

The St. Petersburg Paradox: A Fresh Algorithmic Perspective

Ardavan S. Nobandegani,^{1,3} Thomas R. Shultz^{2,3}

¹Department of Electrical & Computer Engineering, McGill University

²School of Computer Science, McGill University

³Department of Psychology, McGill University
Montreal, QC, Canada

Abstract

The St. Petersburg paradox is a centuries-old puzzle concerning a lottery with infinite expected payoff on which people are only willing to pay a small amount to play. Despite many attempts and several proposals, no generally-accepted resolution is yet at hand. In a recent paper, we show that this paradox can be understood in terms of the mind optimally using its limited computational resources (Nobandegani et al. 2019). Specifically, we show that the St. Petersburg paradox can be accounted for by a variant of normative expected-utility valuation which acknowledges cognitive limitations: *sample-based expected utility* (Nobandegani et al. 2018). SbEU provides a unified, algorithmic explanation of major experimental findings on this paradox. We conclude by discussing the implications of our work for algorithmically understanding human cognition and for developing human-like artificial intelligence.

Originally proposed in 1713 by Nicolas Bernoulli, the St. Petersburg paradox is a famous economic puzzle concerning a risky gamble on which people are invited to place a bid. The gamble goes as follows: The house offers to flip a coin until it comes up heads; the house pays \$1 if heads appears on the first trial (aka initial seed); otherwise the payoff doubles each time tails appears, with this compounding stopping and payment being given at the first heads. The St. Petersburg gamble is outlined in Table 1.

Despite the expected value (EV) of the St. Petersburg gamble being infinite (see Table 1), people are typically willing to place only small bids on this gamble (e.g., Bottom, Bontempo, and Holtgrave 1989; Rivero, Holtgrave, Bontempo, and Bottom 1990; Kroll and Vogt 2009; Cox, Sadiraj, and Vogt 2009; Hayden and Platt 2009). Under the normative stance that people should prefer gambles with higher EVs, this paradox calls human rationality into question: Given that the EV of the gamble is infinite, people should therefore be willing to place arbitrarily large bids on this gamble, but this is far from what experimental evidence indicates.

In a recent paper, we show that the St. Petersburg paradox can be understood in terms of the mind optimally using its

Trial	1	2	3	...	n	...
Event	H	TH	TTH	...	$\overline{TT}\dots\overline{TH}$ <small>$(n-1)$ tails</small>	...
Payoff	\$1	\$2	\$4	...	$\$2^{(n-1)}$...

Table 1: The St. Petersburg gamble. A fair coin is flipped until the first heads appears. On the n th trial of the gamble, corresponding to the event of having the first heads appear on the n th coin flip, the house pays $\$2^{(n-1)}$ to the bidder and the game ends. The expected value (EV) of this gamble is infinite: $EV = \$1 \times (\frac{1}{2}) + \$2 \times (\frac{1}{4}) + \$4 \times (\frac{1}{8}) + \$8 \times (\frac{1}{16}) + \$16 \times (\frac{1}{32}) + \dots = \$\frac{1}{2} + \$\frac{1}{2} + \$\frac{1}{2} + \$\frac{1}{2} + \$\frac{1}{2} + \dots = +\infty$.

limited computational resources (Nobandegani et al. 2019). Specifically, we show that this paradox can be accounted for by a variant of normative expected-utility valuation which acknowledges cognitive limitations: *sample-based expected utility* (SbEU, Nobandegani et al. 2018). Importantly, SbEU adheres to a new mode of inquiry for studying cognition at the algorithmic level of analysis, called *Rational Minimalist Program* (RMP, Nobandegani 2017). RMP holds that, in pursuing optimality, the mind strives to use the minimal amounts of resources, i.e., to take the most economical route to its goal. As such, SbEU has a firm rational basis, which acknowledges the cognitive limitations people are faced with.

