# Efficient Coding and Risky Choice*

Cary Frydman and Lawrence J. Jin

**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 modulates the strength of valuation biases. As we experimentally increase the frequency of large payoffs, we find that subjects perceive the upside of a risky lottery more accurately and take greater risk. 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, our findings highlight the importance of imprecise and efficient coding in economic decision-making.

*JEL classification:* G02, G41
*Keywords:* efficient coding, perception, risky choice, neuroeconomics

*We are grateful to Andrei Shleifer (the Editor), four anonymous referees, Nicholas Barberis, Nicola Gennaioli, Katrin Gödker, Alex Imas, Shimon Kogan, John O'Doherty, Stavros Panageas, Antonio Rangel, Elke Weber, Michael Woodford, and seminar participants at Caltech, the Chinese University of Hong Kong, Harvard University, the London School of Economics, the National University of Singapore, the Ohio State University, Tsinghua University, the University of Hong Kong, the University of Mannheim, 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 Pennsylvania, the University of Southern California, the University of Texas at Dallas, the University of Warwick, the University of Zurich, 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, the Society for Neuroeconomics Conference, and the Sloan-Nomis Workshop on the Cognitive Foundations of Economic Behavior for helpful comments. Frydman thanks the NSF for financial support; Jin thanks the Linde Institute at Caltech for financial support.

# 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 of

[1]For experimental evidence consistent with efficient coding in sensory perception, see Girshick, Landy, and Simoncelli (2011), Wei and Stocker (2015, 2017), Heng, Woodford, and Polanía (2020), and Payzan-LeNestour and Woodford (2021). See also the evidence from Polanía, Woodford, and Ruff (2019) on efficient coding in choice between food items.

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 KLW, our framework generates stochastic choice between a risky lottery and a certain option; but, crucially, by adding the efficient coding mechanism from HWP, we can assess how the environment modulates the probability of choosing the risky lottery.

To build intuition, consider a $DM$ who chooses 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 the perceptual system must now "cover more ground" with its limited resources, the $DM$ finds it more difficult to distinguish between payoffs in the range [\$15, \$25].

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 find it easy 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. As such, 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.[2] This prediction is inconsistent with most standard economic models of risky choice in which val-

[2]This prediction is derived under the assumption that coding resources do not change with the volatility of the environment. Later in the paper, we discuss implications for the case when coding resources can change with the volatility of the environment.

uation 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 science (Stewart, Chater, and Brown, 2006), theories of normalization from neuroscience (Rangel and Clithero, 2012; Carandini and Heeger, 2012; Louie, Glimcher, and Webb, 2015), and alternative specifications of efficient coding (Wei and Stocker, 2015; KLW).

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. We manipulate the range of payoffs across a high volatility condition and a low volatility condition, and crucially, we include a set of 30 "common trials" that are presented in both conditions. These common trials allow us to cleanly compare behavior across conditions and identify the effect of the prior distribution. Another important component of our 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, Conen, Cai, and Padoa-Schioppa, 2017; Ma and Woodford, 2020).[3] 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 payoff of the certain option is associated with an 18.6% increase in the frequency 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 from each condition, and the effect is significant both between and within subjects. We also find that subjects execute decisions significantly faster in the low volatility condition, and therefore our results cannot be driven by an alternative hypothesis where subjects in the low volatility condition choose to process information for a longer period of time.

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

[3]As 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 HWP, 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).

incentivize subjects to 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.

When viewed through the lens of efficient coding, the results from our first experiment indicate that noisier perception generates noisier choice. At the same time, valuation remains largely unbiased, in the sense that the average perception of a payoff is approximately the same as the true payoff. The "bias towards the prior" effect, which is commonly associated with Bayesian models, does not arise in our model when the prior is uniform—as it is in our first experiment. For other priors, however, efficient coding can give rise to 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 will fluctuate over time as an optimal response to changes in the environment. For example, an environment with a decreasing distribution of payoffs will induce the familiar value function with diminishing sensitivity as payoffs become larger; an environment with an increasing distribution of payoffs will generate diminishing sensitivity as payoffs become smaller. Both types of value functions are a consequence of the perceptual system's limited capacity to discriminate between payoffs that occur infrequently (Robson, 2001; Rayo and Becker, 2007; Netzer, 2009; Payzan-LeNestour and Woodford, 2021).

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, a subject exhibits greater risk taking when she is 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 large payoffs; this resource allocation enables the $DM$ to accurately perceive the large and attractive upside of the risky lottery, and thus she chooses to take risk. In contrast, when the subject is adapted to a decreasing distribution, she has difficulty recognizing the large upside of the risky lottery, and hence is unwilling to take risk. Our results are consistent with those from a recent perceptual choice experiment by Payzan-LeNestour and Woodford (2021), who find that "outlier" stimuli are perceived less accurately than frequently occurring stimuli. More generally, we provide novel evidence consistent with the 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. Indeed, it is plausible that sensory perception is governed by efficient coding while a different decision system is activated when a subject is presented with a decision concerning monetary risk. 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; Woodford, 2020; KLW; Enke and Graeber, 2021).

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 data provide evidence consistent with the proposition that efficient coding provides a normative foundation for the value function assumed in prospect theory (Woodford, 2012a,b).[4]

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

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

# 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 KLW 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 prior distribution, which they use to derive novel implications for risky choice. We build on KLW by integrating their model with the efficient coding model of HWP. By combining these two models, the theory is able to generate predictions about how noisy coding—and the probability of risky choice—systematically changes for any payoff distribution to which the $DM$ has adapted.[5]

## *II.A. 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 probability $p$ and zero dollars with probability $1-p$. The $DM$'s task is to choose between these two options.

Under Expected Utility Theory, a $DM$ with utility $U(\cdot)$ and no background wealth chooses the risky lottery over the certain option if and only if

$$p \cdot U(X) + (1-p) \cdot U(0) > U(C). \tag{1}$$

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

We now present the KLW framework of noisy coding, which departs from Expected Utility Theory 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. Before observing the choice set that contains $X$ and $C$, the $DM$ holds prior beliefs about $X$ and $C$, denoted by $f(X, C)$. We further assume the $DM$ believes that $X$ and $C$ are

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

drawn independently: $f(X, C) = f(X)f(C)$. 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 $C$.[6] Each of the two 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 Section II.B.

Given the prior beliefs and noisy signals, the $DM$ follows Bayes' rule to generate posterior distributions about $X$ and $C$, which we denote by $f(X|R_x)$ and $f(C|R_c)$. As in KLW, we assume that the $DM$ has linear utility.[7] Under this assumption, the optimal decision rule depends only on the conditional means of the posterior distributions, $\mathbb{E}[\tilde{X}|R_x]$ and $\mathbb{E}[\tilde{C}|R_c]$.[8] 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 occurs under the following condition: $p \cdot \mathbb{E}[\tilde{X}|R_x] > \mathbb{E}[\tilde{C}|R_c]$.

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 $DM$. In other words, even after the $DM$ is presented with a choice set, she still faces uncertainty about the payoff values of $X$ and $C$. As such, Bayesian inference takes place at the level of a single choice set, and it characterizes how the $DM$'s prior belief shifts after observing a noisy signal of the true payoff. As we show next, the distribution of the noisy signal drives the main predictions of our model.

### *II.B. Likelihood function*

*1. Deriving the optimal likelihood functions.* We depart from KLW by allowing the $DM$ to choose the optimal likelihood functions according to the 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 $f(R_c|C)$. We also retain the assumption from KLW that the $DM$ encodes

[6] In the choice environment we study here, our interpretation is that the noisy signals are generated unconsciously; they are not the outcome of deliberate and conscious information acquisition in the sense of Stigler (1961). For more discussion on this point, see Ma and Woodford (2020).

[7] Although we focus on how imperfect perception affects risky choice, our model does not preclude the more traditional source of risk aversion that operates through diminishing marginal utility of wealth. In Section III.E, we extend the model by allowing for both imperfect perception and intrinsic risk aversion.

[8] We use the notations $\tilde{X}$ and $\tilde{C}$ to emphasize that, when the $DM$ forms the posterior means $\mathbb{E}[\tilde{X}|R_x]$ and $\mathbb{E}[\tilde{C}|R_c]$, she does not directly observe $X$ and $C$; she treats these payoffs as random variables.

$X$ and $C$ independently; this is a natural assumption given the $DM$'s prior belief that $X$ and $C$ are drawn independently. In addition, we assume the $DM$'s performance objective is to maximize expected financial gain, defined as

$$\begin{aligned}\text{Expected financial gain} \equiv & \iint pX \cdot \mathbb{P}\text{rob(choose the risky lottery}|X, C) \cdot f(X)f(C) \cdot dXdC \\ & + \iint C \cdot \mathbb{P}\text{rob(choose the certain option}|X, C) \cdot f(X)f(C) \cdot dXdC.\end{aligned} \tag{2}$$

We focus on the performance objective in (2) because it is commonly used in economic settings—although we discuss alternative objectives below.

When maximizing expected financial gain, the $DM$ faces an information processing constraint. Specifically, HWP assume that the $DM$ encodes $X$ 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, and each neuron takes the value 1 with probability $\theta(X)$ and 0 with the remaining probability $1 - \theta(X)$. The encoded value of $X$ is therefore represented by an output vector of $0s$ and $1s$, with length $n$. Given that the neurons are mutually independent, a sufficient statistic for the output vector is the sum across the $n$ output values, which we denote by $R_x$. Thus, the noisy signal $R_x$ can take on integer values from 0 to $n$. The likelihood function of $X$ can then be defined by

$$f(R_x|X) = \binom{n}{R_x}(\theta(X))^{R_x}(1 - \theta(X))^{n-R_x}. \tag{3}$$

We do not interpret the encoding process described above as a literal description of brain function; rather, we share with HWP the view that equation (3) is useful in analyzing how a system with limited resources may transmit information through a set of $n$ binary signals. As such, $n$ can be thought of as representing an individual-specific 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 $DM$ encodes $C$ using a process that is identical to that for encoding $X$. Each neuron takes the value 1 with probability $\theta(C)$ and 0 with the remaining probability $1 - \theta(C)$. The sum across the $n$ output values is denoted as $R_c$. The likelihood function of $C$ is defined by

$$f(R_c|C) = \binom{n}{R_c}(\theta(C))^{R_c}(1-\theta(C))^{n-R_c}. \tag{4}$$

Equations (3) and (4) show that the likelihood functions are driven by the coding rules, $\theta(X)$ and $\theta(C)$, which map $X$ and $C$ into the probability that a 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.

When $n$ is large, and when $pX$ and $C$ are i.i.d, HWP show that the coding rules that maximize expected financial gain are given by[9]

$$\theta(X) = \left[\sin\left(\frac{\pi}{2}\frac{\int_{-\infty}^{X} f(x)^{2/3}dx}{\int_{-\infty}^{\infty} f(x)^{2/3}dx}\right)\right]^2 \tag{5}$$

and

$$\theta(C) = \left[\sin\left(\frac{\pi}{2}\frac{\int_{-\infty}^{C} f(c)^{2/3}dc}{\int_{-\infty}^{\infty} f(c)^{2/3}dc}\right)\right]^2. \tag{6}$$

Importantly, we see that the coding rules depend explicitly on the prior distributions $f(X)$ and $f(C)$, which gives the model its inherent context-dependence.

As in HWP, we also consider two alternative performance objectives besides maximizing expected financial gain, which we use to assess the robustness of the model's predictions. The first alternative objective is to maximize mutual information between $X$ and its noisy signal, $R_x$,

$$\max_{\theta(X)} I(X, R_x), \tag{7}$$

where the mutual information $I(X, R_x)$ is defined as the difference between the marginal entropy of $R_x$ and the entropy of $R_x$ conditional on $X$. Similarly, the $DM$ is assumed to maximize mutual information between $C$ and $R_c$. The second alternative objective is to maximize the probability of

[9]The coding rules described in this section are derived when $n$, which parameterizes the capacity constraint, is sufficiently large. Appendix 7 of HWP shows that, for any finite $n$ that is greater than or equal to 10, the coding rules remain approximately optimal. When illustrating the model's implications in Section II, we set $n$ to 10.

an accurate choice, given by

$$\begin{aligned}\mathbb{P}\text{rob}_{accurate} \equiv &\iint \left(\mathbb{P}\text{rob}(R_x > R_c | \theta(X) > \theta(C)) \cdot \mathbb{1}_{\theta(X)>\theta(C)}\right) \cdot f(X) f(C) \cdot dX dC \\ &+ \iint \left(\mathbb{P}\text{rob}(R_x < R_c | \theta(X) < \theta(C)) \cdot \mathbb{1}_{\theta(X)<\theta(C)}\right) \cdot f(X) f(C) \cdot dX dC.\end{aligned} \tag{8}$$

Interestingly, there exists a class of priors for which the three performance objectives in (2), (7), and (8) all lead to the same optimal coding rules (as $n$ goes to infinity). Specifically, in Online Appendix A, we show that, when the following two conditions

$$\begin{gathered}(i)\ pX \text{ and } C \text{ are i.i.d.} \\ \text{and } (ii)\ pX \text{ and } C \text{ are uniformly distributed}\end{gathered} \tag{9}$$

are satisfied, the coding rules under all three objectives reduce to

$$\theta(X) = \left[\sin\left(\frac{\pi}{2}\frac{X - X_l}{X_u - X_l}\right)\right]^2 \tag{10}$$

and

$$\theta(C) = \left[\sin\left(\frac{\pi}{2}\frac{C - C_l}{C_u - C_l}\right)\right]^2, \tag{11}$$

where $X_l$, $X_u$, $C_l$ and $C_u$ are the bounds of the uniform distributions, with $X_l < X_u$ and $C_l < C_u$. This result is useful because it enables tests of efficient coding which are robust to changing the assumption that the $DM$ maximizes expected financial gain. In our main risky choice experiment, we build a design that satisfies the conditions in (9).

