Title
Sampling-based probability construction explains individual differences in risk preference

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors
Ankoju, Bhanu Prakash
Srivastava, Nisheeth

Publication Date
2022

Peer reviewed

Sampling-based probability construction explains individual differences in risk preference

Bhanu Prakash Ankoju (ankojubhan20@iitk.ac.in)
Dept of Cognitive Science, IIT Kanpur, India

Nisheeth Srivastava (nsrivast@iitk.ac.in))
Dept of Computer Science and Engineering, IIT Kanpur, India

Abstract

Contemporary models of subjective probability distortions assume that distortions arise during probability encoding. However, such assumptions are inconsistent with the ability of humans to retrieve probabilities veridically in some elicitation formats. We present a sampling-based model of probability judgment for risky prospects that assumes that probability distortions occur because people read out probability judgments as biased averages from working memory contents. Simulations demonstrate that this model shows the classic inverse-S shaped distortion of probability judgments using only retrieval-stage assumptions. The model further predicts that observers with greater working memory capacity would show larger probability distortions on average, which should lead to a particular fourfold pattern of risk preference as a function of working memory capacity. Using cognitive ability measurements as a proxy for working memory capacity, we conducted an experiment with human participants and found results consistent with the model's predictions as well as previous empirical studies. Our results support a role for sampling during assessment of risky prospects, which in turn explains differences in probability distortions seen across different elicitation methods.

Keywords: risk preference, cognitive ability, computational modeling, prospect theory, probability judgments

Introduction

People systematically overweight low probabilities and underweight high ones, demonstrating an inverse-S shaped relationship between objective probabilities and their subjective estimation in simple frequency estimation tasks as well as in risky decisions (Tversky & Kahneman, 1992). Since many household-level financial decisions involve probability judgments, microeconomic models are beginning to benefit from accommodating this stylized fact about peoples' behavior in combination with observations about loss and risk aversion, in the form of prospect theory (Barberis, 2013).

Given recent advances in our ability to estimate individual-level parameters for prospect theory (Nilsson, Rieskamp, & Wagenmakers, 2011), it is striking to note that prospect theory parameters show high inter-temporal consistency within individuals, suggesting that they correspond to stable individual differences in cognition (Glöckner & Pachur, 2012).

However, we are only beginning to understand the cognitive processes that prospect theory parameters map on to. Recent work using process tracing has shown that the relative extent of attention paid to gains and losses is significantly related to participants' estimated loss aversion (Pachur, Schulte-Mecklenbeck, Murphy, & Hertwig, 2018). Specifically with reference to probability distortions, Zhang and Maloney (2012) have shown that assuming a linear log odds representation of probability in the brain is sufficient to account for probability distortions seen across a wide variety of studies. However, this representational claim is consistent with a large number of theoretical possibilities, (Fox & Tversky, 1998; Fox & Rottenstreich, 2003; Martins, 2006) and thus does offers limited process-level understanding.

Furthermore, the view that probabilities are *encoded* in a distorted manner in the brain is inconsistent with evidence that people are actually able to reproduce probabilities veridically when these are elicited using graphical methods (Goldstein & Rothschild, 2014) and motor movements (Trommershäuser, Maloney, & Landy, 2003) . Taken in conjunction with classic studies showing that frequency encoding in humans is significantly veridical (Hasher & Zacks, 1984), such findings suggest that cognitive processes during *retrieval* may be more likely to produce probability distortions.

Outside the specific theoretical frame of prospect theory, multiple studies have sought to characterize individual differences in risk aversion profiles (Frederick, 2005; Burks, Carpenter, Goette, & Rustichini, 2009; Dohmen, Falk, Huffman, & Sunde, 2010). A common observation across these studies is that people with higher cognitive ability have a high-risk appetite in certainty-equivalence experiments with low probable gains (Frederick, 2005; Burks et al., 2009) and are risk-averse in high probability gains (Frederick, 2005; Dohmen et al., 2010). As a paradigmatic example, experiment participants who scored high on Frederick's Cognitive Reflection Test also showed a greater propensity to accept risky choices leading to gains, both when a simple expected utility calculation favored the risky option, but also crucially, when it did not (Frederick, 2005).

