Risk Aversion as a Perceptual Bias[*]

Mel Win Khaw
Columbia University

Ziang Li
Columbia University

Michael Woodford
Columbia University

March 20, 2017

Abstract

The theory of expected utility maximization (EUM) proposed by Bernoulli explains risk aversion as a consequence of diminishing marginal utility of wealth. However, observed choices between risky lotteries are difficult to reconcile with EUM: for example, in the laboratory, subjects' responses on individual trials involve a random element, and cannot be predicted purely from the terms offered; and subjects often appear to be too risk averse with regard to small gambles (while still accepting sufficiently favorable large gambles) to be consistent with any utility-of-wealth function. We propose a unified explanation for both anomalies, similar to the explanation given for related phenomena in the case of perceptual judgments: they result from judgments based on imprecise (and noisy) mental representation of the decision situation. In this model, risk aversion is predicted without any need for a nonlinear utility-of-wealth function, and instead results from a sort of perceptual bias — but one that represents an optimal Bayesian decision, given the limitations of the mental representation of the situation. We propose a specific quantitative model of the mental representation of a simple lottery choice problem, based on other evidence regarding numerical cognition, and test its ability to explain the choice frequencies that we observe in a laboratory experiment.

One of the most commonplace observations about economic life is that people often appear to be risk averse: they are unwilling to accept fair bets, and indeed pass up opportunities that would offer them a higher expected monetary reward for the sake of reduced uncertainty about the outcome. The standard explanation for such behavior, dating back to Bernoulli (1954 [originally 1738]), proposes that people do not choose so as to maximize their expected wealth, but instead to maximize expected utility, where utility is hypothesized to be a strictly concave (rather than linear) function of wealth. The consequences of this theory for choice under risk are now a staple element of undergraduate pedagogy, and the cornerstone of the modern theory of finance.

The theory of expected utility maximization (EUM), however, fails to account for a number of robust features of observed behavior, clearly documented by laboratory studies of choices between small monetary gambles.[1] For one, EUM implies that choice should be a deterministic function of the monetary payoffs offered and their associated probabilities. In the laboratory, instead, choices appear to be random, in the sense that the same subject will not always make the same choice when offered the same set of simple gambles on different occasions (Hey and Orme, 1994; Hey, 1995, 2001). This was evident (though little remarked upon) already in Mosteller and Nogee (1951), one of the earliest experimental studies of the empirical support for EUM. Figure 1 (reproduced from their paper) plots the responses of one of their subjects to a series of questions of a particular type. In each case, the subject was offered a choice of the form: are you willing to pay five cents for a gamble that will pay an amount x with probability 1/2, and zero with probability 1/2? The figure shows the fraction of trials on which the subject accepted the gamble, in the case of each of several different values of x. The authors used this curve to infer a value of x for which the subjects would be indifferent between accepting and rejecting the gamble, and then proposed to use this value of x to identify a point on the subject's utility function.

The fact that the indifference point is at a value of x greater than 10 cents (the case of a fair bet) is taken by the authors to indicate a concave utility function. But in fact, *no* utility function is consistent with the data shown in the figure; for EUM implies that the probability of acceptance should be zero for all values of x below the indifference point, and one for all values above it. Instead one observes probabilistic choice for a range of values of x, with the probability of acceptance increasing monotonically with x. This randomness is often de-emphasized in discussions of the experimental evidence for particular types of preferences over risky gambles, by simply focusing on modal or median responses. Studies that model the randomness in individual responses (e.g., Loomes and Sugden, 1995; Ballinger and Wilcox, 1997; Holt and Laury, 2002; Loomes, 2005; Wilcox, 2008) typically treat the randomness as something that can be specified independently of a "core" deterministic model of preference over lotteries (such as EUM), which is supposed to explain subjects'

---

[1]See Friedman *et al.* (2014) for an overview of the literature.

risk attitudes. Here we propose instead that the randomness of responses may be intimately connected with apparent risk attitude.

A second problem, stressed by Rabin (2000), is that people are often much too averse to small bets for this to be consistent with EUM, under any reasonable utility function. For example, if we consider only the "indifference point" indicated in Figure 1, it implies that in the case of a gamble that pays off with probability 1/2, the subject requires a possibility of gaining more than 5.5 cents to be willing to risk losing 5 cents. This would require a utility-of-wealth function $U(W)$ with the property that $U(W + 5.5) - U(W) < U(W) - U(W - 5)$, where W is the subject's existing wealth (in cents) at the time of considering the gamble. But a utility function with this property for all values of W would be one that also implied that the subject should decline a bet offering a 50 percent chance of gaining unbounded wealth, if there were also a 50 percent chance of losing 84 cents[2] — a highly implausible degree of risk aversion.[3]

We propose that both of these puzzles for the standard theory have a single explanation. Our proposed explanation relies on an analogy between judgments about the value of a risky prospect and perceptual judgments in various sensory domains. In fact, both of the phenomena just discussed are analogous to much-studied features of perceptual judgments. First, it has been extensively documented, since the classic work of Weber (1978 [originally 1834]), that judgments about the relative magnitudes of sensory stimuli are random, rather than perfectly consistent across repetitions using identical stimuli. It is common to plot the probability of a judgment that one stimulus is greater in magnitude (heavier, longer, brighter, etc.) than the other as a function of the true difference in their magnitudes, and obtain a continuously increasing function (called a "psychometric function") rather than a step function,[4] like the function shown in Figure 1. Indeed, it was doubtless due to familiarity with such figures in the literature on sensory perception that Mosteller and Nogee saw little need to remark on the stochasticity of their data, and found it natural to plot their data in the way that they did.[5]

Perceptual judgments are also subject to many well-documented systematic biases. For example, the orientations (degree of tilt) of tilted stimuli are not judged accurately, especially when the stimuli are blurry; not only are estimates of the angle

---

[2] The calculations required to show this are explained in Rabin (2000), though not using these particular numerical values.

[3] We do not, of course, know how Mosteller and Nogee's subject would behave when presented with larger bets, and there has been some dispute about whether people actually display the kind of paradoxical behavior discussed by Rabin (2000). See Cox *et al.* (2013) for examples of experiments in which subjects make choices with respect to both small and large bets that are inconsistent with EUM under any possible concave utility function.

[4] See, for example, Gabbiani and Cox (2010), chap. 25; Gescheider (1997), chap. 3; or Glimcher (2011), chap. 4.

[5] Note that the method used by Mosteller and Nogee to identify an indifference point for their subject corresponds to a standard way of defining a "point of subjective equality" between two different types of sensory stimuli using a psychometric function (see, e.g., Gescheider, 1997, p. 52).

Figure 1: Probability of acceptance of a simple gamble as a function of the amount x that can be won (shown on horizontal axis); from Mosteller and Nogee (1951).

of orientation variable from trial to trial, but they are not correct on average. Instead, orientations are perceived as being farther from vertical (if near-vertical) and farther from horizontal (if near-horizontal) than they truly are, an "oblique bias" that has been noted since Jastrow (1892). Such biases in average judgment can furthermore be attributed, at least in some cases, to the same source as the random variation in judgments: the fact that judgments must be based on an effect that the stimulus has on the subject's nervous system, and that this effect is a random function of the objective properties of the stimulus. When judgments must be based on a subjective representation of the stimulus which involves random noise, then an optimal rule for forming such judgments (one that minimizes some measure of average distortion of the resulting judgments, subject to the constraint that the judgment can only be based on the subjective representation) will generally involve an average bias, as we explain further below. Hence the existence of the bias is actually adaptive, given unavoidable noise in the underlying sensory data on which judgments must be based.

We propose that risk aversion — at least, the kind of risk aversion that is observed in choices with regard to small gambles (as opposed to risks that would make a significant difference for one's standard of living) — can be viewed as a bias of a similar

3

kind.[6] According to our proposal, intuitive estimates of the value of risky prospects (not ones based on symbolic calculations) are based on mental representations of the magnitudes of the available monetary payoffs that are imprecise in roughly the same way that the representations of sensory magnitudes are imprecise, and in particular are similarly random, conditioning on the true payoffs. Intuitive valuations must be some function of these random mental representations. We explore the hypothesis that they are produced by a decision rule that is optimal, in the sense of maximizing the (objective) expected value of the decision maker's expected wealth, subject to the constraint that the decision must be based on the random mental representation of the situation.

Under a particular model of the noisy coding of monetary payoffs, we show that this hypothesis will imply apparently risk-averse choices, in the sense that the expected net payoff of a bet will have to be strictly positive for indifference, in the sense that the subject accepts the bet exactly as often as she rejects it (as in Figure 1). Risk aversion in this sense is consistent with a decision rule that is actually optimal from standpoint of an objective (expected wealth maximization) that involves no "true risk aversion" at all; this bias is consistent with optimality in the same way that perceptual biases (such as the oblique bias in the perception of orientation) can be consistent with Bayesian inference from noisy sensory data. And not only can our theory explain apparent risk aversion without any appeal to diminishing marginal utility, but it can also explain why the "risk premium" required in order for a risky bet to be accepted over a certain payoff does not shrink to zero (in percentage terms) as the size of the bet is made small, contrary to the prediction of EUM.

Section 1 reviews some of the results from perceptual psychology and neuroscience that motivate our theoretical proposal, including evidence relating to the mental representation of numerical magnitudes. Section 2 presents an explicit model of choice between a simple risky gamble and a certain monetary payoff, of the kind that occurs in the experiment of Mosteller and Nogee, and derives predictions for the both the randomness of choice and the degree of apparent risk aversion implied by an optimal decision rule. Section 3 describes a simple experiment in which we are able to test some of the specific quantitative predictions of this model. Section 4 discusses further implications of our theory of risk attitudes, and concludes.

---

[6]Our theory is thus similar to proposals in other contexts (such as Koszëgi and Rabin, 2008) to interpret experimentally observed behavior in terms of mistakes on the part of decision makers — i.e., a failure to make the choices that would maximize their true preferences — rather than a reflection of some more complex type of preferences. More specifically, we follow Woodford (2012), Steiner and Stewart (2016), and Gabaix and Laibson (2017) in proposing that choice biases can reflect optimal Bayesian decision making on the basis of a noisy representation of the decision problem.

1 Noisy Coding and Biased Judgments

We begin by discussing the way in which even optimal judgments will almost inevitably be biased, when they must be based on a noisy mental representation of the decision maker's current situation, and illustrate the application of this idea to explaining perceptual biases. We also review evidence as to the nature of mental representations of numerical magnitudes, in order to motivate a particular model of noisy coding of monetary payoffs.

1.1 Noisy Coding of Sensory Magnitudes

An important recent literature on the neuroscience of perception has argued that biases in perceptual judgments may actually reflect optimal decisions — in the sense of minimizing average error, according to some well-defined criterion, in a particular class of situations that are possible ex ante — given the constraint that the brain can only produce judgments based on the noisy information provided to it by sensory receptors and earlier stages of processing in the nervous system, rather on the basis of direct access to the true physical properties of external stimuli (e.g., Stocker and Simoncelli, 2006; Wei and Stocker, 2015). As an illustration, we briefly discuss a proposed explanation for the "oblique bias" in estimates of the orientation of visual stimuli, already mentioned in the introduction.

The fact that judgments of the orientation of visual stimuli are based on a noisy internal representation of this stimulus feature is indicated by a variety of types of evidence. First, judgments of orientation (both estimates of the orientation of a single stimulus, and comparative judgments of the relative orientation of two stimuli that are presented together) are random; a given subject will not give an identical response each time the same stimulus is presented. This has been thought, since the work of Fechner (1966 [originally 1860]), to reflect randomness in the effect of the physical stimulus on the subject's nervous system. More recently, the celebrated work of Hubel and Wiesel (1959) on the way that differently oriented stimuli affect the production of electrical spikes by neurons in the V1 visual area of the cerebral cortex of the cat launched a large literature that has shown how the pattern of neural activity in this region encodes the orientation of a visual stimulus. What experiments show is that individual neurons produce spikes at a particular average rate per unit time, depending on the orientation of the stimulus. The representation of the stimulus orientation provided by the pattern of neural activity is however a random, rather than a deterministic function of the true orientation, as the spikes produced by each neuron over a given time interval are a random variable (roughly a Poisson variable, with a rate that depends on the orientation in a way summarized by the neuron's "tuning curve").

The degree of noise in this neural representation of orientation can furthermore account for the degree of randomness that is observed in orientation judgments. A

common measure of the accuracy of perceptual judgments is the "discrimination threshold" — the amount by which one stimulus must differ from another (e.g., the number of degrees by which the second must be tilted counterclockwise relative to the first) in order for a subject to correctly judge that the second magnitude is greater than the first (is in fact counterclockwise relative to the first) at least 75 percent of the time. The discrimination threshold should be larger the greater the degree of randomness in the internal representation of each of the stimuli. In the case of orientation, it is well-established both in humans and a number of other animals, that judgments of orientation are more accurate in the case of near-cardinal orientations (i.e., orientations that are nearly vertical or horizontal) than in the case of more oblique orientations; measured discrimination thresholds are about twice as large in the case of maximally oblique orientations as in the case of near-cardinal orientations.[7] This can be accounted for by differences in the degree of randomness of the neural representation of orientation in area V1 in the case of the different orientations that is implied by the degree of inhomogeneity in the number and in the precision of the tuning of the neurons that respond to orientations in different ranges.[8]

This inhomogeneity in the degree of imprecision of the neural coding of stimulus orientation also makes it unsurprising that there should be systematic biases in judgments — because it is actually optimal for such judgments to be biased, given that they must be based on a noisy internal representation of the kind that can not only be shown to exist (by direct measurement of the neural activity associated with processing of particular stimuli) but that is needed in order to account for the degree of imprecision in judgments of orientation. A simple calculation (based on Wei and Stocker, 2015) can illustrate the connection.