*2. 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 $X_u$, and that her prior belief about $C$ is a uniform distribution between $C_l$ and $C_u$. In keeping with the conditions in (9), we further set $C_l = p \cdot X_l$ and $C_u = p \cdot X_u$, so that $pX$ and $C$ are identically distributed. Under these assumptions, the likelihood functions of

$X$ and $C$ are given by

$$
\begin{aligned}
f(R_x|X) =& \binom{n}{R_x}\left(\left[\sin\left(\frac{\pi}{2}\frac{X-X_l}{X_u-X_l}\right)\right]^2\right)^{R_x}\left(1-\left[\sin\left(\frac{\pi}{2}\frac{X-X_l}{X_u-X_l}\right)\right]^2\right)^{n-R_x}, \\
f(R_c|C) =& \binom{n}{R_c}\left(\left[\sin\left(\frac{\pi}{2}\frac{C-C_l}{C_u-C_l}\right)\right]^2\right)^{R_c}\left(1-\left[\sin\left(\frac{\pi}{2}\frac{C-C_l}{C_u-C_l}\right)\right]^2\right)^{n-R_c}.
\end{aligned}
\tag{12}
$$

The expressions in (12) show that the likelihood functions depend directly on the parameters of the prior distributions, $X_l$, $X_u$, $C_l$, and $C_u$. This dependence of the likelihood function on the prior is a signature characteristic of efficient coding.

**[Place Figure I about here]**

Figure I 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 $X_u = 32$). In the low volatility environment, $X$ is distributed uniformly over a narrow range ($X_l = 16$ and $X_u = 24$). Equation (10) implies that these two distributions induce different coding rules $\theta(X)$. Panel B 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 (3). 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 $X = 22$. Thus, as $X$ increases from 18 to 22, the $DM$'s perceptual system can easily detect this change. 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 section, this difference in discriminability has a direct impact on risky choice.

### *II.C. 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 each payoff in the choice set. Under the assumption of linear utility, the $DM$'s optimal decision rule depends only on the conditional means of the posterior distributions. Specifically, the posterior means of $X$ and $C$, conditional on $R_x$ and $R_c$, are given by

$$\mathbb{E}[\tilde{X}|R_x] \equiv \int_{X_l}^{X_u} f(X|R_x)XdX = \frac{\int_{X_l}^{X_u} f(R_x|X)f(X)XdX}{\int_{X_l}^{X_u} f(R_x|X)f(X)dX} \tag{13}$$

and

$$\mathbb{E}[\tilde{C}|R_c] \equiv \int_{C_l}^{C_u} f(C|R_c)CdC = \frac{\int_{C_l}^{C_u} f(R_c|C)f(C)CdC}{\int_{C_l}^{C_u} f(R_c|C)f(C)dC}, \tag{14}$$

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 (12).

Importantly, equation (13) shows that the $DM$'s estimate of $X$ is a random variable, and 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

$$v(X) = \sum\nolimits_{R_x=0}^{n} f(R_x|X) \cdot \mathbb{E}[\tilde{X}|R_x]. \tag{15}$$

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

$$\sigma(X) = \left[\sum\nolimits_{R_x=0}^{n} f(R_x|X)\left(\mathbb{E}[\tilde{X}|R_x]\right)^2 - v^2(X)\right]^{1/2}. \tag{16}$$

Equations (15) and (16) indicate that the curvature of the value function and the randomness in subjective valuation are jointly determined by the $DM$'s prior beliefs and the implied likelihood functions.

**[Place Figure II about here]**

In keeping with the running example from the previous section, Panel A of Figure II plots, for both the high and low volatility environments, 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 environment. This is driven by the greater overlap of likelihood functions in the high volatility environment, compared to the low volatility environment. Because subjective valuation is noisier in the high volatility environment, the model predicts that choices will also be noisier and hence less sensitive to a given change in payoff values.

To formalize this prediction, we compute the probability of choosing the risky lottery—which we refer to from now on as the "probability of risk taking." 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 by comparing the posterior means of $X$ and $C$ in equations (13) and (14). As a result, when fixing $X$, $C$, and the prior distributions, we compute the probability of risk taking as follows:

$$\begin{aligned}\mathbb{P}\text{rob}(\text{risk taking}|X, C) = &\sum_{R_x=0}^{n} \sum_{R_c=0}^{n} \left( \mathbb{1}_{p\cdot\mathbb{E}[\tilde{X}|R_x] > \mathbb{E}[\tilde{C}|R_c]} \cdot f(R_x|X) \cdot f(R_c|C) \right) \\ &+ \sum_{R_x=0}^{n} \sum_{R_c=0}^{n} \left( \mathbb{1}_{p\cdot\mathbb{E}[\tilde{X}|R_x] = \mathbb{E}[\tilde{C}|R_c]} \cdot \frac{1}{2} f(R_x|X) \cdot f(R_c|C) \right).\end{aligned} \tag{17}$$

Equation (17) says that the $DM$ chooses the risky lottery over the certain option when $p\cdot\mathbb{E}[\tilde{X}|R_x] > \mathbb{E}[\tilde{C}|R_c]$, and that the $DM$ randomly chooses between the two options when $p \cdot \mathbb{E}[\tilde{X}|R_x] = \mathbb{E}[\tilde{C}|R_c]$.

**[Place Figure III about here]**

Figure III plots, for both the high and low volatility environments, the probability of risk taking against the difference in expected values between the two options, namely $pX - C$. Naturally, a higher value of $pX - C$ increases the attractiveness of the risky lottery and hence increases the probability of risk taking. Note that, for an expected utility maximizer 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, Figure III shows that under noisy coding, the probability of risk taking has an $S$-shaped relationship with $pX - C$. More importantly, under *efficient* coding, the slope of this function is negatively related to the volatility of the stimulus distribution (for those values of $pX - C$ that do not deliver an extreme probability near 0 or 1). Thus, for a given increase in $X$, the probability of choosing the risky lottery increases more in the low volatility environment. This heightened sensitivity in the low volatility environment can be traced back to the property illustrated in Panel C of Figure I: a given increase in $X$ leads to a larger difference in the distribution of noisy signals in the low volatility environment, compared to the high volatility environment.

### *II.D. Increasing and decreasing payoff distributions*

The results we have presented so far show that payoff volatility affects the dispersion in perceived valuation, $\sigma(X)$. At the same time, the average subjective valuation $v(X)$ largely coincides with $X$ when the prior distribution is uniform. Under other priors, however, efficient coding can induce strong *biases* in valuation, in the sense that $v(X)$ can differ significantly from $X$ (Woodford, 2012a,b; KLW). To illustrate how these strong biases arise, we consider a different environment in which the payoffs of $X$ and $C$ are drawn either from an increasing distribution or from a decreasing distribution. We specify the increasing distribution by

$$f(X; X_l, X_u, X_m^i, h, l) = \begin{cases} l, & \text{if } X_l \leq X \leq X_m^i \\ h, & \text{if } X_m^i < X \leq X_u \end{cases}. \tag{18}$$

And we specify the decreasing distribution by

$$f(X; X_l, X_u, X_m^d, h, l) = \begin{cases} h, & \text{if } X_l \leq X \leq X_m^d \\ l, & \text{if } X_m^d < X \leq X_u \end{cases}. \tag{19}$$

Panel B of Figure II plots, for both the increasing and decreasing prior distributions, the average subjective valuation $v(X)$, as well as its one-standard-deviation bounds $v(X) \pm \sigma(X)$.[10]

[10] When the prior distribution is increasing or decreasing, the conditions in (9) for the equivalence of coding rules no longer hold. Panel B of Figure II presents the subjective valuations based on the coding rule that maximizes the $DM$'s expected financial gain. The results are quantitatively similar if the subjective valuations are instead based on the coding rule that maximizes mutual information.

For the increasing distribution, small values of $X$ occur with low frequency, and the $DM$ therefore allocates little coding resources towards these infrequent and small values. The lack of coding resources dedicated to small values of $X$ gives rise to a positive perceptual bias towards the mean of the distribution: $v(X) > X$ when $X$ is small. Conversely, for the decreasing distribution, the $DM$ allocates little coding resources towards large values of $X$. As a result, the $DM$ is insensitive to high values of $X$ and exhibits a negative perceptual bias towards the mean of the distribution: $v(X) < X$ when $X$ is large. We test implications of these value functions after presenting our main experiment. Specifically, for large values of $X$, we experimentally test whether demand for risk taking is higher in the increasing condition, compared to the decreasing condition.

## III. Experiment 1: Volatility manipulation

In this section, we provide our main experimental test of 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 Online Appendix B for the pre-registration document. Each subject first completes a risky choice task and then a perceptual choice task. This ordering is chosen to minimize any fatigue effects in the risky choice task, which is our main task of interest. Subjects were paid a \$7 participation fee, in addition to earnings from each task.

### *III.A. 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 choose the above design parameters for two reasons. First, because our goal is to isolate the effect of volatility, we keep the mean of each payoff distribution 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 (9): the distributions of $X$ and $C$ are independent, and $pX$ and $C$ are identically and uniformly distributed. These conditions imply that the efficient coding rule is robust to changing the performance objective from maximizing expected financial gain to maximizing mutual information or maximizing the probability of an accurate choice. Thus, our design is optimized to test generic predictions of efficient coding.

**[Place Figure IV about here]**

Figure IV 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. Each condition contains 300 trials, which are broken into two phases: an initial "adaptation" phase with 30 trials, and a subsequent "test" phase with 270 trials. The adaptation phase is intended to allow subjects to adapt to the condition-specific payoff distribution. The test phase contains the trials that we are interested in analyzing.

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; therefore, 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. As such, 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 D.1 in Online Appendix D for exact values. In each condition, the location of a common trial is randomized at the subject level across the 270 possible test trial locations.

We 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. Therefore, for each volatility condition, the distribution of payoffs across all trials accurately reflects the appropriate population distribution. In summary, the common trials simultaneously serve two purposes: they allow for a clean comparison of behavior across conditions, and they also reinforce 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 the distribution of payoffs, subjects may artificially direct their attention to this information, which, in turn, could generate an experimenter demand effect. Furthermore, our design enables us to test for learning effects, which are important when conducting our within subject analyses.

One of the 600 trials was randomly selected for payment and the subject was paid according to their choice on this randomly selected trial. The average earning for the risky choice task was \$10.14. Online Appendix B provides the exact instructions that were given to subjects before the experiment.

### *III.B. Results from the risky choice task*

We produce a large data set that contains 90,000 total observations across all subjects and conditions. As part of our pre-registered data exclusion rule, we drop one subject who chose the certain option on all trials in the first condition. We are then left with 89,400 trials across the first and second experimental conditions, of which we analyze only the 8,940 common trials.

*1. Frequency of risk taking.* We begin our analysis with between subjects tests, which we construct using only trials from the first condition. We find that, on average, subjects choose the risky option on 52.3% of trials (standard error: 2.3%) with an average response time of 2.0 seconds (standard error: 0.085 seconds). Table I provides results from a set of mixed effects linear regressions, which account for heterogeneity across subjects in average levels of risk taking and in sensitivity to $X$ and $C$.[11] Column (1) shows that risk taking increases significantly in $X$ and

[11] We have also estimated an analogous set of mixed effects logistic regressions; however, with random effects on $X$, $C$, and the intercept, we find that these mixed effects logistic regressions do not converge to numerically stable estimates. As an alternative, we estimate a set of logistic regressions without random effects, in which we pool all subjects. Table D.2 in Online Appendix D shows that the results from the logistic regressions are consistent with those presented in Table I. Our preferred specification is the mixed effects linear regression because it accounts for heterogeneity across subjects.

decreases significantly in $C$. The coefficients of interest are those on the interaction terms: the coefficient on $X \times high$ is negative ($p$-value = 0.038) and the coefficient on $C \times high$ is positive ($p$-value = 0.056). This is our first piece of evidence consistent with efficient coding, namely, that a \$1 increase in $X$ and a \$1 decrease in $C$ each leads to a greater increase in the frequency of choosing the risky lottery in the low volatility condition.

**[Place Table I about here]**

Before conducting further analyses, we take one additional step to clean the data. We exclude the trials on which a subject exhibits an excessively fast response time of less than 0.5 seconds, which constitutes 7.6% of the data; the remaining data are referred to as the "restricted sample." This exclusion was not pre-registered, but unsurprisingly these fast decisions are not responsive to the underlying payoff values, and thus we employ this exclusion in all subsequent analyses. Column (2) of Table 2 shows that after excluding fast decisions, the conclusion remains largely the same: the coefficient on $X \times high$ is negative ($p$-value = 0.003) and the coefficient on $C \times high$ is positive ($p$-value = 0.003).

To see a graphical representation of the difference in behavior across conditions, Panel A of Figure V uses the data from Column (2) of Table I to plot the empirical frequency of risk taking as a function of the difference in expected values between the two options.

**[Place Figure V about here]**

The figure shows a striking difference across conditions: a \$1 increase in $pX - C$ leads to a greater increase in the frequency of choosing the risky lottery in the low volatility condition, compared to the high volatility condition. We emphasize that this difference is based on the same exact set of 30 common trials in each condition. Panel B of Figure V presents the data in a different manner, without imposing the assumption that the frequency of risk taking is a function of $pX - C$. Each point represents one of the 30 common trials. The $x$-axis measures the frequency of risk taking in the high volatility condition, while the $y$-axis measures the frequency of risk taking in the low volatility condition. When the frequency of risk taking (in both conditions) is low, most of the data points fall below the 45-degree line. Conversely, when the frequency of risk taking is high, most of the data points are above the 45-degree line. This pattern is consistent with the model: the likelihood functions in the high volatility condition generally lead to less discriminability between payoffs, compared to the low volatility condition. As such, the model tends to predict that the probability of risk taking is less extreme—that is, closer to 50%—in the high volatility condition.

