Efficient Coding and Risky Choice

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

This Version: March, 2019

Abstract: We present a model of risky choice in which the decision maker (*DM*) perceives a lottery payoff with noise due to the brain's limited capacity to represent information. We model perception using the principle of *efficient coding*, which states that stimuli that occur more frequently are perceived more accurately. We show that it is efficient for risk taking to be more sensitive to those payoffs that the *DM* encounters more frequently. Our model also predicts that the *DM*'s value function is malleable and its curvature fluctuates with the recently encountered distribution of payoffs. To test the model, we conduct a laboratory experiment in which we manipulate the distribution of payoffs across 480 choice sets. Consistent with the efficient coding of monetary payoffs, we find that risk taking is more sensitive to those payoffs that are presented more frequently. Moreover, sensitivity to extreme payoffs is initially low, but grows over time after repeated exposure. In a second experiment, we conduct an additional test of the efficient coding mechanism by incentivizing subjects to classify which of two symbolic numbers is larger. We find that accuracy is higher for those numbers that the subject has more frequently observed, providing further evidence that perception of a given numerical quantity varies with the recent environment. Overall, our experimental results suggest that risk taking depends systematically on the payoff distribution to which the *DM*'s perceptual system has recently adapted.

[*] We thank Nicholas Barberis, John O'Doherty, Alex Imas, Antonio Rangel, Andrei Shleifer, Michael Woodford, and seminar participants at Caltech, the Chinese University of Hong Kong, the University of New South Wales, the University of Technology Sydney, the University of Utah, the University of California, San Diego, Washington University in St. Louis, the NBER behavioral finance meeting, and the Sloan-Nomis Workshop on the Cognitive Foundations of Economic Behavior for helpful comments. Frydman thanks the NSF for financial support.

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

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

I. Introduction

When choosing between two lotteries, the decision maker (called "*DM*" hereafter) first perceives the set of payoffs from each lottery and then executes a decision. Because there are constraints on the degree to which the brain can process information, the perception of numerical quantities is inherently noisy (Dehaene, 2011). Understanding precisely how these constraints affect perception of payoffs has the potential to generate new insights about risk taking, and in particular, its instability over time. For example, decades of experiments have shown that one source of instability is the sequence of outcomes that the subject has experienced: past gains and losses have a systematic effect on subsequent risk taking (Thaler and Johnson, 1990; Weber and Camerer, 1998; Imas, 2016). A different potential source of instability is the *DM*'s perception of a payoff, which can vary systematically with the payoffs that she has recently observed.

Why would the *DM*'s perception of a given risky payoff vary across different environments? If the mechanism used for perceiving payoffs is similar to the one used for perceiving sensory stimuli such as light or sound, then it may in fact be optimal to hold different perceptions of the same payoff in different environments. Specifically, a core principle in neuroscience called *efficient coding*, states that the brain should allocate resources so that perception is more sensitive to those stimuli that are expected to occur more frequently (Barlow, 1961; Laughlin, 1981). This principle explains why we are temporarily "blinded" when moving from a dark room to a brightly lit one, because resources have not yet been adjusted for perceiving objects in the new bright environment. If the principle of efficient coding extends to the domain of risky choice, this can provide a normative foundation for the variation in risk taking across environments.

In this paper, we present a model of choice under risk in which the perception of payoffs is governed by efficient coding; we then test the model experimentally to assess whether risk taking varies with the recently encountered payoff distribution. Our model builds on the recent theoretical work of Woodford (2012) and Khaw, Li, and Woodford (2019) (hereafter "KLW"), who assume that the perception of risky payoffs is imperfect and is estimated through Bayesian inference. Our model departs from KLW by allowing the *DM*'s likelihood function to adapt to any continuous distribution of payoffs, while efficiently using all perceptual resources. Specifically, we derive the likelihood function that is efficient in the sense of maximizing mutual information between the noisy signal generated by the *DM*'s perceptual system and the true payoff (Wei and Stocker, 2015). These efficient likelihood functions, when combined with the *DM*'s prior expectations, generate a subjective value function that exhibits several features from prospect theory, including reference dependence and diminishing sensitivity (Kahneman and Tversky, 1979).

Importantly, our model generates novel predictions about how these features of the value function vary with the payoff distribution to which the *DM* has adapted. For example, the degree of diminishing sensitivity over a given range of payoffs is tied directly to how frequently these payoffs occur in the environment. To see this, consider an environment where the upside of a risky lottery is often in the range between $10 and $20, in which case, efficient coding implies that perceptual resources are allocated towards discriminating between payoffs in this range. In this same environment, if the upside is occasionally increased from $30 to $40, then risk taking will not increase much because the *DM*'s perceptual system cannot easily distinguish between these two infrequent amounts. However, if the overall distribution of payoffs changes, so that the upside frequently falls between $30 and $40, then the *DM* can easily perceive this difference, and risk taking will increase substantially when the upside is increased from $30 to $40. Thus, diminishing sensitivity arises from efficient coding, and crucially, the curvature of the value function fluctuates with the payoff distribution to which the *DM* has adapted.

While efficient coding generates a value function that is relatively flat over payoffs that occur infrequently, the value function exhibits its steepest slope around the most frequently occurring payoff. In particular, if we take the "reference point" in prospect theory to be the value for which the *DM* has highest marginal utility, then efficient coding delivers the most frequently occurring payoff as the reference point. More generally, efficient coding implies that marginal utility is high for those payoffs that occur frequently, because these are the payoffs that the *DM* can precisely discriminate between. This generates a strong testable prediction, namely, that risk taking should become more sensitive to a given payoff when shifting to an environment where that payoff occurs more frequently.

To test this prediction, we conduct a laboratory experiment in which subjects make a series of 480 decisions between a risky lottery and a certain option. While such a large number of trials is not typical in economic experiments on risky choice, this design feature enables us to carefully vary the payoff distribution within subjects, over time. We manipulate the distribution of risky payoffs across two conditions: one in which payoffs in the choice set are drawn from a distribution with high volatility, and the other in which the distribution has low volatility. All payoff values in a given block (i.e., a consecutive set of trials) are drawn from either the high volatility or low volatility distribution. Consistent with our model, we find that risk taking is more sensitive to payoffs in the low volatility blocks, compared to the high volatility blocks.

In addition to comparing risk taking across the high volatility and low volatility blocks, we also analyze behavior *within* a block. This allows us to explore the dynamic process by which subjects adapt to a new distribution. We find that over the course of sixty trials in a high volatility block, subjects exhibit a net increase in their sensitivity to extreme values. This occurs because perceptual resources that were previously devoted to intermediate values in the low volatility block, are reallocated towards extreme payoffs that are now more likely to appear in the high volatility block. Our data indicate that this shift in sensitivity occurs relatively quickly, on the order of ten experimental trials, which is consistent with recent work on adaptation in the perceptual domain (Payzan LeNestour and Woodford, 2018).[1]

While we formally present the full model later in the paper, we briefly explain the two basic assumptions and mechanisms here. First, as in KLW, we assume the decision maker encodes each risky payoff with noise. Specifically, when the *DM* is presented with a choice set in which a risky lottery pays X dollars in some state, we assume that the *DM* perceives this payoff

[1] This shift in sensitivity is consistent with a simple model of adaptation that we propose, in which the *DM* uses the last N trials to construct a prior distribution over the payoffs on the current trial. Critically, these last N trials also pin down the *DM*'s likelihood functions through efficient coding. We then use the prior and the implied likelihood functions to generate predictions about how risk taking changes over time; we find that the model provides a good match to the experimental data when assuming that the *DM* uses the last $N = 10$ trials to construct her prior distribution.

as a noisy signal, R_x, which is governed by a probability density function $p(R_x \mid X)$. This assumption captures a fundamental feature of numerical cognition, that our perception of numerical quantities is noisy, even when these quantities are presented as Arabic numerals (Dehaene, 2011). Our second assumption is that the *DM* uses Bayesian inference to compute the optimal estimate of the numerical payoff under consideration.[2] Importantly, both ingredients of Bayesian inference—the prior and the likelihood function—are pinned down by the recent payoff distribution, which provides an extra layer of discipline in the Bayesian framework (Wei and Stocker, 2015). After the *DM* performs Bayesian inference, she chooses the lottery with the maximum estimated expected value.

In our model, we assume that only payoffs are subject to noisy encoding, but that probabilities are perceived without noise. This assumption is for simplicity, and in reality, state probabilities are also likely to be encoded with noise. To further test whether our experimental results are generated by the noisy encoding of payoffs, we run an additional experiment in which the subject still needs to perceive numerical quantities, but there is no need to perceive probabilities or integrate them with payoffs. We run a *riskless* choice experiment where we incentivize subjects to classify whether a number displayed on each trial is above or below a reference number. We find that even in this simpler environment, accuracy depends on the distribution of numbers to which the subject has adapted. For a given number, subjects exhibit greater classification accuracy if the number has occurred more frequently in the recent past. This provides support for our basic model assumption that the perception of symbolic numbers is noisy and changes across environments.

Our model is meant to capture intuitive judgments about choice under risk, such as the judgments between simple gambles that Kahneman and Tversky (1979) sought to explain with prospect theory. Our model does not apply to all decisions under risk, and in particular, it should not be applied to decisions that are based on explicit symbolic calculations. These decisions are instead likely to be governed by a distinct decision-making system (Deheane, 1992). At the same

[2] This assumption is motivated by the literature on sensory perception which finds a tight link between quantitative predictions from a Bayesian framework and data from controlled experiments (Stocker and Simoncelli, 2006; Girshick, Landy, and Simoncelli, 2011; Wei and Stocker, 2015; 2017).

time, our model is not necessarily confined to low stakes decisions, and we believe that it is reasonable to apply in situations similar to those where prospect theory has found success (see Barberis 2013 for a review).

Our paper contributes to a recent literature that examines the effect of imperfect perception and Bayesian inference on economic choice. Gabaix and Laibson (2017) show theoretically that a *DM* with a discount rate of one will appear impatient if payoffs delivered farther in the future are perceived with more noise. Woodford (2012) and KLW provide a framework in which a *DM* with linear utility can appear risk averse if payoffs are encoded with noise. Steiner and Stewart (2016) show that Bayesian inference can generate an overweighting of small probability events, as in prospect theory.[3] Both of our experiments provide evidence that supports the type of perceptual processes proposed in these Bayesian models of economic choice.

More generally, our paper adds to a growing literature that builds cognitive and perceptual foundations for the psychological assumptions in behavioral economics. For example, several behavioral models of financial markets have shown that prospect theory is an important ingredient in explaining puzzling facts such as the high equity premium (see Barberis 2018 for a review). Here, we show that efficient coding can be viewed as one perceptual foundation for prospect theory, which in turn can help explain these facts.[4] Importantly, our theory generates new predictions about the malleability of preferences, which motivates further empirical tests both in the lab and in the field.

Finally, our results also contribute to a literature that uses basic neural computations to constrain patterns of risky choice (Tymula and Glimcher, 2017; Landry and Webb, 2018). A particularly relevant neural computation is that of *normalization,* in which the brain normalizes stimulus values according to the distribution of values in the environment. Several experiments

[3] In related work, Bhui and Gershman (2018) show how efficient coding can provide a normative foundation for a model of multi-attribute decision making called decision by sampling (Stewart, Chater, and Brown, 2006).

[4] Bordalo, Gennaioli, and Shleifer (2012) propose a different perceptual mechanism, called salience theory, that can also generate several features from prospect theory. Later in the paper, we discuss how efficient coding relates to the assumptions in salience theory.

have found evidence consistent with normalization in the brain (Tobler, Fiorillo, and Schultz, 2005; Padoa-Schioppa 2009; Carandini and Heeger, 2012; Rangel and Clithero, 2012; Louie and Glimcher, 2012), but there is less evidence that this process has an associated effect on behavior. Khaw, Glimcher, and Louie (2017) show that the valuation of consumer goods negatively correlates with the average value of recently encountered items, and Polania, Woodford, and Ruff (2019) provide evidence that valuation depends on the entire distribution. Here, we demonstrate that these adaptation effects extend into the domain of risky choice.[5]

The paper proceeds as follows. In Section II, we lay out the basic elements of the model and analyze the model's implications. Section III describes the main experiment of the paper, a risky choice experiment, and discusses its results. Section IV follows with a riskless choice experiment. Section V provides additional discussions. Section VI concludes and suggests directions for future research.

II. The Model

In this section, we develop a static model of risky choice based on efficient coding, following the recent work of KLW and Wei and Stocker (2015, 2017). We derive the *DM*'s value function that characterizes the perception of a payoff, and we show how this perception depends crucially on expectations of the payoff distribution. We then extend the model to incorporate a simple process of adaptation, which generates predictions about the dynamics of choice.

II.1. Choice environment

The *DM* faces a choice set that contains two options: a certain option and a risky lottery. The certain option, denoted as (C, 1), pays C dollars with certainty. The risky lottery, denoted as

[5] See also Payzan-LeNestour, Balleine, Berrada, and Pearson (2016) for experimental evidence on adaptation to variance, and Zimmermann, Glimcher, and Louie (2018) for evidence on adaptive behavior in monkeys in the realm of risky choice.

(X, p; 0, 1 – p), pays X dollars with a probability p and zero dollars with the remaining probability 1 – p. The *DM*'s task is to choose between these two options.

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

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

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

Motivated by the literature on sensory perception, we depart from the expected utility framework by assuming that the *DM* imperfectly perceives the payoffs X and C (Deheane, 2011; Girshick et al., 2011; Wei and Stocker, 2015).[6] We model this imperfect perception in a Bayesian framework, such that, before observing the choice set, the *DM* has a prior distribution over X and C.[7] Upon observing the choice set, the presentations of X and C generate a noisy signal R_x for X, and a noisy signal R_c for C, and each of these noisy signals is randomly drawn from a distinct likelihood function. The *DM* then forms optimal estimates $\mathbb{E}[\tilde{X} | R_x]$ and $\mathbb{E}[\tilde{C} | R_c]$, where $\tilde{X}$ and $\tilde{C}$ are random variables that are associated with the posterior beliefs about X and C, respectively. These estimates are optimal in the sense that they miminize the mean squared error between the estimates and the true values of X and C. As in KLW, we assume that the *DM* has linear utility, and thus chooses the risky lottery if and only if $p \cdot \mathbb{E}[\tilde{X} | R_x] > \mathbb{E}[\tilde{C} | R_c]$.

It is worth noting that the encoding process described above is conditional on the values of X and C, which we assume are perfectly observable to the econometrician but not to the *DM*. That is, even after the *DM* is presented with the choice set and uses the noisy signals to form posterior beliefs, she still faces uncertainty about the payoff values in the current choice set. The inference process therefore takes place at the level of a single choice set, and characterizes how

[6] Further evidence for this assumption comes from recent experimental work which demonstrates that humans have single neurons that selectively and stochastically respond to "preferred" numbers (Kutter, Bostroem, Elger, Mormann, and Nieder, 2018). Such "number neurons" are likely to generate the noisy perception of symbolic numbers.

[7] We assume in this model that the probability p is perceived without noise. In our experimental design (see Section III), we set p to a constant across trials so that, through learning, it is plausible that the *DM* perceives the precise value of p.

the *DM*'s prior beliefs shift after observing a noisy signal of the true payoff. The noisy encoding of payoffs drives our main model predictions. In the next section, we derive this noisy encoding process under efficient coding.

II.2. The likelihood function

We begin our discussion of the likelihood function with a general stimulus value, θ, and its associated noisy signal, m. We further denote the conditional probability density function of the noisy signal for a given stimulus value as $p(m|\theta)$. The likelihood function is then defined as[8]

$$L(\theta \mid m) = p(m \mid \theta). \tag{2}$$

This function governs the likelihood of each stimulus value θ conditional on the noisy signal m. We apply the efficient coding criterion proposed in Wei and Stocker (2015) to constrain the likelihood function. This criterion requires

$$\sqrt{J(\theta)} \propto p(\theta), \tag{3}$$

where Fisher information $J(\theta)$ is given by

$$J(\theta) = \int \left(\frac{\partial \ell\text{n}\, p(m \mid \theta)}{\partial \theta} \right)^2 p(m \mid \theta) dm, \tag{4}$$

and $p(\theta)$ is the true probability density function of the stimulus value θ.

Intuitively, Fisher information $J(\theta)$ measures the amount of coding resources allocated towards perception of a given stimulus value θ. As a result, the efficient coding condition (3) implies that encoding accuracy is greater for stimulus values that occur more frequently. We note that condition (3) is one of several potential definitions of efficient coding in the domain of risky choice, and we use this particular one because of its accuracy in modeling data from perceptual experiments (Wei and Stocker, 2015), and more recently, an economic experiment (Polania, Woodford, and Ruff, 2019). Later in the paper, we discuss alternative coding schemes that could be applied to risky payoffs.

[8] See Chapter 6 of Casella and Berger (2002) for more discussion about the likelihood function.

Proposition 1 characterizes a set of likelihood functions, one for each realization of the noisy signal, that satisfies the efficient coding condition (3).

Proposition 1. Assume the stimulus value has a probability density distribution $p(\theta)$, where θ takes the range of $(-\infty, \infty)$. The cumulative density function is therefore $F(\theta) = \int_{-\infty}^{\theta} p(\xi)d\xi$. The following likelihood function

$$L(\theta \mid m) = \frac{1}{\sqrt{2\pi}\cdot\sigma}\exp\left(-\frac{(F(\theta)-m)^2}{2\sigma^2}\right), \tag{5}$$

satisfies the efficient coding condition of (3), where m also takes the range of $(-\infty, \infty)$.

Proof of Proposition 1. See Appendix A.

This likelihood function contains one free parameter, σ, which represents the noise in the likelihood function. It also has a more fundamental interpretation in our framework of efficient coding: σ is decreasing in the amount of coding resources that are available to the *DM*. If the *DM* is endowed with a small amount of coding resources, then σ is large, and thus the likelihood function will be noisy. Conversely, as the amount of coding resources tends to infinity, σ converges to zero; in this case, the *DM* encodes all payoffs without noise and our model reduces to expected utility theory.

In order to illustrate the implications of Proposition 1 for risky choice, we now specify the stimulus value θ to be X or C, and specify the associated noisy signal m to be R_x or R_c. We further assume that the probability density functions of X and C are lognormal

$$\begin{aligned} p(X;\mu_x,\sigma_x) &= \frac{1}{\sqrt{2\pi}\cdot\sigma_x X}\exp\left(-\frac{(\ell\text{n}X-\mu_x)^2}{2\sigma_x^2}\right), \\ p(C;\mu_c,\sigma_c) &= \frac{1}{\sqrt{2\pi}\cdot\sigma_c C}\exp\left(-\frac{(\ell\text{nC}-\mu_c)^2}{2\sigma_c^2}\right). \end{aligned} \tag{6}$$

This will also be the type of stimulus distribution that we use in the laboratory experiment later in the paper. Under this assumption of lognormality, the likelihood functions are given by:

$$L(X \mid R_x) = p(R_x \mid X) = \frac{1}{\sqrt{2\pi} \cdot \sigma} \exp\left(-\frac{(\Phi((\ell n\, X - \mu_x)/\sigma_x) - R_x)^2}{2\sigma^2} \right),$$
$$L(C \mid R_c) = p(R_c \mid C) = \frac{1}{\sqrt{2\pi} \cdot \sigma} \exp\left(-\frac{(\Phi((\ell n\, C - \mu_c)/\sigma_c) - R_c)^2}{2\sigma^2} \right). \tag{7}$$

The expressions in (7) characterize the likelihood functions of *X* and *C*, and we emphasize that they depend directly on parameters of the stimulus distributions, μ_x, σ_x, μ_c, and σ_c. This dependence on the stimulus distribution is the key source of variation in perception of payoffs.

