

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

A Resource-Rational Process-Level Account of the St. Petersburg Paradox

Permalink

<https://escholarship.org/uc/item/2c90v6k5>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 41(0)

Authors

Nobandegani, Ardavan S.
Castanheira, Kevin da Silva
Shultz, Thomas R.
[et al.](#)

Publication Date

2019

Peer reviewed

A Resource-Rational Process-Level Account of the St. Petersburg Paradox

Ardavan S. Nobandegani^{1,3}, Kevin da Silva Castanheira³, Thomas R. Shultz^{2,3}, & A. Ross Otto³

{ardavan.salehinobandegani, kevin.dasilvacastanheira}@mail.mcgill.ca

{thomas.shultz, ross.otto}@mcgill.ca

¹Department of Electrical & Computer Engineering, McGill University

²School of Computer Science, McGill University

³Department of Psychology, McGill University

Abstract

The St. Petersburg paradox is a centuries-old philosophical puzzle concerning a lottery with infinite expected payoff, on which people are, nevertheless, willing to place only a small bid. Despite many attempts and several proposals, no generally-accepted resolution is yet at hand. In this work, we present the first resource-rational process-level explanation of this paradox, demonstrating that it can be accounted for by a variant of normative expected-utility-maximization which acknowledges cognitive limitations. Specifically, we show that Nobandegani et al.’s (2018) metacognitively-rational model, *sample-based expected utility* (SbEU), can account for major experimental findings on this paradox. Crucially, our resolution is consistent with two empirically well-supported assumptions: (1) people use only a few samples in probabilistic judgments and decision-making, and (2) people tend to overestimate the probability of extreme events in their judgment.

Keywords: St. Petersburg Paradox; bounded rationality; resource-rational process models; expected utility theory; inference by sampling

1 Introduction

Originally proposed in 1738 by Daniel Bernoulli, the St. Petersburg paradox is a famous economic and philosophical puzzle concerning a risky gamble on which people are asked 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.

Trial	1	2	3	...	n	...
Event	H	TH	TTH	...	$\frac{TT\dots TH}{(n-1) \text{ tails}}$...
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: $EV = \$1 \times (\frac{1}{2}) + \$2 \times (\frac{1}{4}) + \$4 \times (\frac{1}{8}) + \$8 \times (\frac{1}{16}) + \$16 \times (\frac{1}{32}) + \dots = \$\frac{1}{2} + \$\frac{1}{2} + \$\frac{1}{2} + \$\frac{1}{2} + \$\frac{1}{2} + \dots = +\infty$.

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, & Holtgrave, 1989; Rivero, Holtgrave, Bontempo, & Bottom, 1990; Kroll & Vogt, 2009; Cox, Sadiraj, & Vogt, 2009; Hayden & Platt, 2009). Under the normative stance that people should prefer gambles with higher EVs, this paradox calls into question human rationality: The EV of the gamble being infinite, people, therefore, should be willing to place *arbitrarily* large bids on this gamble, but this is far from what experimental evidence suggests.

In spite of its innocent appearance, the St. Petersburg paradox occupied the minds of many over the past two centuries, eliciting a variety of reflections and explanations from several notable thinkers, including Daniel and Niklaus Bernoulli, Cramer, de Morgan, Condorcet, Euler, Poisson, and Gibbon, Marschack, Cournot, Arrow, Keynes, Stigler, Samuelson, von Mises, Ramsey and Aumann (see Arrow, 1951; Aumann, 1977; Dutka, 1988; Keynes, 1921; Samuelson, 1960). Nonetheless, no widely accepted explanation of this paradox exists to date.

In this work, we ask whether people’s bids on the St. Petersburg paradox could be understood as an optimal behavior with the mind acting as a cognitive miser. Answering this question in the affirmative, we show that the St. Petersburg paradox can be accounted for by a variant of normative expected-utility-maximization which acknowledges computational and cognitive limitations. Specifically, we demonstrate that Nobandegani, da Silva Castanheira, Otto, and Shultz’s (2018) metacognitively-rational model, *sample-based expected utility* (SbEU), can account for major experimental findings on the St. Petersburg paradox.