From the viewpoint of probability estimation, such behavior is congruent with participants in the experiments in Frederick (2005) over-weighting their estimate of low probability options. For low probability gains, participants with greater cognitive ability appear to be more risk-seeking, consistent with over-weighting of the low probability gain option. For high probability gains, such participants are more risk-avoidant, consistent with over-weighting of the low probability non-gain option. Figure 1 outlines the relationship of

In J. Culbertson, A. Perfors, H. Rabagliati & V. Ramenzoni (Eds.), *Proceedings of the 44th Annual Conference of the Cognitive Science Society*.

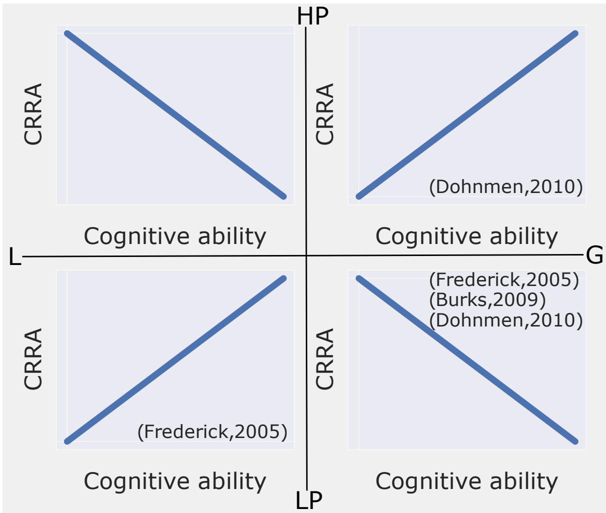

Figure 1: Risk sensitivity increases with measures of cognitive ability, as described in a number of behavioral studies. Multiple measures of risk preference and cognitive ability have been used in different studies. This figure plots the expected variation in the coefficient of relative risk aversion (CRRA) with increase in cognitive ability in the four patterns of behavior observed in the prospect theory view of risk aversion along with references to field studies that support the prediction in the particular quadrant.(HP - High Probable, LP- Low Probable, G- Gain, L- Loss)

risk aversion with cognitive ability we would expect for both gains and losses if the source of risk preference lies in overweighting of low-probability lottery outcomes. The papers referenced in the Figure show evidence consistent with the prediction relevant for each quadrant. Thus, convergent evidence across studies suggests a relationship between cognitive ability and probability distortions. In this paper, we develop a model of probability judgment that illuminates this relationship.

Probability by sampling

The empirical foundations of prospect theory show us that, when given a choice between risky prospects, people behave as if they were constructing a subjective probability estimate $w(p)$ based on the stated prospect risk p. If we take this process hypothesis seriously, we must ask: how do people map p to w? We propose that they do this by sampling from mental simulations, a possibility that has recently proven successful in explaining peoples' understanding of physical situations (Smith & Vul, 2013), as well as biases in probability judgments (Zhu, Sanborn, & Chater, 2020). Focusing on probability judgments for evaluating binary prospects, for simplicity, our probability-by-sampling model assumes that,

1. Observers possess a veridical, possibly noisy, internal probability scale.
2. When asked to reflect on a risky binary prospect, observers sample multiple abstract lotteries parameterized by the prospect risk, as read off the internal scale.
3. The outcomes of these simulated lottery draws are stored in working memory.
4. Observers sample from the lottery until either working memory capacity is reached[1], or both prospects have occurred at least once during sampling.
5. Observers read out the average occurrence of the salient option as their subjective probability estimate for it.

Of these assumptions, #1 follows standard psychophysical premises, #2 is the key sampling assumption of our approach, #3 follows standard assumptions about the role of working memory made in nearly all symbolic cognitive architectures (Ye, Wang, & Wang, 2018), #4 is a novel assumption made based on a recent observation that observers making risky decisions after explicitly sampling them also tend to wait until they have seen all possible prospects at least once before terminating sampling (Srivastava, Müller-Trede, Schrater, & Vul, 2016) and #5 is standard. Thus, the novelty of our model lies in assumptions #2 and #4.