Our design also enables us to examine how coding varies *within* subjects over time when faced with a change in the environment. While our model does not make predictions about adaptation, we can test 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, on trial 301, the payoff distribution switches. Thus, by re-estimating the regression in Column (2) of Table I using data from both 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 those from the between subjects tests. 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 adaptation in the second condition, we exclude trials from the first half of the second condition, namely trials 301 through 450. Column (4) shows that the magnitudes of the interaction effects do indeed get larger.

Columns (5) and (6) further disaggregate the data based on whether the subject experiences the low or high volatility condition first. Column (5) shows that for those subjects who begin with the low volatility condition, the coefficients on $X \times high$ and $C \times high$ are both significant at the 1% level. This result is important because it rules out an alternative theory whereby subjects encode payoffs with noise, and through repeated experience with a payoff, the amount of noise decreases over time. Such a "learning from experience" theory can explain our between subjects results, because subjects in the low volatility condition experience low volatility payoffs more frequently than subjects in the high volatility condition. However, this alternative theory cannot explain the within subject result in Column (5): under this 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. However, we find that the effect has the opposite sign, and is therefore consistent with noisy and *efficient* coding.

Column (6) provides results for those subjects who experience the high volatility condition first. The coefficients on $X \times high$ and $C \times high$ have the predicted sign, but the effects become weaker and are no longer statistically significant. We conjecture that it is easier for subjects to detect a change in the environment when moving from a low volatility to high volatility condition, 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 condition, the information that signals a change in environment is less salient.

*2. Response times.* Not only are subjects more sensitive to payoffs in the low volatility condition, but they also implement decisions more quickly in this condition. Among common trials, subjects take an average of 2.02 seconds in the high volatility condition vs. 1.79 seconds in the low volatility condition ($p$-value = 0.001 for a within subject test). This difference in response times rules out the hypothesis that subjects perceive payoffs more precisely in the low volatility condition because they spend more time on each decision.

Response times can also 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 entire 600 trials during the risky choice task. Figure VI shows the response time data, disaggregated by which condition a subject experienced first; here the figure includes data from non-common trials, in addition to common trials. The upper panel plots, for each of the 600 trials, the response time averaged across subjects who first go 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. This 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 presumably restart the adaptation process.

**[Place Figure VI about here]**

In contrast, the lower panel plots the response time averaged across subjects who first go through the high volatility condition, followed by 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 second condition is in the support of the payoff distribution from the recently experienced first condition. We speculate that this extra difficulty in adapting to the low volatility distribution may be partly responsible for the stronger within subject results in Column (5), compared to those in Column (6) of Table I.

While our theory does not make predictions about response times, the data strongly suggest that response times are systematically related to the *DM*'s prior distribution. A natural way to incorporate response times into our framework would be to allow the *DM* to draw a *sequence* of noisy signals, $\{R_{x,i}\}_{i=1}^{S}$ for a given payoff $X$. The time that the *DM* 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).

### *III.C. 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 $p$ or the \$0 payoff, 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 experimental results discussed above.

To provide a more targeted test of the key efficient coding mechanism, each subject participates 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 expect to find evidence that the perception of $X$ still depends on the recent stimulus distribution.

Our perceptual choice task is informed by work from the literature on 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 experiment. The subjects' task is 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, accuracy decreases and response times increase. 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 $[31, 99]\backslash\{65\}$. In the low volatility condition, $X$ is drawn uniformly from integers in the set $[56, 74]\backslash\{65\}$. 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 VII gives a schematic of the design.

**[Place Figure VII 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, our main analysis focuses on those trials in the test phase for which the stimulus numbers fall in the range of common support across the two conditions, $[56, 74]$.[12] As in the risky choice task, the restriction to common trials is crucial because it allows us to cleanly identify the effect of the prior distribution on behavior.

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 in the perceptual choice task. We incentivize fast responses in this task (but not in the risky choice task) in order to avoid a "ceiling effect" in the choice data where subjects would approach 100% accuracy. While the 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.

[12] For the high volatility condition, we designate 90 out of the 340 test trials as common trials. These 90 common trials are created by sampling each element in the low volatility condition 5 times. 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 $[75, 99]$. For the low volatility condition, we designate the entire 340 test trials as common trials as they all fall in the range $[56, 74]$. This procedure ensures that $X$ is drawn according to its population distribution in both conditions.

### *III.D. Results from the perceptual choice task.*

We begin by reporting results for between subjects tests using all common trials from the first condition. Subjects correctly classify the number on 93.5% of trials (standard error: 0.7%) with an average response time of 0.573 seconds (standard error: 0.012 seconds). 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. Therefore, we exclude them from all subsequent analyses (their average accuracy rates were 51.8% and 55.0%, respectively).[13]

**[Place Figure VIII about here]**

Panel A of Figure VIII plots, for each value of $X$, the proportion of trials that subjects classified $X$ as greater than the reference level of 65. Consistent with previous research on numerical cognition, we see that subjects exhibit errors in classification; moreover, the errors increase as $X$ approaches 65. While it is unsurprising that subjects exhibit errors, the fact that the frequency of errors correlates with $|X - 65|$ provides evidence that coding is noisy and that number comparison is more difficult when the numbers are closer together.

The novel aspect of our design that enables us to test for efficient coding is the manipulation across the 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). This result is analogous to our finding from the risky choice task, in which subjects are more sensitive to payoffs and respond faster in the low volatility condition. Panel B of Figure VIII shows that the average response time for trials on which subjects responded correctly increases as $X$ approaches 65, and that, across the distribution of $X$, the average response time is shorter in the low volatility condition.

Next, we test for a difference in slope between the two choice curves presented in Panel A of Figure VIII. Efficient coding predicts a steeper slope for the choice curve from the low volatility condition. Before proceeding to the test, we note that in the perceptual choice task, the conditions

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

in (9) no longer hold because the $DM$ needs to only encode one payoff, and thus her prior is only over one dimension. We therefore base our analyses for the perceptual choice task on the assumption that the $DM$ maximizes mutual information.[14] To formally test for a difference in slope, Table II presents results from a set of mixed effects logistic regressions. The dependent variable takes the value of one if the subject classifies $X$ as above 65, and zero otherwise.

**[Place Table II about here]**

In Column (1), the coefficient on $X - 65$ is significantly positive, indicating that subjects' propensity to classify $X$ as greater than 65 is increasing in $X$. More importantly, the coefficient on the interaction term, $(X - 65) \times high$, is significantly negative, indicating that choices are noisier on trials in the high volatility condition. Column (2) and Column (3) examine only numbers inside the 60s decade and only numbers outside the 60s decade, respectively; in both cases, the coefficient on the interaction term remains significant. Lastly, Columns (4) to (6) provide within subject tests by pooling data across both conditions for each subject. For each of these three specifications, the coefficient on the interaction term remains significant at the 1% level.

### *III.E. Model estimation*

In this section, we structurally estimate the model to test for a connection between the two tasks. For each subject, we first estimate the model using data from the risky choice task, and then we estimate the model using data from the perceptual choice task.

*1. Estimation of the risky choice task.* Recall that the one free parameter of the model, $n$, 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 the subject's limited perceptual resources: the greater $n$ is, the more precise is the representation of each payoff. We note that our model only accounts for noise in the perception of $X$ and $C$. Additional sources of noise in the decision process that are outside the model—for example, noise in perceiving $p$—may also be captured by the structural estimate of $n$, because it is the only free parameter of the model.

[14] In the perceptual choice task, we incentivize fast responses; therefore, the coding rule from (5) does not necessarily maximize expected financial gain. At the same time, the coding rule from (10) continues to maximize mutual information. For this reason, we opt for the assumption that the $DM$ maximizes mutual information in this task.

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:

$$LL(n|\mathbf{y}) = \sum\nolimits_{t=31}^{300} y_t \cdot \log(\mathbb{P}\mathrm{rob}(y_t|n)) + (1-y_t) \cdot \log(1 - \mathbb{P}\mathrm{rob}(y_t|n)), \tag{20}$$

where $\mathbf{y} \equiv \{y_t\}_{t=31}^{300}$ 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, $\mathbb{P}\mathrm{rob}(y_t|n)$ denotes the model predicted probability of choosing the risky lottery given $n$, $X_t$, and $C_t$; it is computed using (17) from Section II.C. We maximize the log likelihood function in (20) by searching over integer values of $n$ in $[5, 40]$. We find that the average estimate of $n$, across subjects, is 8.9 with a standard deviation of 9.7, indicating substantial heterogeneity.

Our baseline model of efficient coding assumes linear utility; as such, the $DM$'s optimal decision rule depends on $\mathbb{E}[\tilde{X}|R_x]$ and $\mathbb{E}[\tilde{C}|R_c]$. However, the model can easily be integrated with standard nonlinear utility functions. To generalize the baseline model, we assume that the $DM$ maximizes expected utility, with a utility function $U(\cdot) = (\cdot)^{\alpha}$. Under this assumption, the $DM$ chooses the risky lottery if and only if $p \cdot \mathbb{E}[(\tilde{X})^{\alpha}|R_x] > \mathbb{E}[(\tilde{C})^{\alpha}|R_c]$.[15] The optimal coding rules presented in equations (5) and (6) are then replaced by[16]

$$\theta(X) = \left[\sin\left(\frac{\pi}{2}\frac{\displaystyle\int_{X_l}^{X} f(x)^{2/3}(x)^{(\alpha-1)/3}dx}{\displaystyle\int_{X_l}^{X_u} f(x)^{2/3}(x)^{(\alpha-1)/3}dx}\right)\right]^2 \tag{21}$$

and

$$\theta(C) = \left[\sin\left(\frac{\pi}{2}\frac{\displaystyle\int_{C_l}^{C} f(c)^{2/3}(c)^{(\alpha-1)/3}dc}{\displaystyle\int_{C_l}^{C_u} f(c)^{2/3}(c)^{(\alpha-1)/3}dc}\right)\right]^2. \tag{22}$$

[15] As in equation (17), we also assume that the $DM$ randomly chooses between the risky lottery and the certain option when $p \cdot \mathbb{E}[(\tilde{X})^{\alpha}|R_x] = \mathbb{E}[(\tilde{C})^{\alpha}|R_c]$.

[16] The generalization of the coding rules to allow for nonlinear utility follows the analysis of this issue in Payzan-LeNestour and Woodford (2021).

Finally, the probability of risk taking from equation (17) is replaced by

$$\begin{aligned}\mathbb{P}\text{rob}(\text{risk taking}|X,C) = &\sum_{R_x=0}^{n}\sum_{R_c=0}^{n}\left(\mathbb{1}_{p\cdot\mathbb{E}[(\tilde{X})^\alpha|R_x]>\mathbb{E}[(\tilde{C})^\alpha|R_c]}\cdot f(R_x|X)\cdot f(R_c|C)\right)\\ &+\sum_{R_x=0}^{n}\sum_{R_c=0}^{n}\left(\mathbb{1}_{p\cdot\mathbb{E}[(\tilde{X})^\alpha|R_x]=\mathbb{E}[(\tilde{C})^\alpha|R_c]}\cdot\frac{1}{2}f(R_x|X)\cdot f(R_c|C)\right).\end{aligned} \tag{23}$$

We estimate this generalized model using the same maximum likelihood procedure as above, but for each subject we now estimate two parameters, $n$ and $\alpha$. We find that the best fitting parameter pair, averaged across subjects, is $(n, \alpha) = (10.1, 0.93)$. This result indicates that the average subject exhibits a modest degree of intrinsic risk aversion with the small stakes in our experiment. We 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; therefore, our baseline assumption of linear utility is not overwhelmingly restrictive. At the same time, the generalized model with intrinsic risk aversion provides a better fit to a substantial number of subjects.

*2. Estimation of the perceptual choice task.* We now estimate the model for each subject using data from the perceptual choice task. Following the procedure from the previous section, for each subject we maximize the following log likelihood function over $n$, using choice data from the test phase of the first condition:

$$LL(n|\mathbf{z}) = \sum\nolimits_{t=61}^{400} z_t\cdot\log(\mathbb{P}\text{rob}(z_t|n)) + (1-z_t)\cdot\log(1-\mathbb{P}\text{rob}(z_t|n)), \tag{24}$$

where $\mathbf{z}\equiv\{z_t\}_{t=61}^{400}$ 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 $\mathbb{P}\text{rob}(z_t|n)$ denotes the model predicted probability that the subject classifies $X_t$ as greater than 65. We maximize the log likelihood function in (24) by searching over integer values of $n$ in $[5, 40]$. 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 the latent structural parameter $n$ across the 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 $\mathbb{E}[\tilde{X}|R_x]$ or $\mathbb{E}[(\tilde{X})^{\alpha}|R_x]$); as noted above, the structural estimation may then account for these outside sources of noise 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 the estimated $n$ 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 $\alpha = 1$ ($p$-value $= 0.001$). The results are also robust to using Pearson correlations for both models ($p$-value $= 0.002$ for the unrestricted model, and $p$-value $= 0.001$ for the restricted model). These positive correlations demonstrate that, across subjects, variation in perception partly explains variation in risk taking behavior.

## IV. 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 the 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 allows us to test for a bias in perception, which can systematically change the average *level* of risk taking.

### *IV.A. Design of Experiment 2*

As in the risky choice task from Experiment 1, here we design an experiment in which subjects are presented with choice sets of the form $\{(X, 0.5; 0, 0.5); (C, 1)\}$. There are two experimental conditions: in one condition, $X$ is drawn from a weakly decreasing distribution over the range $[2, 8]$; in the other condition, $X$ is drawn from a weakly increasing distribution over the same range.

The specific distributions that we use are shown in Panel B of Figure II.[17]

Our design ensures that large values of $X$—for example, those between \$7 and \$8—are frequent in the increasing condition, but are rare in the decreasing condition. For both experimental conditions, we set the distribution of $C$ to be that of $pX$. This is an important feature of the design because it implies that, under both conditions, the expected value of the risky lottery and the certain option are the same. Thus, while the expected value of the risky lottery is higher in the increasing condition compared to the decreasing condition, the same is true for the certain option. Therefore, any observed difference in risk taking across conditions provides evidence for an endogenous shift in the likelihood function—rather than just a shift in the prior.[18] In other words, our design helps us target a test of *efficient* coding, rather than just noisy coding.