In the present study, our efforts are simultaneously guided by two well-supported observations about human judgment and decision-making under risk: (1) mounting evidence suggests that people often use very few samples in probabilistic judgments and reasoning (e.g., Vul et al., 2014; Battaglia et al., 2013; Lake et al., 2017; Gershman, Horvitz, & Tenenbaum, 2015; Hertwig & Pleskac, 2010; Griffiths et al., 2012; Gershman, Vul, & Tenenbaum, 2012; Bonawitz et al., 2014; Nobandegani et al., 2018; Lieder, Griffiths, Huys, & Goodman, 2018), and (2) people overestimate the probability of extreme events in their judgments (e.g., Tversky & Kahneman, 1972; Ungemach, Chater, & Stewart, 2009; Burns, Chiu, & Wu, 2010; Barberis, 2013; Lieder et al., 2018). As we discuss in the next section, previous explanations of the St. Petersburg paradox fail to respect at least one of these observations.

Our paper is organized as follows. We begin by present-

ing a brief historical overview of major explanations of the St. Petersburg paradox. After providing a brief overview of SbEU, we turn to modeling four major experimental findings on the St. Petersburg paradox. We conclude by discussing the implications of our work for the debate on human rationality.

2 A Brief Historical Overview of the Paradox

In this section, we present a brief overview of major resolutions of the St. Petersburg paradox, followed by notable critiques of them.

It is worth noting that most of the work on the St. Petersburg paradox thus far has been theoretical or philosophical. Comparatively little effort has been directed at providing empirical data on the bids people would be willing to place on the gamble and/or how people's bids would be affected by changing focal characteristics of the gamble, e.g. by varying the initial seed or limiting the number of coin flips in the gamble (e.g., Bottom, Bontempo, and Holtgrave, 1989; Rivero, Holtgrave, Bontempo, and Bottom, 1990; Kroll and Vogt, 2009; Cox, Sadiraj, and Vogt, 2011; Hayden & Platt, 2009; Neugebauer, 2010).

Diminishing marginal utility. Initially presented by Daniel Bernoulli (1738), the diminishing marginal utility explanation of the St. Petersburg paradox argues that, instead of evaluating the expected value (EV) of the gamble (which is infinite, see Table 1), people evaluate the expected utility of the gamble, with the utility function having a concave form (aka diminishing marginal utility).

As this explanation fails to account for super-St. Petersburg paradoxes in which the gamble's payoff increases super-exponentially with every coin flip, recent discussions of this explanation have to make the further assumption that the utility function is bounded from above (e.g., Aumann, 1977; Martin, 2008; Menger, 1934; Samuelson, 1977; Vickrey, 1960).

The diminishing marginal utility explanation has been discredited several times, mainly because it over-predicts bids (Lopes, 1981; Martin, 2008; Menger, 1934; Moritz, 1923; Samuelson, 1960, 1977). (This is not to say that marginal utility does not diminish, just that this factor is insufficient to explain the paradox.) Also, the diminishing marginal utility explanation completely neglects the well-supported observation that people overestimate the probability of extreme events in their judgment (e.g., Tversky & Kahneman, 1972; Ungemach, Chater, & Stewart, 2009; Burns, Chiu, & Wu, 2010; Barberis, 2013; Lieder et al., 2018), mistakenly assuming that the subjective probability of a low-probability extreme event in the St. Petersburg gamble (e.g., to win $\$2^{100}$ with probability $\frac{1}{2^{101}}$) is equal to its objective probability (e.g., $\frac{1}{2^{101}}$). Replacing expected utility with more modern variants which respect the latter observation, e.g. cumulative prospect theory (CPT), does not help either, as empirically fit values strongly over-predict bids in the St. Petersburg paradox (Blavatsky, 2005; Rieger & Wang, 2006; Camerer, 2005).

Finitude of resources. Another classic explanation is that

since the amount of money in the world is finite, the gambler must be skeptical about the ability of the house to pay the large outcomes of the gamble. Relatedly, it has been argued that time is finite, and the gambler, knowing he or she cannot continue playing the game forever, bids less than the expected value of the gamble. This argument has been expressed, in various forms, by several scholars (see Savage, 1954; Tversky & Bar-Hillel, 1983; Vickrey, 1960; Dutka, 1988).

