

A Resource-Rational Mechanistic Account of Human Coordination Strategies

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

{ardavan.salehinobandegani}@mail.mcgill.ca

{thomas.shultz}@mcgill.ca

¹Department of Electrical & Computer Engineering, McGill University

²School of Computer Science, McGill University

³Department of Psychology, McGill University

Abstract

Humans often coordinate their actions in order to reach a mutually advantageous state. These circumstances are chiefly modeled by *coordination games*, a well-known class of games extensively studied in behavioral economics. In this work, we present the first resource-rational mechanistic approach to coordination games, showing that a variant of normative expected-utility maximization acknowledging cognitive limitations can account for several major experimental findings on human coordination behavior in strategic settings. Concretely, we show that Nobandegani et al.'s (2018) rational process model, *sample-based expected utility*, provides a unified account of (1) the effect of time pressure on human coordination, and (2) how systematic variations of risk- vs. payoff-dominance affect coordination behavior. Importantly, Harsanyi and Selten's (1988) theory of equilibrium selection fails to account for (1-2). As such, our work suggests that the optimal use of limited cognitive resources may lie at the core of human coordination behavior. We conclude by discussing the implication of our work for understanding human strategic behavior, moral decision-making, and human rationality.

Keywords: behavioral game theory; one-shot non-cooperative games; coordination games; moral decision-making; resource-rational process models

1 Introduction

On which side of the road should we drive to avoid a head-on collision? How do two people decide on when and where to meet? Having similar or dissimilar preferences, how do two friends decide on a mutual activity? Humans often coordinate their actions in order to reach a mutually advantageous state (Harsanyi & Selten, 1988; Cooper, 1999; Brañas-Garza & Cabrales, 2015; Belloc et al., 2019). These circumstances are chiefly modeled by *coordination games* (CGs), a well-known class of games that are extensively studied in behavioral economics, and particularly, behavioral game theory (Cooper, 1999; Cooper et al., 1992; Cachon & Camerer, 1996; Weber et al., 2001; Camerer, 2011); see Fig. 1.

In addition to several daily life circumstances, CGs capture a broad range of important economic situations, e.g., market entry, macroeconomic policy coordination, choice of product standards, and contract agreement (Poulsen & Sonntag, 2019). Importantly, the question of how rational agents should coordinate has played a profound role in the development of economic theory (Harsanyi & Selten, 1988; Güth & Kalkofen, 1989; Güth, 1992): Given that CGs have multiple Nash equilibria, the problem of strategy selection in CGs

lies at the foundation of Harsanyi and Selten's (1988) Nobel-winning theory of equilibrium selection, which formally delineates how rational economic agents coordinate.

However, despite its normative stance, a substantial body of experimental work has shown that Harsanyi and Selten's theory fails to provide a descriptive account of human coordination (e.g., Straub, 1995; Battalio et al., 2001; Schmidt et al., 2003; Poulsen & Sonntag, 2019; Belloc et al., 2019).

But why should human coordination behavior deviate from the predictions of a rational theory of equilibrium selection? Consistent with bounded rationality (Simon, 1957), we hypothesize that these deviations can be accounted for if we acknowledge the computational and cognitive limitations that people are faced with.

Inspired by the success of Nobandegani et al.'s (2019a) resource-rational model of cooperation in one-shot Prisoner's Dilemma games, we ask whether human coordination can be seen as an optimal behavior with the mind acting as a cognitive miser. As such, we seek to provide an account of human coordination that is both descriptively adequate and normatively sound.

In this work, we present the first resource-rational account of human coordination, demonstrating that Nobandegani et al.'s (2018) rational process model, *sample-based expected utility* (SbEU), provides a unified mechanistic account of (1) the effect of time pressure on human coordination, and (2) how systematic variations of risk- vs. payoff-dominance affect coordination behavior. The two concepts of risk-dominance and payoff-dominance play a critical role in Harsanyi and Selten's (1988) theory of equilibrium selection (for their explanations, see Sec. 3.1). Crucially, Harsanyi and Selten's (1988) theory fails to account for (1-2).

Our paper is organized as follows. After presenting the computational underpinnings of our resource-rational approach, we turn to modeling human coordination behavior. We conclude by discussing the implication of our work for understanding human strategic behavior, moral decision-making, and human rationality.

