

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

A Resource-Rational Process Model of Fairness in the Ultimatum Game

Permalink

<https://escholarship.org/uc/item/7wg1g11r>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 42(0)

Authors

Nobandegani, Ardavan S.

Destais, Constance

Shultz, Thomas R.

Publication Date

2020

Peer reviewed

A Resource-Rational Process Model of Fairness in the Ultimatum Game

Ardavan S. Nobandegani^{1,4}, Constance Destais², & Thomas R. Shultz^{3,4}

{ardavan.salehinobandegani}@mail.mcgill.ca

{constance.destais.16}@ucl.ac.uk

{thomas.shultz}@mcgill.ca

¹Department of Electrical & Computer Engineering, McGill University

²Department of Psychology, University College London

³School of Computer Science, McGill University

⁴Department of Psychology, McGill University

Abstract

Widely regarded as the cornerstone of justice (Rawls, 1971), fairness constitutes one of the pillars of human morality. The Ultimatum Game (UG), extensively studied in behavioral economics, is the canonical task for studying fairness. In sharp contrast to the predictions of normative standards in game theory, people typically reject low offers in UG. In this work, we present the first resource-rational process model of UG. Concretely, by taking into account people's expectations, we show that Nobandegani et al.'s (2018) resource-rational process model, *sample-based expected-utility*, provides a unified account of several experimental findings in UG, namely, the effects of expectation, competition, and time pressure. Assuming that expectation serves as a reference point for subjective valuation of an offer, we show that the rejection of low offers in UG can arise from purely self-interested expected-utility maximization. We conclude by discussing the implication of our work for moral decision-making and, more broadly, human rationality.

Keywords: Ultimatum game; moral decision-making; fairness; rational process models

The fundamental idea in the concept of justice is fairness.

- John Rawls (1971)

1 Introduction

Widely considered as one of the pillars of human morality, fairness is chiefly investigated in the context of the Ultimatum Game (UG, Güth et al., 1982), an extensively studied game in psychology (e.g., Sanfey, 2009; Battigalli et al., 2015; Vavra et al., 2018), neuroscience (Sanfey et al., 2003; Xiang et al., 2013; Chang & Sanfey, 2013), philosophy (Guala, 2008), and behavioral economics (e.g., Güth et al., 1982; Thaler, 1988; Camerer & Thaler, 1995; Fehr & Schmidt, 1999; Sutter et al., 2003; Camerer & Fehr, 2006). UG has a simple design: Two players, Proposer and Responder, have to agree on how to split a sum of money. Proposer makes an offer. If Responder accepts, the deal goes ahead; if Responder rejects, neither player gets anything. In both cases, the game is over.

The normative standards of backward induction and subgame perfect equilibrium predict that Responder will accept any nonzero offer, with the rationale being that any positive amount, even if minuscule, is better than nothing at all (Camerer & Thaler, 1995). Nevertheless, in sharp contrast to the predictions of these normative standards, a substantial body of experimental work has shown that Responders predominantly reject offers below 30% (e.g., Güth et al., 1982;

Thaler, 1988; Güth & Tietz, 1990; Bolton & Zwick, 1995; Nowak et al., 2000; Camerer & Fehr, 2006). The algorithmic foundation of this puzzling behavior has remained largely elusive.

There have been speculations about the role of bounded rationality in UG (e.g., Camerer & Fehr, 2006; Van Damme et al., 2014). Inspired by these speculations, in this work we ask whether Responder's behavior could be understood as optimal behavior with the mind acting as a cognitive miser. Concretely, by taking into account people's expectations, we show that Nobandegani et al.'s (2018) resource-rational process model, *sample-based expected-utility* (SbEU), provides a unified account of several disparate experimental findings in UG, namely, the effects of expectation, competition, and time pressure.