Weaknesses of these arguments have been explicated by several critics. Bertrand argues that, even if the house cannot afford to pay the money, unites of currency can be reasonably replaced by more plentiful stuff, such as grains of sand, inches, or molecules of hydrogen, and the risk aversion still remains (Dutka, 1988). By the same logic, the payment may even be hypothetical or psychological (Martin, 2008; Aumann, 1977).

Ignoring low probabilities. This explanation argues that people consider events whose probability falls below some threshold to be impossible, i.e. they never happen. For example, D'Alembert posited a 1/10,000 threshold, while Niklaus Bernoulli set the cutoff at a more conservative 1/100,000 (Dutka, 1988).

However, there is a serious flaw with this argument: According to the well-known availability bias (Tversky & Kahneman, 1972), people over-represent extreme events, i.e., events whose utility is large (Lieder et al., 2018; Nobandegani et al., 2018). As low-probability events have (exponentially) larger payoffs in the St. Petersburg gamble, people should overestimate those low-probability events, putting more weights on those low-probability events in their valuation of the gamble.

A key contribution of our work is to provide a resource-rational process-level explanation of why people are willing to place only a small bid on the gamble *despite* over-representing extreme events in their judgment and decision-making (see Sec. 3). Particularly, past work has shown that SbEU can account for availability bias (Nobandegani et al., 2018).

Computing the median instead of the mean. Recently, Hayde and Platt (2009) proposed that people report the median (and not the mean) of the distribution associated with the St. Petersburg gamble as their bid. The median of the distribution associated with the St. Petersburg gamble is between \$1 and \$2, and is set by convention at \$1.50 (Weissstein, 2008).

The median explanation of Hayde and Platt (2009) is currently the only model which can simultaneously account for all the major experimental findings on the St. Petersburg gamble. We investigate all these major experimental findings in the present study in Sec. 4.

Nevertheless, despite its quantitative coverage, the median explanation remains too limited to explain the St. Petersburg paradox, markedly detached from the extensive literature on human judgment and decision-making. Similar to the diminishing marginal utility explanation, the median explanation completely neglects the well-supported observation that people

ple overestimate the probability of extreme events in their judgment (e.g., Tversky & Kahneman, 1972; Lieder et al., 2018), mistakenly assuming that the subjective probability of a low-probability extreme event in the St. Petersburg gamble is equal to its objective probability.

In this work, we seek to provide a resource-rational process model of the St. Petersburg paradox that can additionally account for several well-known effects in decision-making under risk; SbEU meets this criterion (see Sec. 3). As such, we seek to understand the St. Petersburg gamble as a particular risky gamble whose process-level explanation should be consistent with a broader process-level model of decision-making under risk.

3 Sample-based Expected Utility Model

Extending the cognitively-rational decision-making model of Lieder, Griffiths, and Hsu (2018) to the realm of metacognition (Cary & Reder, 2002), SbEU is a metacognitively-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 & Handscomb, 1964; Geweke, 1989), with its importance distribution q^* aiming to optimally minimize mean-squared error (MSE):

$$\hat{E} = \frac{1}{\sum_{j=1}^s w_j} \sum_{i=1}^s w_i u(o_i), \quad \forall i: o_i \sim q^*, w_i = \frac{p(o_i)}{q^*(o_i)}, \quad (2)$$

$$q^*(o) \propto p(o)|u(o)| \sqrt{\frac{1 + |u(o)|\sqrt{s}}{|u(o)|\sqrt{s}}}. \quad (3)$$

MSE is a standard normative measure of the quality of an estimator, and is widely adopted in machine learning and mathematical statistics (Poor, 2013). In Eqs. (1-3), 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 can account for availability bias, the tendency to overestimate the probability of events that easily come to mind (Tversky & Kahneman, 1972), and can accurately simulate the well-known fourfold pattern of risk preferences in outcome probability (Tversky & Kahneman, 1992) and in outcome magnitude (Markovitz, 1952; Scholten & 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).

4 Simulation Results

In this section, we show that SbEU can quantitatively account for 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 & Platt, 2009), (2) Bids are strongly increased by repeating the game (Neugebauer, 2010; Hayden & Platt, 2009), (3) Bids are typically lower than twice the smallest payoff (Hayden & Platt, 2009), and (4) Bids depend linearly on the initial seed of the game (Hayden & Platt, 2009).