2 Computational Model

In addition to SbEU, whose mathematical underpinning is discussed in detail in Sec. 2.1, our resource-rational approach to coordination adopts two key ideas. The first is Nobandegani et al.'s (2019a) general framework for conceptualizing




Figure 1: **Examples of coordination games with corresponding payoff matrices.** (a) An instance of the *pure coordination* game, wherein both players prefer the same Nash equilibrium (i.e., (Party, Party)). (b) An instance of the *battle of the sexes* game, wherein players prefer opposing Nash equilibria (the row player prefers (Party, Party) whereas the column player prefers (Home, Home)). (c) An instance of the *stag-hunt* game, wherein both players (hunters) can benefit if they cooperate (hunting a stag). However, cooperation might fail, because each hunter has an alternative which is safer because it does not require cooperation to succeed (hunting a hare).




Figure 2: **Payoff matrix for a generic 2-player, one-shot, non-cooperative game (2ONG).** For example, if Player 1 (Row Player) selects the Top Strategy and Player 2 (Column Player) selects the Left Strategy, Player 1 and Player 2 receive payoffs a and x , respectively.

any n -player, one-shot, non-cooperative game (nONG) as a set of risky gambles. Applying Nobandegani et al.’s (2019a) framework to a generic 2ONG (see Fig. 2), the problem of strategy selection for Player 1 amounts to choosing between the following two risky gambles:

$$T = \begin{cases} a & \text{w.p. } P_l \\ b & \text{w.p. } 1 - P_l \end{cases} \quad (1)$$

$$B = \begin{cases} c & \text{w.p. } P_l \\ d & \text{w.p. } 1 - P_l \end{cases} \quad (2)$$

with gambles T, B corresponding to choosing the Top Strategy and the Bottom Strategy, respectively, and P_l denoting Player 1’s conception of the probability with which Player 2 chooses Left strategy. Similarly, according to Nobandegani et al. (2019a), the problem of strategy selection for Player 2 amounts to choosing between the following two risky gambles:

$$L = \begin{cases} x & \text{w.p. } P_t \\ y & \text{w.p. } 1 - P_t \end{cases} \quad (3)$$

$$R = \begin{cases} v & \text{w.p. } P_t \\ w & \text{w.p. } 1 - P_t \end{cases} \quad (4)$$

with gambles L, R corresponding to choosing the Left Strategy and the Right Strategy, respectively, and P_t denoting Player 2’s conception of the probability with which Player 1 chooses Top strategy.

The second ingredient of our resource-rational approach is the well-supported level- k theory in behavioral economics (Stahl & Wilson, 1994; Nagel, 1995; Costa-Gomes et al., 2001) that posits that (1) players in strategic games base their decisions on their predictions about the likely actions of other players, and (2) players can be categorized by the “depth” of their strategic thought. According to level- k theory, a completely non-strategic level-0 player will choose actions without regard to the actions of other players, i.e., uniformly at random. A level- i player assumes that she is playing against a level- $(i-1)$ player and her action will be the best response consistent with this belief. For example, a level-1 player assumes that she is playing against a level-0 player and her action will be the best response (i.e., the expected-utility maximizing response) consistent with this belief.

Consistent with bounded rationality (Simon, 1957), substantial experimental evidence shows that only a small proportion of players exhibit depths of reasoning of third order or above (e.g., Nagel, 1995; Camerer et al., 2004; Coricelli & Nagel, 2009). Relatedly, recent experimental work shows that depth of reasoning increases with deliberation (e.g., Kocher & Sutter, 2006). Consistent with these findings, and for the sake of parsimony, throughout this paper we assume that, under time pressure, players are level-1 agents, and, when time pressure is not implemented, players are level-2 agents.

Despite acknowledging human bounded rationality, level- k theory still largely rests on perfect rationality as it maintains that the action of higher-level agents will be the best response, in the precise sense of expected-utility maximization, consistent with their beliefs. A satisfying resource-rational, level- k -theory-based account of strategic behavior requires relaxation of the best-response assumption, and, instead, considering a boundedly-optimal response according to which people maximize expected utility “to the best of their abilities.” That is, they optimally maximize expected utility, but this maximiza-

tion is subject to their cognitive limitations. In this work, we assume that people arrive at their boundedly-optimal response using Nobandegani et al.’s (2018) resource-rational process model of risky decision-making: *sample-based expected utility* (SbEU).

