

Efficient Coding and Risky Choice*

Cary Frydman[†] and Lawrence J. Jin[‡]

August 7, 2020

ABSTRACT

We experimentally test a theory of risky choice in which the perception of a lottery payoff is noisy due to information processing constraints in the brain. We model perception using the principle of *efficient coding*, which implies that perception is most accurate for those payoffs that occur most frequently. Across two pre-registered laboratory experiments, we manipulate the distribution from which payoffs in the choice set are drawn. In our first experiment, we find that risk taking is more sensitive to payoffs that are presented more frequently. In a follow-up task, we incentivize subjects to classify which of two symbolic numbers is larger. Subjects exhibit higher accuracy and faster response times for numbers they have observed more frequently. In our second experiment, we manipulate the payoff distribution so that efficient coding induces the decision maker's perceived value function to switch from concave to convex. We find that demand for risk is significantly higher when efficient coding induces a convex value function. Together, our experimental results suggest that risk taking depends systematically on the payoff distribution to which the decision maker's perceptual system has recently adapted. More broadly, we provide novel evidence of the importance of imprecise and efficient coding in economic decision-making.

JEL classification: G02, G41

Keywords: efficient coding, perception, risky choice, neuroeconomics

*We thank Nicholas Barberis, Nicola Gennaioli, Alex Imas, Shimon Kogan, John O'Doherty, Stavros Panageas, Antonio Rangel, Andrei Shleifer, Elke Weber, Michael Woodford, and seminar participants at Caltech, the Chinese University of Hong Kong, the National University of Singapore, the Ohio State University, Tsinghua University, the University of Hong Kong, the University of New South Wales, the University of Notre Dame, the University of Technology Sydney, the University of Utah, the University of California, San Diego, the University of Southern California, the University of Texas at Dallas, the University of Warwick, Washington University in St. Louis, Yale University, the Behavioral Economics Annual Meeting, the Chicago Booth Conference in Behavioral Finance and Decision Making, the Kentucky Finance Conference, the LA Finance Day Conference, the NBER Behavioral Finance Meeting, and the Sloan-Nomis Workshop on the Cognitive Foundations of Economic Behavior for helpful comments. Frydman thanks the NSF for financial support.

[†]Marshall School of Business at USC. cfrydman@marshall.usc.edu, +1 (213) 821-5586.

[‡]Division of the Humanities and Social Sciences at Caltech. lawrence.jin@caltech.edu, +1 (626) 395-4558.

I. Introduction

In nearly all economic models of risky choice, the decision maker (henceforth *DM*) is assumed to make a choice based on a precise representation of available lotteries. Yet a large literature in numerical cognition finds that humans perceive numerical quantities with noise, even when the quantities are clearly presented to the *DM* through Arabic numerals (see [Dehaene, 2011](#), for a review). This basic premise leads to the hypothesis, recently proposed by [Khaw, Li, and Woodford \(2020\)](#) (henceforth *KLW*), that risky choice will also be based on a noisy representation of available lotteries. As *KLW* show theoretically, noisy perception of lottery payoffs can provide a microfoundation for small-stakes risk aversion and stochastic choice.

The idea that perceptual noise drives risk aversion has a variety of important but untested implications. For instance, if perceptual noise systematically varies across environments, so should the *DM*'s appetite for risk ([Woodford, 2012a,b](#)). This implication is particularly relevant because there is evidence that noise in perception of sensory stimuli—such as light or sound—changes *optimally* with the environment. Specifically, a core principle from neuroscience called efficient coding, states that the brain should allocate resources so that perception is relatively more precise for those stimuli that are expected to occur relatively more frequently ([Barlow, 1961](#); [Laughlin, 1981](#)).¹ This principle explains the temporary “blindness” that we experience when moving from a dark room to a brightly lit one, because resources have not yet been adjusted for precisely perceiving objects in the new bright environment. If the principle of efficient coding also governs choice under risk, then the *DM*'s perception of a lottery payoff—and hence her appetite for risk—will vary with the environment.

In this paper, we design and conduct two pre-registered experiments to test the hypothesis that efficient coding operates during risky choice. In each experiment, we measure how the demand for a risky lottery varies as we change the payoff distribution to which a subject has recently adapted. To guide our experimental design, we build a theoretical framework that combines principles from two existing models. First, the foundation of our framework is the *KLW* model, which assumes that the *DM* observes noisy signals of lottery payoffs and subsequently forms optimal estimates

¹For experimental evidence consistent with efficient coding in sensory perception, see [Girshick, Landy, and Simoncelli \(2011\)](#), [Wei and Stocker \(2015, 2017\)](#), [Payzan-LeNestour and Woodford \(2020\)](#), and [Heng, Woodford, and Polanía \(2020\)](#). See also the evidence from [Polanía, Woodford, and Ruff \(2019\)](#) on efficient coding in choice between food items.

of these payoffs through Bayesian inference. Second, we rely on the efficient coding model from [Heng, Woodford, and Polanía \(2020\)](#) (henceforth HWP) to endogenize the conditional distribution of noisy signals—which is called the “efficient code.” As in K LW, our framework generates the probability of choosing a risky lottery over a certain option; but, crucially, by adding the efficient coding mechanism from HWP, we can assess how the probability of choosing the risky lottery varies with the environment.

To build intuition, consider a *DM* who is choosing between a binary risky lottery and a certain option. Furthermore, suppose the *DM* is in a low volatility environment where the upside of the risky lottery is drawn from a narrow range between \$15 and \$25. Efficient coding implies that the brain will allocate its limited resources across this narrow range, allowing the *DM* to easily distinguish between payoffs in the range [\$15, \$25]. Suppose now the volatility increases, so that the upside of the risky lottery is drawn from a wider range between \$5 and \$35. Efficient coding then predicts that resources will be partially reallocated *away* from the narrow range and towards the extremes of the new range. As a consequence, it becomes harder for the *DM* to distinguish between payoffs in the range [\$15, \$25] because the perceptual system must “cover more ground” with the same budget of resources.

The shift in perceptual resources immediately leads to a testable prediction about choice. In the low volatility environment, if we increase the upside of the risky lottery from, say, \$20 to \$21, the *DM* will easily be able to distinguish between the two payoffs, and therefore she can easily perceive the increase in the attractiveness of the risky lottery. As a result, the increase in the risky lottery’s upside payoff will have a large impact on the likelihood that the *DM* accepts the risky lottery. Conversely, in the high volatility environment, perceptual resources are spread across a wider range, and therefore the *DM* will have greater difficulty distinguishing between \$20 to \$21. Thus, the same \$1 increase in the risky lottery payoff will have a *smaller* impact on the *DM*’s likelihood of accepting the risky lottery, compared to that in the low volatility environment.

More generally, the efficient coding model of HWP predicts that perception—and hence behavior—is more sensitive to changes in payoff values when the dispersion of potential payoffs is smaller. This prediction is inconsistent with most standard economic models of risky choice in which valuation is non-stochastic and independent of context. At the same time, the prediction is shared by a broad class of theories including the prominent decision-by-sampling model from cognitive sci-

ence (Stewart, Chater, and Brown, 2006), theories of normalization from neuroscience (Rangel and Clithero, 2012; Carandini and Heeger, 2012; Louie, Glimcher, and Webb, 2015),² and alternative specifications of efficient coding (Wei and Stocker, 2015; Khaw et al., 2020).³

In our first experiment, we test the above prediction by incentivizing subjects to make a series of decisions between a risky lottery and a certain option. A novel aspect of our design, which enables a causal test of efficient coding, is that we manipulate the range of payoffs across a high volatility condition and a low volatility condition. In order to cleanly compare behavior, we carefully select a set of thirty “common trials” that are presented in both conditions. Furthermore, a crucial component of our experimental design is that our tests do not depend on whether the *DM*’s objective is to maximize the precision of her payoff estimate (as in models of sensory perception) or to maximize her expected financial gain (as in models of economic decision-making) (Rustichini et al., 2017; Ma and Woodford, 2020).⁴ Thus, the data we produce can be used to simultaneously test different specifications of efficient coding in choice under risk.

The results from our first experiment provide strong evidence that efficient coding influences the demand for risky lotteries. We find that in the low volatility condition, a \$1 increase in the certain option payoff is associated with an 18.6% increase in the probability of choosing the certain option, compared to a smaller increase of 13.7% in the high volatility condition. These estimates are based on the same exact choice sets in each condition, and the effect is significant both between and within subjects. We also find novel evidence that the certain option payoff and the risky lottery payoff are efficiently coded separately, suggesting that in our experiment, the mechanism operates at the level of individual payoffs. Furthermore, subjects execute decisions significantly faster in the low volatility condition, indicating that the effect of efficient coding on choice is even stronger if it were adjusted for information processing time.

²Several experiments have found evidence consistent with normalization of value signals in the brain (e.g., Tobler, Fiorillo, and Schultz, 2005; Padoa-Schoppa, 2009). For behavioral evidence consistent with normalization, see Soltani, De Martino, and Camerer (2012), Khaw, Glimcher, and Louie (2017), and Zimmermann, Glimcher, and Louie (2018). Recent behavioral economic theories also invoke normalization to explain several prominent patterns of context dependent choice (Glimcher and Tymula, 2019; Landry and Webb, 2019).

³See also related theoretical work on how efficient coding can serve as a normative foundation for decision-by-sampling (Bhui and Gershman, 2018; Heng et al., 2020) and a specific model of normalization (Rustichini, Conen, Cai, and Padoa-Schioppa, 2017).

⁴As emphasized in Ma and Woodford (2020), there are differences in the way that resource constraints are imposed across different models of efficient coding. While we use a specific constraint levied by the Heng et al. (2020) model, the main prediction we test is qualitatively similar to other models of efficient coding in sensory perception that assume different constraints, such as Wei and Stocker (2015).

As an additional test of the mechanism, we present each subject with a “perceptual choice” task following the risky choice task. In this second task, subjects still need to perceive numerical quantities, but they do not need to perceive any probabilities or integrate them with payoffs. We incentivize subjects to simply classify whether a two-digit number displayed on each trial is above or below a reference number. We find that even in this simpler environment, classification accuracy depends strongly on the distribution of numbers to which the subject has adapted. Subjects are significantly more accurate and they respond faster if the number they are classifying was presented more frequently in the recent past.

We then structurally estimate our efficient coding model for each subject, separately for the risky choice and perceptual choice tasks. The estimation indicates that the parameter governing noisy perception is positively correlated across the two tasks, with a rank correlation of 0.30. This finding suggests that risk taking in the first task is partly governed by a perceptual mechanism that can be inferred from riskless choice. Our cross-task evidence complements recent work by [Garcia, Grüschow, Polanía, Woodford, and Ruff \(2018\)](#) who find that a subject’s demand for risk can be predicted from a separate task on numerosity comparison (in which subjects are asked to compare two arrays of dots, rather than two Arabic numerals as in our design).⁵

When viewed through the lens of efficient coding, the results from our first experiment are driven by an increase in noise when shifting from a low to high volatility condition. But efficient coding also predicts strong *biases* in valuation. As emphasized by [Woodford \(2012a,b\)](#), efficient coding can theoretically generate a value function that exhibits several features from prospect theory, including reference dependence and diminishing sensitivity ([Kahneman and Tversky, 1979](#)). Moreover, these features can fluctuate over time as an optimal response to changes in the environment. One particularly stark prediction is that a payoff distribution with a decreasing density will generate a concave value function, while a distribution with an increasing density will generate a convex value function. Both shapes of the value function are a consequence of the perceptual system exhibiting diminishing sensitivity as payoffs become less frequent ([Robson, 2001](#); [Rayo and Becker, 2007](#); [Netzer, 2009](#); [Payzan-LeNestour and Woodford, 2020](#)).⁶

⁵[Schley and Peters \(2014\)](#) also provide evidence of a link between risk taking and imprecise number representations. However, they emphasize that the link between the two domains emanates from a *deterministic* perceptual bias, whereas our parameter identification relies heavily on the stochasticity of perception, which endogenously generates the biases in perception.

⁶For work on how neural computations can generate such an effect, see [Louie et al. \(2015\)](#) and [Webb, Glimcher,](#)

In our second experiment, we test whether risk taking is greater when the *DM* has adapted to an increasing distribution compared to a decreasing distribution. Across two experimental conditions, we manipulate the *shape* of the payoff distribution while holding constant its range. When analyzing identical choice sets across the two conditions, we find evidence consistent with a systematic bias in valuation. As predicted by efficient coding, the probability of choosing the risky lottery is significantly higher when subjects are adapted to the increasing distribution, compared to the decreasing distribution. Intuitively, when the subject is adapted to an increasing distribution, perceptual resources are allocated towards higher payoffs; this resource allocation generates increasing sensitivity to larger payoffs, and thus a convex value function. The convexity gives rise to risk-seeking behavior, and thus a greater demand for risk, compared to the situation in which the *DM* is adapted to a decreasing distribution. Our results are consistent with those from a recent perceptual choice experiment by [Payzan-LeNestour and Woodford \(2020\)](#), who find that “outlier” stimuli are perceived less accurately than more frequently occurring stimuli. More generally, we provide novel evidence consistent with a prominent hypothesis that diminishing sensitivity to payoffs arises, in part, from an optimal allocation of perceptual resources.

We emphasize that the existing evidence of efficient coding, which comes almost exclusively from data on sensory perception, in no way implies that the same mechanisms are deployed during decision-making under risk. Rather, what we test in this paper is precisely the hypothesis that efficient coding and noisy perception are also active in higher-level decision systems that govern risky choice. Our experimental evidence therefore supports a nascent theoretical agenda on the implications of imprecise and efficient coding for economic behavior ([Woodford, 2012a,b](#); [Steiner and Stewart, 2016](#); [Gabaix and Laibson, 2017](#); [Natenzon, 2019](#); [Khaw et al., 2020](#); [Woodford, 2020](#)).

At a broader level, our results contribute to a growing literature that builds cognitive and perceptual foundations for the psychological assumptions in behavioral economics. For instance, several behavioral models of financial markets demonstrate that prospect theory preferences can explain puzzling facts such as the high equity premium of the aggregate stock market (see [Barberis, 2018](#), for a review). Our results provide novel empirical evidence consistent with the proposition that efficient coding provides a normative foundation for the value function assumed in prospect

and Louie (2020).

theory (Woodford, 2012a,b).⁷ As we show in Section IV, the particular efficient coding rule we use can also generate the probability weighting function from prospect theory, which has previously been microfounded on principles of salience (Bordalo, Gennaioli, and Shleifer, 2012), earlier models of noisy coding of probabilities (Steiner and Stewart, 2016; Khaw et al., 2020), and cognitive uncertainty (Enke and Graeber, 2020).

The rest of the paper is organized as follows. In Section II, we present a theory of efficient coding that guides our experimental design. Section III provides experimental tests of the model by studying how the payoff distribution affects choice. Section IV provides additional discussions and Section V concludes.

II. The Model

In this section, we present a theory of efficient coding and risky choice that integrates two existing theoretical models. The foundation of our theory is the Khaw et al. (2020) model of noisy perception of lottery payoffs. In their baseline model, KLW assume a particular form of noisy coding of lottery payoffs and a specific form of the prior, which they use to derive a series of novel implications for risky choice. We build on KLW by integrating it with the Heng et al. (2020) model of efficient coding. By combining these two models, we are able to generate predictions about how the noisy coding rule—and the probability of risky choice—systematically changes for any payoff distribution to which the *DM* has adapted.⁸

1. Choice environment

The *DM* faces a choice set that contains two options: a certain option and a risky lottery. The certain option, denoted as $(C, 1)$, pays $C > 0$ dollars with certainty. The risky lottery, denoted as $(X, p; 0, 1 - p)$, pays $X > 0$ dollars with a probability of p and zero dollars with the remaining probability of $1 - p$. The *DM*'s task is to choose between these two options.

Under expected utility, a *DM* with utility $U(\cdot)$ chooses the risky lottery over the certain option

⁷For alternative approaches to endogenizing the value function in prospect theory, see Friedman (1989) and Denrell (2015).

⁸KLW also provide an extension to their baseline model in which the precision of the noisy coding rule can flexibly change with the volatility of a particular (lognormal) prior distribution. Our framework further generalizes this flexibility by allowing optimal coding rules for any prior distributions.

if and only if

$$p \cdot U(X) + (1 - p) \cdot U(0) > U(C). \quad (1)$$

Conditional on X , C , and p , the *DM*'s choice is non-stochastic.

We now present the KLV framework of noisy coding, which departs from expected utility by assuming that perceptions of X and C are noisy. The noisy coding assumption is motivated by the literature in sensory perception, where a common finding is that, when an identical stimulus is presented on different occasions (e.g., a fixed number of dots presented across different trials of an experiment), experimental subjects judge the stimulus differently across the different occasions (Dehaene, 2011; Girshick et al., 2011; Wei and Stocker, 2015, 2017).⁹ Before observing the choice set containing X and C , the *DM* holds prior beliefs about each payoff, given by $f(X)$ and $f(C)$, respectively. Upon observing the choice set, the *DM*'s perceptual system spontaneously generates a noisy signal, R_x , of X , and a noisy signal, R_c , of C .¹⁰ Each of the two noisy signals is drawn from a distinct conditional distribution: R_x is randomly drawn from $f(R_x|X)$, and R_c is randomly drawn from $f(R_c|C)$. In the language of Bayesian inference, these conditional distributions are the likelihood functions, which we define in the next subsection.

Given the prior beliefs and the noisy signals, the *DM* follows Bayes' rule to generate a posterior distribution for each payoff, which we denote by $f(X|R_x)$ and $f(C|R_c)$. The *DM*'s estimate of each payoff is given by the conditional mean of each posterior distribution: $\mathbb{E}[\tilde{X}|R_x]$ and $\mathbb{E}[\tilde{C}|R_c]$. As in KLV, we assume that the *DM* has linear utility.¹¹ Therefore, the *DM* chooses the risky lottery if and only if the *perceived* expected value of the risky lottery exceeds that of the certain option, which is given by the following condition: $p \cdot \mathbb{E}[\tilde{X}|R_x] > \mathbb{E}[\tilde{C}|R_c]$.¹²

⁹Further evidence for this assumption, particularly in the domain of numerical cognition, comes from recent experimental work which demonstrates that single neurons in the human brain selectively and stochastically respond to a given number (Kutter, Bostroem, Elger, Mormann, and Nieder, 2018). Such “number neurons” are likely to generate the noisy perception of symbolic numbers. For a more comprehensive review of the literature in noisy numerical cognition, see Section 1 of KLV.

¹⁰In the environment of choice between lotteries that we study here, our interpretation is that the noisy signals are generated unconsciously, and are not the outcome of deliberate and conscious information acquisition in the sense of Stigler (1961). For a further discussion on this point, see Ma and Woodford (2020).

¹¹Here we focus on how imperfect perception—rather than intrinsic risk preferences—affects risky choice. We emphasize that the model here does not preclude the more traditional source of risk aversion, through diminishing marginal utility of wealth. In Section III.1.5, we estimate an extended model by allowing for both imperfect perception and intrinsic risk aversion.

¹²We assume for simplicity that (i) the probability p is perceived without noise, and (ii) the probability p is integrated with $\mathbb{E}[\tilde{X}|R_x]$ without noise. In Section III.1, we conduct an experiment on number classification that provides a sharp test of the efficient coding hypothesis without appealing to these two assumptions. In Section IV.3, we further discuss our model's implications for the imperfect perception of probability.

It is worth noting that the encoding process described above—the process that maps X and C to R_x and R_c —is conditional on the values of X and C , which we assume are perfectly observable to the econometrician but not to the DM . That is, even after the DM is presented with a choice set, she still faces uncertainty about the payoff values of X and C . Therefore, Bayesian inference takes place at the level of a single choice set, and characterizes how the DM 's prior belief shifts after observing a noisy signal of the true payoff. The noisy encoding of payoffs drives the main model predictions.¹³

2. Likelihood function and coding optimality

Here we depart from K LW by allowing the likelihood functions to vary for any prior distribution. In this case, the DM still encodes each payoff with noise, but the noise distribution can be optimized to meet a specific performance objective. We rely on the efficient coding model of HWP to endogenize the likelihood functions $f(R_x|X)$ and $f(R_c|C)$. Although there are alternative specifications for the efficient likelihood functions, an attractive feature of HWP is that it is general enough to accommodate multiple different performance objectives—for instance, maximizing expected financial gain or maximizing the mutual information between a stimulus and its noisy signal. This feature is important because it allows us to identify an environment in which predictions are *invariant* to the DM 's performance objective. Such an environment is ultimately used as the basis for our experimental design in Section III.

When maximizing a given performance objective, the DM faces an information processing constraint. HWP assume that information can only be processed through a finite number of n “neurons,” where the output state of each neuron takes the value 0 or 1. The output states of these n neurons are assumed to be mutually independent, with each neuron taking the value 1 with the probability $\theta(X)$ and taking the value 0 with the remaining probability $1 - \theta(X)$. Given that the neurons are mutually independent, a sufficient statistic for the output of the n neurons is the summed output across all neurons, which we denote by R_x (the same is true for R_c). Thus, the

¹³While our focus is on choice between lotteries with non-negative payoffs, the K LW theory can easily be extended to account for negative payoffs. Appendix D of K LW discusses implications of the theory when the magnitude of a loss is encoded with noise, and the DM chooses the lottery with the smallest perceived expected loss.

noisy signal, R_x , can take on integer values from 0 to n , and the likelihood function is defined by

$$f(R_x|X) = \binom{n}{R_x} (\theta(X))^{R_x} (1 - \theta(X))^{n-R_x}. \quad (2)$$

The one free parameter in the encoding process, n , represents a capacity constraint: as n goes to infinity, the random variable R_x/n converges almost surely to its mean $\theta(X)$, and therefore the amount of noise in perceiving X is reduced to zero.

The driving force of the likelihood function is the coding rule, $\theta(X)$, which maps the input value X to the probability that a given neuron emits a value of 1. Intuitively, if the *DM* is particularly concerned about perceiving values of X within a given range, then a good coding rule, $\theta(X)$, should be very sensitive to X over that range.

2.1. Maximizing mutual information

To formalize this notion, HWP consider three different performance objectives. The first objective, which is most often associated with the efficient coding literature in sensory perception, is that the *DM* maximizes mutual information between X and its noisy signal R_x ¹⁴

$$\max_{\theta(X)} I(X, R_x), \quad (3)$$

where the mutual information $I(X, R_x)$ is defined as $I(X, R_x) = H(R_x) - H(R_x|X)$, that is, the difference between the marginal entropy of R_x and the entropy of R_x conditional on X .

HWP show that, when n is large, the optimal coding rule that maximizes the mutual information in (3) is

$$\theta(X) = \left[\sin \left(\frac{\pi}{2} F(X) \right) \right]^2, \quad (4)$$

where $F(X)$ is the cumulative density function of the *DM*'s prior belief $f(X)$. Substituting (4) into (2) therefore pins down the ‘‘efficient’’ likelihood function when the *DM*'s objective is to maximize mutual information.¹⁵

¹⁴For instance, see [Wei and Stocker \(2015, 2017\)](#).