Recent work has provided mounting evidence suggesting that people often use very few samples in probabilistic judgments and reasoning (e.g., Vul et al., 2014; Battaglia et al. 2013; Lake et al., 2017; Gershman, Horvitz, & Tenenbaum, 2015; Hertwig & Pleskac, 2010; Griffiths et al., 2012; Gershman, Vul, & Tenenbaum, 2012; Bonawitz et al., 2014; Nobandegani et al., 2018; Lieder, Griffiths, Huys, & Goodman, 2018). Consistent with this finding, in the present study we assume that bidders draw only one sample ($s = 1$; see Eqs. 2-3) when evaluating their (subject) expected utility of the St. Petersburg gamble.

Concretely, we use the Metropolis–Hastings Markov chain Monte Carlo (MCMC) method—a well-known rational process model for sampling from a probability distribution of interest—to generate a single sample ($s = 1$) from the importance distribution q^* given in Eq. 3. MCMC methods have been successful in simulating important aspects of a wide range of cognitive phenomena, e.g., temporal dynamics of multistable perception (Gershman et al., 2012; Moreno-Bote et al., 2011), developmental changes in cognition (Bonawitz, Denison, Griffiths, & Gopnik, 2014), category learning (Sanborn et al., 2010), and accounting for many cognitive biases (Nobandegani et al., 2018; Dasgupta et al., 2016).

Also, consistent with prospect theory (Kahneman & Tversky, 1979) and cumulative prospect theory (Kahneman & Tversky, 1992), in this paper we assume a standard S-shaped utility function $u(x)$ given by:

$$u(x) = \begin{cases} x^{0.35} & \text{if } x \geq 0, \\ -|x|^{0.45} & \text{if } x < 0. \end{cases} \quad (4)$$

4.1 Bids are weakly affected by truncating the game

In the original St. Petersburg gamble, there is no a priori upper-bound on number of coin flips; theoretically it can continue indefinitely. In a truncated variant of the St. Petersburg gamble, some a priori upper-bound is placed on the number of coin flips. Several experimental studies have shown that bids that people are willing to offer to play the St. Petersburg gamble are only weakly affected by truncating the game (Cox et al., 2007; Cox et al., 2008, 2009; Hayden & Platt, 2009). This finding is generally taken as evidence for people ignoring small-probability events in the game (Neugebauer, 2010).

Recently, Hayden and Platt (2009) investigated bids for the St. Petersburg gamble truncated at 3 flips (maximum payoff: \$8, EV: \$2.50), 5 flips (maximum payoff: \$32, EV: \$3.50),



Figure 1: Hayden and Platt’s (2009) experimental data on the effect of truncation on bids for the St. Petersburg gamble. Adapted from Hayden and Platt (2009).

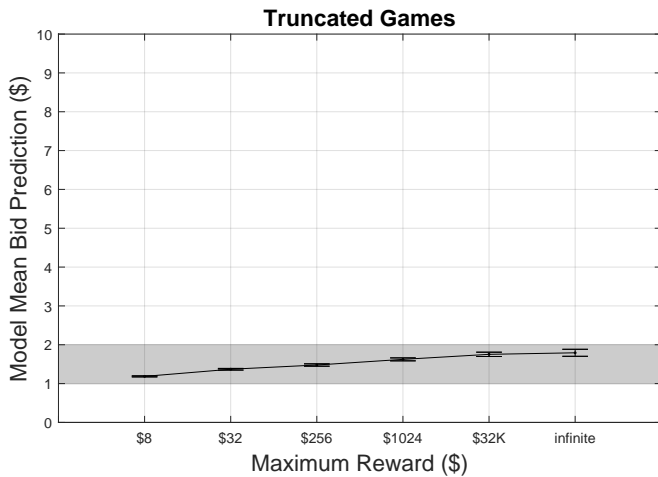


Figure 2: SbEU (Nobandegani et al., 2018) can accurately simulate the experimental data of Hayden and Platt (2009) on the effect of truncation on bids for the St. Petersburg gamble. Error bars indicate ± 1 SEM.