To illustrate this point, in Figure 1 we plot the lognormal stimulus distribution of *X* with two different sets of parameter values, as well as the implied likelihood function $L(X \mid R_x)$ for several values of R_x. For the low volatility distribution, we set $\sigma_x = 0.19$. For the high volatility distribution, we set $\sigma_x = 0.55$ (we also use these volatility levels in our experimental test in Section III). To get a sense of the range of values of R_x, we compute the unconditional distribution of R_x as

$$p(R_x) = \int_0^\infty p(R_x \mid X)\, p(X)\, dX. \tag{8}$$

We find that large values of *X*—values from the right tail of the lognormal distribution $p(X;\, \mu_x,\, \sigma_x)$—tend to generate values of R_x that are close to one. Conversely, small values of *X* tend to generate values of R_x that are close to zero.[9]

[**Place Figure 1 about here**]

Figure 1 highlights some important features of the likelihood function $L(X \mid R_x)$. For a given stimulus distribution, the shape of the likelihood function depends significantly on the value of R_x. In line with the core principle of efficient coding, those values of *X* that occur more frequently tend to generate values of R_x that give rise to likelihood functions with lower dispersion. Moreover, for each value of R_x, changing the stimulus distribution—in particular,

[9] It is easy to check that, with the stimulus distribution in (6) and the likelihood function in (7), $p(R_x)$ does not depend on μ_x or σ_x. More generally, the shape of $p(R_x)$ does not depend on the shape of the underlying stimulus distribution: all continuous stimulus distributions lead to the same $p(R_x)$. Moreover, as σ goes to zero—that is, as the amount of coding resources tends to infinity—$p(R_x)$ converges to a uniform distribution between zero and one. In Appendix A, we plot $p(R_x)$ and provide a more detailed discussion of its properties.

changing σ_x—alters the shape of the likelihood function. For example, a lower standard deviation of the stimulus distribution results in lower dispersion of the likelihood function. Intuitively, when the stimulus distribution has a lower standard deviation, the *DM* allocates her finite coding resources to a narrower range of stimulus values, resulting in higher discriminability between these stimulus values as measured by the lower dispersion of the likelihood function.

II.3. Bayesian inference

The *DM* uses the likelihood function derived in the previous section, in conjunction with her prior belief about the stimulus distribution, to form a posterior belief about each payoff in the choice set. We assume that the *DM* uses the mean of the posterior distribution as her estimate of each payoff. These posterior means for *X* and *C*, conditional on R_x and R_c, are given by

$$\mathbb{E}[\tilde{X}|R_x] = \frac{\int_0^\infty p(R_x \mid X)\, p_0(X)\, X dX}{\int_0^\infty p(R_x \mid X)\, p_0(X)\, dX} \tag{9}$$

and

$$\mathbb{E}[\tilde{C}|R_c] = \frac{\int_0^\infty p(R_c \mid C)\, p_0(C)\, C dC}{\int_0^\infty p(R_c \mid C)\, p_0(C)\, dC}, \tag{10}$$

where $p_0(X)$ and $p_0(C)$ are the *DM*'s prior beliefs about *X* and *C*, respectively. For now, we assume these priors coincide with the true stimulus distributions described in (6)[10]

$$p_0(X) = p(X;\mu_x,\sigma_x), \qquad p_0(C) = p(C;\mu_c,\sigma_c). \tag{11}$$

After the *DM* forms posterior beliefs for each payoff, she chooses the risky lottery if and only if $p \cdot \mathbb{E}[\tilde{X}|R_x] > \mathbb{E}[\tilde{C}|R_c]$. The left hand side of this inequality provides the *DM*'s subjective expected value of the risky lottery, while the right hand side provides her subjective expected

[10] This assumption can be justified if the *DM* has fully adapted to the stimulus distribution, and thus is meant to capture beliefs in a "steady state."

value of the certain option. Later in Section II.6, we present a simple model of adaptation and study the dynamics of choice by relaxing the assumptions in (11).

II.4 Value function

The noisy encoding process described above implies that, the same payoff X will generate different noisy signals when it is presented on different occasions. Because each realized value of the noisy signal R_x maps to a different posterior mean as shown in (9), the *DM* faces a *distribution* of subjective valuations for each value of X. Importantly, the *average* subjective valuation of X will, in general, be different from X itself.

To see this, we compute the subjective valuation of X, averaged over the conditional probability density of R_x, $p(R_x \mid X)$. We denote this average subjective valuation by $v(X)$, which is given by:

$$v(X) = \int_{-\infty}^{\infty} \mathbb{E}[\tilde{X} \mid R_x]\, p(R_x \mid X)\, dR_x. \tag{12}$$

Given the randomness in the noisy signal R_x, we can also compute the standard deviation of the subjective valuation as

$$\sigma(X) = \left[\int_{0}^{\infty} (\mathbb{E}[\tilde{X} \mid R_x])^2\, p(R_x \mid X)\, dR_x - v^2(X) \right]^{1/2}. \tag{13}$$

Figure 2A plots, for both $\sigma_x = 0.19$ (low volatility) and $\sigma_x = 0.55$ (high volatility), the average subjective valuation $v(X)$, as well as its one-standard-deviation bounds $v(X) \pm \sigma(X)$.

[**Place Figures 2A and 2B about here**]

Figure 2A leads to several observations. First, consistent with prospect theory (Kahneman and Tversky, 1979), the lack of discriminability among outliers generates diminishing sensitivity: the marginal utility $v'(X)$ decreases as X becomes very large. At the same time, because very small positive values are also outliers under a lognormal distribution, $v'(X)$ decreases as X gets very small. Second, diminishing sensitivity is more pronounced when stimulus volatility is lower. This is driven by the fact that when subjects are exposed to a narrower range of stimulus values, they perceive a wider range of stimuli to be outliers, which makes the lack of discriminability more severe. Third, both the shape of the value function and the stochasticity in perceived value arise from noisy encoding: for very large values of X, low discriminability leads to both lower marginal utility $v'(X)$ and higher randomness in utility, $\sigma(X)$. Finally, the value of X for which $v(X)$ attains its greatest slope, which typically corresponds to the "reference point" in prospect theory, arises endogenously in our framework. Here, it corresponds to the most frequently occurring stimulus value, and thus has the highest degree of local discriminability.

As mentioned in the previous paragraph, the lack of discriminability among outliers in both tails of the lognormal distribution generates diminishing sensitivity for both very small and very large payoffs. The hump shape of the lognormal stimulus distribution therefore generates an *S*-shaped value function, which is distinct from the globally concave value function for positive payoffs that is often assumed in prospect theory. Interestingly, efficient coding can also generate this more familiar value function when the *DM* can discriminate very well between small payoffs, but has difficulty discriminating between large payoffs. What type of stimulus distribution would lead to this particular pattern of discriminability?

A monotonically decreasing stimulus distribution would produce exactly this. Specifically, when the *DM* faces a distribution in which small payoffs are more probable than large payoffs, then it is optimal for the *DM* to discriminate more precisely between these small numbers at the expense of discriminating precisely between large numbers.[11] In Figure 2B, we plot an example of a monotonically decreasing stimulus distribution—a gamma distribution—and the implied value function. This value function is indeed concave, and hence

[11] The likelihood functions that are generated by efficient coding with a monotonically decreasing prior will exhibit a property similar to "scalar variability," where the likelihood function becomes more dispersed as the stimulus magnitude increases. Evidence for this property is commonly found in experiments on numerical cognition (Deheane, 2011). Moreover, the likelihood functions that are endogenously generated by our model of efficient coding with a monotonically decreasing prior also resemble the logarithmic encoding function that is assumed in KLW. Interestingly, there is evidence that a monotonically decreasing prior is a good approximation for the distribution of naturally occurring numbers (Dehaene and Mehler, 1992). Therefore, our model of efficient coding, when combined with the distribution of naturally occurring numbers, can be seen as providing a plausible microfoundation for the logarithmic encoding function that is assumed in KLW.

provides another example of how the shape of the value function is malleable yet is tied closely to the underlying stimulus distribution.

We note that the implied value function in (12) is general in the sense that it can take any stimulus distribution as an input. Later in the paper when we present our lab experiment, we will vary the stimulus distribution and test its impact on choice. Because our experiment focuses on manipulating the volatility of the stimulus distribution, we retain the lognormality assumption for the rest of the model section.[12]

II.5. Probability of risk taking

The average subjective valuation in (12) can be used to make predictions about risk taking. In particular, given values of X and C, one can compute whether the average subjective value of the risky lottery (that is, averaged over all realized values of the noisy signal R_x) is larger than the average subjective value of the certain option. However, by first computing the average subjective value for each payoff, this strips the model of its rich predictions about stochastic choice.

We can recover these richer predictions by directly computing the probability of choosing the risky option. To do so, first recall that, conditional on X and C, the noisy signals R_x and R_c are drawn from the probability density functions $p(R_x \mid X)$ and $p(R_c \mid C)$. For a given realized value of (R_x, R_c), the *DM* then chooses between the risky lottery and the certain option based on equations (9), (10), and (11). As a result, when holding X, C, and the stimulus distributions fixed, we can compute the probability of risk taking—that is, the probability of choosing the risky lottery—over many realizations of R_x and R_c,

$$\mathbb{P}\text{rob}(risk\ \ taking \mid X, C) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \mathbf{1}_{\{p\cdot\mathbb{E}[\tilde{X}|R_x] > \mathbb{E}[\tilde{C}|R_c]\}}\, p(R_x \mid X)\, p(R_c \mid C)\, dR_x dR_c. \tag{14}$$

[12] One could also examine a stimulus distribution that is monotonically decreasing, and then vary the volatility of this distribution. However, such a manipulation would also strongly affect the mean and skewness of the distribution, which can lead to confounds in testing the effect of efficient coding on choice behavior.

To understand the determinants of the probability of risk taking, Figure 3 plots this probability against the natural logarithm of X over C, $\ell n(X/C)$, for different volatility levels of the stimulus distributions: $\sigma_x = \sigma_c = 0.4$, 0.8, and 1.5. Specifically, for each volatility level, we set C to $\exp(\mu_c + \frac{1}{2}\sigma_c^2)$ while varying the value of X.

[**Place Figure 3 about here**]

Naturally, a higher ratio of X over C increases the attractiveness of the risky lottery and hence increases the probability of risk taking. Notice that, under expected utility theory and with no background wealth, the probability of risk taking should be a step function of $\ell n(X/C)$ with a single step at $\ell n(X/C) = \ell n[U^{-1}((U(C) - (1-p)U(0))/p)/C]$. However, in our model, the probability of risk taking has an S-shaped relationship with $\ell n(X/C)$. Furthermore, the overall slope of this function is *negatively* related to volatility of the lognormal stimulus distribution. That is, risk taking is more sensitive to payoff values in the low volatility condition, compared to the high volatility condition. Intuitively, lower stimulus volatility reduces the range of the stimulus values that subjects are adapted to, and hence increases the encoding accuracy and the discriminability among the stimulus values within this narrower range.[13] This is the main prediction that we test in our experiments.

II.6. Dynamic extension

The model described above is static and thus makes no explicit predictions about how the speed with which the *DM* adapts to new environments affects risk taking. In order to better understand the dynamics of adaptation, we conclude this section by considering a simple

---

[13] More generally, (14) implies that the probability of risk taking is a two-dimensional function of X and C. In the Appendix, Figure B1 plots this probability for two different volatility levels of the stimulus distribution: $\sigma_x = \sigma_c = 0.19$ (low volatility) and $\sigma_x = \sigma_c = 0.55$ (high volatility). Figure B1 makes it obvious that $\ell n(X/C)$ is not a sufficient statistic of the risk taking probability. Instead, X and C jointly affect this probability. For instance, with $\sigma_x = \sigma_c = 0.55$, $\mu_x = 3.05$, and $\mu_c = 2.35$, setting X to = 24.6 and setting C to = 12.2 gives $X/C = 2.01$ and a risk-taking probability of 77.7%. On the other hand, setting X to = 38.7 and setting C to = 19.2 gives the same ratio of $X/C = 2.01$ but a lower risk-taking probability of 74.3%. Thus, our model does not predict that risk taking is scale-invariant, but rather, it depends on levels of both X and C. We return to this point in both the discussion section (Section V) and Appendix B, after presenting our main results.

extension of the model that yields explicit predictions about adaptation, perception, and risk taking.

Suppose that, at time *t*, the *DM* uses the past *N* outcomes of $X, \{X_{t-i}\}_{i=1}^{N}$, to form a sample distribution of *X*, and similarly, uses the past *N* outcomes of $C, \{C_{t-i}\}_{i=1}^{N}$, to form a sample distribution of *C*. We call parameter *N* the length of the lookback window, which allows for sample variation of the *X* and *C* distributions. Before observing X_t and C_t at time *t*, the subject is assumed to hold the following probability density functions as prior beliefs:

$$
\begin{aligned}
p(X;N) &= \frac{1}{N}\left(\sum_{i=1}^{N}\delta(X - X_{t-i})\right),\\
p(C;N) &= \frac{1}{N}\left(\sum_{i=1}^{N}\delta(C - C_{t-i})\right),
\end{aligned}
\tag{15}
$$

where $\delta(\cdot)$ is a standard Dirac delta function. The corresponding cumulative density functions are

$$
F(X;N) = \frac{1}{N}\left(\sum_{i=1}^{N}\mathbf{1}_{X \geq X_{t-i}}\right), \qquad F(C;N) = \frac{1}{N}\left(\sum_{i=1}^{N}\mathbf{1}_{C \geq C_{t-i}}\right). \tag{16}
$$

Equation (5) from Proposition 1 now implies that the likelihood functions become[14]

$$
\begin{aligned}
L(X \mid R_x, N) = p(R_x \mid X, N) &= \frac{1}{\sqrt{2\pi}\cdot\sigma}\exp\left(-\frac{(F(X;N) - R_x)^2}{2\sigma^2}\right),\\
L(C \mid R_c, N) = p(R_c \mid C, N) &= \frac{1}{\sqrt{2\pi}\cdot\sigma}\exp\left(-\frac{(F(C;N) - R_c)^2}{2\sigma^2}\right).
\end{aligned}
\tag{17}
$$

The prior beliefs in (15) and the likelihood functions in (17) jointly imply that, estimates of *X* and *C* conditional on R_x and R_c are now given by

$$
\mathbb{E}[\tilde{X} \mid R_x] = \frac{\sum_{i=1}^{N} X_{t-i}\cdot\exp\left(-(F(X_{t-i};N) - R_x)^2/(2\sigma^2)\right)}{\sum_{i=1}^{N}\exp\left(-(F(X_{t-i};N) - R_x)^2/(2\sigma^2)\right)} \tag{18}
$$

and

[14] Suppose that X_i and X_j ($X_i < X_j$) are two realized outcomes of *X* in the past *N* trials, and that there are no other realized value of *X* between X_i and X_j. Then, all values of *X* in the range $[X_i, X_j)$ map into an identical conditional probability density function of R_x. That is, $p(R_x|X_a) \equiv p(R_x|X_b)$ for any X_a and X_b in $[X_i, X_j)$. In other words, the *DM* cannot discriminate between X_a and X_b.

$$\mathbb{E}[\tilde{C} \mid R_c] = \frac{\sum_{i=1}^{N} C_{t-i} \cdot \exp\left(-(F(C_{t-i};N) - R_c)^2 / (2\sigma^2)\right)}{\sum_{i=1}^{N} \exp\left(-(F(C_{t-i};N) - R_c)^2 / (2\sigma^2)\right)}. \tag{19}$$

As before, the *DM* chooses the risky lottery over the certain option if and only if $p \cdot \mathbb{E}[\tilde{X} \mid R_x]$ is greater than $\mathbb{E}[\tilde{C} \mid R_c]$.

It is worth noting that, when the length of the lookback window, parametrized by *N*, tends to infinity, the sample distribution in (15) will converge to the true stimulus distribution. It follows that in an environment with a stable stimulus distribution and when the *DM* has a long look back window, we can approximate the *DM*'s sample distribution with the true distribution. In our experimental design, we therefore vary the stimulus distribution at a relatively low frequency in order to induce stability of prior beliefs.

We discuss the predictions of this dynamic model in Section III.3, in order to compare them directly with our experimental data.

III. An Experimental Test

In this section, we provide an experimental test of our model. Our experiment is designed specifically to test whether risk taking varies with the payoff distribution that subjects encounter.

III.1. Design

On each trial in the experiment, subjects choose between a risky lottery and a certain option. The risky lottery delivers a positive payoff *X* with probability *p*, and zero otherwise. The certain option delivers a positive payoff *C* with certainty. The experiment consists of eight blocks, with sixty trials in each block. Each subject therefore completes a total of four hundred eighty trials, which we index by $t = 1, 2, \ldots, 480$. At the end of the experiment, subjects are paid according to their decision on one randomly selected trial.

We experimentally manipulated the distribution from which payoffs in the choice set are drawn. On each trial, the values of *X* and *C* were jointly drawn from a lognormal distribution

$$\begin{pmatrix} \ell n\, X \\ \ell n\, C \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \mu_x \\ \mu_c \end{pmatrix}, \begin{pmatrix} \sigma^2 & \rho\sigma^2 \\ \rho\sigma^2 & \sigma^2 \end{pmatrix} \right). \tag{20}$$

We set the mean values to $\mu_x = 3.05$ and $\mu_c = 2.35$, so that on average, the risky lottery offers a higher expected value than the certain option. Our treatment variable is the standard deviation, which we varied across two conditions: high volatility and low volatility. In the high volatility condition, we set $\sigma = 0.55$, and in the low volatility condition, we set $\sigma = 0.19$. The first block of the experiment was a high volatility block, and the blocks alternated deterministically, so that the experimented ended with a low volatility block (Figure 4). We set the correlation between $\ell n(X)$ and $\ell n(C)$ at $\rho = 0.5$. Although this positive correlation is not part of the model we developed in the previous section, it helped to reduce the number of trivial choice sets where $X < C$ (and as a result, the certain option stochastically dominates the risky lottery). The values of X and C were drawn from their associated distribution (high volatility or low volatility) at the subject-trial level, and thus each subject faced a unique path of payoffs during the experiment.

[**Place Figure 4 about here**]

For all trials, we set the probability that the risky lottery paid X to $p = 0.59$. Following KLW, we chose this design feature for two reasons. First, we used a "non-round" number so that subjects could not easily compute the expected value of the risky lottery—which was more likely to happen if we used, for example, $p = 0.5$ or $p = 0.6$. Second, even though our model assumes that the subject does not encode the probability p with noise, in reality, this variable is also likely to be encoded with noise. By presenting the same value of 0.59 on each trial, this increased the plausibility of our simplifying assumption that subjects precisely encoded this particular probability value. Later in the paper, we conduct an additional experiment to directly test the noisy encoding of payoffs, without appealing to any assumptions about probability encoding.