¹⁵As emphasized in HWP, the optimal coding rule $\theta(X)$ is not constrained to take on values in the *open* interval $(0, 1)$. In fact, given the constraint that all information about a payoff must be processed through a set of statistically independent neurons, a deterministic mapping from X to R_x is also feasible. Nonetheless, HWP show that the optimal solution implies that the encoding of value is indeed *stochastic*; that is, for some values of X , $0 < \theta(X) < 1$. The intuition for this result can be briefly described as follows: a deterministic coding rule for each neuron implies

Because we assume that X and C are encoded independently, the optimal coding rule for maximizing the mutual information between C and its noisy signal R_c takes a similar form:

$$\theta(C) = \left[\sin \left(\frac{\pi}{2} F(C) \right) \right]^2, \quad (5)$$

where $F(C)$ is the cumulative density function of the DM 's prior belief $f(C)$.

2.2. Maximizing accuracy and maximizing expected financial gain

HWP consider two alternative performance objectives besides maximizing mutual information. First, they consider how the coding rule would change if the DM were to maximize the probability of an accurate choice, given by:

$$\begin{aligned} \text{Prob}_{\text{accurate}} \equiv & \iint \text{Prob}(R_x > R_c | \theta(X) > \theta(C)) \mathbb{1}_{\theta(X) > \theta(C)} f(X, C) dX dC \\ & + \iint \text{Prob}(R_x < R_c | \theta(X) < \theta(C)) \mathbb{1}_{\theta(X) < \theta(C)} f(X, C) dX dC, \end{aligned} \quad (6)$$

where $f(X, C)$ is the DM 's prior belief about the joint distribution of X and C .¹⁶

An alternative performance objective, and one that is often used in economic settings, is to maximize the DM 's expected financial gain, defined as:

$$\begin{aligned} \text{Expected financial gain} \equiv & \iint pX \cdot \text{Prob}(\text{choose the risky lottery} | X, C) f(X, C) dX dC \\ & + \iint C \cdot \text{Prob}(\text{choose the certain option} | X, C) f(X, C) dX dC. \end{aligned} \quad (7)$$

HWP show that, in general, the optimal coding rule implied by these three different objectives will

that the noisy signal R_x can take on only two values: 0 or n . By allowing for noisy encoding, the set of possible values of R_x expands dramatically to include all integers between 0 to n , which helps increase the mutual information $I(X, R_x)$. Thus, the noise in perception can be interpreted as an optimal response to the particular information processing constraint assumed in HWP. If the information constraint is modified, then it may be optimal to have a deterministic coding rule (Bhui and Gershman, 2018).

¹⁶Formally, an accurate decision, given X and C , is to choose the option with the higher expected payoff observed by the econometrician. That is, when $pX - C > 0$, the DM chooses the risky lottery; when $pX - C < 0$, the DM chooses the safe option; and when $pX - C = 0$, the DM randomly chooses one of the two options. As discussed in Appendix A, the probability of this accurate decision equals the probability that $R_x - R_c$ and $\theta(X) - \theta(C)$ are of the same sign.

be distinct. Importantly, we prove in Appendix A that, when the following conditions

$$\begin{aligned} & (i) \text{ draws of } X \text{ and } C \text{ are independent} \\ & \text{and } (ii) \text{ } pX \text{ and } C \text{ are identically and uniformly distributed} \end{aligned} \tag{8}$$

are satisfied, all three coding rules—one that maximizes mutual information (equation (3)), one that maximizes accuracy (equation (6)), and one that maximizes expected payoff (equation (7))—are equivalent.¹⁷

As mentioned above, this equivalence result is useful because it characterizes a class of environments in which the implications of efficient coding are invariant to the three objectives considered above. In our main experiment, we build a design that satisfies the conditions in (8) in order to provide general tests of efficient coding. Thus, the results we present in the following subsections assume the *DM*'s objective is to maximize mutual information.

2.3. Properties of likelihood function

Here we illustrate how the likelihood function depends explicitly on the *DM*'s prior beliefs. For now, we suppose that the *DM*'s prior belief about X is a uniform distribution between X_l and X_u ($X_l < X_u$), and that her prior belief about C is a uniform distribution between C_l and C_u ($C_l < C_u$). In keeping with the conditions in (8), we further set $C_l = p \cdot X_l$ and $C_u = p \cdot X_u$, so that pX and C are identically distributed. Given these assumptions, the likelihood functions of X and C are

$$\begin{aligned} f(R_x|X) &= \binom{n}{R_x} \left(\left[\sin \left(\frac{\pi}{2} \frac{X - X_l}{X_u - X_l} \right) \right]^2 \right)^{R_x} \left(1 - \left[\sin \left(\frac{\pi}{2} \frac{X - X_l}{X_u - X_l} \right) \right]^2 \right)^{n-R_x}, \\ f(R_c|C) &= \binom{n}{R_c} \left(\left[\sin \left(\frac{\pi}{2} \frac{C - C_l}{C_u - C_l} \right) \right]^2 \right)^{R_c} \left(1 - \left[\sin \left(\frac{\pi}{2} \frac{C - C_l}{C_u - C_l} \right) \right]^2 \right)^{n-R_c}. \end{aligned} \tag{9}$$

The expressions in (9) show that, with a finite n , the likelihood functions depend directly on the parameters of the prior distributions, X_l , X_u , C_l , and C_u .¹⁸ This dependence of the likelihood

¹⁷The coding rules described above are derived when n , the parameter for the capacity constraint, is sufficiently large. With smaller values of n , say, $n = 10$, the coding rules are approximately optimal; for more discussion of this point, see Appendix 7 of HWP. When illustrating the model's implications, we set n to 10.

¹⁸As n goes to infinity, the relation between the "normalized" signal (R_x/n or R_c/n) and the stimulus values (X or C) becomes a deterministic one-to-one mapping. In this limiting case, the parameter values of X_l , X_u , C_l , and

function on the prior is a signature characteristic of efficient coding.

[Place Figure 1 about here]

Figure 1 illustrates the malleability of the likelihood function. Panel A presents two different prior distributions over X , one with high volatility and the other with low volatility. In the high volatility environment, X is distributed uniformly over a wide range ($X_l = 8$ and $X_u = 32$), whereas in the low volatility environment, X is distributed uniformly over a narrow range ($X_l = 16$ and $X_u = 24$). As described by (4), these two distributions induce different coding rules $\theta(X)$. Panel B of Figure 1 shows that the coding rule is steeper for the low volatility distribution, compared to the high volatility distribution. Recall that the coding rule gives the “success probability” of the binomial distribution in (2). Thus, a steeper coding rule implies that the success probability is more sensitive to changes in X . Panel C plots the implied likelihood function for two values, $X = 18$ and $X = 22$, and for each of the two prior distributions. In the low volatility environment, a payoff of $X = 18$ generates a very different distribution of signals $f(R_x|X)$ compared to a payoff of $X = 22$. In the high volatility distribution, however, $X = 18$ and $X = 22$ generate distributions of signals that overlap extensively. The more extensive overlap of the likelihood functions in the high volatility environment leads to less discriminability between the two payoffs, compared to the low volatility environment. As we show in the next subsection, this difference in discriminability has a direct impact on risky choice.

3. Value function and implications for choice

Given the prior and likelihood functions defined above, the *DM* proceeds by using Bayesian inference to compute a posterior distribution of beliefs about each payoff in the choice set. We assume that the *DM* uses the mean of the posterior distribution as her estimate of each payoff. Specifically, the posterior means of X and C , conditional on R_x and R_c , are given by

$$\mathbb{E}[\tilde{X}|R_x] \equiv \frac{\int_{X_l}^{X_u} f(R_x|X)f(X)XdX}{\int_{X_l}^{X_u} f(R_x|X)f(X)dX} \quad (10)$$

C_u do not affect the perceptions of X and C .

and

$$\mathbb{E}[\tilde{C}|R_c] \equiv \int_{C_l}^{C_u} f(C|R_c)C dC = \frac{\int_{C_l}^{C_u} f(R_c|C)f(C)C dC}{\int_{C_l}^{C_u} f(R_c|C)f(C) dC}, \quad (11)$$

where $f(X)$ and $f(C)$ are the *DM's* prior beliefs about X and C , and the likelihood functions $f(R_x|X)$ and $f(R_c|C)$ are from (9).

Importantly, equation (10) shows that the *DM's* estimate of X is a random variable, where the randomness comes from R_x . Therefore, the *DM* faces a *distribution* of perceived values for each X . We now characterize the mean and standard deviation of this distribution. Specifically, we define the value function, $v(X)$, by

$$v(X) = \sum_{R_x=0}^n \mathbb{E}[\tilde{X}|R_x] \cdot f(R_x|X), \quad (12)$$

where $\mathbb{E}[\tilde{X}|R_x]$ is from (10) and $f(R_x|X)$ is from (9). That is, $v(X)$ represents the subjective valuation of X averaged across different values of R_x . Moreover, we define the standard deviation for the subjective valuation, $\sigma(X)$, by

$$\sigma(X) = \left[\sum_{R_x=0}^n (\mathbb{E}[\tilde{X}|R_x])^2 f(R_x|X) - v(X)^2 \right]^{1/2}. \quad (13)$$

Equations (12) and (13), together with equations (4) and (9), indicate that the curvature of the value function and the randomness in subjective valuation are jointly determined by the *DM's* prior belief and the implied likelihood function.

[Place Figure 2 about here]

In keeping with the running example from the previous subsection, Panel A of Figure 2 plots, for both the high volatility condition ($X_l = 8$ and $X_u = 32$) and the low volatility condition ($X_l = 16$ and $X_u = 24$), the average subjective valuation $v(X)$, as well as its one-standard-deviation bounds $v(X) \pm \sigma(X)$.

The figure shows that randomness in utility, $\sigma(X)$, is substantially higher in the high volatility condition. This is a result of the fact that the likelihood functions in the high volatility condition exhibit more overlap with each other, compared to those in the low volatility condition. Because

subjective valuation is noisier in the high volatility condition, the model predicts that choices will also be noisier, and hence less sensitive to a given change in payoff amounts.

To see this explicitly, we compute the probability of choosing the risky lottery—which we refer to from now on as the “probability of risk taking.” To do so, recall that, conditional on X and C , the noisy signals R_x and R_c are drawn from the likelihood functions $f(R_x|X)$ and $f(R_c|C)$. For a given realization of (R_x, R_c) , the *DM* then chooses between the risky lottery and the certain option based on the posterior means of X and C in equations (10) and (11). As a result, when fixing X , C , and the stimulus distributions, we can compute the probability of risk taking as follows:

$$\begin{aligned} \text{Prob}(\textit{risk taking}|X, C) &= \sum_{R_x=0}^n \sum_{R_c=0}^n \left(\mathbb{1}_{p \cdot \mathbb{E}[\tilde{X}|R_x] > \mathbb{E}[\tilde{C}|R_c]} \cdot f(R_x|X) \cdot f(R_c|C) \right) \\ &+ \sum_{R_x=0}^n \sum_{R_c=0}^n \left(\mathbb{1}_{p \cdot \mathbb{E}[\tilde{X}|R_x] = \mathbb{E}[\tilde{C}|R_c]} \cdot \frac{1}{2} f(R_x|X) \cdot f(R_c|C) \right). \end{aligned} \quad (14)$$

Equation (14) says that when $p \cdot \mathbb{E}[\tilde{X}|R_x] > \mathbb{E}[\tilde{C}|R_c]$, the *DM* chooses the risky lottery over the certain option; and when $p \cdot \mathbb{E}[\tilde{X}|R_x] = \mathbb{E}[\tilde{C}|R_c]$, the *DM* randomly chooses between these two options.

[Place Figure 3 about here]

Panel A of Figure 3 plots the probability in (14) against the difference in expected values between the two options, $pX - C$. As before, we examine two volatility conditions: in the high volatility case, we set $X_l = 8$, $X_u = 32$, $C_l = 4$, and $C_u = 16$; and in the low volatility condition, we set $X_l = 16$, $X_u = 24$, $C_l = 8$, and $C_u = 12$. We set p , the probability that the risky lottery pays X dollars, to 0.5; and we set n , the capacity constraint parameter, to 10. Note that, for each volatility level, the conditions in (8) are satisfied, so the coding rules $\theta(X)$ and $\theta(C)$ simultaneously maximize mutual information, the probability of accuracy, and the *DM*’s expected financial gain.

Naturally, a higher value of $pX - C$ increases the attractiveness of the risky lottery and hence increases the probability of risk taking. Note that, under expected utility and with no background wealth, the probability of risk taking should be a step function of $pX - C$ with a single step at $pU^{-1}((U(C) - (1-p)U(0))/p) - C$. However, Panel A of Figure 3 shows that under noisy coding, the probability of risk taking has an *S*-shaped relationship with $pX - C$. Importantly, under *efficient*

coding, the overall slope of this function is negatively related to the volatility of the stimulus distribution. Thus, for a given increase in the payoff X , the probability of choosing the risky lottery increases more in the low volatility condition. This difference in sensitivity to a lottery payoff is inherited directly from the property of the likelihood functions (Panel C of Figure 1), where a given increase in X leads to a larger difference in the distribution of noisy signals in the low volatility condition, compared to the high volatility condition.

Panel A of Figure 3 shows how risk taking varies with $pX - C$. However, an important assumption in the framework we present is that X and C are encoded independently; thus risk taking should also be more sensitive to X and C separately, when they are drawn from the low volatility condition, compared to the high volatility condition. Panel B of Figure 3 demonstrates that the model indeed predicts that X and C should each be efficiently coded separately. In Section III.1.2, we discuss experimental results that are consistent with this notion of “parallel coding.”

4. Increasing and decreasing payoff distributions

So far, we have focused on how payoff volatility affects the dispersion in perception, $\sigma(X)$. However, efficient coding can also induce *biases* in perception, in the sense that $v(X)$ can differ significantly from X (Woodford, 2012a,b; Khaw et al., 2020). Panel A of Figure 2 indicates that with a uniform distribution, $v(X)$ largely coincides with X , although there is evidence of minimal perceptual biases at the extreme ends of each payoff distribution.

Under other payoff distributions, however, efficient coding can induce sizeable perceptual biases. To illustrate how these stronger perceptual biases arise, we consider a different environment in which the payoff distribution of C remains uniform, while the payoff distribution of X is either monotonically increasing or monotonically decreasing. The implications of a monotonically decreasing distribution are particularly interesting, because there is evidence that the naturally occurring distribution of numerical quantities—i.e., the frequency with which number words are printed in written text—follows a monotonically decreasing distribution (Dehaene and Mehler, 1992). Specifically, we set

$$f(X; X_l, X_u, \Delta) = \frac{1}{X_u - X_l} - \Delta + 2\Delta \cdot \frac{X - X_l}{X_u - X_l}, \quad (15)$$

where Δ is a slope parameter: a positive Δ implies an increasing distribution, and a negative Δ implies a decreasing distribution.

Panel B of Figure 2 plots, for both an increasing prior distribution ($X_l = 8$, $X_u = 32$, and $\Delta = 1/30$) and a decreasing prior distribution ($X_l = 8$, $X_u = 32$, and $\Delta = -1/30$), the average subjective valuation $v(X)$, as well as its one-standard-deviation bounds $v(X) \pm \sigma(X)$. For an increasing payoff distribution, the value function is convex overall. In this case, small values of X occur with low frequency, causing the *DM* to allocate less coding resources towards these small values. Thus, the lack of discriminability among small values of X gives rise to a positive perceptual bias: $v(X)$ is greater than X when X is small. Conversely, for a decreasing payoff distribution, the *DM* allocates less coding resources towards large values of X . As a result, the *DM* exhibits a negative perceptual bias: $v(X)$ is less than X when X is large, and hence the implied value function is concave overall.

We test implications of these value functions after presenting our main experiment. In particular, we test whether demand for risk taking is higher in the increasing condition where the prior induces a convex value function, compared to the decreasing condition. In the next section, we turn to our experimental tests of efficient coding.

III. Experimental Tests

In this section, we provide experimental tests of the model. We first examine how the volatility of the payoff distribution affects risky choice. We then analyze how the shape of the payoff distribution affects risky choice.

1. *Experiment 1: Volatility manipulation*

Our first experiment tests the model by manipulating the volatility of payoffs. We pre-register the experiment and recruit 150 students from the University of Southern California to participate in the laboratory (see Appendix B for the pre-registration document). Each subject completes two tasks: a risky choice task and a perceptual choice task. All subjects complete the risky choice task first, followed by the perceptual choice task. We chose this ordering to minimize any fatigue effects in the risky choice task, which is our main task of interest. Subjects were paid a \$7 show-up fee,

in addition to earnings from each task.

1.1. Design of the risky choice task

On each trial, subjects choose between the risky lottery $(X, p; 0, 1 - p)$ and the certain option $(C, 1)$. The probability p is fixed at 0.5 for all trials. The values of X and C are drawn independently, and we manipulate the distribution of each payoff across two volatility conditions. In the high volatility condition, X is drawn uniformly from $[8, 32]$, and C is drawn uniformly from $[4, 16]$. In the low volatility condition, X is drawn uniformly from $[16, 24]$, and C is drawn uniformly from $[8, 12]$.

We chose the above design parameters for two reasons. First, because our goal is to isolate the effect of volatility, we chose parameters such that the mean of each payoff distribution is constant across conditions; the mean of X is fixed at 20 and the mean of C is fixed at 10. Second, our parameter values satisfy the conditions in (8): the distributions of X and C are independent, and pX and C are identically and uniformly distributed. These conditions imply that the efficient coding rules are the same, regardless of whether the subject’s objective is to maximize mutual information, the probability of an accurate choice, or the expected financial gain. This is an important feature of our design. Although many tests of efficient coding in sensory perception assume that the *DM* maximizes mutual information, it is not obvious whether the *DM* has the same objective in economic choice (Ma and Woodford, 2020). Thus, our design is optimized to test generic predictions of efficient coding.

[Place Figure 4 about here]

Figure 4 shows a schematic of the task design. Each subject goes through both the high and low volatility conditions; the order of the two conditions is randomized across subjects. In each condition, there are 300 trials, which are broken into two phases: an initial “adapt” phase (30 trials) and a subsequent “test” phase (270 trials). The adapt phase is intended to allow the subjects to adapt to the condition-specific payoff distribution; the test phase contains the trials that we are interested in analyzing.¹⁹ While such a large number of trials is not typical in economics

¹⁹As argued in Woodford (2020), “Of course, the appropriate prior has to be learned; one should therefore expect the mapping from objective magnitudes to estimates to shift over the course of an experimental session, especially at the beginning.” (footnote 21). Our design explicitly builds in an explicit “initial adaptation” phase for which we are

experiments, this design feature is important for providing a clean test of our hypothesis.

Specifically, in order to generate a clean test of efficient coding, we want to compare decisions on the *same* choice sets across the two volatility conditions. One constraint we face, when designing these “common trials” in the test phase, is that the lottery payoffs must fall in the support of the distribution of both conditions. Our goal is to maximize the number of common trials that satisfy this constraint, while staying faithful to the statistical properties of each payoff distribution.

To do so, we first note that the support of the low volatility distribution is a subset of the support of the high volatility distribution; as a result, payoffs on common trials must fall in the support of the low volatility distribution. Specifically, with $1/9$ probability, a pair (X, C) drawn from the high volatility distribution falls in the support of the low volatility distribution. Therefore, in each condition, we designate 30 of the 270 trials in the test phase as common trials. These common trials are identical across conditions, and we generate them by drawing 30 pairs of (X, C) over approximately equally-spaced grid points of the low volatility distribution (see Table 1 for exact values). At the subject level, the location of each common trial is randomized across the 270 possible test trial locations in each condition.

[Place Table 1 about here]

We then draw the remaining 240 test trials in the low volatility condition from the low volatility distribution. For the remaining 240 test trials in the high volatility distribution, we draw (X, C) uniformly from the high volatility distribution, but critically, we “re-draw” the pair (X, C) if it falls in the support of the low volatility distribution—since this part of the high volatility distribution is already covered by the common trials. As a result, when combined with the 30 common trials, the distribution of payoffs across all trials in each volatility condition accurately reflects the appropriate population distribution. In sum, common trials simultaneously serve two purposes: they allow for a clean comparison of behavior across conditions while also reinforcing the prior on subsequent trials.

Subjects are not explicitly informed about the payoff distributions from which X and C are drawn. We believe that such a design is more natural than telling subjects the payoff distributions that they will experience. In particular, if the experimenter explicitly gives information about

less interested in analyzing behavior.

the distribution of payoffs (which is rarely given in experiments on risky choice), subjects may artificially direct their attention to this information. This, in turn, could generate an experimenter demand effect that threatens the interpretation of our efficient coding tests. Furthermore, our design choice enables us to test for learning effects, which are important when conducting our within-subject analyses.

One of the six-hundred trials was randomly selected for payment and the subject was paid according to their choice on the randomly selected trial. The average earning for the risky choice task was \$10.14. The exact instructions that were given to subjects before the experiment are provided in Appendix B.

1.2. Results from the risky choice task

Reduced form results

We produce a large dataset that contains 90,000 total observations across all subjects and conditions (600 observations per subject). Before proceeding with our analyses, we take two steps to clean the data. First, as part of our pre-registered data exclusion, we drop one subject who chose the certain option on all trials in the first condition. Second, we drop trials for which subjects exhibit an excessively fast response time of less than 0.5 seconds, which constitutes 7.7% of the data. This second step in the data cleaning process was not pre-registered, but unsurprisingly these fast decisions are not responsive to the underlying payoff values. As such, they only add noise to our tests of efficient coding.²⁰ After applying the two data exclusions, we are left with 82,851 trials across the first and second conditions, of which 74,277 trials are in the test phase.

We begin our analysis with between subjects tests by using all trials in the test phase of the first condition—in which subjects are randomized into either the high or low volatility condition. We find that, on average, subjects choose the risky option on 45.8% of trials (standard error: 1.6%). Panel A of Figure 5 plots the probability of risk taking as a function of the difference in expected values between the two options, $pX - C$, for all test trials from the first condition.

[Place Figure 5 about here]

²⁰Changing the exact response time cutoff does not qualitatively change our results. If we increase the response time cutoff, the results we later present in Figure 5 and Table 2 become stronger.

The figure shows a striking difference across conditions: a \$1 increase in $pX - C$ leads to a greater increase in the likelihood of choosing the risky lottery in the low volatility condition, compared to the high volatility condition. This is our first piece of evidence supporting the hypothesis that efficient coding does indeed extend into the domain of risky choice. Moreover, the model we present in Section II assumes that the *DM* first efficiently codes the payoff values of X and C separately, and then compares the perceived expected value of each option. Thus, the effect in Panel A should also be observed when separately analyzing X and C . Panel B of Figure 5 provides evidence that behavior is indeed consistent with efficient coding of each payoff separately. Taken together, Panel A and Panel B of Figure 5 are consistent with the model’s implications in Figure 3.

