THE BOUNDED RATIONALITY OF PROBABILITY DISTORTION

HANG ZHANG[1,2,3*] XIANGJUAN REN[3] LAURENCE T. MALONEY[4,5]

1. School of Psychological and Cognitive Sciences and Beijing Key Laboratory of Behavior and Mental Health, Peking University, Beijing, China
2. PKU-IDG/McGovern Institute for Brain Research, Peking University, Beijing, China
3. Peking-Tsinghua Center for Life Sciences, Beijing, China
4: Department of Psychology, New York University, New York, NY
5: Center for Neural Science, New York University, New York, NY

*Corresponding author:

Hang Zhang, email: hang.zhang@pku.edu.cn

Short title: Bounded log-odds model

ABSTRACT

In decision under risk participants' choices are based on probability values systematically different from those that are objectively correct. Similar systematic distortions are found in tasks involving relative frequency judgments. These distortions limit performance in a wide variety of tasks and an evident question is, why do we systematically fail in our use of probability and relative frequency information? We propose a Bounded Log-Odds Model (BLO) of probability and relative frequency distortion based on three assumptions. The key assumption is that the dynamic range of representation of probability and relative frequency is limited. We tested the model experimentally and found that BLO accounts for individual participants' data better than all previous models in the literature. We also show that subject to the dynamic range limitation, participants' choice of distortion serves to maximize the mutual information between objective and internal values, a form of bounded rationality.

Keywords: probability distortion; frequency judgment; decision under risk; efficient coding; mutual information

In making decisions, we choose among actions whose outcomes are typically uncertain; we can model such choices as choices among lotteries. To specify a lottery L we list all of its possible outcomes $O_1,...,O_n$ and the corresponding probabilities of occurrence $p_1,\cdots,p_n$ that a specific lottery assigns to each outcome. If we knew all the relevant probabilities, we would be engaged in decision under risk [1]. If we can also assign a numerical measure of utility $U\left(O_i\right)$ to each outcome O_i, we could assign an *expected utility* to each lottery,

$$EU\left(L\right)=\sum_{i=1}^{n}p_i U\left(O_i\right) , \tag{1}$$

and a decision maker maximizing expected utility [2, 3] would select the lottery with the highest expected utility among those offered. The probabilities serve to weight the contribution of the utility of each outcome. The Expected Utility Theory (EUT) model is simple, but has a wide range of applications, not just in economic decisions but also in perception [4, 5] and planning of movement [6, 7, 8, 9, 10].

For more than two centuries EUT was treated as an adequate description of human choice behavior in decision under risk until challenged by Allais [11]. In an elegant series of experiments he showed that human decision makers did not weight utilities by the corresponding probabilities of occurrence in choosing among lotteries. In Prospect Theory, Kahneman and Tversky [12] resolved the Allais paradoxes and other shortcomings of EUT by assuming that decision makers use a transformation of probability $\pi\left(p\right)$—a *probability weight* or *decision weight*—in place of probability p in the computation of expected utility. The distortion function in decision under risk $\pi\left(p\right)$ was originally inferred from human choices in experiments and it is often—but not always—an inverted-*S*-shaped function of p [13, 14, 15].

Wu, Delgado, and Maloney [16] compared performance in a “classical” decision under risk task with performance in a mathematically equivalent motor decision task. Each

participant completed both tasks and while the fitted probability distortion functions for the classical task were—as expected—inverted-*S*-shaped, those based on the motor task tended to be better fit by *S*-shaped functions. The same participant could have both the inverted-*S*-shaped and *S*-shaped forms of the distortion function $\pi(p)$ in different decision tasks.

Ungemach, Stewart and Chater [17] found a similar tendency to underweight small probabilities in decisions and overweight large [18, 19, 20]. Probability distortion in the form of inverted-*S*-shaped and *S*-shaped weighting functions is also found in monkeys' choice behavior [21] and is supported by human neuroimaging evidence [22, 23].

Zhang and Maloney [24] reported that both the inverted-*S*-shaped or *S*-shaped distortion functions are found in relative frequency and confidence tasks other than decision-making under risk. For convenience, we will use the term "probability" to include relative frequency and confidence. The same participants had different inverted-*S*-shaped or *S*-shaped probability distortion functions in different experimental conditions even though the trials for the different conditions were randomly interleaved. They concluded that the probability distortion function is not fixed for a participant but *dynamic*, changing systematically with task.

There is increasing evidence that dynamic remapping of representational range occurs along more abstract dimensions, such as value [25, 26, 27, 28, 29], numerosity [30, 31], relative frequency [32], and variance [33].

Zhang and Maloney [24] found that probability distortions could be well fit by linear transformations:

$$\lambda\left[\pi(p)\right] = \gamma\lambda(p) + (1-\gamma)\lambda(p_0)\ , \tag{2}$$

where $\lambda(p) = \log\frac{p}{1-p}$ is the log-odds [34] or logit function [35] and $\gamma > 0$ and $0 < p_0 < 1$ are free parameters. See Figure 1a for examples and Zhang and Maloney [24] for further examples, which include 20 datasets taken from 12 studies involving probability, relative frequency and

confidence, all the studies for which we could recover and analyze data. We caution that these Linear in Log-Odds (LLO) fits to data represent empirical regularities unmotivated by any theory.

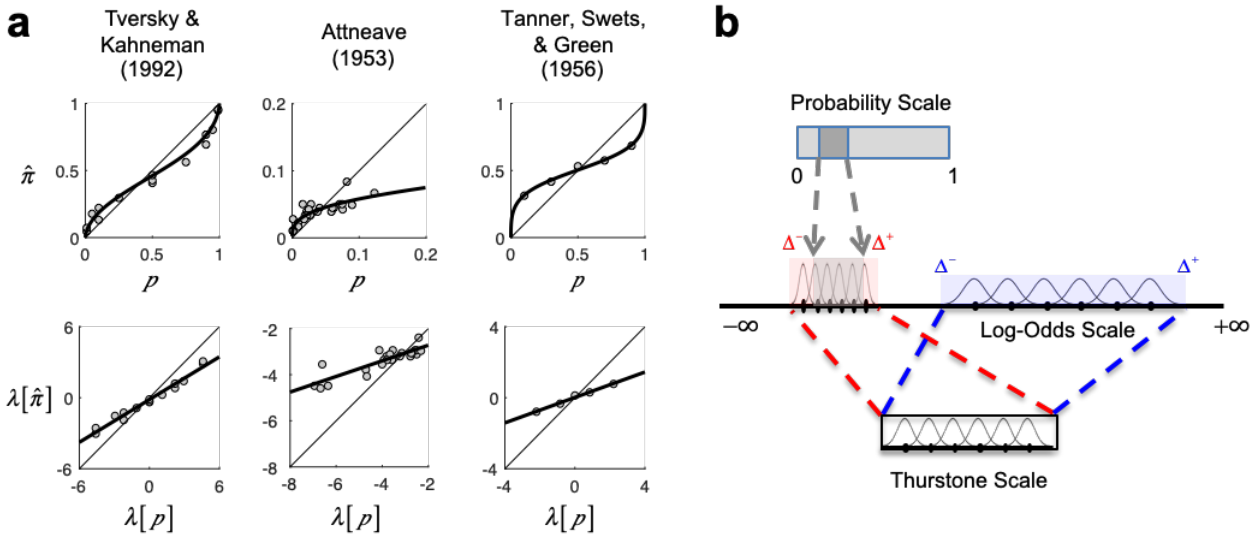

Figure 1. Motivations and intuitions for the bounded log-odds model (BLO). a. Observed probability distortions (top row) can be well captured by a linear fit on the log-odds scale (bottom row). The $\lambda[p]$ and $\lambda[\hat{\pi}]$ respectively denote the log-odds of the objective and subjective probabilities, p and. Circles denote data. Thick curves or lines denote the LLO fits. Tversky and Kahneman (1992): Subjective probability (decision weight) versus objective probability in decision under risk. Attneave (1953): Estimated relative frequency of letters in written English versus actual relative frequency. Tanner, Swets, & Green (1956), c.f. Green and Swets (1966/1974): Estimated probability of signal present versus objective probability in a signal detection task. Adapted from Zhang & Maloney (2012). **b.** Encoding on the Thurstone scale. A selected range of $[\Delta^-, \Delta^+]$ is encoded on the Thurstone scale with limited resolution. The smaller the range, the smaller the encoding variance.

Over the course of this article we will replace Eq. 2 by a new model, Bounded Log-Odds (BLO) based on theoretical considerations. We propose that probability distortion in both decision under risk and in judgment of relative frequency is fundamentally a consequence of a specific limitation on the dynamic range of the neural representation of probability which we identify. As a consequence of this limitation, human performance in a wide variety of tasks

(e.g. the Allais Paradoxes [11]) is necessarily sub-optimal by whatever measure is appropriate to each task.

However, the variation of probability distortion with task hints that—subject to "bounds" on probability representation—individuals might adaptively select the specific form of probability distortion that allows them to perform as well as possible in the task, a form of *bounded rationality* in Herbert Simon's sense [36]. What counts as "as well as possible" varies with task and for the tasks considered here we will identify the appropriate *decision variables.*

BLO is based on three assumptions:

1. log-odds representation

2. representation on a bounded Thurstone scale

3. variance compensation

We will use factorial model comparison [37] to separately test each of the three assumptions against plausible alternatives including LLO. We will then compare the performance of BLO to all previous models of decision under risk currently in the literature. The data used in all model comparisons are taken from a new experiment we report here and data from a previous article by Gonzalez and Wu [14]. Last, we will separately test the Maximum Mutual Information hypothesis that BLO serves to maximize the mutual Shannon information between objective decision variables and their internal representation, a form of bounded rationality.

Results

Assumptions of BLO

Assumption 1: log-odds representation.

In the BLO model probability, p, is internally represented as a linear transformation of log-odds,

$$\lambda(p) = \log\frac{p}{1-p}\ , \tag{3}$$

a one-to-one, increasing transformation of probability. A similar log-odds scale has been introduced by Erev and colleagues [38, 39] to explain the probability distortion in confidence ratings. Such noise-perturbed psychological scales date back to Thurstone [40].

