

Demonstrating the Impact of Prior Knowledge in Risky Choice

Mathew D. Hardy (mdhardy@princeton.edu)

Department of Psychology, Princeton University

Thomas L. Griffiths (tomg@princeton.edu)

Departments of Psychology and Computer Science, Princeton University

Abstract

Bayesian models that optimally integrate prior probabilities with observations have successfully explained many aspects of human cognition. Research on decision-making under risk, however, is usually done through laboratory tasks that attempt to remove the effect of prior knowledge on choice. We ran a large online experiment in which risky options paid out according to the distribution of Democratic and Republican voters in US congressional districts to test the effects of manipulating prior probabilities on participants' choices. We find evidence that people's risk preferences are appropriately influenced by prior probabilities, and discuss how the study of risky choice can be integrated into the Bayesian approach to studying cognition.

Keywords: risk; decision-making; Bayesian modeling

Introduction

You sample six voters from a recent US election in an unknown Florida congressional district. Three voted Democrat and three voted Republican. Would you prefer \$5 for sure, or \$10 if Democrats won that district? What if these voters were from a district in Massachusetts, a district in Oklahoma, or the \$10 pays if the next voter sampled is Republican? Political analysts have to make bets like this all the time, but even casual observers of politics know enough that they could make intelligent choices.

Whether people do so is an open question. When faced with related everyday tasks, people's judgments and perceptions are often best described as inferences and predictions that incorporate prior probabilities with observations in a Bayesian manner (Anderson, 1990; Griffiths & Tenenbaum, 2006). These results contrast with other research showing that people often disregard base-rate information when making laboratory judgments and predictions (Tversky & Kahneman, 1974; Kahneman et al., 1982). Given the prevalence of risk in everyday choice and the breadth of biases that characterize people's preferences in laboratory studies (Friedman et al., 2014), understanding the relationship between prior everyday knowledge and risk preferences is crucial to understanding how these decisions are made.

The importance of prior probabilities is especially notable in decisions from experience, in which people must learn reward structures through sampling, rather than being given this information directly (Hertwig et al., 2004). However, research on decisions from experience is usually done through abstract laboratory tasks, in which people's contextual, domain, and prior knowledge is assumed not to influence

choice. This has enabled results from decisions from experience to be directly compared with those from decisions from description, in which probabilities and rewards are given directly to participants, and has led to a number of notable findings on systematic differences in people's behavior in these domains (Hertwig & Erev, 2009; Rakow & Newell, 2009; Weber et al., 2004; Ert & Haruvy, 2017; Ert & Lejarraga, in press). However, this has come at the cost of testing behavior on choice tasks that are systematically different from those people face in everyday life.

To determine the effects of prior probabilities on people's risk preferences, we ran a large online experiment in which participants chose between safe and risky options. In each case, participants first viewed a sample of Democratic and Republican voters in a congressional district in a US state. While participants knew the state, they did not know the district. The risky option paid out depending on either the identity of the next voter sampled or the majority of voters in the district. By keeping the voter sample constant and varying only the state being sampled from, we were able to measure the impact of prior probabilities on people's risk preferences as well as compare people's choices with those of the optimal Bayesian model.

Background

While the impact of prior knowledge on risk preferences is not well understood, the effects of prior probabilities on judgments and risk on choice are active areas of research. In both, there is debate as to the best way to model these processes, as well as the extent to which people behave optimally.

Bayesian models of human judgment

Research on the role of prior probabilities on people's judgments is traditionally done by studying how people incorporate base-rate knowledge with given information in laboratory tasks. Early work appeared to show that people tend to disregard base-rate information in their judgments and predictions (Tversky & Kahneman, 1974; Kahneman et al., 1982), a phenomenon termed the base rate fallacy. However, other researchers have questioned the robustness of this fallacy. Gigerenzer and Hoffrage (1995) showed that changing the presentation format of statistical information from probabilities to frequencies caused people to make decisions that better incorporated base-rate information. Furthermore, Grif-

fiths and Tenenbaum (2006) showed that people’s conditional predictions of many real-world phenomena are often indistinguishable from those of the optimal Bayesian model. In their setup, participants were not explicitly given prior distribution information but appeared able to accurately incorporate it into their predictions. Thus, in everyday tasks, it appears that people’s judgments can and do incorporate prior information in the optimal manner.