Our efforts are simultaneously guided by two well-supported observations about judgment and decision-making under risk: (1) mounting evidence suggests that people often use very few samples in probabilistic judgments and reasoning (Vul et al. 2014; Battaglia et al. 2013; Lake et al. 2017; Gershman, Horvitz, and Tenenbaum 2015; Herwig and Pleskac 2010; Griffiths et al. 2012; Bonawitz et al. 2014; Lieder et al. 2018a), and (2) people overestimate the probability of extreme events in their judgments (Tversky and Kahneman 1973; Ungemach, Chater, and Stewart 2009; Burns, Chiu, and Wu 2010; Barberis 2013; Lieder et al. 2018b). Unlike SbEU, previous explanations of the

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

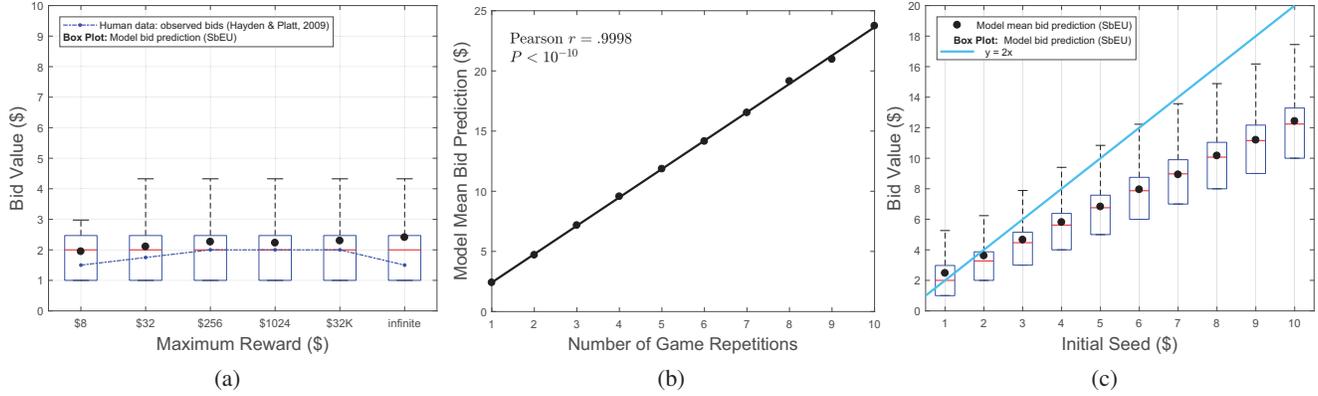


Figure 1: SbEU (Nobandegani et al. 2018) simulating major experimental findings on the St. Petersburg paradox. **(a)** Boxplots of SbEU bid predictions are depicted. SbEU accounts for the experimental finding of Hayden and Platt (2009) showing that bids are only weakly affected by truncating the game. Hayden and Platt truncated the game at 3 flips (maximum payoff: \$8), 5 flips (maximum payoff: \$32), 8 flips (maximum payoff: \$256), 10 flips (maximum payoff: \$1024), and 15 flips (maximum payoff: \$32,768). Black dots show model mean bid predictions. On each box, the central red mark indicates the median, and the bottom and top edges of the box indicate the 25th and 75th percentiles of the data, respectively. On each box, the whisker extends to the most extreme data points not considered outliers. Outliers are not shown. **(b)** SbEU qualitatively simulates the experimental finding of Hayden and Platt (2009) showing that people are willing to place higher bids for a larger number of game repetitions (Pearson $r = .9998$, Kendall $\tau = 1$, Spearman $\rho = 1$, $P_s < .001$). **(c)** Boxplots of SbEU bid predictions are depicted. SbEU accounts for the experimental findings of Hayden and Platt (2009) showing that (1) people’s bids are typically lower than twice the smallest payoff (i.e., initial seed) in the St. Petersburg gamble, and (2) bids depend linearly on the initial seed of the St. Petersburg gamble. The boldfaced blue solid line depicts $y = 2x$.

St. Petersburg paradox fail to respect at least one of these observations.

1 Sample-based Expected Utility Model

SbEU is a rational process model of risky choice that posits that agents rationally adapt their strategies depending on the amount of time available for decision-making (Nobandegani et al. 2018). Concretely, SbEU assumes that an agent estimates expected utility

$$\mathbb{E}[u(o)] = \int p(o)u(o)do, \quad (1)$$

using self-normalized importance sampling (Hammersley and Handscomb 1964; Geweke 1989), with its importance distribution q^* aiming to optimally minimize mean-squared error (MSE):

$$\hat{E} = \frac{1}{\sum_{j=1}^s w_j} \sum_{i=1}^s w_i u(o_i), \quad \forall i: o_i \sim q^*, \quad w_i = \frac{p(o_i)}{q^*(o_i)},$$

$$q^*(o) \propto p(o)|u(o)|\sqrt{\frac{1 + |u(o)|\sqrt{s}}{|u(o)|\sqrt{s}}}. \quad (2)$$

MSE is a standard normative measure of the quality of an estimator, and is widely used in machine learning and mathematical statistics (Poor 2013). In Eqs. (1-2), o denotes an outcome of a risky gamble, $p(o)$ the objective probability of outcome o , $u(o)$ the subjective utility of outcome o , \hat{E}

the importance-sampling estimate of expected utility given in Eq. (1), q^* the importance-sampling distribution, o_i an outcome randomly sampled from q^* , and s the number of samples drawn from q^* .