Assumption 2: representation on a bounded Thurstone scale.

Thurstone [40] proposed several alternative models for representing subjective scales and methods for fitting a wide variety of data to such models. We are not concerned with methods for fitting data to Thurstone scales or their use in constructing attitude scales; we are only interested in Thurstone scales as convenient mathematical structures. We can think of the bounded Thurstone scale [40] as an imperfect neural device capable of storing magnitudes within a fixed range. We can encode a magnitude signal s anywhere in this range and later retrieve it. The retrieved value s, however, is perturbed by Gaussian noise with mean 0 and variance σ^2: we might store 0.5 and retrieve 0.63 or 0.48. The schematic Gaussian distributions in Figure 1b capture this representational uncertainty. For simplicity we assume that Gaussian error is independent, identically-distributed across the scale (Thurstone's Case V).

We can pick any interval on the log-odds scale and map it linearly to the Thurstone device. In Figure 1b we illustrate two choices. One represents a small range of the log-odds

scale using the full range of the Thurstone device, the other represents a larger range also mapped to the full range of the Thurstone device. The row of Gaussians on the two intervals of the log-odds scale symbolize the encoding-decoding uncertainty induced by the Thurstone scale.

The greater the log-odds range that needs to be encoded, the greater the density of the magnitudes along the Thurstone scale, and the greater the chances of confusion of nearby codes and vice versa. The challenge is to choose a transformation that is most beneficial to the organism.

Our concern is with the representation of probability, specifically in the form of log-odds. In mathematical notation we select an interval on the log-odds scale to be mapped to the full range of the Thurstone scale and in effect we confine the representation of log-odds λ to this interval:

$$\Gamma[\lambda]=\begin{cases}\Delta^-, & \lambda<\Delta^- \\ \lambda, & \Delta^-\le\lambda\le\Delta^+ \\ \Delta^+, & \lambda>\Delta^+\end{cases} . \qquad (4)$$

The value Δ^- will be mapped to the minimum on the Thurstone scale, the value Δ^+ to the maximum. The smaller the *half-range* $\Delta=\left(\Delta^+-\Delta^-\right)/2$ the smaller the uncertainty of the encoded and decoded values (Figure 1b) relative to the log-odds scale. We refer to $\Lambda(p)=\Gamma\left(\lambda[p]\right)$ as "truncated log-odds".

Assumption 3: variance compensation.

Besides the random error on the representational scale, there could be an additional uncertainty (variance) associated with the encoding of probability as when estimates of probability are used in place of exact probabilities (see Supplements S1 & S2). Less reliable estimates are in effect allotted less space on the representational scale.

The resulting variance compensation is a non-linear contraction or expansion around a fixed point Λ_0 on the log-odds scale. For any values of probability p ,

$$\Lambda^{\omega}(p) = \tau\left[\ \omega_p \Lambda(p) + (1-\omega_p)\Lambda_0\ \right] , \tag{5}$$

where $\Lambda(p)$ is, as before, truncated log-odds, ω_p is a measure of the reliability of probability that can vary with p, and the scaling factor τ adjusts the bounded interval to the fixed Thurstone scale. This transformation is an example of ***efficient encoding***: the transformation maximizes the information encoded by the scale. There is experimental evidence for variance compensation analogous to efficient coding in perception [41, 42, 43, 44] and its recent applications to value and probability [27, 29, 45]. See especially the review by Simoncelli and Olshausen [43]. The key idea is that a more precise coding of one stimulus at one point in a scale may come at the cost of a coarser coding of another stimulus at another point.

Overview of the experimental tests of BLO

To test BLO, we first performed a new experiment where each participant completed both a decision-making under risk (DMR) task and a judgment of relative frequency (JRF) tasks. We also re-analyzed the data of Gonzalez and Wu's [14] DMR experiment. Objective probabilities in these two representative tasks can be readily manipulated and subjective probabilities precisely estimated.

In Gonzalez and Wu [14], 10 participants were tested on 165 two-outcome lotteries, a factorial combination of 15 value sets by 11 probabilities (see Methods). Participants chose between lotteries and sure rewards so that their *certainty equivalent* (CE)—the value of sure reward that is equally preferred—to each lottery was measured. We refer to Gonzalez and Wu's [14] dataset as GW99, the set of lotteries included in which is large and rich enough to allow for reliable modeling on the individual level—as demonstrated in Gonzalez and Wu [14].

The decision variables for such a task are the certainty equivalents. To the extent that the participant can correctly order certainty equivalents, she can maximize expected utility.

We refer to the new experiment as Experiment JD (see Methods). In the experiment, each of 75 participants completed a DMR task whose procedure and design (Figure S1a) followed that of Gonzalez and Wu (1999) as well as a JRF task (Figure S1b) where participants reported the relative frequency of black or white dots among an array of black and white dots. The same 11 probabilities were used in the two tasks. The decision variables for such a task are the relative frequencies themselves.

By comparing the performance of individuals in two different tasks that involved the same set of probabilities, we hoped to identify the possible common representation of probability and how it may vary with task.

Based on the measured CEs (for DMR) or estimated relative frequencies (for JRF), we performed a non-parametric estimate and model fits for the probability distortion of each participant and each task (see Methods). Similar to previous studies of DMR [14] and JRF [24, 46], we found inverted-*S*-shaped probability distortions for most participants but also marked individual differences in both tasks (Figure 2abc). The DMR results of GW99 (Figure 2a) and Experiment JD (Figure 2b) were similar and were collapsed in further analysis whenever possible.

We used the non-parametric estimates to assess participants' probability distortions and we compared model fits with them. For an average participant (the last panels in Figure 2abc), the LLO and BLO models provided almost equally good fits. However, an examination of individual participants' probability distortions revealed that, compared to the LLO fit, the BLO fit captured observed individual differences considerably better. This observation can be quantified using the mean absolute deviations of the model fits from the non-parametric estimates (Figure 2de), which was significantly smaller for BLO than for LLO at $p = 0.25$ for

DMR (paired *t*-test, $t(84) = 2.40$, $P = 0.019$) and for all 11 *p*'s for JRF (paired *t*-tests, all $t(74) > 2.87$, $p < 0.006$).

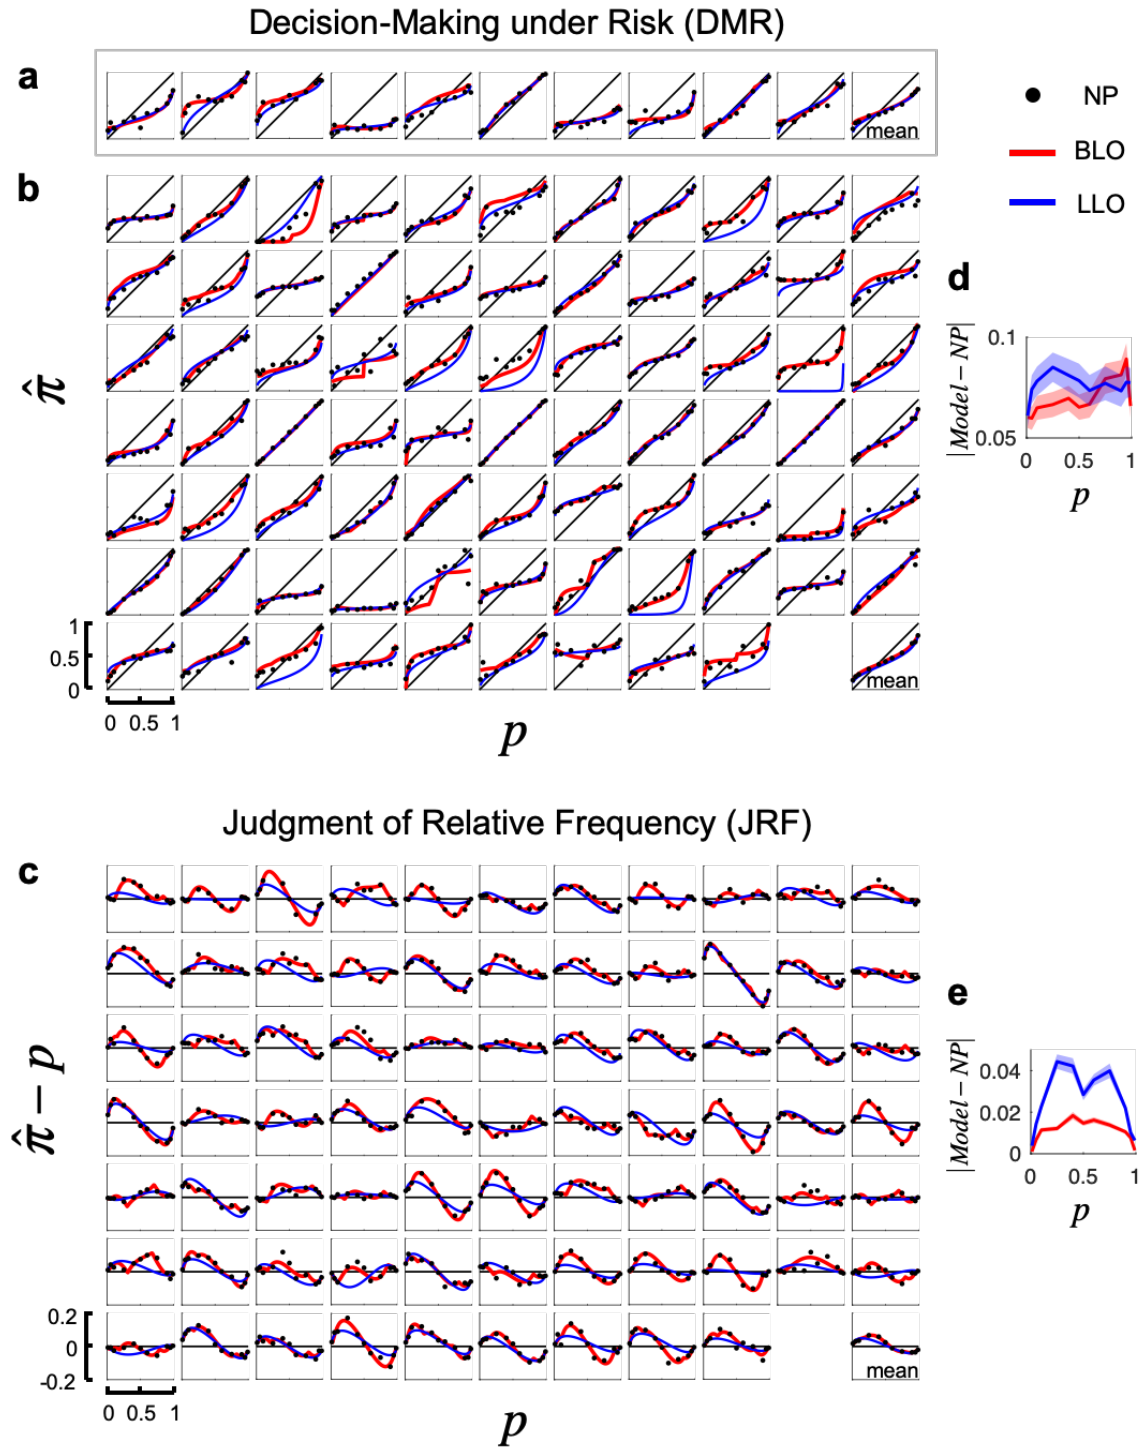

Figure 2. Comparison of model fits to non-parametric estimates of probability distortions. a. Reanalysis of DMR data from Wu & Gonzalez (1999). In the first 10 panels the data $\hat{\pi}(p)$ for