Before the experiment began, subjects were told that they would be asked to choose between two lotteries on each of four hundred eighty trials and these trials would be separated into eight parts. However, subjects were not given any information about the distribution of X or C, nor were they told that these distributions changed across blocks. We chose not to provide subjects with this information because it allowed us to study the adaptation process more generally, without imposing any specific structure on the subject's prior beliefs at the beginning of the experiment. The exact instructions that were given to subjects before the experiment are provided in the Appendix.

III.2. Experimental procedures

We recruited $N = 34$ subjects for this experiment, which was conducted across three sessions at Caltech and USC. Before starting the experiment, subjects went through a set of ten practice trials to become familiar with the task and the software. Figure 4 provides an example trial from the experiment, in which the risky lottery is presented on the left as a colored bar chart, and the value X is displayed at the bottom next to its associated probability of 0.59. The certain option is presented on the right side of the screen. On each trial, subjects were instructed to select the left or right option by pressing one of two keys. The location of the risky lottery was randomized across subjects and trials, and subjects had unlimited time to make their decision on each trial. At the end of each block of sixty trials, a progress screen appeared, which reported how many of the eight blocks the subject had completed.

At the end of the eighth block, the computer randomly selected one of the four hundred eighty trials from the experiment. If the subject chose the risky lottery on this trial, a random number generator determined whether the subject received the payoff of $\$X$ or the payoff of \$0, according to the probabilities associated with these payoffs. If the subject chose the certain option, she received the amount of $\$C$. In addition to the earnings from this randomly selected trial, each subject received a \$7 show-up fee. The average earning, including the show-up fee, was \$25.89.

III.3. Experimental results

III.3.A. Treatment effects

Subjects chose the risky lottery on 40.5% of trials in the low volatility condition and on 42.7% of trials in the high volatility condition. One subject did not exhibit any variation in risk taking in the low volatility condition (choosing the certain option on each trial), and we exclude this subject from all subsequent analyses.

[**Place Figure 5 about here**]

Figure 5 plots the proportion of trials on which subjects chose the risky lottery, as a function of the natural logarithm of X over C, $\ell n(X/C)$. Recall that the probability p stays constant across all trials, and thus $\ell n(X/C)$ provides a good—though insufficient—statistic that summarizes the attractiveness of the risky lottery relative to the certain option. The figure shows that risk taking increases in $\ell n(X/C)$ in both conditions, which provides a basic consistency check on the data. One can also see that the slope of the curve in the low volatility condition appears to be steeper than that in the high volatility condition. This is consistent with a basic prediction of our model: when the stimulus distribution becomes more concentrated, choice sensitivity increases.

To conduct formal empirical tests, we run regressions where the dependent variable takes the value of one (zero) if the subject chose the risky lottery (certain option) on trial t. We pool all 15,840 trials across subjects and conditions, and run a logistic regression. The results in Column (1) of Table 1A show that the regression coefficient on $\ell n(X/C)$, which provides a measure of the sensitivity of risk taking in the low volatility condition, is positive and strongly significant. *high* is a dummy variable that takes the value of one if the trial is in the high volatility condition, and zero otherwise. The coefficient of interest is on the interaction term $\ell n(X/C)\times high$, which is significantly negative, indicating that risk taking becomes less sensitive to $\ell n(X/C)$ in the high volatility condition. This provides formal support for a difference in choice sensitivity between the high and low volatility conditions.

[**Place Table 1 about here**]

Our model predicts that this difference in choice sensitivity stems from different *perceptions* of X and C across conditions. However, the distributions of X and C themselves vary across conditions, and thus the difference we detect may simply be driven by a different response to extreme values of $\ell n(X/C)$. In particular, the range of $\ell n(X/C)$ in the low volatility condition is (0.02, 1.36), but in the high volatility condition, there are many trials for which $\ell n(X/C)$ falls outside this range.

To help address this concern, we restrict our regression to trials with similar levels of $\ell n(X/C)$ across the two volatility conditions. We re-estimate the regression in Column (1) using only trials for which $0.02 < \ell n(X/C) < 1.36$; this represents the set of values of $\ell n(X/C)$ which appear in *both* conditions of our experiment. Column (2) shows that our results are quite similar on this subset of data. Given this, it is unlikely that differences in the current choice set drive the full effect; however, we note that this is not a perfect control because the distribution of $\ell n(X/C)$ still differs within this restricted domain. Columns (3) and (4) show that our main result holds within each half of a block. Column (3) is estimated using data from the first half of each block (the first thirty trials), while column (4) uses data from the second half (the last thirty trials), and the main result holds in both of these subsets of the data.

Column (5) shows our results hold even within the first ten trials of each block. In fact, if we restrict to the *first trial* of each block (Column (6)), we find that the coefficient on the interaction term remains significantly negative (p-value of 0.046). This result is potentially concerning, because it is consistent with a theory in which the subject adapts to the new distribution instantly. However, another more plausible explanation is based on the perception of outliers. On the first trial of each block of the high volatility condition, the subject has just faced sixty choice sets from the low volatility condition, and presumably has adapted to a narrow range of $\ell n(X/C)$. On the first trial of the high volatility block, there is a good chance that $\ell n(X/C)$ falls outside this range (or near the extremes of this range), and thus this value will be perceived as an outlier.[15] Because the subject's coding resources do not adjust immediately following the low volatility block, the perception of the outlier will be noisy.[16]

[15] In our experiment, there is a 21% chance that $\ell n(X/C)$ falls outside the range (0.02, 1.36) on the first trial of a high volatility block.

[16] We note that in each specification in Table 1A, the coefficient on the *high* dummy variable is significantly positive. This is likely because $\ell n(X/C)$ is not a sufficient statistic for risk taking and may contribute to model misspecification. Table 1B presents regression results where we separately enter the X and C regressors, as well as

III.3.B. Adaptation and dynamics

The above logic can be extended more generally to all trials in the high volatility block. Consider a value of $\ell n(X/C)$ that is extreme within the context of the low volatility block, but not within the high volatility block. If the subject encounters this value *early* in a high volatility block, it will be perceived as an outlier, but the presentation of this value should lead to a subsequent adjustment of coding resources through adaptation. When a value of $\ell n(X/C)$ that was perceived as extreme early in a high volatility block is presented a few times, then later in the high volatility block, it will no longer be perceived as an outlier because coding resources have been allocated to a wider range of payoffs. Therefore, risk taking should become more sensitive to extreme values over the course of the high volatility block. Critically, because the total amount of coding resources is conserved, this reallocation of resources must come at the expense of resources used for perception of intermediate values. It follows that risk taking should become *less* sensitive to intermediate values over the course of the high volatility block. Taken together, this implies that the net reallocation of resources from intermediate to outlier payoffs is positive.

To test this prediction, we examine how risk taking varies over time among outliers and intermediate values in the high volatility block. We define an outlier as a value of $\ell n(X/C)$ that is more than three standard deviations from the mean. The standard deviation and mean are computed at the subject-block level, using the sixty trials from the immediately preceding low volatility block (we restrict analysis to blocks 3, 5 and 7, so that each of these high volatility block can be matched to an immediately preceding low volatility block). Under this definition, 30.4% of trials in the high volatility block are considered outliers.

We define a dummy called *outlier* that takes the value of one if the value of $\ell n(X/C)$ is an outlier, and zero otherwise. We also define a dummy called *second* which takes the value of one if the trial takes place in the second half of the block (i.e., in the last thirty trials). We then estimate the following logistic regression, using only data from the high volatility blocks:

their associated interaction terms. We see that the coefficient on the *high* dummy variable is no longer significant in these regressions; moreover, consistent with the result in Table 1A, we also find that risk taking becomes less sensitive to both X and C in the high volatility condition.

$$\begin{aligned}\text{risky}_t &= \alpha + \beta_1 \cdot \ell\text{n}\left(\frac{X_t}{C_t}\right) + \beta_2 \cdot \ell\text{n}\left(\frac{X_t}{C_t}\right) \times \mathit{outlier} \\ &+ \beta_3 \cdot \ell\text{n}\left(\frac{X_t}{C_t}\right) \times \mathit{second} + \beta_4 \cdot \ell\text{n}\left(\frac{X_t}{C_t}\right) \times \mathit{second} \times \mathit{outlier} + \varepsilon_t.\end{aligned} \tag{21}$$

In Column (1) of Table 2, the regression coefficient β_1 on $\ell\text{n}(X/C)$ provides the sensitivity of risk taking among intermediate value (non-outlier) trials in the first half of the high volatility block. As expected, this coefficient is significantly positive. Furthermore, the coefficient β_2 on the interaction $\ell\text{n}(X/C)\times \mathit{outlier}$ is significantly negative. This indicates that in the first half of the block, shortly after experiencing sixty low volatility trials, subjects are less sensitive to outliers than to intermediate values.

[**Place Table 2 about here**]

The next two rows allow these coefficients to differ in the second half of the block. The coefficient of interest is β_4, which measures the change over time in choice sensitivity to outlier payoffs, relative to intermediate values. As described above, we expect this coefficient to be positive because increased exposure to outliers during the first half of the high volatility block should trigger adjustment of coding resources away from intermediate values and towards these extreme values in the second half of the block. In Column (1) we find that the estimated coefficient is indeed positive, although it is not significantly different from zero. One potential reason for this insignificance is that adaptation takes place relatively quickly, and thus subjects hold similar priors in the first and last thirty trials of the block; this, in turn, would make it difficult to detect a difference in behavior between the two halves of the block.

To further investigate this, Columns (2) – (4) use more restricted subsets of the data, in which we use only the first and last M trials of the high volatility block, for M = 20, 10, and 5, respectively. We see that the results become stronger after restricting to more extreme subsets of the data. In particular, after restricting to the first and last five trials of each block in Column (4), we find that β_4 is significantly positive ($p = 0.03$), indicating that there is a net reallocation of resources towards outlier payoffs. This result also holds at the 10% significance level when restricting to the first and last 10 trials of each block (Column (3)). However, it becomes insignificant as we relax the sample restriction to the first and last 20 trials of each block (Column (2)). This pattern provides a clue about the speed of adaptation: it is likely to be shorter than 20 trials, but longer than 10 trials.

To provide a more direct test of the speed of adaptation, we appeal to the dynamic model presented in Section II.6. Recall that in this model extension, the *DM* uses the past N realized values of X and C, and each of these values receives the same weight, $1/N$, in constructing the prior distribution.[17] To generate predictions for our experimental design, we set the lookback window parameter N to 10. We then use the model to generate choice data for a sample of one thousand "pseudo subjects," each characterized by $\sigma = 0.1$, and each facing the same design that our experimental subjects faced.

[**Place Figure 6 about here**]

Using the model generated choice data, we repeat the logistic regression estimation in (21), and again restrict to the first and last 5 trials of each block. The top panel of Figure 6 summarizes the regression results by plotting the average sensitivity to each of the four types of trials: 1st half outliers, 2nd half outliers, 1st half intermediate values, and 2nd half intermediate values. These model generated point estimates confirm the qualitative model predictions developed at the beginning of this section, whereby the sensitivity to outliers increases over time while the sensitivity to intermediate values decreases over time. To provide a direct comparison between model and experimental data, we plot the same chart in the bottom panel of Figure 6 using experimental data.

Our simple model of adaptation offers a good, though imperfect, match with the experimental data when $N = 10$. One difference is that the model predicts a higher choice sensitivity to $\ell n(X/C)$ in all four categories compared to the experimental data. There are at least

[17] We note that our model of adaptation assumes that the *DM* has access to the true values of X and C when forming prior beliefs, yet she can only imperfectly observe the values X and C when making her decision on each trial. Interestingly, Robson and Whitehead (2018) show that with a sufficiently long time horizon, the *DM can* adapt to the true distribution without perfectly observing the stimulus on each trial.

two possible explanations for this difference. First, we assume a particular value of σ (= 0.1) when generating the model-implied choice data. Given that σ reflects the amount of noise in the likelihood functions, a larger value of σ gives rise to noisier estimates of perception, and thus lower sensitivities to payoffs. Second, in addition to assuming a particular value of σ in the model generated choice data, we also assume that each pseudo subject is characterized by the *same* value of σ. However, in our sample of experimental subjects, there is likely to be heterogeneity in σ across subjects. Together, these two factors may contribute to the lower choice sensitivity to payoffs observed in the experimental data (compared to the model predictions). In the next section, we investigate the extent of heterogeneity across the experimental subjects in our sample.

III.3.C. Heterogeneity across subjects

[**Place Figure 7 about here**]

The data presented in Figure 5 are pooled across subjects, and therefore mask any heterogeneity in the change in risk taking across volatility conditions. To investigate the extent of this heterogeneity, for each subject and condition, we run the following logistic regression:

$$\text{risky}_t = \alpha + \beta \cdot \ell\text{n}\left(\frac{X_t}{C_t}\right) + \varepsilon_t. \tag{22}$$

We record the estimates $\hat{\beta}$ for each subject in the high and low conditions, and plot these against each other in Figure 7. We see there is substantial heterogeneity across subjects in the sensitivity to $\ell\text{n}(X/C)$. Moreover, subjects who are more sensitive in the high condition are also more sensitive in the low condition, indicating that σ is relatively stable across conditions, within subjects. Most importantly, we see that, for a majority of subjects, the data lie above the blue forty-five degree line. This confirms that the greater choice sensitivity in the low volatility condition is present within most of our subjects.

III.3.D. Assumptions about noiseless encoding process

All of the results in our theoretical model are driven by the noisy encoding of *X* and *C*. In particular, we make two simplifying assumptions: *i*) there is no noise in encoding the probability *p*; and *ii*) there is no noise in computing the product of *p* and $\mathbb{E}[\tilde{X} \mid R_x]$ (which is used as the *DM*'s estimate of the expected value of the risky lottery). In reality, there is likely to be noise in both of these processes, which could potentially be responsible for some of the above experimental results. However, because the noisy encoding of payoffs is sufficient to generate our main theoretical predictions, we should still find evidence that the perception of *X* depends on the recent stimulus distribution, even when there is no need to perceive *p*. To provide a sharper test of the effect of noisy encoding of payoffs, we run a follow-up experiment in which the subject still needs to perceive *X*, but does not need to perceive *p* or integrate probabilities with payoffs.

IV. Riskless Choice Experiment

IV.1. Experimental design

The design of our second experiment is informed by decades of work from the literature on perception of numerical quantities (Moyer and Landauer, 1967). Our experimental design builds on that of Dehaene, Dupoux, and Mehler (1990), who on each trial of their experiment, present subjects with an Arabic numeral between 31 and 99. The subject's task is simply to classify whether the Arabic numeral presented on the screen is larger or smaller than the reference level of 65. Their main result is that as the stimulus numeral gets closer to the reference level, response times increase and classification accuracy decreases. These results are consistent with the noisy encoding of Arabic numerals, which lies at the foundation of our model of risk taking.

One notable feature of the Dehaene, Dupoux, and Mehler (1990) experiment is that the stimulus distribution is held constant throughout the experiment. Here, we exogenously vary the stimulus distribution, in much the same way that we varied the distribution of monetary amounts in our previous experiment. We have a high volatility distribution (uniform over integers in the range [31, 99]) and a low volatility distribution (uniform over integers in the range [51, 79]). Subjects are incentivized to correctly classify whether each Arabic numeral, which we denote by X, is larger or smaller than 65, over sixteen blocks of trials. The blocks alternate between the high volatility condition and the low volatility condition. Each block consists of eighty trials, for a total of 1,280 trials per subject (Figure 8).

[**Place Figure 8 about here**]

We pay subjects as a function of both their accuracy and their speed. In addition to a $7 participation fee, subjects earn a payoff of $(20×*accuracy* – 10×*avgseconds*), where *accuracy* is the percentage of correctly classified trials, and *avgseconds* is the average response time (in seconds) across all trials in the experiment. In this design, the subject still needs to perceive the value X, but there are no probabilities to encode, nor any need to integrate probabilities with payoffs. Therefore, this design provides a clean setting in which we can test whether the perception of an Arabic numeral, X, depends on the recently observed stimulus distribution.

IV.2. Experimental procedures

We recruited an additional $N = 13$ subjects from Caltech for this experiment. Before the first block, subjects went through a set of ten practice trials to become familiar with the task. On each trial, the stimulus numeral was displayed in white font against a black background, on the center of the screen (Figure 8). Subjects were instructed to press one of two keys to indicate whether the stimulus was smaller or greater than 65. After responding on each trial, a white fixation cross appeared for 500 milliseconds, followed by the stimulus from the next trial. At the end of each block of eighty trials, a progress screen appeared, which reported how many of the sixteen blocks remained. The progress screen was self-paced, and subjects were given the opportunity to take a break during this screen. The average earning, including the show-up fee, was $20.58.

IV.3. Experimental results

Subjects accurately classified the stimuli on 90.4% of trials with an average response time of 0.45 seconds. Figure 9 shows the proportion of trials that subjects classified the stimulus as greater than 65, for each value of X. If subjects had accurately classified all stimuli, the figure would generate a step function, with a single step at $X = 65$. Instead, the figure replicates results from several previous experiments in the literature, which show that errors decrease in the distance between the two numbers under comparison (Moyer and Landauer, 1967; Dehaene, Dupoux, and Mehler, 1990). To be clear, while it is unsurprising that subjects make errors in general, the more important result is that the error rate is correlated with the distance between the stimulus number and the reference level of 65. It is also worth noting that the average subject from Caltech has very high mathematical aptitude, and thus the error rates reported here are likely to be close to a lower bound for the error rates among other samples.[18]

[**Place Figure 9 about here**]

Turning to a comparison of our two experimental conditions, we find that subjects correctly classify stimuli on 91.4% of trials in the high volatility condition, and on 89.4% of trials in the low volatility condition. A more informative statistic is the difference in accuracy between conditions, when restricting to stimuli that are common to both conditions: $51 \leq X \leq 79$. This controls for the fact that, on average, trials in the high volatility condition are "easier," in the sense that the average distance to the reference level is greater than in the low volatility condition. We find that accuracy among these trials in the high volatility condition is 86.5%, which is significantly lower than the 89.4% accuracy in the low volatility condition (p-value $=$ 0.004). This is consistent with the efficient coding hypothesis: in the low volatility condition, subjects adapt and devote more coding resources to the concentrated range $51 \leq X \leq 79$. In the high volatility condition, subjects need to "spread" these coding resources over a wider range,

[18] There is evidence that accurate perception of *non-symbolic* representations of numbers (e.g., a visual array of dots) is positively correlated with mathematical aptitude (Halberda, Mazzocco, and Feigenson, 2008), though it is unclear whether this correlation extends to tasks like ours that use symbolic numerical representations.

which leads to increased noise when encoding stimuli in the concentrated range (relative to the low volatility condition).

A sharper test of the efficient coding hypotheses is to compare the slopes in Figure 9. As in our previous experiment, we expect a steeper slope in the low volatility condition. The figure provides suggestive visual evidence for a difference in slopes, and to formally test this, we run a series of logistic regressions. The dependent variable in our logistic regression takes on the value of one if the subject classified X as above 65, and zero otherwise. Column (1) of Table 3 shows that the coefficient on $\ell n(X/65)$ is significantly positive, which indicates that subjects' propensity to classify X as greater than 65 is increasing in $\ell n(X/65)$. More importantly, we find that the coefficient on the interaction term, $\ell n(X/65)\times high$, is significantly negative, indicating that choices are noisier on trials in the high volatility condition.