Recently, Nobandegani et al. (2018) showed that SbEU simulates availability bias, the tendency to overestimate the probability of events that easily come to mind (Tversky and Kahneman 1973), and the well-known fourfold pattern of risk preferences in outcome probability (Tversky and Kahneman 1992) and in outcome magnitude (Markowitz 1952; Scholten and Read 2014). Notably, SbEU is the first rational process model to score near-perfectly in optimality, economical use of limited cognitive resources, and robustness, all at the same time (Nobandegani et al. 2018; Nobandegani et al. 2019a).

2 Simulation Results

In this section, we show that SbEU provides a unified, algorithmic explanation of four major experimental findings on the St. Petersburg paradox: (1) Bids are only weakly affected by truncating the game (e.g., Cox et al. 2007; Neugebauer 2010; Hayden and Platt 2009), (2) Bids are strongly increased by repeating the game (Neugebauer 2010; Hayden and Platt 2009), (3) Bids are typically lower than twice the smallest payoff (Hayden and Platt 2009), and (4) Bids depend linearly on the initial seed of the game (Hayden and Platt 2009).

As shown in Fig. 1, SbEU accounts for these four experimental findings. In Fig. 1, we simulate $N = 1000$ partic-

ipants, with $s = 1$ (see Eq. 2), and use the utility function $\forall x \in \mathbb{R}^{\geq s_0} u(x) = (x - s_0)^{0.35} + s_0$, with s_0 denoting the initial seed of the St. Petersburg gamble. Among plausible utility functions, this utility function stands as a rational choice as it grants that every possible outcome of the St. Petersburg gamble is worth at least s_0 dollars (thus the whole game is subjectively worth at least s_0 dollars), with the utility of higher payoffs of the gamble increasing sublinearly.

3 General Discussion

The St. Petersburg paradox stands among the oldest philosophical puzzles of human decision-making. In this work, we provide an algorithmic-level account of major experimental findings on this paradox. Specifically, we show that a single parameterization of Nobandegani et al.'s (2018) model, SbEU, provides a unified process-level explanation of why (1) bids are only weakly affected by truncating the game, (2) people are willing to place higher bids for a larger number of game repetitions, (3) bids are typically lower than twice the smallest payoff of the game (aka initial seed), and (4) bids depend linearly on the initial seed of the game. As such, Items (1-4) can be understood as optimal behavior subject to cognitive limitations.

Recent work shows that SbEU provides a resource-rational mechanistic account of (ostensibly irrational) cooperation in one-shot Prisoner's Dilemma games, thus successfully bridging between game-theoretic decision-making and risky decision-making (Nobandegani, da Silva Castanheira, Shultz, and Otto 2019b). SbEU also accounts for violation of betweenness in risky choice (Nobandegani, da Silva Castanheira, Shultz, and Otto 2019c) and provides a rational process-level explanation of several contextual effects in risky and value-based decision-making (da Silva Castanheira, Nobandegani, Shultz, and Otto 2019; Nobandegani et al. 2019c). There is also experimental confirmation of a counterintuitive prediction of SbEU: Deliberation leads people to move from one well-known bias, framing effect, to another well-known bias, the fourfold pattern of risk preferences (da Silva Castanheira, Nobandegani, and Otto 2019). Importantly, SbEU is the first, and thus far the only, rational process model that bridges between risky, value-based, and game-theoretic decision-making.

SbEU adheres to a new mode of inquiry for studying cognition at the algorithmic level of analysis: Rational Minimalist Program (RMP, Nobandegani 2017). RMP maintains that, in pursuing optimality, the mind strives to use the minimal amounts of resources, i.e., to take the most economical route to its goal. In addition to the realm of human decision-making, recent work has shown that RMP-inspired models successfully simulate important aspects of a wide range of cognitive phenomena, e.g., developmental shift in infant information processing (Nobandegani 2017, Chap. 2; Nobandegani and Psaromiligkos 2015), causal reasoning (Nobandegani 2017, Chap. 3), action selection in causal domains (Nobandegani 2017, Chap. 4), probabilistic independence judgment (Nobandegani 2017, Chap. 5), and human discriminative and generative abilities (Nobandegani 2017, Chap. 6). Relatedly, RMP bridges between computer science and cognitive science by making contact with a range of core

topics in computer science, e.g., design and analysis of algorithms, data structures, parameterized complexity theory, and distributed computing (Nobandegani 2017, Chap. 7).