each participant is plotted versus p as black circles. The LLO fit to the participant's data is drawn as a blue contour, the BLO fit as a red contour. The last panel is the mean across participants. **b. DMR data from our experiment.** The format is identical with data and model fits for 75 participants. The last panel is the mean across participants. **c. JRF data from our experiment.** For each of the 75 participants we plot the residuals $\hat{\pi}(p) - p$ versus p to illustrate the small but patterned probability distortions found. We also plot the fits of LLO (blue) and BLO (red) to the residuals. Corresponding panels in **b** and **c** are for the same participant. Compared to the LLO fits (blue curves), the BLO fits (red curves) were overall in better agreement with the non-parametric estimates of probability distortions. **d. e.** Mean absolute deviations of the model fits from the non-parametric estimates are plotted against p, separately for DMR (**d**) and JRF (**e**). Shadings denote SE.

Factorial model comparisons

BLO is built on three assumptions: log-odds representation, boundedness, and variance compensation. To test these assumptions, we used *factorial model comparison* [37] and constructed 12 models whose assumptions differ in the following three "dimensions" (see Methods for details).

D1: scale of transformation. The scale of transformation can be the log-odds scale, the Prelec scale [47], or the linear scale based on the *neo-additive family* [48, 49, 50, 51, 52, 53, 54, 55].

D2: bounded versus bounds-free.

D3: variance compensation. The variance to be compensated can be the encoding variance (denoted $V(p)$) or constant (denoted $V = const$).

The models we considered are not all nested (see below) nor does factorial model comparison [37] require nested models. Both BLO and LLO are special cases of the 12 models, respectively corresponding to [log-odds, bounded, $V(p)$] and [log-odds, bounds-free, $V = const$].

For each participant, we fit each of the 12 models to the participant's CEs (for DMR) or estimated relative frequencies (for JRF) using maximum likelihood estimation (see

Supplement S3 for details). The Akaike information criterion with a correction for sample sizes, AICc [56, 57], was used for model selection. For a specific model, the ΔAICc was computed for each participant and each task as the difference of AICc between the model and the minimum AICc among the 12 models. A higher value of ΔAICc indicates a worse model fit.

For both DMR and JRF, BLO was the model of the lowest summed ΔAICc across participants (Figure 3ab). The results were similar for participants in different experiments (Figure S3). To see how well each of BLO's assumptions behaves compared to its alternatives, we divided the 12 models into model families by their assumptions on D1, D2, or D3 (e.g. the bounded family and the bounds-free family). We first calculated for each model the number of participants best fit by the model (lowest ΔAICc) and the exceedance probability from the group-level Bayesian model selection [58], which is an omnibus measure of the probability that the model is the best model among the 12 models. The summed number of best-fit participants is then plotted for each model family in Figure 3cd. For both DMR and JRF, the assumptions of BLO outperformed the alternative assumptions on each of the three dimensions, with the summed exceedance probability approaching 1.

We also performed model comparisons separately for participants with inverted *S*-shaped and participants with *S*-shaped distortions (Figure S4), and tested a range of additional models of decision under risk outside the framework of Cumulative Prospect Theory (CPT) we currently used (Figure S5). Again, the BLO model outperformed all alternative models (see Supplement S4 for details).

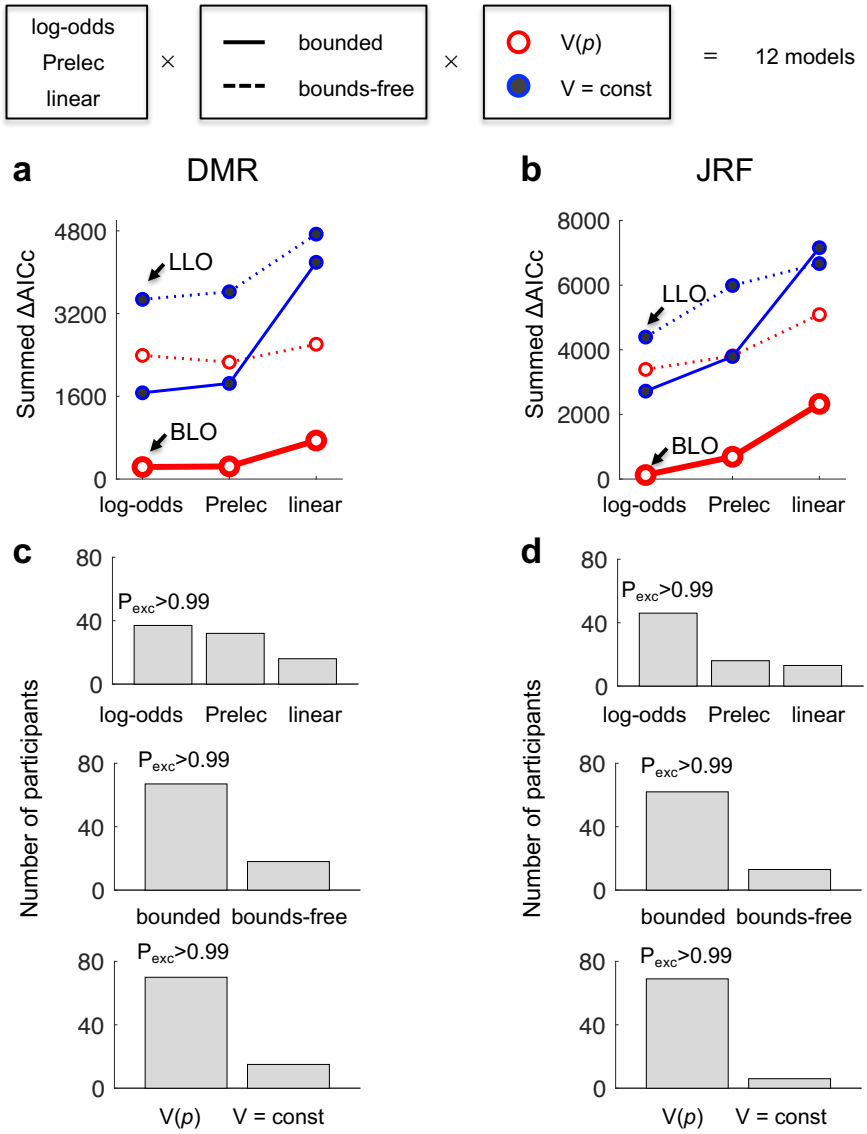

Figure 3. Results of factorial model comparison. We compared 12 models that differ on three dimensions ("factors") of assumptions: scale of transformation (log-odds, Prelec, or linear), boundedness (bounded or bounds-free) and variance compensation ($V(p)$ or $V = const$). BLO corresponds to [log-odds, bounded, $V(p)$]. LLO corresponds to [log-odds, bounds-free, $V = const$]. The summed ΔAICc across participants is plotted for each model, separately for DMR (**a.** 85 participants) and JRF (**b.** 75 participants). Lower values of ΔAICc are better. BLO outperformed the alternative models in both tasks. **c, d**. Each assumption of BLO (log-odds, bounded, and $V(p)$) also outperformed the alternative assumptions on its dimension. Each panel is for comparisons across one dimension, separately for DMR (**c**) and JRF (**d**). For a family of models with a specific assumption, shaded bars denote the number of participants best accounted by the model family. The P_{exc} above the highest bar denotes the summed exceedance probability of the corresponding model family.

Invariances across time and tasks

The scaling parameter τ maps the bounded range of $\left[\Delta^{-},\Delta^{+}\right]$ to the Thurstone scale of fixed length. An exploration of the estimated parameters of BLO (Supplement S5) shows that the product of τ and Δ was indeed close to a constant across participants, whose median estimate was 2.03 for DMR and 3.89 for JRF (Figure 4a). We call $\Psi \equiv \tau\Delta$ the ***Thurstone invariant***, the observed invariance of which suggests that the length of the Thurstone scale is similar for different individuals.

One might question whether the invariance of $\tau\Delta$ might be a consequence of model redundancy, that is, whether the apparent two parameters are effectively a single parameter and the variation observed is random. Instead, if Thurstone invariant Ψ is a personal signature, we should expect it to vary slightly from individual to individual but be positively correlated for the same individual across time and tasks.

Among the 75 participants of Experiment JD, 51 participants completed two sessions on two different days, which allowed us to evaluate the consistency of Thurstone invariant across time. As expected, we found significant positive correlation between Session 1's and Session 2's Ψ (Figure S6) for both DMR (r_s = 0.57, one-tailed $P < 0.001$) and JRF (r_s = 0.80, one-tailed $P < 0.001$).

We also found a modest but significant positive correlation between a participant's Ψ in JRF and the participant's Ψ in DMR (Figure 4b): r_s = 0.21, one-tailed P = 0.039. Given that the two tasks involve entirely different responses and processing of probability information, such across-task correlation is surprising and suggests a common constraint underlying the probability representations in different tasks. This constraint is probably tighter in tasks that

demand more working memory, corresponding to a smaller Ψ for DMR (when value as well as probability needs to be represented) than that for JRF.