Panel B of Figure II shows that, under efficient coding, the decreasing distribution of $X$ generates perception that is insensitive and biased downward for large values of $X$. On the other hand, the increasing distribution generates perception that is approximately unbiased for large values of $X$. This systematic difference in perception of large values of $X$ across conditions will affect the $DM$'s appetite for risk.

To test the predicted difference in risk taking, we again create a set of common trials that subjects face in both conditions. Specifically, we create 8 common trials, where we fix $C$ at \$2.70 and vary $X$ from \$7.13 to \$7.99 in approximately \$0.12 increments; see Table D.3 in Online Appendix D for exact values.[19]

There are 300 trials per experimental condition, and subjects are randomized into whether they first experience the increasing condition or the decreasing condition. In each condition, subjects first make choices on 60 "adaptation trials." Subjects then go through 8 consecutive test blocks, where each test block contains 30 trials. At the end of each test block, we insert a common trial. For example, the first common trial from the first condition occurs on trial 90, and the last common trial occurs on trial 300. The order of the 8 common trials is randomized at the subject-condition

[17]The increasing and the decreasing distributions take the form of (18) and (19) from Section II.D. The parameter values are: $X_l = 2$, $X_u = 8$, $h = \frac{7}{25}$, $l = \frac{1}{125}$, $X_m^i = 4.5$, and $X_m^d = 5.5$.

[18]Online Appendix A provides a brief proof of this statement.

[19]Our common trials focus on large values of $X$ because these values lead to a difference in risk taking across the two experimental conditions that remains substantial even when subjects exhibit a strong degree of intrinsic risk aversion. By contrast, small values of $X$ lead to a difference in risk taking across the two conditions that diminishes when subjects' degree of risk aversion is sufficiently high. Moreover, we fix the value of $C$ at \$2.70 so that it has a high density in both conditions. As a result, the difference in the perception of $C$ across the two conditions is minimal and only has a small impact on the difference in risk taking across conditions.

level. On non-common trials, we draw payoffs according to the increasing or decreasing payoff distribution, depending on the condition to which the trial belongs.[20]

Our main testable prediction is that a subject's average demand for the risky lottery on common trials is greater in the increasing condition, compared to the decreasing condition. Recall that on common trials, $X$ takes values between \$7 and \$8. For the increasing condition, these values of $X$ are frequent outcomes because they come from the high density part of the increasing prior distribution. As such, the perceptual bias in $X$ is minimal: $v(X) \approx X$. However, in the decreasing condition, these large values of $X$ are rare outcomes. They come from the long right tail of the decreasing prior distribution and hence lead to a perceptual bias that is substantially negative: $v(X) < X$. Efficient coding therefore predicts that a subject will perceive the risky lottery on each common trial to be less attractive in the decreasing condition, compared to when it is presented in the increasing condition.

We pre-register the experiment and recruit 200 subjects from Prolific, an online data collection platform; see Online Appendix B for the pre-registration document. As outlined in our pre-registration document, all of our statistical tests in Experiment 2 are conducted within subjects.[21] Given that Experiment 2 is conducted online, we impose a 10-second time limit on each trial in order to promote engagement. If a subject does not respond within the 10-second time limit, the computer randomly chooses one of the two options. Each subject completes the task and is paid according to one randomly selected trial, in addition to a \$6.50 participation fee. The average earning for this task, including the participation fee, was \$9.27. The experimental instructions for Experiment 2 are provided in Online Appendix B.

### *IV.B. Results from Experiment 2*

We begin our analysis by applying the following four exclusion criteria outlined in the pre-registration document. First, we exclude those subjects who failed to correctly answer at least one of the two comprehension quiz questions. Second, we exclude those trials for which subjects failed

[20] The design in this task is similar to that of Payzan-LeNestour and Woodford (2021) who insert a "test trial" every 40 trials, although their design is implemented in the perceptual choice where a subject is incentivized to discriminate between shades of grey.

[21] We choose a within subject design because of the added statistical power that it provides, which is useful for two reasons. First, in pilot testing on Prolific, we observed substantial variation across subjects in the frequency of choosing the risky lottery. Second, in this experiment we restrict our analysis to only 8 common trials per condition; in comparison, the risky choice task from Experiment 1 contains 30 common trials per condition.

to respond within the 10-second time limit. Third, we exclude those subjects who chose the risky option on less than 2.5% or more than 97.5% of non-common trials.[22] Finally, we exclude those trials for which subjects exhibited excessively fast response times, defined as less than 0.5 seconds. After applying these four exclusion criteria, 151 subjects remain with a total 85,703 trials, of which 2,278 are common trials. All regression results presented below are robust to using the full sample without applying the above exclusion criteria.

Our main hypothesis involves testing whether, for a fixed value of $C$, a subject's appetite for risk is higher when large values of $X$ are more frequent. Table III presents results from mixed effects linear regressions in which the dependent variable takes the value of one if the subject chooses the risky lottery, and zero otherwise. All regressions in Table III include only common trials. Column (1) shows that the frequency of choosing the risky lottery is 7.5% higher in the context of the increasing distribution, compared to the decreasing distribution ($p$-value = 0.001). As in Experiment 1, this result is based on a comparison where the choice sets are fixed and only the context varies across conditions. We emphasize that the significant difference in risk taking *cannot* be driven by the mere fact that $X$ has a higher mean in the increasing condition, because this effect is exactly offset by a higher mean of $C$ in the increasing condition. Thus, we interpret the empirical finding in Column (1) as a consequence of an endogenous shift in the likelihood function, rather than an exogenous shift in the prior mean of $X$.

**[Place Table III about here]**

The significant difference in risk taking also holds when we add a linear control for $X$ across the 8 common trials, as shown in Column (2). Note that we do not include any controls for $C$ because it is fixed at \$2.70 across all common trials. Columns (3) and (4) demonstrate that the main result in Column (1) continues to hold when we condition only on trials in the first half or the second half of each condition. Overall, the systematic increase in risk taking as we shift from the decreasing to increasing distribution demonstrates that diminishing sensitivity can arise, in part, as a consequence of an optimal allocation of perceptual resources.

[22]This exclusion criterion is motivated by the observation that, in pilot testing on Prolific, a small fraction of subjects exhibited almost zero variation in behavior across the entire 600 trials, indicating a strategy that does not depend on the values of $X$ and $C$. So as not to select on the dependent variable, this criterion is based on choices from non-common trials, which never enter our main analyses.

Our results from this experiment complement recent empirical evidence from a perceptual decision-making task on "outlier blindness" (Payzan-LeNestour and Woodford, 2021). Those authors provide data that subjects are less accurate in classifying shades of grey when the stimuli under consideration are outliers according to the subjects' prior. Analogously, we show that when a subject faces an outlier payoff—a large payoff from the decreasing distribution—the perceptual system is "blinded" and cannot accurately perceive such a high value. Crucially, efficient coding provides a directional prediction about the misperception of outliers; our data are consistent with this prediction of a negative perceptual bias for large payoffs in the decreasing condition.

# V. Discussion

### *V.A. Adaptation dynamics*

The theoretical framework 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. Most empirical tests of efficient coding in sensory perception assume full adaptation to the prior distribution (Laughlin, 1981; Wei and Stocker, 2015), and this assumption has also been recently invoked in papers on efficient coding in value-based decisions (Polanía et al., 2019; Rustichini et al., 2017). Following this literature, we have assumed that subjects in our experiments are fully adapted to the population distribution after completing an initial set of pre-registered "adaptation trials." Yet we emphasize 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 as a result, the $DM$ must learn from the history of *perceived* payoffs (Robson and Whitehead, 2018; Młynarski and Hermundstad, 2019; Aridor, Grechi, and Woodford, 2020).

Moreover, even if we assume that subjects can fully adapt to the environment through experience, this leaves open an important question regarding how much experience is needed to learn the prior distribution. For example, previous experimental work from HWP shows that convergence takes place after about 200 trials. To explore these dynamics in our data, we analyze the patterns of risky choice over the course of Experiment 1. Specifically, we compare behavior across the first and second half of the first experimental condition; that is, we compare behavior on trials 31 to

165 with behavior on trials 166 to 300. For each of these two subsamples, we estimate the mixed effects linear regression specified in Table I. If the treatment effect becomes stronger in the second half of the first condition, this would provide evidence that adaptation is not complete by trial 165.

Table D.4 in Online Appendix D shows that the estimated coefficients on $X \times high$ and $C \times high$ are significant at the 5% level in both subsamples, indicating that the treatment effect is present in both halves of the first condition. Nonetheless, we do not detect any significant difference in the strength of the treatment effect between the two subsamples. Thus, we cannot rule out the possibility that full adaptation has taken place by trial 165.

An alternative way to assess the speed of adaptation is by looking at the time series of response times from Figure VI. One can see signatures of learning dynamics as response times fall sharply at the beginning of the experiment; importantly, we find that the decrease in response times continues after the first 30 adaptation trials. In fact, response times asymptote around trial 200, consistent with the experimental evidence from HWP. While the decrease in response times may reflect learning about the task in general—rather than about the prior exclusively—the spike in average response time in the upper panel of Figure VI (at trial 301) strongly suggests that response times capture information about the adaptation process. We also emphasize that not all past observations are likely to receive the same weight when subjects form a prior. For example, if subjects are more likely to recall past payoffs that were experienced more recently or payoffs that are more similar to those in the current choice set, then they may overweight these recent or similar stimuli when forming a prior. As such, research from the memory literature is likely to have important implications for efficient coding of economic stimuli (Kahana, 2012; Bordalo, Gennaioli, and Shleifer, 2020; Wachter and Kahana, 2021).

### *V.B. Multi-dimensional efficient coding*

The efficient coding model we present in Section II assumes that the capacity constraint parameter $n$ is fixed and does not vary with the prior distribution. It is conceivable, however, that the $DM$ optimally chooses to allocate a larger capacity $n$ towards the dimensions of $X$ and $C$ when the payoff volatility along these dimensions is higher. Intuitively, when the payoff volatility along the two dimensions of $X$ and $C$ becomes higher, these dimensions could receive more resources from a "third" dimension, which we can think of as representing other task demands besides perceiving $X$

and $C$. Indeed, this is an implication from the multi-dimensional efficient coding models of Woodford (2012a,b) and Dewan (2020)—although the performance objectives in those models slightly differ from what we assume in Section II.

Interestingly, our experimental data are consistent with this implication: when separately estimating the capacity constraint parameter $n$ across the low and high volatility conditions using the risky choice data from Experiment 1, we find $n = 7.05$ for the low volatility condition and $n = 11.66$ for the high volatility condition ($p$-value $< 0.001$ for a within subject test). Moreover, the rank correlation between the values of $n$ estimated from the two volatility conditions is 0.53 ($p$-value $< 0.001$), indicating that $n$ is a persistent trait at the subject level.

**[Place Figure IX about here]**

While we do find that subjects allocate a larger capacity towards perceiving $X$ and $C$ in the high volatility condition, this does not imply that choice sensitivity is higher in the high volatility condition, compared to the low volatility condition. Indeed, Figure IX shows that our main theoretical result, as presented earlier in Figure III, continues to hold when we set $n = 7$ for the low volatility condition and $n = 12$ for the high volatility condition (these values of $n$ are chosen to match the results from the structural estimation).

How can the larger capacity parameter $n$ that we estimate in the high volatility condition generate less choice sensitivity displayed in Figure IX? As we have done throughout the paper, here we analyze only those payoffs that are drawn from the common support of the high and low volatility distributions. This restriction is crucial for identifying the effect of context, which we implement in our experiment through the use of common trials. But importantly, the optimal coding capacity allocated to the common trials in the high volatility condition is only a small fraction of the overall capacity; a large remaining fraction of the coding capacity is consumed by perceiving payoffs in the tails of the high volatility distribution. Therefore, while the $DM$ chooses a larger overall capacity in the high volatility condition, this effect is more than offset by the fact that we restrict our analysis to only those payoffs in the common support of the two volatility distributions.

### *V.C. 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 choice 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).[23] A particularly relevant class 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 Experiment 1 strongly support the predictions of range normalization models. In particular, our results highlight the interpretation that normalization can implement normative principles of efficient coding.[24] It is worth noting that Experiment 2 offers a test between our efficient coding model and range normalization models. In Experiment 2, we hold constant the range of payoffs across experimental conditions. Therefore, our finding of greater risk taking in the increasing condition cannot be explained by range normalization models, suggesting that other statistics of the distribution besides the range do affect coding. See Online Appendix C for more discussion.

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

[23] Several 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, 2021).

[24] Not 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 idea that the $DM$ puts less weight on attributes that exhibit greater variability is a key *assumption* in the relative thinking model by Bushong, Rabin, and Schwartzstein (2020).

core mechanism in DbS.[25]

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 coding share the feature that expectations shape the $DM$'s perception of a lottery payoff. Yet an important distinction exists between the two models. In efficient coding, the driving force of the model is the $DM$'s expectation of a payoff value, after it has been presented in the choice set. In KR, however, the driving force is the $DM$'s expectation over which payoff value she will receive as a future, unrealized outcome from the lottery. Therefore, conditional on a choice set, KR does not predict a change in behavior as the prior distribution varies. We further elaborate this point in Online Appendix C.

Salience theory is an alternative model of risky choice that is also grounded in principles of perception and delivers context-dependent behavior (Bordalo, Gennaioli, and Shleifer, 2012). Salience theory and models of efficient coding are fundamentally linked. Under salience theory, attention is drawn to payoffs that are very different from a reference payoff. Consistent with this assumption, the multi-dimensional efficient coding models of Woodford (2012a,b) and Dewan (2020) imply that, when the payoff volatility along a dimension increases, more coding resources will flow towards that dimension, and therefore extreme payoffs will receive more weight in the decision process.