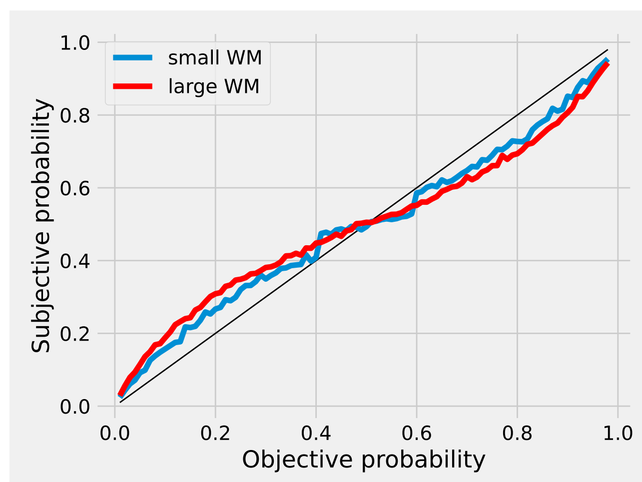

Figure 2: Subjective probability judgments extracted from the probability-by-sampling model for cohorts of 1000 observers sampled from low (blue) and high (red) working memory capacity pools.

Formally,

$$w(p) = \frac{1}{|\mathcal{M}|}\sum_{m}^{\mathcal{M}} I_m, \tag{1}$$

where I_m is an indicator function that takes the value 1 if the low probability outcome is sampled in the m^{th} memory slot,

[1] If memory sampling fails to retrieve a sample of the low probability outcome by the time capacity is reached, the model returns a probability of 0.01 for the low probability outcome.

and 0 otherwise. Also, $\mathcal{M}$ represents the set of memory slots[2] in working memory filled up at the time $w(p)$ is read out (up to maximum capacity), which in turn is determined by the number of samples it takes to see two distinct outcomes during sampling. For each memory sample,

$$I_m \sim N(p, \sigma_p), \tag{2}$$

where σ_p is noise in the internal Thurstonian magnitude scale. In the simulation results reported below, we use $\sigma_p = 0.5 * p$.

Figure 2 shows indicative results from an *in silico* experiment using our probability-by-sampling model. We conducted the experiment by sampling 1000 probability-by-sampling observers with working memory capacities sampled from normal distributions with means $\mu_{low} = 5, \mu_{high} = 10$ and SD = 1. Observers from both low and high WM capacity groups then responded to binary prospects across all possible probability values (quantized in steps of 0.01), producing subjective probability estimates for all these values. Figure 2 plots the average of these estimates for both WM size groups.

Two observations are salient. One, probability-by-sampling observers produce an inverse-S shaped distortion of probabilities (Tversky & Kahneman, 1992), based on assumptions about how they are retrieved. This is consistent with the fact that it is possible to elicit probabilities veridically in some elicitation formats (Trommershäuser et al., 2003; Goldstein & Rothschild, 2014). Two, we note that the high WM group shows greater probability distortion than the low WM group. These observations remain constant across multiple numeric values of our simulation parameters, but with probability distortions fading away for working memory sizes larger than 12.

The explanation for these observations is straightforward. Since observers sample simulated outcomes until they have seen both outcomes at least once, and then average over the outcomes sampled so far to read out the lower probability, there are two main possibilities. They will either not sample the low probability option at all, and read out zero, or sample the low probability option once and terminate sampling. In the latter case, the read out probability will be inflated by the small number of samples drawn. For example, suppose a lottery has an objective probability of 0.2 to pay out. The probability that the low-probability outcome will not be sampled even once in a working memory of size 4 is $0.8^4 = 0.41$, so the read out probability will be zero less than half the time. However, in the majority of cases that the outcome is sampled, the read out probability will be heavily inflated, e.g. if it is sampled on the second simulation, the probability will be read out as 0.5. Averaged across the population, this asymmetry yields probability over-weighting. For a larger working memory, say of size 8, the probability of not encountering a single sample of the low probability outcome reduces still further to $0.8^8 = 0.17$. In 67% of cases (1 - $0.8^5 = 0.67$), the observer will sample the low-probability outcome at least once, and read out a subjective probability estimate equal or greater than the objective probability. The probability will be read out as zero in much fewer instances for high WM observers than for low WM observers. Averaged across observers, this leads to greater probability over-weighting for high WM observers.

[2] While we use an explicit fixed slot interpretation of working memory in our exposition, probability-by-sampling is consistent with richer representations also (Ma, Husain, & Bays, 2014).