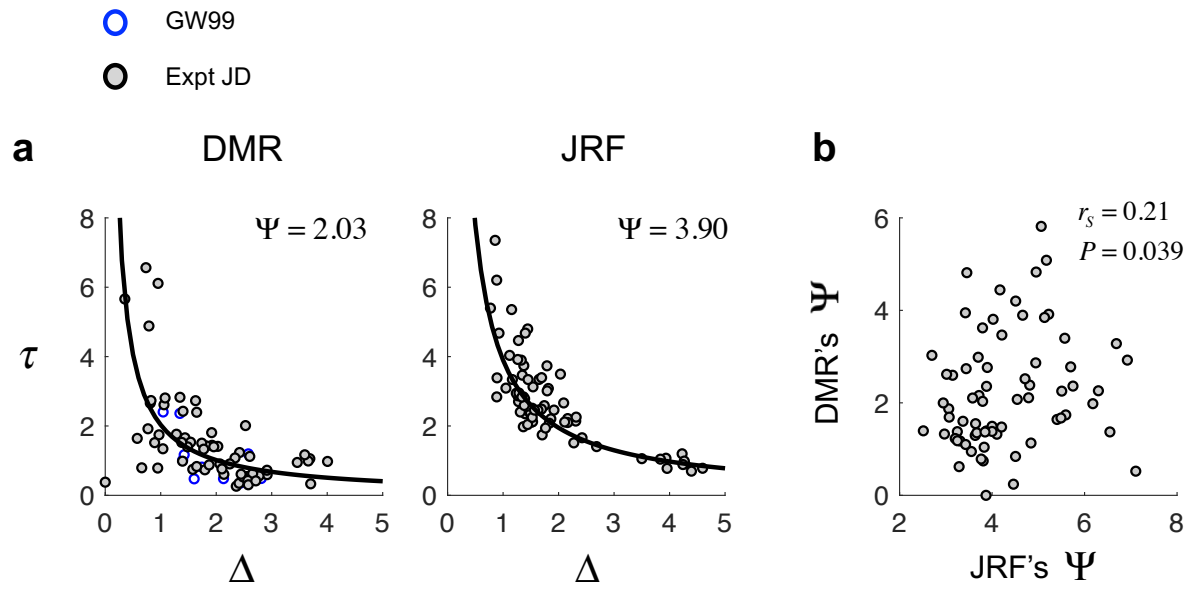

Figure 4. The Thurstone invariant Ψ. Each circle is for one participant. **a.** The product of τ and Δ —the Thurstone invariant Ψ —was nearly invariant across participants. The black curve denotes $\Psi \equiv \tau\Delta = C$, where *C* was the median value across participants. (11/85 data points are outside the plot range of the DMR panel.) **b.** The estimated Ψ was positively correlated across tasks, for participants who completed both the JRF and DMR tasks. (5/75 data points are outside the plot range.) The r_s on the plot refers to Spearman's correlation coefficient, which is robust to outliers, and P is right-tailed.

Maximizing mutual information

Participants had a bounded log-odds representation $\left[\Delta^-,\Delta^+\right]$ that corresponds to a probability range far narrower than the range of objective probabilities ($\left[0.01,0.99\right]$). As we will see below, this choice of $\left[\Delta^-,\Delta^+\right]$ conforms with the Maximum Mutual Information hypothesis.

For our purposes, the efficiency of coding can be quantified by the mutual information between stimuli $s_1,\cdots,s_n$ and responses $r_1,\cdots,r_n$:

$$I_m = \sum_{i=1}^{n} P(s_i,r_i)\log_2 \frac{P(s_i,r_i)}{P(s_i)P(r_i)}, \quad (6)$$

where $P(s_i)$ denotes the probability of occurrence of a specific stimulus s_i, $P(r_i)$ denotes the probability of occurrence of a specific response r_i, and $P(s_i,r_i)$ denotes the conjoint probability of the co-occurrence of the two. Stimuli and responses refer to objective and subjective relative frequencies in JRF and refer to expected values and certainty equivalents in DMR. For a specific task and BLO parameters, we could use the BLO model to generate simulated responses and computed expected mutual information using a Monte Carlo method (see Methods).

For a virtual participant endowed with median parameters, we evaluated how the expected mutual information in DMR or JRF varied with Δ^+ and Δ^-, given the invariance of Ψ. We found that the expected mutual information varied non-monotonically with the values of Δ^+ and Δ^- (Figure 5a) and there is evidently a unique maximum over the range considered. For both DMR and JRF, the observed median values of Δ^+ and Δ^- (marked by solid red circles) were close to the values maximizing the expected mutual information: the mutual information associated with the observed Δ^+ and Δ^- were lower than maximum only by 0.14% for DMR and 0.53% for JRF. In contrast, if no bounds had been imposed on the probability range of $[0.01,0.99]$ (i.e. $\Delta^- = -4.6$, $\Delta^+ = 4.6$), the mutual information would be 6.23% and 14.1% lower than maximum, respectively for DMR and JRF.

The observed Δ^+ and Δ^- of individual participants were highly symmetric around 0 (i.e. symmetric around 0.5 on the probability scale) in the JRF task but more variant in the DMR task. As Figure 5b shows, this difference may also be driven by mutual information

maximization: To achieve no less than 95% of maximum mutual information, Δ^+ and Δ^- could tolerate a much larger deviation from symmetry in DMR than in JRF.

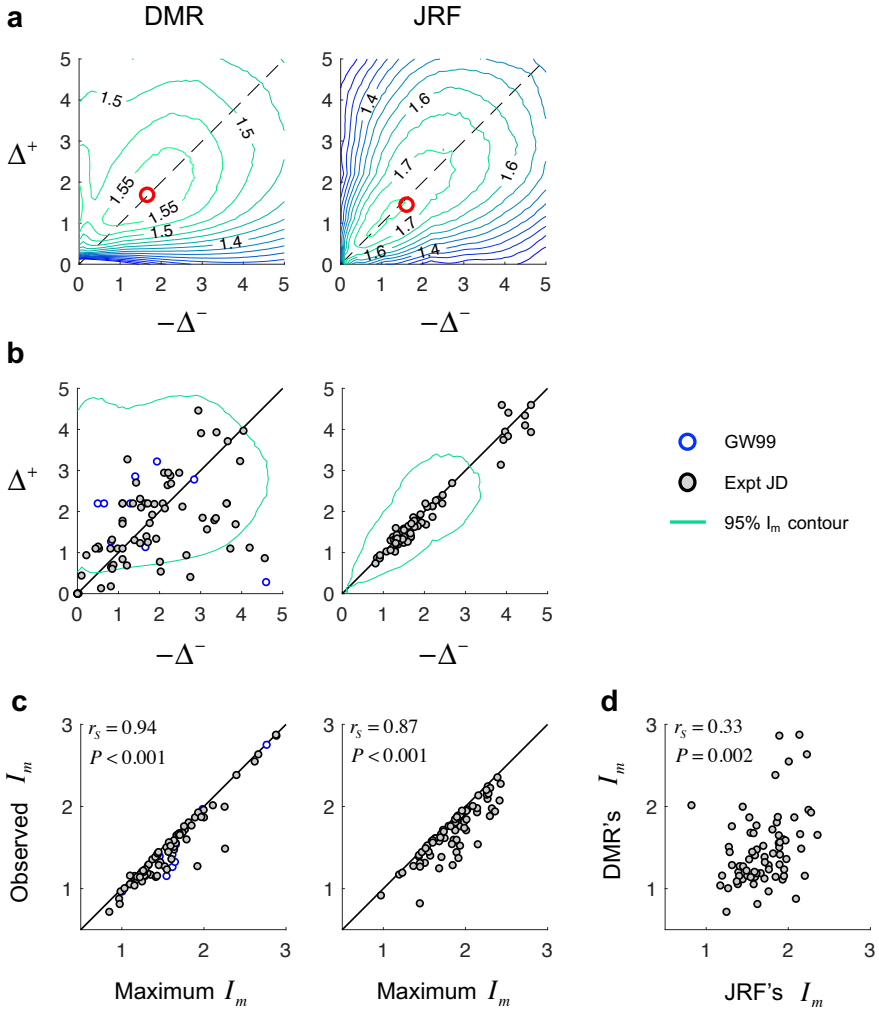

Figure 5. Choice of Δ^- and Δ^+ as mutual information maximization. a. Expected mutual information between stimuli and responses (in bits) is plotted against $-\Delta^-$ and Δ^+ as contour map, separately for DMR and JRF. Higher values are coded as more reddish and lower values as more bluish. For both tasks, the observed median value of $\left(-\Delta^-,\Delta^+\right)$ (marked by the red circle) was close to maximizing the expected mutual information. **b.** Observed Δ^+ versus $-\Delta^-$ for individual participants, compared to the contour of 95% maximum mutual information (green curve). The observed $\left(-\Delta^-,\Delta^+\right)$ had a more pronounced deviation from symmetry in DMR than in JRF. **c.** Observed mutual information was positively correlated with maximum mutual information (6/85 data points are outside the range of the DMR plot.) The values are clustered just below the identity lines in both plots. **d.** Observed mutual information was positively correlated across tasks. (3/75 data points are outside the plot range.) The r_s on the plot refers to Spearman's correlation coefficient and P is right-tailed.

We further computed the maximum mutual information for each participant based on the participant's BLO parameters, given that Δ^+ and Δ^- are allowed to vary. The maximum mutual information was positively correlated with the observed mutual information for both DMR (r_s = 0.94, one-tailed P < 0.001) and JRF (r_s = 0.87, one-tailed P < 0.001), consistent with the mutual information maximization hypothesis (Figure 5c). The diagonal lines in both plots of Figure 5c are bounds on possible performance and the data in both plots are clustered below them.