Decision-making under risk

Somewhat at odds with the colloquial concept, the riskiness of an option in the behavioral sciences is generally defined as the variance of its expected outcomes. Research in this area has shown that people’s choices often reflect risk aversion. That is, people prefer safe options – those with low or zero outcome variance – to more risky options, all else being equal. This is most often experimentally studied in the simplest case, in which people make choices between a risky option that pays off only with some probability p , and a safe option than pays off with certainty.

This aversion to risk led Bernoulli (1954/1738) to propose a model of diminishing marginal utility of wealth. In this model, later expanded and formalized as Expected Utility Theory (Neumann & Morgenstern, 1953), people maximize their expected utility of wealth, rather than their expected wealth, where utility is a monotonic but concave function of wealth. Under these assumptions, it can be shown that a utility-maximizing decision-maker will prefer an option with low outcome variance to another with equal expected value but higher variance.

Others have questioned both the descriptive validity of this model and the ubiquity of risk aversion across decision contexts (Friedman et al., 2014). Perhaps most notably, Kahneman and Tversky (1979, 1992) developed Prospect Theory as an alternative descriptive model of decision-making under uncertainty. Rather than a purely concave function a wealth, utility in Prospect Theory is defined with respect to a reference point; outcomes above this reference point are gains, and outcomes below it are losses. For gains, the classic utility function holds, but for losses this curve is convex. In addition to this value function, Prospect Theory models the effect of probability on choice through a nonlinear weighting function that overweights low probabilities and underweights high probabilities. The combination of these utility and probability weighting functions account for what Kahneman and Tversky called the fourfold pattern of risk attitudes; people are risk averse for gains and risk seeking for losses with high probability, but risk seeking for gains and risk averse for losses with low probability.

Decisions from description and experience

Both Expected Utility Theory and Prospect Theory were developed to explain people’s preferences in decisions from description, in which people are given both the probabilities and outcomes for every option. However, the predictions of these models do not always generalize to tasks in which proba-

bilities and outcomes have to be learned through sampling (Barron & Erev, 2003). For example, in these decisions from experience people tend to overweight high-probability events and underweight low-probability events, contradicting the experimental findings in decisions from description. These discrepancies have lead some to call for separate models for the two paradigms (Rakow & Newell, 2009), while others have argued that traditional models can be adapted to explain behavior in both (?) or have developed new models that do so (Lieder et al., 2018).

While learning outcome information through sampling appears to have systematic effects on people’s risk preferences, the impact of people’s expectations on these preferences is less clear. Ert and Trautmann (2014) found that people’s estimates of payout probabilities after sampling tended to lie between their initial estimates and the observed frequencies, as would be predicted by a Bayesian model. However, sampling impacted choice in ways that this model does not fully account for. After sampling payouts from a choice with an unknown payout probability, participants were then given a choice between taking another sample from this option and a sample from a distribution with a known payout probability. Participants tended to prefer continuing to sample from this distribution when payoff probabilities were high, and sample from the known probability when payoff probabilities were low. This behavior contradicts both the predictions of a Bayesian updating model and the probability judgments of other participants on an identical task. Sampling may therefore affect people’s risk preferences in ways that are not explained by the optimal Bayesian model.

A Bayesian approach to modeling risk

To test the effects of varying prior probabilities on participants’ risk preferences, we need to estimate the effects these probabilities have on an option’s value. In this section, we review how these values are traditionally estimated as well as how to do so in a Bayesian framework.