An Experimental Test

As we note above, the key novelty of the probability-by-sampling account of probability distortions is the assumption that observers mentally simulate lottery outcomes until they have seen at least one instance of both lottery prospects. We see in the simulation results above that this assumptions leads to a clear prediction relating working memory capacity to risk preferences - greater working memory capacity should lead to greater over-weighting of small probabilities. It is well-known that working memory capacity is strongly correlated with general cognitive ability, as measured by progressive matrices tests (Fukuda, Vogel, Mayr, & Awh, 2010). Therefore, treating cognitive ability as an empirical proxy for working memory capacity, our model predicts the specific relationship between cognitive ability and risk preference shown in Figure 1. While previous studies partially support the existence of the fourfold pattern illustrated in Figure 1, differences in protocols, analysis methods and operationalization of both independent and dependent variables make it difficult to assess the net weight of the evidence. To address this concern, we conducted an experiment to measure risk aversion as the CRRA coefficient of isoelastic utility functions in certainty equivalence problems selected to represent each of the four quadrants for participants with different cognitive ability levels, as measured by RSPM. We expected to see our dependent variable show the specific pattern of behavior predicted in Figure 1 as the outcome of this experiment.

Subjects

We solicited participants via email and social media. 103 participants (41 female, 62 male) responded and provided consent for participation. Out of 103 participants who appeared for the IQ test, 80 participants (32 female, 48 male) expressed interest to participate in the online risk-preference study.The mean age of the participants was 23.83 years. Since this was a between-subject design, participants were assigned to one of the four quadrants randomly at the time of experiment participation. All experimental protocols were approved by an Institutional Review Board. Participants signed a consent form describing all experimental procedures before participating in the study. Each participant was compensated for their time.

Measuring cognitive ability

To measure cognitive ability, we used Raven's Standard Progressive(SPM) Matrices (Raven, Court, & Raven, 1989) containing 60 questions and designed a website to administer the test online. Participants were shown puzzles from SPM one by one on the screen with corresponding options. They had to answer the puzzles by clicking one of the options. The raw scores (number of correct responses) obtained for each participant were converted to standard SPM percentiles using the SPM manual. There was no time limit for the test. Out of 80 participants, SPM's standard score for three participants was 5. These were excluded from analysis since their test duration was less than four minutes for 60 questions, suggesting random responding leaving us with 77 participants(30 female, 47 male) for the risk preference experiment, assigned randomly to the four quadrants of the experiment. The average time to complete the IQ test by the participants was 35.3min and our sample's average SPM percentile score was 62.4, suggesting that it was representative.

Measuring risk preference

We measured risk preference for each participant using choice table, which had 20 rows. Every participant provided their preference for each row of the table.(Dohmen et al., 2010). The choice tables used follow the ones used in Dohmen et al. (2010).

For gain-based problems, we asked the participants to choose whether to buy a lottery ticket that could fetch them lottery money with some uncertainty or accept the safe amount. Each quadrant had a set of five choice tables, with different lottery amounts and payoffs. The lottery amounts and payoffs for all the choice tables were derived from Frederick study 1 and then converted to equivalent local currency by considering Purchasing Power Parity in 2005 and inflation (2005-2021). The choice tables in these quadrants represent either high probable or low probable gain conditions. The order of the choice table presented to each participant was randomized. In a choice table, the lottery amount remained the same while the safe option increased systematically for every row; a rational agent would be willing to take risks until the safe amount is less than the expected value of the gamble and then switches to the safe option. We presented participants one row at a time and asked to choose whether 'to buy the lottery ticket (risky)', which can fetch them a lottery amount with some uncertainty, or 'Not to buy the lottery ticket (safe)' and accept the safe amount.

Once the participant switched from risky option to safe option, the algorithm asked the participant whether they would accept all higher safe amounts or not (see also (Dohmen et al., 2010)) if they responded as yes, the algorithm considers all other safe options in the table as their preferences, and the participant was progressed to the new choice table. Otherwise, the participant had to decide for the rest of the table manually and then be presented with a new choice table. Following Dohmen et al. (2010), we also informed participants that one row from one of the five choice tables would be randomly selected, and they would be rewarded with the amount proportional to the choice they made in that selected row, to encourage participants to choose according to their true preferences for each row.