To provide formal tests of efficient coding, we conduct a series of regressions in which we account for heterogeneity across subjects in baseline sensitivity to X and C . Specifically, we conduct mixed effects linear regressions in which there is a random effect on both X and C (in addition to a random intercept). The dependent variable takes the value of one if the subject chooses the risky lottery, and zero otherwise. Column (1) of Table 2 shows that risk taking increases significantly in X and decreases significantly in C among all trials in the test phase of the first condition. The critical coefficients are those on the interaction terms, $X \times high$ and $C \times high$, which are significantly negative and positive, respectively, each at the 1% level.²¹

[Place Table 2 about here]

Importantly, our design enables us to provide a cleaner test of efficient coding by restricting to only the common trials in the test phase, which are identical across volatility conditions. Column (2) shows that the interaction terms remain significant at the 1% level. This finding represents strong support for efficient coding: for the same choice sets, a given increase in X or C leads to a larger change in risk taking when payoffs are drawn from the low volatility distribution, compared to the high volatility distribution. The results in Columns (1) and (2) provide results from between subjects tests, as they only use data from the first condition that subjects experience. Our design also enables us to examine how coding varies *within* subjects over time when faced with

²¹We also estimated mixed effects logistic regressions. However, with three random effects, mixed effects logistic regressions do not converge to numerically stable estimates. As an alternative, we estimated a series of logistic regressions without random effects, in which we pool subjects and cluster standard errors by subject. Table E1 shows that the results are consistent with those presented in Table 2. Our preferred specification is the mixed effects linear regression because it accounts for heterogeneity across subjects.

a change in the environment. While our model does not make predictions about adaptation, we can provide tests of whether behavior changes in the direction predicted by efficient coding as subjects experience a shift in the environment. Recall that halfway through the risky choice task, the payoff distribution switches. Thus, by re-estimating the regression in Column (2) using data from both volatility conditions for each subject, we can measure the effect within subjects.

Column (3) shows that the coefficients on the interaction terms have the predicted sign, though the effects are weaker compared to the between subjects tests. The coefficient on the interaction term $X \times high$ is significant at the 10% level, while the coefficient on $C \times high$ is significant at the 5% level. One reason for these weaker effects is that, at the beginning of the second condition, subjects may still be adapted to the first condition. To allow for longer adaption in the second condition, we restrict to only trials from the last half of the second condition (trials 451 through 600). Column (4) shows that the magnitudes of the interaction effects do indeed get larger.

In Columns (5) and (6) we further disaggregate the data based on whether the subject experiences the low or high volatility condition first. Column (5) shows that when subjects begin with the low volatility condition, the coefficients on $X \times high$ and $C \times high$ remain significant at the 5% level. This result is important because it rules out an alternative theory whereby subjects encode payoffs with noise, and through experience with a payoff, the payoff-specific perceptual noise decreases over time. Such a theory could explain our between subjects results, since subjects in the low volatility condition experience low volatility payoffs more frequently than subjects in the high volatility condition. The crucial insight from the within subject test in Column (5) is that under this alternative “learning from experience” theory, behavior in the second (high volatility) condition should be less noisy, as subjects experience the same set of 30 common trials for a second time. Yet we find that the effect has the opposite sign, which is instead consistent with noisy and *efficient* coding.

Column (6) provides results for those subjects who experience the high volatility condition first, and we see that the results have the predicted sign, but the effects become weaker. The coefficient on $C \times high$ remains significantly greater than zero at the 5% level, but the coefficient on $X \times high$ is no longer significantly different from zero. We speculate that it is easier for subjects to detect a change in the environment when shifting from low volatility to high volatility, because “outlier payoffs” that are never experienced in the first condition begin to appear in the second condition.

In contrast, when moving from a high volatility to low volatility environment, the information that signals a change in environment is less salient.²²

Response times

Not only are subjects more sensitive to payoffs in the low volatility condition, but they also implement decisions more quickly in the low volatility condition. Among common test trials, subjects in the high volatility condition take an average of 2.20 seconds vs. 1.97 seconds in the low volatility condition (p -value = 0.01 for difference in the mean $\log(\text{response time})$ estimated with a mixed effects linear regression). The fact that response times are significantly shorter in the low volatility condition suggests that our results on risk taking would become stronger if we account for sensitivity in choice *per unit time*. We provide more details on response times in Appendix C.

1.3. Design of the perceptual choice task

Recall that all the implications of our model are driven by the noisy encoding of X and C . In particular, we make two simplifying assumptions: (i) there is no noise in encoding the probability p , and (ii) there is no noise in computing the product of p and $\mathbb{E}[\tilde{X}|R_x]$. In reality, there is likely to be noise in both of these processes, which could potentially be responsible for some of the above experimental results.

To provide a more targeted test of the key efficient coding mechanism, subjects participate in a second “perceptual choice task.” In this task, subjects still need to perceive X , but do not need to perceive the probability p or integrate probabilities with perceived payoffs. Given that the noisy encoding of payoffs is sufficient to generate our main theoretical predictions in Section II, we should still find evidence that the perception of X depends on the recent stimulus distribution even when there is no need to perceive the probability p .

Our perceptual choice task is informed by work from the literature in perception of symbolic numbers (Moyer and Landauer, 1967). We build on the design of Dehaene, Dupoux, and Mehler (1990), who present subjects with an Arabic number between 31 and 99 on each trial of their

²²Some evidence that supports this conjecture comes from the time series of response times. We find that when subjects shift from the low to high volatility condition, there is a spike in response time which may be triggered by payoffs that were previously not experienced (see the upper panel of Figure C2). However, when subjects switch from high to low volatility condition, there is no corresponding spike, since there are no outliers (see the lower panel of Figure C2). We return to discussing the topic of adaptation in Section IV.1 and Appendix C.

experiment. The subject’s task is simply to classify whether the Arabic numeral presented on the screen is larger or smaller than the reference level of 65. Dehaene et al. (1990) find that as the stimulus gets closer to the reference level, response times increase and accuracy decreases. These results are consistent with the noisy encoding of Arabic numerals, which lies at the foundation of the model of risky choice we present in Section II.

One notable feature of the Dehaene et al. (1990) experiment is that the stimulus distribution is held constant throughout the experiment. Here, we exogenously vary the stimulus distribution across two conditions: a high volatility condition and a low volatility condition. In the high volatility condition, subjects are presented with an Arabic numeral, which we denote by X , that is drawn uniformly from integers in the set $[31, 99] \setminus \{65\}$. In the low volatility condition, X is drawn uniformly from integers in the set $[56, 74] \setminus \{65\}$. In each condition, subjects are asked to classify whether X is above or below the reference level of 65. Each subject completes both conditions, and we randomize the order of conditions across subjects. Figure 6 gives a schematic of the design.

[Place Figure 6 about here]

In all other respects, the perceptual choice task design follows closely the design of the risky choice task. In each condition, there is an initial set of 60 trials which are intended to allow subjects to adapt to a given distribution. As outlined in our pre-registration document, we only analyze behavior after the adaptation phase in the subsequent 340 test trials. To generate a clean comparison across conditions, we focus our main analysis on those trials in the test phase for which the stimulus numbers fall in the range of common support across the two conditions, $[56, 74]$.²³ As in the risky choice task, the restriction to common stimuli is crucial because it allows us to identify the effect of efficient coding by varying only the distribution of non-common trials in the test phase.

We pay subjects based on both the accuracy and speed of their classifications. Specifically, subjects earn a payoff of $\$(15 \times accuracy - 10 \times avgseconds)$, where *accuracy* is the percentage of correctly classified trials, and *avgseconds* is the average response time (in seconds) across all trials

²³Similar to the structure of our risky choice task, we designate 90 out of the 340 test trials as common trials. In both conditions, the 90 common trials are created by sampling each element in the low volatility condition 5 times. The common trials are randomly placed among the 340 test trials. The remaining 250 trials in the low volatility condition are drawn from the low volatility distribution. The remaining 250 trials in the high volatility condition are drawn with 50% probability from a uniform distribution over $[31, 55]$ and with 50% probability from a uniform distribution over $[75, 99]$. This procedure guarantees that the empirical distribution of X matches the population distribution in both conditions.

in the perceptual choice task. We incentivize fast responses in this task (but not in the risky choice task) in order to avoid “ceiling effects” in the choice data where subjects would approach 100% accuracy.²⁴ While a ceiling effect is not problematic on its own, it would cause difficulty in detecting any differences in the choice data across experimental conditions. The average earning for the perceptual choice task was \$8.70.

1.4. Results from the perceptual choice task

We begin by reporting results for between subjects tests using all test trials from the first condition. Subjects correctly classify the number on 95.0% of trials (standard error: 0.1%) with an average response time of 0.557 seconds (standard error: 0.001). Two out of the 150 subjects exhibit an average response time of only 0.05 and 0.10 seconds, which indicates that they used a guessing strategy, and thus we exclude them from all subsequent analyses (their average accuracy rates were 51.8% and 55.0%, respectively).²⁵

[Place Figure 7 about here]

Panel A of Figure 7 plots the proportion of trials that subjects classified X as larger than the reference level of 65, for each value of X . Consistent with previous research on numerical cognition, we see that subjects exhibit errors in classification, and more importantly, the errors increase as X approaches 65. To be clear, while it is unsurprising that subjects exhibit errors, the fact that error frequency correlates with X , provides evidence consistent with noisy coding. The novel aspect of our design that enables us to test for *efficient* coding, is the manipulation across two volatility conditions.

Among trials for which $X \in [56, 74]$, we find that subjects exhibit significantly greater accuracy in the low volatility condition, compared to the high volatility condition (95.0% vs. 92.3%, with p -value < 0.001 under a mixed effects linear regression). Not only are subjects more accurate in the low volatility condition, they also respond significantly faster (0.576 seconds vs. 0.611 seconds, with p -value = 0.03 under a mixed effects linear regression for difference in the mean $\log(\text{response time})$). Panel B of Figure 7 shows that response times (on trials where subjects classify the stimulus

²⁴Earlier experiments on number classification, e.g., [Dehaene et al. \(1990\)](#), typically instruct subjects to answer as quickly as possible, though the instructions do not contain financial incentives as they do in our experiment.

²⁵Including these two subjects in our subsequent analyses does not affect any of the main results.

correctly) become longer as X approaches 65, and they are shorter in the low volatility condition across the distribution of X .²⁶

Next, we compare the slopes across the two conditions from Panel A of Figure 7. Efficient coding predicts a steeper slope in the low volatility condition. Before proceeding to the test, we note that in this task, the conditions in (8) are no longer satisfied because there is only one variable that the DM needs to encode; thus, our tests in this section assume that the DM maximizes mutual information. To formally test for the difference in slopes, Table 3 presents results from a series of mixed effects logistic regressions. The dependent variable in our regression takes on the value of one if the subject classifies X as above 65, and zero otherwise, and we only use data for which $X \in [56, 74]$. Column (1) shows that the coefficient on X is significantly positive, indicating that subjects' propensity to classify X as greater than 65 is increasing in X . More importantly, we find that the coefficient on the interaction term, $X \times high$, is significantly negative, indicating that choices are noisier on trials in the high volatility condition. The next two columns indicate that the effect remains significant when examining only numbers in the 60s decade (Column 2) and examining only numbers outside the 60s decade (Column 3).²⁷

[Place Table 3 about here]

In Columns (4)-(6) of Table 3, we pool data across both conditions for each subject, which provides within subject tests of efficient coding. We see that, for each specification across these three columns, the coefficient on the interaction term remains significant at the 1% level. We also note that the coefficient on the *high* dummy variable is significantly positive in all specifications except Column (3). Moreover, there appears to be an asymmetry in the choice curve in Panel A of Figure 7, where accuracy is higher for numbers above the reference level compared to those below the reference level. Previous literature has observed the same asymmetry in an almost identical task, which may be driven by early stages of perceptual processing that are unrelated to the

²⁶The fact that both error rates and response times increase as X approaches the reference level of 65 is consistent with the “distance effect” in Moyer and Landauer (1967) and Dehaene et al. (1990).

²⁷We also find substantial heterogeneity in accuracy and response times across the 147 subjects we analyze in this task. (Note that we examine 147 out of 150 total subjects because we exclude 1 subject based on performance in the risky choice task and 2 subjects based on performance in the perceptual choice task.) Moreover, average accuracy and average response time are significantly positively correlated across subjects (p -value = 0.01 and p -value = 0.02, in the high and low volatility condition, respectively). The positive correlation is consistent with the well-known speed accuracy tradeoff in perceptual decision-making.

decision-making system (Dehaene, 1989; Hinrichs, Yurko, and Hu, 1981; Dehaene et al., 1990).²⁸

1.5. Model estimation

In this section, we structurally estimate the model in order to assess its quantitative fit to the experimental data.

Estimation of the risky choice task

We first estimate the model for each subject using data from the risky choice task. Recall that the model’s one free parameter, n , denotes the number of binary readings that are used to generate the noisy signal R_x (and R_c). That is, n represents the amount of “perceptual resources” the brain has: the greater n is, the more precise is the representation of risky payoffs.

To estimate n , we use maximum likelihood. Specifically, for each subject, we maximize the following log likelihood function over n , using choice data from the test phase of the first condition:

$$LL(n|\mathbf{y}) = \sum_{t=31}^{300} y_t \cdot \log(\mathbb{P}\text{rob}(y_t|n)) + (1 - y_t) \cdot \log(1 - \mathbb{P}\text{rob}(y_t|n)), \quad (16)$$

where $\mathbf{y} = \{y_t\}_{t=31}^{300}$, and y_t denotes the subject’s choice on trial t ; $y_t = 1$ if the subject chooses the risky lottery, and $y_t = 0$ if the subject chooses the certain option. In addition, $\mathbb{P}\text{rob}(y_t|n)$ denotes the model predicted probability of choosing the risky lottery given n , X_t , and C_t ; it is computed using equation (14) from Section II.3. We maximize the log likelihood function in (16) by searching over integer values of n in $[5, 40]$. We find that the average n across subjects is 8.5 with a standard deviation of 9.7, indicating substantial heterogeneity.

Recall that our model of efficient coding assumes linear utility; the DM ’s valuation of a payoff X is given by $\mathbb{E}[\tilde{X}|R_x]$. Here we also consider a more general model in which the DM first forms an optimal estimate of the payoff that is offered on a given trial, and then converts this estimate into subjective utility by applying a nonlinear transformation. We modify the baseline model by assuming that the DM assesses the value of a payoff X by $(\mathbb{E}[\tilde{X}|R_x])^\alpha$, where $\alpha \leq 1$ captures

²⁸Dehaene (2008) notes that there is a discrete jump in response times for numbers in the range $[60, 69]$, which he attributes to the extra attention needed to focus on the second unit of the number to distinguish it from the reference level. Indeed, Dehaene (2008) estimates a random walk model and finds that the “non-decision time”—which corresponds to the time needed to implement early sensory processes and motor preparation—is higher for those numbers in the range $[60, 69]$. Because our objective is to test for efficient coding and thus compare behavior *across* conditions, our tests do not depend critically on the mechanism that generates the asymmetry *within* a condition.

intrinsic risk aversion. Under this model, the coding rule presented in equation (4) is still optimal if the *DM*'s objective is to maximize the mutual information between the stimulus and its noisy signal, but the coding rule is no longer optimal when the *DM*'s objective is to maximize expected payoff.

With this caveat in mind, we estimate the modified model in which the *DM* chooses the risky lottery if and only if $p \cdot \mathbb{E}([\tilde{X}|R_x])^\alpha > \mathbb{E}([\tilde{C}|R_c])^\alpha$.²⁹ We use the same maximum likelihood procedure as above, but for each subject we now estimate two parameters (n, α) . We find that the best fitting parameter pair, averaged across subjects, is $n = 10.0$ and $\alpha = 0.93$. These results indicate that the average subject exhibits a modest degree of intrinsic risk aversion with the small stakes in our experiment. We can also quantitatively assess the validity of our linear utility assumption by running an Akaike information criterion (AIC) test at the subject level. For each subject, we compare the AIC across the baseline model in which we constrain $\alpha = 1$ and the generalized model. We find that 54% of subjects are best fit using the restricted model with $\alpha = 1$. Thus, our baseline assumption of linear utility is not overwhelmingly restrictive. At the same time, the model provides a better fit to a substantial number of subjects when estimated payoffs are subsequently passed through a nonlinear utility function.

To assess the model's out-of-sample fit, we re-estimate the model using a training set and then measure its fit on a distinct validation dataset. Specifically, for each subject, we record the best fitting (n, α) by maximizing the log likelihood function in (16), but using data only from the *even*-numbered trials in the test phase of the first condition. We then use the estimated (n, α) to compute the theoretical probability of choosing the risky lottery for each *odd*-numbered trial in the test phase of the first condition. Lastly, we bin trials according to their theoretical probabilities, and for each bin, we compute the proportion of trials on which subjects choose the risky lottery.

[Place Figure 8 about here]

Figure 8 presents the results of out-of-sample tests across four separate cuts of the data. In the top panel, we split the data based on whether the first condition is the high or low volatility condition. The red forty-five degree line provides the prediction of the model, conditional on subject-

²⁹As in equation (14), we also assume that when $p \cdot \mathbb{E}([\tilde{X}|R_x])^\alpha = \mathbb{E}([\tilde{C}|R_c])^\alpha$, the *DM* randomly chooses between the two options.

specific parameters of (n, α) . If the model perfectly fits the data, then all blue dots would lie on the forty-five degree line; visually, we see a good match between the model-implied probabilities and the data across the entire probability domain of $[0, 1]$. The bottom left panel pools the data together from the two top panels.

In order to quantitatively explore the core mechanism of efficient coding—which predicts that perception of a given monetary amount depends on the environment—we also perform a counterfactual analysis. Our goal here is to understand how the model’s prediction varies across environments, holding *constant* the choice set. To do so, we analyze how a subject’s perception of a choice set in the low volatility condition would change if the same choice set was instead presented in the high volatility condition. In particular, we repeat the analysis used to generate the top right panel of Figure 8—which uses data only from the low volatility condition—but we counterfactually assume that the subject’s prior is the high volatility payoff distribution. In the bottom right panel of Figure 8, we plot the relationship between data on choices in the low volatility condition and the counterfactual probabilities. We see that this counterfactual analysis results in a systematic mismatch between model predicted probabilities and data, both at high and low ends of the probability domain. A formal AIC test demonstrates that the model which assumes a low volatility prior provides a better fit to the data compared to the counterfactual model (AIC value for the model which assumes a low volatility prior is 9,729; and AIC value for the model which assumes a high volatility prior is 9,803). These results demonstrate that the environmental distribution to which the subject has adapted, provides important information about the subject’s perception of the choice set, and hence, her probability of risk taking.³⁰

Estimation of the perceptual choice task

We now estimate the model for each subject using data from the perceptual choice task. The maximum likelihood procedure we implement is nearly identical to that from (16). For each subject, we maximize the following log likelihood function over n , using choice data from the test phase of

³⁰It is not feasible to provide the analogous analysis for data in the high volatility condition, since a majority (8/9) of stimuli in that condition fall outside the support of the low volatility distribution, and thus the model would not make any reasonable predictions in such a case.

the first condition:

$$LL(n|\mathbf{z}) = \sum_{t=61}^{400} z_t \cdot \log(\mathbb{P}\text{rob}(z_t|n)) + (1 - z_t) \cdot \log(1 - \mathbb{P}\text{rob}(z_t|n)), \quad (17)$$

where $\mathbf{z} = \{z_t\}_{t=61}^{400}$, and z_t denotes the subject’s choice on trial t ; $z_t = 1$ if the subject classifies the stimulus X_t on trial t as greater than 65, and $z_t = 0$ if the subject classifies X_t as less than 65. The term $\mathbb{P}\text{rob}(z_t|n)$ is the model predicted probability that the subject perceives X_t to be greater than 65. We maximize (17) over integer values of n in the range $[5, 40]$. The best fitting value of n , averaged across all subjects, is 15.8, with a standard deviation of 13.1.

Given that each subject completes both the risky choice task and the perceptual choice task, our design enables us to compare how the latent structural parameter n compares across these two tasks. Two observations are worth noting. First, the average value of n is lower in the risky choice task than in the perceptual choice task. This difference is likely driven by the fact that the risky choice task is more complex, and hence additional sources of noise enter the decision process (e.g., encoding the probability p and integrating p with $\mathbb{E}[\tilde{X}|R_x]$), which the model accounts for through a lower value of n .

Second, we test for a correlation between the estimated n from each task, across subjects. This test is important because it allows us to assess whether errors in numerical discrimination from the perceptual choice task can explain variation in the risky choice task. We find a modest but significant rank correlation of 0.30 between parameters from each task when allowing $\alpha \leq 1$ (p -value < 0.001). The correlation remains significant at 0.26 when using the estimated n from the restricted model where $\alpha = 1$ for all subjects (p -value = 0.001). The results are also robust to using Pearson correlations for both models (p -value = 0.002 for unrestricted model; p -value = 0.001 for restricted model). These positive correlations demonstrate that, across subjects, variation in perception partly explains variation in risk taking behavior.

2. Experiment 2: Shape manipulation

In the experiments reported in the previous section, we focused on manipulating the range (and hence the volatility) of the payoff distribution while holding its mean constant. In this section, we investigate whether manipulating the *shape* of the payoff distribution, while holding the range

constant, affects risk taking in the manner predicted by efficient coding. Such a manipulation is important for at least two reasons. First, it allows us to further distinguish between efficient coding models and alternative theories in which behavior is affected only by the range of stimuli encountered (and the current choice set). Second, such an experiment provides a fundamental test of Weber’s law.

As in the risky choice task from Experiment 1, we design an experiment in which subjects are presented with choice sets of the form $\{(X, 0.5; 0, 0.5); (C, 1)\}$. In one condition, X is drawn from a linearly decreasing distribution over the range $[8, 32]$. In another condition, X is drawn from a linearly increasing distribution over the same range.³¹ In both conditions, C is drawn from a uniform distribution over $[2, 18]$. We manipulate only the distribution of X across the two conditions, in order to isolate the induced shift in the value function from a single variable.

As shown in Panel B of Figure 2, under efficient coding, the decreasing payoff distribution of X leads to a concave value function $v(X)$. Conversely, the increasing distribution gives rise to a convex $v(X)$. In this case, efficient coding predicts a type of “anti-Weber’s” law: perception becomes *more* precise as payoff values increase, and small payoffs are biased upwards towards the mean of the payoff distribution. A provocative implication of such a convex value function, of course, is that it induces risk seeking behavior.

To test this implication, we create a set of common trials that are presented in both experimental conditions, which we use to generate the data for our main tests. Specifically, we create 10 common trials, where we fix C at \$6.09 and vary X from \$9.04 to \$18.06 in approximately \$1 increments, as shown in Table 4; we focus our design on small values of X , where the value function is most locally convex for the increasing distribution.