Rational decision-making

Risk preferences on decisions from experience are estimated by comparing the proportion of people that choose a risky option O_R over a safe option O_S . This risky option pays out V_R with some unknown probability θ and 0 otherwise, and the safe option pays out V_S with certainty. Before making a decision, the decision maker is given these payoff values as well as n independent samples from the risky option $D = \{d_1, d_2, \dots, d_n\}$, $D \sim B(n, \theta)$. In this sample are z successes – that is observations that would cause the risky option to pay – and $n - z$ failures. Observations that are successes are encoded as 1 and failures as 0 so that $\sum_i^n d_i = z$. Computing the expected value of O_R after observing the data, $E[O_R|D]$ thus requires estimating the probability $p(d_{n+1} = 1) = \theta$. This is usually done using the maximum likelihood estimate of θ given by $\frac{z}{n}$. In this case, the rational choice is to choose O_R if:

$$E[O_R|D] = \frac{zV_R}{n} > V_S \quad (1)$$

and choose the safe option otherwise. These values are often equal, and if participants on average choose O_S more than O_R they are said to be risk averse for the choice. We will refer to the maximum likelihood estimate of the value of the risky option as the sample-based estimate.

This conditional expected value $E[O_R|D]$ can also be estimated using a Bayesian model. This approach combines the maximum likelihood estimate used above with prior probabilities for possible values of θ . While these prior probabilities may not be useful in many laboratory studies, they are crucial for solving many real-world tasks optimally (Anderson, 1990).

One of the strengths of Bayesian models is their ability to perform both inference and prediction. Inference tasks involve questions about the true parameters of a distribution, while prediction tasks involve estimating the distribution of future samples. While in most studies on risky choice estimating the value of the risky option requires prediction, we believe that both are important in the risky choices people make in everyday life. We first show how traditional prediction-based risky choices can be modeled in a Bayesian framework, and then how inference-based choices can be modeled as extensions of this setup.

Prediction

The prediction-based setup is identical to the one described above, but the methods used for estimating $E[O_R|D]$ differ. Instead of estimating a single value of θ , we assume that there is a set of M possible values θ can take on and estimate the probabilities of each. In this case, the risky option should be chosen if:

$$E[O_R|D] = \sum_i^M [p(\theta_i|D)E[O_R|\theta_i]] > V_S \quad (2)$$

To make this choice, both $p(\theta_j|D)$, the posterior distribution of θ_j as well as $E[O_R|\theta_j]$, the expected value of the risky option conditional on a certain value θ_j need to be estimated for all M possible values of θ . We show how to do both below.

Estimating the posterior distribution The posterior distribution of θ_j – its distribution conditional on sample D – can be computed using Bayes' rule:

$$p(\theta_j|D) = \frac{p(D|\theta_j)p(\theta_j)}{p(D)} \quad (3)$$

Because we assume that D is a binomial sample and that observations are independent given the true value of θ , Equation (3) can be rewritten as follows:

$$p(\theta_j|D) = \frac{\theta_j^z(1-\theta_j)^{n-z}p(\theta_j)}{\sum_i^M [\theta_i^z(1-\theta_i)^{n-z}p(\theta_i)]} \quad (4)$$

Estimating the expected value of the risky option The probability that the next sample will pay given observations D can be estimated using the posterior predictive distribution. This gives a distribution of possible future samples conditioned on the data already observed. In our setup, the risky option pays V_R if the next sampled value equals 1, and 0 otherwise. This can be estimated as follows:

$$E[O_R|\theta_j] = p(d_{n+1} = 1|\theta_j)V_R = \theta_j V_R \quad (5)$$

Equation (5) can then be combined with the posterior probability estimates from Equation (4) to estimate the expected value of the risky option after observing data D :

$$E[O_R|D] = \sum_i^M [p(\theta_i|D)\theta_i] V_R \quad (6)$$

Inference

Inference-based risky choices can be constructed similarly to the prediction setup described above, except that option O_R pays depending on the true value of θ . Thus, the payoff of the risky option is deterministic, but stochastic from the perspective of the decision-maker and modeler. In this case, the risky option O_R pays V_R if the true value of θ is in some set C . This can be used to compute the expected value of the risky option given probability θ_j :

$$E[O_R|\theta_j] = \mathbb{1}_C(\theta_j)V_R \quad (7)$$

where $\mathbb{1}_C(\theta_j)$ is the indicator function that returns 1 if $\theta_j \in C$ and 0 otherwise. This value can be combined with the posterior distribution estimates from Equations (3) and (4) to estimate the expected value of the risky option conditional on observations D :

$$E[O_R|D] = \sum_i^M [\mathbb{1}_C(\theta_i)p(\theta_i|D)] V_R \quad (8)$$

Testing the model predictions

We test the predictions of the Bayesian model above by studying risk preferences when D is presented as a random sample of Democratic and Republican voters. This sample is from a known state but unknown congressional district in that state. This allows us to test the model predictions in a setting in which many Americans have relatively strong prior knowledge, and in which the true prior distributions can be obtained from election data. Furthermore, this setup allows prior probabilities to be easily manipulated by changing the state D is sampled from.