Let the true stimulus magnitude be measured by a real number θ, and let the internal representation of the stimulus magnitude also be summarized by a single real number r.[9] Suppose furthermore that conditional on the value of θ, the internal representation is a random draw from a probability distribution

$$r \sim N(F(\theta), \nu^2),$$

where $F(\theta)$ is a smooth, monotonically increasing function that maps the real line onto itself, and the parameter ν^2 measures the degree of noise in the internal repre-

---

[7]See Appelle (1972) for an early review, and Girshick *et al.* (2011) for a recent discussion.

[8]See Girshick *et al.* (2011) and Ganguli (2012) for references on the neurophysiology of orientation perception, and for quantitative demonstrations that a theoretical model of discrimination based on noisy neural coding, parameterized to match the measured distribution of neural tuning curves, can account well for the relative size of measure discrimination thresholds across different orientations.

[9]Here for convenience we let the state θ be a value on the real line, following the exposition by Wei and Stocker, even though orientation should really take values on a circle. Note that this allows us to use the mathematics of normal random variables in the discussion below, and also provides an introduction to the calculations needed below for our discussion of the coding of numerical information. Note also that it is not necessary for the conclusions reached below that the internal representation be summarized by a single number; Wei and Stocker offer a more general treatment, which is simplified here.

sentation. If we suppose that the magnitude of stimulus θ_2 is perceived as greater than that of stimulus θ_1 if and only if the corresponding internal representations satisfy $r_2 > r_1$, then the probability of such a judgment will be an increasing function of $[F(\theta_2) - F(\theta_1)]/\nu$. Hence the discrimination threshold near a given magnitude will be inversely proportional to the rate of increase of $F(\theta)$ near that value of θ. We can account for the facts cited above about the accuracy of orientation discrimination if we assume that $f(\theta) \equiv F'(\theta)$ is a periodic, positive-valued function, with $f(\theta + \pi/2) = f(\theta)$ for all θ, achieving its maximum for values of θ near integral multiples of $\pi/2$ (cardinal orientations), and its minimum for values that are near odd multiples of $\pi/4$ (maximally oblique orientations).

Let us next consider the estimate of the stimulus magnitude that a subject should make, if it must be based on the internal representation r. The mean squared error will be minimized if the estimate is given by the Bayesian posterior mean,[10]

$$\hat{\theta}(r) \equiv \mathrm{E}[\theta|r],$$

computed using a prior distribution over values of θ. Girshick *et al.* (2011) show (by numerical analysis of the edges contained in a database of natural scenes) that horizontal and vertical orientations occur more often than do oblique orientations; in fact, they show that the frequency of occurrence of different orientations in the world is roughly inversely proportional to measured discrimination thresholds. We can therefore follow Wei and Stocker, and assume a prior for θ with a density function given by $f(\theta)$, the function used above to characterize the inhomogeneity in the discriminability of nearby orientations.[11] This means that if we work in terms of the transformed state variable $\tilde{\theta} \equiv F(\theta)$, the prior distribution for $\tilde{\theta}$ is uniform.

In this case, the estimate implied by an internal representation r will be given by

$$\hat{\theta}(r) = \mathrm{E}[F^{-1}(r + \epsilon)],$$

where the expectation is over realizations of ϵ, an independent draw from the distribution $N(0, \nu^2)$. It then follows that for any true magnitude θ, the mean estimate should be given by

$$\mathrm{E}[\hat{\theta}|\theta] = g(\theta) \equiv \mathrm{E}[F^{-1}(F(\theta) + \tilde{\epsilon})], \tag{1.1}$$

where $\tilde{\epsilon} \sim N(0, 2\nu^2)$.

An optimal estimate need not be unbiased: in general, (1.1) implies that $\mathrm{E}[\hat{\theta}|\theta] \neq \theta$.[12] In particular, our assumptions above about $f(\theta)$ imply that $F^{-1}(\tilde{\theta})$ will have the smallest slope at values of $\tilde{\theta}$ corresponding to cardinal orientations, and the greatest

---

[10]Qualitatively similar biases, though not quantitatively identical, can be predicted under other assumptions about how a subject's estimate is derived from the internal representation, as shown by Girshick *et al.* (2011).

[11]If θ takes values on the real line, this implies an improper prior. However, the posterior implied by any representation r remains well-defined.

[12]Bayesian rationality requires that $\mathrm{E}[\theta|\hat{\theta}] = \hat{\theta}$, but not that $\mathrm{E}[\hat{\theta}|\theta] = \theta$.

slope at values corresponding to maximally oblique orientations. Thus for values of $\tilde{\theta}$ near some cardinal orientation $\tilde{\theta}^* \equiv F(\theta^*)$, $F^{-1}(\tilde{\theta})$ will have a decreasing slope (will be concave) for $\tilde{\theta} < \tilde{\theta}^*$, and will have an increasing slope (will be convex) for $\tilde{\theta} > \tilde{\theta}^*$. By Jensen's inequality, (1.1) then implies that $\mathrm{E}[\hat{\theta}|\theta] < \theta$ for $\theta < \theta^*$, while $\mathrm{E}[\hat{\theta}|\theta] > \theta$ for $\theta > \theta^*$, at least in the case of a small enough noise variance ν^2.

Estimates are thus predicted to be biased away from the cardinal orientations; a similar argument can be made near any of the maximally oblique orientations, to show that estimated orientations should be biased toward the maximally oblique orientations. Thus this theory predicts exactly the kind of "oblique bias" that had been noted since Jastrow (1892). Tomassini *et al.* (2010) and Wei and Stocker (2015) show that Bayesian models of this kind, calibrated to match the measured degree of inhomogeneity in the neural coding of orientation in area V1, can account quantitatively for the size of experimentally measured biases in estimates of orientation. This explanation implies that the existence of the biases is intimately bound up with the noise in the internal representation, which also results in randomness of estimates and discriminations. In fact, experiments also show, as the theory would predict, that when the randomness of the internal representation is increased (by presenting lower-contrast stimuli, so that the orientation is perceived less sharply), the size of the oblique bias increases (De Gardelle, 2010).

It should be clear that in a theory of this kind, the predicted biases in judgment depend crucially on the nature of the inhomogeneity in the precision with which magnitudes are mentally coded, over different ranges of value for the magnitude in question. (Equation (1.1) only predicts a biased estimate to the extent that $F(\theta)$ is non-linear, which is to say, to the extent that $f(\theta)$ is non-uniform.) Hence we must consider what it would be reasonable to assume about the imprecision in mental coding of the quantities that are relevant to economic decisions like the one faced by the subjects of Mosteller and Nogee (1951). In this regard, it will be useful to review what is known about imprecision in the mental representation of numbers.

1.2 Noisy Coding of Numerical Magnitudes

The relevance of these observations about perceptual judgments for economic decision making might be doubted. Some may suppose that the kind of imprecision in mental coding just discussed matters for the way in which we perceive our environment through our senses, but that an intellectual consideration of hypothetical choices is an entirely different kind of thinking. Moreover, it might seem that typical decisions about whether to accept gambles in a laboratory setting, such as the experiment of Mosteller and Nogee (1951), involve only numerical information that is presented to the subjects in an exact (symbolic) form, offering no obvious opportunity for imprecise perception. However, we have reason to believe that reasoning about numerical information often involves imprecise mental representations of a kind directly analogous to those involved in sensory perception.

Despite an inability to use symbols or language, both animals and human infants are able to perceive the number of items in a set, and to carry out simple calculations — for example, not only estimating whether one set has more items than another, but whether the sum of the numbers of items in two sets is greater than the number of items in some third set (Dehaene, 2011, chaps. 2 and 3). The same is true of Amazonian tribesmen whose language lacks words for numbers larger than five, and have no training in arithmetic calculations (Pica *et al.*, 2004). It appears that they use a non-verbal system — a so-called "number sense" — which represents numbers analogically and approximately, rather than exactly, in a way that resembles the mental representation of sensory magnitudes (Dantzig, 2007; Dehaene, 2011). This same capacity appears also to be used by adults from cultures with complex number systems and formal training in arithmetic, under certain circumstances, such as when required to estimate numerical magnitude and undertake numerical reasoning on the basis of numerical information that is not presented symbolically.

For example, people are able to estimate the number of dots present in a visual display of a random cloud of dots, without counting them.[13] The estimate given for any presented visual array is random, just as with estimates of sensory magnitudes such as length or orientation. And just as in the case of sensory magnitudes, the randomness in judgments can be attributed to randomness in the neural coding of numerosity, resulting from the width of the "tuning curves" of neurons that selectively respond to arrays with greater or smaller numbers of dots.[14]

We can learn about how the degree of randomness of the mental representation of a number varies with its size from the frequency distribution of errors in estimation of numerosity. A well-established finding is that when subjects must estimate which of two numerosities is greater, or whether two arrays contain the same number of dots, the accuracy of their judgments is a function of the ratio of the two numbers (but independent of their absolute magnitudes) — a "Weber's Law" for the discrimination of numerosity analogous to the one observed to hold in many sensory domains (Ross, 2003; Cantlon and Brannon, 2006; Nieder and Merten, 2007; Nieder, 2013). Moreover,

---

[13]Of course, adults from numerate cultures are able to count large numbers of dots, one by one, when given sufficient time and willing to exert the effort. However, they may also form an estimate using the "number sense" that they share with infants and animals, when time is short or precision is not too important, in order to economize on cognitive effort. For example, when Dewan and Neligh (2017) allow their subjects to spend as long as they like before announcing their estimate of the number of dots present, they find that some subjects appear to switch their cognitive strategy depending on the monetary incentive for a correct answer — taking much more time and making many fewer errors (presumably by counting the individual dots) when the incentive exceeds a certain threshold, answering more quickly but with more errors (while still doing better than by pure guessing) when the incentive is lower.

[14]The tuning curves of "number neurons" have been measured using single-cell recording techniques in the brains of both cats and macaques (Thompson *et al.*, 1970; Nieder and Merten, 2007; Nieder and Dehaene, 2009). While similar methods cannot be used with humans, more indirect evidence suggests the existence of "number neurons" in the human brain as well (Piazza *et al.*, 2004; Nieder, 2013).

when subjects must report an estimate of the number of dots in a visual array, the standard deviation of the distribution of estimates grows in proportion to the mean estimate, with both the mean and standard deviation being larger when the true number is larger (Izard and Dehaene, 2008; Kramer *et al.*, 2011); and similarly, when subjects are required to produce a particular number of responses (without counting them), the standard deviation of the number produced varies in proportion to the target number (and to the mean number of responses produced) — the property of "scalar variability" (Whalen *et al.*, 1999; Cordes *et al.*, 2001).

All of these observations are consistent with a theory according to which such non-symbolic computations are based on a "semantic" mental representation of numbers which is stochastic, and can be represented mathematically by a quantity that is proportional to the logarithm of the numerical value that is being encoded, plus a random error the variance of which is independent of the numerical value that is encoded (van Oeffelen and Vos, 1982; Izard and Dehaene, 2008).[15] Let the number n be represented by a real number r that is drawn from the distribution $N(\log n, \nu^2)$, where ν^2 is a parameter independent of n. Then the degree of overlap between the distributions of possible subjective representations (which should determine the frequency of errors in telling the two numbers apart) depends on the difference of their logarithms, or equivalently on the ratio of the two numbers.

If an estimate of the number must be produced based on this subjective representation, and the estimate $\hat{n}$ is a number whose logarithm is an affine function of r,[16] then $\hat{n}$ will be log-normally distributed; specifically, $\log \hat{n}$ will have the distribution $N(\hat{\mu}(n), \hat{\sigma}^2)$, where $\hat{\mu}(n)$ is an affine function of $\log n$, and $\hat{\sigma}^2$ is independent of n. It thus follows that the estimate $\hat{n}$ will be drawn from a distribution with mean and standard deviation

$$\mathrm{E}[\hat{n}|n] \;=\; m(n) \;\equiv\; e^{\hat{\mu}(n)+(1/2)\hat{\sigma}^2}, \qquad \mathrm{SD}[\hat{n}|n] \;=\; m(n) \cdot \sqrt{e^{\hat{\sigma}^2} - 1}$$

respectively. Hence the standard deviation is a positive multiple of the mean, and the property of scalar variability is verified as well.

Human subjects' estimates of numerosity are not only variable, but generally biased as well (in the sense that $m(n) \neq n$). The mean estimate $m(n)$ is often found to be well fit (for numbers of dots greater than five[17]) by a concave power law; that

---

[15]Buckley and Gillman (1974) had earlier proposed a similar model to explain behavior in experiments involving magnitude comparisons between numbers represented by Arabic numerals; these related experiments are discussed below.

[16]We provide a justification for an estimation rule of this form below. Note that we abstract from the requirement that the estimate be an integer, in order to simplify our calculations, which should be regarded as only an approximation to the predictions of a more exact model of numerosity estimation, like the ones presented by van Oeffelen and Vos (1982) and Izard and Dehaene (2008).

[17]Authors beginning with Kaufman *et al.* (1949) have argued that in the case of visual arrays, a distinct cognitive process, "subitizing," is used to quickly apprehend the number of dots; this process is faster, more accurate, and allows greater confidence than the method of estimation that must be used for larger numbers of dots, which the model of logarithmic coding presented in the text is intended to describe.

is,

$$m(n) = An^{\beta} \tag{1.2}$$

for some $A > 0, 0 < \beta < 1$ (Krueger, 1972, 1984; Indow and Ida, 1977). It follows that the number of dots is over-estimated, on average, in the case of small enough numbers of dots (no more than 10, in the studies of Kaufman *et al.* (1949) and Indow and Ida (1977); less than 25, in the studies reviewed by Krueger (1984); but all numbers less than 130 dots, in the study of Hollingsworth *et al.* (1991)), and instead under-estimated on average when the number of dots is larger.