Importantly, a series of behavioral (Sanfey, 2009; Battigalli et al., 2015; Vavra et al., 2018) and neuroimaging studies (Xiang et al., 2013; Chang & Sanfey, 2013) have recently revealed that expectation plays a crucial role in UG, with Responders adapting their behavior based on the amount of offer they expect. Particularly, Sanfey (2009) and Vavra et al. (2018) experimentally showed that when Responders expect to see low offers, they are more likely to accept these low offers than when their prior expectations are higher, thus providing supporting evidence for the hypothesis that Responders' prior expectation serves as a reference point for subjective valuation of offers.

This paper is organized as follows. After elaborating on the computational underpinnings of SbEU (Sec. 2), we turn to modeling the effects of expectation, competition, and time pressure on Responders' behavior in UG (Sec. 3). We conclude by discussing the implication of our work for human strategic behavior, moral decision-making, and, more broadly, rationality.

2 Computational Model

We now present our resource-rational process model of Responders' behavior in UG. We assume that, adopting their probabilistic expectation as a reference point for subjective valuation, Responders use SbEU to provide a resource-rational response. That is, Responders optimally maximize their expected utility, but this maximization is subject to their cognitive limitations.

More specifically, we assume that Responder uses SbEU

to estimate the expected-utility gap between their expectation and the offer, i.e., $\mathbb{E}[u(\text{offer}) - u(\text{expectation})]$, where $u(\cdot)$ denotes Responder’s utility function. If this estimate is positive — indicating that the offer made is, on average, higher than Responder’s expectation — Responder accepts the offer; otherwise, Responder rejects the offer.

2.1 Sample-based Expected Utility Model

SbEU is a metacognitively-rational process model of risky decision-making that posits that an agent rationally adapts 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 measure of estimation quality, widely used in decision theory 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^* .

SbEU assumes that, when choosing between a pair of risky gambles A, B , people choose depending on whether the expected value of the utility difference $\Delta u(o)$ is negative or positive (w.p. stands for “with probability”):

$$A = \begin{cases} o_A & \text{w.p. } P_A \\ 0 & \text{w.p. } 1 - P_A \end{cases} \quad (4)$$

$$B = \begin{cases} o_B & \text{w.p. } P_B \\ 0 & \text{w.p. } 1 - P_B \end{cases} \quad (5)$$

$$\Delta u(o) = \begin{cases} u(o_A) - u(o_B) & \text{w.p. } P_A P_B \\ u(o_A) - u(0) & \text{w.p. } P_A(1 - P_B) \\ u(0) - u(o_B) & \text{w.p. } (1 - P_A)P_B \\ 0 & \text{w.p. } (1 - P_A)(1 - P_B) \end{cases} \quad (6)$$

In Eq. (6), $u(\cdot)$ denotes the subjective utility function of a decision-maker. In this paper, we assume the same utility function $u(x)$ used by Nobandegani et al. (2018, 2019a) to

explain both the fourfold pattern of risk preferences and cooperation in one-shot Prisoner’s Dilemma games:

$$u(x) = \begin{cases} x^{0.85} & \text{if } x \geq 0, \\ -|x|^{0.95} & \text{if } x < 0. \end{cases} \quad (7)$$

As such, in this work we do *not* fine-tune the utility function to maximize descriptive power.

Nobandegani et al. (2018) showed that SbEU can account for people’s tendency to overestimate the probability of events that easily come to mind (availability bias, Tversky & Kahneman, 1973), and can 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 (see Nobandegani et al., 2018; Nobandegani et al., 2019b).

Relatedly, recent work has shown that SbEU provides a resource-rational mechanistic account of cooperation in one-shot Prisoner’s Dilemma games (Nobandegani et al., 2019a) and human coordination strategies in coordination games (Nobandegani & Shultz, 2020a), thus successfully bridging between game-theoretic and risky decision-making. SbEU can also account for violation of betweenness in risky choice (Nobandegani et al., 2019c), the centuries-old St. Petersburg paradox in human decision-making (Nobandegani & Shultz, 2020b, 2020c), and provides a resource-rational process-level explanation of several contextual effects in risky and value-based decision-making (da Silva Castanheira, Nobandegani, Shultz, & Otto, 2019; Nobandegani et al., 2019c). There is also experimental confirmation of a counterintuitive prediction of SbEU: Deliberation leads people to move from one cognitive bias, the framing effect, to another, the fourfold pattern of risk preferences (da Silva Castanheira; Nobandegani, & Otto, 2019). Importantly, SbEU is the first, and thus far the only, resource-rational process model that bridges between risky, value-based, and game-theoretic decision-making.