[**Place Table 3 about here**]

To control for the difference in distributions of X, we re-estimate the regression using data only in the range $51 \leq X \leq 79$. When restricting to this range, the distribution of X on the current trial is the same across conditions, and the only difference is the distribution of previously encountered stimuli. Column (2) provides these estimation results, and we find the slope remains steeper in the low volatility condition (though the difference in slopes is smaller compared to the estimates using the full sample in Column (1)).

One assumption we make in interpreting the results in these first two columns, is that there is no "external stimulus" noise: the stimulus number is displayed clearly on the screen and the font is easy to read (as opposed to, e.g., fuzzy text). We assume that the noise that corrupts the mental representation of the stimulus is based on the internal noise in the subject's nervous system. Nonetheless, it is plausible that comparing 59 with 65 may be easier than comparing 60 with 65, not because of the distance, but because the first digits are visually distinct. To address this, we re-estimated the regression in Column (1) using only trials for which the first digit differs from the first digit of the reference level: $\{X < 60, X \geq 70\}$. Column (3) shows that the slope in the low volatility condition remains steeper, indicating that such a “first-digit” effect cannot explain the full extent of the shift in slope.

To summarize this experiment, we find that the accuracy of classifying an Arabic numeral is affected by *i*) the distance to a reference level, and *ii*) the distribution of previously encountered stimuli. The latter result provides useful evidence supporting a basic assumption of our model of risky choice. Specifically, in an experimental task where there is no need to encode probabilities or integrate with payoffs, we find that choice sensitivity depends on the distribution of previously encountered stimuli.

V. Discussion

V.1. Definition of efficiency and the encoding function

In our model, we assume that the encoding process is efficient in the sense that mutual information between the stimulus value and its noisy representation is maximized. This definition of efficiency is taken from recent work in theoretical neuroscience (Wei and Stocker, 2015), and we choose to maintain this assumption in our model because it generates predictions that closely match choice data in both perceptual and economic experiments (Polania, Woodford, and Ruff, 2019). This assumption can also be justified on normative grounds if the *DM*’s objective is to minimize the mean squared error of her estimate of the true stimulus value.

However, our assumed definition of efficiency is not necessarily the only plausible one. For example, an alternative definition is that the *DM* encodes payoffs in the choice set to maximize the expected financial gain on each trial. This type of objective is closer in spirit to models of rational inattention (Sims, 2011; Woodford, 2012), in which the *DM* chooses a signal structure that maximizes expected utility. Interestingly, KLW show that, within a specific class of logarithmic encoding functions, maximizing financial gain implies a very similar prediction to the one we obtain in our model, namely, that the sensitivity of risk taking to payoffs will decrease in the variance of the prior distribution.[19]

We emphasize that one difference between our model and KLW is that the *shape* of the encoding function in our model depends heavily on the shape of the prior distribution. As described in Section II, for a monotonically decreasing prior, both our model and KLW predict that the *DM* encodes smaller stimulus values more accurately. However, for priors in which there is a greater chance of observing a large stimulus value compared to a small one, KLW continue to predict that smaller stimulus values are encoded more accurately, while our model now predicts that larger stimulus values are encoded more accurately. This pattern stands in stark contrast to Weber's law—which states that larger values are discriminated less accurately than smaller values. Yet, recent experimental work from Polania, Woodford, and Ruff (2019) provides evidence that there is indeed less variability in the subjective ratings of *high* value food items, compared to low value food items. The authors interpret this empirical pattern as arising from efficient coding and a particular prior distribution in which the expected value of the prior is shifted towards higher value items. They also make the important observation that such an empirical pattern does not invalidate Weber's law, but rather, it suggests that the strength of Weber's law depends on the prior distribution to which the *DM* has adapted.[20]

V.2. Insensitivity to outliers: A comparison with salience theory

The results from our risky choice experiment suggest that risk taking is less sensitive to outlier payoffs compared to those payoffs that subjects encounter more frequently. This result is consistent with a basic prediction of efficient coding, namely that the ability to discriminate between outlier payoffs is weaker than the ability to discriminate between frequently occurring payoffs. This result may seem at odds with alternative theories of decision making in which

[19] When the *DM* holds a prior distribution that is monotonically decreasing, our model will generate a set of encoding functions that resemble the logarithmic encoding functions assumed in KLW. This offers an alternative interpretation for some of the experimental results in KLW. In their experimental design, KLW sample values of X and C from a distribution that induce an approximately monotonically decreasing prior distribution for both X and C (since the distance between successively larger payoffs is increasing in their design). In such an experiment, our model would thus predict that the encoding function should be similar to the one assumed in KLW.

[20] In the Appendix, we provide further analysis of our model's implications for Weber's law and the "scale invariance" of risk taking.

extreme payoffs capture attention, and are then overweighed at the expense of typical payoffs, as in Bordalo, Gennaioli, and Shleifer (2012) and Koszegi and Szeidl (2013). A closer look at these theories, however, suggests that their core mechanism is complementary to the one studied here, rather than contradictory.

To see this, note that in our model of efficient coding, there are two layers of uncertainty. In the first layer, the *DM* is uncertain about the payoff values of *X* and *C* that characterize each lottery—she only draws noisy signals R_x and R_c of *X* and *C*—even after the choice set is presented to her. In the second layer, conditional on her optimal estimates of the payoff values, $\mathbb{E}[\tilde{X} \mid R_x]$ and $\mathbb{E}[\tilde{C} \mid R_c]$, she then faces the more standard source of uncertainty about whether the risky lottery will deliver its upside of *X* (with probability *p*), or its downside of 0 (with probability $1 - p$). We can apply salience theory to both layers of uncertainty, one at a time. We give a detailed treatment of this in Appendix C, but briefly report the intuition here.

When applying salience theory to the first layer uncertainty, we can imagine there are a continuum of states, and each state is characterized by a different pair of draws from the joint lognormal distribution of *X* and *C* specified in equation (20) for our experimental setup. Note that in this specification of the state space, the *DM* is choosing between the priors that generate the draws of *X* and *C* (which then characterize the choice set that subjects actually face in our experiment.) In this case, we show in Appendix C that salience theory predicts that large values of *X* are overweighted, and hence distort choices. Yet this intuition does not apply in our experiment because subjects are not choosing between these two lognormal priors.

If instead we apply salience theory to the choice set that subjects actually face, we find that the theory does a good job explaining our data on most dimensions (Figure C2). The one dimension on which salience cannot explain our experimental data is the dependence of risk taking on past choice sets. Of course, a natural way to incorporate this dependence is to allow the salience function to depend on past payoffs. This is in the spirit of more recent versions of salience theory, in which memory of past experiences shapes perception of the current choice set (Bordalo, Gennaioli, and Shleifer, 2017).[21]

V.3. Implications for eliciting risk attitudes

At a methodological level, our results have important implications for future experimental work on measuring risk preferences. Most of the current methods for eliciting risk preferences in economics are grounded in the assumption that there is a single and stable utility function that governs choices over the entire experiment (Charness, Gneezy, and Imas, 2013). Efficient coding implies that the subject's utility function is malleable even over the course of a single experimental session, and therefore provides some guidance for improving measurement.

For example, our experimental results indicate that the order in which choice sets are presented has a systematic effect on choice behavior. Controlling for such order effects may therefore improve the precision of preference parameter estimates. Recently, a set of clever experimental procedures has been proposed to efficiently measure risk preferences, in the sense of minimizing the number of questions the experimenter must ask in order to obtain a given level of estimation precision (Imai and Camerer, 2018; Chapman et al., 2018). In these "dynamically optimized sequential experimentation" algorithms, the choice set on the current trial depends on the subject's history of choices, so as to maximize information gain for the experimenter. Because efficient coding implies that the order in which choice sets are presented has a systematic effect on behavior, our data suggest that it may be possible to achieve even more precise parameter estimation by conditioning on the specific *path* of questions presented to subjects.

VI. Conclusion

In this paper, we derive the implications for risk taking when the perception of payoffs is noisy and governed by efficient coding. We show that the *DM*'s value function is malleable, and

[21] The model of Bordalo, Gennaioli, and Shleifer (2017) focuses mainly on applications to riskless choice. However, our paper focuses on risky choices.

its shape fluctuates with the distribution of recent payoffs. In particular, because the *DM* has difficulty discriminating between those payoffs that do not occur frequently, efficient coding generates diminishing sensitivity, which itself can change across environments. Earlier work by Woodford (2012) provides a model of efficient coding that shares some of our predictions, but in that model, imperfect perception is applied to the net gain of a payoff. The model of KLW that we build on here instead assumes that the *DM* encodes the absolute value of a symbolic number. This is a more realistic assumption, as the perceptual system responsible for noisy encoding of numerosity is unlikely to support negative numbers (Feigenson, Dehaene, and Spelke, 2004).[22]

To test our model, we conduct two laboratory experiments in which we find evidence consistent with efficient coding of risky payoffs. Specifically, risk taking becomes more sensitive to those payoffs that appear more frequently in the choice set. Such adaptation takes place relatively quickly, on the order of approximately ten to twenty experimental trials. In our second experiment, where subjects need only classify whether a symbolic number is larger than a reference level, we find that classification accuracy systematically changes with the distribution of recently experienced numbers. This provides evidence supporting our basic model assumption, that payoffs which are encountered more frequently are perceived more precisely.

Finally, to demonstrate how the sensitivity of risk taking evolves over time, we proposed a simple model of adaption. In this model, the *DM* forms a prior distribution by equally weighting the payoff values from the past *N* choice sets. In reality, however, the adaption process is likely to be much more complex. For example, the amount of past data that subjects use to form prior beliefs is itself likely to vary with the rate of environmental change (Behrens et al., 2007) and across different timescales (Zimmermann et al., 2018). Therefore, an important next step is to combine the model of efficient coding that we present here with a more realistic model of the adaption process. This will generate richer predictions about learning and choice dynamics, which can then be tested in future experiments.

[22] As in KLW, negative numbers can be accommodated in our model by first encoding their absolute value, and then multiplying the encoded value by minus one.

References

Barberis, Nicholas, "Thirty Years of Prospect Theory in Economics: A Review and Assessment," *Journal of Economic Perspectives*, 27 (2013), 173–196

Barberis, Nicholas, "Psychology-based Models of Asset Prices and Trading Volume," in *Handback of Behavioral Economics*, (eds. Bernheim, D., DellaVigna, S., Laibson, D.) North Holland, Amsterdam, 2018.

Barlow, Horace, "Possible Principles Underlying the Transformations of Sensory Messages," in *Sensory Communication* (ed. Rosenblith, W.A.) 217–234, MIT Press, 1961.

Behrens, Timothy, Mark Woolrich, Mark Walton, and Matthew Rushworth, "Learning the Value of Information in an Uncertain World,' *Nature Neuroscience*, 10, (2007), 1214-1221.

Bhui, Rahul and Samuel Gershman, "Decision by Sampling Implements Efficient Coding of Psychoeconomic Functions," *Psychological Review,* 125, (2018), 985-1001.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Salience Theory of Choice Under Risk," *Quarterly Journal of Economics*, 127 (2012), 1243–1285.

Carandini, Matteo, and David J. Heeger, "Normalization as a Canonical Neural Computation," *Nature Reviews Neuroscience,* 13 (2012), 51–62.

Casella, George, and Roger L. Berger, *Statistical Inference*, Pacific Grove, CA: Duxbury, 2002.

Chapman, Jonathan, Erik Snowberg, Stephanie Wang, and Colin Camerer, "Loss Attitudes in the U.S. Population: Evidence from Dynamically Optimized Sequential Experimentation (DOSE)", NBER working paper No. 25072, (2018).

Charness, Gary, Uri Gneezy, and Alex Imas, "Experimental Methods: Eliciting Risk Preferences," *Journal of Economic Behavior & Organization*, 87 (2013), 43-51.

Dehaene, Stanislas, "Varieties of Numerical Abilities," *Cognition* 44 (1992), 1–42.

Dehaene, Stanislas, *The Number Sense*, Oxford: Oxford University Press, 2011.

Dehaene, Stanislas, Emmanuel Dupoux, and Jacuqes Mehler, "Is Numerical Comparison Digital? Analogical and Symbolic Effects in Two-Digit Number Comparison," *Journal of Experimental Psychology: Human Perception and Performance*, 16 (1990), 626–641.

Dehaene, Stanislas and Jacque Mehler, "Cross-Linguistic Regularities in the Frequency of Number Words," *Cognition,* 43 (1992), 1–29.

Feigenson, Lisa, Stanislas Dehaene, and Elizabeth Spelke, “Core Systems of Number,” *Trends in Cognitive Science*, 8 (2004), 307–314.

Gabaix, Xavier, and David Laibson, “Myopia and Discounting,” NBER working paper No. 23254, (2017).

Girshick, Ahna R., Michael S. Landy, and Eero P. Simoncelli, “Cardinal Rules: Visual Orientation Perception Reflects Knowledge of Environmental Statistics,” *Nature Neuroscience*, 14 (2011), 926–932.

Halberda, Justin, Michele Mazzocco, and Lisa Feigenson, “Individual Differences in Non-Verbal Number Acuity Correlate with Maths Achievement,” *Nature,* 455 (2008), 665–668.

Imai, Taisuke, and Colin Camerer, “Estimating Time Preferences from Budget Set Choices Using Optimal Adaptive Design,” Working paper, (2018).

Imas, Alex, “The Realization Effect: Risk-Taking after Realized versus Paper Losses,” *American Economic Review*, 106 (2016), 2086–2109.

Kahneman, Daniel, and Amos Tversky, “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, 47 (1979), 263–292.

Khaw, Mel Win, Paul W. Glimcher, and Kenway Louie, “Normalized Value Coding Explains Dynamic Adaptation in the Human Valuation Process,” *Proceedings of the National Academy of Sciences* 114, (2017), 12696–12701.

Khaw, Mel Win, Ziang Li, and Michael Woodford, “Cognitive Imprecision and Small-Stakes Risk Aversion,” NBER working paper No. 24978, (2019).

Koszegi, Botond, and Adam Szeidl, “A Model of Focusing in Economic Choice,” *Quarterly Journal of Economics*, 128 (2013), 53–104.

Kutter, Esther, Jan Bostroem, Christian Elger, Florian Mormann, and Andreas Nieder, “Single Neurons in the Human Brain Encode Numbers,” *Neuron* 100, (2018), 1–9.

Landry, Peter, and Ryan Webb, “Pairwise Normalization: A Neuroeconomic Theory of Multi-Attribute Choice,” SSRN working paper, (2018).

Laughlin, Simon, “A Simple Coding Procedure Enhances a Neuron’s Information Capacity,” *Zeitschrift fur Naturforschung* 36, (1981), 9–10.

Louie, Kenway, and Paul W. Glimcher, “Efficient Coding and the Neural Representation of Value,” *Annals of the New York Academy of Sciences* 1251, (2012), 13–32.

Moyer, Robert S., and Thomas K. Landauer, “Time Required for Judgments of Numerical Inequality,” *Nature* 215, (1967), 1519–1520.

Padoa-Schoppa, Camillo, “Range-Adapting Representation of Economic Value in the Orbitofrontal Cortex,” *Journal of Neuroscience* 29 (2009), 14004-14014.

Payzan-LeNestour, Elise, Bernard Balleine, Tony Berrada, and Joel Pearson, “Variance After-Effects Distort Risk Perception in Humans,” *Current Biology* 26, (2016), 1500–1504.

Payzan-LeNestour, Elise, and Michael Woodford, “‘Outlier Blindness’: Efficient Coding Generates an Inability to Represent Extreme Values,” SSRN working paper, (2018).

Polania, Rafael, Michael Woodford, and Christian Ruff, “Efficient Coding of Subjective Value,” *Nature Neuroscience* 22, (2019), 134–142.

Rangel, Antonio, and John A. Clithero, “Value Normalization in Decision Making: Theory and Evidence,” *Current Opinion in Neurobiology* 22, (2012), 970–981.

Robson, Arthur J., and Lorne A. Whitehead, “Adaptive Hedonic Utility,” Working paper, (2018).

Sims, Christopher A., “Rational Inattention and Monetary Economics,” in *Handback of Monetary Economics*, (eds. Friedman, B. M., Woodford, M.) Amsterdam: Elsevier, 2011.

Steiner Jakub, and Colin Stewart, “Perceiving Prospects Properly,” *American Economic Review* 106, (2016), 1601–1631.

Stewart, Neil, Nick Chater, and Gordon Brown, “Decision by Sampling,” *Cognitive Psychology* 53, (2006), 1–26.

Stocker Alan A., and Eero P. Simoncelli, “Noise Characteristics and Prior Expectations in Human Visual Speed Perception,” *Nature Neuroscience* 9, (2006), 578–585.

Thaler, Richard, and Eric Johnson, “Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice,” *Management Science* 36, (1990), 643–660.

Tobler Philippe N., Christopher D. Fiorillo, and Wolfram Schultz, “Adapting Coding of Reward Value by Dopamine Neurons,” *Science* 307, (2005), 1642–1645.

Tymula, Agnieszka, and Paul Glimcher, “Expected Subjective Value Theory (ESVT): A Representation of Decision Under Risk and Certainty,” SSRN working paper, (2017).

Weber, Martin, and Colin Camerer, “The Disposition Effect in Securities Trading: An Experimental Analysis,” *Journal of Economic Behavior and Organization* 33, (1998), 167–184.

Wei, Xue-Xin, and Alan A. Stocker, “A Bayesian Observer Model Constrained by Efficient Coding Can Explain ‘Anti-Bayesian’ Percepts,” *Nature Neuroscience* 18, (2015), 1509–1517.

Wei, Xue-Xin, and Alan A. Stocker, “Lawful Relation between Perceptual Bias and Discriminability,” *Proceedings of the National Academy of Sciences* 114, (2017), 10244–10249.

Woodford, Michael, “Prospect Theory as Efficient Perceptual Distortion,” *American Economic Review Papers & Proceedings* 102, (2012), 41–46.

Zimmermann, Jan, Paul Glimcher, and Kenway Louie, “Multiple Timescales of Normalized Value Coding Underlie Adaptive Choice Behavior,” *Nature Communications* 9, (2018), 3206.