The model of noisy logarithmic coding of numerosity also predicts this type of bias, under a simple hypothesis about how estimates are produced. Suppose that the estimate $\hat{n}$ must be based on the mental representation r of the number of dots in a visual array. If we suppose further that the decision rule $\hat{n}(r)$ is the one that minimizes the mean squared error of the estimate (that is, that minimizes $\mathrm{E}[(\hat{n} - n)^2]$), given some prior distribution over the possible numbers of dots that might be presented, then it follows that the decision rule should be given by $\hat{n}(r) = \mathrm{E}[n|r]$. That is, the estimate $\hat{n}$ should correspond to the posterior mean of the possible values for n, where the posterior can be derived from the prior distribution and the likelihood function implied by the noisy logarithmic coding, using Bayes' rule.[18]

If the prior is also log-normal ($\log n \sim N(\mu, \sigma^2)$), then under the above model of noisy coding, the posterior distribution corresponding to any mental representation r is a log-normal distribution, $\log n|r \sim N(\mu_{post}(r), \sigma^2_{post})$, where

$$\mu_{post}(r) \equiv (1 - \beta)\mu + \beta r, \qquad \sigma^2_{post} \equiv (1 - \beta)\sigma^2, \tag{1.3}$$

in which expressions

$$\beta \equiv \frac{\sigma^2}{\sigma^2 + \nu^2}, \tag{1.4}$$

so that $0 < \beta < 1$. It then follows that the optimal (minimum-MSE) decision rule is given by

$$\hat{n}(r) = \mathrm{E}[n|r] = e^{\mu_{post}(r)+(1/2)\sigma^2_{post}},$$

so that $\log \hat{n}(r)$ is an affine function of r, $\alpha + \beta r$, where

$$\alpha \equiv (1 - \beta) \left[\mu + \frac{1}{2}\sigma^2 \right] \tag{1.5}$$

and β is defined in (1.4).

---

[18]As in our discussion above of a Bayesian model of judgments of orientation, such a hypothesis does not imply that subjects in numerosity estimation experiments consciously calculate anything using Bayes' rule; only that, in some way or another, their intuitive judgments have come to be calibrated so as to be optimal for a certain environment.

We thus obtain a quantitative model of magnitude estimation consistent with the property of scalar variability. Moreover, the predicted average estimate, conditional on the true number of dots presented, is given by

$$m(n) \;=\; \mathrm{E}[\hat{n}(r)|n] \;=\; e^{\alpha + \beta \log n + (1/2)\beta^2 \nu^2},$$

which is a power law of the form (1.2), with

$$A \;\equiv\; e^{\alpha + (1/2)\beta^2 \nu^2}$$

and an exponent β given by (1.4). Thus the model can explain the observation of a power law, and both the over-estimation of small numerosities and the under-estimation of large ones.

More precise quantitative predictions are obtainable only with some independent source of evidence as to the appropriate parameterization of the prior. We shall not make a proposal about this, but note that the fact that the values of both A and β in (1.2) are predicted to depend on the prior provides a possible explanation for the differing findings of different studies with regard to the quantitative specification of the $m(n)$ curve. Izard and Dehaene (2008) point out that the study of Hollingsworth *et al.* (1991), in which $m(n) > n$ was found for much larger values of n than in the classic earlier studies, was also a study in which subjects were presented with a range of numerosities including larger values (arrays containing up to 650 dots), and suggest that "the estimation pattern seems to be influenced by the range of stimuli tested" (p. 1222).[19]

This is exactly what our theory of optimal estimation based on a noisy mental representation would imply: if the prior to which the subject's estimation rule is adapted is determined by the frequency with which different numbers of dots are presented in the experiment, then an experiment in which numbers are drawn from a log-normal distribution with a larger value of μ (but the same value of σ) should result in a shift up of the $m(n)$ curve (with $\log m(n)$ increased for each n by an amount proportional to the increase in μ), and a corresponding increase in the critical number at which the sign of the estimation bias changes. Our theory also predicts, for a given prior, that increased imprecision in mental coding (a larger value of ν^2) should result in a lower value of β, and hence a more concave relationship between the true and estimated numerosities; and this is what Anobile *et al.* (2012) find when subjects' cognitive load is increased, by requiring them to perform another perceptual classification task in addition to estimating the number of dots present.

---

[19]This can be clearly seen in the results of Anobile *et al.* (2012) for numerosity estimation under conditions of increased cognitive load. These authors required subjects to report their estimate of the numerosity of a field of dots by moving a slider on a number line that ranged either from 1 to 10, from 1 to 30, or from 1 to 100, with corresponding variation in the range of numbers of dots presented under the three conditions; the number at which subjects switched from over-estimation to under-estimation progressively increased as the allowed scale of responses (and the range of numbers of dots actually presented) increased — see panel B of their Figure 3.

The evidence just summarized relates to estimation of the number of dots in a visual array, and other non-symbolic presentations of numerical information. It might be thought that even if a "number sense" analogous to the faculties used in sensory perception is used in such cases, this would have no obvious implication for the cognitive processes involved with numerical information is presented using number words and symbols, as in the classic experiments of Mosteller and Nogee (1951), Kahneman and Tversky (1979), or the experiments described below. However, there is a significant body of evidence suggesting that a faculty that allows simple arithmetic judgments to be made on the basis of approximate, non-symbolic number representations is also used by numerate adults, even when answering certain kinds of questions using information that has been presented verbally or symbolically (Dehaene, 2011, chaps. 3, 5, 10).

For example, Moyer and Landauer (1967) presented subjects with two numerals, and required them to press one of two keys to indicate which numeral indicated the larger number. They found that both the fraction of incorrect responses and the time required to decide were decreasing functions of the numerical distance between the two numbers referred to by the numerals; these findings are analogous to the way that both error rates and response times vary with the magnitude difference between two sensory stimuli in experiments where a subject must determine which of two stimuli is greater in magnitude (the louder sound, the longer line, and so on). Moyer and Landauer conclude that "the displayed numerals are converted [by the mind] to analogue magnitudes, and a comparison is then made between those magnitudes in much the same way that comparisons are made between physical stimuli" (p. 1520).

Exact arithmetic calculations appear to use a distinct brain system (connected to language processing) from that used for approximate magnitudes. For example, Dehaene and Cohen (1991) report a patient with brain damage who was severely impaired at finding exact solutions to even the simplest problems (the patient judged 2+2=3 and 2+2=4 to be equally plausible), but who could nonetheless perform approximate calculations (and so could readily reject 2+2=9 as not plausible). Bilingual adults typically perform exact arithmetic calculations (including exact counting) only in one language (the one in which they were taught arithmetic); but Spelke and Tsivkin (2001) found that bilingual students were able to recall information about approximate numbers equally rapidly in either language (while being faster in their native language in the case of exact numerical information), suggesting that approximate quantitative information (even when presented verbally) is represented in some analog fashion, not tied to particular symbols.

Moreover, there is evidence that the mental representation of numerical information used for approximate calculations involves the same kind of logarithmic compression as in the case of non-symbolic numerical information, even when the numerical magnitudes have originally been presented symbolically. Moyer and Landauer (1967), Buckley and Gillman (1974), and Banks et al. (1976) find that the reaction time required to judge which of two numbers (presented as numerals) is larger varies with

the distance between the numbers on a compressed, nonlinear scale — a logarithmic scale, as assumed in the model of the coding of numerosity sketched above, or something similar — rather than the linear (arithmetic) distance between them.[20] In an even more telling example, Dehaene and Marques (2002) showed that in a task where people had to estimate the prices of products, the estimates produced exhibited the property of scalar variability, just as with estimates of the numerosity of a visual display, even though both the original information people had received about prices and the responses they produced involved symbolic representations.[21]

Our hypothesis is that when people make judgments about whether a risky prospect (offering either of two possible monetary amounts as the outcome) is worth more or less than another monetary amount (that could be obtained with certainty), they use the same mental faculty as is involved in judging whether the sum of two numbers is greater or less than some other number. If this is approached as an approximate judgment rather than an exact calculation (as will often be the case, even with numerate subjects), such a judgment is made on the basis of mental representations of the monetary amounts that are approximate and analog, rather than exact and symbolic; and these representations involve a random location of the amount on a logarithmically compressed "mental number line."[22] We shall show that the fact the judgment must be based on imprecise representations of this kind can account for both the randomness of the choices made by experimental subjects, and the fact that they appear to be surprisingly risk averse even when offered very small gambles.[23]

---

[20]Buckley and Gillman (1974) develop an extension of the model of noisy logarithmic coding of numerical magnitudes sketched above that explicitly models the dynamic process of comparison between two magnitudes, and show that the model predicts not only that the frequency of correct ranking should depend on the ratio of the two numbers (as discussed above) but that the mean time required to decide should depend on this ratio as well, as they find in their experiment. (See also Dehaene, 2008, for a related model.) The dynamic version of the model is not needed for our purposes here.

[21]This example is of particular relevance for our purposes, as it involves the mental representation of monetary amounts.

[22]Note that we do not assume that all decisions involving money are made in this way. If someone is asked to choose between $20 and $22, either of which can be obtained with certainty, we do not expect that they will sometimes choose the $20, because of noise in their subjective sense of the size of these two magnitudes. The question whether $20 is greater or smaller than $22 can instead be answered reliably (by anyone who remembers how to count) using exact knowledge about symbolic quantities.

[23]Schley and Peters (2014) also propose that a compressive nonlinear mapping of symbolically presented numbers into mental magnitudes can give rise to additional risk aversion, alongside the risk aversion that can be attributed to diminishing marginal utility; but as we discuss in section 4.2 below, their theory differs in important respects from the one that we propose here.

2 A Model of Noisy Coding and Risky Choice

We now consider the implications of the model of noisy mental representation of numerical magnitudes sketched above for choices between simple lotteries, of the kind that subjects are presented with in experiments like that of Mosteller and Nogee (1951). We assume a situation in which a subject is presented with a choice between two options: receiving a monetary amount $C > 0$ with certainty, or receiving the outcome of a lottery, in which she will have a probability $0 < p < 1$ of receiving a monetary amount $X > 0$. We wish to consider how decisions should be made if they must be based on imprecise mental representations of the quantities X and C rather than their exact values.

In line with the evidence discussed in the previous section regarding the mental representation of numerical magnitudes, we assume that the quantities X and C have mental representations r_x and r_c respectively, each a random draw from a probability distribution of possible representations, with distributions

$$r_x \sim N(\log X, \, \nu^2), \qquad r_c \sim N(\log C, \, \nu^2). \tag{2.1}$$

Here $\nu^2 > 0$ is a parameter that measures the degree of imprecision of the mental representation of such quantities (assumed to be the same regardless of the monetary amount that is represented); we further assume that r_x and r_c are distributed independently of one another. We wish to consider possible decision rules that specify whether the subject should choose the risky lottery or the certain payment, on the basis of the mental representations r_x and r_c specifying the available options on the particular occasion. We treat the parameter p as known (it does not vary across trials in the experiment described below), so that the decision rule can (and indeed should) depend on this parameter as well.[24]

In order to determine an optimal decision rule, it is necessary to specify what is to be optimized. We assume an objective of maximizing the mathematical expectation of the subject's wealth; thus there is assumed to be no true risk aversion (resulting from diminishing marginal utility of additional wealth, as proposed by Bernoulli), at least in the case of modest gambles of the size with which we are concerned. This assumption has the advantage of allowing us to show that apparently risk averse choices can be justified even in the absence of diminishing marginal utility. It also allows us to make much sharper quantitative predictions: there are no free parameters associated with the specification of a nonlinear utility function, and the predicted probability of choosing the risky lottery is independent of contextual variables such as the subject's existing wealth or other sources of income risk in the subject's life.

---

[24]We leave the implications of imprecise mental representation of probabilities for future work. An assumption that p also has a noisy mental representation, and that the decision rule can be based only on that representation, would affect the precise formulas describing both the bias and randomness associated with an optimal decision rule; but it would not change our conclusion that an optimal rule implies both bias and randomness, nor would it lessen the importance of the consequences (for both bias and randomness) of the logarithmic coding of monetary payoffs treated here.

In order to determine the optimal decision in the case of a pair of mental representations $\mathbf{r} = (r_x, r_c)$, it is also necessary to specify the set of possible decision situations in which those representations can have arisen. Note that our theory is not one in which the mental representation r_x of some amount of money represents a particular belief about the amount of money in question — the representation should not be thought of as a "mental picture" of a pile of coins, which the subject regards as a literal depiction of the truth, even if it is not — and so we cannot settle the question of what the decision should be by asking what it would be right to decide if the amounts of money offered were the ones depicted in the subject's mind. These representations are simply mental states (patterns of neural activation), on which the decision may depend; and the answer to what decision will best serve the objective of maximizing expected wealth depends not on what the representations "say" but on what possible objective situations they may have arisen in.

This in turn depends not only on the specification (2.1) of the noisy coding, conditional on the true magnitudes, but also on the relative ex ante likelihood of different possible decision situations, which we specify by a prior probability distribution over possible values of (X, C). We can then consider the optimal decision rule from the standpoint of Bayesian decision theory. It is easily seen that the subject's expected wealth is maximized by a rule under which the risky lottery is chosen if and only if

$$p \cdot \mathrm{E}[X|\mathbf{r}] \ > \ \mathrm{E}[C|\mathbf{r}], \tag{2.2}$$

which is to say if and only if the expected payoff from the risky lottery exceeds the expected value of the certain payoff C.[25] Here $\mathrm{E}[\cdot]$ indicates an expectation under the joint distribution for X, C, and $\mathbf{r}$ implied by the prior probability distribution over decision situations and the conditional probabilities (2.1) for different possible mental representations.