8 flips (maximum payoff: \$256, EV: \$5), 10 flips (maximum payoff: \$1024, EV: \$6) and 15 flips (maximum payoff: \$32,768, EV: \$8.50); their experimental data are shown in Fig. 1.

Fig. 2 shows that SbEU can account for the experimental data of Hayden and Platt (2009). In Fig. 2, we simulate $N = 1000$ participants, with $s = 1$.

4.2 Bids rise with repetitions of the game

Recently, Hayden and Platt (2009) experimentally showed that bids to play the (un-truncated) St. Petersburg gamble are strongly affected by repeating the game, with people willing

to place higher bids with a larger number of game repetitions.

Fig. 3 shows that SbEU can qualitatively simulate people’s tendency to place higher bids for a larger number of game repetitions, as experimentally shown by Hayden and Platt (2009). In Fig. 3, we simulate $N = 1000$ participants, with $s = 1$.

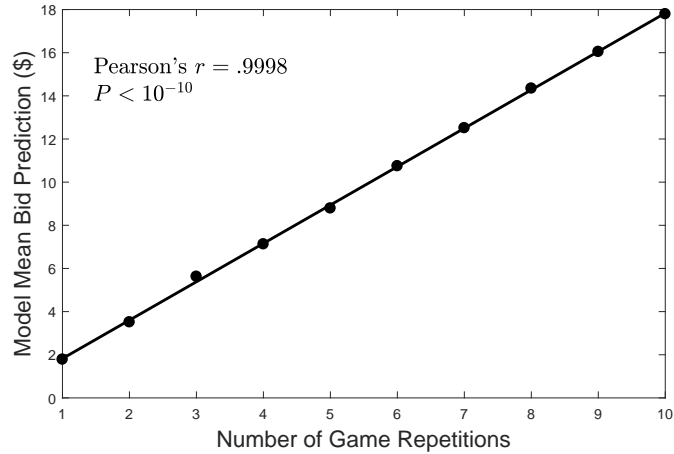


Figure 3: SbEU (Nobandegani et al., 2018) can account for the experimental finding of Hayden and Platt (2009) showing that people willing to place higher bids for a larger number of game repetitions (Pearson’s $r = .9998$, Kendall’s $\tau = 1$, Spearman’s $\rho = 1$, $P_s < .001$).

4.3 Bids are typically lower than twice the smallest payoff

In their recent work, Hayden and Platt (2009) showed that bids to play the (un-truncated) St. Petersburg gamble are typically lower than twice the smallest payoff of the game.

Fig. 4 shows that SbEU can account for this experimental finding of Hayden and Platt (2009). In Fig. 4, we simulate $N = 1000$ participants, with $s = 1$.

4.4 Bids depend linearly on the initial seed

Interestingly, Hayden and Platt (2009) showed that bids to play the (un-truncated) St. Petersburg gamble depend linearly on the initial seed of the game, thus providing a quantitatively well-characterized criterion for evaluating a computational account.

Fig. 5 shows that SbEU can accurately account for this experimental finding of Hayden and Platt (2009) (Pearson’s $r = .9758$, Kendall’s $\tau = 0.9556$, Spearman’s $\rho = .9879$, $P_s < .001$). In Fig. 5, we simulate $N = 1000$ participants, with $s = 1$.

5 General Discussion

The St. Petersburg paradox (Bernoulli, 1738) stands among the oldest philosophical puzzles of human decision-making, and has played a pivotal role in the emergence of the concept of the subjective utility curve, a central concept in economics