The two models, however, differ with respect to their primitive assumptions. Bordalo et al. (2012) appeal to Weber's law of diminishing sensitivity, in part, as a justification for their definition of salience; 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. It is also notable that for prior distributions that are increasing in payoff values, efficient coding predicts an "anti-Weber's" law. This difference leads the two models to generate different predictions in many environments. For example, salience theory does not predict that risk taking will

[25] Our 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 their 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 payoff distribution 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.

increase when the $DM$ is adapted to an increasing payoff distribution, as we find in Experiment 2. Nor does it deliver stochastic choice, where the degree of stochasticity changes systematically with the prior—as we observe in Experiment 1. At the same time, there are extant empirical patterns in the literature that salience theory can explain, which the baseline efficient coding model from Section II cannot, such as the dependence of risk taking on the correlation between mutually exclusive lotteries.

### *V.D. 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 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 preference elicitation 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 subjects with choice sets that maximize the information gain for estimating preference parameters from prospect theory. Efficient coding suggests 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.

## VI. Conclusion

We have experimentally tested the hypothesis that efficient coding, a core principle from neuroscience, 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. Our experimental data are consistent with the noisy perception of lottery payoffs, and moreover, we find that the noise distribution 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 find that the level of risk taking changes with the shape of the payoff distribution, which highlights the role that efficient coding plays in generating perceptual biases. Together, our 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 of course depends on the $DM$'s prior beliefs about payoffs—which we manipulate in our experiments—but it also depends on higher order priors about the rate at which the environment changes. For example, 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, Heasly, and Gold, 2012). Theory is already being developed along this direction, but future experimental evidence of the adaptation process will be critical in guiding further 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 shed light on the relevant timescale for forming prior beliefs, and hence perceptions. Institutional factors will also likely shape the relevant timescale for adaptation. For example, in financial markets, a stock's price distribution over the past 52 weeks is typically salient to investors, and therefore may be a good candidate for investors' 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.

University of Southern California

California Institute of Technology

# 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.

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, Salience 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.

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.

Denrell, Jerker C., 2015, Reference-dependent risk sensitivity as rational inference, *Psychological Review* 122, 461–484.

Dewan, Ambuj, 2020, Costly multidimensional information, Working paper.

Enke, Benjamin, and Thomas Graeber, 2021, 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.

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.

Hébert, Benjamin, and Michael Woodford, 2019, Rational inattention when decisions take time, Working paper.

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

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.

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, 2021, Pairwise normalization: A neuroeconomic theory of multi-attribute choice, *Journal of Economic Theory* 193, 105221.

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

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.

Młynarski, 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 Heasly, 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, 2021, Outlier blindness: A neurobiological foundation for neglect of financial risk, *Journal of Financial Economics* forthcoming.

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.

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.

Wachter, Jessica A., and Michael J. Kahana, 2021, 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.

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.

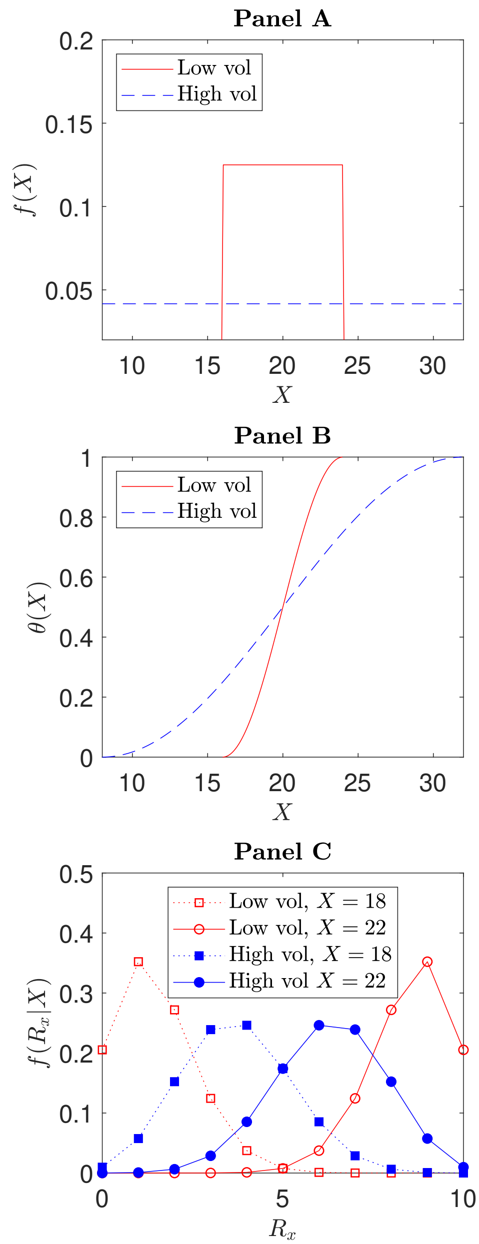


Figure I

Prior distributions, coding rules, and the optimal likelihood functions

Panel A plots two uniform prior distributions for $X$, one with low volatility ($X_l = 16$ and $X_u = 24$) and the other with high volatility ($X_l = 8$ and $X_u = 32$). Panel B plots the coding rule $\theta(X)$, defined in equation (10), for both volatility environments. Panel C plots the implied likelihood function $f(R_x|X)$, defined in equation (12), for two values, $X = 18$ and $X = 22$, and for each of the two prior distributions. The capacity constraint parameter $n$ is set to 10.

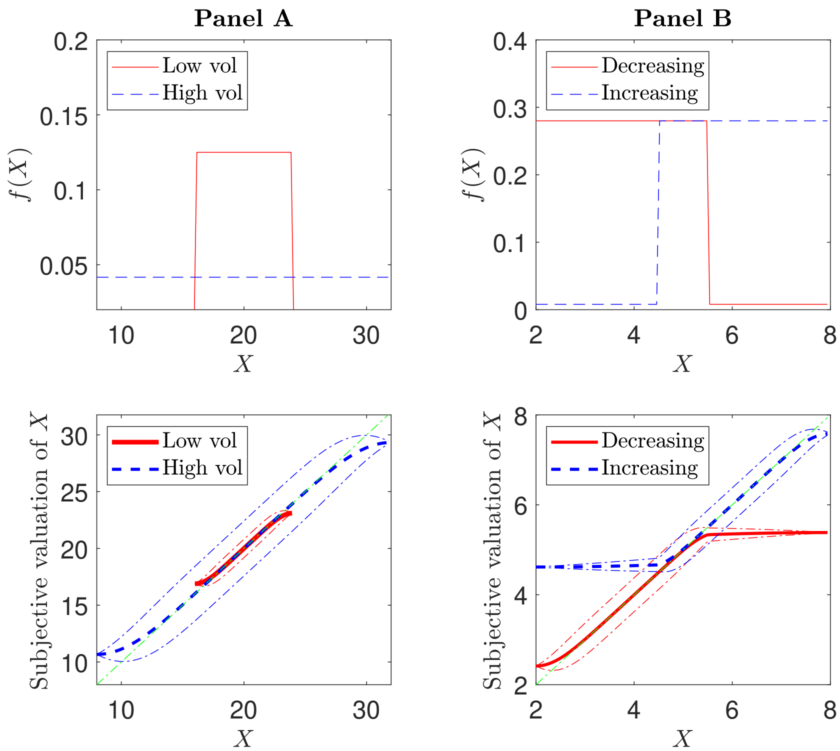


Figure II
Prior distributions and value functions

Panel A: the upper graph plots two uniform prior distributions for $X$, one with low volatility ($X_l = 16$ and $X_u = 24$) and the other with high volatility ($X_l = 8$ and $X_u = 32$). 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 prior distributions, one increasing and one decreasing. The increasing distribution is characterized by

$$f(X; X_l, X_u, X_m^i, h, l) = \begin{cases} l, & \text{if } X_l \leq X \leq X_m^i \\ h, & \text{if } X_m^i < X \leq X_u \end{cases},$$

where $X_l = 2$, $X_u = 8$, $X_m^i = 4.5$, $h = \frac{7}{25}$, and $l = \frac{1}{125}$. The decreasing distribution is characterized by

$$f(X; X_l, X_u, X_m^d, h, l) = \begin{cases} h, & \text{if } X_l \leq X \leq X_m^d \\ l, & \text{if } X_m^d < X \leq X_u \end{cases},$$

where $X_l = 2$, $X_u = 8$, $X_m^d = 5.5$, $h = \frac{7}{25}$, and $l = \frac{1}{125}$. The lower graph plots the subjective valuations implied by efficient coding, $v(X)$, and their one-standard-deviation bounds $v(X) \pm \sigma(X)$. For both panels, the capacity constraint parameter $n$ is set to 10. In the lower graph of each panel, the green dash-dot line is the forty-five degree line.

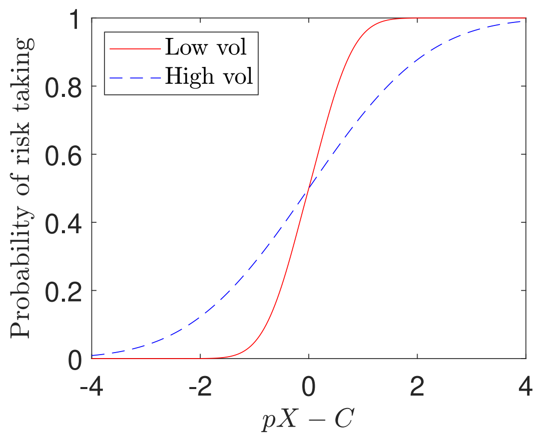


Figure III
Model predicted probability of choosing the risky lottery

The graph plots the probability of risk taking, defined in equation (17), for each of the two volatility environments (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 prior distributions for $X$ and $C$ are uniform. The probability $p$ that the risky lottery pays $X$ dollars is set to 0.5. The capacity constraint parameter $n$ is set to 10. For each volatility environment, we draw $X$ uniformly from $[16, 24]$ and $C$ uniformly from $[8, 12]$; that is, we draw the payoffs from the common support of the low volatility distribution and the high volatility distribution. We then compute, for a given $X$ and $C$, the probability of risk taking. Finally, we aggregate these probabilities for each level of $pX - C$.

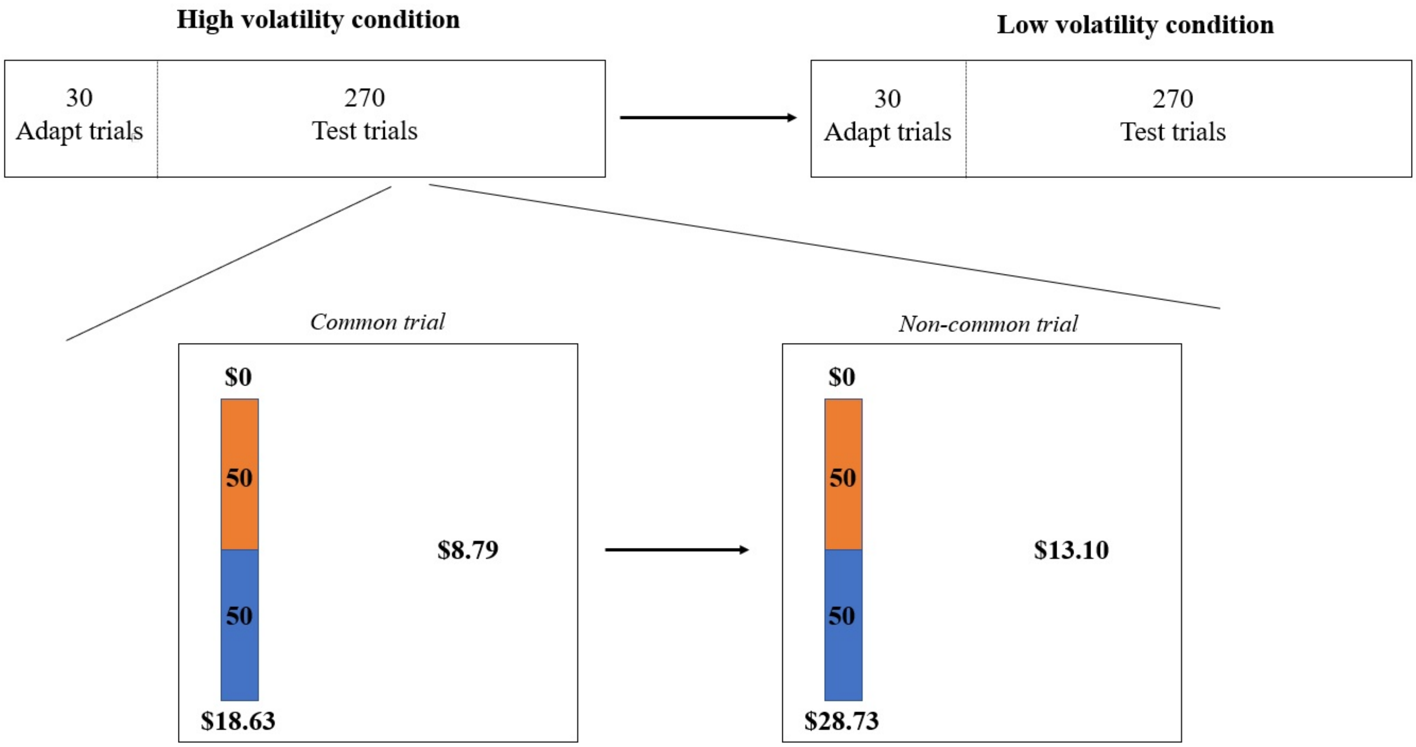


Figure IV
Design of the risky choice task in Experiment 1

The task consists of two blocks of trials: one block contains trials from the high volatility condition, and the other block contains trials from the low volatility condition. The order of the two blocks is randomized across subjects. Each block begins with 30 "adaptation 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, subjects have unlimited time to decide which of the two options they prefer.