In this setup, the prior probability distribution $p(\theta)$ for a state is the distribution of Democratic and Republican voters by district in that state. θ_j thus denotes the proportion of Democratic voters (among only Democratic and Republican voters) in district j for the M total districts in the state. Inference-based risky tasks pay out depending on the true distribution of Republican and Democratic voters in the district, and prediction-based risky tasks pay out according to

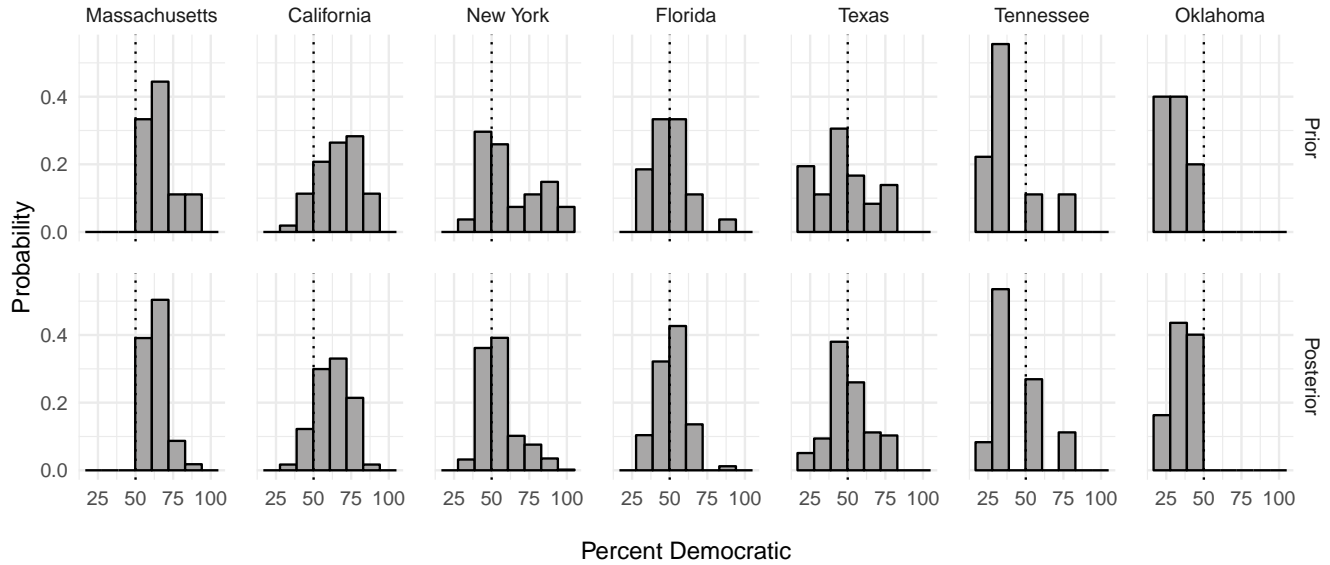


Figure 1: Distribution of the percentage of Democratic voters in the 2016 Presidential Election by congressional district. The top row shows prior distributions, and the second row posterior distributions after observing three Democratic and three Republican voters. These percentages are among only Democratic and Republican voters.

the identity of the next voter sampled. Histograms of prior and posterior distributions for the seven US states included in our study are given in Figure 1. All voter data used in the experiment and results are from the 2017 Cook Political Report Partisan Voter Index.