The implications of our logarithmic model of noisy coding are simplest to calculate if (as in the previous section) we assume a log-normal prior distribution for possible monetary quantities. To reduce the number of free parameters in our model, we assume that under the prior X and C are assumed to be independently distributed, and furthermore that the prior distributions for both X and C are the same (some ex ante distribution for possible payments that one may be offered in a laboratory experiment). It is then necessary only to specify the parameters of this common prior:

$$\log X, \log C \ \sim \ N(\mu, \sigma^2) \tag{2.3}$$

where $\sigma^2 > 0$. Under the assumption of a common prior for both quantities, the common prior mean μ does not affect our quantitative predictions about choice behavior;

---

[25]Note that while the payoff C is certain, rather than random, once one knows the decision situation (which specifies the value of C), it is a random variable ex ante (assuming that many different possible values of C might be offered), and continues to be random even conditioning on a subjective representation of the current decision situation, assuming that mental representations are noisy as assumed here.

the value of σ^2 does instead matter, as this influences the ex ante likelihood of X being sufficiently large relative to C for the gamble to be worth taking. The model thus has two free parameters, to be estimated from subjects' behavior: σ^2, indicating the degree of ex ante uncertainty about what the payoffs might be, and ν^2, indicating the degree of imprecision in the mental coding of information that is presented about those payoffs on a particular trial.

2.1 Predicted Frequency of Acceptance of a Gamble

Under this assumption about the prior, the posterior distributions for both X and C are log-normal, as in the model of numerosity estimation in the previous section. The posterior distribution for X is given by

$$\log X|\mathbf{r} \;=\; \log X|r_x \;\sim\; N(\mu_{post}(r_x), \sigma^2_{post}),$$

where the function $\mu_{post}(r)$ and the value of σ^2_{post} are again the ones defined in (1.3), with the sensitivity factor β again defined by (1.4). Similarly, the posterior distribution for C is given by

$$\log C|\mathbf{r} \;=\; \log X|r_c \;\sim\; N(\mu_{post}(r_c), \sigma^2_{post}),$$

using the same definitions. It follows that the posterior means of these variables are given by

$$\mathrm{E}[X|\mathbf{r}] \;=\; e^{\alpha+\beta r_x}, \qquad \mathrm{E}[C|\mathbf{r}] \;=\; e^{\alpha+\beta r_c},$$

with α again defined by (1.5).

Taking the logarithm of both sides of (2.2), we see that this condition will be satisfied if and only if

$$\log p \;+\; \beta r_x \;>\; \beta r_c,$$

which is to say, if and only if the mental representation of the decision situation satisfies

$$r_x - r_c \;>\; \beta^{-1}\log p^{-1}. \tag{2.4}$$

Under our hypothesis about the mental coding, r_x and r_c are independently distributed normal random variables (conditional on the true decision situation), so that

$$r_x - r_c \;\sim\; N(\log X/C, \, 2\nu^2).$$

It follows that the probability of (2.4) holding, and the risky gamble being chosen, is given by

$$\mathrm{Prob}[\text{accept risky}|X, C] \;=\; \Phi\left(\frac{\log X/C - \beta^{-1}\log p^{-1}}{\sqrt{2}\nu}\right), \tag{2.5}$$

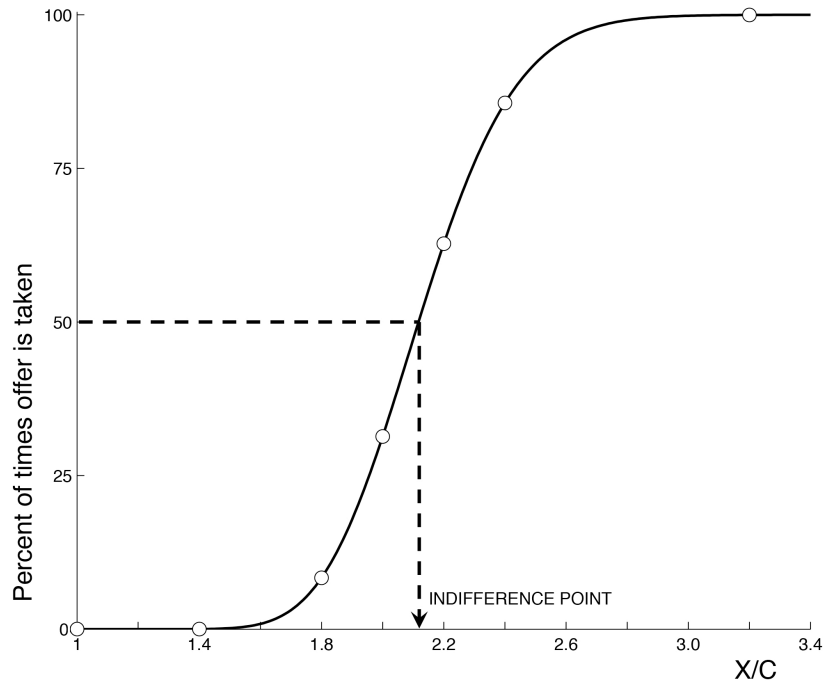where $\Phi(z)$ is the cumulative distribution function of a standard normal random variable.

Figure 2: Theoretically predicted probability of acceptance of a simple gamble, as a function of X/C. Circles show the data from Figure 1, in which $C = 5$ cents.

Equation (2.5) is the behavioral prediction of our model. It implies that choice in a problem of this kind should be stochastic, as observed by Mosteller and Nogee (1951). Furthermore, it implies that across a set of gambles in which the values of p and C are the same in each case, but the value of X varies, the probability of acceptance should be a continuously increasing function of X, as shown in Figure 1. Figure 2 gives an example of what this curve is predicted to be like, in the case that $\sigma = 0.25$ and $\nu = 0.08$. Note that these values allow a reasonably close fit to the choice frequencies plotted in the figure from Mosteller and Nogee.

Moreover, the parameter values required to fit the data are fairly reasonable ones. The value $\nu = 0.08$ for the so-called "Weber fraction" is only half as large as the value of 0.17 in the logarithmic coding model that best fits human performance in comparisons of the numerosity of different fields of dots (Dehaene, 2008, p. 540); on the other hand, Dehaene (2008, p. 552) argues that one should expect the Weber fraction to be smaller in the case of numerical information that is presented symbolically (as in the experiment of Mosteller and Nogee) rather than non-symbolically (as in the numerosity comparison experiments). Hence this value of ν is not an implausible degree of noise to assume in the mental representations of numerical magnitudes used in approximate calculations.[26]

---

[26]In the experiment reported below, our subjects' choices are best fit by values of ν larger than

The value of σ for the degree of dispersion of the prior over possible monetary rewards implies that if under the prior, the median value of X was expected to be 10 cents in the experiment, then the subject should have expected X to fall within a range between 6 cents and 16 cents 95 percent of the time — and this is more or less the range of values offered to the subject, as shown in Figure 1. Hence a prior with this degree of uncertainty would be fairly well calibrated to the subject's actual situation.

2.2   Explaining the Rabin Paradox

Our model explains not only the randomness of the subject's choices, but also his apparent risk aversion, in the sense that the indifference point (a value of X around 10.7 cents in Figure 1) corresponds to a gamble that is better than a fair bet. This is a general prediction of the model, since the indifference point is predicted to be at $X/C = (1/p)^{\beta^{-1}} > 1/p$, where the latter quantity would correspond to a fair bet. The model predicts risk neutrality (indifference when $X/C = 1/p$) only in the case that $\beta = 1$, which according to (1.4) can occur only in the limiting cases in which $\nu = 0$ (perfect precision of the mental representation of numerical magnitudes), or σ is unboundedly large (radical uncertainty about the value of the payoff that may be offered, which is unlikely in most contexts).

The model furthermore explains the Rabin (2000) paradox: the fact that the compensation required for risk does not become negligible in the case of small bets. According to EUM, the value of X required for indifference in a decision problem of the kind considered above should be implicitly defined by the equation

$$pU(W_0 + X) + (1 - p)U(W_0) \;=\; U(W_0 + C), \tag{2.6}$$

where $U(W)$ is the utility associated with wealth W after the outcome of the gamble, and W_0 is the subject's wealth before being offered the two choices. For any increasing, twice continuously differentiable utility function $U(W)$ with $U'' < 0$, if $0 < p < 1$, condition (2.6) implicitly defines a solution $X(C; p)$ with the property that $pX(C; p)/C > 1$ for all $C > 0$, implying risk aversion. However, this solution is such that for small C,

$$\frac{pX(C; p)}{C} \;=\; 1 + \frac{1 - p}{2p}\left(-\frac{U''(W_0)}{U'(W_0)}\right) \cdot C + \mathcal{O}(C^2).$$

Hence the ratio pX/C required for indifference exceeds 1 (the case of a fair bet) only by an amount that becomes arbitrarily small in the case of a small enough bet. It is not possible for the required size of pX to exceed the certain payoff even by 7 percent (as in the case shown in Figure 1), in the case of a very small value certain

---

this — in fact, more similar to the Weber fraction obtained in the study of numerosity comparisons.

payoff, unless the coefficient of absolute risk aversion $(-U''/U')$ is very large — which would in turn imply an implausible degree of caution with regard to large bets.

In our model, instead, the ratio pX/C required for indifference should equal $\Lambda \equiv p^{-(\beta^{-1}-1)}$, which is greater than 1 (except in the limiting cases mentioned above) by the same amount, regardless of the size of the gamble. As discussed above, the degree of imprecision in mental representations required for Λ to be on the order of 1.07 is one that is quite consistent with other evidence. Hence the degree of risk aversion indicated by the choices in Figure 1 is wholly consistent with a model that would predict only a modest degree of risk aversion in the case of gambles involving thousands of dollars.

It is also worth noting that our explanation for apparent risk aversion in decisions about small gambles does not rely on loss aversion, the explanation proposed by Rabin. The hypothesis of loss aversion proposes that the utility assigned to a prospective eventual wealth W is a function $U(W; \tilde{W})$ that evaluates W relative to a "reference point" $\tilde{W}$. It is further assumed that U is not a differentiable function of W at the value $W = \tilde{W}$; specifically, the left derivative (the limiting value of $[U(\tilde{W} + \Delta; \tilde{W}) - U(\tilde{W}; \tilde{W})]/\Delta$ as Δ approaches zero from below) is assumed to be a larger positive number than the right derivative (the limit of the same expression as Δ approaches zero from above). If in the problem considered above, the reference point is assumed to be the amount that it would be possible to obtain with certainty, $\tilde{W} = W_0 + C$, and at this point the left derivative is $\lambda > 1$ times as large as the right derivative, then the value of pX/C required for indifference will approach $\lambda > 1$ for all small enough values of C. Thus a modest degree of loss aversion (a value $\lambda = 1.07$) would suffice to explain the indifference point in Figure 1.

This calculation, however, depends on assuming that the reference point should be the amount that the subject could obtain with certainty,[27] $W_0 + C$, rather than his wealth W_0 prior to being offered the choice. In the case of the latter reference point, loss aversion would yield the same prediction as EUM, since only the properties of the utility function for values $W \geq \tilde{W}$ would be relevant.

The explanation that we offer, instead, does not assume that subjects should have a different attitude to prospective gains in excess of C than to gains that fall short of C. Our model of the mental representation of prospective gains assumes that the coding and decoding of the risky payoff X are independent of the value of C, so that there is no way that small increases in X above C can have a materially different effect than small decreases of X below C.[28] Thus the explanation offered for aversion

---

[27]Alternatively, one can obtain a similar conclusion by assuming that the reference point is given by the subject's expected eventual wealth (as proposed by Koszegi and Rabin, 2006, 2007), which would equal $W_0 + pX$ in the case that the subject has expected to face this opportunity at the time that the reference point is determined, and that the subject chooses to accept the gamble. But it is important that the reference point not be given by the subject's wealth W_0 prior to being offered the choice.

[28]Our theory can be extended so as to predict loss aversion (see footnote 41 below); but in this case, the distinction is between gains and losses relative to the prior wealth W_0, not gains and losses

to small gambles does not involve an assumed kink in the valuation function, as in explanations based on loss aversion.

Instead, in our theory the EUM result that the compensation for risk must become negligible in the case of small enough gambles fails for a different reason. In our theory, the risky gamble is chosen more often than not if and only if $p \cdot m(X) > m(C)$,[29] where $m(\cdot)$ is the function defined in (1.2). The EUM result would still obtain in such a theory if the derivative $m'(0)$ were equal to some finite positive value, but instead, in our theory, $m'(\Delta)$ becomes unboundedly large as Δ approaches zero. This is, however, not an arbitrary assumption on our part, but an inevitable consequence of the hypothesis of logarithmic coding of numerical magnitudes, and more specifically of the existence of a power law for estimates of numerical magnitudes based on noisy mental coding of the numbers in question — a law for which we have seen that there is experimental support in other contexts.

3   An Experimental Test

In order to test the predictions of our model in more detail — in particular, to consider the validity of its prediction that the degree of compensation required for risk (in percentage terms) should be independent of the size of the stakes — we conducted an experiment of our own, in which we varied the magnitudes of both X and C. We recruited 20 subjects from the student population at Columbia University,[30] each of whom was presented with a sequence of several hundred trials, with each independent trial involving a choice between a certain monetary payment and a risky payment.

Figure 3 illustrates the screen observed by one of our subjects on a single trial. The two sides of the screen indicate the two options available on that trial; the subject must indicate whether she chooses the left or right option (by pressing the corresponding key). On the left side of the screen, the dollar amount shown is the quantity C that can be obtained with certainty if "left" is chosen. The right side of the screen shows the possible payoffs if the subject chooses the risky lottery instead. The amounts at the top and bottom of the right side indicate the two possible monetary prizes; the colored rectangular regions in the center indicate the respective probabilities of these two outcomes, if the lottery is chosen. The relative areas of the two rectangular regions provide a visual indication of the relative probabilities of the two outcomes; in addition, the number printed in each region indicates the probability (in percent) that that outcome will occur. (Thus on the trial shown, the subject must choose between a certain payment of \$5.55 and a lottery in which there would be a 58 percent chance of receiving \$15.70, but a 42 percent chance of receiving nothing.) The subject is told that at the end of the experiment, one trial will

---

relative to some expected increment to that wealth.