3 Modeling Ultimatum Game

Having presented the computational underpinnings of SbEU, we now turn to modeling the effects of expectation, competition, and time pressure on Responder behavior.

Importantly, 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; Nobandegani et al., 2019a, Nobandegani & Shultz, 2020a). Consistent with this finding, throughout this paper we assume that Responder draws very few samples ($s = 1$; see Eqs. (2-3)) when deciding if they should accept or reject the offer — except for Sec. 3.3 in which we directly investigate the effect of number of samples s on Responder’s behavior.

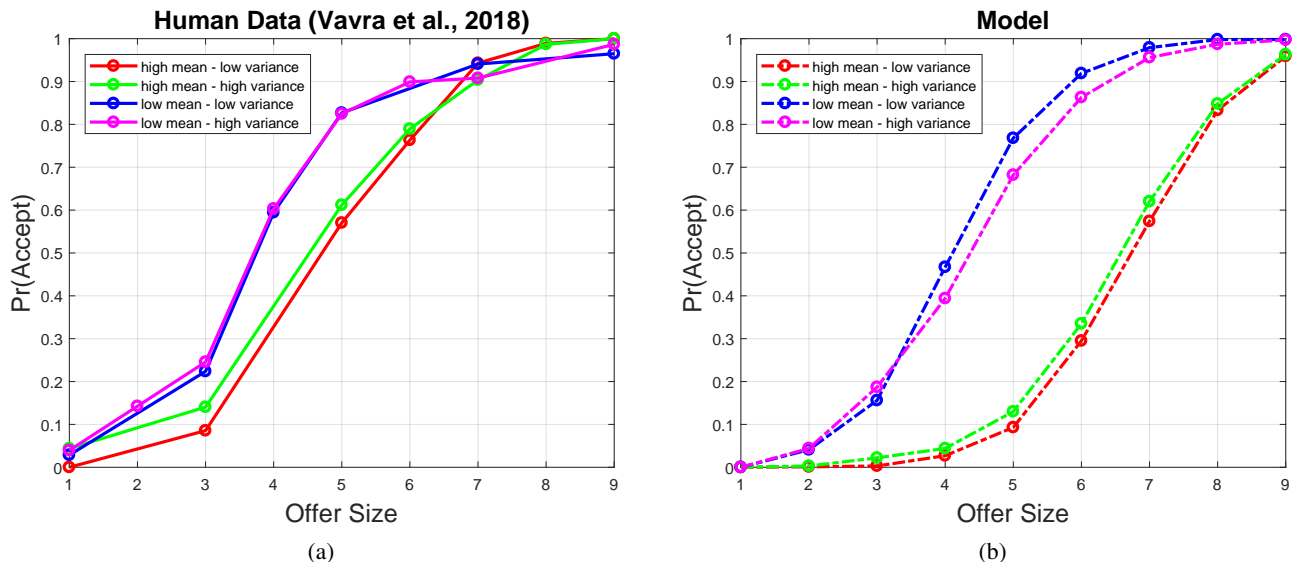


Figure 1: **Simulating the effect of manipulating Responder’s offer expectation.** Horizontal and vertical axes correspond to the amount offered by Proposer and the probability of Responder accepting that offer, respectively. (a) Experimental data of Vavra et al. (2018) for the four treatments: High-Mean Low-Variance (red curve), High-Mean High-Variance (green curve), Low-Mean Low-Variance (blue curve), and Low-Mean High-Variance (magenta curve). (b) Simulation results. Our expectation-based, resource-rational process model accurately captures the qualitative trend of Vavra et al.’s (2018) data (High-Mean Low-Variance: Spearman $\rho = 1$, High-Mean High-Variance: Spearman $\rho = 1$, Low-Mean Low-Variance: Spearman $\rho = 1$, The Low-Mean High-Variance: Spearman $\rho = 1$, $ps < .005$). We simulated $N = 1000$ Responders, with $s = 1$.