Participants

We recruited 1,384 participants from Amazon Mechanical Turk, limiting our study to those living in the United States. Participants earned \$0.50 for completion of the experiment, with the ability to earn an additional bonus of either \$0.25 or \$0.50. 24 participants were excluded for failing a basic attention check, and an additional 134 participants for failing to identify the region of the US containing the state they were assigned to. This left a total of 1,226 participants whose choices we report.

Procedure

Participants were assigned randomly to either the district condition or the individual condition. Estimating the value of the risky choice in the district condition was an inference task and paid off depending on the party affiliation of the majority of the voters in the district. Estimating this value in the individual condition was a prediction task, and the risky option paid off depending on the next sampled voter’s party affiliation. This risky option paid off either \$0.50 or \$0, and in both conditions participants made a choice between this risky option and a safe option that always paid \$0.25. In total, of the 1,226 participants whose data we report, 619 were assigned to the district condition, and 607 to the individual condition. In both conditions, the political party that determined whether the risky option paid off was randomized across participants.

Participants were instructed that they would see a random sample of six voters from the 2016 US Presidential election in a certain congressional district. They were told that all of these voters would be sampled from a single congressional district in a US state, and while they would know the state, they would not know the district. They were also told that each district had the same chance of being selected. Participants then saw a map of the congressional districts in the state they were assigned to as well as the number of districts in that state before moving onto the choice phase. Seven total states were possible: Massachusetts, California, New York, Florida, Texas, Tennessee, and Oklahoma. These states were chosen to cover as much of the prior distribution space as possible.

After a practice round, participants completed a single test trial. Participants were assigned to the same condition (district or individual) for the practice and test trial, but different states were always assigned for both. The six sampled voters on the test trial were always three Democrats and three Republicans. Thus, the sample-based estimate of the risky option’s value and the value of the safe option were always equal, and changes in participants’ preferences in different states would be determined only by their prior knowledge. If the risky option was chosen, it then paid out with its true posterior probability.

Results

A scatter plot of the proportion of participants who chose the risky option for both conditions and for every state is given in Figure 2. In both the district and the individual conditions, a majority of participants made the optimal choice for 13 of the 14 possible state-party combinations. These proportions