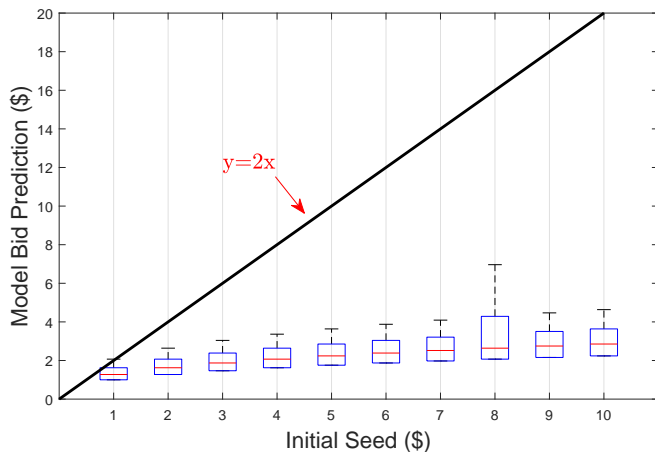


Figure 4: Boxplots of the model's bids. SbEU can account for the experimental finding of Hayden and Platt (2009) showing that people's bids are typically lower than twice the smallest payoff (i.e. initial seed) in the St. Petersburg gamble. On each box, the central red mark indicates the median, and the bottom and top edges of the box indicate the 25th (denoted by q_1) and 75th (denoted by q_3) percentiles of the data, respectively. On each box, the whisker extends to the most extreme data points not considered outliers. Outliers are data points that lie outside the interval $[q_1 - 1.5 \times (q_3 - q_1), q_3 + 1.5 \times (q_3 - q_1)]$, and are not shown in this plot. The boldfaced black solid line depicts $y = 2x$.

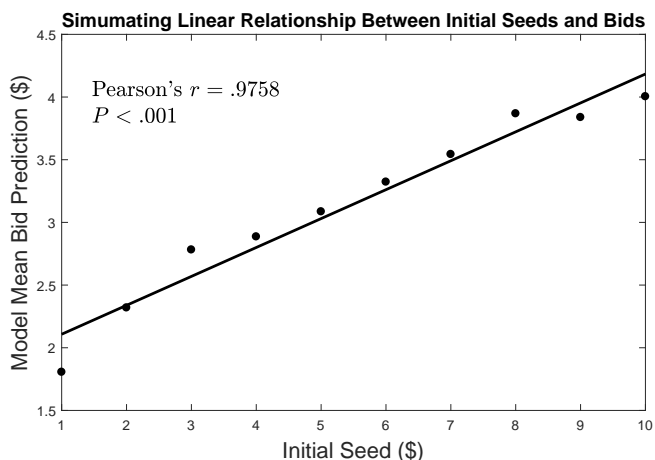


Figure 5: SbEU (Nobandegani et al., 2018) can account for the experimental finding of Hayden and Platt (2009) showing that bids depend linearly on the initial seed of the St. Petersburg gamble (Pearson's $r = .9758$, Kendall's $\tau = 0.9556$, Spearman's $\rho = .9879$, $P_s < .001$).

(Dutka, 1988). Despite occupying the minds of many important thinkers, eliciting many attempts and several proposals, no generally-accepted resolution is yet at hand.

In this work, we provide an algorithmic-level account of major experimental findings on the St. Petersburg para-

dox. Specifically, we show that a single parameterization of Nobandegani et al.'s (2018) metacognitively-rational model, SbEU, provides a unified, resource-rational, process-level explanation of (1) why bids are only weakly affected by truncating the game, (2) why people are willing to place higher bids for a larger number of game repetitions, (3) why bids are typically lower than twice the smallest payoff of the game (aka initial seed), and (4) why 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.

As opposed to the competing median explanation of Hayden and Platt (2009) that is too specific to the St. Petersburg paradox, our work provides a resource-rational process model of the St. Petersburg paradox that can additionally account for several well-known effects in decision-making under risk (Nobandegani et al., 2018), 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, & Brown, 2006; Sanborn & Chater, 2016).

Recent work has shown 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, & Otto, 2019b). 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, & Otto, 2019).

Crucially, our explanation retains the well-supported assumption that people overestimate the probability of extreme events in their judgment and decision-making (Tversky & Kahneman, 1972; Lieder et al., 2018; Nobandegani et al., 2018), and is fully in line with mounting evidence suggesting that people use only a few samples in probabilistic judgments and reasoning (e.g., Vul et al., 2014; Battaglia et al. 2013; Lake et al., 2017; Gershman, Horvitz, & Tenenbaum, 2015; Hertwig & Pleskac, 2010; Griffiths et al., 2012; Gershman, Vul, & Tenenbaum, 2012; Bonawitz et al., 2014; Nobandegani et al., 2018; Lieder, Griffiths, Huys, & Goodman, 2018).