Accordingly, a systematic pursuit of RMP in domains that are of great importance for both human cognition and AI would be an effective way toward deepening our understanding of the algorithmic foundation of human cognition and developing human-like AI systems. Future work should investigate this possibility.

The median explanation of Hayden and Platt (2009), that people report the median (and not the mean) of the distribution associated with the St. Petersburg gamble, is currently the only model which can simultaneously account for the four major experimental findings on the St. Petersburg gamble (see Sec. 2). In sharp contrast to the competing median explanation of Hayden and Platt (2009) that is too specific to the St. Petersburg paradox, our work provides a rational process model of this paradox that additionally accounts for several well-known effects in risky, value-based, and game-theoretic decision-making (Nobandegani et al. 2018; da Silva Castanheira et al. 2019; Nobandegani et al. 2019b; Nobandegani et al. 2019c), and is fully in line with the much broader process-level understanding of human probabilistic judgment and reasoning based on sampling (e.g., Stewart, Chater, and Brown 2006; Sanborn and Chater 2016).

There have been several recent studies attempting to show that many well-known (purportedly irrational) behavioral effects and cognitive biases can be understood as optimal behavior subject to computational and cognitive limitations (see Lieder and Griffiths 2018, for a review). Our work contributes to this line of research by showing that SbEU, a rational model of risky choice that adheres to RMP, provides a rational process-level account of a centuries-old puzzle concerning human decision-making. Future work should investigate whether other long-standing paradoxes of human decision-making, e.g., the Ellsberg paradox, could be also understood as optimal behavior subject to cognitive limitations.

References

- Barberis, N. 2013. The psychology of tail events: Progress and challenges. *American Economic Review* 103(3):611–16.
- Battaglia, P. W.; Hamrick, J. B.; and Tenenbaum, J. B. 2013. Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences* 110(45):18327–18332.
- Bernoulli, D. 1738. Specimen theoriae novae de mensura sortis (exposition of a new theory on the measurement of risk). *Comentarii Acad Scient Petropolis* (Translated to English in *Econometrica*) 5(22):23–36.
- Bonawitz, E.; Denison, S.; Griffiths, T. L.; and Gopnik, A. 2014. Probabilistic models, learning algorithms, and response variability: sampling in cognitive development. *Trends in cognitive Sciences* 18(10):497–500.
- Bottom, W. P.; Bontempo, R. N.; and Holtgrave, D. R. 1989. Experts, novices, and the st. petersburg paradox: Is one solution enough? *Journal of Behavioral Decision Making* 2(3):139–147.
- Burns, Z.; Chiu, A.; and Wu, G. 2010. Overweighting of small probabilities. *Wiley Encyclopedia of Operations Research and Management Science*, John Wiley & Sons, Inc.