2.1 Sample-based Expected Utility Model

SbEU is a metacognitively-rational process model of risky choice 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

$$\mathbb{E}[u(o)] = \int p(o)u(o)do, \quad (5)$$

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

$$\hat{E} = \frac{1}{\sum_{j=1}^s w_j} \sum_{i=1}^s w_i u(o_i), \quad \forall i: o_i \sim q^*, w_i = \frac{p(o_i)}{q^*(o_i)}, \quad (6)$$

$$q^*(o) \propto p(o)|u(o)|\sqrt{\frac{1+|u(o)|\sqrt{s}}{|u(o)|\sqrt{s}}}. \quad (7)$$

MSE is a standard measure of estimation quality, and is commonly used in mathematical statistics (Poor, 2013). In Eqs. (5-7), 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. (5), 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”):

$$A = \begin{cases} o_A & \text{w.p. } P_A \\ 0 & \text{w.p. } 1 - P_A \end{cases} \quad (8)$$

$$B = \begin{cases} o_B & \text{w.p. } P_B \\ 0 & \text{w.p. } 1 - P_B \end{cases} \quad (9)$$

$$\Delta u(o) = \begin{cases} u(o_A) - u(o_B) & \text{w.p. } P_A P_B \\ u(o_A) - u(0) & \text{w.p. } P_A(1 - P_B) \\ u(0) - u(o_B) & \text{w.p. } (1 - P_A)P_B \\ 0 & \text{w.p. } (1 - P_A)(1 - P_B) \end{cases} \quad (10)$$

In Eq. (10), $u(\cdot)$ 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:

$$u(x) = \begin{cases} x^{0.85} & \text{if } x \geq 0, \\ -|x|^{0.95} & \text{if } x < 0. \end{cases} \quad (11)$$

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

Recently, Nobandegani et al. (2018) showed that SbEU can account for the availability bias, the tendency to overestimate the probability of events that come easily to mind (Tversky & Kahneman, 1973), and can also 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 unified, resource-rational mechanistic account of cooperation in one-shot Prisoner’s Dilemma games and inequality-aversion in the Ultimatum game, thus successfully bridging between game-theoretic decision-making and risky decision-making (Nobandegani et al., 2019a; Nobandegani, Destais, & Shultz, 2020). 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, 2020a, 2020b), 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 makes people move from one bias, the framing effect, to another bias, 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 One-shot Coordination Games

In this section, we model human coordination behavior in the context of 2-player, one-shot coordination games (2OCGs). We particularly focus on two major topics in human coordination: (1) how systematic variations of risk- vs. payoff-dominance, as two concepts playing a critical role in Harsanyi and Selten’s (1988) theory of equilibrium selection, affect coordination behavior (e.g., Straub, 1995; Battalio et al., 2001; Schmidt et al., 2003) and (2) how time pressure affects human coordination (e.g., Poulsen & Sonntag, 2019; Belloc et al., 2019).

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., 2020). Consistent with this finding, throughout this paper we assume that a player draws very few samples ($s = 1$; see Eqs. (6-7))




Figure 3: Modeling how systematic variations of risk- vs. payoff-dominance affect coordination behavior. Horizontal and vertical axes correspond, respectively, to human participants’ probability of choosing a strategy when playing coordination games, and SbEU model prediction for the probability of choosing that strategy. **(a)** Simulating Games 1-4 and Game 7 from Straub (1995) and the original game from Cooper et al. (1989). **(b)** Simulating the seven coordination games collectively studied in Battalio et al. (2001) and Schmidt et al. (2003). **(c)** Simulating aggregate experimental data. SbEU can account for the collection of experimental data simulated separately in **(a)** and **(b)** (Pearson $r = .8846$, $p < 10^{-4}$); two datapoints overlap.

when choosing their strategy.

3.1 Systematic Variations of Risk- vs. Payoff-Dominance

In their seminal work, Harsanyi and Selten (1988) formally presented two main concepts for understanding equilibrium selection and coordination: *risk-dominance* and *payoff-dominance*. Broadly, risk-dominance recommends the strategy corresponding to the Nash equilibrium that is the least risky of all equilibria, and hence the safest to play. (For a formal characterization of risk-dominance, see Harsanyi and Selten, 1988.) On the other hand, payoff-dominance recommends the strategy corresponding to the Nash equilibrium yielding the highest payoff for all players (aka Pareto-optimal equilibrium).

Several experimental studies have systematically investigated the role of payoff- and risk-dominance in human coordination behavior (e.g., Straub, 1995; Battalio et al., 2001; Schmidt et al., 2003). Next, we show that SbEU can account for the experimental data of several of these studies, namely, Straub (1995), Battalio et al. (2001), and Schmidt et al. (2003).