The same procedure was used for loss-based problems, except that the problems were framed as a choice to buy insurance costing a small fixed amount or retain a small probability of suffering a larger loss.

Out of 385 instances(77 participants * 5 choice tables), there were 16 incidents, where participants switched from risk option to safe option multiple times. The sixteen instances can be classified into two scenarios. Scenario-1, they selected safe options consecutively. One example, a participant switched from risky option to safe option at 10^{th} row and selected safe option again in next rows(11,12) and then moved to the next choice table. In this case, we considered the first switch (in this example, 10^{th} row) as their risk preference. And in scenario -2, they switched from risky option to safe option and again selected the risky option then switched to the safe option. One example, a participant switched from risky option to safe option at 4^{th} row and then selected risky option in the 5^{th} row and continued risky option till row 10 and switched to safe option at 11^{th} row. Here we considered the latest switch(here row 11) as their risk preference. And in another example, a participant switched from risky option to safe option at 4^{th} row and then selected risky option in the 5^{th} row and continued risky option till end of the table. Here we considered last row(20^{th} row) as their risk preference.

The coefficient of relative risk aversion was calculated from an individual's utility function (Burks et al., 2009) . We follow Burks et al. (2009) in assuming that the participant's utility for the lottery would be at the midpoint of $safe_i$ and $safe_j$. (where 'i' and 'j' refers to the steps when participant prefers to take the risk at $safe_i$, but switches to the safe option at $safe_j$.). The individual's utility function is then given by,

$$u(c) = \frac{c^{1-\sigma}}{1-\sigma}, \tag{3}$$

where σ is the CRRA coefficient we are interested in measuring.

Following Burks et al. (2009) and assuming expected utility maximization, the equation below holds when a participant switches their lottery preference between cells i and j of the table, and is solved analytically for lottery utility and then numerically for σ to obtain the coefficient of relative risk aversion,

$$p\, u(lottery) = 0.5\, u(safe_i) + 0.5\, u(safe_j), \tag{4}$$

where p corresponds to the stated probability of winning the lottery. The same procedure was used to estimate CRRA in loss conditions as well.

Results

The mean CRRA estimates for all five choice tables seen by participants in each quadrant are shown in Figure 3. For each choice table, we found the best fit line relating CRRA to IQ. To obtain a summary measure of the trend across choice tables for each quadrant, we shifted the CRRA points from each choice table to a common intercept (the average intercept across the best fit lines). We then replotted the points using individual slope values from the table-wise best fit lines.

Finally, we fitted a linear regression to the combined CRRA estimates (see rightmost column in Figure 3).

Table 1: Average slope in all quadrants.

Quadrant	Sign Prediction	Coefficient	p	f^2
HPG	+	0.31	0.03	0.17
LPG	-	-0.03	0.03	0.17
LPL	+	-0.21	0.09	0.1
HPL	-	-2.41	0.000008	1.06

Table 1 documents the coefficient of the IQ variable in the combined regression for all four quadrants, alongside the predicted sign of the coefficient, as seen in Figure 1. We note that the measured coefficients are directionally consistent with our predictions in three of four quadrants. Results for three quadrants (high probability gains, low probability gains and high probability losses) statistically significant at the traditional 0.05 alpha-error level and displaying medium effect sizes ($f^2 > 0.15$) (Cohen, 1988). For the low probable loss quadrant, we see small effect sizes ($f^2 > 0.02$), with the relationship failing to meet statistical significance.

To verify that the observed relationships between cognitive ability and risk preference are not an artifact of our data pooling procedure across choice tables, we fit a hierarchical linear regression model for every quadrant separately. We model the relationship between CRRA and IQ in each quadrant as follows

$$
\begin{aligned}
CRRA_i &= slope_i * IQ + intercept_i + \varepsilon \\
slope_i &\sim \mathcal{N}(\mu_{slope}, \sigma^2{}_{slope}) \\
intercept_i &\sim \mathcal{N}(\mu_{intercept}, \sigma^2{}_{intercept}) \\
\varepsilon &\sim HalfCauchy(5)
\end{aligned}
$$

where $slope_i$, $intercept_i$ are the slope and intercept parameters for the choice-table 'i' in a quadrant and ε is noise. We used Gaussian and half-Gaussian priors, respectively, for our two mean and two standard deviation hyperparameters.