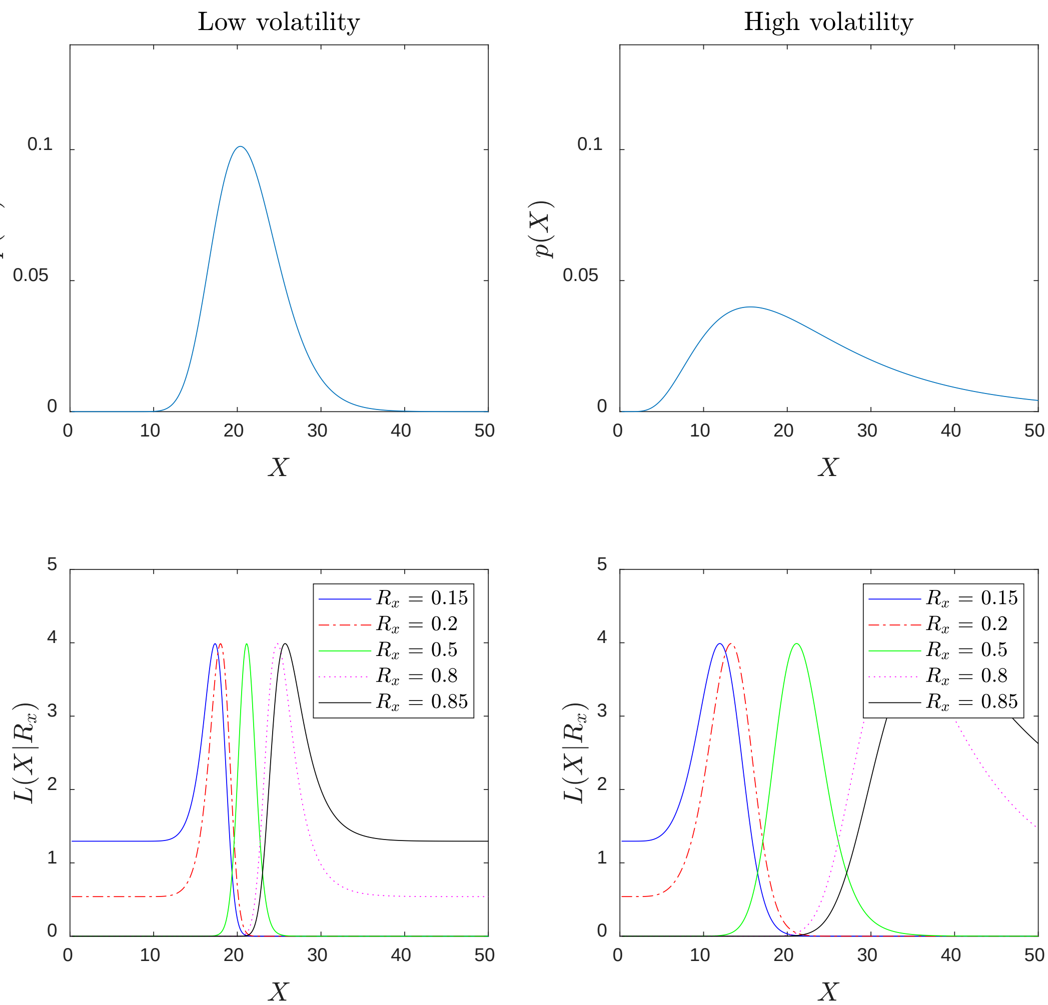

Figure 1. The stimulus distribution and the implied likelihood functions under efficient coding. Each of the two upper graphs plots a different lognormal stimulus distribution of X (left: low volatility; right: high volatility). The two lower graphs plot the implied likelihood functions (left: low volatility; right: high volatility), for five different values of the noisy signal, $R_x = 0.15$, 0.2, 0.5, 0.8, and 0.85. (Large values of X tend to generate R_x close to one, while small values of X tend to generate R_x that is close to zero; see Appendix A for more details about the distribution of R_x.) The lower graphs demonstrate that the likelihood functions depend on the volatility of the stimulus distribution. The parameter values for the two stimulus distributions are $\sigma_x = 0.19$ (left) and $\sigma_x = 0.55$ (right); for both stimulus distributions we set $\mu_x = 3.05$ and $\sigma = 0.1$.

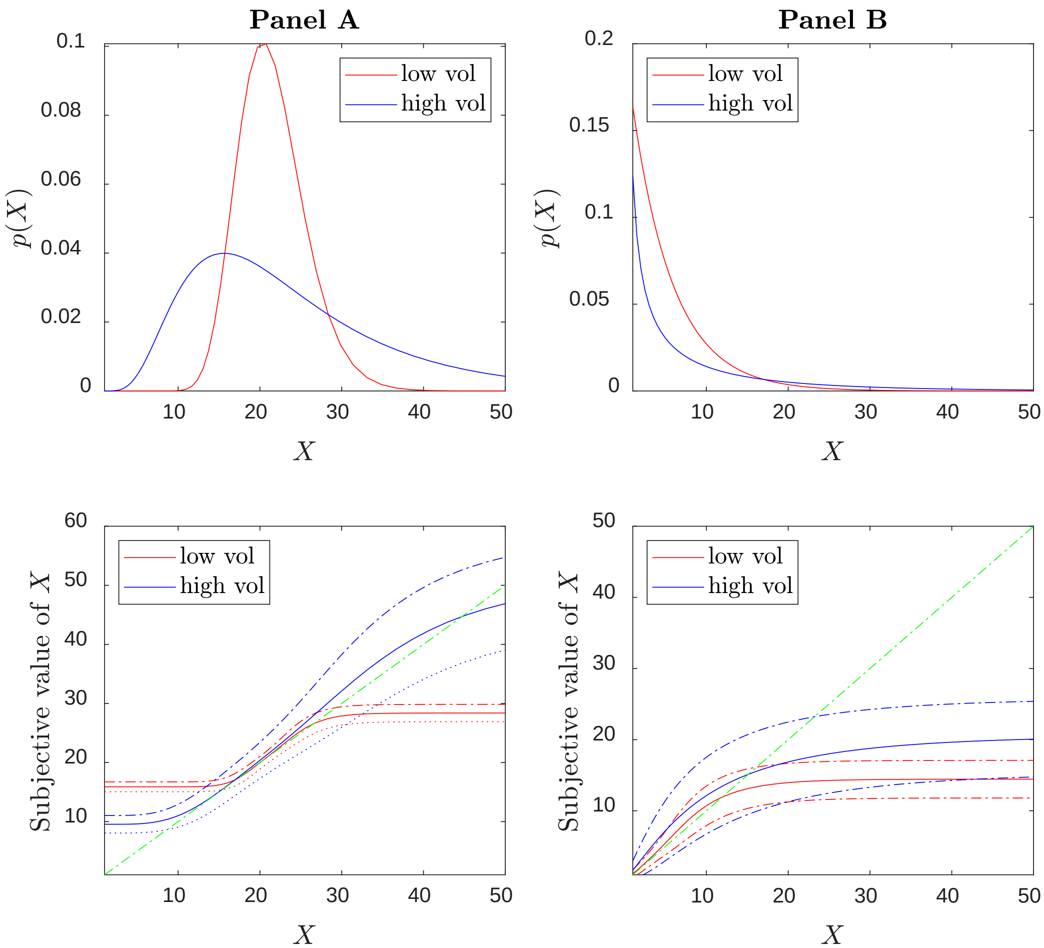

Figure 2. The subjective value function depends on the shape of the stimulus distribution. Panel A: the upper graph plots two lognormal stimulus distributions for X (low volatility distribution in red is given by $\sigma_x = 0.19$; high volatility distribution in blue is given by $\sigma_x = 0.55$). The bottom graph plots the implied subjective value functions, $v(X)$, and their one-standard-deviation bounds $v(X) \pm \sigma(X)$. The other parameter values are $\mu_x = 3.05$ and $\sigma = 0.1$. Panel B: the upper graph plots two monotonically decreasing stimulus distributions for X, which are characterized by a gamma distribution,

$$p(X;k,\theta) = \frac{1}{\Gamma(k)\theta^k} X^{k-1} e^{-X/\theta},$$

where $\Gamma(\cdot)$ is the gamma function (low volatility distribution in red is given by $\theta = 5$ and $k = 1$; high volatility distribution in blue is given by $\theta = 20$ and $k = 0.25$). The bottom graph plots the implied subjective value functions, $v(X)$, and their one-standard-deviation bounds $v(X) \pm \sigma(X)$. The other parameter value is $\sigma = 0.1$. In both bottom panels, the green dashed line is the forty-five degree line.

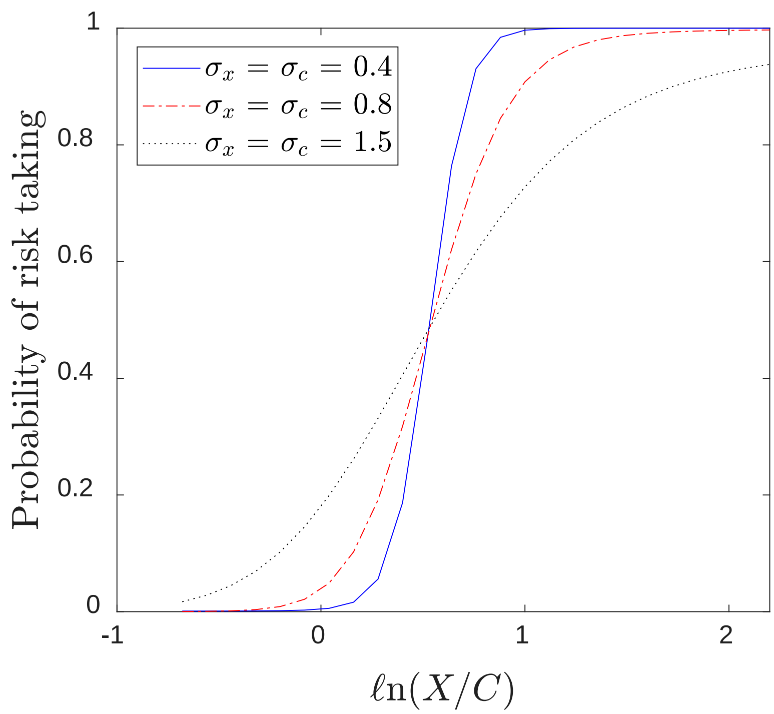

Figure 3. The relationship between the probability of risk taking and ℓn(*X*/*C*). The figure plots, for each volatility level of the stimulus distribution, $\sigma_x = \sigma_c = 0.4$, 0.8, and 1.5, the probability of risk taking computed in (14) against the natural logarithm of *X* over *C*, ℓn(*X*/*C*). For each volatility level, we set *C* to $\exp(\mu_c + \frac{1}{2}\sigma_c^2)$ while we vary the value of *X*. The other parameter values are: $\mu_x = 3.05$, $\mu_c = 2.35$, $p = 0.59$, and $\sigma = 0.1$.

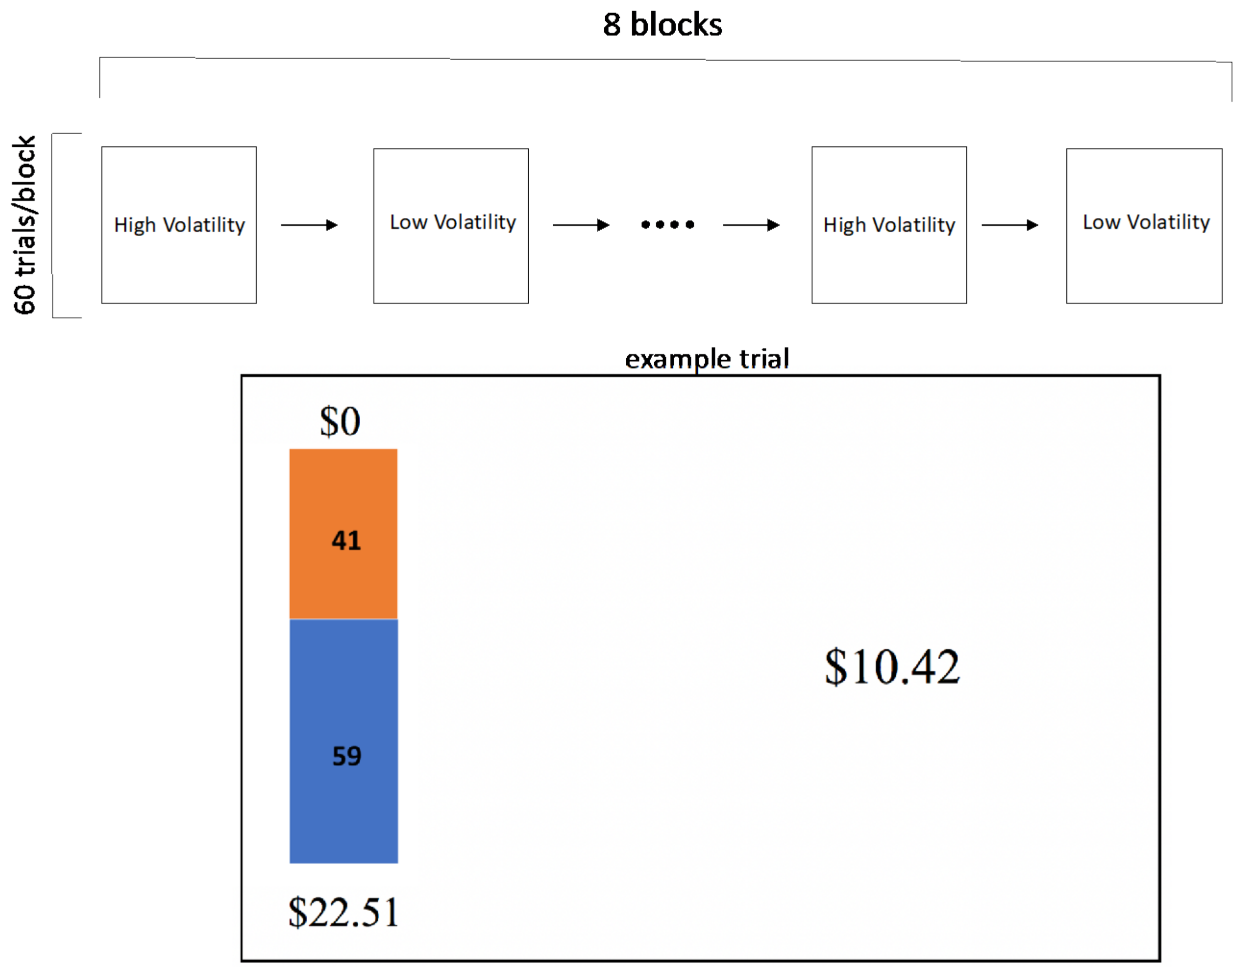

Figure 4. Experimental design for risky choice task. The task consists of eight blocks, which alternated between a high volatility condition and a low volatility condition. In the example trial screenshot above, the risky lottery is shown on the left, and the certain option is shown on the right. In each trial, the subject has unlimited time to decide which of the two options she prefers. At the end of each block, the subject is allowed to take a self-paced break, after which the next block begins.

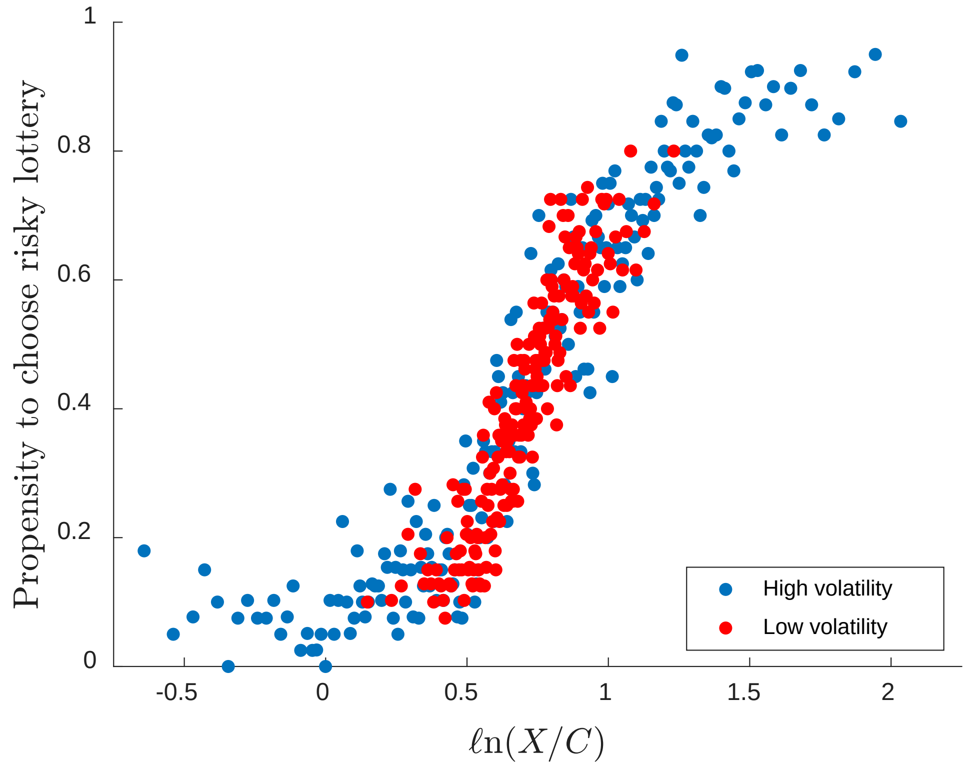

Figure 5. Average levels of risk taking across conditions. The figure plots the proportion of trials on which subjects choose the risky lottery as a function of ℓn(X/C), for both the high volatility condition and the low volatility condition. Data are pooled across trials and subjects. For each of the two experimental conditions, we bin the ℓn(X/C) variable into two-hundred bins such that each bin has an equal number of trials.

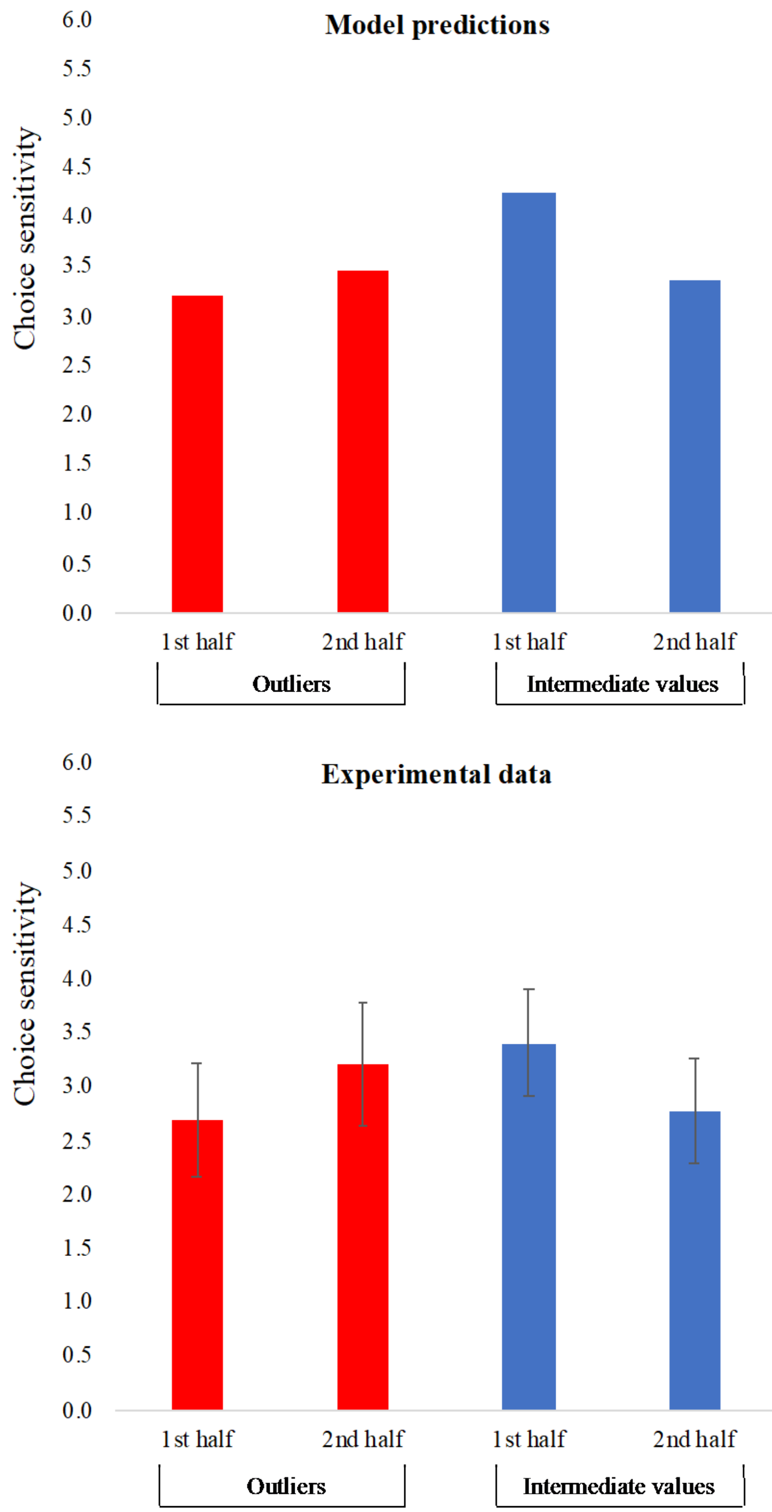

Figure 6. Adaptation dynamics. The top panel shows model predictions, and the bottom panel shows experimental data. Data are taken from blocks 3, 5, and 7, and only include the first and last 5 trials of each block. Outliers are defined as values of $\ell n(X/C)$ that fall outside of three standard deviations of the mean value of the preceding low volatility block. For the model generated choice data, we set σ to 0.1, and set the lookback window N to 10.