3.1.1 Straub (1995) In a series of coordination games, which include a replication attempt of a coordination game studied by Cooper et al. (1989), Straub (1995) showed that people predominantly rely on risk-dominance (as opposed to payoff-dominance) when coordinating in strategic settings. This finding undermines Harsanyi and Selten’s (1988) theory of equilibrium selection. We next show that, SbEU, as a resource-rational process model, can account for the experimental data of Straub (1995).

For ease of analysis, in our simulation we focus on the symmetric coordination games studied in Straub (1995), i.e., Games 1-4, and Game 7; we also include in our simula-

tion the original experiment of Cooper et al. (1989).¹ As shown in Fig. 3(a), SbEU can account for these experimental data (Pearson $r = .8489$, $p = .0325$). We have simulated $N = 100,000$ participants, with $s = 1$.

3.1.2 Battalio et al. (2001) and Schmidt et al. (2003) To experimentally test Harsanyi and Selten’s (1988) theory of equilibrium selection, Battalio et al. (2001) and Schmidt et al. (2003) systematically investigated the role of risk- and payoff-dominance (and possible interaction between the two) in human coordination. As Battalio et al. (2001) and Schmidt et al. (2003) show, Harsanyi and Selten’s (1988) theory provides a poor account of their experimental data.

Battalio et al. (2001) tested participants in three coordination games; Schmidt et al. (2003) tested participants in four coordination games. Using SbEU, we next simulate the experimental data of Battalio et al. (2001) and Schmidt et al. (2003). As shown in Fig. 3(b), SbEU can accurately account for these experimental data (Pearson $r = .9320$, $p = .0022$). We simulated $N = 100,000$ participants, with $s = 1$.

Collecting the experimental data simulated separately in Sec. 3.1.1 and Sec. 3.1.2, Fig. 3(c) shows that SbEU can accurately account for the aggregate experimental data (Pearson $r = .8846$, $p < 10^{-4}$). We have simulated $N = 100,000$ participants, with $s = 1$.

3.2 The Effect of Time Pressure on Coordination

The effect of time pressure on human coordination behavior is largely understudied, but a series of recent studies are beginning to shed light on it (Poulsen & Sonntag, 2019; Belloc et al., 2019).

Next, we show that SbEU (together with level- k hypoth-

¹Game 8 from Straub (1995) is not included in our simulation, as experimentally observed behavior is inconsistent with the findings of Schmidt et al. (2003). We simulate the experimental data of Schmidt et al. (2003) in Sec. 3.1.2.

esis) can account for the main experimental findings of Poulsen and Sonntag (2019) and Belloc et al. (2019) on the effect of time pressure on coordination, thereby providing the first resource-rational account of these findings.

We assume that, under time pressure, players are level-1 agents, and, when time pressure is not implemented, players are level-2 agents (see Sec. 2 for supporting evidence).

It is worth noting that, being purely static (as opposed to dynamic), Harsanyi and Selten's (1988) theory fails to account for how coordination behavior changes over time.

Poulsen and Sonntag (2019) In a recent study, Poulsen and Sonntag (2019) experimentally investigated the effect of time pressure on both symmetric pure coordination games and the Battle of the Sexes games (which capture coordination in the presence of conflict of interest). In the case of symmetric pure coordination games, Poulsen and Sonntag (2019) showed that time pressure has no effect on people's strategic behavior. However, in the case of the Battle of the Sexes, Poulsen and Sonntag (2019) showed that deliberation makes a player move from choosing the strategy that corresponds to the equilibrium state yielding the highest self-payoff, to the strategy that corresponds to the equilibrium state yielding the highest payoff for the other player. This is moving from self-regarding behavior to other-regarding behavior.

SbEU, together with level- k theory, provides a mechanistic account of this experimental finding of Poulsen and Sonntag (2019). Simulating the symmetric pure coordination game from Poulsen and Sonntag (2019), our resource-rational model predicts that time pressure has no effect on people's strategic behavior ($\chi^2_{(1)} = .0168, p = .8968$). However, simulating the Battle of the Sexes game from Poulsen and Sonntag (2019), our resource-rational model predicts that deliberation makes a player move from self-regarding behavior to other-regarding behavior ($\chi^2_{(1)} = 3754.9, p < 10^{-3}$). We have simulated $N = 100,000$ participants, with $s = 1$.