Recently, Blavatsky (2005) showed that conventional parameterizations of cumulative prospect theory (CPT; Kahneman & Tversky, 1992) do not explain the St. Petersburg paradox. As we demonstrate in this work, assuming a standard S-shaped utility function, as advocated by CPT, suffices for explaining the St. Petersburg paradox with SbEU (see Eq. 4).

There have been several recent studies (see Lieder & Griffiths, 2018, for a review) 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 (*rational minimalist program*, Nobandegani, 2017; Griffiths, Lieder, & Goodman, 2015). The present study contributes to this line of work by providing a resource-rational process-level account of a

centuries-old puzzle concerning human decision-making.

Future work should investigate whether other long-standing paradoxes of human judgment and decision-making, e.g., the Ellsberg paradox (Ellsberg, 1961), could be also understood as optimal behavior subject to cognitive limitations. We see our work as a step in this direction.

Acknowledgments This work is supported by an operating grant to TRS from the Natural Sciences and Engineering Research Council of Canada. We would like to thank Constance Destais, Ashley Stendel, Marcel Montrey, and Peter Helfer for helpful comments on an earlier draft of this work.

References

- Arrow, K. J. (1951). Alternative approaches to the theory of choice in risk-taking situations. *Econometrica: Journal of the Econometric Society*, 404–437.
- Aumann, R. J. (1977). The st. petersburg paradox: A discussion of some recent comments. *Journal of Economic Theory*, 14(2), 443–445.
- Barberis, N. (2013). The psychology of tail events: Progress and challenges. *American Economic Review*, 103(3), 611–16.
- Battaglia, P. W., Hamrick, J. B., & 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 in *Econometrica*), 5(22), 23–36.
- Blavatsky, P. R. (2005). Back to the st. petersburg paradox? *Management Science*, 51(4), 677–678.
- Bonawitz, E., Denison, S., Griffiths, T. L., & 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., & 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., & Wu, G. (2010). Overweighting of small probabilities. *Wiley Encyclopedia of Operations Research and Management Science*.
- Camerer, C. (2005). Three cheers—psychological, theoretical, empirical—for loss aversion. *Journal of Marketing Research*, 42(2), 129–133.
- Cary, M., & Reder, L. M. (2002). Metacognition in strategy selection. In *Metacognition: Process, Function and Use* (pp. 63–77). Springer.
- Cox, J., Sadiraj, V., & Vogt, B. (2009). On the empirical relevance of st. petersburg lotteries. *Economics Bulletin*, 29(1), 214–220.
- Cox, J., Sadiraj, V., Vogt, B., Dasgupta, U., et al. (2007). Is there a plausible theory for risky decisions? *Georgia State University*.
- Cox, J. C., & Sadiraj, V. (2008). Risky decisions in the large and in the small: Theory and experiment. In *Risk aversion in experiments: Research in experimental economics* (Vol. 12, pp. 9–40). Emerald Group Publishing Limited.
- Dasgupta, I., Schulz, E., & Gershman, S. J. (2016). Where do hypotheses come from? Center for Brains, Minds and Machines (CBMM) Memo No. 056.
- da Silva Castanheira, K., Nobandegani, A. S., & 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.
- Dutka, J. (1988). On the st. petersburg paradox. *Archive for History of Exact Sciences*, 39(1), 13–39.
- Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics*, 75(4), 643–669.
- Gershman, S. J., Horvitz, E. J., & Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, 349(6245), 273–278.
- Gershman, S. J., Vul, E., & Tenenbaum, J. B. (2012). Multistability and perceptual inference. *Neural Computation*, 24(1), 1–24.
- 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., & Sanborn, A. N. (2012). Bridging levels of analysis for probabilistic models of cognition. *Current Directions in Psychological Science*, 21(4), 263–268.
- Hammersley, J., & Handscomb, D. (1964). *Monte carlo methods*. London: Methuen & Co Ltd.
- Hayden, B. Y., & Platt, M. L. (2009). The mean, the median, and the st. petersburg paradox. *Judgment and Decision Making*, 4(4), 256–272.
- Hertwig, R., & Pleskac, T. J. (2010). Decisions from experience: Why small samples? *Cognition*, 115(2), 225–237.
- Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive psychology*, 3(3), 430–454.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–291.
- Keynes, J. M. (1921). *A treatise on probability*. Courier Corporation.
- Kroll, E. B., & Vogt, B. (2009). The st. petersburg paradox despite risk-seeking preferences. an experimental study.
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40.