[29]See section 4.2 for the derivation of this consequence of the model presented above.

[30]Our procedures were approved by the Columbia University Institutional Review Board, under protocol IRB-AAAQ2255.

Figure 3: The computer screen during a single trial of our experiment. The two sides of the screen show the two options available on this trial.

be selected at random, and the subject will actually receive the certain or random monetary reward chosen on that trial, in addition to a fixed monetary payment for participation in the experiment.

The probability p of the non-zero outcome under the lottery was 0.58 on all of our trials, as we were interested in exploring the effects of variations in the magnitudes of the monetary payments, rather than variations in the probability of rewards, in order to test our model of the mental coding of monetary amounts. Maintaining a fixed value of p on all trials, rather than requiring the subject to pay attention to the new value of p associated with each trial, also made it more plausible to assume (as in the model above) that the value of p should be known precisely, rather than having to be inferred from an imprecisely coded observation on each occasion. We chose a probability of 0.58, rather than a round number (such as one-half, as in the Mosteller and Nogee experiment discussed above), in order not to encourage our subjects to approach the problem as an arithmetic problem that they should be able to solve exactly ("Which is larger, 10 dollars or half of 22 dollars?"). We expect Columbia students to be able to solve simple arithmetic problems using methods of exact mental calculation that are unrelated to the kind of approximate judgments about numerical magnitudes with which our theory is concerned, but did not want to test this in our experiment. We chose dollar magnitudes for C and X on all trials that were not round numbers, either, for the same reason.

The value of the certain payoff C varied across trials, taking on the values \$5.55, \$7.85, \$11.10, \$15.70, \$22.20, or \$31.40. (Note that these values represent a geometric series, with each successive amount $\sqrt{2}$ times as large as the previous one.) The non-zero payoff X possible under the lottery option was equal to C multiplied by a factor $2^{m/4}$, where m took one of the values 0, 1, 2, 3, 4, 5, 6, 7, or 8. There were thus

22

only a finite number of decision situations (defined by the values of C and X) that ever appeared, and each was presented to the subject several times over the course of a session. This allowed us to check whether a subject gave consistent answers when presented repeatedly with the same decision, and to compute the probability of acceptance of the risky gamble in each case, as in the experiment of Mosteller and Nogee. The order in which the various combinations of C and X were presented was randomized, in order to encourage the subject to treat each decision as an independent problem, with the values of both C and X needed to be coded and encoded afresh, and with no expectations about these values other than a prior distribution that could be assumed to be the same on each trial.

Our experimental procedure thus differed from ones often used in decision-theory experiments, where care is taken to present a sequence of choices in a systematic order, so as to encourage the subject to express a single consistent preference ordering. We were instead interested in observing the randomization that, according to our theory, should occur across a series of genuinely independent reconsiderations of a given decision problem; and we were concerned to simplify the context for each decision by eliminating any obvious reason for the data of one problem to be informative about the next.

We also chose a set of possible decision problems with the property that each value of X could be matched with the same geometric series of values for C, and vice versa, so that on each trial it was necessary to observe the values of both C and X in order to recognize the problem, and neither value provided much information about the other (as assumed in our theoretical model). At the same time, we ensured that the ratio X/C, on which the probability of choosing the lottery should depend according to our model, always took on the same finite set of values for each value of C. This allowed us to test whether the probability of choosing the lottery would be the same when the same value of X/C recurred with different absolute magnitudes for X and C.

3.1 Testing Scale-Invariance

Figure 4 shows how the frequency with which our subjects chose the risky lottery varied with the monetary amount X that was offered in the event that the gamble paid off, for each of the five different values of C. (For the analysis in this section, we pool the data from all 20 subjects.) Each data point in the figure (shown by a circle) corresponds to a particular combination (C, X).

In the first panel, the horizontal axis indicates the value of X, while the vertical axis indicates the frequency of choosing the risky lottery on trials of that kind [$Prob(Risky)$]. The different values of C are indicated by different colors of circles, with the darker circles corresponding to the lower values of C, and the lighter circles the higher values. (The six successively higher values of C are the ones listed above.) We also fit a sigmoid curve to the points corresponding to each of the different values

Figure 4: The probability of choosing the risky lottery, plotted as a function of the risky payoff X (data pooled from all 20 subjects). (a) The probability plotted as a function of X, for each of the different values of C (indicated by darkness of lines). (b) The same figure, but plotted against $\log X$ for each value of C.

of C, where the color of the curve again identifies the value of C. Each curve has an equation of the form

$$\text{Prob(Risky)} \;=\; \Phi(\delta_C + \gamma_C \log X), \tag{3.1}$$

where $\Phi(z)$ is again the CDF of the standard normal distribution, and the coefficients (δ_C, γ_C) are estimated separately for each value of C so as to maximize the likelihood of the data corresponding to that value of C. Note that for each value of C, we obtain a sigmoid curve similar to the one in Figure 2, though the fit is less perfect (at least partly because here, unlike in Figure 2, we are pooling the data from 20 different subjects).

The similarity of the curves obtained for different values of C can be seen more clearly if we plot them as a function of $\log X$, rather than on a scale that is linear in X, as shown in the second panel of Figure 4. (The color coding of the curves corresponding to different values of C is again the same.) The individual curves now resemble horizontal shifts of one another. The elasticity γ_C is similar for each of the values of C (with the exception of the highest value, $C = \$31.40$), and the value of $\log X$ required for indifference increases by a similar amount each time C is multiplied by another factor of $\sqrt{2}$.
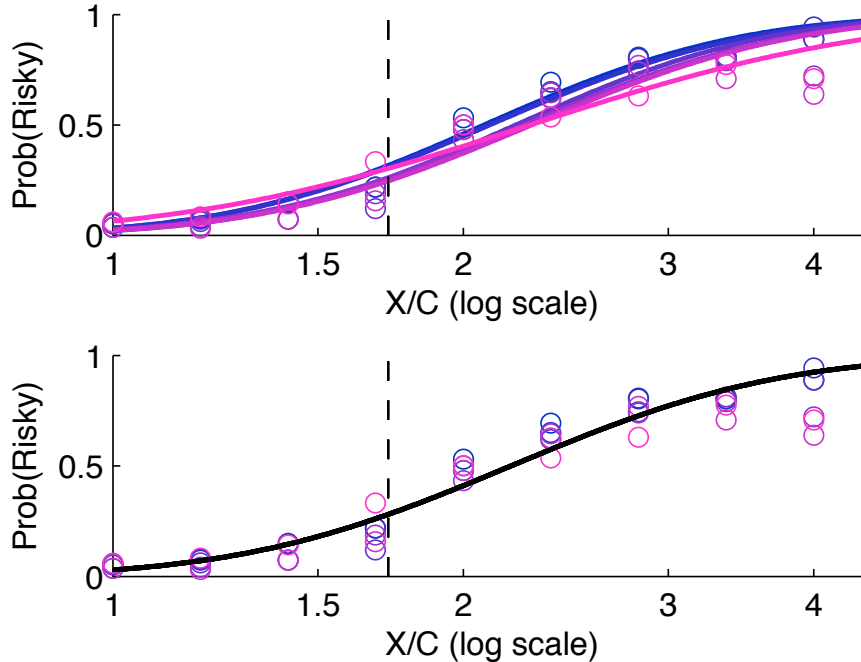
24

Figure 5: The same data as in Figure 4, now plotted as a function of $\log X/C$. (a) A separate choice curve estimated for each value of C, as in Figure 4. (b) A single choice curve, with parameters estimated to maximize the likelihood of the pooled data.

These observations are exactly what we should expect, according to the proposed model of logarithmic mental coding of the data. Condition (2.5) implies that a relationship of the form

$$\text{Prob(Risky)} \;=\; \Phi(\delta + \gamma \log(X/C)) \tag{3.2}$$

should hold for all values of C, meaning that in equation (3.1), γ_C should be the same for each value of C, and that the value of $\log X$ required for indifference should equal a constant plus $\log C$. We can see more clearly the extent to which these precise predictions hold by plotting the curves in Figure 4(b) as functions of $\log(X/C)$, rather than as functions of $\log X$; this is done in the first panel of Figure 5. The six different curves come close to falling on top of one another, as predicted by the model (although, again, the curve for $C = \$31.40$ is somewhat out of line with the others). If we instead simply estimate parameters (δ, γ) to maximize the likelihood of the pooled data under the model (3.2), we obtain the single choice curve shown in the second panel of Figure 5. This fits the data for the different values of X/C slightly worse than the individual choice curves shown in the previous panel, but not by much.

The maximum-likelihood parameter estimates for the different choice curves (estimates of (3.1) for each of the individual values of C, and the estimate of (3.2) using

C	N_{obs}	δ	γ	LL	BIC
\$5.50	1476	-6.16	2.52	-681.3	1377.1
\$7.85	1476	-6.93	2.47	-685.8	1386.3
\$11.10	1476	-8.15	2.54	-654.6	1323.8
\$15.70	1476	-8.56	2.40	-674.6	1363.8
\$22.20	1476	-9.52	2.42	-666.8	1348.2
\$31.40	696	-7.87	1.84	-366.7	746.4
All	8076	-1.88	2.39	-3751.5	7521.0

Table 1: Maximum-likelihood estimates of choice curves for each of the values of C considered separately, and when data from all values of C are pooled. (In each case, data from all subjects are pooled.)

the pooled data) are shown in Table 1. For each estimated model, the table also indicates the number of observations N_{obs} used to estimate the parameters, and the maximized value of the log-likelihood of the data, LL. We can use this information to compute a Bayes information criterion (BIC) statistic for each model, defined as[31]

$$BIC \; \equiv \; -2LL \, + \, 2 \log N_{obs},$$

since each model has two free parameters.

We can consider quantitatively the extent to which our data are more consistent with the more flexible model (3.1) than with the more restrictive predictions of our theory, using the BIC to penalize the use of additional free parameters. If we consider as one possible model of our complete data set a theory according to which there is a curve of the form (3.1) for each value of C, with parameters that may differ (in an unrestricted way) for different values of C, then the BIC associated with this theory (with 12 free parameters) is the sum of the BIC statistics shown in the last column of Table 1 for the individual values of C, equal to 7545.5.[32] The BIC associated with our more restrictive theory (with only two free parameters) is instead only 7521.0, as reported in the bottom row of the table.[33]

The more restrictive model is therefore preferred under the BIC: it leads to a lower value of the BIC, since the increase in the log-likelihood of the data allowed by the additional free parameters is not large enough to offset the penalty for additional free

---

[31]Here, as elsewhere in the paper, "log" refers to the natural logarithm.

[32]Here the BIC is equal to minus 2 times the log-likelihood of the complete data set under the optimized parameters, plus a penalty of $N_{obs}(\theta)$ for each free parameter θ, where $N_{obs}(\theta)$ is the number of observations that are fit using the parameter θ. In the present application, this is the sum of the BICs reported for the models fit to the data for individual values of C.

[33]Our theory implies not only that choice probabilities should be given by a relationship of the form (3.2), but also that the parameters must satisfy conditions (3.3) stated below. However, the unrestricted maximum of the likelihood is attained by parameter values (shown in the bottom line of Table 1) that satisfy these restrictions, so that the best-fitting parameter estimates consistent with our theory, and the associated BIC, are the ones given in the table.

parameters. In fact, the BIC for our more restrictive model is lower by 24.6 points, which implies that the data increase the relative posterior probability of the restrictive model being the correct one, over whatever prior probability may have been assigned to this, by a factor of more than 200,000.[34] Thus our subjects' behavior (pooling the data for the different subjects) exhibits the scale-invariance predicted by our model of noisy mental coding.

Our data fail to be completely scale-invariant in one respect: the estimated choice curve in the case $C = \$31.40$ is not a perfect horizontal translation of the others, and instead is somewhat flatter.[35] This may indicate inaccuracy of the assumption of a log-normal prior (2.3), used in our theoretical calculations above for convenience. Under the assumption of a log-normal prior, $\log \mathrm{E}[X|\mathbf{r}]$ is a linearly increasing function of r_x, with constant slope β. But if people instead form correct inferences based on a prior under which monetary payments greater than \$50 are less likely than a log-normal prior would allow (as was actually the case in our experiment, since we never offered lotteries involving $X/C > 4$), then $\log \mathrm{E}[X|\mathbf{r}]$ would increase less rapidly with further increases in r_x, for values of r_x above $\log 50$. (Under the prior, such large values of r_x would more likely result from mis-coding of a payment of less than \$50 than from a large value of X that has been correctly coded.) This would result in a frequent failure to recognize how attractive the risky lottery truly is when X exceeds \$50, and hence less frequent acceptance of the risky lottery in such cases than the scale-invariant model would predict, as can be observed in Figure 4. (We leave for future work more detailed consideration of the extent to which our data may be better explained by a more subtle account of subjects' prior beliefs.)

Holt and Laury (2002) also obtain nearly perfect scale-invariant choice curves (see their Figure 1), when the amounts offered in hypothetical gambles are scaled up by a factor as large as 90 times those used in their small-stakes gambles. They find, however, that their subjects' apparent degree of risk aversion increases when the scale of the gambles is increased, in the case of gambles for real money (their Figure 2). It is unclear whether this difference from our results (which also involve real money) reflects a difference in the kind of gambles presented to their subjects, or the fact that their large gambles involved greater amounts of money than even our largest gambles (hundreds of dollars rather than mere tens of dollars). Further studies would be desirable to clarify this.

---

[34]The ratio of the posterior probability of the restrictive model being correct to the posterior probability of the flexible model is approximately equal to the ratio of the prior probabilities of the two models being correct, times a factor equal to $e^{\Delta BIC/2}$, where ΔBIC is the amount by which the BIC of the flexible model exceeds that of the restrictive model. See, for example, Burnham and Anderson (2002), p. 303.