Belloc et al. (2019) Relatedly, Belloc et al. (2019) systematically investigated the effect of time pressure on Stag-Hunt games, a well-known class of coordination games which model choice between a safe (aka Hare) vs. risky (aka Stag) strategy; see Fig. 1(c).

Belloc et al. (2019) tested participants in four Stag-Hunt games, under two experimental conditions. In the Time Pressure condition, subjects were given 10 s to choose their strategy in each game. In the Control condition, no time pressure was implemented. Belloc et al. (2019) showed that people's tendency to choose the safe strategy increased with deliberation.

SbEU, together with level- k theory, provides a process-level account of the experimental finding of Belloc et al. (2019). Simulating the four Stag-Haunt games from Belloc et al. (2019), our resource-rational model predicts that the tendency to choose the safe strategy increases with deliberation (Game 1: $\chi^2_{(1)} = 7263.2$, Game 2: $\chi^2_{(1)} = 6882.3$, Game 3: $\chi^2_{(1)} = 3379.6$, Game 4: $\chi^2_{(1)} = 8806.8, ps < 10^{-4}$). We

have simulated $N = 100,000$ participants, with $s = 1$.

4 General Discussion

Humans often opt to coordinate their actions in order to reach a mutually advantageous state (e.g., Cooper, 1999; Brañas-Garza & Cabrales, 2015; Belloc et al., 2019). In this work, we presented the first resource-rational process model of human coordination. As we demonstrated, a single-parameterization of our resource rational model provides a unified account of (1) the effect of time pressure on human coordination, and (2) how systematic variations of risk- vs. payoff-dominance affect coordination. Crucially, Harsanyi and Selten's (1988) Nobel-winning theory of equilibrium selection fails to account for (1-2).

Investigating the effect of time pressure on human coordination in the Stag-Hunt game, Belloc et al. (2019) showed that people's tendency to choose the safe strategy markedly increases with deliberation. Because Stag can be interpreted as a more collaborative and trusting strategy than Hare, Belloc et al. (2019) interpreted their finding as evidence for a dual-process model of human strategic decision-making, with intuition promoting prosociality while increased deliberation discourages prosociality. Our work provides a completely new interpretation of Belloc et al.'s (2019) experimental finding. In sharp contrast to a dual-process perspective, our work presents the first, and thus far the only, *single-process* model of the effect of time pressure on human coordination, providing a more parsimonious account. According to our resource-rational account, it is the optimal use of limited cognitive resources that underlies deliberation diminishing prosociality. More specifically, according to our single-process model, a boundedly-rational agent who selfishly maximizes their expected utility while optimally using their limited cognitive resources should show prosociality as an intuitive response (by choosing Stag more often) and increasingly move away from prosociality by choosing Hare. Thus, according to our work, intuitive prosociality is the effect of selfishly maximizing expected utility while optimally using limited cognitive resources.

Importantly, our work is fully consistent with the recent work showing that humans' intuitive response being to cooperate in one-shot Prisoner's Dilemma games is also the effect of selfishly maximizing expected utility while optimally using limited cognitive resources (Nobandegani et al., 2019a). Relatedly, recent modeling work on the Ultimatum game has shown that Responder's intuitive response being pro-fairness can be also accounted for by selfishly maximizing expected utility while optimally using limited cognitive resources (Nobandegani et al., 2020).

Our work suggests that the optimal use of limited cognitive resources may lie at the core of human coordination behavior. In that light, our work is fully in line with the recent work by Nobandegani et al. (2019a) showing that the optimal use of limited cognitive resources can explain ostensibly irrational cooperation in one-shot Prisoner's Dilemma games.

As such, 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). This line of work demonstrates that a wide range of human behaviors are rational, provided that the computational and cognitive limitations of the mind are taken into consideration (Simon, 1957).

