have successfully completed the comprehension quiz.

To begin the survey, please click the button below. Remember to answer each question carefully as you might be chosen to receive a bonus according to your answer on that question.

→