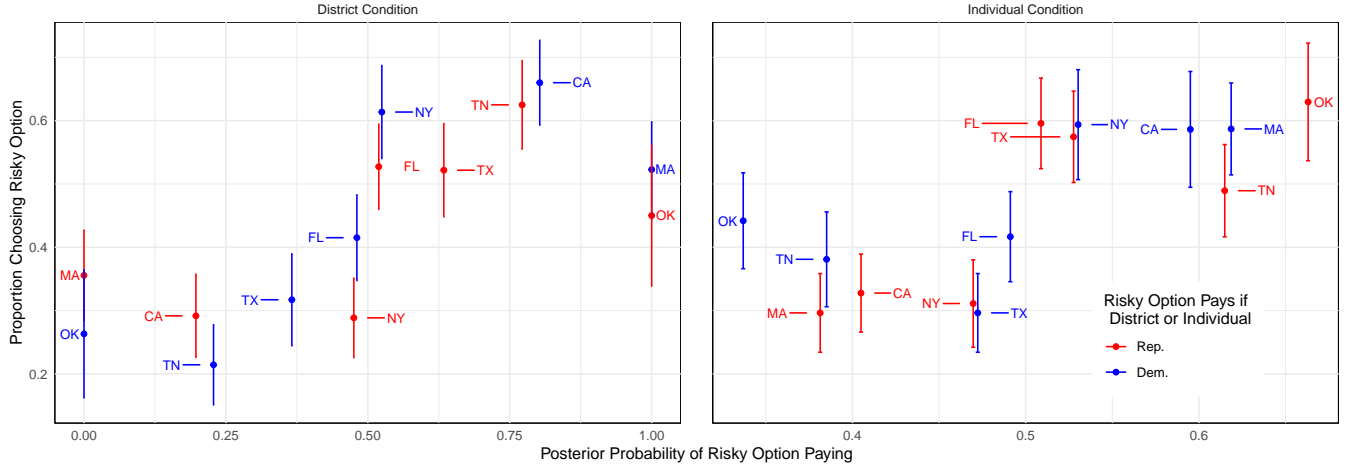


Figure 2: A comparison of the posterior value of the risky choice and the proportion of participants who chose it. Red points indicate that the risky option paid if the district was majority Republican or the next voter sampled was Republican, and blue points indicate the same for Democratic districts and voters. States are indicated with two-letter abbreviations, and vertical lines show standard error estimates. In each case, the safe option paid \$0.25 with certainty, and the risky option paid \$0.50 or \$0. The maximum likelihood estimate of the probability of the risky option paying was 0.5 for all points. Note that in the 2016 US Presidential Election, all of Massachusetts’ districts voted Democratic and all of Oklahoma’s districts voted Republican.

were significantly greater than 0.5 when the risky choice was optimal in both conditions ($p < 0.017$ for both), as well as significantly less than 0.5 when the safe choice was optimal ($p < 0.0001$ for both). Thus, given the prior probabilities, participants’ risk preferences were on average consistent with the Bayesian model.

While participants on average chose optimally, a higher proportion did so when the safe option was optimal than when the risky option was optimal. This difference was significant in the district condition (difference = 0.117, $p < 0.01$), but not in the individual condition (difference = 0.077, $p = 0.056$). Thus, although participants adapted their choice strategies to changes in the prior probabilities, there is some evidence that they did this more often when the optimal choice was the safe one.

To model the effects of the posterior value of the risky option on choice, we model utility for participant k as a latent variable U_k . Every participant thus has utility $U_{R,k}$ for the risky option and $U_{S,k}$ for the safe option. These utilities are defined as follows:

$$U_{R,k} = \alpha_R + \beta E[O_R|D] + \varepsilon_{R,k} \quad (9)$$

$$U_{S,k} = \alpha_S + \beta V_S + \varepsilon_{S,k} \quad (10)$$

That is, participant k ’s utility for an option is a sum of the value of the option and an unobserved component ε_k . Higher values of β indicate that a larger proportion of participants’ utility is explained by the value of the option, and α_R and α_S denote the expected utility when the unobserved component is zero and the risky and safe options have expected value of zero. Under the assumption that these unobserved

components have independent Type I extreme value distributions, it can be shown that for participant k , the probability $p(O_R|E[O_R|D])$ of choosing the risky option conditional on the data can be written as:

$$p(O_R|E[O_R|D]) = \frac{1}{1 + e^{-[\alpha_R - \alpha_S] - \beta[E[O_R|D] - V_S]}} \quad (11)$$

This is the standard equation for the logit model with $E[O_R|D] - V_S$ as a predictor of choice (McFadden, 1973). If participants estimate $E[O_R|D]$ using a sample-based estimate, their utility for the risky and safe option will be equal, and thus β will not be a significant predictor of choice. However, if people use prior probabilities in their evaluations, options with higher posterior value will have higher utility, and so β will be positive.

To test this model, we ran a logistic regression of the posterior value of the risky option minus the value of the safe option on choice for both the district and individual conditions. In both regressions, the coefficient β was significant and positive (district condition estimate = 3.0062, individual condition estimate = 7.6859, $p < 0.0001$ for both), confirming our hypothesis that people’s choices are influenced by the prior probabilities in the optimal way. That is, increasing the posterior value of the risky option by changing the prior distribution caused people to choose it significantly more often due to its positive effect on utility U .

Discussion

Consistent with previous work on human judgment and perception, we find strong evidence that people’s behavior on

a risky choice task is influenced by prior probabilities in a way that is consistent with Bayesian inference. While there is some evidence that participants chose optimally more often when the safe option was optimal, this bias is relatively small when compared to the effects of changes in prior probabilities. In contrast with most previous work combining real-world data with Bayesian modeling, however, our results are notable as they are not based on a task people commonly face. That people were still influenced by the prior probabilities in the optimal way suggests that people can use their knowledge in decisions in relatively task-independent ways. We argue that this provides additional strength to the case that the study of risky choice can be integrated into the Bayesian approach to studying cognition.