3.1 Expectation Manipulation

In a recent experimental study, Vavra et al. (2018) investigated the effect of manipulating Responders’ expectation in UG. Vavra et al. (2018) altered Responders’ offer expectations by showing them histograms of previous offers made by Proposer. Using a two-by-two design, Vavra et al. (2018) independently manipulated both mean (low vs. high) and variance (low vs. high) of these histograms. In the low-mean and high-mean conditions, participants were explicitly informed that the average offer had been around 4 and 7 Euros, respectively. Participants were not explicitly informed about the variance of the past offers and they would have to visually estimate it from the histograms pictorially presented to them, suggesting substantial differences across the variances used by individual Responders to form their expectations. The experimental data of Vavra et al. (2018) are shown in Fig. 1(a).

Next, we simulate the four experimental conditions of Vavra et al. (2018); see Fig. 1(b). Because the histograms of the past offers used by Vavra et al. (2018) were approximately normally distributed, we assume that Responders’ expectation is normally distributed with mean 4, for the low-mean conditions, and mean 7, for the high-mean conditions. The variance used to simulate the high-variance condition was chosen to obtain a relatively well-fitting curve in the low-mean condition, and subsequently kept for the high-mean condition. The variance used to simulate the low-variance condition was also chosen to obtain a relatively well-fitting curve in the low-mean condition, and subsequently

kept for the high-mean condition. Importantly, the qualitative trends accurately captured by simulation results are robust across a wide range of variance parameterizations. Our simulation of the four experimental conditions of Vavra et al. (2018) is presented in Fig. 1(b). As Fig. 1(b) shows, our expectation-based, resource-rational process model can qualitatively account for the experimental data of Vavra et al. (2018) (High-Mean Low-Variance condition: Spearman $\rho = 1$, High-Mean High-Variance condition: Spearman $\rho = 1$, Low-Mean Low-Variance condition: Spearman $\rho = 1$, Low-Mean High-Variance condition: Spearman $\rho = 1$, $ps < .005$). We have simulated $N = 1000$ Responders, with $s = 1$.

3.2 The Effect of Competition Among Responders

Fischbacher et al. (2003) investigated the effect of increased competition among Responders in a variant of UG. In this version, Proposer makes an offer. Then all Responders simultaneously accept or reject this offer. If all of the Responders reject the offer, nobody gets anything. If only a single Responder accepts the offer, Proposer and the single Responder each take their respective share (all other Responders earn nothing). If several Responders accept the offer, a single Responder is randomly selected to be the one who trades. Fischbacher et al. (2003) tested participants in three conditions: 1-Responder (standard UG), 2-Responders (one Proposer and two Responders), and 5-Responders (one Proposer and five Responders). The experimental data of Fischbacher et al. (2003) are shown in Fig. 2(a).