**Panel A**

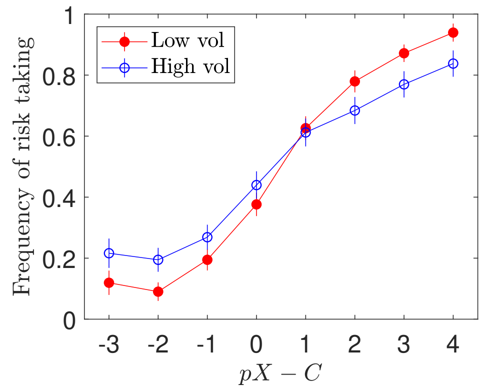


**Panel B**

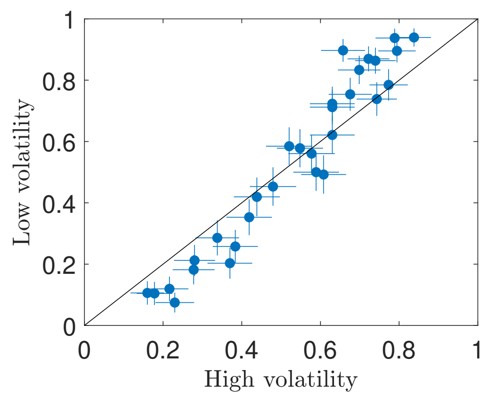


Figure V
Frequency of risk taking across volatility conditions

Panel A: the graph plots, for each volatility condition, the empirical frequency of risk taking against the difference in expected values between the risky lottery and the certain option, namely $pX - C$. The frequency of risk taking is computed as the proportion of trials on which subjects choose the risky lottery. Data are pooled across subjects over all common trials in the first condition, and thus represent between subjects comparisons. For each volatility condition, we bin the running variable, $pX - C$, to its nearest integer value, 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. Panel B: each point represents one of the 30 common trials in the first condition. The $x$-axis measures the frequency of risk taking in the high volatility condition, while the $y$-axis measures the frequency of risk taking in the low volatility condition. Inside each data point, the length of the vertical bar denotes two standard errors of the mean frequency of risk taking in the low volatility condition. The length of the horizontal bar denotes two standard errors of the mean frequency of risk taking in the high volatility condition.

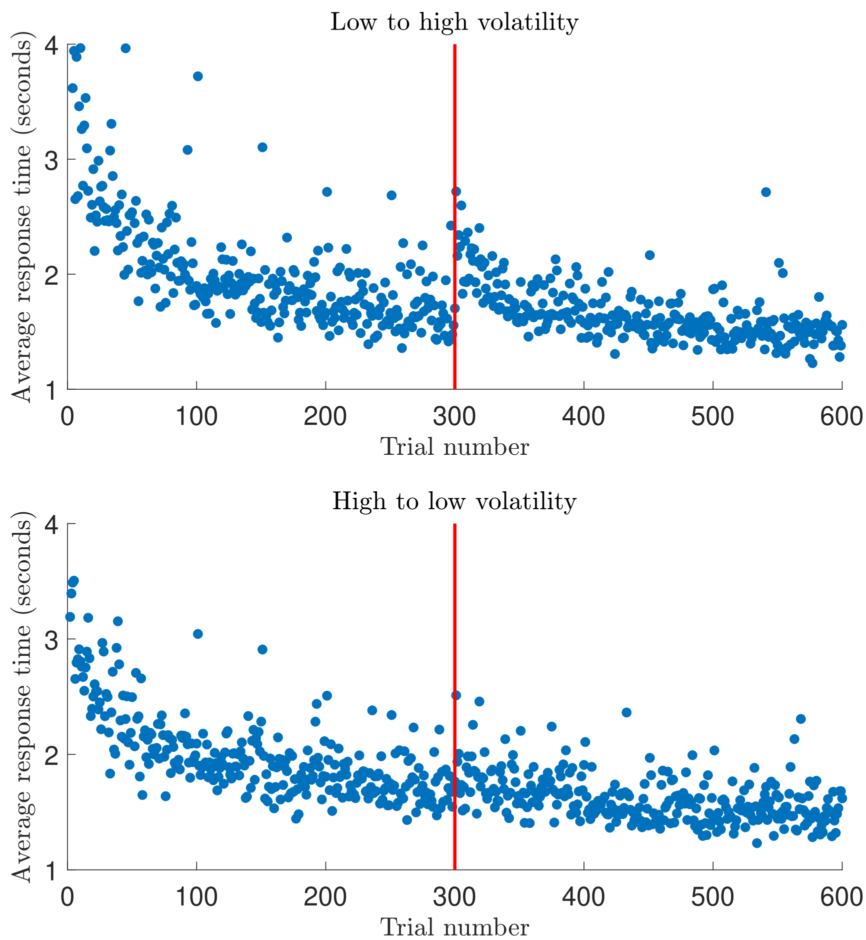


Figure VI
Evolution of response times over the 600 trials from the risky choice task (Experiment 1)

Each 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 vertical line denotes the onset of the change in environment. Data include both common and non-common trials.

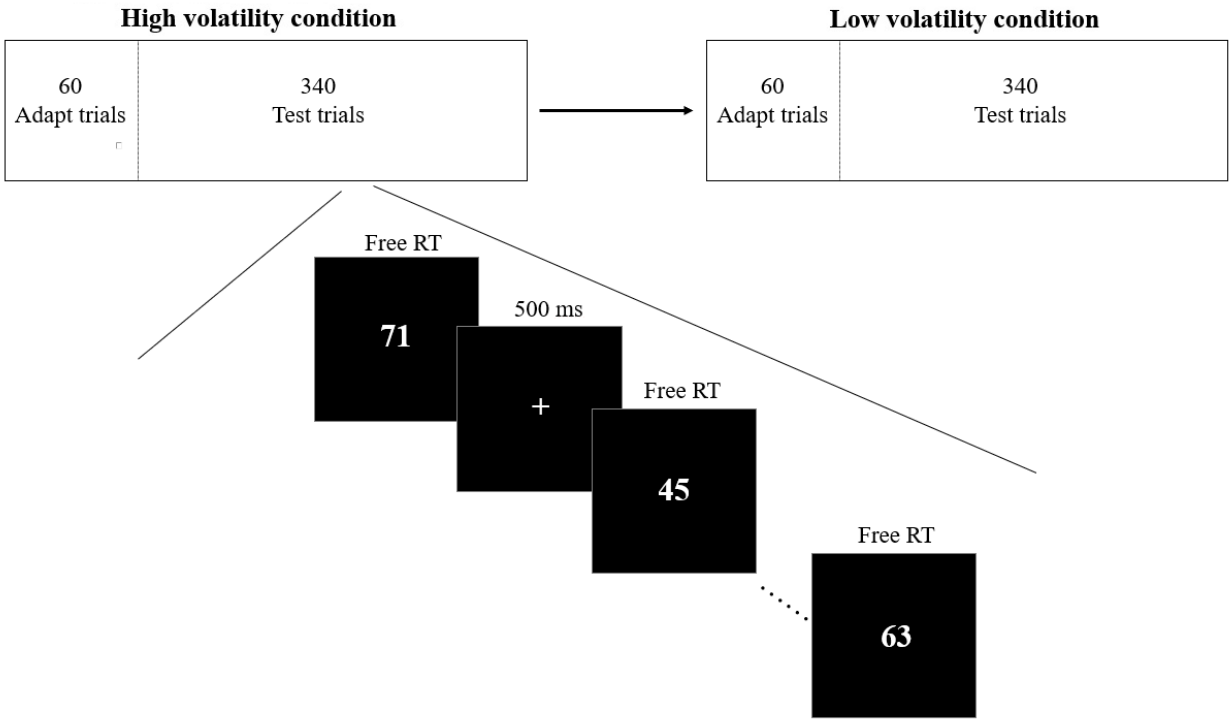


Figure VII
Design of the perceptual choice task in Experiment 1

The task consists of two blocks of trials: one block contains trials from the high volatility condition, and the other block contains trials from the low volatility condition. The order of the blocks is randomized across subjects. Each block begins with 60 “adaptation trials,” followed by 340 “test trials.” On each trial, subjects are incentivized to classify whether the number shown on the screen is greater or less than the reference level of 65.

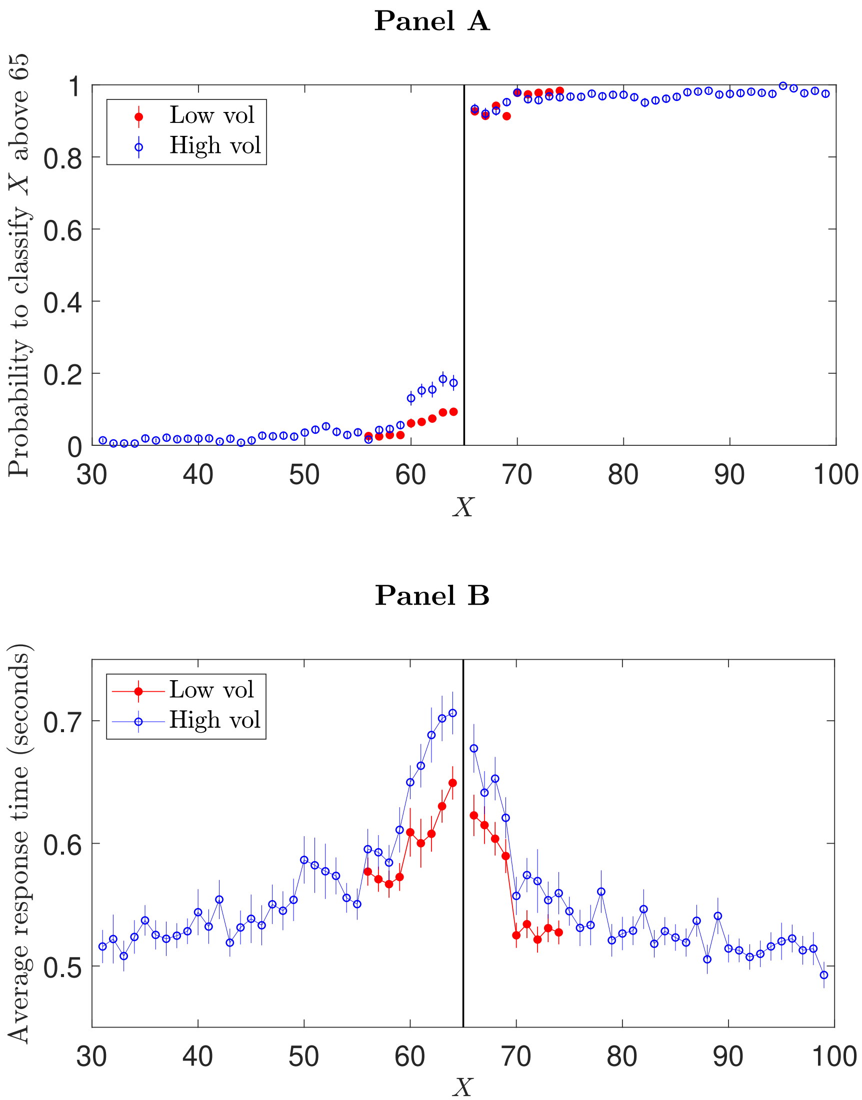


Figure VIII
Classification accuracy and response times for the perceptual choice task

Panel A: the $x$-axis denotes the integer $X$ that is presented on each trial, and the $y$-axis denotes the proportion of trials on which subjects classified $X$ as greater than 65. Panel B: the $y$-axis denotes the average response time for subjects to execute a decision, for trials on which subjects 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.

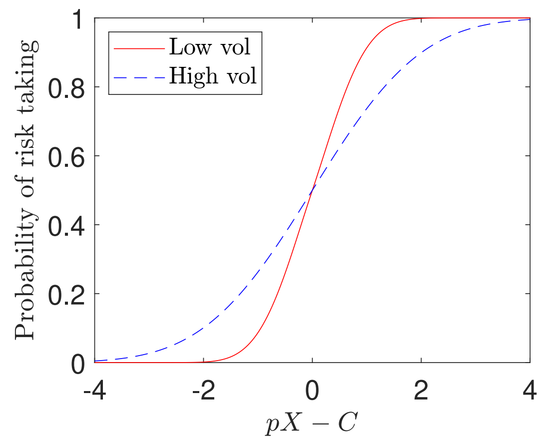


Figure IX
Model predicted probability of choosing the risky lottery when the capacity constraint parameter $n$ changes across volatility environments

The graph plots, for each of the two volatility environments (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 risk taking, defined in equation (17). The prior distributions for $X$ and $C$ are uniform. The probability $p$ that the risky lottery pays $X$ dollars is set to 0.5. The capacity constraint parameter $n$ is set to 7 for the low volatility environment and 12 for the high volatility environment. For each volatility environment, we draw $X$ uniformly from $[16, 24]$ and $C$ uniformly from $[8, 12]$. We then compute, for a given $X$ and $C$, the probability of risk taking. Finally, we aggregate these probabilities for each level of $pX - C$.