[35]Note however that this curve is also less well estimated than the others shown in Figure 4, as a number of our subjects were not presented with trials including values of C this large, so that the N_{obs} for this case is smaller, as indicated in Table 1.

3.2 Heterogeneity in the Precision of Mental Coding

Our data analysis in the previous section pools the data from all subjects, as if a single choice curve should describe each of our subjects. In fact, however, if we estimate separate curves for each subject, we find that they are not identical — in particular, that there is considerable variation in the sensitivity of the choice frequencies to the value of X/C. We can interpret this, in terms of our model, as an indication of heterogeneity across subjects in the degree of precision of the mental coding of prospective monetary payoffs (measured by the parameter ν^2).

The choice curves of individual subjects are on the whole consistent with the theoretical predictions of our model, if we allow the model parameters σ^2 and ν^2 to vary across subjects. Note that the model predicts that there should be a curve of the form (3.2) for each subject, albeit with coefficients γ and δ that may be subject-specific. In addition, the coefficients for each subject should satisfy the inequalities

$$\gamma \geq 0, \qquad -\frac{\delta}{\gamma} \geq \log p^{-1}, \tag{3.3}$$

which are required in order for there to exist values of σ^2 and ν^2 that will imply a choice curve with those coefficients.

We can test the prediction of scale-invariance for each individual subject's behavior using the method employed above in the case of the pooled data. For the individual subject, we again estimate a curve of the form (3.1) for each individual value of C, and also estimate a single curve of the form (3.2), pooling the subject's data for all values of C. The scatter plot in Figure 6 shows for each subject, on the vertical axis, the amount by which the BIC is larger for the model that allows a separate choice curve for each value of C, relative to the model that imposes scale-invariance.[36] We see that not only is the BIC larger for the unrestricted model when we estimate choice curves using the pooled data (as discussed in the previous section), but that for the majority of our subjects, we reach the same conclusion when estimating subject-specific choice curves. The value of ΔBIC is instead negative for a few of our subjects, but in all cases but one, the degree to which the more flexible model is preferred remains modest. Only in the case of subject 9 (the data point toward the bottom of the figure) do we find that the model that assumes scale-invariance fits very poorly.

We next consider the scale-invariant choice curves estimated for the other 19 subjects. We see from Figure 6 that in all but one case, the estimated value of γ is positive, in accordance with the theoretical prediction of our model; however, the estimated values of γ are quite different across subjects. Even in the case of the subject for whom the best-fitting parameter values involve $\gamma < 0$, the value of γ is

---

[36]The vertical axis actually plots the difference in the values of the BIC for the two models divided by N_{obs}, the number of observations used to compute the BIC for that subject. Because N_{obs} is not the same for each of our subjects, the value of $\Delta BIC/N_{obs}$ is more directly comparable across subjects.

Figure 6: Heterogeneity in individual subjects' choice curves. The vertical axis shows for each subject the statistic used to test whether a scale-invariant choice curve better accounts for that subject's data. The horizontal axis shows the individual subject's estimated value of γ, assuming a scale-invariant choice curve.

only slightly negative; and since a likelihood of acceptance of the risky lottery that is genuinely decreasing in X would be difficult to interpret, we presume that this represents sampling error, and treat this subject as having a γ of zero.

In what follows, we estimate a scale-invariant choice curve (3.2) for each of the 19 subjects other than subject 9, now also imposing the constraint that $\gamma \geq 0$ from (3.3). (This last constraint affects our estimates for only one subject.) We next wish to determine whether the other constraint in (3.3) is satisfied as well for each subject. In comparing the choice curves of the different subjects, it is useful to parameterize them not by γ and δ, but instead by the values of γ and

$$\pi \equiv e^{\delta/\gamma}.$$

In terms of this alternative parameterization, the constraints (3.3) can be written:

$$\gamma \geq 0, \qquad \pi \leq p. \tag{3.4}$$

Note that when $\delta/\gamma \leq 0$ (as required by our theoretical model), $0 \leq \pi \leq 1$, and π has the interpretation of a "risk-neutral probability": the subject's indifference point

29

Figure 7: Heterogeneity in subjects' choice curves. Each shaded region indicates the credible region for an individual subject's parameters γ and π, with an open circle at that subject's maximum-likelihood values. The dashed line shows the theoretical relationship between γ and π that should exist if all subjects share a common prior, under which $\sigma = 0.35$.

is the same as that of a risk-neutral, optimizing decision maker who believes that the probability of the non-zero lottery payoff is π (rather than the true probability p). The prediction that $\pi \leq p$ is another way of saying that our theoretical model predicts apparent risk aversion; and the degree to which a subject's estimated π is less than p provides a simple measure of the degree of apparent risk aversion.

The scatter plot in Figure 7 shows the estimated values of γ and π for each of the 19 subjects for whom it is not grossly inaccurate to estimate a scale-invariant choice curve (now imposing the theoretical constraint that $\gamma \geq 0$). For each subject, an open circle indicates the parameter values that maximize the likelihood of the data for that subject (the ML estimate), and the surrounding shaded region indicates the set of parameter values for which the log likelihood of the data is no more than 2 points lower than at the maximum.[37] Thus the shaded region indicates a Bayesian

---

[37]The estimated coefficients of two subjects are shown by pink regions rather than blue ones. These are subjects for whom the accuracy of single estimated choice curves may be doubted, on the ground that their behavior is fairly different during the second half of the session relative to the first half.

credible region for the parameter estimates, under the assumption of a uniform prior for those parameters.[38]

The largest value of π that would be consistent with prediction (3.4) is indicated by the horizontal dotted line in Figure 7; we see that for all but one of the 19 subjects, the credible region includes points consistent with this prediction. Thus the individual choice curves of 18 out of our 20 subjects are reasonably consistent with both the model prediction of scale invariance and with the coefficient constraints (3.4).

Accounting for the choice curves of all subjects in this way, however, requires us to allow different subjects to have different priors (more specifically, different values for σ^2). If instead we assume a common log-normal prior (2.3) for all subjects, but allow the precision of mental coding of monetary amounts (i.e., the parameter ν^2) to vary across subjects, then the values of γ and π estimated for each subject are predicted by the model to be linked by a relationship of the form

$$\pi = p^{1+(2\sigma^2\gamma^2)^{-1}}, \tag{3.5}$$

where σ^2 is a parameter common to all subjects.[39] This is an upward-sloping relationship, of the kind illustrated by the dashed curve in Figure 7, which graphs equation (3.5) in the case that $\sigma = 0.35$.[40] Here ν^2 is decreasing (the precision of mental coding is increasing) as one moves up and to the right along the dashed curve.

If we estimate a choice curve for each subject without imposing the restriction of a common σ^2, the estimated coefficients do not all line up perfectly on a curve consistent with (3.5); nonetheless, there is a strong positive correlation between the ML estimates of γ and π for the various subjects, as can be seen in Figure 7. That is, the degree of apparent risk aversion (measured by the degree to which π is less than p) is generally greater for those subjects whose choices are less sensitive to variation in X/C (measured by the size of γ). The fact that these two features of behavior go hand in hand is consistent with our theory, which attributes both to greater imprecision in the mental coding of monetary payoffs (a larger value of ν^2).

---

[38]The boundary of each maximum-posterior density credible region is chosen according to a criterion which, in the case of a Gaussian posterior for a single variable, would report the interval corresponding to the mean estimate plus or minus two standard deviations.

[39]Equation (3.5) can be derived by using (1.4) and (2.5) to obtain an equation for γ as a function of σ^2 and ν^2; inverting this to obtain the value of ν^2 implied by any subject's value for γ; and then using this result to eliminate ν^2 from the model prediction for the value of π.

[40]This particular value is chosen because it is the one that maximizes the likelihood of the data for the 17 subjects other than subject 9 (omitted from Figure 7) and the two subjects who are indicated by X's in Figure 7. It represents a somewhat low assumed degree of prior uncertainty, given the variability of the monetary amounts actually presented in the experiment; if we pool the data from all trials, the standard deviation of $\log C$ is 0.55, and the standard deviation of $\log X$ is 0.68.

4   Discussion

We have shown that it is possible to give a single unified explanation for the observed randomness in choices by subjects evaluating risky income prospects on the one hand, and the apparent risk aversion that they display on average on the other, as natural consequences of people's intuitions about the value of gambles being based on imprecise mental representations of the monetary amounts that are offered. Our proposed explanation not only explains the possibility of risk aversion without any need to appeal to diminishing marginal utility of wealth, but can also explain the fact (demonstrated in our experiment) that the degree of risk aversion, as measured by the percentage by which the expected value of a random payoff must exceed the certain payoff in order for a subject to be indifferent between them, is relatively independent of the size of the stakes (as long as these remain small), contrary to what should be found if risk aversion were due to diminishing marginal utility.

4.1   Further Implications of the Theory

Our model of noisy mental coding of monetary amounts can also account for further anomalous features of subjects' choices with regard to small gambles documented by Kahneman and Tversky (1979). For example, Kahneman and Tversky report that if subjects must choose between a risky loss and a certain loss — with similar probabilities and monetary quantities as in the kind of problem considered above,[41] but with the *signs* of the monetary payoffs reversed — risk *seeking* is observed more often than risk aversion (something they call the "reflection effect"). The coexistence of both risk-averse choices and risk-seeking choices by the same subject, depending on the nature of the small gambles that are offered, is a particular puzzle for the EUM account of risk attitudes, since a subject should be either risk averse or risk seeking (depending whether the subject's utility of wealth is concave or convex) regardless of the sign of the gambles offered.

The explanation of risk aversion for small gambles offered here instead naturally implies that the sign of the bias (i.e., of the apparent risk attitude) should switch if the signs of the monetary payoffs are switched. Consider instead the case of a choice between a risky gamble that offers a probability p of losing an amount X (but losing nothing otherwise), and the option of a certain loss of an amount C. If we assume that the quantities X and C are mentally represented according to the same model of noisy logarithmic coding as above,[42] regardless of whether they represent gains or

---

[41]The results referred to here are for cases in which, as in the experiments discussed above, the probability of the non-zero payoff (whether positive or negative) is substantial, on the order of one-half or more. Kahneman and Tversky find that risk attitudes change again if instead the probability of the non-zero payoff is small. This kind of effect is also consistent with our theory, if we allow for noisy mental coding of the probability in addition to the noisy mental coding of the monetary amount that can be obtained; but we leave this extension of the theory for a separate treatment.

[42]Note that we assume that the absolute value of each of the monetary payoffs is coded, rather

losses, then in the case of losses, the subject's expected wealth is maximized by a rule under which the risky lottery is chosen if and only if

$$p \cdot \mathrm{E}[X|\mathbf{r}] \; < \; \mathrm{E}[C|\mathbf{r}], \tag{4.1}$$

reversing the sign in (2.2).

The set of representations $\mathbf{r}$ for which this holds will be the complement of the set discussed earlier, so that the model predicts

$$\mathrm{Prob}[\text{accept risky}|X, C] \; = \; \Phi\left(\frac{\beta^{-1}\log p^{-1} - \log X/C}{\sqrt{2}\nu}\right). \tag{4.2}$$

Indifference again will require $pX > C$, but this will now count as *risk-seeking* behavior; when $pX = C$, the risky loss should be chosen more often than not.

Kahneman and Tversky (1979) further show that subjects' preferences between a risky and a safe outcome can be flipped, depending whether the options are presented as involving gains or losses. In one of their problems, subjects are asked to imagine being given a substantial monetary amount $2M$,[43] and then being presented with a choice between (a) winning an additional M with certainty, or (b) a gamble with a 50 percent chance of winning another $2M$ and a 50 percent chance of winning nothing. In a second problem, the initial amount was instead $4M$, and the subsequent choice was between (a) losing M with certainty, and (b) a gamble with a 50 percent chance of losing $2M$ and a 50 percent chance of losing nothing.

These two problems are equivalent, in the sense that in each case the subject chooses between (a) ending up with $3M$ more than their initial wealth with certainty, or (b) a gamble under which they have an equal chance of ending up with $2M$ or $4M$ more than their initial wealth. Nonetheless, a substantial majority of their subjects chose (a) in the first problem, while a substantial majority chose (b) in the second. This contradicts any theory (not just EUM) under which people should have a consistent preference ranking of probability distributions over final wealth levels.

Our theory easily explains this finding. If the initial gift is denoted G, and the monetary amounts G, X, and C defining the decision problem must each be independently represented in the fashion postulated above, then in the first problem, an expected wealth-maximizing decision rule will choose (b) if and only if

$$\mathrm{E}[G|\mathbf{r}] \, + \, p \cdot \mathrm{E}[X|\mathbf{r}] \; > \; \mathrm{E}[G|\mathbf{r}] \, + \, \mathrm{E}[C|\mathbf{r}],$$

than a signed magnitude. This is in accordance with the model of approximate numerical cognition proposed by authors such as Dehaene (2008), which assumes that all numerical quantities are coded as positive amounts, making use of brain circuits originally developed to represent information about the numerosity of sets of items in one's environment. The information whether the quantity in question represents a monetary gain or loss must also be represented, but is assumed to be coded separately, and without error.

[43]In their experiment, conducted in the 1970s, $2M$ was equal to 1000 Israeli shekels, a substantial fraction of a typical monthly income at the time.

which is equivalent to (2.2), while in the second problem it will choose (b) if and only if

$$\mathrm{E}[G|\mathbf{r}] - p \cdot \mathrm{E}[X|\mathbf{r}] > \mathrm{E}[G|\mathbf{r}] - \mathrm{E}[C|\mathbf{r}],$$

which is equivalent to (4.1). We then get different probabilities of choosing (b) in the two cases, given by equations (2.5) and (4.2) respectively.