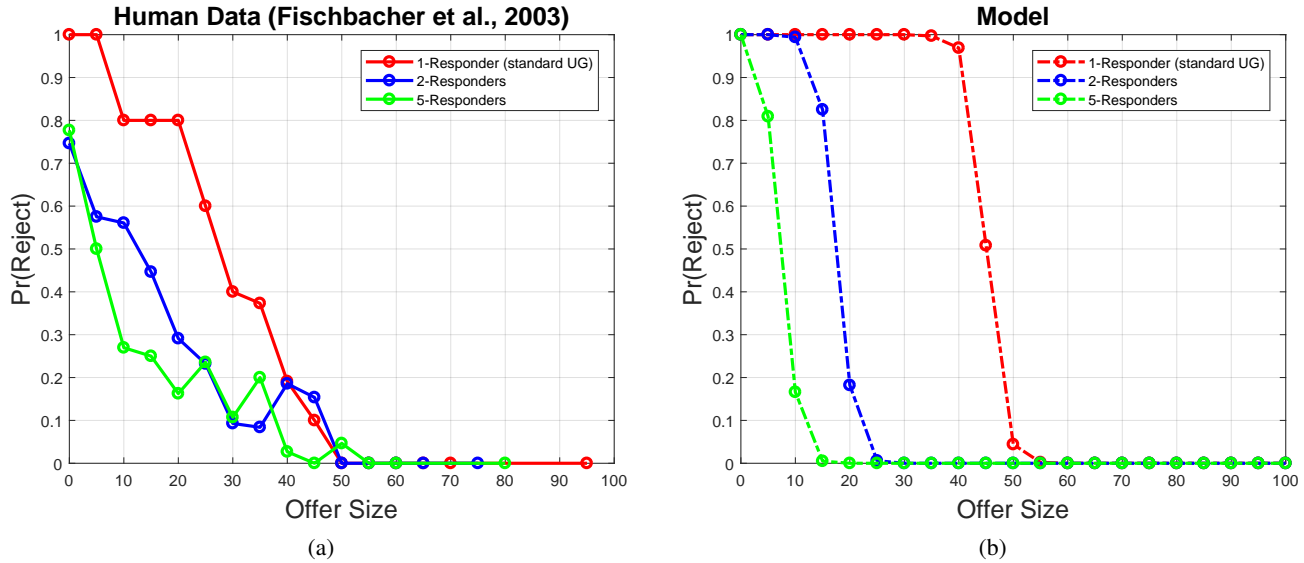


Figure 2: **Simulating the effect of introducing competition among Responders.** Horizontal and vertical axes correspond to the amount offered by Proposer and the probability of Responder rejecting that offer, respectively. (a) Experimental data of Fischbacher et al. (2003) for the three treatments: 1-Responder (red curve), 2-Responders (blue curve), and 5-Responders (green curve). (b) Simulation results. Our expectation-based, resource-rational process model captures the qualitative trend of Fischbacher et al.’s (2003) data (1-Responder condition: Spearman $\rho = .9396$, 2-Responders condition: Spearman $\rho = .9016$, 5-Responders condition: Spearman $\rho = .8073$, $ps < .001$). We simulated $N = 1000$ Responders, with $s = 1$.

Fischbacher et al. (2003) showed that the probability of accepting low offers increased as the number of Responders competing for the offer increased.

Next, we simulate the three experimental conditions of Fischbacher et al. (2003); see Fig. 2(b). In their paper, Fischbacher et al. (2003) supplied the observed distribution of offers made by Proposer in each condition; they were approximately normally distributed. Since we are interested in providing a rational, expectation-based basis for Responder’s behavior in UG, in our simulation of each condition, we assume that Responder’s offer expectation closely approximates the distribution of offers made by Proposer in that condition. (This is clearly a cognitively demanding assumption on the part of Responder, and the descriptive power of our model markedly improves if we relax this assumption by positing that Responder’s offer expectation only roughly approximates the distribution of offers made by Proposer in a condition.) Our simulation of the three experimental conditions of Fischbacher et al. (2003) is presented in Fig. 2(b). As Fig. 2(b) shows, our model can qualitatively account for the experimental data of Fischbacher et al. (2003) (1-Responder condition: Spearman $\rho = .9396$, 2-Responders condition: Spearman $\rho = .9016$, 5-Responders condition: Spearman $\rho = .8073$, $ps < .001$). We have simulated $N = 1000$ Responders, with $s = 1$.

3.3 The Effect of Time Pressure

Perhaps the most puzzling finding on UG is the effect of time pressure on Responder’s behavior. Several experimental stud-

ies have shown that, with increased deliberation, the threshold above which all offers are accepted decreases (e.g., Sutter et al., 2003; Cappelletti et al., 2008), revealing that deliberation brings Responder’s behavior closer to the predictions of the normative standards of game theory (i.e., that Responder accepts any positive offer, however minuscule). For example, Sutter et al. (2003) compares the acceptance rates of responders under time pressure with those of Responders under no time pressure (10 s and 100 s to make the accept/reject decision, respectively). The total sum to be split between Proposer and Responder was 10 Euros. While all offers above 45% of the total sum were accepted by Responders under time pressure, this threshold decreases to 15% for Responders under no time pressure.