Individual participants' mutual information was positively correlated across time (Figure S7, DMR: r_s = 0.53, one-tailed P < 0.001; JRF: r_s = 0.53, one-tailed P < 0.001) and across tasks (Figure 5d, r_s = 0.33, one-tailed P = 0.002). Since there was little across-task correlation in noise variance (r_s = 0.066, one-tailed P = 0.29), the across-task correlation in mutual information could not be due to correlated levels of response noises. Instead, it implies a common information processing capacity for different tasks.

Discussion

We presented a model (Bounded Log-Odds) of probability and relative frequency distortion and tested it experimentally in an experiment with two conditions. In one condition participants made judgments of relative frequency (JRF) and, in the other, decisions under risk (DMR). Each participant completed both conditions, allowing us to compare performance in the two tasks within participant. We also reported a reanalysis of a data set from a DMR experiment carried out by Gonzalez and Wu [14].

The BLO Model is intended to model performance in both tasks and it is the first model that attempts to do so. It is based on three assumptions: log-odds representation, boundedness, and variance compensation. We independently tested each of these assumptions using factorial model comparison to verify that they are all essential to fitting human data. If we replace any assumption by the alternatives we considered, the resulting model is strictly inferior to BLO. We then compared BLO with all of the other models in the literature intended to account for probability distortion. BLO outperformed all these models in accounting for our experimental results as well as the data of Gonzalez and Wu [14].

We further show that human performance comes close to maximizing **the mutual information** between decision variables in a task and their imperfect internal representations. The decision variables in the DMR task are certainty equivalents (CEs). To the extent that the participant can correctly order the distorted CEs, she will maximize expected value. The decision variables in the JRF task are the relative frequencies themselves.

In the four experiments we report, participants chose probability distortions consistent with BLO and also with maximizing mutual information. Two recent articles use the same criterion (maximum mutual information) to model human encoding of value [29] or to re-interpret the context effects of decision under risk [45]. These articles taken together are consistent with a claim, supported by considerable experimental data, that many observed failures in DMR can

be viewed as attempts to compensate for immutable limits in cognitive processing in order to preserve Shannon information, a form of bounded rationality [36].

There are many theoretical models intended to account for inverted-*S*- or *S*-shaped probability distortion: the power model of proportion judgment [59, 60], the support theory model of probability judgment [61, 62], the calibration model [63], the stochastic model of confidence rating [38, 39], and the adaptive probability theory model of decision under risk [64]. However, almost all these models were proposed for one specific type of task and are not intended as general explanations for observed distortion of probability and relative frequency. Neither do they explain why participants exhibit different probability distortions in different tasks and task conditions. There was even a belief, at least in decision under risk, that the parameters of distortion should be specific to each participant but constant across all tasks [65].

In contrast, BLO models a common mechanism underlying all probability distortion, where we identified two constraints—boundedness and compensation for representational uncertainty (variance)—that are pervasive in models of cognitive and perceptual tasks [37, 44, 66, 67]. We found that BLO can be used to estimate an individual's probability distortion in one task and to some extent predict the same individual's performance in another task. BLO also accounts for variation in the "slope" of the distortion function with task or different settings of the same task. We next describe some of the implications of BLO.

Discontinuities at p = 0 and p = 1

BLO and any model based on the boundedness assumption predict that $\pi(0) > 0$ and $\pi(1) < 1$, that is, probability distortion with discontinuities at $p = 0$ and $p = 1$. Such discontinuities are also found in the neo-additive family of weighting functions [54], but are not found in other, widely accepted families of probability distortion such as LLO [14, 24] and Prelec's

family [47]. Kahneman and Tversky's original Prospect theory [12] included similar discontinuities in probability weighting functions (see their Figure 4).

The bounded ranges of probability represented on the Thurstone scale according to the BLO model fits are fairly limited, approximately [0.15, 0.85]. Given that the occurrence of probabilities as extreme as 0.05 and 0.95—or even 0.01 and 0.99—is not uncommon in laboratory tasks or real life, bounding is likely to exert detectable influences on probability representation and performance under many circumstances. Indeed, there are clues indicating boundedness in previous studies. For example, Yang and Shadlen [68] studied monkeys' probabilistic inference and found that the strength of a specific evidence perceived by the monkey was, in general, proportional to the objective log-odds of the evidence. But for "sure evidence" that corresponded to minus or plus infinity in log-odds, the subjective log-odds were bounded, equivalent to [0.15, 0.81] and [0.30, 0.64] in probability for the two tested monkeys.

For the DMR task, where probability is explicitly defined and no explicit sampling process seems to be involved, we still found that the slope of probability distortion relies on a $p(1-p)$ term, varying with p. It is as if people are compensating for the variation of a virtual sampling process [69], or for the variation caused by Gaussian noise on the Thurstonian log-odds scale Supplement S1, see also [70]. Lebreton et al. [70] show that a generalized form of $p(1-p)$ is correlated with the confidence of value or probability perception and is automatically encoded in the ventromedial prefrontal cortex (vmPFC) of the human brain. Under certain circumstances, such variance compensation may result in counterintuitive non-monotonic probability distortion that is indeed empirically observed (see Supplement S6).

The Thurstone invariant

According to BLO, a bounded range $\left[\Delta^{-},\Delta^{+}\right]$ of the log-odds scale is mapped to a Thurstone scale of fixed length by a scaling factor τ. Indeed, for both JRF and DMR, we found $\tau\Delta$, which corresponds to the half length of the Thurstone scale, was almost invariant across participants. We defined $\Psi \equiv \tau\Delta$ and termed it the Thurstone invariant.

Though individual differences in Ψ are small, an individual's Ψ at one time point is positively correlated with the individual's Ψ at another time point. The individual's Ψ is also positively correlated across tasks. Thus the Thurstone invariant proves to be a personal signature that imposes a common constraint on the representation of probability distortion in different tasks.

Predicting the slope of probability distortion

The invariance of the Thurstone invariant Ψ implies that when the encoded range Δ is narrower, the scaling parameter τ would be greater, leading to a greater slope of probability distortion. Meanwhile, mutual information maximization requires Δ to scale with the range of probabilities in the stimuli. Thus BLO predicts that the narrower the probability range of the stimuli, the greater the slope of distortion.

We performed the following meta-analysis on previous DMR studies to test this prediction. Fox and Poldrack [71, Table A.3] summarized the results of a number of decision-making studies that were modeled in the framework of Prospect Theory. In Fox and Poldrack's list, we identified the studies where the gamble set was explicitly defined and each gamble consisted of two outcomes that could be denoted $\left(x_1, p; x_2, 1-p\right)$ (see Supplemental Table S3 for the 12 studies included). Though different functional forms—LLO, Prelec's [47] one-parameter and two-parameter functions, and Tversky and Kahneman's [13] weighting function—had been assumed in different studies, all had a parameter for the slope of probability distortion that is roughly

equivalent to the γ in LLO. For each study, we computed the standard deviation of $\lambda(p)$'s distribution as a measure of the probability range of stimuli. Consistent with the BLO prediction, we found this measure was significantly negatively correlated with the slope of probability distortion (Figure 6), r_s = –0.65, one-tailed P = 0.010.

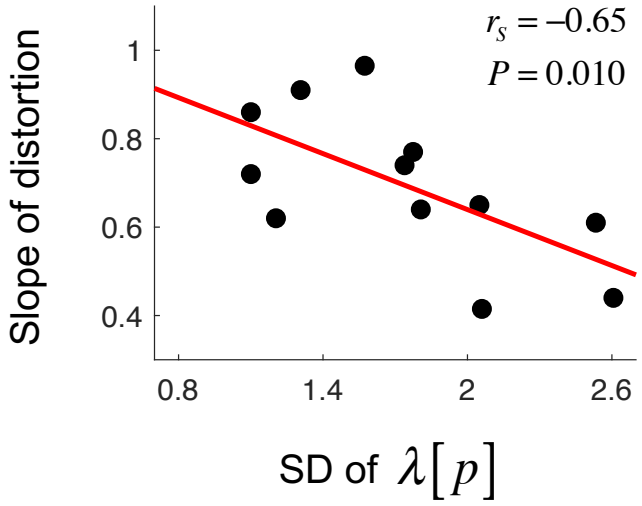

Figure 6. Meta-analysis of previous studies supporting BLO's prediction on the slope of probability distortion in decision under risk. The estimated slope of probability distortion ($\hat{\gamma}$) is plotted against the standard deviation of the objective log-odds ($\lambda(p)$) of the gamble set, where p denotes the probability for the higher outcome of a two-outcome gamble, $(x_1, p; x_2, 1-p)$. Each data point is for one published study. The red line denotes the regression line. The correlation is negative and significant. We describe the selection of studies in the text. See Table S3 for a full list of the studies. That the slope of distortion decreases with the standard deviation of $\lambda(p)$ is consistent with the prediction of BLO.

The crossover point

A puzzle we did not address earlier concerns the crossover point of probability distortion (i.e. the point on the distortion curve where overestimation changes into underestimation or the reverse). It has been frequently observed that the crossover point is near 0.5 for the JRF task [24] but approximately 0.37 for the DMR task [47]. That is, the probability distortion is symmetric around 0.5 in the former but asymmetric in the latter. There are plausible reasons to have

symmetry, but why asymmetry? Here we conjecture that the asymmetry is also driven by the maximization of mutual information, which, for the DMR task, means to have the CEs of different gambles as discriminable as possible. Following conventions [14, 71] and for parsimony, we had assumed a uniform Gaussian noise on the CE scale. However, larger CEs may tend to be associated with higher variances, according to Weber's law [72]. To compensate for this, more of the representational scale should be devoted to larger probabilities and thus to the larger CEs associated with them. Indeed, adopting a smaller crossover point (i.e. less than 0.5) would map larger probabilities to a longer range of subjective probabilities and effectively implement such a strategy of probability representation.

Open questions and future directions

The judgment of relative frequency and decision under risk are the only two tasks where BLO and its assumptions have been tested, but these two tasks together represent a vast body of previous research. The model may be applied to a wider range of tasks involving frequency and probability. It will likely shed light on the common and distinctive mechanisms of probability distortion in different tasks.

What determine the slope and crossover point of probability distortion in a specific task? Why may the parameters of probability distortion change from task to task and from individual to individual? In the present study we have provided a tentative answer: they change because the brain actively compensates for its own fixed limitations.