[Place Table 4 about here]

There are 300 trials per condition, and each of the 10 common trials is presented every 30 trials. Except for the 10 common trials, the choice sets in each condition are drawn from their respective payoff distribution (either increasing or decreasing).³²

³¹Specifically, both the decreasing distribution and the increasing distribution take the form of (15) in Section II. We set X_l to 8 and X_u to 32 for both conditions; and we set Δ to $-1/30$ for the decreasing distribution and to $1/30$ for the increasing distribution.

³²The design in this task is similar to that of Payzan-LeNestour and Woodford (2020) who insert a “test trial” every 40 trials, though their design is implemented in the domain of a perceptual task where a subject is incentivized

Our main testable prediction is that demand for the risky lottery among choice sets on common trials will be higher in the increasing condition compared to the decreasing condition. We pre-register the experiment and recruit 40 students from the University of Southern California to participate in the laboratory, none of whom participated in the previous volatility experiment (see Appendix B for the pre-registration document). Each subject completes the task and is paid according to one randomly selected trial, in addition to a \$7 show-up fee. The average total earning for the task, including the show up fee, is \$16.63. The experimental instructions for this task can be found in Appendix B.

2.1. Results from Experiment 2

We begin by cleaning the data and removing trials for which subjects responded in less than 0.5 seconds, which constitutes 4.9% of the dataset. Of the remaining 22,835 trials, subjects choose the risky lottery on 41.3% of trials (standard error: 2.6%).

As outlined in our pre-registration document, our tests for this experiment are all conducted within subjects. Table 5 presents results from mixed effects linear regressions, where the dependent variable takes the value of one if the subject chooses the risky lottery, and zero otherwise. In Column (1) we use all data from the experiment, and find that controlling for X and C , subjects are 2.5% more likely to choose the risky lottery in the context of the increasing distribution, compared to the decreasing distribution (p -value < 0.001). This result, however, relies on the assumption that the linear controls for X and C fully account for differences across current choice sets.

[Place Table 5 about here]

To provide a cleaner test of the hypothesis that risk taking is greater in the increasing condition, in Column (2), we restrict to only the 10 common trials in each condition. Recall that for these common trials, C is fixed at \$6.09. Hence, we remove the control for C . Compared to Column (1), we find that the effect size doubles and subjects are 5.4% more likely to take risk on common trials from the increasing distribution (p -value = 0.02). This result implies that for the same ten choice sets, the same subject is more willing to take risk after being recently exposed to an increasing distribution of X .

to discriminate between shades of grey.

We interpret this result as a direct effect of efficient coding. Because C is fixed on each common trial—and because the distribution of C is fixed across conditions—all variation in risk taking must come through X . As discussed earlier, Panel B of Figure 2 shows that efficient coding generates a convex value function when the prior distribution is increasing. In this case, low values of X occur infrequently, resulting in an upward perceptual bias which induces risk taking. In contrast, when the prior distribution is decreasing, the implied value function is concave. In this case, low values of X occur frequently, and thus they are not biased. Overall, the difference in perceived valuation of X gives rise to higher risk taking when common trials are drawn from the increasing distribution, consistent with the findings reported in Column (2).

While the difference in risk taking is predicted to be higher for small values of X , the precise value of X at which the predicted difference switches sign depends on the capacity parameter n . To get a more quantitative handle on exactly where this cutoff is located, we compute the cutoff implied by the distribution of (n, α) that we estimated from the 147 subjects in the risky choice task from Experiment 1; the implied cutoff is $X = 15.79$.³³ In Column (3) of Table 5, we re-estimate the regression restricting to test trials for which $X < 15.79$, and we find that our results get slightly stronger, as the average difference in risk taking across conditions is 6.8% (p -value = 0.02).

IV. Discussion

1. Dynamics across trials and within trials

The theoretical framework we present in Section II is a static model of risky choice, and hence does not tackle the important question regarding *how* the DM learns the prior distribution. For simplicity, we assume that, in our experiments, subjects are fully adapted to the population distribution after completing an initial set of pre-registered “adapt trials.” We emphasize that this assumption is not trivial, particularly because the DM ’s learning problem is more complex than in standard settings, where Bayesian inference would typically generate convergence. The additional layer of complexity is due to the DM ’s inability to observe the sequence of *objective* payoffs, and hence the DM must learn from the history of *perceived* payoffs (Robson and Whitehead,

³³Because the subjects from each of the two experiments are recruited from the same subject pool, the distribution of (n, α) should be largely similar.

2018; Młynarski and Hermundstad, 2019; Aridor, Grechi, and Woodford, 2020).

In an extension to their main model, HWP consider an environment in which it is costly to acquire and store information about the entire history of the stimulus distribution. They argue that such a constraint is most relevant in situations where the *DM* expects the distribution of values to change frequently. Under this criterion, the relevance of a large cost to acquiring and storing information about the prior is fairly minimal in our experiments, because the payoff distribution changes only once over 600 trials, and never changes in our between subjects tests. Nonetheless, understanding how the *DM* learns the prior in more general environments is clearly an important area for future work (Payzan-LeNestour and Woodford, 2020), and the mechanisms through which past stimuli are recalled from memory are likely to have important implications for efficient coding of economic stimuli (Kahana, 2012; Bordalo, Gennaioli, and Shleifer, 2020; Wachter and Kahana, 2020).

The above discussion is in regard to the dynamics across trials in our experiment. There is also a dynamic component to a single decision *within* a trial, which we have analyzed empirically. Specifically, we find that response times on common trials are significantly longer in the high volatility condition compared to the low volatility condition of our first experiment. There is also a striking similarity between the response time patterns in the risky choice and perceptual choice tasks (see Panel A of Figure C1 and Panel B of Figure 7); the similarity in response times provides another piece of evidence in favor of a common mechanism that governs both types of choices. A natural way to extend our framework to address variation in response times would be to allow the *DM* to draw a *sequence* of noisy signals, $\{R_{X,i}\}_{i=1}^N$ for a given payoff X . The time it takes to execute a decision would then reflect the number of signals drawn, which is a common interpretation of sequential sampling models from mathematical psychology (Ratcliff, 1978; Bogacz, Brown, Moehlis, Holmes, and Cohen, 2006; Krajbich, Armel, and Rangel, 2010), and more recently, from economics (Woodford, 2014; Fudenberg, Strack, and Strzalecki, 2018; Hébert and Woodford, 2019).

2. Comparison with alternative theories

In this section, we discuss alternative models of behavior and how their predictions relate to our experimental results. As noted in the Introduction, the main prediction we test in Experiment 1—that sensitivity to payoff values increases when the dispersion of potential values decreases—

is also shared by models of “normalization” (Rangel and Clithero, 2012; Louie et al., 2015). A particularly relevant set of normalization models are those in which value is normalized based on the *range* of potential stimuli (Soltani et al., 2012; Rustichini et al., 2017). Under this class of models, the subjective value of a payoff depends only on the payoff itself and the range of potential payoffs. The results from our first experiment strongly support the predictions of range normalization models, and specifically, they highlight the interpretation that normalization can implement normative principles of efficient coding.³⁴ See Appendix D for more detail.

At the same time, range normalization cannot explain all our experimental results. In Experiment 2, we hold constant the range of payoff values while changing the shape of the payoff distribution across conditions. Our data show systematic differences in the demand for risk across conditions, whereas models with range normalization do not predict any differences. Thus, range normalization does not offer a complete picture of how the distribution of values affects behavior.

In the decision-by-sampling (DbS) model by Stewart et al. (2006), the *DM*’s subjective value of a stimulus is given by its rank within a distribution of values recalled from memory. To the extent that the distribution of recalled values is related to the prior distribution that we focus on in this paper, DbS and efficient coding models make qualitatively similar predictions. In fact, Bhui and Gershman (2018) show that efficient coding can serve as a normative foundation for DbS. We interpret the data from both of our risky choice experiments as novel evidence consistent with the core mechanism in DbS.³⁵

Kőszegi and Rabin (2007) (KR) offer a model of risky choice where the reference point is given by rational expectations about outcomes from a reference lottery. At a basic level, KR and efficient

³⁴Not all models of normalization are grounded in principles of optimization; some are instead developed to describe the decision process and its outcome. For example, in the Soltani et al. (2012) model, range normalization is assumed, and its implications are shown to provide a good description of decoy effects in risky choice (though normalization takes place over values on a *single* experimental trial, rather than over the history of trials experienced). In more recent models, such as Rustichini et al. (2017), normalization is the outcome of an optimization procedure. Relatedly, the prediction of greater sensitivity to attributes with a smaller range is a key assumption in the relative thinking model by Bushong, Rabin, and Schwartzstein (2020).

³⁵Our evidence also speaks to a recent controversy in interpreting tests of the DbS model. Stewart, Reimers, and Harris (2015) and Walasek and Stewart (2015) claim to find supporting evidence for DbS by manipulating the distribution of payoffs across choice sets, similar to the manipulation in our design. However, a re-analysis of the experimental evidence finds that neither paper can be interpreted as supporting DbS (Alempaki, Canic, Mullett, Skylark, Starmer, Stewart, and Tufanod, 2019; André and de Langhe, 2020). The issue arises from the fact that behavior was analyzed on different choice sets across experimental conditions. In contrast, our design has the important advantage of presenting a collection of choice sets that are common to both conditions, and the common choice sets reflect the statistical properties of the environment assumed in the theory. Thus, our results should help restore faith in the empirical validity of DbS. Moreover, our design provides a template for future experimental tests of the DbS theory.

coding share the feature that expectations shape the *DM*'s perception of a lottery payoff. Yet an important distinction between the two models is that in efficient coding, the driving force of the model is the *DM*'s expectation of a payoff value, once it has been presented in the choice set. In KR, however, the driving force is expectations over which payoff value she is likely to receive as a future outcome from the lottery. In Appendix D.1, we examine KR's implications for our shape manipulation task (Experiment 2). In most common specifications of expectations, KR does not predict a difference in behavior as the shape of the payoff distribution varies.³⁶

Saliency theory is an alternative model of risky choice that is also grounded in principles of perception and delivers context-dependent behavior (Bordalo et al., 2012). Under saliency theory, attention is drawn to those payoffs that are very different in percentage terms from a reference payoff. Bordalo et al. (2012) appeal to Weber's law of diminishing sensitivity, in part, as a justification for their definition of saliency. Importantly, in their model, Weber's law is an exogenous assumption. In contrast, Weber's law arises endogenously under efficient coding for prior distributions that are decreasing, but an "anti-Weber's" law will arise when prior distributions are increasing in payoff values. In this sense, saliency theory and efficient coding differ with respect to their primitive assumptions.

Naturally, this leads the two models to generate different predictions in many environments. For example, saliency theory does not predict that risk taking will increase when the *DM* is adapted to an increasing payoff distribution, as we find in our second experiment. Nor does it deliver stochastic choice, where the degree of stochasticity changes systematically with the prior—as we observe in our first experiment. At the same time, there are extant empirical patterns in the literature that saliency theory can explain, which the efficient coding model in Section II cannot, such as the dependence of risk taking on the correlation between mutually exclusive lotteries.

3. *Probability weighting*

The efficient coding model we present in Section II can also be used to encode probabilities. As we now show, one interesting implication of the efficient coding of probabilities, using the HWP model, is that it generates the classic probability weighting function from prospect theory. Suppose

³⁶Similarly, the related models of disappointment aversion (Bell, 1985; Loomes and Sugden, 1986) do not predict a difference in behavior across the two conditions in our shape manipulation task.

the *DM* holds prior beliefs about the payoff probability, given by the distribution, $f(p)$. Also suppose that, upon observing the probability p , the *DM*'s perceptual system generates a noisy signal, R_p , of p , based on a conditional distribution $f(R_p|p)$.

Following HWP, we assume that the conditional distribution $f(R_p|p)$ is given by

$$f(R_p|p) = \binom{n}{R_p} (\theta(p))^{R_p} (1 - \theta(p))^{n-R_p}, \quad (18)$$

where n represents the *DM*'s coding capacity, and $\theta(p)$ is the coding rule that maps the value of p to the probability that a given neuron emits a value of 1. When the performance objective is to maximize the mutual information between p and R_p ,

$$\max_{\theta(p)} I(p, R_p), \quad (19)$$

HWP show that the optimal coding rule is

$$\theta(p) = \left[\sin \left(\frac{\pi}{2} F(p) \right) \right]^2, \quad (20)$$

where $F(p)$ is the cumulative density function of the *DM*'s prior belief $f(p)$.

Given the *DM*'s prior, the conditional distribution $f(R_p|p)$ from (18), and the coding rule from (20), the value function of the probability p and the standard deviation for the subjective valuation are

$$v(p) = \sum_{R_p=0}^n \mathbb{E}[\tilde{p}|R_p] \cdot f(R_p|p) \quad (21)$$

and

$$\sigma(p) = \left[\sum_{R_p=0}^n (\mathbb{E}[\tilde{p}|R_p])^2 f(R_p|p) - v(p) \right]^{1/2}, \quad (22)$$

respectively, where $\mathbb{E}[\tilde{p}|R_p]$, the posterior mean of p , conditional on R_p , is given by

$$\mathbb{E}[\tilde{p}|R_p] = \frac{\int_0^1 f(R_p|p) f(p) p dp}{\int_0^1 f(R_p|p) f(p) dp}. \quad (23)$$

Importantly, the shape of $v(p)$ depends crucially on the shape of the *DM*'s prior distribution

$f(p)$. What is a good description of the naturally occurring distribution of probabilities, that the *DM* would use as a prior? [Stewart et al. \(2006\)](#) present evidence that, in the English language, probability phrases occur with higher frequency for extreme probabilities than for mid-range probabilities, and large probabilities occur more often than small probabilities.³⁷ Figure 9 presents a numerical example, in which $f(p)$ takes an asymmetric *U*-shape, as suggested by [Stewart et al. \(2006\)](#). Interestingly, the resulting value function $v(p)$ gives rise to both the overweighting of small probabilities—that is, $v(p) > p$ when p is close to zero—and the asymmetry in $v(\cdot)$ —that is, $v(0.5) < 0.5$. Both are important features of the probability weighting function in prospect theory ([Tversky and Kahneman, 1992](#)).

[Place Figure 9 about here]

An important observation is worth noting. While previous behavioral economics models have been used to microfound the probability weighting function ([Bordalo et al., 2012](#); [Steiner and Stewart, 2016](#); [Khaw et al., 2020](#); [Enke and Graeber, 2020](#)), the mechanism discussed here differs from these models by explicitly linking the *DM*'s prior distribution to the shape of the probability weighting function. This link leads to a testable prediction: changes in $f(p)$ cause systematic variation in the degree of probability weighting. Conducting such a test—perhaps by using a similar technique that we use in our two experiments to manipulate the prior distribution—can help to validate efficient coding as a source of probability weighting. We leave this test for future work.

4. *Instability of preference parameter estimates*

Our main experimental results are related to, but fundamentally distinct from, much work in experimental economics that documents how risk taking depends systematically on the realized lottery *outcomes* from previous choices ([Thaler and Johnson, 1990](#); [Weber and Camerer, 1998](#); [Imas, 2016](#)). Importantly, we find that, even when lottery outcomes are not presented to subjects, the distribution of previous *choice sets* still causes systematic variation in behavior. As a result, efficient coding may provide a distinct source of variation of behavior in typical lab experiments in which

³⁷[Stewart et al. \(2006\)](#) also highlight that, in lab experiments designed specifically to study the shape of the probability weighting function, the probabilities shown to experimental subjects are consistent with a distribution that puts higher density on extreme probabilities. For instance, in Figure 8 of [Stewart et al. \(2006\)](#), the authors reproduce the probability distribution used in [Gonzalez and Wu \(1999\)](#), which does put more weight on extreme probabilities.

preference parameters are elicited by presenting subjects with a sequence of choice sets (Broomell and Bhatia, 2014).

The causal effect of past choices sets on risk taking is particularly relevant for newer methods of eliciting preferences in which the ordering of choice sets is tailored in real time to a subject’s history of choices. For example, Toubia, Johnson, Evgeniou, and Delquié (2013) and Chapman, Snowberg, Wang, and Camerer (2019) present a sequence of choice sets to subjects that maximizes the information gain for estimating parameters of the value function from prospect theory. Efficient coding predicts that the optimal choice set to present to a subject should condition not only on the history of the subject’s choices, but also on the history of the presented *choice sets*. Conditioning on this extra aspect of the subject’s past experience should therefore aid in further optimizing experimental designs.

V. Conclusion

We have conducted an experimental test of whether a core principle from neuroscience—efficient coding—is a driving force in decision-making under risk. Our results provide strong evidence that the *DM*’s willingness to take risk depends systematically on the payoff distribution to which she has recently adapted. The results are consistent with the noisy perception of lottery payoffs, and moreover, we find that the particular distribution from which noise is drawn varies as an optimal response to a change in the environment. In Experiment 1, we show that risky choice becomes noisier as the volatility of the payoff distribution increases. In Experiment 2, we show that the likelihood of taking risk increases as the payoff distribution changes from monotonically decreasing to increasing. Together, the data indicate that risk taking is systematically unstable across environments, in a manner that closely mimics the instability of sensory perception.

Our results raise a number of important directions for future work. There is a strong need to understand how the *DM* adapts to a given environment based on the history of perceived payoffs. This mechanism will of course depend on the *DM*’s prior expectations about payoffs—which we manipulate in our experiments—but also on higher order priors about the rate of change of the environment. For instance, if the *DM* expects the environmental distribution to change rapidly, then adaptation will also likely take place at a fast pace (Behrens, Woolrich, Walton, and

Rushworth, 2007; Nassar, Rumsey, Wilson, Parikh, Heasley, and Gold, 2012). Theory is already being developed along this dimension, but future experimental evidence of the adaptation process will be critical in guiding the development of such theory (Robson and Whitehead, 2018; Młynarski and Hermundstad, 2019; Aridor et al., 2020).

Another important direction for future research is to test the implications of efficient coding outside the laboratory. A challenge here is to measure the prior distribution to which the *DM* has adapted. A more refined theory of adaptation will be integral for guiding empirical work in the field, as it will help shed light on the timescale which is relevant for forming prior expectations, and hence perceptions. It is also likely that institutional factors will shape the relevant timescale for adaptation. For instance, in financial markets, the distribution of a stock's price over the past 52 weeks is typically salient to investors, and therefore may be a good candidate for the investor's prior distribution. We ourselves expect that future progress on the topic of efficient coding will benefit from the close interplay between theory, experimental tests, and empirical validations in the field.

REFERENCES

- Alempaki, Despoina, Emina Canic, Timothy L. Mullett, William J. Skylark, Chris Starmer, Neil Stewart, and Fabio Tufanod, 2019, Reexamining how utility and weighting functions get their shapes: A quasi-adversarial collaboration providing a new interpretation, *Management Science* 65, 4841–4862.
- André, Quentin, and Bart de Langhe, 2020, No evidence for loss aversion disappearance and reversal in Walasek and Stewart (2015), *Journal of Experimental Psychology: General* forthcoming.
- Aridor, Guy, Francesco Grechi, and Michael Woodford, 2020, Adaptive efficient coding: A variational auto-encoder approach, Working paper.
- 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).
- Barlow, Horace, 1961, Possible principles underlying the transformations of sensory messages, in Walter A. Rosenblith, ed., *Sensory Communication*, 217–234 (MIT Press).
- Behrens, Timothy, Mark Woolrich, Mark Walton, and Matthew Rushworth, 2007, Learning the value of information in an uncertain world, *Nature Neuroscience* 10, 1214–1221.
- Bell, David E., 1985, Disappointment in decision making under uncertainty, *Operations Research* 33, 1–27.
- Bhui, Rahul, and Samuel Gershman, 2018, Decision by sampling implements efficient coding of psychoeconomic functions, *Psychological Review* 125, 985–1001.
- Bogacz, Rafal, Eric Brown, Jeff Moehlis, Philip Holmes, and Jonathan D. Cohen, 2006, The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks, *Psychological Review* 113, 700–765.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2012, Saliency theory of choice under risk, *Quarterly Journal of Economics* 127, 1243–1285.

- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2020, Memory, attention, and choice, *Quarterly Journal of Economics* 135, 1399–1442.
- Broomell, Stephen B., and Sudeep Bhatia, 2014, Parameter recovery for decision modeling using choice data, *Decision* 1, 252–274.
- Bushong, Benjamin, Matthew Rabin, and Joshua Schwartzstein, 2020, A model of relative thinking, *Review of Economic Studies* forthcoming.
- Carandini, Matteo, and David J. Heeger, 2012, Normalization as a canonical neural computation, *Nature Reviews Neuroscience* 13, 51–62.
- Chapman, Jonathan, Erik Snowberg, Stephanie Wang, and Colin Camerer, 2019, Loss attitudes in the u.s. population: Evidence from dynamically optimized sequential experimentation (DOSE), Working paper.
- Clithero, John A., 2018, Response times in economics: Looking through the lens of sequential sampling models, *Journal of Economic Psychology* 69, 61–86.
- Dehaene, Stanislas, 1989, The psychophysics of numerical comparison: A reexamination of apparently incompatible data, *Perception & Psychophysics* 45, 557–566.
- Dehaene, Stanislas, 2008, Symbols and quantities in parietal cortex: elements of a mathematical theory of number representation and manipulation, in Patrick Haggard, Yves Rossetti, and Mitsuo Kawato, eds., *Attention and Performance XXII. Sensori-motor foundations of higher cognition*, 527–574 (Oxford University Press).
- Dehaene, Stanislas, 2011, *The Number Sense* (Oxford University Press, Oxford, United Kingdom).
- Dehaene, Stanislas, Emmanuel Dupoux, and Jacques Mehler, 1990, Is numerical comparison digital? analogical and symbolic effects in two-digit number comparison, *Journal of Experimental Psychology: Human Perception and Performance* 16, 626–641.
- Dehaene, Stanislas, and Jacques Mehler, 1992, Cross-linguistic regularities in the frequency of number words, *Cognition* 43, 1–29.

- Denrell, Jerker C., 2015, Reference-dependent risk sensitivity as rational inference, *Psychological Review* 122, 461–484.
- Enke, Benjamin, and Thomas Graeber, 2020, Cognitive uncertainty, Working paper.
- Friedman, Daniel, 1989, The S-shaped value function as a constrained optimum, *American Economic Review* 79, 1243–1248.
- Fudenberg, Drew, Philipp Strack, and Tomasz Strzalecki, 2018, Speed, accuracy, and the optimal timing of choices, *American Economic Review* 108, 3651–3684.
- Gabaix, Xavier, and David Laibson, 2017, Myopia and discounting, NBER working paper No. 23254.
- Garcia, Miguel, Marcus Grüschow, Rafael Polanía, Michael Woodford, and Christian C. Ruff, 2018, Predicting risk attitudes from the precision of mental number representation, poster presented at the 16th Annual Meeting of the Society for Neuroeconomics.
- 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.
- Glimcher, Paul, and Agnieszka Tymula, 2019, Expected subjective value theory (ESVT): A representation of decision under risk and certainty, SSRN working paper.
- Gonzalez, Richard, and George Wu, 1999, On the shape of the probability weighting function, *Cognitive Psychology* 38, 129–166.
- Heng, Joseph, Michael Woodford, and Rafael Polanía, 2020, Efficient sampling and noisy decisions, Working paper.
- Hinrichs, James V., Dale S. Yurko, and Jing-Mei Hu, 1981, Two-digit number comparison: Use of place information, *Journal of Experimental Psychology: Human Perception and Performance* 7, 890–901.
- Hébert, Benjamin, and Michael Woodford, 2019, Rational inattention when decisions take time, Working paper.