Crucially, this finding serves as an ideal test-bed for a resource-rational model of UG. If Responders are rationally using their limited cognitive resources, as resource-rational models posit, the effect of having more time to decide should be explainable by assuming that Responder is deploying more resources (i.e., more samples in our case) when deciding if they should accept or reject the offer.

The role of cognitive limitations in UG is further corroborated by studies establishing a causal link between cognitive load and Responder behavior (e.g., De Neys et al., 2011). For example, De Neys et al. (2011) reported a general increase in rejection rates when Responders’ cognitive load is increased by a simultaneous memory task. De Neys et al. (2011) also showed that Responders with low acceptance thresholds (hence more “rational”) performed better on

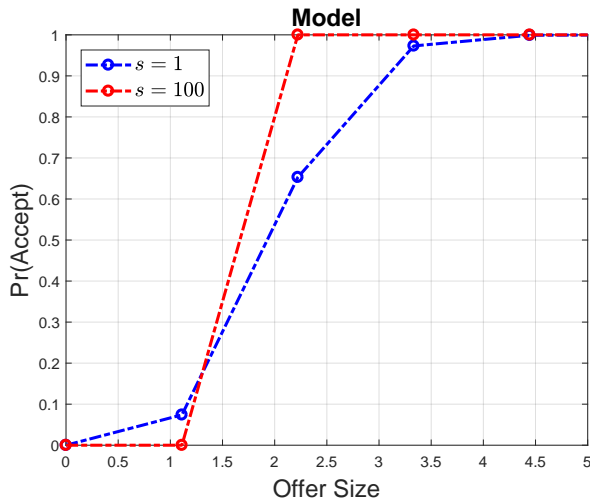


Figure 3: **Investigating the effect of manipulating the number of samples s .** Horizontal and vertical axes correspond to the amount offered by Proposer and the probability of Responder accepting that offer, respectively. Increased deliberation (operationalized by drawing a larger number of samples) brings Responder’s behavior closer to the predictions of the normative standards of game theory (i.e., that Responder accepts any positive offer, however minuscule). For $s = 1$ (blue curve), the model predicts that all offers above 45% of the total sum are accepted by Responder. However, for $s = 100$ (red curve), corresponding to increased deliberation time, the model predicts that all offers above 22.22% of the total sum are accepted by Responder.

a cognitive control test than Responders with high acceptance thresholds.

Next, we simulate the two experimental conditions of Sutter et al. (2003); see Fig. 3. Consistent with our past simulations (see Secs. 3.1 and 3.2), here again we assume that Responder’s offer expectation is normally distributed.¹ It is important to note that the qualitative trend of our simulation results is robust across a wide range of parameterization of this normally distributed expectation. As Fig. 3 shows, our resource-rational process model can qualitatively account for the experimental data of Sutter et al. (2003).

4 General Discussion

Widely considered as one of the pillars of human morality, fairness is chiefly investigated in the context of the Ultimatum Game (UG, Güth et al., 1982), an extensively studied game in psychology (e.g., Sanfey, 2009), neuroscience (e.g., Sanfey et al., 2003), and behavioral economics (e.g., Camerer

¹Sutter et al. (2003) did not inform Responders about the distribution of offers made by Proposer. In their experimental setup, and unbeknown to Responders, offers were selected between 0 and 5 Euros, uniformly at random. But this offer distribution radically differs from the distribution of offers typically made by a human Proposer, which is approximately normally distributed (see Fischbacher et al., 2003; Camerer & Fehr, 2006).

& Thaler, 1995; Fehr & Schmidt, 1999). In this work, we present the first resource rational process model of UG, providing a unified account of several disparate experimental findings, namely, the effects of expectation, competition, and time pressure on UG Responder.