Important questions for future research also include: How may probability distortion change from trial to trial? We conjecture that the human representation of probability can adapt to the environment, in the spirit of efficient coding [41, 42, 43]. The current version of BLO is a stationary model, whose prediction will not change with time or experience. In contrast, non-

stationarity has been identified in probability distortion for both the judgment of relative frequency [24] and decision under risk [73].

Can the BLO model apply to decisions among complex gambles with more than two non-zero outcomes? In theory, BLO just specifies the probability weighting function in Cumulative Prospect Theory [13] and can apply to any circumstances where Prospect Theory applies. But it is still an empirical question whether probability distortion for gambles with more than two non-zero outcomes can be predicted by BLO.

We chose not to test "decision from experience" [20]—another important form of decision-making—because the decision from experience task does not require that the decision maker estimate the frequency of items [19, 74]. The decision maker may estimate the multinomial distribution of rewards in a card deck—or she may simply register reward and punishment and base her decision on a form of reward averaging or reinforcement learning. The results of the comprehensive model competition of Erev, Ert [75] are consistent with this claim. More recently, there has been neuroimaging evidence that human decisions from experience may be based on the retrieval of individual samples from past experience [76, 77]. If the decision maker does not estimate relative frequency then BLO does not apply.

A final note: Kahneman and Tversky's original Prospect Theory contained the assumption that decision makers would first interpret ("edit") available information [12]. In this initial *editing* stage they might, for example, convert the probability 0.31317 to the more tractable 1/3. Only then would they assign prospect values to lotteries in the second, *evaluation* stage. In presenting the BLO Model we focus on evaluation. Still, nothing about the theory would preclude adding an editing phase or discretizing the representation of probability if justified by empirical results.

Methods

Experiment

Experiment JD was approved by the Institutional Review Board of School of Psychological and Cognitive Sciences at Peking University. All participants gave written informed consent in accordance with the Declaration of Helsinki. Each participant performed two tasks: Decision-Making under Risk (DMR) and Judgment of Relative Frequency (JRF).

The procedures and designs of the DMR task were the same as those of Gonzalez and Wu [14], except that payoffs in the gambles were in RMB instead of in USD. On each trial (Figure S1a), participants were presented with a two-outcome gamble $(x_1, p; x_2, 1-p)$ and tables of sure amounts of rewards. They were asked to check on each row of the tables whether they preferred the gamble or the sure amount. The range of the sure amounts started with $[x_2, x_1]$, and was narrowed down in the second table so that we could estimate participants' certainty equivalent (CE) for the gamble. There were 15 possible outcome pairs (x_1, x_2) : (25, 0), (50, 0), (75, 0), (100, 0), (150, 0), (200, 0), (400, 0), (800, 0), (50, 25), (75, 50), (100, 50), (150, 50), (150, 100), (200, 100), (200, 150). There were 11 possible probabilities: 0.01, 0.05, 0.1, 0.25, 0.4, 0.5, 0.6, 0.75, 0.9, 0.95, 0.99. A full combination of them resulted in 165 different gambles used in the experiment.

The stimuli and procedures of the JRF task followed Zhang and Maloney (2012). On each trial (Figure S1b), participants were presented with an array of black and white dots and reported their estimate of the relative-frequency of black or white dots by clicking on a horizontal bar with tick marks from 0 to 100%. Each participant was randomly assigned to report the relative frequency either for the black or for the white dots. The objective relative frequency of JRF was chosen from the same 11 possible values as its counterpart in DMR. The total number of dots (numerosity) in a trial was varied across trials, which could be 200,

300, 400, 500, or 600. The dots in each display were distributed within a circular area of 12° diameter or a square area of 17°×17° diameter.

Experiment JD (a total of 75 participants) consisted of two sub-experiments, JDA (51 participants, 20 male, aged 18 to 29) and JDB (24 participants, 10 male, aged 18 to 27). Six additional participants failed to complete the experiment for technical or personal reasons. Each session had 11 (probability) × 15 (outcome pair) = 165 DMR trials and 11 (probability) × 5 (numerosity) × 6 = 330 JRF trials, which took approximately two hours. In Experiment JDA, each participant completed two sessions on two different days, so that we could evaluate the consistency of their performance. Trials from the two tasks were randomly interleaved. In Experiment JDB, each participant completed only one session, during which one task preceded the other, with DMR first for half of the participants and JRF first for the other half. Similar patterns of probability distortions (Figure 2bc, first 51 panels for Experiment JDA and last 24 panels for Experiment JDB) and results of model comparisons (Figure S3) were found for participants in the two sub-experiments. Thus we collapsed the two sub-experiments in our analysis whenever applicable.

The relationship between BLO and LLO

We compare LLO (Eq. 2)

$$\lambda\left[\pi(p)\right]=\gamma\lambda(p)+(1-\gamma)\lambda(p_0)$$

to BLO rewritten:

$$\begin{aligned}\Lambda^{\omega}(p)&=\tau\left[\ \omega_p\Lambda(p)+(1-\omega_p)\Lambda_0\ \right]\\&=\tau\omega_p\Lambda(p)+(1-\tau\omega_p)\frac{\tau(1-\omega_p)}{1-\tau\omega_p}\Lambda_0\end{aligned}\quad . \tag{7}$$

Over the range $\left[\Delta^{-},\Delta^{+}\right]$ truncated log-odds coincides with log-odds: $\Lambda(p)=\lambda(p)$.

Comparing the equations above, over the range $\left[\Delta^{-},\Delta^{+}\right]$ $\tau\omega_p$ replaces γ and $\frac{\tau(1-\omega_p)}{1-\tau\omega_p}\Lambda_0$ replaces $\lambda(p_0)$. In LLO, both γ and p_0 are fixed but in BLO, however, the reliability parameter ω_p may vary with the value of p depending on the model of variance appropriate to a given task. If a specific dataset is generated by BLO but fitted by LLO, we would expect that the estimates of γ and p_0 would change with experimental conditions as predicted by BLO. Consequently, we can fit LLO to data and look for the pattern of deviations in the fitted coefficients predicted by BLO, a test of the BLO Model (e.g. Figure 6).

Applying BLO to JRF

We need additional assumptions when applying BLO to the JRF experiments. One of the key assumptions of BLO is variance compensation and, to apply BLO, we need to specify a model of the participant's sampling process and the variance of the resulting estimates. First, we assume that humans may not have access to all the tokens presented briefly in a display or in a sequence, due to perceptual and cognitive limits [78, 79]. Instead, they take samples from the population and are thus subject to the randomness associated with sampling. Within BLO, probability distortion arises in part from a compensation for the sampling noise captured in our model by the reliability parameter ω_p.

Denote the total number of dots in a display as N and the relative frequency of black dots as p. Suppose a sample of n_s dots is randomly drawn from the display. We assume that the sampling is without replacement. That is, the same dot will not be drawn twice during one sampling, which is reasonable in our case. As a result, the variance of $\hat{p}$ requires a correction for finite population [80] (see Supplement S2 for the derivation):

$$V(\hat{p}) = \frac{p(1-p)}{n_s}\frac{N-n_s}{N-1} \tag{8}$$

The finite population correction is intuitive: the larger the sample size relative to the population, the smaller the variance. When $n_s = N$, i.e. when the whole population is included in the sample, we should have $\hat{p} = p$ for each sample and thus $V(\hat{p}) = 0$. At the other extreme, when $n_s = 1$, sampling without replacement is equivalent to sampling with replacement, the familiar $p(1-p)$. The BLO variance correction is a weighted mixture of an estimate based on the sample and an "anchor" Λ_0 which may also be stochastic, with its own variance $V(\Lambda_0)$. The optimal weight for combining the two is

$$\omega_p = \frac{V(\hat{p})^{-1}}{V(\hat{p})^{-1} + V(\Lambda_0)^{-1}} \ . \tag{9}$$

Denote $\kappa = V(\Lambda_0)^{-1}$, the *anchor precision parameter*. This equation can be rewritten as

$$\omega_p = \frac{1}{1 + \kappa V(\hat{p})} \ . \tag{10}$$

Finally, we assume that encoded values $\pi(p)$ are perturbed by additive Gaussian error (Thurstone, 1927), updating Eq. 5 to

$$\lambda[\pi(p)] = \tau\left[\ \omega_p \Lambda(p) + (1-\omega_p)\Lambda_0\ \right] + \varepsilon_\lambda \ , \tag{11}$$

where ε_λ is Gaussian error on the log-odds scale with mean 0 and variance σ_λ^2.

Applying BLO to DMR

To model $\pi(p)$, BLO's assumptions for different tasks are the same, except that encoding variance is task-specific. Probability is described explicitly in DMR and there need be no uncertainty about its value. Participants' choices suggested, however, that they were

still compensating for some kind of encoding uncertainty that varies with the value of probability. Gaussian encoding noise on the log-odds scale, when transformed back to the probability scale, results in variance that is approximately proportional to $p(1-p)$ (see Supplement S1 for proof). The reliability parameter in Eq. 5 is thus:

$$\omega_p = \left[1+\kappa p(1-p)\right]^{-1}, \quad (12)$$

where κ is a free parameter. This same equation can be reached if, alternatively, we assume that participants were compensating for a virtual sampling process (the $\frac{1}{n_s}\frac{N-n_s}{N-1}$ term in Eq. 8 can be assimilated into κ for constant N and n_s). Compensation for virtual sampling was also assumed in some previous theories on probability distortion [64, 81].