Table I

Frequency of choosing the risky lottery in Experiment 1 (volatility manipulation)

| | Between subjects tests | | Within subject tests | | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Dependent variable: "Choose risky lottery" | Unrestricted sample | Restricted sample | Restricted sample | Restricted sample (w/out trials 301-450) | Restricted sample —low volatility first (w/out trials 301-450) | Restricted sample —high volatility first (w/out trials 301-450) |
| *high* | 0.023 | 0.001 | –0.023 | –0.061 | 0.016 | –0.164 |
| | (0.191) | (0.200) | (0.097) | (0.113) | (0.168) | (0.139) |
| $X$ | 0.066*** | 0.074*** | 0.062*** | 0.064*** | 0.073*** | 0.052*** |
| | (0.006) | (0.006) | (0.004) | (0.004) | (0.006) | (0.006) |
| $C$ | –0.167*** | –0.186*** | –0.163*** | –0.170*** | –0.186*** | –0.155*** |
| | (0.013) | (0.012) | (0.009) | (0.009) | (0.012) | (0.014) |
| $X \times high$ | –0.017** | –0.023*** | –0.006 | –0.009** | –0.015*** | –0.002 |
| | (0.008) | (0.008) | (0.003) | (0.004) | (0.006) | (0.005) |
| $C \times high$ | 0.033* | 0.049*** | 0.013* | 0.025*** | 0.030*** | 0.019 |
| | (0.017) | (0.017) | (0.008) | (0.009) | (0.011) | (0.013) |
| Constant | 0.776*** | 0.787*** | 0.826*** | 0.850*** | 0.801*** | 0.943*** |
| | (0.166) | (0.179) | (0.106) | (0.123) | (0.180) | (0.135) |
| Observations | 4,470 | 4,170 | 8,257 | 6,411 | 3,125 | 3,286 |

*Notes.* The table reports results from mixed effects linear regressions in which 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 common trials are included. There are random effects on the independent variables $X$, $C$, and the intercept. Standard errors of the fixed effect estimates are clustered at the subject level and reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

Table II

Frequency of classifying $X$ as greater than 65 in the perceptual choice task

| | Between subjects tests | | | Within subject 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*** | 1.096*** | 0.578*** | 0.792*** | 1.051*** | 0.558*** |
| | (0.047) | (0.061) | (0.019) | (0.028) | (0.038) | (0.012) |
| $(X-65) \times high$ | –0.209*** | –0.279*** | –0.081*** | –0.147*** | –0.236*** | –0.077*** |
| | (0.051) | (0.069) | (0.025) | (0.014) | (0.023) | (0.011) |
| $high$ | 0.288*** | 0.405*** | –0.041 | 0.259*** | 0.300*** | 0.146** |
| | (0.052) | (0.070) | (0.101) | (0.037) | (0.047) | (0.064) |
| Constant | 0.065** | 0.071** | 0.269*** | 0.036* | 0.077*** | 0.119** |
| | (0.030) | (0.036) | (0.067) | (0.020) | (0.023) | (0.048) |
| Observations | 31,230 | 15,522 | 15,708 | 63,210 | 31,580 | 31,630 |

*Notes.* The table reports results from mixed effects logistic regressions in which the dependent variable takes the value of one if the subject classifies the integer $X$ as larger than 65, and zero otherwise. The integer $X$ is drawn uniformly from the set $[31,99]\backslash\{65\}$ in the high volatility condition, while it is drawn uniformly from the set $[56,74]\backslash\{65\}$ 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 independent variable $X-65$ and the intercept. Standard errors of the fixed effect estimates are clustered at the subject level and reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

Table III

Frequency of choosing the risky lottery in Experiment 2 (shape manipulation)

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| Dependent variable: "Choose risky lottery" | All common trials | All common trials | First half of each condition | Second half of each condition |
| *increasing prior* | 0.075*** | 0.075*** | 0.088*** | 0.070*** |
| | (0.023) | (0.023) | (0.029) | (0.024) |
| $X$ | | 0.045* | | |
| | | (0.024) | | |
| Constant | 0.690*** | 0.351* | 0.684*** | 0.695*** |
| | (0.029) | (0.188) | (0.032) | (0.030) |
| Observations | 2,278 | 2,278 | 862 | 1,416 |

*Notes.* The table reports results from mixed effects linear regressions in which 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. Only data from common trials are included. The variable $C$ is constant among common trials, and therefore is not included in the regressions as a control variable. There are random effects on the independent variable $X$ and the intercept. Standard errors of the fixed effect estimates are clustered at the subject level and reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

# Online Appendix

## A. Theoretical Derivations

*A.1. Equivalence of coding rules*

HWP have derived coding rules under three different performance objectives: one that maximizes expected financial gain, one that maximizes mutual information, and one that maximizes accuracy. In this section, we follow HWP and prove that, when the conditions in equation (9) of the main text

$$\begin{gathered}(i)\ pX \text{ and } C \text{ are i.i.d.}\\ \text{and } (ii)\ pX \text{ and } C \text{ are uniformly distributed}\end{gathered} \tag{A.1}$$

are satisfied, the three coding rules are equivalent. First, if the performance objective is to maximize expected financial gain, then the coding rules $\theta(X)$ and $\theta(C)$ in equations (5) and (6) of the main text come directly from HWP.

Next, if the performance objective is to maximize mutual information between a payoff ($X$ or $C$) and its noisy signal ($R_x$ or $R_c$), then HWP derive the following coding rules

$$\theta(X) = \left[\sin\left(\frac{\pi}{2}F(X)\right)\right]^2 \quad \text{and} \quad \theta(C) = \left[\sin\left(\frac{\pi}{2}F(C)\right)\right]^2, \tag{A.2}$$

where $F(X)$ and $F(C)$ are the cumulative distribution function of $f(X)$ and $f(C)$, respectively. When $X$ and $C$ are uniformly distributed, the coding rules in (A.2) reduce to equations (5) and (6). That is, maximizing mutual information and maximizing expected financial gain lead to the same coding rules. Finally, 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

$$\begin{aligned}&(i)\ \mathbb{E}[p\tilde{X}|R_x], \text{ viewed as a function of } R_x, \text{ and } \mathbb{E}[\tilde{C}|R_c], \text{ viewed as a function of } R_c,\\ &\text{are identical,}\\ &\text{and } (ii) \text{ the optimal coding rules are related: } \theta(X) = \theta(C = pX),\ \forall X \in [X_l, X_u].\end{aligned} \tag{A.3}$$

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

$$\mathbb{P}\text{rob}_{error} \equiv \int_{C_l}^{C_u} dC \int_{X_l}^{X_u} \mathbb{P}\text{rob}\left(error|\theta(X), \theta(C)\right) \cdot f(X)f(C) \cdot dX, \tag{A.4}$$

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

Combining (A.4) with the two conjectures, we observe that when $n$ is large,

$$\begin{aligned} &\mathbb{P}\text{rob}(error|\theta(X), \theta(C)) \\ &= \mathbb{P}\text{rob}(error|\theta(Y), \theta(C)) \approx \Phi\left(-\frac{|\theta(Y)-\theta(C)|}{\sqrt{\frac{\theta(Y)(1-\theta(Y)+\theta(C)(1-\theta(C)))}{n}}}\right), \end{aligned} \tag{A.5}$$

where $Y \equiv pX$, so $Y$ and $C$ are independently and identically distributed. Moreover, (A.4) can be written as

$$\begin{aligned} \mathbb{P}\text{rob}(error|\theta(X), \theta(C)) &= \int_{C_l}^{C_u} dC \int_{X_l}^{X_u} f(X) f(C) \cdot \mathbb{P}\text{rob}\left(error|\theta(X), \theta(C)\right) \cdot dX \\ &= \int_{C_l}^{C_u} dC \int_{C_l}^{C_u} f(Y) f(C) \cdot \mathbb{P}\text{rob}\left(error|\theta(Y), \theta(C)\right) \cdot dY. \end{aligned} \tag{A.6}$$

Given (A.6), 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.3) is straightforward. ■

### *A.2. Theoretical prediction of efficient coding in Experiment 2*

In this section, we prove an analytical result which justifies the claim that Experiment 2 in the main text targets a specific test of efficient coding, rather than a general test of noisy coding.

**Proposition**: Assume the design of Experiment 2. Further assume ($i$) $pX$ and $C$ are independently and identically coded, ($ii$) $f(R_x|X)$ is identical across the two experimental conditions (no efficient coding), and ($iii$) $f(R_c|C)$ is also identical across the two experimental conditions (no efficient coding). Then, for all values of $X$ and $C$, $\mathbb{P}\text{rob}(\text{risk taking}|X, C)$ is identical across the two conditions.

**Proof**: Given assumption ($i$), $\mathbb{E}[p\tilde{X}|R_x] = \mathbb{E}[\tilde{C}|R_c]$ when $R_x = R_c$. It is easy to show that this function of $R$ is increasing in $R$. As such, the probability of risk taking is

$$\begin{aligned} &\sum_{R_x=0}^{n} \sum_{R_c=0}^{n} \left( \mathbb{1}_{p \cdot \mathbb{E}[\tilde{X}|R_x] > \mathbb{E}[\tilde{C}|R_c]} \cdot f(R_x|X) \cdot f(R_c|C) + \mathbb{1}_{p \cdot \mathbb{E}[\tilde{X}|R_x] = \mathbb{E}[\tilde{C}|R_c]} \cdot \frac{1}{2} f(R_x|X) \cdot f(R_c|C) \right) \\ &= \sum_{R_x=0}^{n} \sum_{R_c=0}^{n} \left( \mathbb{1}_{R_x > R_c} \cdot f(R_x|X) \cdot f(R_c|C) \right) + \sum_{R=0}^{n} \frac{1}{2} f(R|X) \cdot f(R|C). \end{aligned} \tag{A.7}$$

Given assumptions ($ii$) and ($iii$), the last line in (A.7) is identical across the two conditions. In other words, under assumption ($i$), any difference in risk taking for a given $(X, C)$ across the two conditions serves as evidence that the likelihood functions, $f(R_x|X)$ and $f(R_c|C)$, respond endogenously to changes in the prior distribution (i.e., a violation of assumptions ($ii$) and ($iii$)). ■

# B. Experimental Instructions and Pre-Registration Documents

## *B.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:

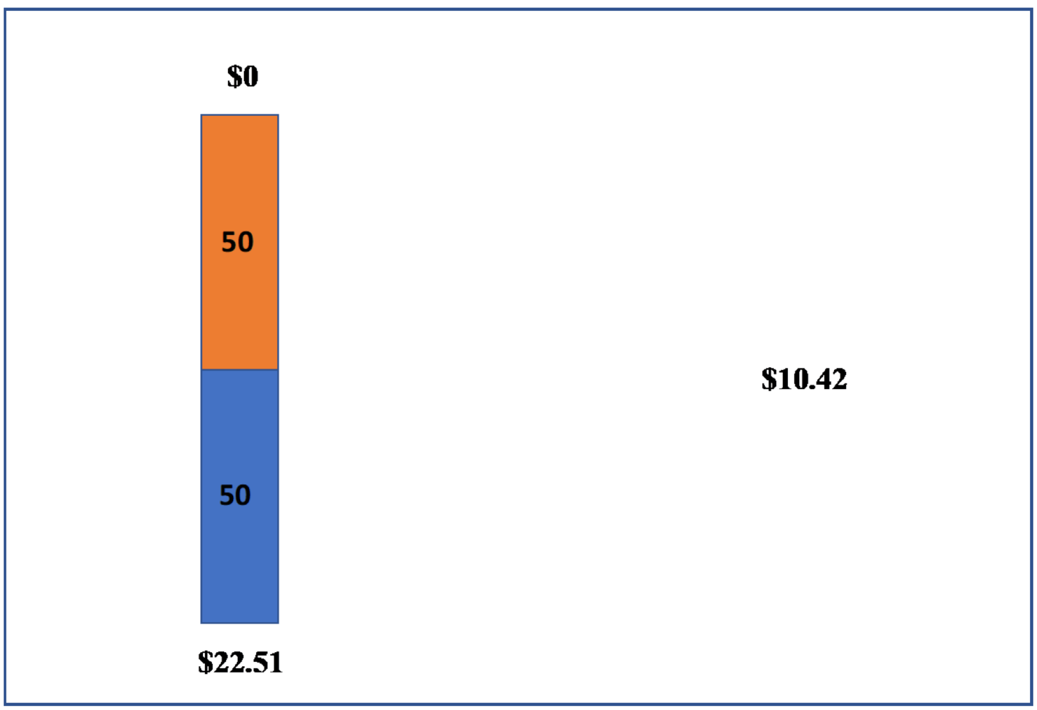


In the above example, the risky gamble is on the LEFT side of the screen and the sure thing is on the RIGHT side of the screen. The risky gamble pays $22.51 with 50% chance, and $0 with 50% chance. The sure thing pays $10.42 with 100% chance. For each question, you will be asked to select one of the 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.*

## *B.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:

Payout = \$ (15 x accuracy – 10 x 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 x 100% - 10 * 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 x 70% - 10* 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.*

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

# Welcome to the Experiment

Thank you for participating in this study!

Before we begin, please close all other applications. This study will last approximately 40 minutes.

For completing the experiment, you will earn $6.50, and you will also have the opportunity to earn an additional **bonus payment** depending on your answers during the experiment.

Click **Next** to see the instructions.

Next

# Instructions

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. Below is an example screenshot from the experiment:

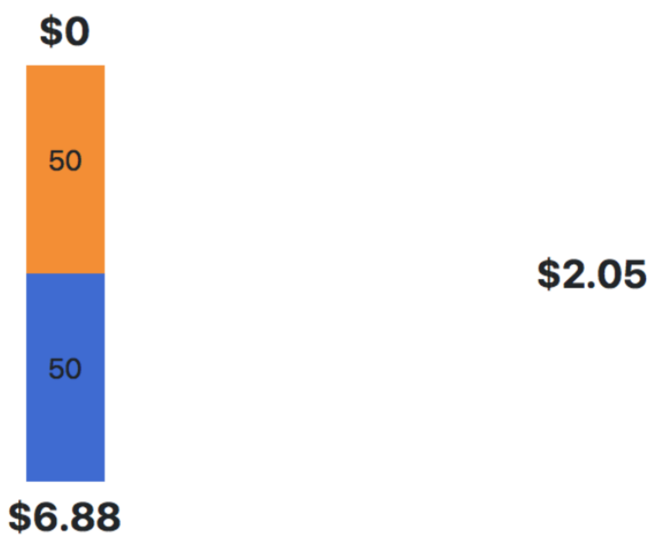


In the above example, the risky gamble is on the LEFT and the sure thing is on the RIGHT. The risky gamble pays $6.88 with 50% chance, and $0 with 50% chance. The sure thing pays $2.05 with 100% chance. For each question, you will be asked to select one of the two options by pressing either the "a" key for the LEFT option or the "k" key for the RIGHT option. For some questions, the risky gamble will be on the LEFT, and for the other questions it will be on the RIGHT. The dollar amounts will change on each question, so please pay attention to each question carefully.