The assumptions of the Bayesian model are especially notable in their similarity to and differences from those of Prospect Theory. For example, the Bayesian model can clarify aspects of Prospect Theory that were not rigorously defined, such as the origins and importance of reference points, which we model as prior probability distributions. Using these distributions gives both a hypothesis of what someone's expectations could be – not just in terms of base-rates, but also in terms of variance and distribution type – but also a rational basis of what these expectations should be. However, the relationship between reference points in Prospect Theory and in the Bayesian model is not entirely straightforward. In our model, we make the assumption that participants have no expectations about the possible payoff values of an option, but do have expectations about probabilities. In Prospect Theory, however, reference points are defined in terms of payouts and not probabilities.

Part of this confusion is due to the fact that in both models, the payoff amount of the risky option has no relationship with its probability. These interactions may not be important in laboratory tasks where people do not have strong prior knowledge, but they may be crucial in the everyday risky decisions that people make. Previous work has shown that the distribution of rewards and probabilities in real-world domains has a general structure that is absent in most studies on the topic (Pleskac & Hertwig, 2014). Furthermore, people use this structure to predict probabilities from outcomes, and outcomes from probabilities. Further generalizations of the Bayesian model in which expectations exist and interact for both payouts and probabilities may help to disentangle these effects and enable clearer and more robust behavioral predictions.

Conclusion

While the study of decision making under risk and uncertainty is a core focus of research in behavioral economics and psychology, it has yet to be unified with Bayesian models of cognition. These models have proved successful in explaining a wide range of psychological and behavioral phenomena, and we show that they can be applied to explain the effects of prior probabilities on risky choice. We believe that

this approach will be crucial to better understanding decision-making under uncertainty in the future.

References

- Anderson, J. R. (1990). *Studies in cognition. the adaptive character of thought*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, 16, 215-233.
- Bernoulli, D. (1954). Exposition of a new theory on the measurement of risk. *Econometrica*, 22(1), 23–36.
- Ert, E., & Haruvy, E. (2017). Revisiting risk aversion: Can risk preferences change with experience? *Economics Letters*, 151, 91 - 95.
- Ert, E., & Lejarraga, T. (in press). The effect of experience on context-dependent decisions. *Journal of Behavioral Decision Making*, 31(4), 535-546.
- Ert, E., & Trautmann, S. (2014). Sampling experience reverses preferences for ambiguity. *Journal of Risk and Uncertainty*, 49(1), 31-42.
- Friedman, D. R., Isaac, M., James, D., & Sunder, S. (2014). *Risky curves: On the empirical failure of expected utility theory*. London: Routledge.
- Gigerenzer, G., & Hoffrage, U. (1995). How to improve bayesian reasoning without instruction: Frequency formats. *Psychological Review*, 102(4), 684-704.
- Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological Science*, 17(9), 767-773.
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15(8), 534-539.
- Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences*, 13(12), 517-523.
- Kahneman, D., Slovic, P., & Tversky, A. e. (1982). *Judgment under uncertainty: heuristics and biases*. Cambridge, England: Cambridge University Press.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263-91.
- Lieder, F., Griffiths, T. L., & Hsu, M. (2018). Over-representation of extreme events in decision making reflects rational use of cognitive resources. *Psychological Review*, 125(1), 1-32.
- McFadden, D. (1973). Conditional logit analysis of qualitative choice behaviour. In P. Zarembka (Ed.), *Frontiers in econometrics* (p. 105-142). New York, NY, USA: Academic Press New York.
- Neumann, J. v., & Morgenstern, O. (1953). *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.
- Pleskac, T. J., & Hertwig, R. (2014). Ecologically rational choice and the structure of the environment. *Journal of Experimental Psychology*, 143(5), 2000 - 2019.
- Rakow, T., & Newell, B. R. (2009). Degrees of uncertainty: An overview and framework for future research on experience-based choice. *Journal of Behavioral Decision Making*, 23(1), 1-14.
- The Cook Political Report. (2017). *2017 cook political report partisan voter index*. (<https://www.cookpolitical.com/introducing-2017-cook-political-report-partisan-voter-index>, Last accessed on 2019-02-01)
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4), 297–323.
- Weber, E. U., Shafir, S., & Blais, A.-R. (2004). Predicting risk-sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review*, 111, 430-445.