Any lottery in GW99 or Experiment JD can be written as $(x_1,p;x_2,1-p)$, which offers the value x_1 with probability p and otherwise x_2, with $x_1 > x_2 \geq 0$. For each participant, we modeled the certainty equivalent (CE) of each lottery using Cumulative Prospect Theory (CPT) [13] and assumed a Gaussian error term on the CE scale, as in Gonzalez and Wu [14]:

$$CE = U^{-1}\left[\pi(p)U(x_1)+\left(1-\pi(p)\right)U(x_2)\right]+\varepsilon_{CE}, \quad (13)$$

where $U(.)$ denotes the utility function, $U^{-1}(.)$ denotes the inverse of $U(.)$, $\pi(p)$ denotes the probability distortion function, and ε_{CE} is a Gaussian random variable with mean 0 and variance σ^2_{CE}. The utility function for non-negative gains alone (none of the lotteries involved losses) was assumed to be a power function with parameter $\alpha > 0$:

$$U(x) = x^{\alpha}. \quad (14)$$

Non-parametric estimation of probability distortion

A non-parametric estimation of probability distortion is plotted in Figure 2 for each participant and each task. For JRF, where participants explicitly estimated their subjective relative frequency, the non-parametric estimation $\hat{\pi}_{NP}(p)$ for a specific p was simply the participant's mean estimate across trials, averaged on the log-odds scale:

$$\hat{\pi}_{NP}(p) = \lambda^{-1}\left[\frac{1}{m}\sum_{t=1}^{m}\lambda\left[\hat{\pi}_t(p)\right]\right], \quad (15)$$

where $\hat{\pi}_t(p)$ denotes the participant's estimate of relative frequency on trial *t*, $t = 1,2,...,m$.

For the $\hat{\pi}_{NP}(p)$ of DMR, we modeled participants' CE in the framework of CPT as we did for BLO and LLO fits (Eq. 13) except that no functional form was assumed for the probability distortion function. Instead, the $\hat{\pi}_{NP}(p)$ for each of the 11 *p*'s was fitted as a free parameter. The same power functional form was assumed for the utility function in the non-parametric estimation as in the model fits to minimize possible differences irrelevant to probability distortion. This procedure was different from the non-parametric method of Gonzalez and Wu [14], where no functional forms were assumed for either probability distortion or utility. We verified in the GW99 dataset that our non-parametric estimation of probability distortion led to similar results as Gonzalez and Wu [14] (Figure S2).

Factorial model comparisons

We used *factorial model comparison* [37] to separately test the assumptions of BLO, comparing alternative models that differ in the following three "dimensions".

D1: scale of transformation. We considered two alternatives to the log-odds scale: the "Prelec scale" and the linear scale.

The Prelec scale is derived from Prelec's function [47]:

$$\pi(p) = \exp\left(-\beta(-\log p)^{\eta}\right) \tag{16}$$

with η and β as free parameters. LLO and the Prelec families both are among the probability weighting functions that typically fit best to data [82, 83]. They are difficult to distinguish empirically [15, 82].

Taking logarithms and negating twice on both sides of Prelec's function, we can see that Prelec's function is equivalent to a linear transformation

$$\lambda'\left[\pi(p)\right] = \eta\lambda'[p] - \log\beta \tag{17}$$

on the Prelec scale

$$\lambda'(p) = -\log(-\log p) \ . \tag{18}$$

The functional form of the linear scale is based on the *neo-additive family* [48, 49, 50, 51, 52, 53, 54, 55] which refers to a linear transformation of probability except that it may have discontinuities at the extremes to ensure $\pi(p)$ is within [0,1]:

$$\pi(p) = \begin{cases} 0, & if\ p = 0\ or\ kp + b < 0 \\ 1, & if\ p = 1\ or\ kp + b > 1 \\ kp + b, & otherwise \end{cases} \tag{19}$$

where k and b are free parameters. The linear scale is accordingly defined as

$$\lambda^{\circ}(x) = \begin{cases} 0, & if\ x < 0 \\ 1, & if\ x > 1 \\ x, & otherwise \end{cases} \ . \tag{20}$$

For models that use the Prelec scale, we simply replaced the log-odds and its reverse transformation with $\lambda'(p)$ and its reverse. For models that use the linear scale, the log-odds transformation was replaced by $\lambda^{\circ}(p)$ but there was no inverse transformation, because $\lambda^{\circ}(p)$ is not invertible, and no need for an inverse

transformation, because $\lambda^{\circ}(p)$ corresponds to probability. The noise term ε_{λ} was always added to the log-odds of $\pi(p)$ to maintain a fair comparison between models. D1 does not influence the number of free parameters of the model.

D2: bounded or bounds-free. We assumed a bounding operation in BLO (Eq. 4), where probabilities outside specific boundaries are truncated to the boundaries. We considered alternative models that are bounds-free. Bounds-free models would not include the bounds parameters, Δ^{-} and Δ^{+}.

D3: variance compensation. In BLO, ω_p is inversely related to $V(\hat{p})$ (Eq. 10) so that the encoding variance is appropriately compensated in the framework of Bayesian inference. Alternatively, we considered ω_p as a constant that does not change with p, as if the compensated variance is constant. When ω_p is constant, the potentially non-linear transformation of BLO (Eq. 5) is reduced to the linear transformation of LLO, as a re-parameterization would reveal (with one free parameter reduced). The two forms of variance compensation will be referred to as $V(p)$ and $V = const$.

The three dimensions—D1, D2, and D3—correspond to the three assumptions of BLO.

We did not list the presence or absence of the scaling factor τ as a possible dimension for both theoretical and practical reasons. On one hand, τ is required to map the bounded interval to the fixed Thurstone scale. On the other hand, the absence of scaling would preclude a greater-than-one $\hat{\gamma}$. There is evidence from several laboratories other than our own experiments that $\hat{\gamma}$ can be greater than 1 [16, 84, 85, 86].

The three dimensions are independent of each other, analogous to the different factors manipulated in an experiment with a factorial design. In total, we tested 3 (D1: log-odds,

Prelec, linear) × 2 (D2: bounded, bounds-free) × 2 (D3: $V(p)$, $V = const$) = 12 different models. LLO is among the 12 models (when D1=log-odds, D2=bounds-free, D3=$V = const$).

Computation of mutual information

For a specific real or virtual participant in a specific task, we used the BLO model to generate simulated responses for the stimuli of the experiment and then computed the expected mutual information between the stimuli and responses using a Monte Carlo method. When Δ^+ and Δ^- were varied, the value of τ was determined by the invariance $\tau = \Psi/\Delta$. To obtain a stable estimate of the expected mutual information, we repeated the stimulus set of each task to produce 198,000 trials for JRF and 3,300,000 trials for DMR. The s_i and r_i in Eq. 6 refer to objective and subjective relative frequencies in JRF and refer to expected values and certainty equivalents in DMR. In the numerical computation of mutual information, continuous variables need to be quantized. For JRF, the objective and subjective probabilities were quantized by rounding to the 2nd decimal, and for DMR, the expected values and certainty equivalents were quantized by rounding to the closest integer.

Acknowledgments

We would like to thank Richard Gonzalez and George Wu for sharing their dataset with us. H.Z. was supported by grants 31571117 and 31871101 from National Natural Science Foundation of China and funding from Peking-Tsinghua Center for Life Sciences. L.T.M. was supported by grant EY019889 from the National Institutes of Health, the Humboldt Research Prize of the Alexander v. Humboldt Foundation, a Guggenheim Fellowship from the John Simon Guggenheim Foundation and a Fellowship from the Institut des Etudes Avancées de Paris. A previous version of this work was published as a preprint on bioRxiv.

Author contributions: H.Z. and L.T.M. conceived the study. H.Z. and X.R. designed the experiment. X.R. performed the experiment. H.Z. and L.T.M. developed the theory. H.Z. and X.R. analyzed the data. H.Z. and L.T.M. wrote the manuscript.

The authors declare no conflicts of interest.

References

1. Knight FH. *Risk, uncertainty and profit*. Houghton Mifflin (1921).

2. Bernoulli D. Exposition of a new theory on the measurement of risk. *Econometrica* **22**, 23-36 (1738/1954).

3. von Neumann J, Morgenstern O. *Theory of games and economic behavior*, 3rd edn. Princeton University Press (1944/1953).

4. Maloney LT, Zhang H. Decision-theoretic models of visual perception and action. *Vision Res* **50**, 2362-2374 (2010).

5. Green DM, Swets JA. *Signal detection theory and psychophysics*. Wiley (1966/1974).

6. Trommershäuser J, Maloney LT, Landy MS. Decision making, movement planning and statistical decision theory. *Trends Cogn Sci* **12**, 291-297 (2008).

7. Wolpert DM, Landy MS. Motor control is decision-making. *Curr Opin Neurobiol* **22**, 996-1003 (2012).

8. Wu S-W, Delgado MR, Maloney LT. Motor decision-making. In: *Brain Mapping: An Encyclopedic Reference* (ed^(eds Toga A). Elsevier Science & Technology (2015).

9. Trommershäuser J, Maloney LT, Landy MS. Statistical decision theory and the selection of rapid, goal-directed movements. *J Opt Soc Am A* **20**, 1419-1433 (2003).

10. Trommershäuser J, Maloney LT, Landy MS. Statistical decision theory and trade-offs in the control of motor response. *Spat Vis* **16**, 255-275 (2003).

11. Allais M. Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école Américaine (The behavior of a rational agent in the face of risk: Critique of the postulates and axioms of the American school). *Econometrica* **21**, 503-546 (1953).

12. Kahneman D, Tversky A. Prospect theory: An analysis of decision under risk. *Econometrica* **47**, 263-291 (1979).

13. Tversky A, Kahneman D. Advances in prospect theory: Cumulative representation of uncertainty. *J Risk Uncertainty* **5**, 297-323 (1992).

14. Gonzalez R, Wu G. On the shape of the probability weighting function. *Cogn Psychol* **38**, 129-166 (1999).

15. Luce RD. *Utility of gains and losses: Measurement-theoretical and experimental approaches*. Lawrence Erlbaum (2000).

16. Wu S-W, Delgado MR, Maloney LT. Economic decision-making under risk compared with an equivalent motor task. *Proc Natl Acad Sci U S A* **106**, 6088-6093 (2009).

17. Ungemach C, Chater N, Stewart N. Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)? *Psychol Sci* **20**, 473-479 (2009).

18. Barron G, Erev I. Small feedback-based decisions and their limited correspondence to description-based decisions. *J Behav Decis Mak* **16**, 215-233 (2003).

19. Fox CR, Hadar L. Decisions from experience = sampling error + prospect theory: Reconsidering Hertwig, Barron, Weber & Erev (2004). *Judgm Decis Mak* **1**, 159-161 (2006).

20. Hertwig R, Barron G, Weber EU, Erev I. Decisions from experience and the effect of rare events in risky choice. *Psychol Sci* **15**, 534-539 (2004).