By assuming that Responders use their expectation as a reference point for subjective valuation of an offer, a hypothesis well supported by a series of recent behavioral and neuroimaging studies (e.g., Sanfey, 2009; Battigalli et al., 2015; Xiang et al., 2013; Chang & Sanfey, 2013; Vavra et al., 2018), our work counterintuitively shows that the rejection of low offers in UG can arise from purely self-interested expected-utility maximization.

A recent line of experimental work has shown that intuition favors cooperation, fairness, and prosociality while deliberation promotes selfishness (e.g., Rubinstein, 2007; Rand, Greene, & Nowak, 2012; Rand, 2016). The current widely-accepted explanation for this behavioral shift is dual process theory (Evans, 2003). Our work offers a completely new way of understanding this experimental finding — both qualitatively and quantitatively.

On the quantitative front, in sharp contrast to a dual-process perspective, our work presents the first, and thus far the only, *single-process* model of Responder’s behavior in UG, providing a resource-rational mechanistic explanation of why deliberation makes Responders act more selfishly by lowering the threshold above which offers are invariably accepted. According to our work, it is the optimal use of limited cognitive resources that underlies deliberation promoting selfishness in UG Responder.

On the qualitative side, our work offers a radically different interpretation of Responder’s behavior in UG than the one provided by the classical dual-process account. From a dual-process perspective, intuition (moderated by System 1) is good and pro-fairness while deliberation (moderated by System 2) is evil and anti-fairness. However, according to our single-process model, a boundedly-rational agent that simultaneously (1) adopts her expectation as a reference point for subjective valuation of an offer, and (2) selfishly maximizes expected utility while optimally using her limited cognitive resources, should lower the threshold above which all offers are invariably accepted with increased deliberation. As such, Responder’s intuitive response being pro-fairness, is still, quite counter-intuitively, the effect of selfishly maximizing expected utility while optimally using limited cognitive resources.

Accordingly, our work contributes to an emerging line of work explaining human cognition as an optimal use of limited cognitive resources (*rational minimalist program*, Nobandegani, 2017; Griffiths, Lieder, & Goodman, 2015). Our work suggests that this emerging perspective has great potential to shed new light on the computational foundation of moral decision-making.

Acknowledgments This work is supported by an operating grant to TRS from the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