We fit this model using PyMC3's NUTS sampler using 2 chains of 2000 draw iterations with 1000 tuning steps. The key parameter of interest for us is the mean of the distribution of μ_{slope} from which slopes for different choice sets are sampled. Figure 4 plots the quadrant-wise posterior distributions for μ_{slope} from the fitted model. The key observation is that the MAP estimates of μ_{slope} reliably track the average slope estimates we obtained in our pooled analysis, suggesting that the pattern seen in the previous analysis is not an artifact of the data pooling procedure.

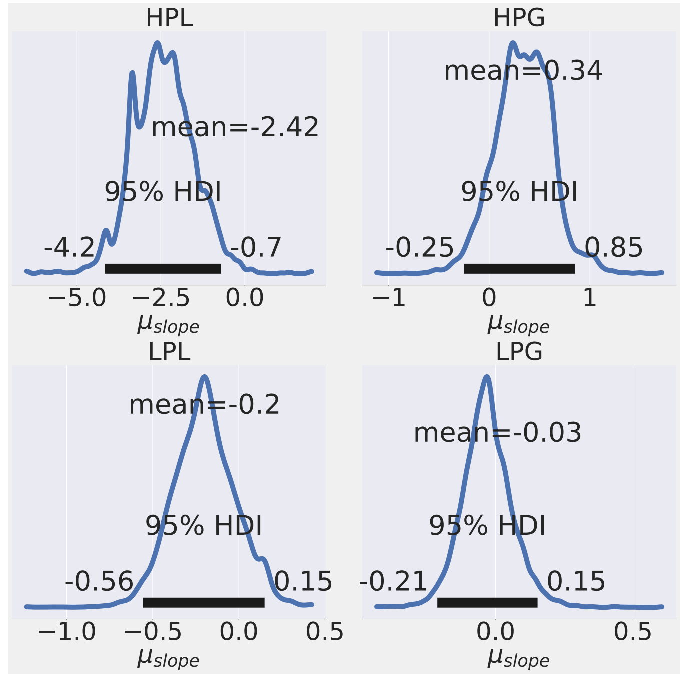

Figure 4: This figure plots the posterior distribution for the mean distribution from which slopes of all quadrants are sampled with 95% credible intervals.

Discussion

In this paper, we have presented a sampling-based model of probability judgment for risky prospects and demonstrated that it shows the classic inverse-S shaped distortion of probability judgments. The model predicted a specific pattern of correlations between cognitive ability and risk aversion, which we tested using an experiment with human subjects. The pattern of results is consistent with the model's predictions in three of the four quadrants, as well as with earlier empirical studies of the relationship between cognitive ability and risk preference (Frederick, 2005; Burks et al., 2009; Dohmen et al., 2010). Andersson, Holm, Tyran, and Wengström (2016) proposed that the relationship between cognitive ability and risk aversion is spurious, and the direction of correlation depends on the behavioural noise and the biased risk elicitation method. In the current study, all the choice tables in every quadrant are biased in the same direction, although we still see both positive and negative correlations between cognitive ability and risk aversion. This suggests that the relationship between cognitive ability and risk aversion does not depend solely on the bias in the risk elicitation method and the behavioural noise. However, since results for low probable loss quadrant, failed to meet statistical significance, further work is needed to verify it.

While previous work has proposed general models of probability distortions (Gershman & Wilson, 2010; Zhang & Maloney, 2012) as well as models of probability distortion that use retrieval-specific assumptions (Fox & Tversky, 1998; Fox