Note that our theory assumes that the decision rule is in all cases the one that maximizes expected final wealth, so that only the sum of the initial gift and the additional gain or loss from the further option is assumed to matter to the decision maker; there is no intrinsic interest assumed in gains or losses relative to what one had in the past or what one expected to have. The relevance of the sequence of gains and losses by which one arrives at a given final wealth comes not from the decision maker's assumed objective, but from the need to mentally represent the quantities used in the description of the options, in the form in which they are presented, before integrating the separate pieces of information in further calculations. If this representation were possible with infinite precision (and subsequent operations could also be perfectly precise), then different ways of presenting information that imply the same possible final wealth levels would indeed be equivalent, and lead to the same choices. But when the precision with which each monetary amount can be represented is limited, mathematically equivalent problems are not processed in identical ways, and the resulting behavior can be different as a result, despite its optimality in each case (conditional on the mental representation).

Thus far, we have discussed implications of our model, taking the precision of coding (parameterized by ν^2) to be fixed. But the model also makes predictions about the effects of varying ν^2, which might be subject to predictable variation for a variety of reasons. For example, one might well suppose that increased time pressure, distraction or cognitive load should reduce the cognitive resources used to represent the monetary magnitudes that define a particular decision problem, and that this should correspond, in our model, to an increase in ν^2. According to our model, this should result in both decreased sensitivity of a subject's decisions to variations in the risky payoff X that is offered (i.e., a lower value of γ) and increased apparent risk aversion (a value of π that is lower relative to p).[44] Interestingly, Whitney *et al.* (2008) find that subjects make less risk-averse choices between risky and certain gains when cognitive load is increased, by requiring them to concurrently maintain a list of random letters in memory. However, they also find that their subjects make less risk-seeking choices between risky and certain losses (or more precisely, reductions of the amount that they expect to gain from the experiment), which would not be predicted by our model, under the explanation proposed above for risk-seeking in the domain of losses. We hope to further examine these implications of our model in future work.

---

[44]Recall that the dashed curve in Figure 7 shows the effect on both γ and π of varying ν^2, while holding fixed the prior distribution over possible values of X and C.

A number of studies have found that subjects who score better on tests of numeracy or reasoning ability exhibit less apparent risk aversion in their choices between small gambles — and more specifically, are more likely to make choices that maximize expected monetary reward (e.g., Frederick, 2005; Cokely and Kelley, 2009; Benjamin *et al.,* 2013; Weller *et al.,* 2013; and Schley and Peters, 2014). Such results are often interpreted as evidence for two alternative mental "systems" that may be used, with the more numerate subjects (for whom explicit calculation has lower cost) more likely to solve the problem by calculation of expected values, while others instead rely on an intuitive valuation system that produces biased judgments. Our theory is consistent with a view of this kind; the model presented above offers an explanation of the bias in intuitive valuations, while not denying that educated adults with some facility for arithmetic should be able to value simple gambles more accurately through calculations based on symbolic number representations, rather than the approximate representations modeled here.

But our theory also makes possible another interpretation of the association between numeracy and less risk-averse choices. Even when people rely upon intuitive valuations, some may devote more processing capacity (for example, more space in working memory) to the task than others, allowing them more precise approximate representations of the monetary payoffs (a smaller ν^2, in our model); according to our theory, this should result in less randomness in choice and closer conformity to expected-value maximization. The fact that our subjects display a range of degrees of apparent risk aversion, as well as a range of degrees of randomness in their choices, as shown in Figure 7 — rather than simply two clusters corresponding to the users of two very different mental systems — is more easily explained under this latter view.

4.2 Comparison with Related Theories

The predictions of our theory developed above, apart from those relating to the randomness of subjects' choices, are also predictions of prospect theory (Kahneman and Tversky, 1979). In fact the explanation that we offer is not so much an alternative to the prospect-theoretic explanation as it is an explanation for why the sort of distorted perceptions proposed in prospect theory should exist; in particular, we offer an explanation for the kind of biases implied by the nonlinear "value function" proposed by Kahneman and Tversky, that assigns values to prospective gains or losses relative to some "reference point." Kahneman and Tversky motivate their assumptions about the shape of the value function on the basis of an analogy with sensory perception;[45]

---

[45]As justification for the proposed shape of their value function, Kahneman and Tversky note: "Many sensory and perceptual dimensions share the property that the psychological response is a concave function of the magnitude of physical change.... We propose that this principle applies [as well] to the evaluation of monetary changes" (p. 278). They also justify their assumption that the value function is applied to the gain or less relative to a reference point, rather than to the absolute level of wealth that would prospectively be reached, by analogy with adaptation effects in sensory perception (p. 277).

we explain more fully the biases in choices between risky prospects that are predicted if one takes such an analogy seriously.

The predictions of our theory about the cases considered above can in fact be expressed in a form similar to the one proposed by Kahneman and Tversky. Our theory assumes that each risky option is evaluated on the basis on the expected value of the final wealth W that would be obtained, if the estimate of expected final wealth is based on a mental representation $\mathbf{r}$ of the features of the option. This subjective expected value can be written as

$$\mathrm{E}[W|\mathbf{r}] \;=\; \mathrm{E}[W_0|\mathbf{r}] \;+\; \sum_i p_i \mathrm{E}[\Delta_i|\mathbf{r}],$$

where W_0 is the subject's wealth prior to additional gains or losses offered by the risky prospect, the quantities $\{\Delta_i\}$ represent the additional net gains offered by the risky prospect (with the different possible outcomes indexed by i), and the $\{p_i\}$ are the respective probabilities of these outcomes.[46] If we further assume that the mental representation $\mathbf{r}$ has independent components r_0 and $\{r_i\}$ representing the initial wealth W_0 and the possible net gains $\{\Delta_i\}$ respectively, with each component being drawn from a probability distribution of possible values that depends only the value of the corresponding component of the true state, and that the prior distribution over possible true states is one in which each of these components is distributed independently of the others, then the subjective expected value is given by a term

$$V(\mathbf{r}) \;\equiv\; \sum_i p_i \mathrm{E}[\Delta_i|r_i]$$

that depends only on the potential net gains associated with the risky prospect, plus a term $\mathrm{E}[W_0|r_0]$ that is independent of the risky prospect. The prospect that is chosen on a given trial should then be the one with the larger value of $V(\mathbf{r})$.

The mean value of this subjective valuation, integrating over the possible realizations of the random mental representation $\mathbf{r}$, conditional on the true potential net gains, is in turn given by

$$EV \;=\; \sum_i p_i\, v(\Delta_i), \tag{4.3}$$

where $v(\Delta)$ is equal to the function $m(\Delta)$ defined in (1.2) if $\Delta \geq 0$, and equal to $-m(-\Delta)$ if $\Delta \leq 0$. Note that this has the form of an expected value with a nonlinear "value function" for net gains, as assumed in prospect theory; moreover, the value function is concave for gains, and convex for losses, as assumed by Kahneman

---

[46]For simplicity, in this paper we treat the probabilities as precisely understood. Noisy coding of this information as well (which would be entirely consistent with our general argument) can give rise to additional biases, similar to those resulting from nonlinear probability weighting in prospect theory. Because the experiments reported in this paper do not involve comparisons between different probabilities of gain or loss, we leave such an extension of our theory for development elsewhere.

and Tversky.[47] In our case, however, the nonlinearity of the valuation function is derived from an explicit model of decision making on the basis of a noisy mental representation.

In the case that (as in the problems considered above) each option involves only one possible non-zero net gain, then $V(\mathbf{r})$ will be a log-normal random variable, conditional on the true net gains $\{\Delta_i\}$.[48] In fact, we can show that

$$\log |V| \ \sim \ N(\log |EV| - (1/2)\beta^2\nu^2, \ \beta^2\nu^2)$$

for each option of this simple kind (with the sign of V also necessarily equal to the sign of EV). The option that is predicted to be chosen more often is then the option with the higher value of EV,[49] so that modal behavior is predicted by the ranking of alternative risky prospects under a bilinear form (4.3), as in prospect theory.

While the predictions of our theory are thus (at least in certain simple cases discussed in this paper) the same as those of prospect theory, we believe that the account given here nonetheless adds to that provided by Kahneman and Tversky (1979). First, our theory predicts not only the modal choices, but the probability of making each choice, as a function of the parameters of the decision problem; Kahneman and Tversky instead do not seek to model the randomness of their subjects' behavior.

In addition, our theory provides a functional explanation for the existence of a nonlinear value function. If one supposes that the value function of prospect theory represents a deterministic distortion of the data presented to the subjects, the result of which is available to subjects with perfect precision, one might well wonder why their brains should have evolved to produce such a distorted report (rather than a

---

[47]Kahneman and Tversky assume that their value function has a kink at $\Delta = 0$, reflecting loss aversion. If we assume an identical model of noisy coding for gains and losses, and also identical priors for the potential true sizes of gains and losses, then the model derived here implies that the value function should exhibit perfect antisymmetry ($v(-\Delta) = -v(\Delta)$), and hence no loss aversion. However, if we assume that the prior distribution for losses differs from that for gains, this will not be the case. We will then have different functions $m^{gain}(\Delta)$ and $m^{loss}(\Delta)$, each defined as in (1.2) but with different values of A and β in the case of gains and losses, owing to the different values of μ and σ^2 in the two cases; and (depending on how the prior distributions differ) the asymmetry can imply loss aversion. The idea that loss aversion results from a difference in the distributions of gains and losses which subjects expect to encounter is supported by the results of Walasek and Stewart (2015), who are able to change subjects' degree of loss aversion by manipulating the distributions of gains and losses presented in experimental trials. It is also consistent with the finding of Ert and Erev (2008) that the degree to which subjects' choices exhibit loss aversion depends on how the options are presented to them, in a ways that might well lead to different priors regarding the distribution of likely offers. Note that both of these latter results are inconsistent with the view that loss aversion reflects a difference in the degree to which decision makers care about losses as opposed to gains.

[48]As in the model of numerosity estimation discussed in section 1.2, the logarithm of $|V(\mathbf{r})|$ will be an affine function of the subjective representation r_i of the non-zero payoff, which is normally distributed.

[49]Note that this would not always be true in the case of more complex gambles.

correct one), and why their decisions treat the distorted value as if it were a true value (when instead it is inaccurate in a perfectly predictable way, and an alternative decision rule based on the same information — namely, one that inverts each of the valuations $v(\Delta_i)$ to recover the true valuation Δ_i before weighting and summing to estimate the expected value — would yield better outcomes on average). In our interpretation, instead, the available mental representation of the decision problem is instead imprecise (an inevitable result of the use of limited mental resources for such representations), and the decision rule is actually the one that achieves the highest possible expected wealth for the decision maker, among all rules that use no more information than that contained in the imprecise representation.

Schley and Peters (2014) similarly offer an explanation for apparent risk aversion which is based on the idea that the perception of the numerical magnitudes of prospective monetary payoffs is biased, and more specifically that perceived magnitudes are an increasing, strictly concave function of the magnitudes. Like us, they base their proposal on limitations on people's general ability to accurate represent numbers mentally, rather than on the true utility obtained from having more money (as in the EUM explanation of risk aversion) or a theory of distorted valuations that is specific to the domain of value-based decision making (as with prospect theory); and in fact they show that subjects who less accurately represent numbers for other purposes also exhibit greater apparent diminishing marginal utility of income and greater apparent risk-aversion in choices between risky gambles.[50]

However, their discussion assumes that less capacity for symbolic number mapping results in a deterministic distortion of perceived numerical magnitudes (a true quantity X is always perceived as exactly $\hat{X} = AX^\beta$), rather than in a more random mental representation as in our theory. This means that they do not explore the connection between the randomness of subjects' choices and apparent risk aversion, as we do here; and their theory provides no explanation for why people should value lotteries according to the average value of the perceived payoffs $\hat{X}_i$ instead of, say, according to the average value of $\hat{X}_i^{1/\beta}$ — a criterion that would reliably maximize expected wealth, taking into account the perceptual distortion.

A theory more similar to ours is the model of risk-sensitive foraging by animals (such as starlings) proposed by Kacelnik and Abreu (1998). These authors are concerned with how animals choose between options that they repeatedly face (alternative possible locations for foraging), on the basis of past experience of the probability distribution of possible outcomes associated with each, and their theory accordingly turns on the way that previously experienced outcomes are represented in memory, and the way in which the distribution represented in memory is drawn upon at the

---

[50]More precisely, they fit each of their subjects' choices to a prospect-theoretic valuation formula, where the value function for either gains or losses is assumed to be of the power-law form (1.2), and estimate a value of the exponent β for each subject. They find that subjects who score higher on a test of ability to accurately locate symbolically presented numbers on a spatial number line have values of β closer to 1.

time of a new prospective choice; it is not a theory of the representation of numerical descriptions of available options (which are not available to foraging starlings). Moreover, the variability in the choices made across repeated presentations of the same options is attributed to variation in the random samples drawn on each occasion from a fixed mental representation of the distribution of possible outcomes under each option (with the situation being recognized by the organism as a repetition of the same situation as before), rather than randomness in new representations of the payoffs that are assumed in our theory to be formed each time the decision problem is presented again (and not recognized as having been previously encountered).

Nonetheless, the implications of their theory — which like ours is based on randomness in the representation of rewards that conforms to "Weber's Law," and derives predictions for both choice probabilities and apparent risk aversion — are similar to those of ours, while not mathematically identical.[51] The fact that a similar model can successfully explain animal behavior of the kind that Kacelnik and Abreu review provides further reason, in our view, to consider our proposed explanation for intuitive judgments by humans a plausible one.

Finally, the model of Woodford (2012) is similar to ours, in that it derives both risk aversion with respect to gains and risk seeking with respect to losses from a model of noisy coding of prospective net gains, with a decision rule that maximizes the subject's expected wealth. However, the model of noisy coding is different: net gain is coded as a single variable (which may be of either sign), rather than gains and losses being coded separately, and the assumed inhomogeneity in the precision of coding of net gains makes the mean Bayesian estimate of net gain an S-shaped function of the true net gain, rather than there being separate concave functions for the mean estimates of gains and losses as above. We feel that the model of mental coding proposed here is more realistic, because in the experiments that we seek to explain, the prospective outcomes are described to subjects in terms of positive quantities of money that can be gained or lost, rather than in terms of a signed net gain. Whether the kind of inhomogeneity in the precision of coding relied upon here can be justified as an efficient use of finite processing resources, as in the model proposed in Woodford (2012), is an important topic for further investigation.