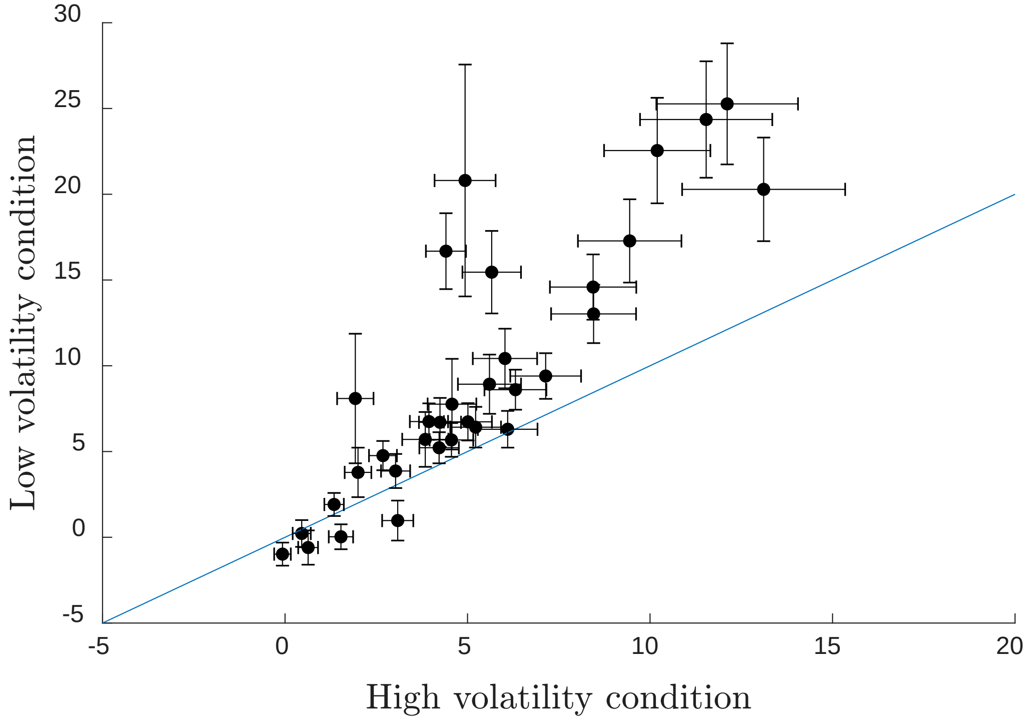

Figure 7. Individual subject estimated treatment effects. For each subject, and each condition, we run a logistic regression of the form: $\text{risky}_t = \alpha + \beta\cdot\ell\text{n}(X_t/C_t)+\varepsilon_t$. The *x*-axis measures the estimated β in the high volatility condition, while the *y*-axis measures the estimated β in the low volatility condition. Each point represents a single subject, and the length of each black bar denotes two standard errors of the mean. The blue line is the forty-five degree line.

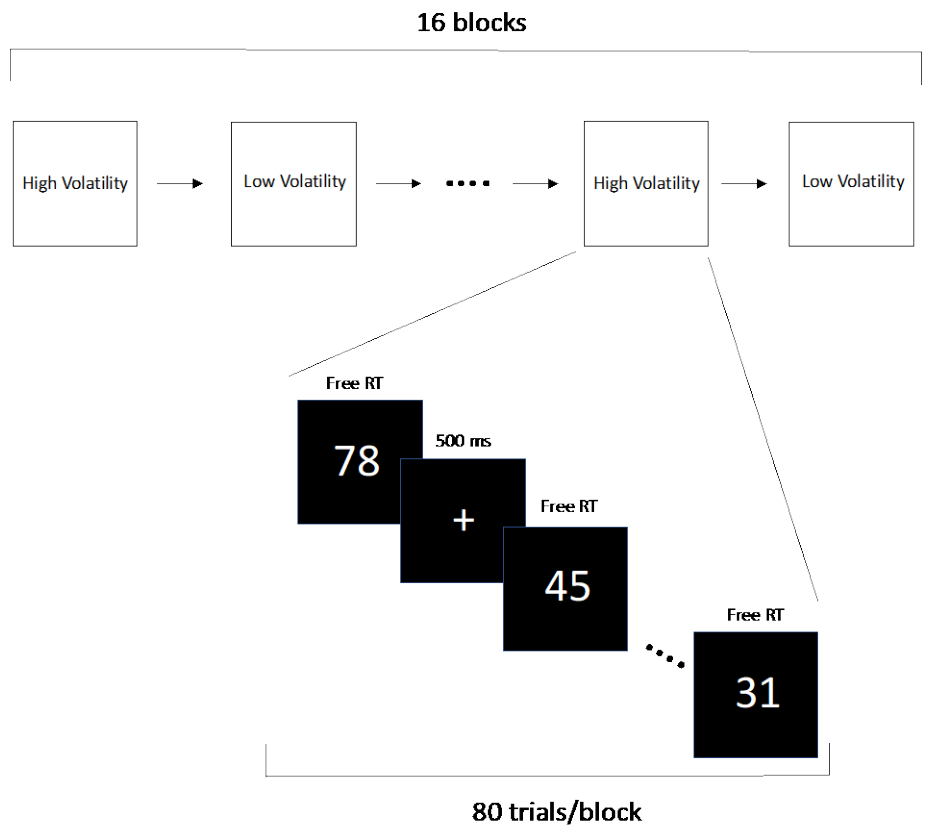

Figure 8. Experimental design for riskless choice task. The task consists of sixteen blocks, which alternated between a high volatility condition and a low volatility condition. On each trial, the subject is incentivized to classify as quickly and accurately as possible, whether the stimulus integer is larger or smaller than the number 65. In the high volatility condition, the integers are drawn uniformly from [31, 99], while in the low volatility condition, the integers are drawn uniformly from [51, 79].

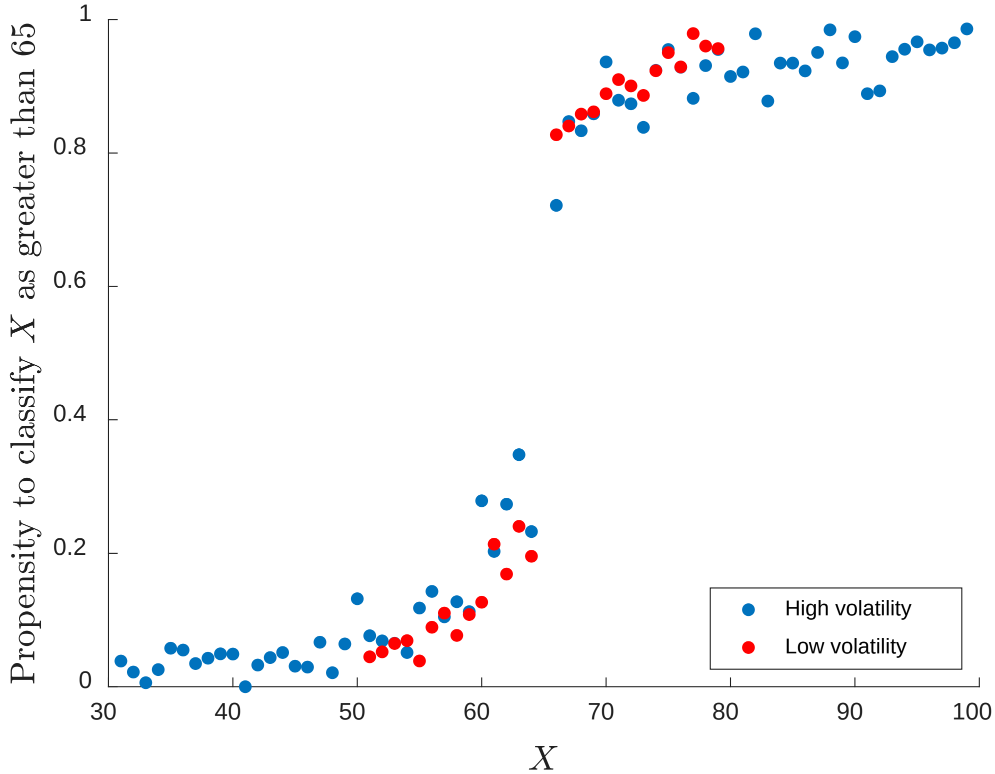

Figure 9. Classification of numbers in riskless choice task. The *x*-axis denotes the integer *X* that is presented on each trial. The *y*-axis denotes the proportion of trials for which the subject classified the integer *X* as greater than 65. Data are disaggregated by high and low volatility condition. In the high volatility condition, the integers are drawn uniformly from [31, 99], while in the low volatility condition, the integers are drawn uniformly from [51, 79].

Panel A

Dependent variable: "Choose risky lottery"	(1) All data	(2) $0.02 < \ell n(X/C) < 1.36$	(3) First half of block	(4) Second half of block	(5) First 10 trials of each block	(6) First trial of each block
high	0.99***	0.68***	0.84***	1.15***	1.10***	2.05**
	(0.28)	(0.26)	(0.30)	(0.32)	(0.29)	(1.04)
$\ell n(X/C)$	4.21***	4.21***	4.05***	4.37***	3.94***	5.05***
	(0.62)	(0.62)	(0.62)	(0.65)	(0.70)	(1.27)
$\ell n(X/C) \times high$	−1.35***	−0.92***	−1.11***	−1.59***	−1.14***	−2.43**
	(0.38)	(0.32)	(0.40)	(0.42)	(0.41)	(1.22)
Constant	−3.38***	−3.38***	−3.26***	−3.51***	−3.26***	−4.15**
	(0.47)	(0.47)	(0.47)	(0.50)	(0.50)	(0.96)
Pseudo R-squared	0.17	0.13	0.18	0.17	0.17	0.19
Observations	15,840	14,101	7,920	7,920	2,640	264

Panel B

Dependent variable: "Choose risky lottery"	(1) All data	(2) First half of block	(3) Second half of block	(4) First 10 trials of each block	(5) First trial of each block
high	−0.10	−0.20	−0.01	0.03	−0.42
	(0.29)	(0.30)	(0.33)	(0.49)	(1.18)
X	0.18***	0.19***	0.17***	0.18***	0.21***
	(0.03)	(0.03)	(0.03)	(0.03)	(0.07)
C	−0.41***	−0.42***	−0.41***	−0.39***	−0.54***
	(0.06)	(0.07)	(0.06)	(0.07)	(0.14)
$X \times high$	−0.08***	−0.09***	−0.08***	−0.07***	−0.13**
	(0.02)	(0.02)	(0.02)	(0.03)	(0.06)
$C \times high$	0.18***	0.19***	0.16***	0.17***	0.34**
	(0.04)	(0.04)	(0.04)	(0.04)	(0.14)
Constant	0.03	−0.05	0.12	−0.19	0.59
	(0.42)	(0.43)	(0.45)	(0.53)	(1.19)
Pseudo R-squared	0.15	0.15	0.15	0.15	0.16
Observations	15,840	7,920	7,920	2,640	264

Table 1. Logistic regressions of probability of risk taking. In both panels A and B, the dependent variable takes the value of one if the subject chose the risky lottery, and zero if the subject chose the certain option. 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. Standard errors are clustered at the subject level, and ***, **, * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Dependent variable: "Choose risky lottery"	**(1)** All data	**(2)** First and last 20 trials	**(3)** First and last 10 trials	**(4)** First and last 5 trials
$\ell n(X/C)$	3.07***	3.10***	3.41***	3.40***
	(0.46)	(0.46)	(0.49)	(0.50)
$\ell n(X/C) \times outlier$	–0.39**	–0.38**	–0.61***	–0.70*
	(0.18)	(0.17)	(0.23)	(0.41)
$\ell n(X/C) \times second$	–0.08	–0.08	–0.53**	–0.63**
	(0.09)	(0.13)	(0.21)	(0.27)
$\ell n(X/C) \times second \times outlier$	0.15	0.02	0.57*	1.14**
	(0.18)	(0.20)	(0.31)	(0.53)
Constant	–2.51***	–2.47***	–2.53***	–2.43***
	(0.37)	(0.36)	(0.35)	(0.38)
Pseudo R-squared	0.25	0.25	0.27	0.27
Observations	5,940	3,960	1,980	990

Table 2. Adaptation and sensitivity to outliers. Logistic regression results using only data from the high volatility blocks (except for the first block, which does not have an immediately preceding low volatility block). The dependent variable *outlier* takes the value of one if the value of $\ell n(X/C)$ is more than three standard deviations from the mean, where these statistics are calculated using the sample moments from the sixty trials in the immediately preceding low volatility block. The dummy variable *second* takes the value of one if the trial belongs to the second half of the block (trials 31-60), and zero if it belongs to the first half of the block (trials 1-30). Standard errors are clustered at the subject level, and ***, **, * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Dependent variable: "Classifiy as greater than 65"	(1) All data	(2) $51 \leq X \leq 79$	(3) $X < 60$ or $X \geq 70$
high	0.00	0.04	–0.07
	(0.02)	(0.05)	(0.04)
$\ell\mathrm{n}(X/65)$	19.53***	19.53***	17.45***
	(2.09)	(2.09)	(1.58)
$\ell\mathrm{n}(X/65)\times high$	–9.71***	–3.56***	–7.98***
	(1.11)	(0.95)	(0.69)
Constant	0.20***	0.20***	0.30***
	(0.07)	(0.07)	(0.08)
Pseudo R-squared	0.53	0.48	0.63
Observations	16,640	11,892	12,807

Table 3. Classification in riskless choice task. Logistic regression where the dependent variable takes the value of one if the subject classified the stimulus, X, as larger than 65, 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. In the high volatility condition, the integer X is drawn uniformly from [31, 99], while in the low volatility condition, the integer is drawn uniformly from [51, 79]. Standard errors are clustered at the subject level, and ***, **, * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Appendix A: Theoretical Derivations

A.1. Proof of Proposition 1

To find a set of likelihood functions for a given stimulus distribution $p(\theta)$ that satisfies the efficient coding condition in (3), we follow Wei and Stocker (2015) and transform the stimulus space into a "sensory space" through a change of variable $\tilde{\theta} = F(\theta)$, where $F(\theta) = \int_{-\infty}^{\theta} p(\xi)d\xi$. is the cumulative density function of θ.

We first show that the efficient coding condition is satisfied in the sensory space if the transformed likelihood function *location-independent*:

$$L(\tilde{\theta} \mid m) = p(m \mid \tilde{\theta}) = g(\tilde{\theta} - m), \tag{A1}$$

where $g(\cdot)$ is some smooth density function that integrates to one. For this proof, note that equation (A1) allows us to write the Fisher information $J(\tilde{\theta})$ as

$$\begin{aligned} J(\tilde{\theta}) &= \int_{-\infty}^{\infty} \left(\frac{\partial \ell n\, p(m \mid \tilde{\theta})}{\partial \tilde{\theta}} \right)^2 p(m \mid \tilde{\theta}) dm \\ &= \int_{-\infty}^{\infty} (g'(\alpha))^2 g^{-1}(\alpha) d\alpha = C, \end{aligned} \tag{A2}$$

where $C > 0$ is a constant, and $\alpha \equiv \tilde{\theta} - m.$ In this sensory space, $\tilde{\theta}$ is uniformly distributed between zero and one. As a result,

$$\sqrt{J(\tilde{\theta})} = \sqrt{C} \propto p(\tilde{\theta}) = 1 \tag{A3}$$

for any $\tilde{\theta} \in [0,1]$. That is, the efficient coding condition is satisfied in the sensory space.

Next, we show that (A1) implies that the efficient coding condition is also satisfied in the original stimulus space. By (A1), the likelihood function in the stimulus space is

$$L(\theta \mid m) = p(m \mid \theta) = g(F(\theta) - m). \tag{A4}$$

The Fisher information $J(\theta)$ is

$$\begin{aligned} J(\theta) &= \int_{-\infty}^{\infty} \left(\frac{\partial \ell n\, p(m \mid \theta)}{\partial \tilde{\theta}} \right)^2 p(m \mid \theta) dm \\ &= p^2(\theta) \int_{-\infty}^{\infty} (g'(\alpha))^2 g^{-1}(\alpha) d\alpha = C \cdot p^2(\theta). \end{aligned} \tag{A5}$$

Therefore

$$\sqrt{J(\theta)} = \sqrt{C} \cdot p(\theta) \propto p(\theta) \tag{A6}$$

for any $\theta \in (-\infty, \infty)$. That is, the efficient coding condition is satisfied in the stimulus space.

In Proposition 1, we present the likelihood function $L(\theta \mid m)$ with a specific function for the $g(\cdot)$ function: we assume g is a normal density function with mean 0 and variance σ^2. The likelihood function then becomes

$$L(\theta \mid m) = p(m \mid \theta) = \frac{1}{\sqrt{2\pi} \cdot \sigma} \exp\left(-\frac{(F(\theta) - m)^2}{2\sigma^2} \right) \tag{A7}$$

for any given $m \in (-\infty, \infty)$, same as Equation (5) in the main text.