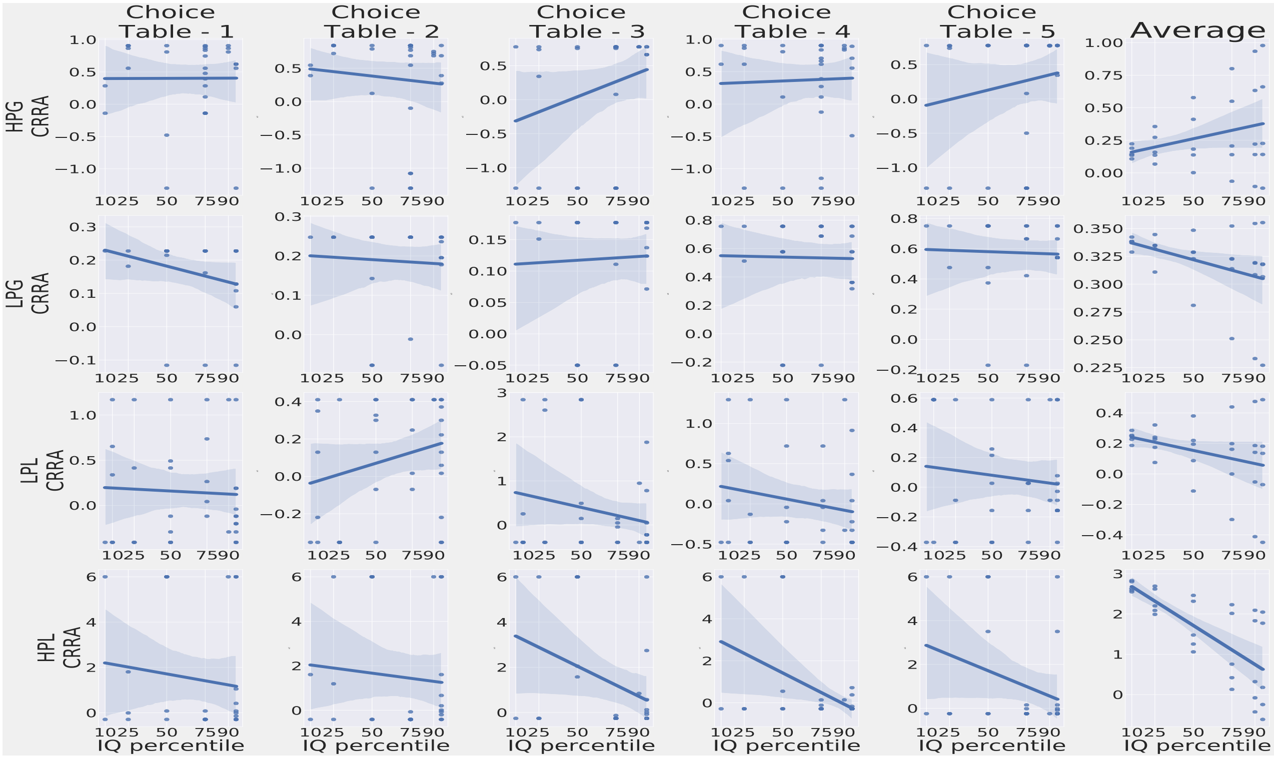

Figure 3: This figure plots the relation between CRRA vs IQ for every choice table in all the quadrants, and the average plot to show the overall trend in a quadrant.

& Rottenstreich, 2003), these proposals have so far been mutually exclusive, in the sense that the former category produce generalized models by making encoding-based assumptions, and the latter category produce context- and task-specific models. The probability-by-sampling model, while specialized to prospect risk in our current presentation, can be extended to other tasks easily. For example, for frequency estimation, we simply need assume that observers sample tokens until they sample the one they are estimating the frequency of once, and then average across the token count to produce a frequency estimate.

Probability-by-sampling is thus a task-general retrieval-based model of probability distortions. It is therefore, able to accommodate the possibility of veridical encoding of frequency information (Hasher & Zacks, 1984) and the possibility of near-veridical retrieval of probability information using non-symbolic elicitation procedures (Trommershäuser et al., 2003; Goldstein & Rothschild, 2014), which are problematic for encoding-based accounts of probability distortions. Probability-by-sampling is also consistent with recently documented evidence for the use of sampling in probabilistic judgments in other studies (Zhu et al., 2020).

We note with interest a number of theoretical connections between probability-by-sampling, and a recent improvement upon the linear-log-odds model (Zhang, Ren, & Maloney, 2020). In addition to assuming a linear log odds representation of probability, Zhang et al. (2020) show that human frequency and probability judgments are better explained if we further assume that the distorted probabilities are mapped dynamically to a quantized internal Thurstonian scale, with the noise of the scale subject to variance compensation. The quantization implicit in the former assumption maps nicely onto the discrete nature of memory sampling in probability-by-sampling. Even more interestingly, (Zhang et al., 2020) show that Gaussian encoding noise on a Thurstone scale in log-odds, when transformed back into probability is approximately proportional to the variance of a binomial distribution parameterized by the probability value. Since probability-by-sampling involves a sequence of Bernoulli trials parameterized by the probability value, the signature of scale noise in our model would also be exactly binomial. Exploring synergies and differences between the two models presents a clear direction for future work.