---

[51] A key mathematical difference is that Kacelnik and Abreu do not model random coding in conformity with Weber's Law in the same way that we do; they assume a truncated normal distribution for the mental representations, with a standard deviation proportional to the mean, rather than a log-normal distribution.

References

[1] Anobile, Giovanni, Guido Marco Cicchini, and David C. Burr, "Linear Mapping of Numbers onto Space Requires Attention," *Cognition* 122: 454-459 (2012).

[2] Appelle, Stuart, "Perception and Discrimination as a Function of Stimulus Orientation: The 'Oblique Effect' in Man and Animals," *Psychological Bulletin* 78: 266-278 (1972).

[3] Ballinger, T. Parker, and Nathaniel T. Wilcox, "Decisions, Error, and Heterogeneity," *Economic Journal* 107: 1090-1105 (1997).

[4] Banks, William P., Milton Fujii, and Fortunee Kayra-Stewart, "Semantic Congruity Effects in Comparative Judgments of Magnitudes of Digits," *Journal of Experimental Psychology: Human Perception and Performance* 2: 435-447 (1976).

[5] Benjamin, Daniel J., Sebastian A. Brown, and Jesse M. Shapiro, "Who is Behavioral? Cognitive Ability and Anomalous Preferences," *Journal of the European Economics Association* 11: 12311255 (2013).

[6] Bernoulli, Daniel, "Exposition of a New Theory on the Measurement of Risk," *Econometrica* 22: 22-36 (1954). [English translation of Latin original, 1738.]

[7] Buckley, Paul B., and Clifford B. Gillman, "Comparisons of Digits and Dot Patterns," *Journal of Experimental Psychology* 103: 1131-1136 (1974).

[8] Burnham, Kenneth P., and Anderson, David R., *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach,* 2d. ed., New York: Springer, 2002.

[9] Cantlon, Jessica F., and Elizabeth M. Brannon, "Shared System for Ordering Small and Large Numbers in Monkeys and Humans," *Psychological Science* 17: 401-406 (2006).

[10] Cordes, Sara, Rochel Gelman, Charles R. Gallistel, and John Whalen, "Variability Signatures Distinguish Verbal from Non-Verbal Counting for Both Large and Small Numbers," *Psychonomic Bulletin and Review* 8: 698-707 (2001).

[11] Cox, James C., Vjollca Sadiraj, Bodo Vogt, and Utteeyo Dasgupta, "Is There a Plausible Theory for Decision Under Risk? A Dual Calibration Critique," *Economic Theory* 54: 305-333 (2013).

[12] Dantzig, Tobias, *Number: The Language of Science,* London: Plume, 2007.

[13] DeGardelle, Vincent, Sid Kouider, and Jerome Sackur, "An Oblique Illusion Modulated by Visibility: Non-Monotonic Sensory Integration in Orientation Processing," *Journal of Vision* 10 (10), article 6 (2010).

[14] Dehaene, Stanislas, "Symbols and Quantities in Parietal Cortex: Elements of a Mathematical Theory of Number Representation and Manipulation," in P. Haggard, Y. Rossetti, and M. Kawato, eds., *Sensorimotor Foundations of Higher Cognition,* Oxford: Oxford University Press, 2008.

[15] Dehaene, Stanislas, *The Number Sense,* revised and updated edition, Oxford: Oxford University Press, 2011.

[16] Dehaene, Stanislas, and Laurent Cohen, "Two Mental Calculation Systems: A Case Study of Severe Acalculia with Preserved Approximation," *Neuropsychologia* 29: 1045-1074 (1991).

[17] Dehaene, Stanislas, and J. Frederico Marques, "Cognitive Euroscience: Scalar Variability in Price Estimation and the Cognitive Consequences of Switching to the Euro," *Quarterly Journal of Experimental Psychology A* 55: 705-731 (2002).

[18] Dewan, Ambuj, and Nate Neligh, "Estimating Information Cost Functions in Models of Rational Inattention," unpublished, Columbia University, January 2017.

[19] Ert, Eyal, and Ido Erev, "The Rejection of Attractive Gambles, Loss Aversion, and the Lemon Avoidance Heuristic," *Journal of Economic Psychology* 29: 715-723 (2008).

[20] Fechner, Gustav, *Elements of Psychophysics,* New York: Holt, Rinehart and Winston, 1966. (English translation of the German original, 1860.)

[21] Frederick, Shane, "Cognitive Reflection and Decision Making," *Journal of Economic Perspectives* 19: 25-42 (2005).

[22] Friedman, Daniel, R. Mark Isaac, Duncan James, and Shyam Sunder, *Risky Curves: On the Empirical Failure of Expected Utility,* London: Routledge, 2014.

[23] Gabaix, Xavier, and David Laibson, "Myopia and Discounting," NBER Working Paper no. 23254, March 2017.

[24] Gabbiani, Fabrizio, and Steven J. Cox, *Mathematics for Neuroscientists,* Amsterdam: Academic Press, 2010.

[25] Ganguli, Deep, "Chapter 3: Efficient Sensory Coding Predicts the Heterogeneities of Neural Populations and Perception," in *Efficient Coding and Bayesian Estimation with Neural Populations,* Ph.D. dissertation, Department of Neural Science, New York University, 2012.

[26] Gescheider, George A., *Psychophysics: The Fundamentals,* 3d ed., Mahwah, NJ: Lawrence Erlbaum Associates, 1997.

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

[28] Glimcher, Paul W., *Foundations of Neuroeconomic Analysis,* Oxford: Oxford University Press, 2011.

[29] Hey, John D., "Experimental Investigations of Errors in Decision Making under Risk," *European Economic Review* 39: 633-640 (1995).

[30] Hey, John D., "Does Repetition Improve Consistency?" *Experimental Economics* 4: 5-54 (2001).

[31] Hey, John D., and Chris Orme, "Investigating Parsimonious Generalizations of Expected Utility Theory using Experimental Data," *Econometrica* 62: 1291-1329 (1994).

[32] Hollingsworth, Walter H., J. Paul Simmons, Tammy R. Coates, and Henry A. Cross, "Perceived Numerosity as a Function of Array Number, Speed of Array Development, and Density of Array Items," *Bulletin of the Psychonomic Society* 29: 448-450 (1991).

[33] Holt, Charles A., and Susan K. Laury, "Risk Aversion and Incentive Effects," *American Economic Review* 92: 1644-1655 (2002).

[34] Hubel, David H., and Torsten N. Wiesel, "Receptive Fields of Single Neurones in the Cat's Striate Cortex," *Journal of Physiology* 148: 574-591 (1959).

[35] Indow, Tarow, and Masashi Ida, "Scaling of Dot Numerosity," *Perception and Psychophysics* 22: 265-276 (1977).

[36] Izard, Véronique, and Stanislas Dehaene, "Calibrating the Mental Number Line," *Cognition* 106: 1221-1247 (2008).

[37] Jastrow, Joseph, "Studies from the University of Wisconsin: On the Judgment of Angles and Positions of Lines," *American Journal of Psychology* 5: 214-248 (1892).

[38] Kacelnik, Alex, and Fausto Brito e Abreu, "Risky Choice and Weber's Law," *Journal of Theoretical Biology* 194: 289-298 (1998).

[39] Kahneman, Daniel, and Amos Tversky, "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica* 47: 263-291 (1979).

[40] Kaufman, E.L., M.W. Lord, T.W. Reese, and J. Volkmann, "The Discrimination of Visual Number," *American Journal of Psychology* 62: 498-525 (1949).

[41] Koszëgi, Botond, and Matthew Rabin, "A Model of Reference-Dependent Preferences," *Quarterly Journal of Economics* 121: 1133-1165 (2006).

[42] Koszëgi, Botond, and Matthew Rabin, "Reference-Dependent Risk Attitudes," *American Economic Review* 97: 1047-1073 (2007).

[43] Koszëgi, Botond, and Matthew Rabin, "Revealed Mistakes and Revealed Preferences," in A. Caplin and A. Schotter, eds., *The Foundations of Positive and Normative Economics,* Oxford: Oxford University Press, 2008.

[44] Kramer, Peter, Maria Grazia De Bono, and Marco Zorzi, "Numerosity Estimation in Visual Stimuli in the Absence of Luminance-Based Cues," *PLoS ONE* 6(2): e17378 (2011).

[45] Krueger, Lester E., "Perceived Numerosity," *Perception and Psychophysics* 11: 5-9 (1972).

[46] Krueger, Lester E., "Perceived Numerosity: A Comparison of Magnitude Production, Magnitude Estimation, and Discrimination Judgments," *Perception and Psychophysics* 35: 536-542 (1984).

[47] Loomes, Graham, "Modeling the Stochastic Component of Behaviour in Experiments: Some Issues for the Interpretation of Data," *Experimental Economics* 8: 301-323 (2005).

[48] Loomes, Graham, and R. Sugden, "Incorporating a Stochastic Element into Decision Theories," *European Economic Review* 39: 641-648 (1995).

[49] Mosteller, Frederick, and Philip Nogee, "An Experimental Measurement of Utility," *Journal of Political Economy* 59: 371-404 (1951).

[50] Moyer, Robert S., and Thomas K. Landauer, "Time Required for Judgements of Numerical Inequality," *Nature* 215: 1519-1520 (1967).

[51] Nieder, Andreas, "Coding of Abstract Quantity by 'Number Neurons' in the Primate Brain," *Journal of Comparative Physiology A* 199: 1-16 (2013).

[52] Nieder, Andreas, and Stanislas Dehaene, "Representation of Number in the Brain," *Annual Review of Neuroscience* 32: 185-208 (2009).

[53] Nieder, Andreas, and Katharina Merten, "A Labeled-Line Code for Small and Large Numerosities in the Monkey Prefrontal Cortex," *Journal of Neuroscience* 27: 5986-5993 (2007).

[54] Piazza, Manuela, Véronique Izard, Philippe Pinel, Denis Le Bihan, and Stanislas Dehaene, "Tuning Curves for Approximate Numerosity in the Human Intraparietal Sulcus," *Neuron* 44: 547-555 (2004).

[55] Pica, Pierre, Cathy Lemer, Véronique Izard, and Stanislas Dehaene, "Exact and Approximate Arithmetic in an Amazonian Indigene Group," *Science* 306: 499-503 (2004).

[56] Rabin, Matthew, "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," *Econometrica* 68: 1281-1292 (2000).

[57] Ross, John, "Visual Discrimination of Number without Counting," *Perception* 32: 867-870 (2003).

[58] Schley, Dan R., and Ellen Peters, "Assessing 'Economic Value': Symbolic-Number Mappings Predict Risky and Riskless Valuations," *Psychological Science* 25: 753-761 (2014).

[59] Spelke, Elizabeth S., and Sanna Tsivkin, "Language and the Brain: A Bilingual Training Study," *Cognition* 78: 45-88 (2001).

[60] Steiner, Jakub, and Colin Stewart, "Perceiving Prospects Properly," *American Economic Review* 106: 1601-1631 (2016).

[61] Stocker, Alan A., and Eero P. Simoncelli, "Noise Characteristics and Prior Expectations in Human Visual Speed Perception," *Nature Neuroscience* 9: 578-585 (2006).

[62] Thompson, Richard F., Kathleen S. Mayers, Richard T. Robertson, and Charlotte J. Patterson, "Number Coding in Association Cortex of the Cat," *Science* 168: 271-273 (1970).

[63] Tomassini, Alessandro, Michael J. Morgan, and Joshua A. Solomon, "Orientation Uncertainty Reduces Perceived Obliquity," *Vision Research* 50: 541-547 (2010).

[64] van Oeffelen, Michiel P., and Peter G. Vos, "A Probabilistic Model for the Discrimination of Visual Number," *Perception and Psychophysics* 32: 163-170 (1982).

[65] Walasek, Lukasz, and Neil Stewart, "How to Make Loss Aversion Disappear and Reverse: Tests of the Decision by Sampling Origin of Loss Aversion," *Journal of Experimental Psychology: General* 144: 7-11 (2015).

[66] Weber, Ernst H., *The Sense of Touch,* New York: Academic Press, 1978. [English translation of German original, 1834.]

[67] Wei, Xue-Xin, and Alan A. Stocker, "A Bayesian Observer Model Constrained by Efficient Coding Can Explain 'Anti-Bayesian' Percepts," *Nature Neuroscience* 18: 1509-1517 (2015).

[68] Weller, Joshua A., Nathan F. Dieckmann, Martin Tusler, C.K. Mertz, William J. Burns, and Ellen Peters, "Development and Testing of an Abbreviated Numeracy Scale: A Rasch Analysis Approach," *Journal of Behavioral Decision Making* 26: 198-212 (2013).

[69] Whalen, J., Charles R. Gallistel, and Rochel Gelman, "Non-Verbal Counting in Humans: The Psychophysics of Number Representation," *Psychological Science* 10: 130-137 (1999).

[70] Whitney, Paul, Christa A. Rinehart, and John M. Hinson, "Framing Effects Under Cognitive Load," *Psychonomic Bulletin and Review* 15: 1179-1184 (2008).

[71] Wilcox, Nathaniel T., "Stochastic Models for Binary Discrete Choice Under Risk: A Critical Primer and Econometric Comparison," in J.C. Cox and G.W. Harrison, eds., *Research in Experimental Economics, vol. 12: Risk Aversion in Experiments,* Bingley, UK: Emerald Group Publishing, 2008.

[72] Woodford, Michael, "Prospect Theory as Efficient Perceptual Distortion," *American Economic Review* 102(3): 1-8 (2012).