- Imas, Alex, 2016, The realization effect: Risk-taking after realized versus paper losses, *American Economic Review* 106, 2086–2109.
- Kahana, Michael J., 2012, *Foundations of Human Memory* (Oxford University Press).
- Kahneman, Daniel, and Amos Tversky, 1979, Prospect theory: An analysis of decision under risk, *Econometrica* 47, 263–291.
- Khaw, Mel Win, Paul W. Glimcher, and Kenway Louie, 2017, Normalized value coding explains dynamic adaptation in the human valuation process, *Proceedings of the National Academy of Sciences* 114, 12696–12701.
- Khaw, Mel Win, Ziang Li, and Michael Woodford, 2020, Cognitive imprecision and small-stakes risk aversion, *Review of Economic Studies* forthcoming.
- Krajbich, Ian, Carrie Armel, and Antonio Rangel, 2010, Visual fixations and the computation and comparison of value in simple choice, *Nature Neuroscience* 13, 1292–1298.
- Kutter, Esther, Jan Bostroem, Christian Elger, Florian Mormann, and Andreas Nieder, 2018, Single neurons in the human brain encode numbers, *Neuron* 100, 1–9.
- Kőszegi, Botond, and Matthew Rabin, 2006, A model of reference-dependent preferences, *Quarterly Journal of Economics* 121, 1133–1165.
- Kőszegi, Botond, and Matthew Rabin, 2007, Reference-dependent risk attitudes, *American Economic Review* 97, 1047–1073.
- Landry, Peter, and Ryan Webb, 2019, Pairwise normalization: A neuroeconomic theory of multi-attribute choice, SSRN working paper.
- Laughlin, Simon, 1981, A simple coding procedure enhances a neuron’s information capacity, *Zeitschrift fur Naturforschung* 36, 9–10.
- Loomes, Graham, and Robert Sugden, 1986, Disappointment and dynamic consistency in choice under uncertainty, *Review of Economics Studies* 53, 271–282.
- Louie, Kenway, Paul W. Glimcher, and Ryan Webb, 2015, Adaptive neural coding: from biological to behavioral decision-making, *Current Opinion in Behavioral Sciences* 5, 91–99.

- Ma, Wei Ji, and Michael Woodford, 2020, Multiple conceptions of resource rationality, *Behavioral and Brain Sciences* 43, 1–60.
- Moyer, Robert S., and Thomas K. Landauer, 1967, Time required for judgments of numerical inequality, *Nature* 215, 1519–1520.
- Mlynarski, Wiktor, and Ann M. Hermundstad, 2019, Adaptability and efficiency in neural coding, Working paper.
- Nassar, Matthew R., Katherine M. Rumsey, Robert C. Wilson, Kinjan Parikh, Benjamin Heasley, and Joshua I. Gold, 2012, Rational regulation of learning dynamics by pupil-linked arousal systems, *Nature Neuroscience* 15, 1040–1046.
- Natenzon, Paulo, 2019, Random choice and learning, *Journal of Political Economy* 127, 419–457.
- Netzer, Nick, 2009, Evolution of time preferences and attitudes toward risk, *American Economic Review* 99, 937–955.
- Padoa-Schoppa, Camillo, 2009, Range-adapting representation of economic value in the orbitofrontal cortex, *Journal of Neuroscience* 29, 14004–14014.
- Payzan-LeNestour, Elise, and Michael Woodford, 2020, “Outlier blindness”: Efficient coding generates an inability to represent extreme values, SSRN working paper.
- Polanía, Rafael, Michael Woodford, and Christian Ruff, 2019, Efficient coding of subjective value, *Nature Neuroscience* 22, 134–142.
- Rangel, Antonio, and John A. Clithero, 2012, Value normalization in decision making: Theory and evidence, *Current Opinion in Neurobiology* 22, 970–981.
- Ratcliff, Roger, 1978, A theory of memory retrieval, *Psychological Review* 85, 59–108.
- Rayo, Luis, and Gary S. Becker, 2007, Evolutionary efficiency and happiness, *Journal of Political Economy* 115, 302–337.
- Robson, Arthur J., 2001, The biological basis of economic behavior, *Journal of Economic Literature* 39, 11–33.

- Robson, Arthur J., and Lorne A. Whitehead, 2018, Adaptive hedonic utility, Working paper.
- Rustichini, Aldo, Katherine E. Conen, Xinying Cai, and Camillo Padoa-Schioppa, 2017, Optimal coding and neuronal adaptation in economic decisions, *Nature Communications* 8, 1–14.
- Schley, Dan R., and Ellen Peters, 2014, Assessing “economic value”: Symbolic-number mappings predict risky and riskless valuations, *Psychological Science* 25, 753–761.
- Soltani, Alireza, Benedetto De Martino, and Colin Camerer, 2012, A range-normalization model of context-dependent choice: A new model and evidence, *PLoS Computational Biology* 8, 1–15.
- 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.
- Stewart, Neil, Stian Reimers, and Adam J. L. Harris, 2015, On the origin of utility, weighting, and discounting functions: How they get their shapes and how to change their shapes, *Management Science* 61, 687–705.
- Stigler, George J., 1961, The economics of information, *Journal of Political Economy* 69, 213–225.
- Thaler, Richard, and Eric Johnson, 1990, Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice, *Management Science* 36, 643–660.
- Tobler, Philippe N., Christopher D. Fiorillo, and Wolfram Schultz, 2005, Adapting coding of reward value by dopamine neurons, *Science* 307, 1642–1645.
- Toubia, Olivier, Eric Johnson, Theodoros Evgeniou, and Philippe Delquié, 2013, Dynamic experiments for estimating preferences: An adaptive method of eliciting time and risk parameters, *Management Science* 59, 613–640.
- Tversky, Amos, and Daniel Kahneman, 1992, Advances in prospect theory: Cumulative representation of uncertainty, *Journal of Risk and Uncertainty* 5, 297–323.
- Wachter, Jessica A., and Michael J. Kahana, 2020, A retrieved-context theory of financial decisions, Working paper.

- Walasek, Lukasz, and Neil Stewart, 2015, How to make loss aversion disappear and reverse: Tests of the decision by sampling origin of loss aversion, *Journal of Experimental Psychology: General* 144, 7–11.
- Webb, Ryan, Paul W. Glimcher, and Kenway Louie, 2020, The normalization of consumer valuations: Context-dependent preferences from neurobiological constraints, *Management Science* forthcoming.
- Weber, Martin, and Colin Camerer, 1998, The disposition effect in securities trading: An experimental analysis, *Journal of Economic Behavior and Organization* 33, 167–184.
- 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.
- Wei, Xue-Xin, and Alan A. Stocker, 2017, Lawful relation between perceptual bias and discriminability, *Proceedings of the National Academy of Sciences* 114, 10244–10249.
- 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.
- Woodford, Michael, 2014, Stochastic choice: An optimizing neuroeconomic model, *American Economic Review* 104, 495–500.
- Woodford, Michael, 2020, Modeling imprecision in perception, valuation and choice, *Annual Review of Economics* forthcoming.
- Zimmermann, Jan, Paul W. Glimcher, and Kenway Louie, 2018, Multiple timescales of normalized value coding underlie adaptive choice behavior, *Nature Communications* 9, 3206.

Figure 1. Prior distributions, coding rules and the optimal likelihood functions

Panel A plots two uniform stimulus distributions of X , one with high volatility ($X_l = 8$ and $X_u = 32$) and the other with low volatility ($X_l = 16$ and $X_u = 24$). Panel B plots the coding rule $\theta(X)$, defined in equation (4) of the main text, for both volatility conditions. Panel C plots the implied likelihood function $f(R_x|X)$, defined in equation (9) of the main text, for two values, $X = 18$ and $X = 22$, and for each of the two stimulus distributions. The capacity constraint parameter n is set to 10.

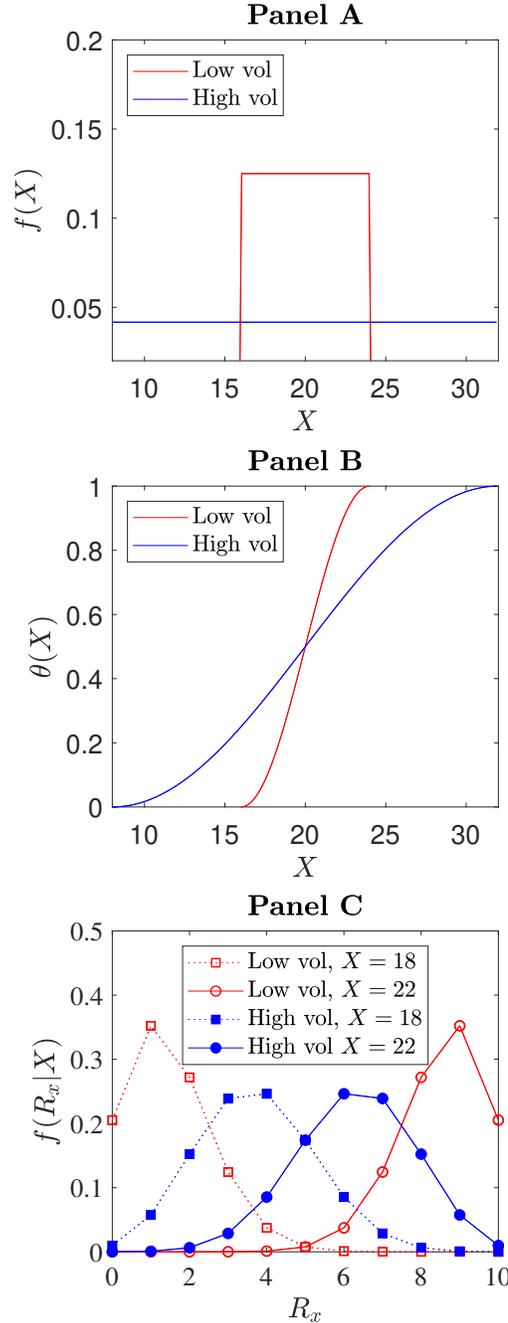


Figure 2. Value functions and the underlying stimulus distributions

Panel A: the upper graph plots two uniform stimulus distributions for X , one with high volatility ($X_l = 8$ and $X_u = 32$) and the other with low volatility ($X_l = 16$ and $X_u = 24$). The lower graph plots the subjective valuations implied by efficient coding, $v(X)$, and their one-standard-deviation bounds $v(X) \pm \sigma(X)$. Panel B: the upper graph plots two stimulus distributions, one monotonically increasing and one monotonically decreasing. These distributions are characterized by

$$f(X; X_l, X_u, \Delta) = \frac{1}{X_u - X_l} - \Delta + 2\Delta \cdot \frac{X - X_l}{X_u - X_l},$$

where $X_l = 8$ and $X_u = 32$. For the increasing distribution, Δ is set to $1/30$; and for the decreasing distribution, Δ is set to $-1/30$. The lower graph plots the subjective valuations implied by efficient coding, $v(X)$, and their one-standard-deviation bounds $v(X) \pm \sigma(X)$. The capacity constraint parameter n for both panels is set to 10. In the lower graph of each panel, the green dash-dot line is the forty-five degree line.

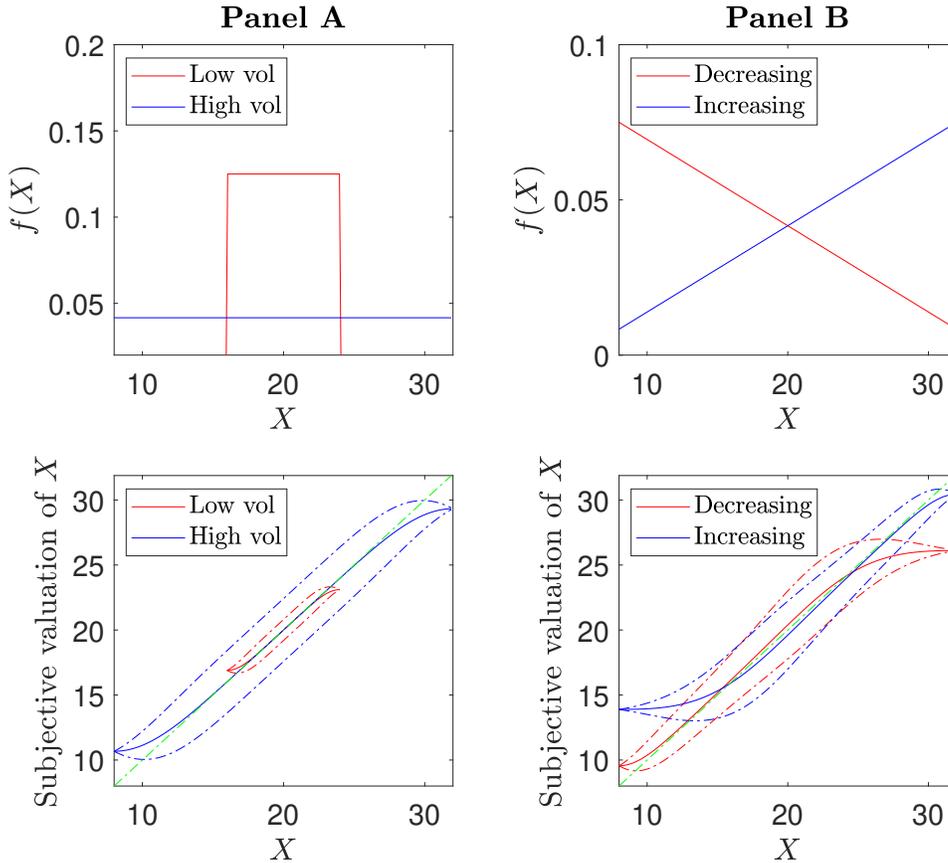


Figure 3. Model predicted probability of choosing the risky lottery

Panel A: the graph plots, for each of the two volatility levels (low volatility: $X_l = 16$, $X_u = 24$, $C_l = 8$, and $C_u = 12$; high volatility: $X_l = 8$, $X_u = 32$, $C_l = 4$, and $C_u = 16$), the probability of choosing the risky lottery (equation (14) of the main text). The stimulus distributions for X and C are uniform. The probability p for the risky lottery to pay X dollars is set to 0.5. The capacity constraint parameter n is set to 10. For each volatility condition, we draw X and C from their respective uniform distributions, and then compute the probability of risk taking for each level of $pX - C$. Panel B: for each of the two volatility conditions described above, the left graph plots the probability of choosing the risky lottery against each given level of X , integrated over values of C ; and the right graph plots the probability of choosing the risky lottery against each given level of C , integrated over values of X .

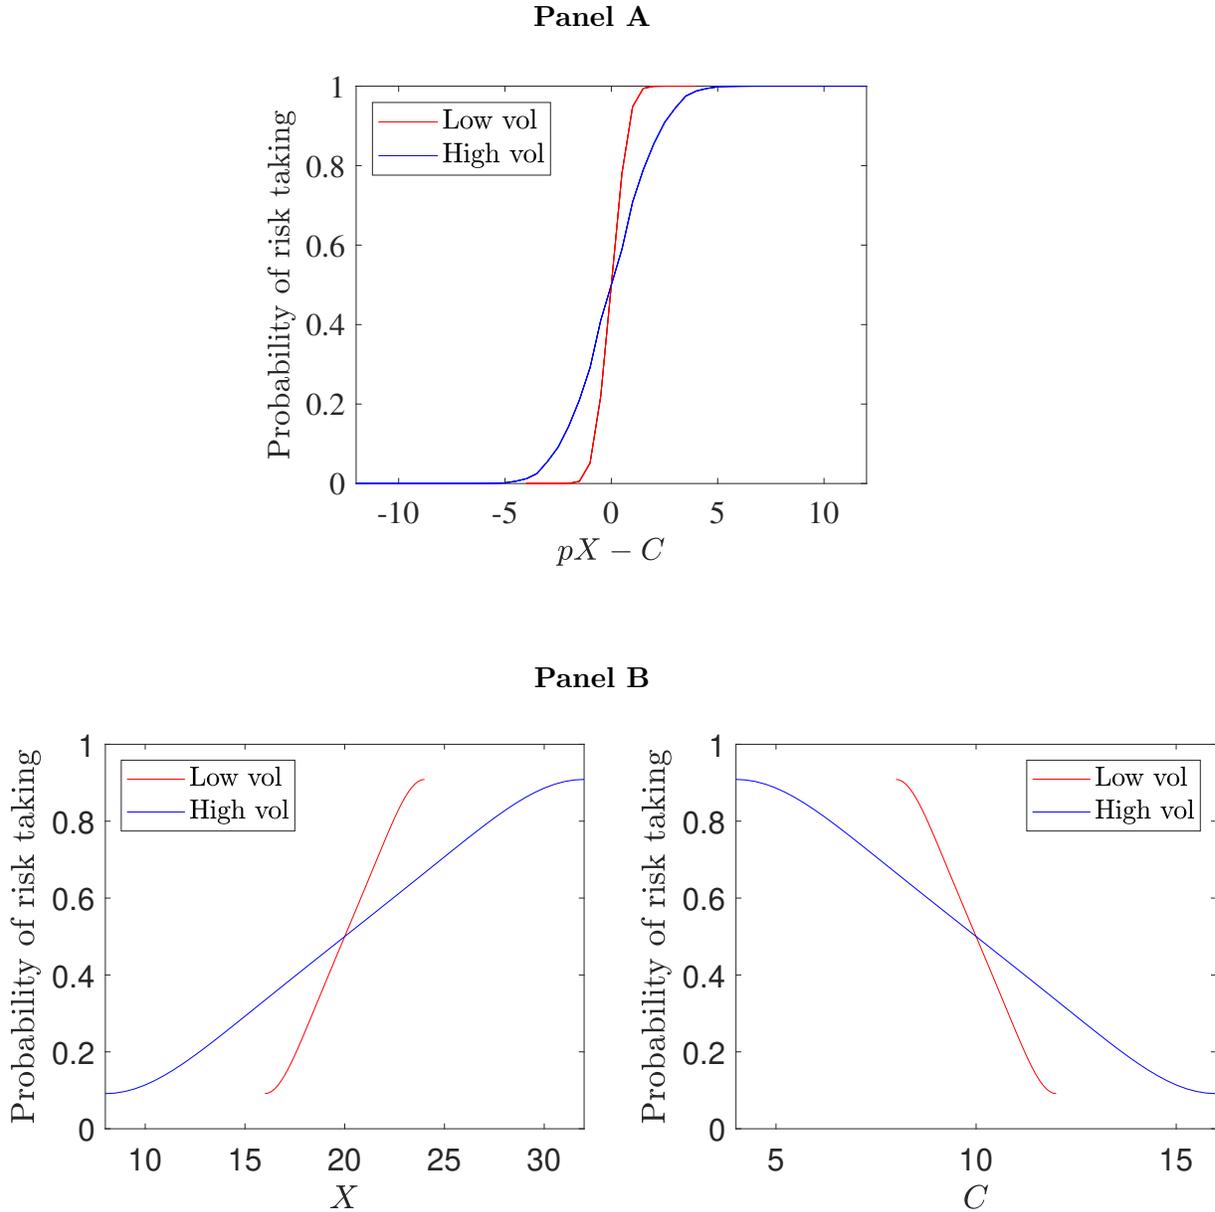


Figure 4. Experimental design for the risky choice task in Experiment 1

The task consists of two blocks of trials: one block contains trials from the high volatility condition while the other block contains trials from the low volatility condition. The order of the blocks is randomized across subjects. Each block begins with 30 “adapt trials,” followed by 270 “test trials.” Among the 270 test trials, we designate 30 “common trials” that are identical across both volatility conditions and are the basis of our main tests. On each trial, the subject has unlimited time to decide which of the two options she prefers.

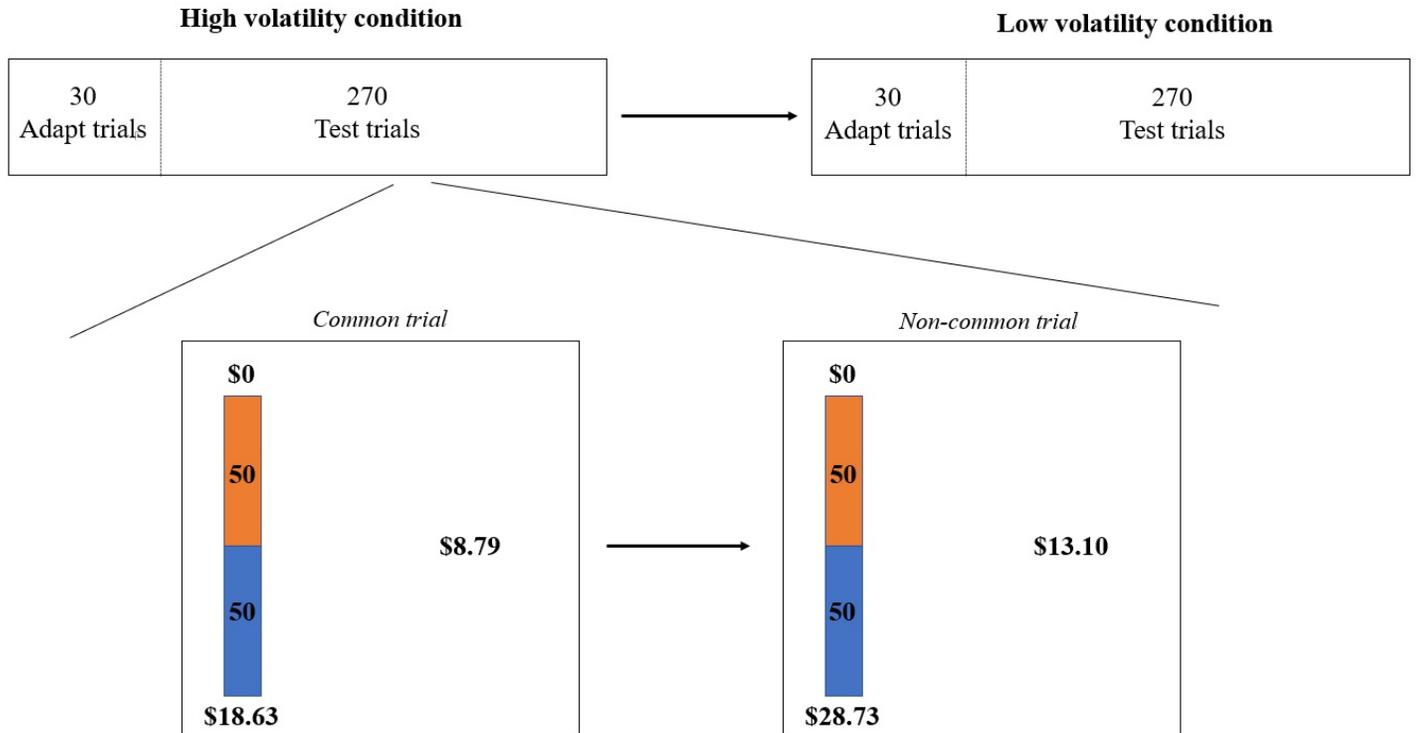


Figure 5. Average probability of risk taking across volatility conditions

Panel A: the probability of choosing the risky lottery against the difference in expected values between the risky lottery and the certain option, $pX - C$. Panel B: the probability of choosing the risky lottery against values of X (left) and values of C (right). The probability of choosing the risky lottery is computed as the proportion of trials on which subjects choose the risky lottery. Data are pooled across subjects over all test trials in the first condition, and thus represent between subjects comparisons. For each volatility condition, we bin the running variables— $pX - C$ in Panel A, and X or C in Panel B—to their nearest integer values, and plot the mean for each bin. The length of the vertical bar inside each data point denotes two standard errors of the mean. Standard errors are clustered by subject.