21. Stauffer WR, Lak A, Bossaerts P, Schultz W. Economic Choices Reveal Probability Distortion in Macaque Monkeys. *J Neurosci* **35**, 3146-3154 (2015).

22. Tobler PN, Christopoulos GI, O'Doherty JP, Dolan RJ, Schultz W. Neuronal distortions of reward probability without choice. *J Neurosci* **28**, 11703-11711 (2008).

23. Hsu M, Krajbich I, Zhao C, Camerer CF. Neural response to reward anticipation under risk is nonlinear in probabilities. *J Neurosci* **29**, 2231-2237 (2009).

24. Zhang H, Maloney LT. Ubiquitous log odds: a common representation of probability and frequency distortion in perception, action and cognition. *Frontiers in Neuroscience* **6**, 1 (2012).

25. Rustichini A, Conen KE, Cai X, Padoa-Schioppa C. Optimal coding and neuronal adaptation in economic decisions. *Nature Communications* **8**, 1208 (2017).

26. Tobler PN, Fiorillo CD, Schultz W. Adaptive Coding of Reward Value by Dopamine Neurons. *Science* **307**, 1642-1645 (2005).

27. Louie K, Khaw MW, Glimcher PW. Normalization is a general neural mechanism for context-dependent decision making. *Proceedings of the National Academy of Sciences* **110**, 6139-6144 (2013).

28. Kobayashi S, Carvalho OPD, Schultz W. Adaptation of reward sensitivity in orbitofrontal neurons. *J Neurosci* **30**, 534–544 (2010).

29. Polanía R, Woodford M, Ruff CC. Efficient coding of subjective value. *Nat Neurosci* **22**, 134-142 (2019).

30. Burr D, Ross J. A Visual Sense of Number. *Curr Biol* **18**, 425-428 (2008).

31. Cicchini GM, Anobile G, Burr DC. Compressive mapping of number to space reflects dynamic encoding mechanisms, not static logarithmic transform. *Proceedings of the National Academy of Sciences*, 201402785 (2014).

32. Ren X, Wang M, Zhang H. Context Effects in the Judgment of Visual Relative-Frequency: Trial-by-Trial Adaptation and Non-linear Sequential Effect. *Frontiers in Psychology* **9**, (2018).

33. Payzan-LeNestour E, Balleine Bernard W, Berrada T, Pearson J. Variance After-Effects Distort Risk Perception in Humans. *Curr Biol* **26**, 1500-1504 (2016).

34. Barnard GA. Statistical inference. *J R Stat Soc Series B* **11**, 115-149 (1949).

35. Berkson J. Application of the logistic function to bio-assay. *J Amer Statistical Assoc* **39**, 357-365 (1944).

36. Simon HA. *Models of bounded rationality: Economic Analysis and Public Policy*. MIT press (1982).

37. van den Berg R, Awh E, Ma WJ. Factorial comparison of working memory models. *Psychol Rev* **121**, 124–149 (2014).

38. Erev I, Wallsten TS, Budescu DV. Simultaneous over- and underconfidence: The role of error in judgment processes. *Psychol Rev* **101**, 519-527 (1994).

39. Wallsten TS, Budescu DV, Erev I, Diederich A. Evaluating and combining subjective probability estimates. *J Behav Decis Mak* **10**, 243-268 (1997).

40. Thurstone LL. A law of comparative judgment. *Psychol Rev* **34**, 273–286 (1927).

41. Attneave F. Some informational aspects of visual perception. *Psychol Rev* **61**, 183-193 (1954).

42. Barlow HB. Possible principles underlying the transformations of sensory messages. In: *Sensory communication* (ed^(eds Rosenblith WA). MIT Press (1961).

43. Simoncelli EP, Olshausen BA. Natural image statistics and neural representation. *Annu Rev Neurosci* **24**, 1193-1216 (2001).

44. Wei X-X, Stocker AA. A Bayesian observer model constrained by efficient coding can explain 'anti-Bayesian' percepts. *Nat Neurosci* **18**, 1509–1517 (2015).

45. Bhui R, Gershman SJ. Decision by sampling implements efficient coding of psychoeconomic functions. *Psychol Rev* **125**, 985–1001 (2018).

46. Varey CA, Mellers BA, Birnbaum MH. Judgments of proportions. *J Exp Psychol Hum Percept Perform* **16**, 613-625 (1990).

47. Prelec D. The probability weighting function. *Econometrica* **66**, 497-527 (1998).

48. Bell DE. Disappointment in decision making under uncertainty. *Oper Res* **33**, 1-27 (1985).

49. Chateauneuf A, Eichberger J, Grant S. Choice under uncertainty with the best and worst in mind: Neo-additive capacities. *J Econ Theory* **137**, 538-567 (2007).

50. Cohen M. Security level, potential level, expected utility: A three-criteria decision model under risk. *Theory Dec* **33**, 101-134 (1992).

51. Gilboa I. A combination of expected utility and maxmin decision criteria. *J Math Psychol* **32**, 405-420 (1988).

52. Loomes G, Moffatt PG, Sugden R. A microeconometric test of alternative stochastic theories of risky choice. *J Risk Uncertainty* **24**, 103-130 (2002).

53. Teitelbaum JC. A unilateral accident model under ambiguity. *The Journal of Legal Studies* **36**, 431-477 (2007).

54. Wakker PP. *Prospect theory: For risk and ambiguity*. Cambridge University Press (2010).

55. Birnbaum MH, Stegner SE. Measuring the importance of cues in judgment for individuals: Subjective theories of IQ as a function of heredity and environment. *J Exp Soc Psychol* **17**, 159–182 (1981).

56. Akaike H. A new look at the statistical model identification. *Automatic Control, IEEE Transactions on* **19**, 716-723 (1974).

57. Hurvich CM, Tsai C-L. Regression and time series model selection in small samples. *Biometrika* **76**, 297-307 (1989).

58. Stephan KE, Penny WD, Daunizeau J, Moran RJ, Friston KJ. Bayesian model selection for group studies. *NeuroImage* **46**, 1004-1017 (2009).

59. Hollands JG, Dyre BP. Bias in proportion judgments: The cyclical power model. *Psychol Rev* **107**, 500-524 (2000).

60. Spence I. Visual psychophysics of simple graphical elements. *J Exp Psychol Hum Percept Perform* **16**, 683-692 (1990).

61. Fox CR, Tversky A. A belief-based account of decision under uncertainty. *Manage Sci* **44**, 879-895 (1998).

62. Fox CR, Rottenstreich Y. Partition priming in judgment under uncertainty. *Psychol Sci* **14**, 195-200 (2003).

63. Smith M, Ferrell WR. The effect of base rate on calibration of subjective probability for true-false questions: Model and experiment. In: *Analysing and Aiding Decision Processes* (ed^(eds Humphreys P, Svenson O, Vari A). North Holland (1983).

64. Martins ACR. Probability biases as Bayesian inference. *Judgm Decis Mak* **1**, 108-117 (2006).

65. Birnbaum MH. Three New Tests of Independence That Differentiate Models of Risky Decision Making. *Manage Sci* **51**, 1346-1358 (2005).

66. Nassar MR, Helmers JC, Frank MJ. Chunking as a Rational Strategy for Lossy Data Compression in Visual Working Memory. *Psychol Rev* **125**, 486–511 (2018).

67. Orhan AE, Jacobs RA. A probabilistic clustering theory of the organization of visual short-term memory. *Psychol Rev* **120**, 297–328 (2013).

68. Yang T, Shadlen MN. Probabilistic reasoning by neurons. *Nature* **447**, 1075-1080 (2007).

69. Sanborn AN, Chater N. Bayesian Brains without Probabilities. *Trends Cogn Sci* **20**, 883-893 (2016).

70. Lebreton M, Abitbol R, Daunizeau J, Pessiglione M. Automatic integration of confidence in the brain valuation signal. *Nat Neurosci* **18**, 1159–1167 (2015).

71. Fox CR, Poldrack RA. Prospect theory and the brain. In: *Neuroeconomics* (ed^(eds Glimcher PW, Fehr E). 2nd edn. Academic Press (2014).

72. Gibbon J. Scalar expectancy theory and Weber's law in animal timing. *Psychol Rev* **84**, 279–325 (1977).

73. Zeisberger S, Vrecko D, Langer T. Measuring the time stability of Prospect Theory preferences. *Theory Dec* **72**, 359-386 (2012).

74. Hadar L, Fox CR. Information asymmetry in decision from description versus decision from experience. *Judgm Decis Mak* **4**, 317-325 (2009).

75. Erev I*, et al.* A choice prediction competition: Choices from experience and from description. *J Behav Decis Mak* **23**, 15-47 (2010).

76. Bornstein AM, Khaw MW, Shohamy D, Daw ND. Reminders of past choices bias decisions for reward in humans. *Nature Communications* **8**, 15958 (2017).

77. Bornstein AM, Norman KA. Reinstated episodic context guides sampling-based decisions for reward. *Nat Neurosci* **20**, 997–1003 (2017).

78. Sperling G. The information available in brief visual presentations. *Psychological monographs: General and applied* **74**, 1-29 (1960).

79. Miller GA. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychol Rev* **63**, 81-97 (1956).

80. Cochran WG. *Sampling techniques.* John Wiley & Sons (1977).

81. Fennell J, Baddeley R. Uncertainty plus prior equals rational bias: An intuitive Bayesian probability weighting function. *Psychol Rev* **119**, 878-887 (2012).

82. Cavagnaro DR, Pitt MA, Gonzalez R, Myung JI. Discriminating Among Probability Weighting Functions Using Adaptive Design Optimization. *J Risk Uncertainty* **47**, 255-289 (2013).

83. Stott HP. Cumulative prospect theory's functional menagerie. *J Risk Uncertainty* **32**, 101-130 (2006).

84. Shuford EH. Percentage estimation of proportion as a function of element type, exposure time, and task. *J Exp Psychol* **61**, 430-436 (1961).

85. Pitz GF. The sequential judgment of proportion. *Psychon Sci* **4**, 397-398 (1966).

86. Brooke JB, MacRae AW. Error patterns in the judgment and production of numerical proportions. *Percept Psychophys* **21**, 336-340 (1977).