Click **Next** to see how your bonus payment will be computed.

Next

# Instructions


**$2.05**

The experiment is broken down into 12 parts, for a total of 600 questions. At the end of the experiment, one question will be randomly selected by the computer, and **you'll be paid a bonus** according to your decision on that question. For example, if the above question was randomly chosen for the bonus question, and you selected the sure thing, you would be paid a **bonus** of $2.05. If instead you chose the risky gamble, you'd be paid a **bonus** of either $0 or $6.88, depending on which outcome the computer randomly selects.

**You will have a 10-second time limit to make each choice. If you fail to enter a response within 10 seconds, the computer will randomly select a decision for you and it will advance to the next question.** At the end of each part, you can take a short break.

Before you begin, you will go through a short comprehension check and two practice questions.

Click **Next** to start the comprehension check.

Next

## Comprehension Check

Please answer each question to the best of your ability.

**Question 1:**

**True or False: The possible dollar amounts you can win will change on every question.**

- True
- False

**Question 2:**

**True or False: One of the 600 questions will be randomly selected by the computer, and you will be paid a bonus depending on your answer to the randomly selected question.**

- True
- False

When you have answered the above questions, press **Next** to begin the 2 practice questions.

Next

## *B.4. Pre-registration document for Experiment 1*

**CONFIDENTIAL - FOR PEER-REVIEW ONLY**

ASPREDICTED

**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)

(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.

## B.5. *Pre-registration document for Experiment 2*

ASPREDICTED

Wharton CREDIBILITY LAB UNIVERSITY of PENNSYLVANIA

**CONFIDENTIAL - FOR PEER-REVIEW ONLY**

**EC Risk Taking with Monotonic Priors - Within Subjects on Prolific (#58963)**

**Created:** 02/21/2021 08:23 PM (PT)
**Shared:** 03/01/2021 05:56 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 hypothesize that the shape of the distribution from which risky payoffs are drawn affects a subject's level of 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 key 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" condition and a "decreasing" condition. In the increasing condition:

1) The payoff X will be drawn from a mixture of two uniform distributions: with probability 49/50, X will be drawn uniformly from [4.5, 8]; with probability 1/50, X will be drawn uniformly from [2, 4.5).

2) The payoff C will be drawn from a mixture of two uniform distributions: with probability 49/50, C will be drawn uniformly from [2.25, 4]; with probability 1/50, C will be drawn uniformly from [1, 2.25).

In the decreasing condition:

1) The payoff X will be drawn from a mixture of two uniform distributions: with probability 49/50, X will be drawn uniformly from [2, 5.5]; with probability 1/50, X will be drawn uniformly from (5.5, 8].

2) The payoff C will be drawn from a mixture of two uniform distributions: with probability 49/50, C will be drawn uniformly from [1, 2.75]; with probability 1/50, C will be drawn uniformly from (2.75, 4].

The design is within subjects, and the order of the two conditions is randomized across subjects. There are 300 trials in each condition, and we will insert 8 "common test trials" in each condition. The set of common test trials are the same across conditions, and the values of (X, C) for the 8 common test trials are given by:

(7.13, 2.70)
(7.26, 2.70)
(7.37, 2.70)
(7.49, 2.70)
(7.62, 2.70)
(7.76, 2.70)
(7.87, 2.70)
(7.99, 2.70)

The common test trials will be inserted on trials 90, 120, 150, 180, 210, 240, 270, and 300 of each condition. The order of the 8 common test trials is randomized across subjects and conditions.

**5) Specify exactly which analyses you will conduct to examine the main question/hypothesis.**
Our main analysis will be to compare the probability of choosing the risky lottery on the common test trials across the increasing and decreasing conditions. We therefore restrict our analysis to trials 90, 120, 150, 180, 210, 240, 270, 300 in each condition (for a total of 16 trials per subject). We will regress the probability of choosing the risky lottery on a condition dummy variable, which takes the value 1 if the trial belongs to the increasing condition, and 0 otherwise. Our main test is whether the estimated coefficient on the condition dummy variable is greater than zero.

**6) Describe exactly how outliers will be defined and handled, and your precise rule(s) for excluding observations.**
We will apply multiple rules for excluding observations, which we implement in the following order:

1) We will exclude any subject who fails to answer either of the two questions on a comprehension quiz, which will be given after the experimental instructions and before any risky choice decisions are made.

2) We will then exclude trials for which a subject fails to enter a response within the 10 second time limit.

3) We will then exclude subjects who exhibit sufficiently minimal variation in choice behavior on non-common test trials. That is, for those trials other than 90, 120, 150, 180, 210, 240, 270, and 300 in each condition, we compute the probability that each subject chooses the risky lottery. If this probability is strictly less than 2.5% or strictly greater than 97.5%, we exclude the subject.

4) We will then exclude any trial on which a subject responds in less than 0.5 seconds.

**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 data on Prolific until we obtain N=200 subjects who have completed the experiment.

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

# C. Predictions of Alternative Models

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

We consider predictions of the model 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, $x_1, x_2, \ldots, x_N$; 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

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

where

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

$\eta$ measures the relative importance of the gain-loss utility, and $\lambda$ 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's payoff distribution as the reference point distribution, 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.

## *C.2. Normalization models and decision-by-sampling models*

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

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

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

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

Panel A of Figure C.1 plots, for the two volatility conditions specified in Section III.A, the probability of choosing the risky lottery from equation (C.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 C.1 about here]**

As mentioned in Section V.C, our results from Experiment 2 cannot be explained by range normalization models. It is plausible that another form of normalization could explain the results, although the model would need to operate through more than just the range of payoff values. For example, Louie et al. (2015) discuss a model of divisive normalization, in which the value of an option is normalized by a function—not just the range—of all past values. Depending on the functional form of normalization, which is a large degree of freedom, such a model could also explain the difference in average levels of risk taking we observe in Experiment 2.

Lastly, we turn to the decision-by-sampling model. As discussed in HWP, DbS implies the following coding rules under the resource constraints in HWP:

$$\theta(X) = F(X), \qquad \theta(C) = F(C). \tag{C.5}$$

Panel B of Figure C.1 shows that DbS also predicts that a given increase in $X$ or a given decrease in $C$ leads to a larger increase 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 a 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.

**Panel A**

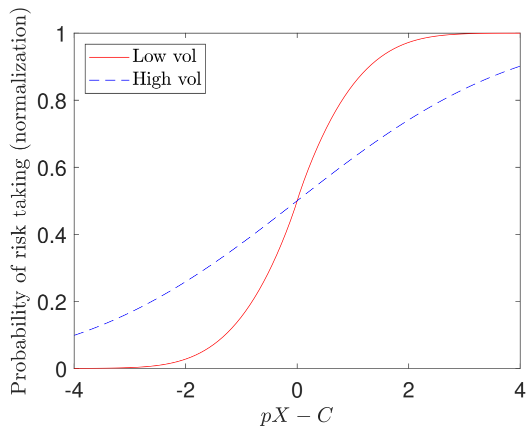


**Panel B**

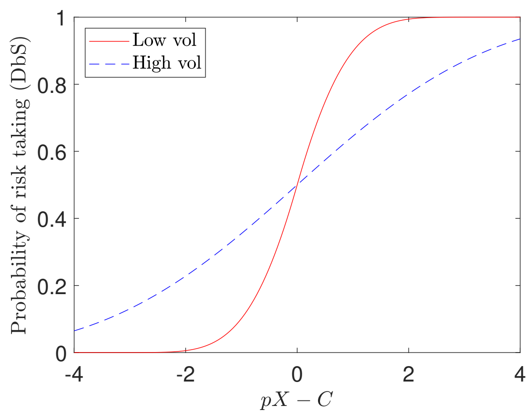


Figure C.1
Probability of choosing the risky lottery under alternative models

Panel A: the graph plots, for each of the two volatility environments (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 risk taking implied by the range normalization model of Rustichini et al. (2017). The prior 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 $\overline{v}/\chi$ is set to 15; as in Rustichini et al. (2017), this ratio is a measure of coding capacity. Panel B: the graph plots, for each of the two volatility environments described above, the probability of risk taking implied by decision-by-sampling (Stewart et al., 2006; Bhui and Gershman, 2018). The capacity constraint parameter $n$ is set to 10. In each panel and for each volatility environment, we draw $X$ uniformly from [16, 24] and $C$ uniformly from [8, 12]. We then compute, for a given $X$ and $C$, the probability of risk taking. Finally, we aggregate these probabilities for each level of $pX - C$.

# D. Additional Tables

Table D.1
Payoff values on the 30 common trials in Experiment 1

| $X$ | $C$ | $X$ | $C$ |
|---|---|---|---|
| \$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 |

*Notes.* The table presents the payoff values that comprise the 30 common trials in the risky choice task of Experiment 1. The set of 30 common trials is presented to subjects once in each volatility condition; this amounts to a total of 60 test trials per subject. The order of the common trials is randomized at the subject-condition level.

Table D.2

Frequency of choosing the risky lottery in Experiment 1 using logistic regressions

| | Between subjects tests | | Within subject tests | | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Dependent variable: "Choose risky lottery" | Unrestricted sample | Restricted sample | Restricted sample | Restricted sample (w/out trials 301-450) | Restricted sample —low volatility first (w/out trials 301-450) | Restricted sample —high volatility first (w/out trials 301-450) |
| *high* | 0.063 | –0.051 | –0.013 | –0.088 | 0.507 | –0.628 |
| | (0.856) | (0.934) | (0.440) | (0.541) | (0.777) | (0.660) |
| $X$ | 0.332*** | 0.407*** | 0.331*** | 0.353*** | 0.407*** | 0.270*** |
| | (0.039) | (0.041) | (0.027) | (0.029) | (0.041) | (0.040) |
| $C$ | –0.825*** | –0.987*** | –0.830*** | –0.894*** | –0.987*** | –0.749*** |
| | (0.088) | (0.094) | (0.061) | (0.067) | (0.094) | (0.089) |
| $X \times high$ | –0.108** | –0.171*** | –0.058*** | –0.111*** | –0.150*** | –0.034 |
| | (0.049) | (0.051) | (0.020) | (0.027) | (0.033) | (0.035) |
| $C \times high$ | 0.222** | 0.368*** | 0.117** | 0.239*** | 0.263*** | 0.129* |
| | (0.108) | (0.114) | (0.048) | (0.057) | (0.071) | (0.077) |
| Constant | 1.188 | 1.283 | 1.365*** | 1.512*** | 1.283 | 1.860*** |
| | (0.754) | (0.848) | (0.486) | (0.584) | (0.850) | (0.639) |
| Observations | 4,470 | 4,170 | 8,257 | 6,411 | 3,125 | 3,286 |

*Notes.* The table reports results from logistic regressions in which 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. Data are pooled across all subjects and standard errors are clustered at the subject level. Only data from common trials are included. Standard errors are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

Table D.3
Payoff values on the 8 common trials in Experiment 2

| $X$ | $C$ |
|---|---|
| \$7.13 | \$2.70 |
| \$7.26 | \$2.70 |
| \$7.37 | \$2.70 |
| \$7.49 | \$2.70 |
| \$7.62 | \$2.70 |
| \$7.76 | \$2.70 |
| \$7.87 | \$2.70 |
| \$7.99 | \$2.70 |

*Notes.* The table presents the payoff values that comprise the 8 common trials in Experiment 2. The set of 8 common trials is presented to subjects once in each condition, on trials 90, 120, 150, 180, 210, 240, 270, and 300; this amounts to a total of 16 test trials per subject. The order of the common trials is randomized at the subject-condition level.

Table D.4
Frequency of choosing the risky lottery in Experiment 1: Comparison between early and late trials

| | (1) | (2) | (3) |
|---|---|---|---|
| Dependent variable: "Choose risky lottery" | Trials 31-165 | Trials 166-300 | Trials 31-300 |
| *high* | –0.007 | 0.012 | –0.025 |
| | (0.241) | (0.247) | (0.243) |
| $X$ | 0.077*** | 0.073*** | 0.075*** |
| | (0.007) | (0.007) | (0.007) |
| $C$ | –0.189*** | –0.182*** | –0.189*** |
| | (0.013) | (0.016) | (0.013) |
| $X \times high$ | –0.024*** | –0.022** | –0.023** |
| | (0.009) | (0.009) | (0.010) |
| $C \times high$ | 0.054*** | 0.045** | 0.054*** |
| | (0.108) | (0.020) | (0.018) |
| *second* | | | –0.010 |
| | | | (0.224) |
| $high \times second$ | | | 0.057 |
| | | | (0.287) |
| $X \times second$ | | | –0.002 |
| | | | (0.008) |
| $C \times second$ | | | 0.007 |
| | | | (0.014) |
| $X \times high \times second$ | | | 0.001 |
| | | | (0.010) |
| $C \times high \times second$ | | | –0.009 |
| | | | (0.018) |
| Constant | 0.755*** | 0.768*** | 0.790*** |
| | (0.204) | (0.215) | (0.205) |
| Observations | 2,147 | 2,023 | 4,170 |

*Notes.* The table reports results from the first condition of the risky choice task in Experiment 1. We present mixed effects linear regressions in which 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. The dummy variable, *second*, takes the value of one if the trial belongs to the second half of the first condition (trials 166-300), and zero if it belongs to the first half (trials 31-165). Only data from common trials are included. There are random effects on the independent variables $X$, $C$, and the intercept. Standard errors of the fixed effect estimates are clustered at the subject level and reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.