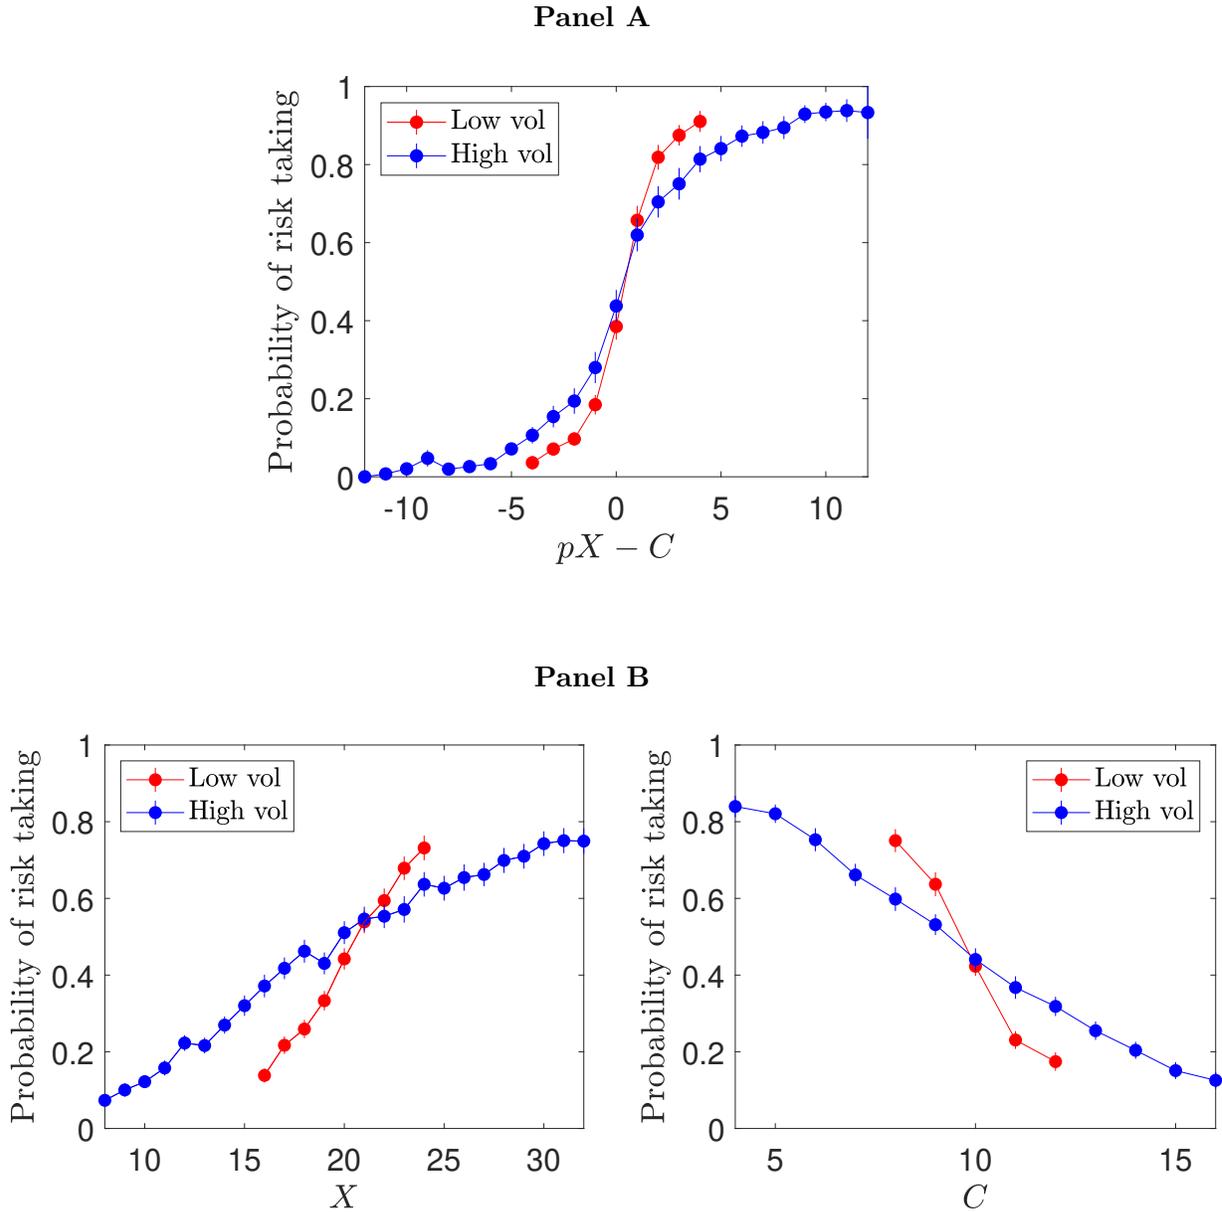


Figure 6. Experimental design for the perceptual choice task in Experiment 1

The task consists of two blocks of trials: one block contains trials from the high volatility condition while the other block contains trials from the low volatility condition. The order of the blocks is randomized across subjects. Each block begins with 60 “adapt trials,” followed by 340 “test trials.” Among the 340 test trials, we designate 90 “common trials” that are identical across both volatility conditions and are the basis of our main tests. On each trial, the subject is incentivized to classify whether the number shown on the screen is greater than or less than a reference level of 65.

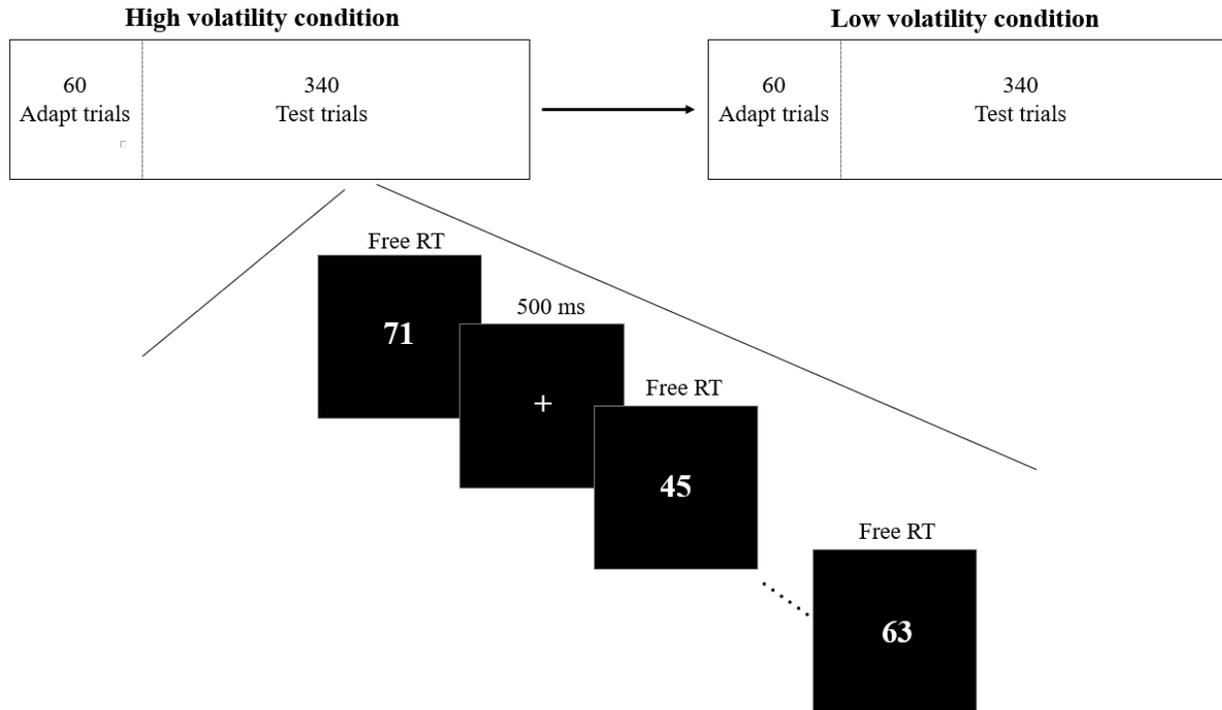


Figure 7. Classification performance and response time for the perceptual choice task

Panel A: the x -axis denotes the integer X that is presented on each trial. The y -axis denotes the proportion of trials for which subjects classified X as greater than 65. Panel B: the y -axis denotes the average response time for the subject to execute a decision, for those trials on which the subject responded correctly. Data are pooled across subjects over all test trials in the first condition, and thus represent between subjects comparisons. The length of the vertical bar inside each data point denotes two standard errors of the mean. Standard errors are clustered by subject.

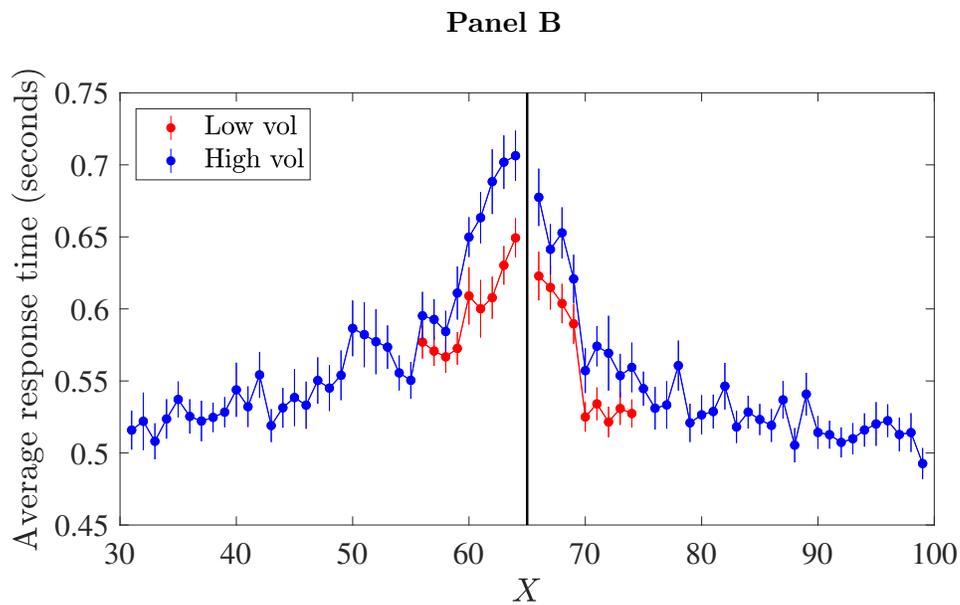
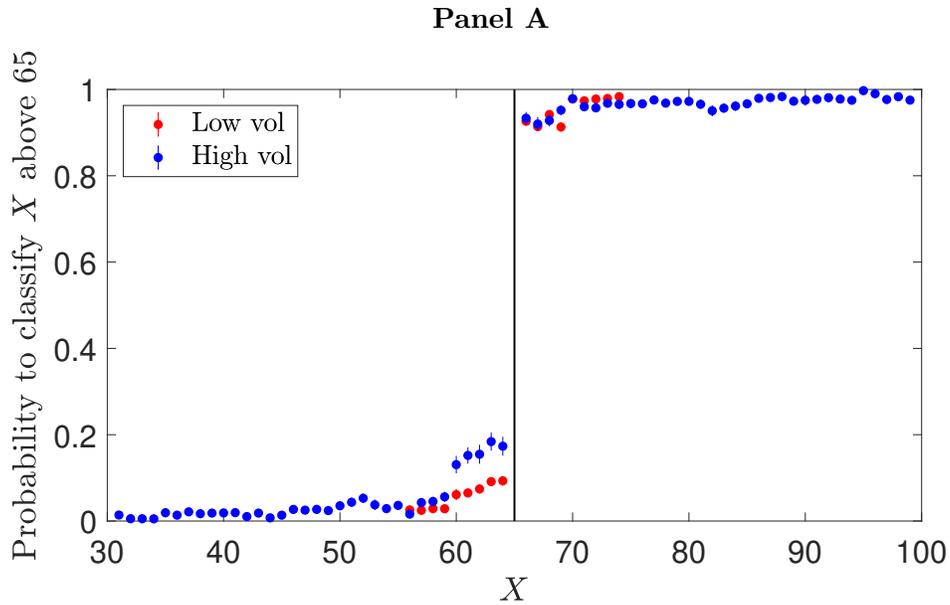


Figure 8. Out-of-sample model fits for the risky choice task in Experiment 1

The x -axis in each of the four panels is the model predicted probability of choosing the risky lottery. In all four panels, we plot data only from odd-numbered trials. In the upper two panels, we compute the predicted probabilities using parameters estimated from even-numbered trials, and data are either from the high volatility condition (left) or from the low volatility condition (right). The lower left panel pools data from both the high and low volatility condition. The lower right panel reports results from a counterfactual analysis where choice data are identical to those in the upper right panel, except the model predicted probability is generated under the counterfactual assumption that the prior is the high volatility distribution.

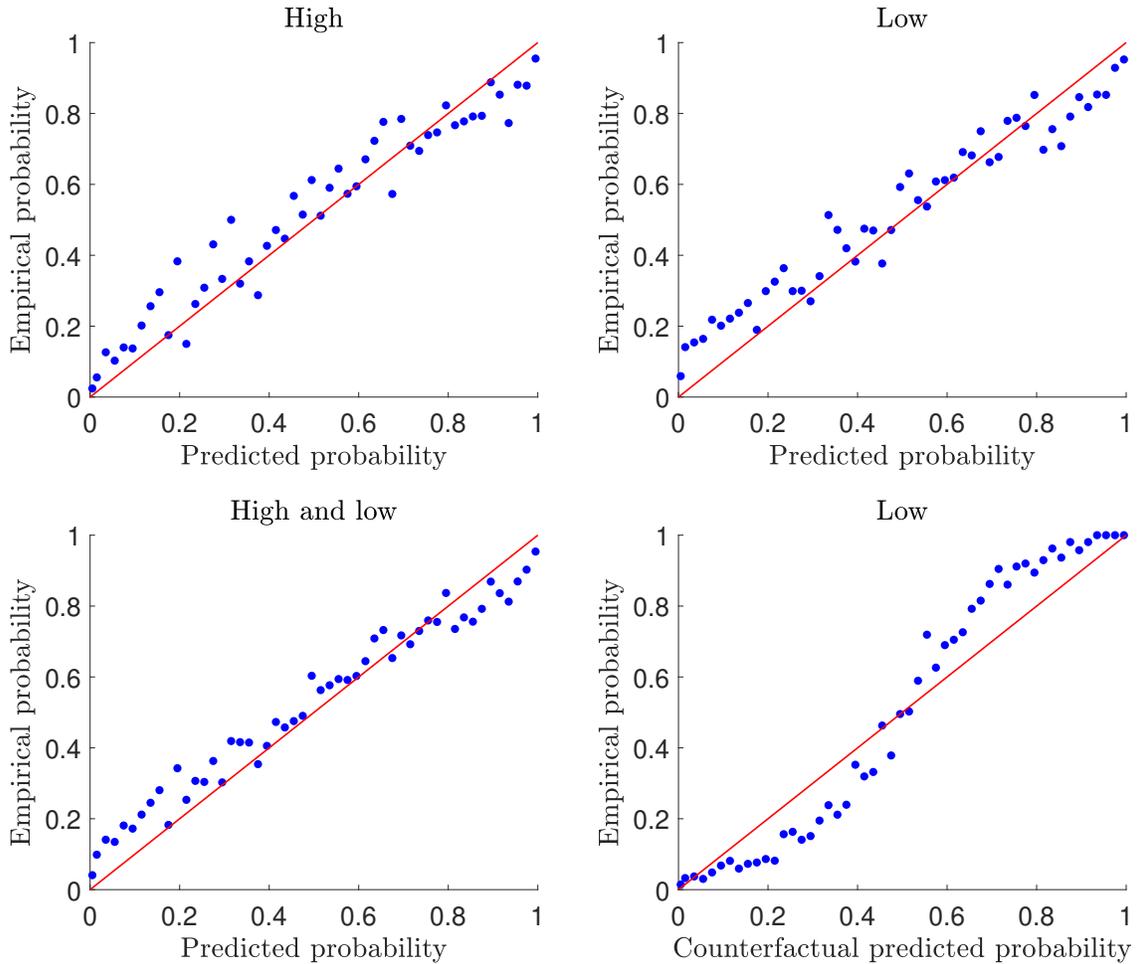


Figure 9. Efficient coding of probability

The upper graph plots the following distribution of probability

$$f(p; \lambda_1, \lambda_2, w) = K^{-1}[w \cdot e^{-\lambda_1 p} + (1 - w) \cdot e^{-\lambda_2(1-p)}],$$

where K , a normalization factor, is $(w/\lambda_1)(1 - e^{-\lambda_1}) + ((1-w)/\lambda_2)(1 - e^{-\lambda_2})$. The lower graph plots the implied probability weighting function, $v(p)$, and its one-standard-deviation bounds $v(p) \pm (p)$. The parameter values are: $\lambda_1 = 7.5$, $\lambda_2 = 8.5$, $w = 0.25$, and $n = 10$. The green dash-dot line is the forty-five degree line.

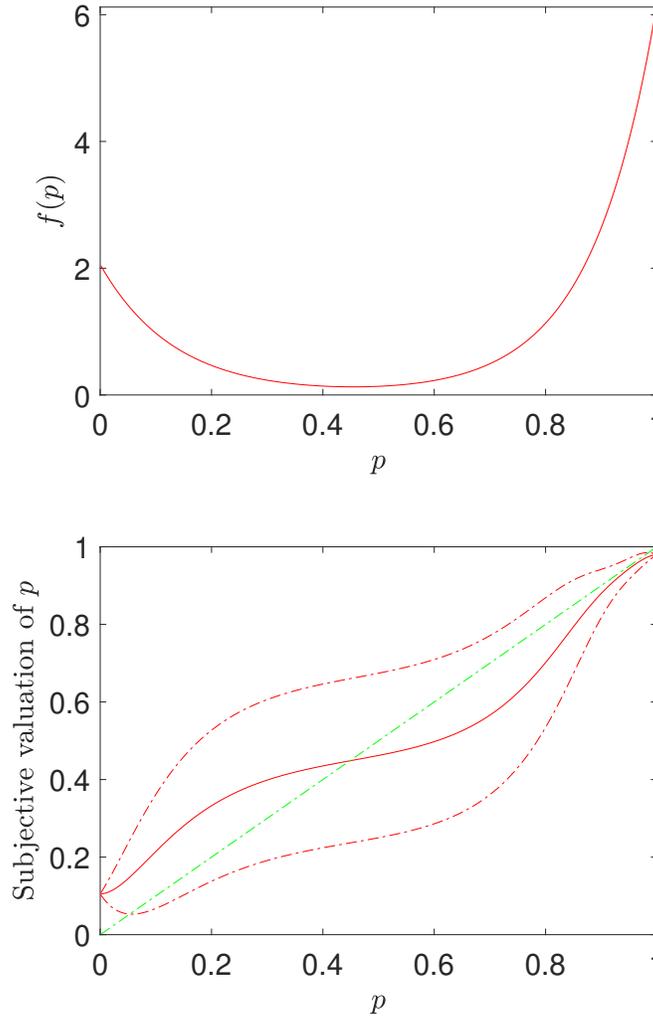


Table 1. Payoff values on the 30 common test trials (Experiment 1)

The table presents the payoff values that comprise the 30 common test trials in the risky choice task of Experiment 1. The set of 30 common test trials is presented to subjects once in each volatility condition. The order of the test trials is randomized across subjects.