- Cox, J.; Sadiraj, V.; Vogt, B.; Dasgupta, U.; et al. 2007. Is there a plausible theory for risky decisions? Georgia State University Workin Paper 2010-06.
- Cox, J.; Sadiraj, V.; and Vogt, B. 2009. On the empirical relevance of st. petersburg lotteries. *Economics Bulletin* 29(1):214–220.
- da Silva Castanheira, K.; Nobandegani, A. S.; Shultz, T. R.; and Otto, A. R. 2019. Contextual effects in value-based decision making: A resource-rational mechanistic account [Abstract]. In: *Proceedings of the 41st Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- da Silva Castanheira, K.; Nobandegani, A. S.; and Otto, A. R. 2019. Sample-based variant of expected utility explains effects of time pressure and individual differences in processing speed on risk preferences. In: *Proceedings of the 41st Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Gershman, S. J.; Horvitz, E. J.; and Tenenbaum, J. B. 2015. Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science* 349(6245):273–278.
- Geweke, J. 1989. Bayesian inference in econometric models using monte carlo integration. *Econometrica: Journal of the Econometric Society* 1317–1339.
- Griffiths, T. L.; Vul, E.; and Sanborn, A. N. 2012. Bridging levels of analysis for probabilistic models of cognition. *Current Directions in Psychological Science* 21(4):263–268.
- Hammersley, J., and Handscomb, D. 1964. *Monte Carlo Methods*. London: Methuen & Co Ltd.
- Hayden, B. Y., and Platt, M. L. 2009. The mean, the median, and the st. petersburg paradox. *Judgment and Decision Making* 4(4):256–272.
- Hertwig, R., and Pleskac, T. J. 2010. Decisions from experience: Why small samples? *Cognition* 115(2):225–237.
- Kroll, E. B., and Vogt, B. 2009. The St. Petersburg paradox despite risk-seeking preferences: An experimental study. *Individual Decision Making Under Risk and Uncertainty*. No. F15-V1.
- Lake, B. M.; Ullman, T. D.; Tenenbaum, J. B.; and Gershman, S. J. 2017. Building machines that learn and think like people. *Behavioral and Brain Sciences* 40:1–72.
- Lieder, F.; Griffiths, T. L.; Huys, Q. J.; and Goodman, N. D. 2018a. Empirical evidence for resource-rational anchoring and adjustment. *Psychonomic Bulletin & Review* 25(2):775–784.
- Lieder, F.; Griffiths, T. L.; and Hsu, M. 2018b. Overrepresentation of extreme events in decision making reflects rational use of cognitive resources. *Psychological Review* 125(1):1–32.
- Markowitz, H. 1952. The utility of wealth. *Journal of Political Economy* 60(2):151–158.
- Nobandegani, A. S., and Psaromiligkos, I. N. 2015. Multi-context models for reasoning under partial knowledge: Generative process and inference grammar. In *Proceedings of the 31st Conference on Uncertainty in Artificial Intelligence* (pp. 642–651).
- Nobandegani, A. S.; da Silva Castanheira, K.; Otto, A. R.; and Shultz, T. R. 2018. Over-representation of extreme events in decision-making: A rational metacognitive account. In: *Proceedings of the 40th Annual Conference of the Cognitive Science Society* (pp. 2391–2396). Austin, TX: Cognitive Science Society.
- Nobandegani, A. S.; da Silva Castanheira, K.; Shultz, T. R.; and Otto, A. R. 2019. A resource-rational process-level account of the St. Petersburg paradox. In *Proceedings of the 41st Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Nobandegani, A. S.; da Silva Castanheira, K.; O'Donnell, T. J.; and Shultz, T. R. 2019a. On robustness: An undervalued dimension of human rationality. In *Proceedings of the 17th International Conference on Cognitive Modeling*. Montreal, QC.
- Nobandegani, A. S.; da Silva Castanheira, K.; Shultz, T. R.; and Otto, A. R. 2019b. A resource-rational mechanistic approach to one-shot non-cooperative games: The case of Prisoner's Dilemma. In: *Proceedings of the 41st Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Nobandegani, A. S.; da Silva Castanheira, K.; Shultz, T. R.; and Otto, A. R. 2019c. Decoy effect and violation of betweenness in risky decision making: A resource-rational mechanistic account. In *Proceedings of the 17th International Conference on Cognitive Modeling*. Montreal, QC.
- Nobandegani, A. S. 2017. *The Minimalist Mind: On Minimality in Learning, Reasoning, Action, & Imagination*. McGill University, PhD Dissertation.
- Poor, H. V. 2013. *An Introduction to Signal Detection and Estimation*. Springer Science & Business Media.
- Rivero, J. C.; Holtgrave, D. R.; Bontempo, R. N.; and Bottom, W. P. 1990. The st. petersburg paradox: Data at last. *Commentary*, 8 (3–4) 46–51.
- Sanborn, A. N., and Chater, N. 2016. Bayesian brains without probabilities. *Trends in Cognitive Sciences* 20(12):883–893.
- Scholten, M., and Read, D. 2014. Prospect theory and the “forgotten” fourfold pattern of risk preferences. *Journal of Risk and Uncertainty* 48(1):67–83.
- Stewart, N.; Chater, N.; and Brown, G. D. 2006. Decision by sampling. *Cognitive Psychology* 53(1):1–26.
- Tversky, A., and Kahneman, D. 1973. Availability: A heuristic for judging frequency and probability. *Cognitive Psychology* 5(2):207–232.
- Tversky, A., and Kahneman, D. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5(4):297–323.
- Ungemach, C.; Chater, N.; and Stewart, N. 2009. Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)? *Psychological Science* 20(4):473–479.
- Vul, E.; Goodman, N.; Griffiths, T. L.; and Tenenbaum, J. B. 2014. One and done? optimal decisions from very few samples. *Cognitive Science* 38(4):599–637.