A.2. Properties of $p(R_x)$

Given $p(X; \mu_x, \sigma_x)$ in (6) and $p(R_x \mid X)$ in (7), the distribution of R_x can be derived as

$$\begin{aligned} p(R_x) &= \int_0^\infty p(R_x \mid X) p(X) dX \\ &= \frac{1}{2\pi\sigma} \int_0^\infty \exp\left(-\frac{(\Phi((\ell\text{n}\, X - \mu_x)/\sigma_x) - R_x)^2}{2\sigma^2} \right) \frac{1}{\sigma_x X} \exp\left(-\frac{(\ell\text{n}\, X - \mu_x)^2}{2\sigma_x^2} \right) dX \\ &= \frac{1}{2\pi\sigma} \int_0^\infty \exp\left(-\frac{(\Phi(y) - R_x)^2}{2\sigma^2} \right) \exp\left(-\frac{y^2}{2} \right) dy. \end{aligned} \tag{A8}$$

Notice that this expression does not depend on distribution parameter μ_x and σ_x. Furthermore, this "invariance" result is a general statement that is independent of the specific assumption of lognormal distribution for $p(X)$: all continuous stimulus distributions lead to the same $p(R_x)$. To see this, we write in general

$$p(R_x \mid X) = g\left(\int_{-\infty}^{X} p(y) dy - R_x\right), \tag{A9}$$

where, in the case of (7), $g(\cdot)$ is a normal density function with mean zero and standard deviation σ. Equations (A8) and (A9) then imply

$$\begin{aligned} p(R_x) &= \int_0^\infty p(R_x \mid X) p(X) dX \\ &= \int_0^\infty g\left(\int_{-\infty}^{X} p(y) dy - R_x\right) p(X) dX = \int_0^1 g(z - R_x) dz, \end{aligned} \tag{A10}$$

$$\text{where } z \equiv \int_{-\infty}^{X} p(y) dy.$$

This equation makes it clear that not only is the case that $p(R_x)$ does not depend on μ_x and σ_x, it does not depend on the entire *shape* of the stimulus distribution $p(X)$. A sufficient condition for this "invariance" result is that *i*) the likelihood function is location-independent in the sensory space (as we assume in equation (A1)), and *ii*) the transformation function from stimulus space to the sensory space is the cumulative density function of the stimulus value.

[Place Figure A1 about here]

Next, we look at the asymptotic behavior of $p(R_x)$ as σ goes to zero. From (A10) we know that

$$p(R_x) = \int_0^1 g(z - R_x) dz \xrightarrow[\delta \to 0]{} \int_0^1 \delta(z - R_x) dz = \begin{cases} 1 & 0 < R_x < 1 \\ 0 & \text{otherwise} \end{cases}, \tag{A11}$$

where $\delta(\cdot)$ represents the Dirac delta function. That is, in the limiting “noiseless” case, $p(R_x)$ converges to a uniform distribution between zero and one. To illustrate this point, Figure A1 plots $p(R_x)$ for $\sigma = 0.15$, 0.1, and 0.05; as σ decreases, $p(R_x)$ indeed converges toward a uniform distribution.

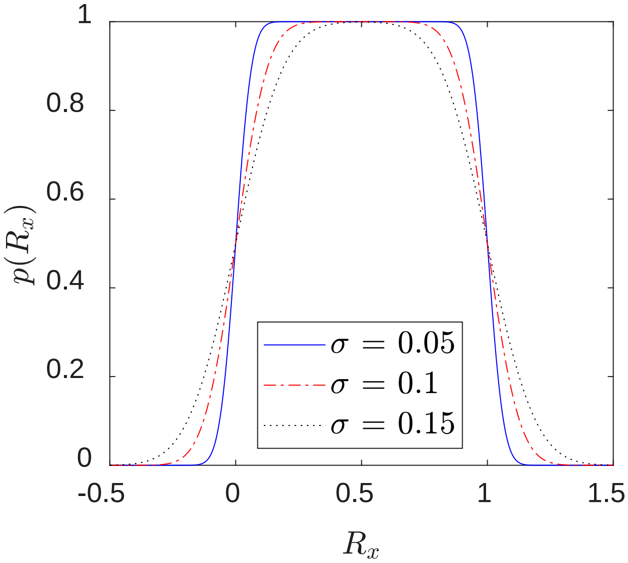

Figure A1. The unconditional probability density function of the noisy signal R_x. This figure plots the unconditional distribution of the noisy signal $p(R_x)$ according to (8), for $\sigma = 0.05$, 0.1, and 0.15.

Appendix B: Additional Analyses on Scale Invariance

In the discussion section of the main text, we noted that Weber's law does not hold universally in our model, but it can arise under a particular type of prior distribution—one in which there is greater mass on higher value stimuli, compared to lower value stimuli. In our choice environment, Weber's law would predict a type of "scale invariance," in which risk taking depends only on the ratio, X/C, but not the separate levels of X and C. In fact, this is a key prediction of the KLW model, which stems, in part, from the logarithmic compression in the encoding function that is assumed for both X and C. In our model, when the prior distribution is lognormal, efficient coding does *not* imply logarithmic compression in the encoding functions.

To see how this breaks the prediction of scale invariance in our model, consider a choice set in which X and C are drawn from their respective lognormal distributions. Suppose further that, in this choice set, C takes the mean value of its distribution, but X is drawn from the left tail of its lognormal distribution. In this case, efficient coding implies that the *DM* will be able to finely discriminate between C and nearby values, but she will not be able to discriminate well between X and *its* nearby values. As a consequence, if we present the *DM* with a new choice set by multiplying the original values of X and C by a common constant that is greater than one, then this multiplication will cause a substantial increase in the *DM*'s perceived valuation of C, but not of X. Therefore, the *perceived* ratio (X/C) decreases on average, causing a reduction in risk taking, and hence, a violation of scale invariance.

A key ingredient of this argument, however, is that X must be drawn from one of the tails of its distribution, so that the multiplier generates a smaller increase in the perceived value of X, compared to that of C. If instead, the values of X and C are close to their respective mean values under a lognormal prior, then multiplying both values by a common constant should lead to a similar increase in the perception of each value, and therefore, only a small change in risk taking. Our model therefore predicts that large violations of scale invariance are not expected to occur often, since these large violations occur only when one of the payoffs is drawn from the tails of its distribution.

To assess this prediction quantitatively, in the upper panel of Figure B2, we plot the model predicted probability of accepting the risky lottery as a function of $\ell n(X/C)$, for three different regions of X: 1) low values of X (bottom 30% of its distribution); 2) intermediate values of X (between 30^{th} and 70^{th} percentile of its distribution); and 3) high values of X (top 30% of its distribution). Each point in the figure represents a single choice set and its model predicted probability of choosing the risky lottery. The values of X and C are generated over equally spaced grid points of $[X_{min}, X_{max}] \times [C_{min}, C_{max}] = [4, 50] \times [2.5, 25]$.

The figure clearly shows violations of scale invariance: for a given value of $\ell n(X/C)$, the model predicted probability depends on whether X is drawn from the left tail (in blue), intermediate region (in orange), or right tail (in grey). As expected, these violations are most severe for choice sets in which X is drawn from the left tail, and specifically, when $\ell n(X/C)$ is relatively large. In order for X to be drawn from the left tail and for $\ell n(X/C)$ to be relatively large, C must *also* be drawn from its left tail, which makes the occurrence of these particular choice sets very rare. Thus, the largest violations of scale invariance occur for those choice sets that the *DM* is rarely presented. In contrast, the points in orange, which represent the most likely values of X, fall roughly on the same curve, which is consistent with the model predicting an approximate scale invariance among these likely values of $\ell n(X/C)$.

To assess whether our experimental data are consistent with scale invariance, we plot our risky choice data by also cutting the data into the three subsamples of X described above. In the bottom panel of Figure B2 we plot the data using 20 bins, where each bin contains an equal number of trials, for each of the three different subsamples of X; this binning procedure is important because it controls for the frequency that each value of $\ell n(X/C)$ is presented in our experiment. The data are noisy since we cut the data into three subsamples and analyze only the high volatility condition, but the important feature to note is that for each value of $\ell n(X/C)$ in the figure—which is weighted by its frequency of being presented in the experiment—the probability of risk taking varies only slightly with the level of X.

In summary, our model predicts that risk taking will depend on the level of X and C, but this dependence is concentrated mainly among outlier values, for which the *DM* has difficulty discriminating between nearby values. Among payoffs that the *DM* is most likely to observe in a choice set, our model predicts that risk taking should be approximately scale invariant, which is broadly consistent with our data.

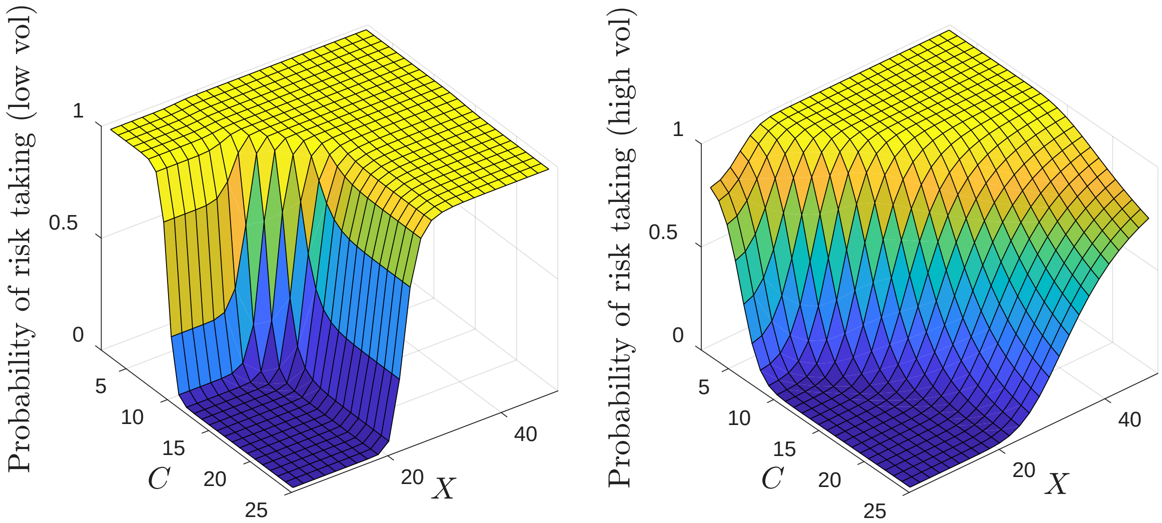

Figure B1. Probability of risk taking as a function of *X* and *C*. The figure plots the probability of risk taking computed in equation (14) in the main text, for two different volatility levels of the stimulus distribution, $\sigma_x = \sigma_c = 0.19$ (low volatility) and $\sigma_x = \sigma_c = 0.55$ (high volatility), the probability of risk taking computed in (14) as a function of both *X* and *C*. The parameter values are: $\mu_x = 3.05$, $\mu_c = 2.35$, $p = 0.59$, and $\sigma = 0.1$.

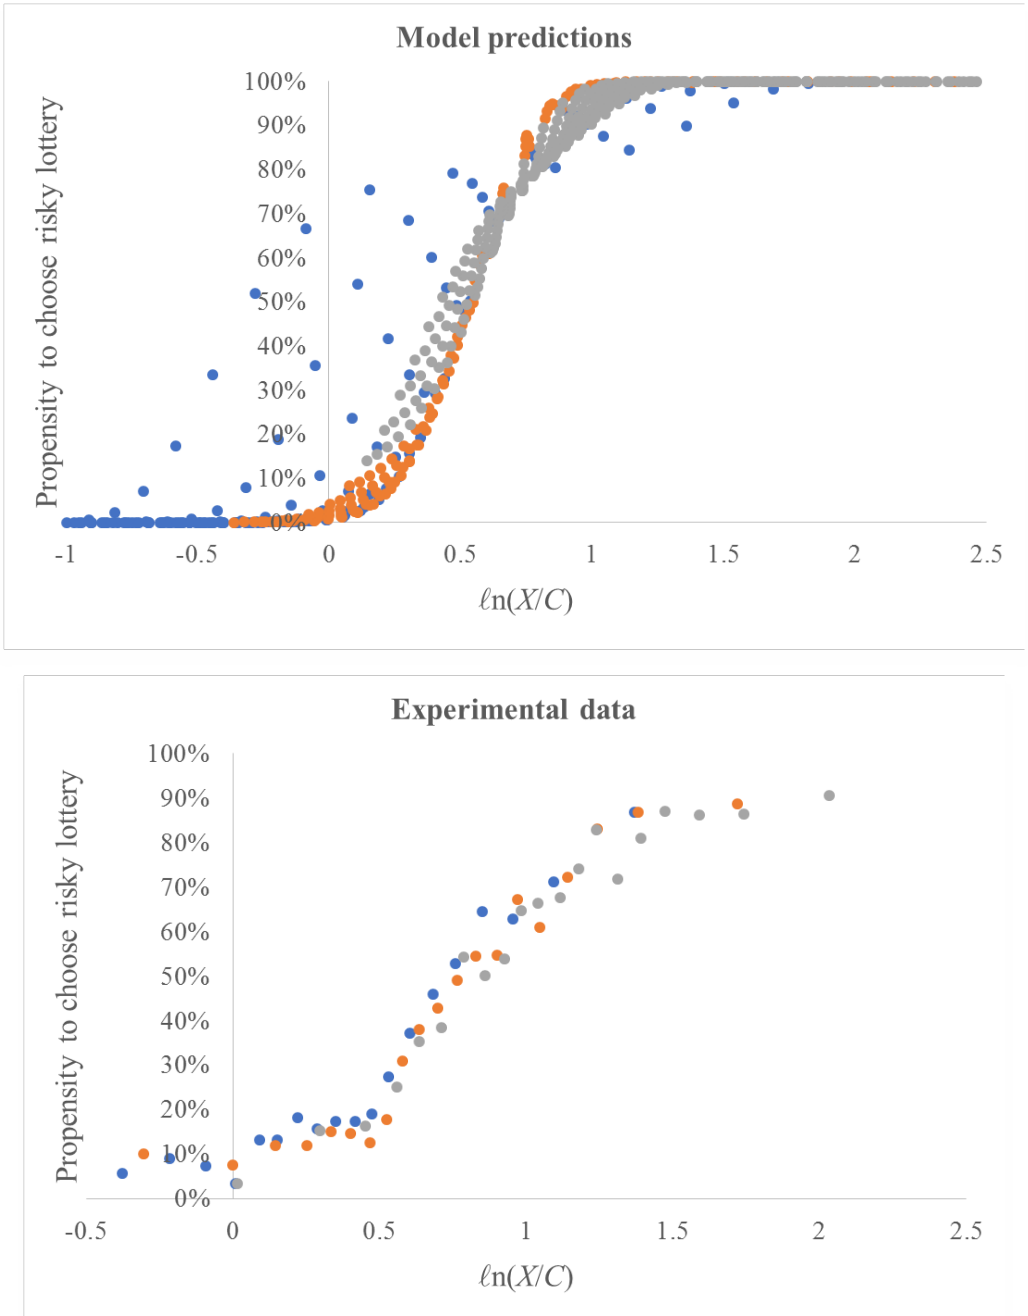

Figure B2. Probability of risk taking for different levels of *X*. The upper panel plots the model predicted probability of risk taking in the high volatility condition ($\sigma_x = \sigma_c = 0.55$), for the three different intervals of *X*: 1) *Blue*: low values of *X* (bottom 30% of distribution); 2) *Orange*: intermediate values of *X* (between 30th and 70th percentile of distribution); and 3) *Grey*: high values of *X* (top 30% of distribution). Each point in the figure represents a single choice set and its model predicted probability of choosing the risky lottery. The bottom panel plots the experimental data from the (same) high volatility condition. For each of the three different intervals of *X*, we place the values of $\ell n(X/C)$ into 20 equally sized bins, and compute the average value of $\ell n(X/C)$ within this bin and the associated average levels of risk taking. For the upper panel, we set $\sigma = 0.1$.

Appendix C: Comparing Efficient Coding with Salience Theory (Bordalo, Gennaioli, and Shleifer, 2012)

Here we provide a comparison of our model with salience theory of choice under risk (Bordalo, Gennaioli, and Shleifer, 2012) (henceforth BGS). Under salience theory, attention is drawn towards salient payoffs, which are then overweighted relative to their objective probabilities. According to BGS, "a lottery payoff is salient if it is very different in percentage terms from the payoffs of other available lotteries" (pg. 1244). This definition captures two important psychological properties: (i) perception is based on the difference between stimuli; and (ii) it becomes more difficult to detect a given difference between two stimuli as the *level* of each stimulus increases.

While both efficient coding and salience theory provide models of distorted perception in risky choice, the mechanisms that generate this distortion are distinct. In our experiment, there are two potential sources of uncertainty: (i) the *DM* has imperfect perception and is therefore uncertain about the payoff values that characterize each lottery (X and C in our experiment), even after the choice set is presented; and (ii) the *DM* is exposed to the more standard source of uncertainty whereby, if she chooses the risky lottery, she is uncertain about whether she will receive X (with probability p) or 0 (with probability $1 - p$). In our model of efficient coding, only the first source of uncertainty triggers the mechanism that distorts perception. Conditional on the distorted perception of the payoff values, the *DM* then handles the second source of risk in the standard way by selecting the lottery with the highest expected value.

In contrast, salience theory assumes that only the second source of uncertainty—regarding the outcome of the risky lottery (X, p; 0, $1 - p$)—is relevant. There is no uncertainty about the payoff values that characterize each lottery; the *DM* observes X and C perfectly. Instead, these payoffs differentially attract the *DM*'s attention, which distorts the decision weight attached to each payoff.

It is worth noting that, even though salience theory assumes that payoffs are perceived without noise, the theory can easily be applied to the case in which the *DM*'s only source of uncertainty is about the payoff values that are presented in the choice set. In this case, we can think of there being a continuous set of states, and each state delivers a pair (X, C). In this framework, the *DM*'s choice would then be between the two *priors* that generate X and C. Notice that this decision is *not* the one that subjects in our experiments face, but nonetheless, salience theory makes sharp predictions about which payoffs are overweighted. In particular, we show below that the *DM* overweights outlier values of X , which distorts the perceived value of each prior.

After analyzing the effect of salience when the *DM*'s only source of uncertainty is about the payoff values in the choice set, we examine the second source of uncertainty. When examining this second source, we assume that the *DM* has perfect information about the payoff values in the choice set, but now faces uncertainty about which outcome the risky lottery will deliver. This second application of salience theory is more closely aligned with our experimental setup, and we find that the theory does a good job explaining most dimensions of our data.

C.1 Salience Model

Salience theory operates by distorting the weight on each state payoff for a given lottery. Here we consider the case in which the choice set contains only two lotteries, L^1 and L^2. We define a state space S, where each $s \in S$ occurs with probability p_s, and lottery L^i delivers payoff x_s^i in state s. As in our model of efficient coding, we assume the *DM* uses a linear value function $v(z) = z$. Without any salience distortions, the value of lottery L^i is given by:

$$V(L^i) = \sum_{s\in S} p_s v(x_s^i). \tag{C1}$$

The salience model departs from this valuation equation by assuming that the *DM* does not use the set of objective probabilities $\{p_s\}$, but instead uses a set of decision weights for lottery L^i denoted by $\{\pi_s^i\}$, where $\pi_s^i = p_s\omega_s^i$ for each state *s*. The distortion factor, ω_s^i, distorts the objective probability, and is defined as:

$$\omega_s^i = \frac{\delta^{-\sigma(x_s^i, x_s^{-i})}}{\sum_{r\in S} \delta^{-\sigma(x_r^i, x_r^{-i})} p_r}. \tag{C2}$$

In the above equation, $\delta \in (0, 1]$ captures the degree to which salience distorts objective probabilities. $\sigma(x_s^i, x_s^{-i})$ is the *salience function* which maps state payoffs delivered by L^i and L^{-i} into a salience measure. This function formalizes features of human perception; the two most important properties of the salience function (for our experimental setup) are ordering and diminishing sensitivity. Ordering implies that states in which there is a larger difference between payoffs (across lotteries) are more salient. This captures the intuition that attention is drawn towards attributes with larger differences. Diminishing sensitivity implies that adding a constant to all payoffs in a state will decrease the salience of that state.[23] This property is closely linked to the diminishing sensitivity in efficient coding, which also predicts that perception is less sensitive among outlier values.