<i>X</i>	<i>C</i>	<i>X</i>	<i>C</i>
\$17.30	\$11.20	\$21.29	\$8.82
\$17.31	\$9.62	\$21.30	\$10.39
\$17.32	\$8.04	\$21.34	\$11.23
\$17.34	\$8.75	\$21.34	\$9.58
\$17.38	\$10.38	\$21.38	\$8.00
\$18.63	\$8.79	\$22.62	\$9.62
\$18.64	\$11.16	\$22.66	\$8.79
\$18.68	\$9.61	\$22.67	\$8.04
\$18.68	\$8.03	\$22.69	\$11.21
\$18.72	\$10.44	\$22.71	\$10.42
\$19.96	\$9.62	\$23.96	\$9.55
\$19.98	\$8.01	\$23.97	\$11.80
\$19.99	\$8.81	\$23.98	\$10.37
\$20.03	\$10.40	\$23.98	\$8.03
\$20.04	\$11.22	\$23.99	\$8.84

Table 2. Probability of choosing the risky lottery in Experiment 1 (volatility manipulation)

The table reports results from mixed effects linear regressions where the dependent variable takes the value of one if the subject chooses the risky lottery, and zero otherwise. The dummy variable, *high*, takes the value of one if the trial belongs to the high volatility condition, and zero if it belongs to the low volatility condition. Only data from test trials are included. There are random effects on the independent variables *X*, *C*, and the intercept. Standard errors of the fixed effect estimates are reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

	Between subjects tests		Within subjects tests			
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable: “Choose risky lottery”	All data	Common trials only	Common trials only	Common trials only (w/out trials 301-450)	Common trials only—low vol first (w/out trials 301-450)	Common trials only—high vol first (w/out trials 301-450)
<i>high</i>	-0.156 (0.114)	0.001 (0.135)	-0.023 (0.092)	-0.061 (0.108)	0.016 (0.168)	-0.164 (0.155)
<i>X</i>	0.074*** (0.004)	0.074*** (0.004)	0.062*** (0.004)	0.064*** (0.004)	0.073*** (0.004)	0.053*** (0.006)
<i>C</i>	-0.173*** (0.007)	-0.186*** (0.007)	-0.164*** (0.008)	-0.170*** (0.007)	-0.186*** (0.008)	-0.155*** (0.011)
<i>X</i> × <i>high</i>	-0.044*** (0.005)	-0.023*** (0.005)	-0.006* (0.003)	-0.009** (0.004)	-0.015** (0.006)	-0.002 (0.006)
<i>C</i> × <i>high</i>	0.110*** (0.010)	0.049*** (0.010)	0.013** (0.007)	0.025*** (0.008)	0.030** (0.012)	0.019* (0.011)
Constant	0.668*** (0.084)	0.787*** (0.100)	0.827*** (0.081)	0.850*** (0.086)	0.801*** (0.104)	0.943*** (0.127)
Observations	37,612	4,170	8,257	6,411	3,125	3,286

Table 3. Probability of classifying X as greater than 65 in perceptual choice task

The table reports results from mixed effects logistic regressions where the dependent variable takes the value of one if the subject classifies the integer X as larger than 65, and it takes the value of zero otherwise. The integer X is drawn uniformly from the set $[31, 99] \setminus \{65\}$ in the high volatility condition, while X is drawn uniformly from the set $[56, 74] \setminus \{65\}$ in the low volatility condition. The dummy variable, *high*, takes the value of one if the trial belongs to the high volatility condition, and zero if it belongs to the low volatility conditions. There are random effects on the variable $X - 65$ and the intercept. Standard errors of the fixed effect estimates are reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

	Between subjects tests			Within subjects tests		
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable: “Classify X as greater than 65”	$56 \leq X \leq 74$	$60 \leq X \leq 69$	$56 \leq X \leq 59$ or $70 \leq X \leq 74$	$56 \leq X \leq 74$	$60 \leq X \leq 69$	$56 \leq X \leq 59$ or $70 \leq X \leq 74$
$X - 65$	0.855*** (0.037)	1.096*** (0.050)	0.578*** (0.018)	0.792*** (0.024)	1.051*** (0.033)	0.558*** (0.012)
$(X - 65) \times high$	-0.209*** (0.051)	-0.279*** (0.070)	-0.081*** (0.026)	-0.147*** (0.010)	-0.237*** (0.018)	-0.077*** (0.011)
<i>high</i>	0.288*** (0.054)	0.405*** (0.068)	-0.041 (0.112)	0.259** (0.035)	0.300** (0.043)	0.146** (0.072)
Constant	0.065** (0.029)	0.071* (0.037)	0.269*** (0.063)	0.036** (0.020)	0.077*** (0.024)	0.119*** (0.046)
Observations	31,230	15,522	15,708	63,210	31,580	31,630

Table 4. Payoff values on the 10 common test trials (Experiment 2)

The table presents the payoff values that comprise the 10 common test trials in Experiment 2. The set of 10 common test trials is presented to subjects once in each condition. The order of the test trials is randomized across subjects.

<i>X</i>	<i>C</i>
\$9.04	\$6.09
\$10.09	\$6.09
\$11.05	\$6.09
\$12.02	\$6.09
\$13.08	\$6.09
\$14.01	\$6.09
\$15.04	\$6.09
\$16.07	\$6.09
\$17.04	\$6.09
\$18.06	\$6.09

Table 5. Probability of choosing the risky lottery in Experiment 2 (shape manipulation)

The table reports results from mixed effects linear regressions where the dependent variable takes the value of one if the subject chooses the risky lottery, and zero otherwise. The dummy variable, *increasing prior*, takes the value of one if the trial belongs to the increasing prior condition, and zero if it belongs to the decreasing prior condition. The variable C is constant among common trials, and thus we do not include it as a control in Columns (2) and (3). Only data from test trials are included. There are random effects on the variables X , C , and the intercept. Standard errors of the fixed effect estimates are reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)
Dependent variable: "Choose risky lottery"	All data	Common trials only	Common trials only and $X < 15.79$
<i>increasing prior</i>	0.025*** (0.005)	0.054** (0.024)	0.068** (0.029)
X	0.022*** (0.001)	0.060*** (0.006)	0.073*** (0.008)
C	-0.063*** (0.002)		
Constant	0.594*** (0.322)	-0.400*** (0.080)	-0.545*** (0.096)
Observations	22,835	764	538

Appendices

A. Theoretical Derivations

HWP have derived coding rules under three different performance objectives: one that maximizes mutual information, one that maximizes accuracy, and one that maximizes expected financial gain. In this section, we closely follow HWP and prove that, when the conditions in equation (8) of the main text

$$\begin{aligned} & (i) \text{ draws of } X \text{ and } C \text{ are independent} \\ & \text{and } (ii) \text{ } pX \text{ and } C \text{ are identically and uniformly distributed} \end{aligned} \tag{A.1}$$

are satisfied, the three coding rules are equivalent. First, if the performance objective is to maximize mutual information between a payoff and its noisy signal, then the coding rules $\theta(X)$ and $\theta(C)$ in equations (4) and (5) of the main text come directly from HWP.

Next, if the performance objective is to maximize accuracy—in the sense of choosing the option with the highest objective expected value—then, under the conditions in (A.1), we make two conjectures

$$\begin{aligned} & (i) \mathbb{E}[p\tilde{X}|R_x], \text{ viewed as a function of } R_x, \text{ and } \mathbb{E}[\tilde{C}|R_c], \text{ viewed as a function of } R_c, \\ & \text{are identical functions,} \\ & \text{and } (ii) \text{ the optimal coding rules are related: } \theta(X) = \theta(C = pX), \forall X \in [X_l, X_h]. \end{aligned} \tag{A.2}$$

We observe that maximizing accuracy is equivalent to minimizing the probability of error

$$\text{Prob}_{error} \equiv \int_{C_l}^{C_h} dC \int_{X_l}^{X_h} f(X, C) \mathbb{P}\text{rob}(error|\theta(X), \theta(C)) dX, \tag{A.3}$$

where $\mathbb{P}\text{rob}(error|\theta(X), \theta(C))$ represents the probability that the *DM* chooses the option with the lower expected payoff observed by the econometrician. Given the two conjectures from (A.2), this probability error equals the probability that $R_x - R_c$ and $\theta(X) - \theta(C)$ are of the opposite sign. Combining (A.3) with the two conjectures, we observe that when n is large,

$$\begin{aligned} & \mathbb{P}\text{rob}(error|\theta(X), \theta(C)) \\ & = \mathbb{P}\text{rob}(error|\theta(Y), \theta(C)) \approx \Phi \left(-\frac{|\theta(Y) - \theta(C)|}{\sqrt{\frac{\theta(Y)(1-\theta(Y) + \theta(C)(1-\theta(C)))}{n}}} \right), \end{aligned} \tag{A.4}$$

where $Y \equiv pX$, so Y and C are independently and identically distributed. Moreover, (A.3) can be

written as

$$\begin{aligned} \mathbb{P}\text{rob}(\text{error}|\theta(X),\theta(C)) &= \int_{C_l}^{C_h} dC \int_{X_l}^{X_h} f(X)f(C)\mathbb{P}\text{rob}(\text{error}|\theta(X),\theta(C)) dX \\ &= \int_{C_l}^{C_h} dC \int_{C_l}^{C_h} f(Y)f(C)\mathbb{P}\text{rob}(\text{error}|\theta(Y),\theta(C)) dY. \end{aligned} \tag{A.5}$$

Given (A.5), the derivation of the optimal coding rules $\theta(X)$ and $\theta(C)$ follow directly from the Appendix of HWP; these coding rules are identical to those when the performance objective is to maximize mutual information. Given the coding rules, verifying the two conjectures from (A.2) is straightforward.

Finally, if the performance objective is to maximize expected financial gain, then with the two conjectures from (A.2), derivations from HWP imply the following coding rules

$$\theta(X) = \left(\sin \left[\frac{\pi}{2} \frac{\int_{X_l}^X f(x)^{2/3} dx}{\int_{X_l}^{X_h} f(x)^{2/3} dx} \right] \right)^2, \quad \theta(C) = \left(\sin \left[\frac{\pi}{2} \frac{\int_{C_l}^C f(c)^{2/3} dc}{\int_{C_l}^{C_h} f(c)^{2/3} dc} \right] \right)^2. \tag{A.6}$$

Verifying the two conjectures is again straightforward. Moreover, when X and C are uniformly distributed, the coding rules in (A.6) reduce to equations (4) and (5) of the main text. That is, maximizing expected financial gain and maximizing mutual information give rise to the same coding rules. ■

B. Experimental Instructions and Pre-Registration Documents

1. Instructions for the risky choice task in Experiment 1

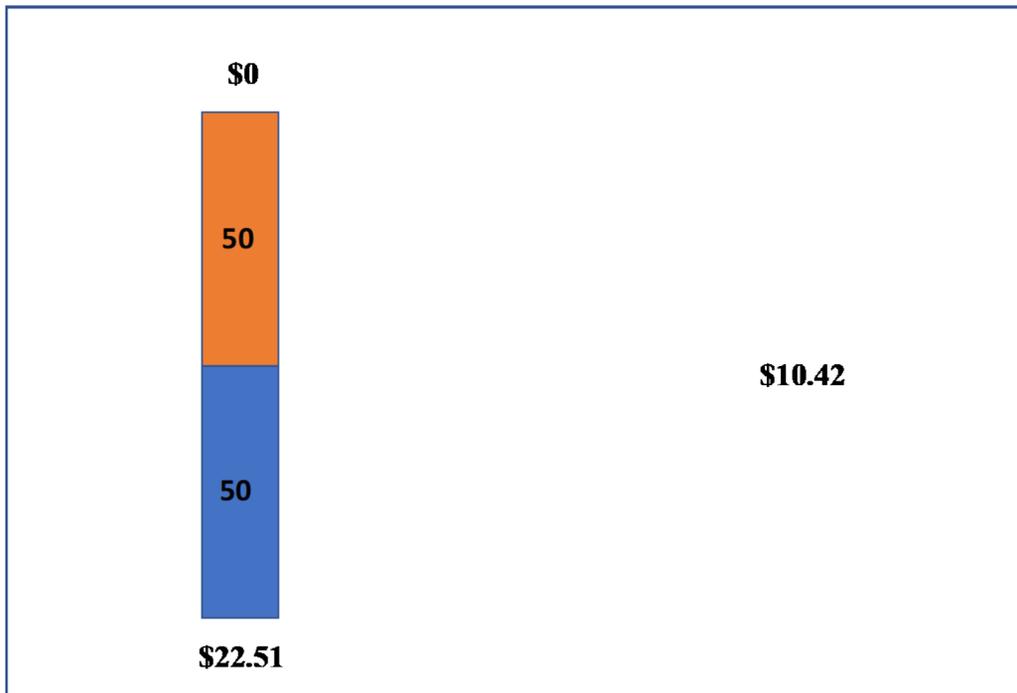
Experiment Instructions

Thank you for participating in this experiment. Before we begin, please turn off all cell phones and put all belongings away. For your participation, you have already earned \$7, and you will have the opportunity to earn more money depending on your answers during the experiment.

The experiment consists of **two phases**. The instructions for Phase I are given below. After you go through Phase I, you will be given a new set of instructions for Phase II.

Phase I

In Phase 1 of the experiment, you will be asked to make a series of decisions about choosing a “risky gamble” or a “sure thing”. The risky gamble will pay a positive amount with 50% chance, and \$0 with 50% chance. The amount shown for the sure thing will be paid with 100% chance, if chosen. Below is an example screen from the experiment:



In the above example, the risky gamble is on the LEFT side of the screen and the sure thing is on the RIGHT side of the screen. The risky gamble pays \$22.51 with 50% chance, and \$0 with 50% chance. The sure thing pays \$10.42 with 100% chance. For each question, you will be asked to select one of the two options for each question by pressing either the “z” key for the LEFT option or the “?” key for the RIGHT option. On some questions, the risky gamble will be on the LEFT, and other questions it will be on the RIGHT. **Phase I of the experiment is broken down into 12 parts, and each part contains 50 questions.**

At the end of the experiment (after both Phase I and Phase II are completed), one trial will be randomly selected, and you’ll be paid according to your decision on that trial. For example, if the above trial was

chosen, and you selected the sure thing you would be paid \$10.42. If instead you chose the risky gamble, you'd be paid either \$0 or \$22.51, depending on which outcome the computer randomly selects.

Therefore, you should choose the option on each question that you prefer, since it may end up being the question that you are actually paid for. Remember that all earnings for Phase I and Phase II will be added to your \$7 show-up fee.

Before you begin Phase I, you will see a set of 10 practice trials so you can become familiar with the software and have a chance to ask any questions. These 10 practice trials will not count towards your actual payment.

When you are ready to begin the practice trials, press "Enter" on the computer ONCE. When you are finished with Phase I, please wait quietly and raise your hand. An experimenter will then come give you instructions for Phase II.

If you have any questions during the experiment, please raise your hand quietly.

2. *Instructions for the perceptual choice task in Experiment 1*

Phase II

In Phase II of the experiment, you will see a series of numbers and will be asked to classify whether each number is greater than or smaller than 65. If the number displayed is less than 65, press the “z” key. If the number displayed is greater than 65, press the “?” key. At the end of the experiment, you will be paid depending on the speed and accuracy of your classifications (in addition to your earnings from Phase I and the show-up fee). Specifically, you will be paid:

$$\text{Payout} = \$ (15 \times \text{accuracy} - 10 \times \text{avgseconds}),$$

where “accuracy” is the percentage of trials where you correctly classified the number as larger or smaller than 65. “avgseconds” is the average amount of time it takes you to classify a number throughout the experiment, in seconds. For example, if you correctly classified the number on all trials and it took you 0.3 seconds to respond to each question, you would earn $\$(15 \times 100\% - 10 \times 0.3) = \12.00 . If instead you only classified 80% of the questions correctly and took 0.8 seconds to respond to each question, you would be paid $\$(15 \times 70\% - 10 \times 0.8) = \2.50 . Therefore, you will make the most money by answering as quickly and as accurately as possible.

The experiment will be separated into sixteen parts, and each part will contain 50 trials. In between each part, you can take a short break, and then continue at your own pace.

When you are ready to begin Phase II, press “Enter” on the computer ONCE.

When you finish all sixteen parts, please quietly raise your hand and an experimenter will come give you payment instructions.

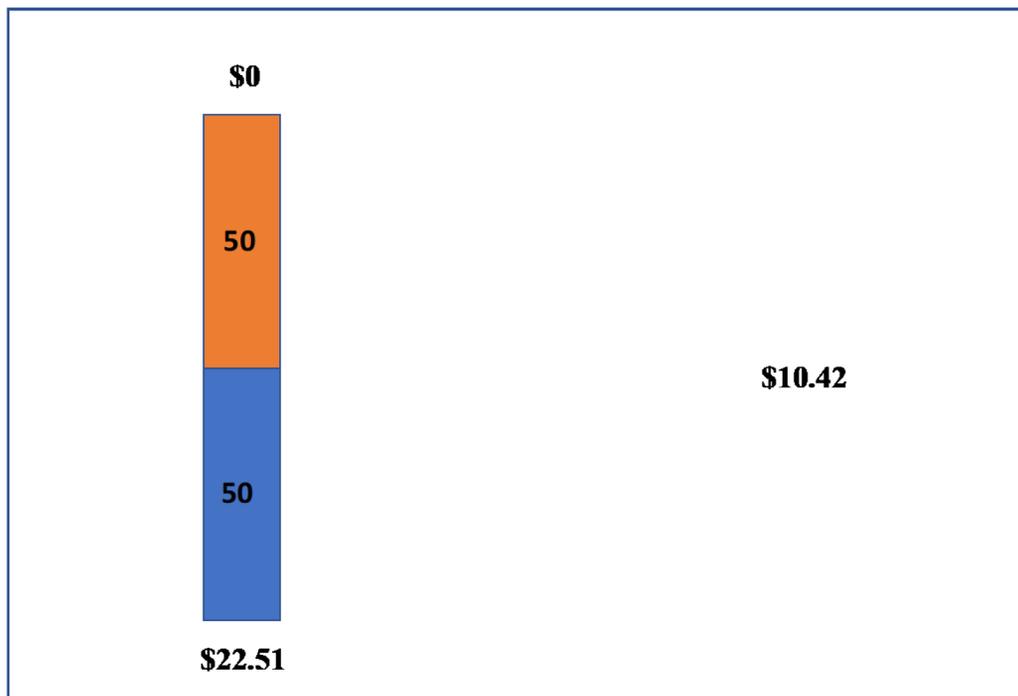
If you have any questions during the experiment, please raise your hand quietly.

3. *Instructions for the risky choice task in Experiment 2*

Experiment Instructions

Thank you for participating in this experiment. Before we begin, please turn off all cell phones and put all belongings away. For your participation, you have already earned \$7, and you will have the opportunity to earn more money depending on your answers during the experiment.

In this experiment you will be asked to make a series of decisions about choosing a “risky gamble” or a “sure thing”. The risky gamble will pay a positive amount with 50% chance, and \$0 with 50% chance. The amount shown for the sure thing will be paid with 100% chance, if chosen. Below is an example screen from the experiment:



In the above example, the risky gamble is on the LEFT side of the screen and the sure thing is on the RIGHT side of the screen. The risky gamble pays \$22.51 with 50% chance, and \$0 with 50% chance. The sure thing pays \$10.42 with 100% chance. For each question, you will be asked to select one of the two options for each question by pressing either the “z” key for the LEFT option or the “?” key for the RIGHT option. On some questions, the risky gamble will be on the LEFT, and other questions it will be on the RIGHT. **The experiment is broken down into 12 parts, and each part contains 50 questions.**

At the end of the experiment, one trial will be randomly selected, and you’ll be paid according to your decision on that trial. For example, if the above trial was chosen, and you selected the sure thing you would be paid \$10.42. If instead you chose the risky gamble, you’d be paid either \$0 or \$22.51, depending on which outcome the computer randomly selects. **Therefore, you should choose the option on each question that you prefer, since it may end up being the question that you are actually paid for.** Remember that your earnings will be added to your \$7 show-up fee.

Before you begin, you will see a set of 10 practice trials so you can become familiar with the software and have a chance to ask any questions. These 10 practice trials will not count towards your actual payment.

When you are ready to begin the practice trials, press “Enter” on the computer ONCE. **When you are finished with the experiment, please wait quietly and raise your hand. An experimenter will then come give you instructions for payment.**

If you have any questions during the experiment, please raise your hand quietly.

4. Pre-registration documents for Experiment 1

CONFIDENTIAL - FOR PEER-REVIEW ONLY



EC Risk Taking and Numerical Comparison - Feb 2020 (#35331)

Created: 02/09/2020 11:09 PM (PT)

Shared: 07/17/2020 03:26 PM (PT)

This pre-registration is not yet public. This anonymized copy (without author names) was created by the author(s) to use during peer-review. A non-anonymized version (containing author names) will become publicly available only if an author makes it public. Until that happens the contents of this pre-registration are confidential.

1) Have any data been collected for this study already?

No, no data have been collected for this study yet.

2) What's the main question being asked or hypothesis being tested in this study?

We investigate whether human subjects make decisions about monetary lotteries and number comparisons in a manner that is consistent with theories of efficient coding.

3) Describe the key dependent variable(s) specifying how they will be measured.

There are two main tasks in the experiment. There is a "Risky choice" task in which subjects will choose between a risky option and a certain option. In this task, the dependent variable is the decision to choose the risky lottery. Second, there is a "Number classification" task in which the subject classifies whether a number is greater than or less than the number "65." In this task, the dependent variable is whether the subject accurately classified the number on a given trial. We will also collect response times for each trial and each task.

4) How many and which conditions will participants be assigned to?

In both tasks, subjects will be assigned to two conditions: a "high volatility" and a "low volatility" condition. "High" and "Low" refer to the volatility of the distribution from which monetary amounts or numerical quantities are drawn in each of the two tasks. We will randomize, at the subject level, whether the first condition is the "high" or "low" volatility condition; this enables us to test for both within and between subjects variation.

In the risky choice task, the high and low volatility distributions are uniform with the same mean, but the high volatility distribution has larger volatility. Each condition begins with 30 "initial adapt" trials, and for each condition we only analyze data after these "initial adapt" trials. Our main focus of analysis will be on 30 choice sets that are identical across conditions, which have the form (X, 0.5; 0, 0.5) vs. (C, 1). We call these "test trials," and the 30 different values of (X,C) are given by:

(17.30, 11.20)
(17.31, 9.62)
(17.32, 8.04)
(17.34, 8.75)
(17.38, 10.38)
(18.63, 8.79)
(18.64, 11.16)
(18.68, 9.61)
(18.68, 8.03)
(18.72, 10.44)
(19.96, 9.62)
(19.98, 8.01)
(19.99, 8.81)
(20.03, 10.4)
(20.04, 11.22)
(21.29, 8.82)
(21.3, 10.39)
(21.34, 11.23)
(21.34, 9.58)
(21.38, 8.00)
(22.62, 9.62)
(22.66, 8.79)
(22.67, 8.04)
(22.69, 11.21)
(22.71, 10.42)
(23.96, 9.55)
(23.97, 11.18)
(23.98, 10.37)



AS PREDICTED

(23.98, 8.03)

(23.99, 8.84)

The design for the number classification task is nearly identical, except we use 60 "initial adapt" trials in each condition, and the test trials are the integers that fall in the support of the low volatility distribution, [56, 74].

5) Specify exactly which analyses you will conduct to examine the main question/hypothesis.

For each of the two tasks, we will test whether responses on test trials in the high volatility condition exhibit more noise compared to responses on test trials in the low volatility condition. Because responses are binary in both tasks, we will use logistic regressions where the main independent variables are X and C in the risky choice task and the main independent variable in the number classification task is (X-65). Between subjects tests will be conducted only on test trials in the first condition. Within subjects test will be conducted on all test trials, across conditions.

6) Describe exactly how outliers will be defined and handled, and your precise rule(s) for excluding observations.

We will exclude any subject who exhibits no variation in choice behavior in either of the two tasks.

7) How many observations will be collected or what will determine sample size? No need to justify decision, but be precise about exactly how the number will be determined.

We will collect N=150 subjects, where each subject completes the risky choice and number classification task.

8) Anything else you would like to pre-register? (e.g., secondary analyses, variables collected for exploratory purposes, unusual analyses planned?)

As an additional analysis, we will test for a correlation in behavior across the number comparison and risky choice tasks. We will also test whether response times are longer for more "difficult" decisions, and whether response times are longer on test trials in the high volatility condition.

5. Pre-registration documents for Experiment 2

CONFIDENTIAL - FOR PEER-REVIEW ONLY



EC Risk Taking with monotonic distributions - Feb 2020 (#36412)

Created: 02/27/2020 11:16 AM (PT)

Shared: 07/17/2020 03:28 PM (PT)

This pre-registration is not yet public. This anonymized copy (without author names) was created by the author(s) to use during peer-review. A non-anonymized version (containing author names) will become publicly available only if an author makes it public. Until that happens the contents of this pre-registration are confidential.

1) Have any data been collected for this study already?

No, no data have been collected for this study yet.

2) What's the main question being asked or hypothesis being tested in this study?

Does the shape of the distribution from which payoffs are drawn affect risk taking, in a manner consistent with efficient coding?

3) Describe the key dependent variable(s) specifying how they will be measured.