- Lieder, F., & Griffiths, T. L. (2018). Resource-rational analysis: understanding human cognition as the optimal use of limited computational resources. Available on Researchgate.
- Lieder, F., Griffiths, T. L., & Hsu, M. (2018). Overrepresentation of extreme events in decision making reflects rational use of cognitive resources. *Psychological Review*, *125*(1).
- Lieder, F., Griffiths, T. L., Huys, Q. J., & Goodman, N. D. (2018). Empirical evidence for resource-rational anchoring and adjustment. *Psychonomic Bulletin & Review*, *25*(2), 775–784.
- Lopes, L. L. (1981). Decision making in the short run. *Journal of Experimental Psychology: Human Learning and Memory*, *7*(5), 377.
- Markowitz, H. (1952). The utility of wealth. *Journal of Political Economy*, *60*(2), 151–158.
- Martin, R. (2017). The st. petersburg paradox. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2017 ed.). Stanford University.
- Menger, K. (1934). Das unsicherheitsmoment in der wertlehre. *Z Nationaloeken*, *51*:(459–485).
- Moreno-Bote, R., Knill, D. C., & Pouget, A. (2011). Bayesian sampling in visual perception. *Proceedings of the National Academy of Sciences*, *108*(30), 12491–12496.
- Moritz, R. E. (1923). Some curious fallacies in the study of probability. *The American Mathematical Monthly*, *30*(2), 58–65.
- Neugebauer, T., et al. (2010). Moral impossibility in the petersburg paradox: a literature survey and experimental evidence. *Luxembourg School of Finance Research Working Paper Series*(10-14).
- Nobandegani, A. S. (2017). The Minimalist Mind: On Minimality in Learning, Reasoning, Action, & Imagination. McGill University, PhD Dissertation.
- Nobandegani, A. S., da Silva Castanheira, K., O'Donnell, T. J., & 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., Otto, A. R., & 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., Otto, A. R., & Shultz, T. 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.
- Poor, H. V. (2013). *An Introduction to Signal Detection and Estimation*. Springer Science & Business Media.
- Rieger, M. O., & Wang, M. (2006). Cumulative prospect theory and the st. petersburg paradox. *Economic Theory*, *28*(3), 665–679.
- Rivero, J. C., Holtgrave, D. R., Bontempo, R. N., & Bottom, W. P. (1990). The st. petersburg paradox: Data at last. *Commentary*, *8* (3–4), 46–51.
- Samuelson, P. A. (1960). The st. petersburg paradox as a divergent double limit. *International Economic Review*, *1*(1), 31–37.
- Samuelson, P. A. (1977). St. petersburg paradoxes: De-fanged, dissected, and historically described. *Journal of Economic Literature*, *15*(1), 24–55.
- Sanborn, A. N., & Chater, N. (2016). Bayesian brains without probabilities. *Trends in Cognitive Sciences*, *20*(12), 883–893.
- Sanborn, A. N., Griffiths, T. L., & Navarro, D. J. (2010). Rational approximations to rational models: alternative algorithms for category learning. *Psychological Review*, *117*(4), 1144.
- Savage, L. J. (1954). *The foundations of statistics*. Courier Corporation.
- Scholten, M., & 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., & Brown, G. D. (2006). Decision by sampling. *Cognitive Psychology*, *53*(1), 1–26.
- Tversky, A., & Bar-Hillel, M. (1983). Risk: The long and the short. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *9*, 713–717.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, *5*(4), 297–323.
- Ungemach, C., Chater, N., & Stewart, N. (2009). Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)? *Psychological Science*, *20*(4), 473–479.
- Vickrey, W. (1960). Utility, strategy, and social decision rules. *The Quarterly Journal of Economics*, *74*(4), 507–535.
- Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? optimal decisions from very few samples. *Cognitive Science*, *38*(4), 599–637.
- Weisstein, E. W. (2008). Statistical median. *From MathWorld – A Wolfram Web Resource* <http://mathworld.wolfram.com/StatisticalMedian.html>.