We will use a particular functional form of the salience function that satisfies these two properties, given by:

$$\sigma(x_s^i, x_s^{-i}) = \frac{|x_s^i - x_s^{-i}|}{|x_s^i| + |x_s^{-i}| + \theta}, \tag{C3}$$

where $\theta > 0$ controls the degree of diminishing sensitivity. Conditional on this salience function, the valuation of lottery L^i under salience theory is given by:

$$V(L^i) = \sum_{s\in S} \pi_s^i v(x_s^i). \tag{C4}$$

It follows that a payoff is overweighted, relative to its objective probability, p_s, if and only if $\omega_s^i > 1$. The *DM* then chooses lottery L^1 if and only if $V(L^1) > V(L^2)$.

C.2 Case 1: Applying Salience Theory When *DM* is Only Uncertain about Payoff Values

Recall that there are two potential sources of uncertainty that the *DM* faces in our experiment. First, the *DM* can be uncertain about which payoffs (X, C) characterize the risky lottery and the certain option in the choice set. Second, conditional on knowing these values of X and C, the *DM* is uncertain about which outcome will obtain from the risky lottery: X (with probability p) or 0 (with probability $1 - p$). In this section, we apply salience theory to the first source of uncertainty, in which the *DM* is uncertain only about the values of X and C. We can then think of L^1 as delivering a random payoff X and L^2 as delivering a random payoff C. As in our experiment, we assume that the values of X and C are jointly drawn from a lognormal distribution:

[23] As BGS put it, "the intensity with which payoffs in a state are perceived increases as the state's payoffs approach the status quo of zero…" (BGS 2012, pg. 1254).

$$\begin{pmatrix} \ell n X \\ \ell n C \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \mu_x \\ \mu_c \end{pmatrix}, \begin{pmatrix} \sigma_x^2 & \rho\sigma_x\sigma_c \\ \rho\sigma_x\sigma_c & \sigma_c^2 \end{pmatrix} \right). \tag{C5}$$

As a result, in state s in which L^1 delivers X_s, L^2 will on average deliver[24]

$$\mathbb{E}[\tilde{C} \mid X_s] = \exp\left(\mu_c + \frac{\rho\sigma_c}{\sigma_x}(\ell n X_s - \mu_x) + \frac{1}{2}\sigma_c^2(1-\rho^2) \right). \tag{C6}$$

The salience measure for L^1 in state s is then given by:

$$\sigma(X_s, \mathbb{E}[\tilde{C} \mid X_s]) = \frac{\left| X_s - \exp\left(\mu_c + \frac{\rho\sigma_c}{\sigma_x}(\ell n X_s - \mu_x) + \frac{1}{2}\sigma_c^2(1-\rho^2) \right) \right|}{X_s + \exp\left(\mu_c + \frac{\rho\sigma_c}{\sigma_x}(\ell n X_s - \mu_x) + \frac{1}{2}\sigma_c^2(1-\rho^2) \right) + \theta}. \tag{C7}$$

For each X_s, we can then compute its distortion factor for L^1 as

$$\omega^1(X_s) = \frac{\delta^{-\sigma(X_s, \mathbb{E}[\tilde{C}|X_s])}}{\int_0^\infty p(X; \mu_x, \sigma_x)\delta^{-\sigma(X, \mathbb{E}[\tilde{C}|X])} dX}, \tag{C8}$$

where $p(X;\mu_x,\sigma_x) = \frac{1}{X\sigma_x\sqrt{2\pi}}\exp\left(-\frac{(\ell n X - \mu_x)^2}{2\sigma_x^2} \right)$ is a lognormal probability density function.

Below we plot $\omega^1(X)$ as a function of X, for the parameter values we used in our experimental design: $\mu_x = 3.05$, $\mu_c = 2.35$, $\sigma_x = \sigma_c = 0.55$, $\rho = 0.5$. For salience-based parameters, we set $\delta = 0.7$ and $\theta = 0.1$, consistent with the point estimates in BGS.

[Place Figure C1 about here]

We see that very large values of X draw the *DM*'s attention, and these payoffs are overweighted (their distortion factors are greater than 1). To be clear, large values of X are not salient because they themselves are large. Rather, large values of X are salient, because when they are compared to the average payoff of $L^2, \mathbb{E}[\tilde{C} \mid X]$, the *difference* is large. By the same logic, very small values of X are also salient, and thus overweighted.[25]

Without any salience distortions, it is clear that the *DM* would always prefer L^1 over L^2 because the lotteries are the same except L^1 has a strictly higher mean. Indeed, for $\delta = 1$, $V(L^1) - V(L^2) > 12.37$. However, for $\delta = 0.5$ and $\theta = 0.1$, we find that $V(L^1) - V(L^2) = 14.03 > 12.37$. That is, lottery L^1 becomes

[24] There are two ways to justify comparing X and the average realization of C when computing the salience function for X. First, one can think that for each value of X, there is a continuum of states each for a realization of C. Alternatively, one can think that for each value of X, there is a continuum of other assets with different realizations of C.

[25] With the parameter values used in Figure C1, as X and $\mathbb{E}[\tilde{C} \mid X]$ get larger, the ordering feature of the salience function $\sigma(\cdot, \cdot)$ dominates diminishing sensitivity. However, if we set $\mu_c = 3.55 > \mu_x = 3.05$, and set $\rho = 0.9$, then diminishing sensitivity dominates in this case. That is, ω^1 would *decrease* in X as X becomes very large.

relatively more attractive because its extreme payoffs are overweighted. This prediction is thus consistent with the intuition that outlier payoffs are overweighted and this drives choice in favor of the lottery L^1.

We emphasize that with our experimental data, we cannot test the above prediction. The reason is that subjects are *not* given the choice between the two "prior" lotteries. Rather, they are given the choice between two lotteries that are characterized by *draws* from these two priors. In the next section, we examine the predictions of salience theory, conditional on these realized draws.

C.3 Case 2: Applying Salience Theory When *DM* is Certain about Payoff Values

Now consider the more standard case in which the *DM* has no uncertainty about the payoff values in the choice set. The only source of uncertainty in this case is about the outcome of the risky lottery, (X, p; 0, $1 - p$). There are only two possible states of the world, $s \in \{up, down\}$. In the *up* state (with probability p), the risky lottery delivers payoff X; and in the *down* state (with probability $1 - p$), the risky lottery delivers a payoff of 0. On the other hand, the certain option delivers a payoff of C in both states of the world.

To examine the basic explanatory power of the salience model, we set $\delta = 0.7$ and $\theta = 0.1$ and plot the probability of choosing the risky lottery against V(risky lottery) – V(certain option).[26] Figure C2 clearly shows that the probability of choosing the risky lottery increases in the difference between the two subjective lottery values.[27] This is, however, not a particularly strong test of salience theory, because for $\theta = 0.1$, the *down* state is salient for all choice sets in our experiment.[28] Thus, the observed variation in behavior across trials is driven in large part by variation in payoffs, without much variation in their salience. However, the more noteworthy feature of the figure is that the *slope* of the probability of risk taking curve is greater for the low volatility condition compared to the high volatility condition, which resembles Figure 5 in the main text.

[Place Figure C2 about here]

In the basic version of the salience model we outlined above, risk taking is independent of whether the choice set is presented in the low volatility or high volatility condition. This invariance result from the basic version of the salience theory is driven by the assumption that only payoffs in the current choice set affect perception. Efficient coding, on the other hand, predicts that perception depends systematically on the payoff distribution to which the subject has recently adapted. In particular, efficient coding allows the degree of diminishing sensitivity to fluctuate based on past payoffs, which helps to generate the different slopes observed in Figure C2. In order for salience theory to account for this pattern of behavior, one could generalize the salience function to allow *past* payoffs to affect the salience of current payoffs, which would explicitly introduce context dependence in the time series, in addition to the one in the cross-section. Some guidance for this approach of using past experiences to form a perception

[26] The basic predictions are robust to the value of $\delta \in (0, 1]$.

[27] The salience model does not predict stochastic choice, and thus the model predicts only choice probabilities of 1 or 0. To generate intermediate probabilities, one could add noise to each payoff, thus generating a "salience + logit" or "salience + probit" model. Alternatively, because we pool all data across subjects, heterogeneity in δ can help generate a choice curve with intermediate probabilities of risk taking.

[28] This is because of the salient zero payoff delivered by the risky option in each trial. Things do not change significantly as we vary θ. For example, if we reduce the degree of diminishing sensitivity by setting $\theta = 10$, then the *down* state remains salient in 93% of trials. Moreover, while we model distortions as a smooth function of the difference in salience between the two states, there is still a relatively small amount of variation in the difference in salience measures, across trials.

of the current choice set is given in Bordalo, Gennaioli, and Shleifer (2017), though applications in that model focus mainly on riskless choice.

C.4 Summary

A key intuiton from salience theory is that outlier values are salient, and these values attract the *DM*'s attention and distort the weights attached to these values upwards. Indeed, when analyzing the first source of uncertainty about the payoff values of X and C that characterize the choice set, we showed (in Section C.2) that extreme values of X (both very small and very large values) are overweighted. However, these implications apply to an environment in which the *DM* is choosing between the two lognormal priors. These are not the lotteries presented to subjects in our experiment.

Instead, our experiment presents subjects with lotteries $\{(X, p; 0, 1 - p), (C,1)\}$, in which X and C are the realized values from the two lognormal priors. This implies there are only two states. When applying salience theory to this choice problem, the theory does a good job explaining the data, except for the difference in slopes across our two experimental conditions (Figure C2).

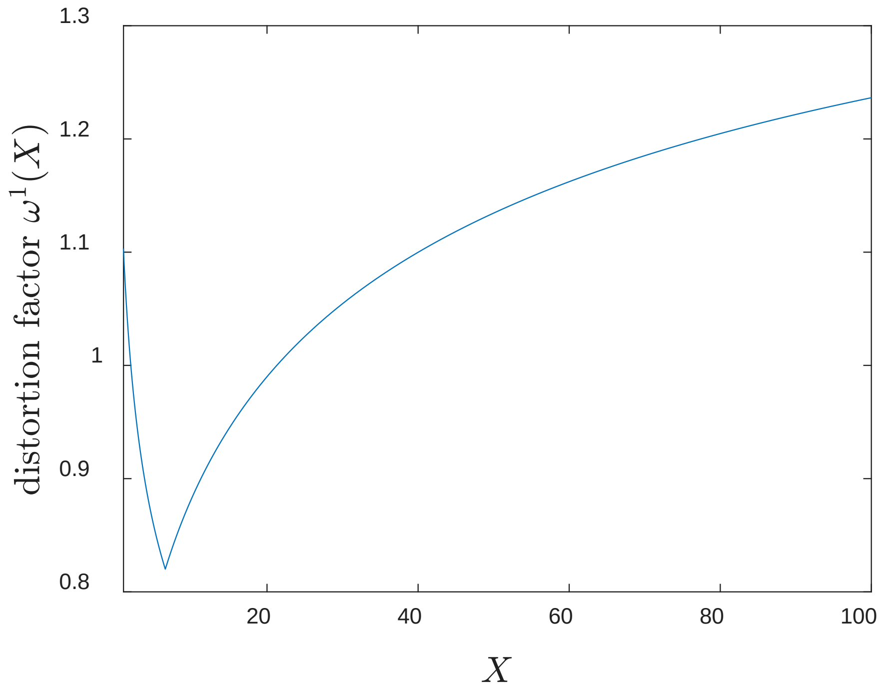

Figure C1. Distortion factor ω¹ for lottery *L*¹ as a function of payoff *X*. We plot the distortion factor $\omega^1(X)$ from equation (C8) as a function of X. The parameter values are: $\mu_x = 3.05$, $\mu_c = 2.35$, $\sigma_x = \sigma_c = 0.55$, $\rho = 0.5$, $\delta = 0.7$, and $\theta = 0.1$.

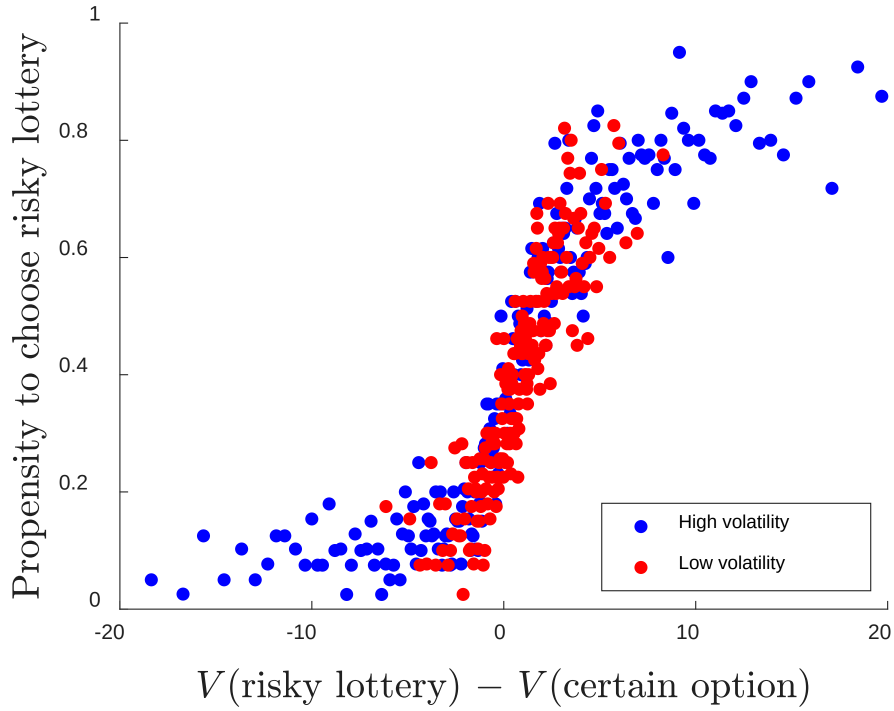

Figure C2. Probability of risk taking as a function of salience implied lottery valuations. We first compute the valuations of the risky lottery and certain option, *V*(risky lottery) and *V*(certain option), both under salience theory. We use the salience function in (C3) and the distortion factor in (C2), with $\delta = 0.7$, $\theta = 0.1$, and $s \in \{up, down\}$. We then plot the proportion of trials on which subjects choose the risky lottery as a function of *V*(risky lottery) – *V*(certain option), for both the high volatility condition ($\sigma_x = \sigma_c = 0.55$) and the low volatility condition ($\sigma_x = \sigma_c = 0.19$). Data are pooled across trials and subjects. For each of the two experimental conditions, we bin the *V*(risky lottery) – *V*(certain option) variable into two-hundred bins such that each bin has an equal number of trials. The figure is analogous to Figure 5 in the main text.

Appendix D: Experimental Instructions

D.1. Instructions for Risky Choice Task

Experiment Instructions

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

In 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 59% chance, and $0 with 41% chance. The amount shown for the sure thing will be paid with 100% chance, if chosen. Below is an example screen from the experiment:

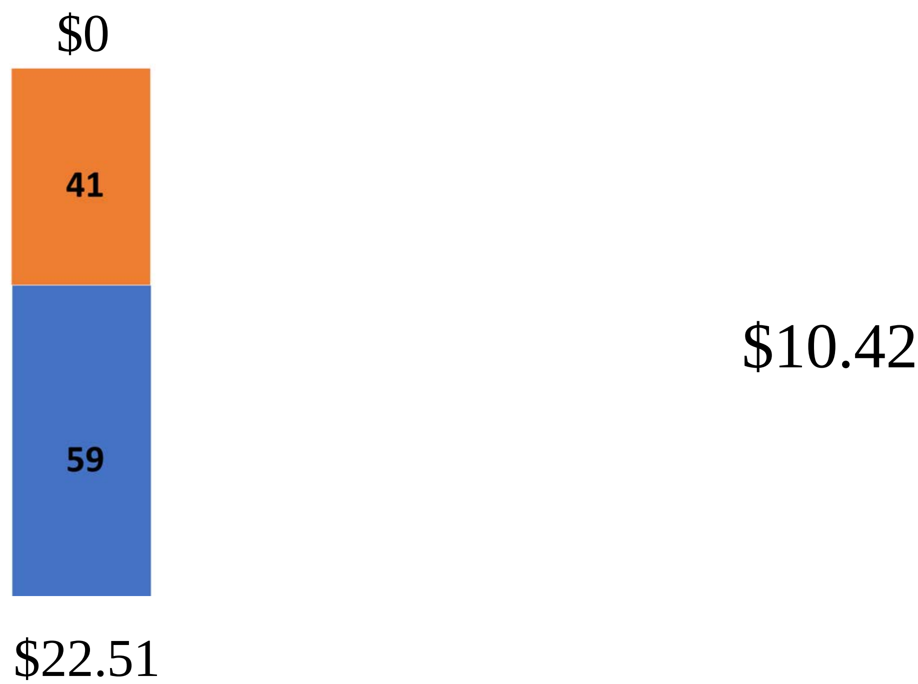

In this example, the risky gamble pays $22.51 with 59% chance, and $0 with 41% chance. The sure thing pays $10.42 with 100% chance. You will be asked to select one of the two options for each question in the experiment. The experiment is broken down into eight parts, and each part contains sixty questions.

At the end of the experiment, one trial will be randomly selected, and you’ll be paid according to your decision on that trial. For example, if the above trial was chosen, and you selected the sure thing you would be paid a total of $10.42 + $7 = $17.42. If instead you chose the risky gamble, you’d be paid either $7 or ($22.51 + $7) = $29.51, depending on which outcome the computer randomly selects. Before we begin, you will see 10 practice trials to familiarize yourself with the software. These 10 practice trials will not count towards the real experiment.

D.2. Instructions for Riskless Choice Task

Experiment Instructions

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

In the experiment, you will see a series of numbers and will be asked to classify whether the number is larger or smaller than the number "65". If the number is larger than 65, press the "?" key, and if it is smaller than 65, press the "z" key. At the end of the experiment, you will be paid depending on the speed and accuracy of your classifications. Specifically, you will be paid:

$$\text{Payout} = \$(20\times \mathit{accuracy} - 10\times \mathit{avgseconds}),$$

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 all trials and it took you 0.3 seconds to respond to each question, you would earn $\$(20\times 100\% - 10\times 0.3) = \$17.00$ (plus the \$7 show-up fee). If instead you only answer 75% of the questions accurately and took 1 second to respond to each question, you would be paid $\$(20\times 75\% - 10\times 1) = \$5.00$ (plus the \$7 show-up fee). 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 80 trials. In between each part, you can take a short (~1 minute) break, and then continue at your own pace. When you finish all sixteen parts, please raise your hand and do not disturb other subjects.

Before you begin the experiment, you will go through 10 practice trials to familiarize yourself with the software. These 10 practice trials will not be counted when computing your final payout.