References

Andersson, O., Holm, H. J., Tyran, J.-R., & Wengström, E. (2016). Risk aversion relates to cognitive ability: Preferences or noise? *Journal of the European Economic Association*, *14*(5), 1129–1154.

Barberis, N. C. (2013). Thirty years of prospect theory in economics: A review and assessment. *Journal of Economic Perspectives*, *27*(1), 173–96.

Burks, S. V., Carpenter, J. P., Goette, L., & Rustichini, A. (2009). Cognitive skills affect economic preferences, strategic behavior, and job attachment. *Proceedings of the National Academy of Sciences*, *106*(19), 7745–7750.

Cohen, J. (1988). The effect size. *Statistical power analysis for the behavioral sciences*, 77–83.

Dohmen, T., Falk, A., Huffman, D., & Sunde, U. (2010). Are risk aversion and impatience related to cognitive ability? *American Economic Review*, *100*(3), 1238–60.

Fox, C. R., & Rottenstreich, Y. (2003). Partition priming in judgment under uncertainty. *Psychological Science*, *14*(3), 195–200.

Fox, C. R., & Tversky, A. (1998). A belief-based account of decision under uncertainty. *Management science*, *44*(7), 879–895.

Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic perspectives*, *19*(4), 25–42.

Fukuda, K., Vogel, E., Mayr, U., & Awh, E. (2010). Quantity, not quality: The relationship between fluid intelligence and working memory capacity. *Psychonomic bulletin & review*, *17*(5), 673–679.

Gershman, S., & Wilson, R. (2010). The neural costs of optimal control. *Advances in neural information processing systems*, *23*, 712–720.

Glöckner, A., & Pachur, T. (2012). Cognitive models of risky choice: Parameter stability and predictive accuracy of prospect theory. *Cognition*, *123*(1), 21–32.

Goldstein, D. G., & Rothschild, D. (2014). Lay understanding of probability distributions. *Judgment & Decision Making*, *9*(1).

Hasher, L., & Zacks, R. T. (1984). Automatic processing of fundamental information: the case of frequency of occurrence. *American psychologist*, *39*(12), 1372.

Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nature neuroscience*, *17*(3), 347–356.

Martins, A. C. (2006). Probability biases as bayesian inference. *Judgment and Decision Making*, *1*(2), 108.

Nilsson, H., Rieskamp, J., & Wagenmakers, E.-J. (2011). Hierarchical bayesian parameter estimation for cumulative prospect theory. *Journal of Mathematical Psychology*, *55*(1), 84–93.

Pachur, T., Schulte-Mecklenbeck, M., Murphy, R. O., & Hertwig, R. (2018). Prospect theory reflects selective allocation of attention. *Journal of Experimental Psychology: General*, *147*(2), 147.

Raven, J. C., Court, J. H., & Raven, J. E. (1989). *Standard progressive matrices*. Australian Council for Educational Research Limited.

Smith, K. A., & Vul, E. (2013). Sources of uncertainty in intuitive physics. *Topics in cognitive science*, *5*(1), 185–199.

Srivastava, N., Müller-Trede, J., Schrater, P. R., & Vul, E. (2016). Modeling sampling duration in decisions from experience. In *Cogsci.*

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003). Statistical decision theory and the selection of rapid, goal-directed movements. *JOSA A*, *20*(7), 1419–1433.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, *5*(4), 297–323.

Ye, P., Wang, T., & Wang, F.-Y. (2018). A survey of cognitive architectures in the past 20 years. *IEEE transactions on cybernetics*, *48*(12), 3280–3290.

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

Zhang, H., Ren, X., & Maloney, L. T. (2020). The bounded rationality of probability distortion. *Proceedings of the National Academy of Sciences*, *117*(36), 22024–22034.

Zhu, J.-Q., Sanborn, A. N., & Chater, N. (2020). The bayesian sampler: Generic bayesian inference causes incoherence in human probability judgments. *Psychological review*, *127*(5), 719.