- 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.
- Battigalli, P., Dufwenberg, M., & Smith, A. (2015). Frustration and anger in games.
- Bolton, G. E., & Zwick, R. (1995). Anonymity versus punishment in ultimatum bargaining. *Games and Economic Behavior*, *10*(1), 95–121.
- 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.
- Camerer, C. F., & Fehr, E. (2006). When does “economic man” dominate social behavior? *Science*, *311*(5757), 47–52.
- Camerer, C. F., & Thaler, R. H. (1995). Anomalies: Ultimatums, dictators and manners. *Journal of Economic Perspectives*, *9*(2), 209–219.
- Cappelletti, D., Güth, W., & Ploner, M. (2008). *Being of two minds: An Ultimatum experiment investigating affective processes* (Tech. Rep.). Jena Economic Research Papers.
- Chang, L. J., & Sanfey, A. G. (2013). Great expectations: neural computations underlying the use of social norms in decision-making. *Social Cognitive and Affective Neuroscience*, *8*(3), 277–284.
- 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.
- da Silva Castanheira, K., Nobandegani, A. S., Shultz, T. R., & 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.
- De Neys, W., Novitskiy, N., Geeraerts, L., Ramautar, J., & Wagemans, J. (2011). Cognitive control and individual differences in economic ultimatum decision-making. *PLoS One*, *6*(11).
- Evans, J. S. B. (2003). In two minds: dual-process accounts of reasoning. *Trends in Cognitive Sciences*, *7*(10), 454–459.
- Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The quarterly journal of economics*, *114*(3), 817–868.
- Fischbacher, U., Fong, C. M., & Fehr, E. (2003). Fairness and the power of competition. *Working paper/Institute for Empirical Research in Economics*, 133.
- 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., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science*, *7*(2), 217–229.
- 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.
- Guala, F. (2008). Paradigmatic experiments: The ultimatum game from testing to measurement device. *Philosophy of Science*, *75*(5), 658–669.
- Güth, W., Schmittberger, R., & Schwarze, B. (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization*, *3*(4), 367–388.
- Güth, W., & Tietz, R. (1990). Ultimatum bargaining behavior: A survey and comparison of experimental results. *Journal of Economic Psychology*, *11*(3), 417–449.
- Hammersley, J., & Handscomb, D. (1964). *Monte Carlo Methods*. London: Methuen & Co Ltd.
- Hertwig, R., & Pleskac, T. J. (2010). Decisions from experience: Why small samples? *Cognition*, *115*(2), 225–237.
- 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*.
- Markowitz, H. (1952). The utility of wealth. *Journal of Political Economy*, *60*(2), 151–158.
- 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. (2019b). 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., Shultz, T. R., & Otto, A. R. (2019a). 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., & 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., & Shultz, T. R. (2020a). A resource-rational mechanistic account of human coordination strategies. In *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Nobandegani, A. S., & Shultz, T. R. (2020b). A resource-rational, process-level account of the St. Petersburg paradox. *Topics in Cognitive Science*, 12(1), 417–432.
- Nobandegani, A. S., & Shultz, T. R. (2020c). The St. Petersburg paradox: A fresh algorithmic perspective. In *Proc. of the 34th Conference on Artificial Intelligence (AAAI)*.
- Nowak, M. A., Page, K. M., & Sigmund, K. (2000). Fairness versus reason in the ultimatum game. *Science*, 289(5485), 1773–1775.
- Poor, H. V. (2013). *An Introduction to Signal Detection and Estimation*. Springer Science & Business Media.
- Rand, D. G. (2016). Cooperation, fast and slow: Meta-analytic evidence for a theory of social heuristics and self-interested deliberation. *Psychological Science*, 27(9), 1192–1206.
- Rand, D. G., Greene, J. D., & Nowak, M. A. (2012). Spontaneous giving and calculated greed. *Nature*, 489(7416), 427.
- Rawls, J. (1971). *A Theory of Justice*. Harvard University Press.
- Rubinstein, A. (2007). Instinctive and cognitive reasoning: A study of response times. *The Economic Journal*, 117(523), 1243–1259.
- Sanfey, A. G. (2009). Expectations and social decision-making: Biasing effects of prior knowledge on ultimatum responses. *Mind & Society*, 8(1), 93–107.
- Sanfey, A. G., Rilling, J. K., Aronson, J. A., Nystrom, L. E., & Cohen, J. D. (2003). The neural basis of economic decision-making in the ultimatum game. *Science*, 300(5626), 1755–1758.
- Scholten, M., & Read, D. (2014). Prospect theory and the forgotten fourfold pattern of risk preferences. *Journal of Risk and Uncertainty*, 48(1), 67–83.
- Simon, H. A. (1957). *Models of Man*. Wiley.
- Sutter, M., Kocher, M., & Strauß, S. (2003). Bargaining under time pressure in an experimental ultimatum game. *Economics Letters*, 81(3), 341–347.
- Thaler, R. H. (1988). Anomalies: The ultimatum game. *Journal of Economic Perspectives*, 2(4), 195–206.
- Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2), 207–232.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4), 297–323.
- Van Damme, E., Binmore, K. G., Roth, A. E., Samuelson, L., Winter, E., Bolton, G. E., ... others (2014). How werner güth's ultimatum game shaped our understanding of social behavior. *Journal of Economic Behavior & Organization*, 108, 292–318.
- Vavra, P., Chang, L. J., & Sanfey, A. G. (2018). Expectations in the ultimatum game: distinct effects of mean and variance of expected offers. *Frontiers in Psychology*, 9, 992.
- 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.
- Xiang, T., Lohrenz, T., & Montague, P. R. (2013). Computational substrates of norms and their violations during social exchange. *Journal of Neuroscience*, 33(3), 1099–1108.