Subjects will choose between a risky option and a certain option, where each choice set has the form $(X, 0.5; 0, 0.5)$ vs. $(C, 1)$. The dependent variable is the decision to choose the risky lottery. We will also collect response times for each decision.

4) How many and which conditions will participants be assigned to?

There will be two conditions, an "increasing" and a "decreasing" condition. In the increasing (decreasing) condition, the payoff X will be drawn from a distribution which is linearly increasing (decreasing) over $[8, 32]$. The distribution of C will be uniform over $[2, 18]$ in both conditions. Each subject will answer questions in both conditions, and thus our test will be within subjects. Each condition contains three hundred trials, and in each condition, we will insert 10 "test trials" which have the following values for (X, C) :

(9.04, 6.09)
(10.90, 6.09)
(11.05, 6.09)
(12.02, 6.09)
(13.08, 6.09)
(14.01, 6.09)
(15.04, 6.09)
(16.07, 6.09)
(17.04, 6.09)
(18.06, 6.09)

5) Specify exactly which analyses you will conduct to examine the main question/hypothesis.

Our main analysis will be to compare the average level of choosing the risky lottery in the increasing condition versus the decreasing condition.

6) Describe exactly how outliers will be defined and handled, and your precise rule(s) for excluding observations.

None.

7) How many observations will be collected or what will determine sample size? No need to justify decision, but be precise about exactly how the number will be determined.

We will collect $N=40$ subjects.

8) Anything else you would like to pre-register? (e.g., secondary analyses, variables collected for exploratory purposes, unusual analyses planned?)

None

Addendum to Preregistration: "EC Risk Taking with monontonic distributions – Feb 2020"

In our response to question (4) on the pre-registration on the previous page, we explicitly listed the payoff values for the 10 common test trials. There is a typo in the 2nd line: the value for X was inadvertently listed as "10.90" when it should have been listed as "10.09".

C. Response Times in the Volatility Experiment

Here we present data on response times from the risky choice task of Experiment 1, our volatility manipulation experiment. As indicated in the main text, we find that on common trials, response times are significantly shorter in the low volatility condition, compared to the high volatility condition. Panel A of Figure C1 shows that response times in both conditions become longer as the absolute difference in expected values, $|pX-C|$, approaches zero. This result is consistent with a large literature on sequential sampling models in which response times become longer as the absolute difference between values of two alternatives gets smaller (Krajbich et al., 2010; Clithero, 2018).

[Place Figure C1 about here]

More importantly, for our tests of efficient coding, we find that response times are shorter across the distribution of $pX-C$ in the low volatility condition, compared to the high volatility condition. This suggests that the “sensitivity” effects we observe from the choice data—Panel B of Figure C1 reproduces Panel A of Figure 5—represent a conservative estimate of efficient coding. In other words, if we adjust for information processing time, which is a fraction of the total response time, then the difference in choice sensitivity across conditions would likely be larger.

Not only are response times helpful for interpreting the choice data, but they can be used to understand how subjects adapt to a new payoff distribution. Specifically, we examine the time series evolution of response times over the course of the risky choice task. Recall that in the risky choice task, there are 300 trials in the first condition and 300 trials in the second condition. The order of the conditions is randomized across subjects. Figure C2 shows the response time data, disaggregated by which condition a subject experienced first. The upper panel plots, for each of the 600 trials, the response time averaged across subjects who first went through the low volatility condition, followed by the high volatility condition. We see a spike in response time at trial 301, which is the beginning of the high volatility condition. As we mentioned in the main text when presenting the results from Table 2, the spike in response time may be due to the fact that subjects begin to experience novel and hence salient payoffs that they had not seen in the first 300 trials. As a result, these payoffs signal a change in environment and cause subjects to take longer to make decisions in the new environment.

[Place Figure C2 about here]

In contrast, the lower panel plots the response time averaged across subjects who first experience the high volatility condition and then go through the low volatility condition. In this cut of the data, we find no corresponding spike in response time at the beginning of the second condition. Here, subjects do not observe salient “outlier” payoffs: every payoff in the (low volatility) second condition is in the support of the payoff distribution from the recently experienced (high volatility) first condition. We speculate that this extra difficulty in adapting to the low volatility distribution

may be partly responsible for our within subject results in Table 2: the coefficients on the interaction terms, $X \times high$ and $C \times high$, have smaller magnitudes in Column (6) of Table 2, compared to those in Column (5).

Figure C1. Response times and choice data for the risky choice task

Panel A: the x -axis denotes the difference in expected values between the risky lottery and the certain option, $pX - C$. The y -axis denotes the average response time (in seconds). Data are pooled across subjects over all test trials in the first condition; the initial 30 adapt trials are excluded. Panel B: the probability of choosing the risky lottery against $pX - C$. This panel is identical to Panel A of Figure 5.

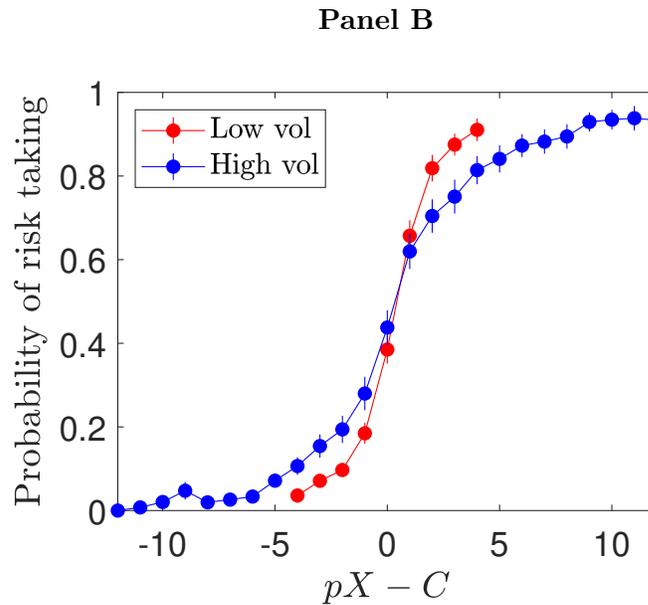
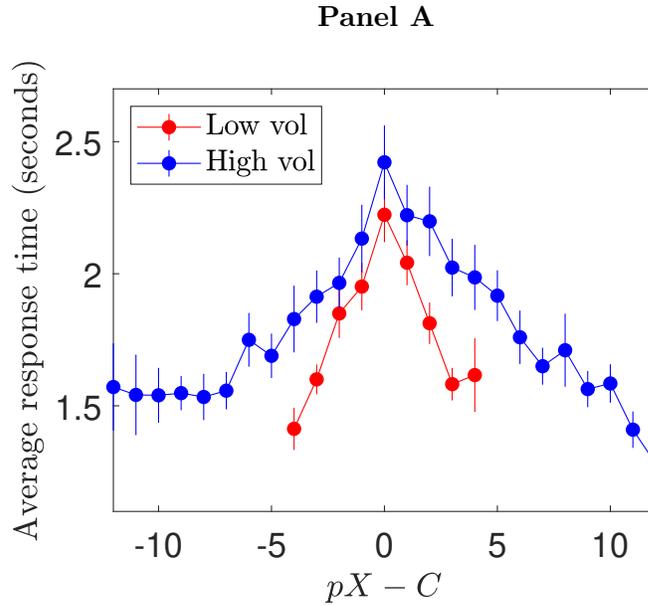
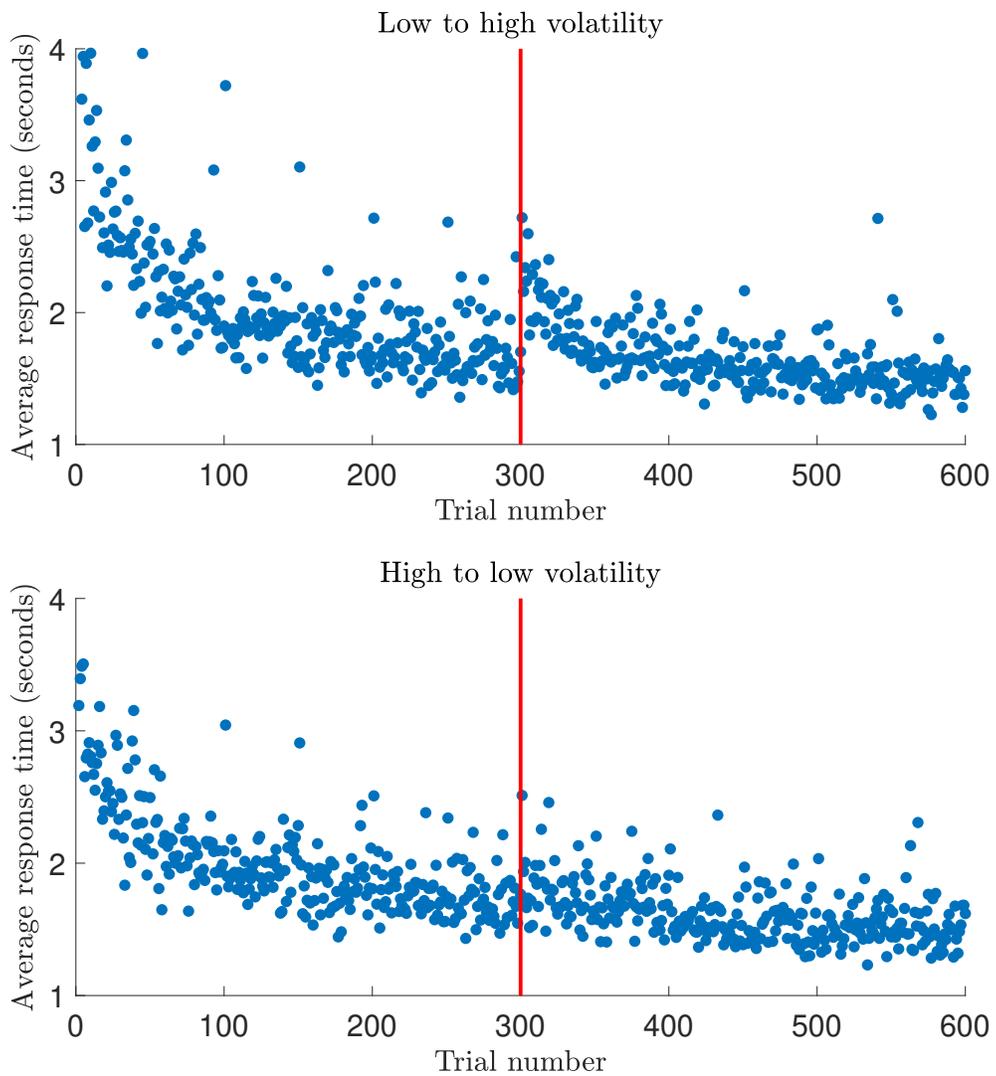


Figure C2. Average response time across the 600 trials from the risky choice task

Each blue dot represents a trial-specific response time averaged across subjects. The upper panel presents data from subjects who experienced the low volatility condition first. The lower panel presents data from subjects who experienced the high volatility condition first. The red vertical line denotes the onset of the change in environment.



D. Predictions of Alternative Models

1. Expectations-based reference points: *Kőszegi and Rabin (2006, 2007)*

We consider predictions of expectations-based reference points as proposed by [Kőszegi and Rabin \(2006, 2007\)](#) (KR). In this model, the *DM* evaluates a lottery by comparing each of its individual outcomes to a (possibly stochastic) reference payoff. Specifically, suppose the *DM* expects the reference point distribution to be $G(r)$, and suppose the lottery F the *DM* evaluates has N potential outcomes, x_1, x_2, \dots, x_N ; outcome x_n is associated with probability p_n . Then, the utility of payoff x_n relative to the stochastic reference point G can be written as

$$v(x_n|G) = \int (x_n + \mu(x_n - r))dG(r), \tag{D.1}$$

where

$$\mu(y) = \begin{cases} \eta \cdot y & \text{if } y \geq 0 \\ (\eta\lambda) \cdot y & \text{if } y < 0 \end{cases}, \tag{D.2}$$

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

KR propose that expectations are given by the *DM*'s rational expectations. As a first pass, one can assume that whenever the *DM* considers a lottery F , she views the lottery itself as the reference point, and thus the overall utility of F is $V(F|F) = \sum_{n=1}^N p_n v(x_n|F)$. KR denote such a reference point specification as the ‘‘choice acclimating personal equilibrium.’’ Under this specification, the valuation of a lottery—and hence risk taking behavior—depends only on the payoff distribution of that particular lottery, which is held constant across different experimental conditions (in both Experiment 1 and Experiment 2). Thus, this specification of KR cannot explain our main experimental results.

Below we further consider two alternative versions of the reference point distribution under KR:

1. The reference point for the risky lottery is a mixed prior distribution of X and zero: with 50% probability, the reference point is drawn from the prior distribution of X , and with the remaining 50% probability, it takes the value of zero. The reference point for the certain option is drawn from the uniform distribution $[C_l, C_h]$.
2. The reference point for X is drawn from the prior distribution of X . The reference point for zero takes the value of zero. The reference point for the certain option is drawn uniformly from $[C_l, C_h]$.

[Place Table [D1](#) about here]

Here we focus on testing whether KR can make sense of the main result from Experiment 2—that demand for the risky lottery is higher in the increasing condition, compared to the decreasing

condition.³⁸ Table D1 computes, for each of the 10 common trials from Experiment 2 and for each of the two experimental conditions, the difference between the KR value of the risky lottery and the KR value of the certain option. If this difference is positive, the *DM* chooses the risky lottery; conversely, if this difference is negative, the *DM* chooses the certain option. For these computations, we set $\eta = 1$ and $\lambda = 2$.

Table D1 shows that, under both versions of the reference point distribution, the relative valuation of the risky lottery is lower in the increasing condition, which is the *opposite* of our main result from Experiment 2. In the increasing condition, the reference point is more likely to be above the payoff value of X . Thus, loss aversion tends to lower the overall utility of the risky lottery.

2. Range normalization models and decision-by-sampling models

We now examine range normalization models and the decision-by-sampling (DbS) model, with an emphasis on how their predictions relate to our experiments. First, we analyze the range normalization model of Rustichini et al. (2017). This model gives rise to the following probability of risk taking

$$\mathbb{P}\text{rob}(\text{risk taking}|X, C) = \Phi \left(\frac{K_X t_X (X - X_l) - K_C t_C (C - C_l)}{\sqrt{\chi (K_X^2 t_X (X - X_l) + K_C^2 t_C (C - C_l))}} \right), \quad (\text{D.3})$$

where $t_X = \bar{v}/(X_u - X_l)$ and $t_C = \bar{v}/(C_u - C_l)$. Here \bar{v} and χ are parameters of coding capacity. Also, when $p(X_u - X_l) = C_u - C_l$, $K_X = K_C$. Equation (D.3) can be further simplified as

$$\mathbb{P}\text{rob}(\text{risk taking}|X, C) = \Phi \left(\sqrt{\frac{\bar{v}}{\chi}} \cdot \frac{(X - X_l)/(X_u - X_l) - (C - C_l)/(C_u - C_l)}{\sqrt{(X - X_l)/(X_u - X_l) + (C - C_l)/(C_u - C_l)}} \right). \quad (\text{D.4})$$

Panel A of Figure D1 plots, for the two volatility conditions specified in Section III.1.1, the probability of choosing the risky lottery from equation (D.4). This result shows that models of efficient coding and models of range normalization both give rise to the main prediction from Experiment 1—that sensitivity to payoff values increases when the dispersion of potential values decreases.

[Place Figure D1 about here]

At the same time, equation (D.4) shows that, holding the ranges— $X_u - X_l$ and $C_u - C_l$ —and X and C constant, the probability of risk taking in models of range normalization does not vary with the shape of the payoff distribution. Therefore, these models cannot explain the results from Experiment 2—that demand for the risky lottery is higher in the increasing condition compared to the decreasing condition.

Lastly, we turn to the decision-by-sampling model. As discussed in HWP, DbS implies the

³⁸The main result from Experiment 1 relates to stochastic choice, which is outside the deterministic model of KR.

following coding rules under the resource constraints in HWP:

$$\theta(X) = F(X), \quad \theta(C) = F(C). \quad (\text{D.5})$$

Panel B of Figure D1 shows that DbS also predicts that a given increase in X or C leads to a larger change in risk taking when payoffs are drawn from the low volatility distribution, compared to the high volatility distribution (as in our Experiment 1). Moreover, DbS implies higher demand for the risky lottery when risky payoffs are drawn from an increasing distribution, compared to a decreasing distribution (as in our Experiment 2). Thus, efficient coding and DbS generate qualitatively similar predictions. Indeed, [Bhui and Gershman \(2018\)](#) show that efficient coding can serve as a normative foundation for DbS.

Figure D1. Probability of choosing the risky lottery under alternative models

Panel A: the graph plots, for each of the two volatility levels (low volatility: $X_l = 16$, $X_u = 24$, $C_l = 8$, and $C_u = 12$; high volatility: $X_l = 8$, $X_u = 32$, $C_l = 4$, and $C_u = 16$), the probability of choosing the risky lottery implied by the range normalization model of Rustichini et al. (2017). The stimulus distributions for X and C are uniform. The probability p for the risky lottery to pay X dollars is set to 0.5. The ratio \bar{v}/χ is set to 15, as in Rustichini et al. (2017), which is a measure of coding capacity. Panel B: the graph plots, for each of the two volatility levels described above, the probability of choosing the risky lottery implied by decision-by-sampling (Stewart et al., 2006; Bhui and Gershman, 2018). The capacity constraint parameter n is set to 10. For each volatility condition, we draw X and C from their respective uniform distributions, and then compute the probability of risk taking for each level of $pX - C$.

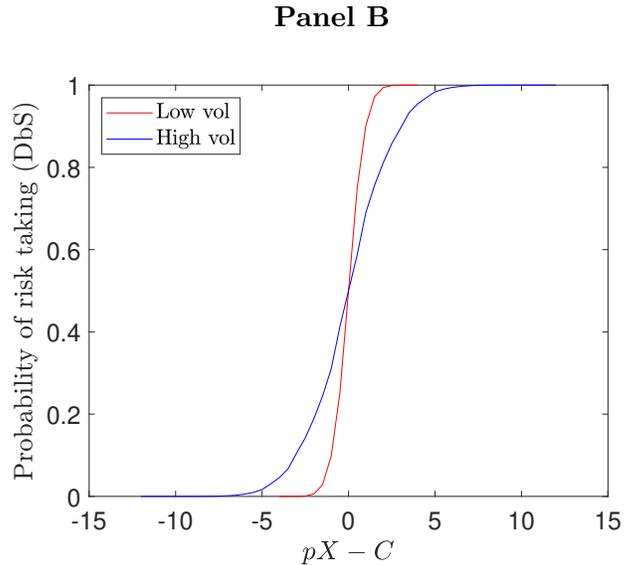
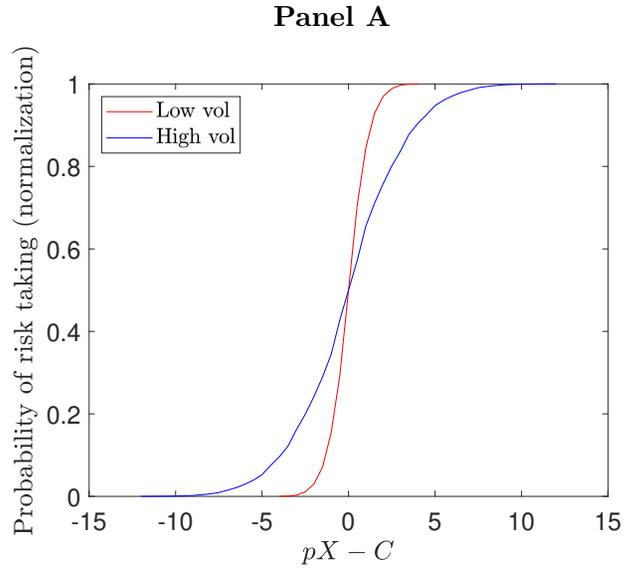


Table D1. Predictions of [Kőszegi and Rabin \(2006, 2007\)](#)

The table computes, for each of the 10 common trials from Experiment 2 and for each of the two experimental conditions, the difference between the KR value of the risky lottery and the KR value of the certain option. Version 1 and version 2 correspond to two different specifications of the reference point distribution, as discussed in Appendix D.1. Parameter values are: $\eta = 1$ and $\lambda = 2$.

Trial	X	C	Version 1		Version 2	
			Increasing	Decreasing	Increasing	Decreasing
1	9.04	6.09	-9.65	-3.26	-7.39	-1.01
2	10.09	6.09	-8.34	-1.97	-5.82	0.51
3	11.05	6.09	-7.15	-0.82	-4.40	1.86
4	12.02	6.09	-5.95	0.33	-2.97	3.19
5	13.08	6.09	-4.64	1.58	-1.41	4.62
6	14.01	6.09	-3.50	2.65	-0.06	5.84
7	15.04	6.09	-2.24	3.83	1.43	7.16
8	16.07	6.09	-0.99	4.99	2.90	8.46
9	17.04	6.09	0.18	6.07	4.27	9.65
10	18.06	6.09	1.40	7.20	5.70	10.88

Table E1. Probability of choosing the risky lottery in Experiment 1 using logistic regressions

The table reports results from logistic regressions where the dependent variable takes the value of one if the subject chooses the risky lottery, and zero otherwise. Data are pooled across all subjects and standard errors are clustered at the subject level. The dummy variable, *high*, takes the value of one if the trial belongs to the high volatility condition, and zero if it belongs to the low volatility condition. Only data from test trials are included. Standard errors are reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

	Between subjects tests		Within subjects tests			
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable: “Choose risky lottery”	All data	Common trials only	Common trials only	Common trials only (w/out trials 301-450)	Common trials only—low vol first (w/out trials 301-450)	Common trials only—high vol first (w/out trials 301-450)
<i>high</i>	-0.464 (0.694)	-0.051 (0.934)	-0.013 (0.440)	-0.088 (0.541)	0.507 (0.777)	-0.628 (0.660)
<i>X</i>	0.456*** (0.038)	0.407*** (0.041)	0.331*** (0.027)	0.353*** (0.030)	0.407*** (0.041)	0.270*** (0.040)
<i>C</i>	-1.003*** (0.082)	-0.988*** (0.094)	-0.830*** (0.061)	-0.894*** (0.067)	-0.988*** (0.094)	-0.749*** (0.089)
<i>X</i> × <i>high</i>	-0.255*** (0.042)	-0.171*** (0.051)	-0.058*** (0.020)	-0.111*** (0.027)	-0.150*** (0.033)	-0.034 (0.036)
<i>C</i> × <i>high</i>	0.568*** (0.090)	0.368*** (0.114)	0.117** (0.048)	0.239*** (0.057)	0.263*** (0.071)	0.129* (0.077)
Constant	0.487 (0.673)	1.283 (0.848)	1.365*** (0.486)	1.512*** (0.584)	1.283 (0.851)	1.860*** (0.639)
Observations	37,612	4,170	8,257	6,411	3,125	3,286