In 1994, Nash, Harsanyi and Selten jointly received the Nobel prize in economics; Nash for formally developing the concept of Nash equilibrium, and, Harsanyi and Selten for developing a rational theory of equilibrium selection. A substantial body of experimental work has shown that both Nash equilibrium and Harsanyi and Selten's (1988) theory of equilibrium selection fail to provide a descriptive account of human strategic behavior (e.g., Fehr & Gächter, 2000; Keser & van Winden, 2000; Brandts & Schram, 2001; Straub, 1995; Battalio et al., 2001; Schmidt et al., 2003; Belloc et al., 2019). The work presented here, together with Nobandegani et al.'s (2019a) resource-rational account of cooperation in one-shot Prisoner's Dilemma, begins to provide a completely fresh perspective on the Nash equilibrium and equilibrium selection — two foundational topics in behavioral economics. Taken together, it suggests that understanding human strategic behavior in terms of optimal use of limited cognitive resources is not only a promising approach to account for experimentally-documented deviations from Nash equilibrium (cooperation in one-shot Prisoner's Dilemma being a prominent example), but can also offer a rational process level explanation of strategy selection in the presence of multiple equilibria. As such, our work suggests that understanding human strategic behavior in terms of optimal use of limited cognitive resources (Nobandegani, 2017; Griffiths, Lieder, & Goodman, 2015) is a fresh, promising approach to behavioral economics.

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.
- Battalio, R., Samuelson, L., & Van Huyck, J. (2001). Optimization incentives and coordination failure in laboratory stag hunt games. *Econometrica*, *69*(3), 749–764.
- Belloc, M., Bilancini, E., Boncinelli, L., & DAlessandro, S. (2019). intuition and deliberation in the stag hunt game. *Scientific Reports*, *9*(1), 1–7.
- 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.
- Brañas-Garza, P., & Cabrales, A. (2015). *Experimental economics: Volume i: Economic decisions*. Springer.
- Brandts, J., & Schram, A. (2001). Cooperation and noise in public goods experiments: applying the contribution function approach. *Journal of Public Economics*, *79*(2), 399–427.
- Cachon, G. P., & Camerer, C. F. (1996). Loss-avoidance and forward induction in experimental coordination games. *The Quarterly Journal of Economics*, *111*(1), 165–194.
- Camerer, C. F. (2011). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.
- Camerer, C. F., Ho, T.-H., & Chong, J.-K. (2004). A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, *119*(3), 861–898.
- Cooper, R. (1999). *Coordination games*. Cambridge University Press.
- Cooper, R., DeJong, D. V., Forsythe, R., & Ross, T. W. (1989). Communication in the battle of the sexes game: some experimental results. *The RAND Journal of Economics*, 568–587.
- Cooper, R., DeJong, D. V., Forsythe, R., & Ross, T. W. (1992). Communication in coordination games. *The Quarterly Journal of Economics*, *107*(2), 739–771.
- Coricelli, G., & Nagel, R. (2009). Neural correlates of depth of strategic reasoning in medial prefrontal cortex. *Proceedings of the National Academy of Sciences*, *106*(23), 9163–9168.
- Costa-Gomes, M., Crawford, V. P., & Broseta, B. (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica*, *69*(5), 1193–1235.
- 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.
- Fehr, E., & Gächter, S. (2000). Cooperation and punishment in public goods experiments. *American Economic Review*, *90*(4), 980–994.
- 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.
- Güth, W. (1992). Equilibrium selection by unilateral deviation stability. In *Rational interaction* (pp. 161–189). Springer.
- Güth, W., & Kalkofen, B. (1989). *Unique solutions for strategic games: equilibrium selection based on resistance avoidance* (Vol. 328). Springer Science & Business Media.
- Hammersley, J., & Handscomb, D. (1964). *Monte Carlo Methods*. London: Methuen & Co Ltd.
- Harsanyi, J. C., & Selten, R. (1988). *A General Theory of Equilibrium Selection in Games*. MIT Press.
- Hertwig, R., & Pleskac, T. J. (2010). Decisions from experience: Why small samples? *Cognition*, 115(2), 225–237.
- Kocher, M. G., & Sutter, M. (2006). Time is moneytime pressure, incentives, and the quality of decision-making. *Journal of Economic Behavior & Organization*, 61(3), 375–392.
- 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.
- Nagel, R. (1995). Unraveling in guessing games: An experimental study. *The American Economic Review*, 85(5), 1313–1326.
- 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 *Proc. of the 17th International Conference on Cognitive Modeling*.
- 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).
- 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., Destais, C., & Shultz, T. R. (2020). A resource-rational process model of fairness in the Ultimatum game. In *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Nobandegani, A. S., & Shultz, T. R. (2020a). 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. (2020b). The St. Petersburg paradox: A fresh algorithmic perspective. In *Proc. of the 34th Conference on Artificial Intelligence (AAAI)*.
- Poor, H. V. (2013). *An Introduction to Signal Detection and Estimation*. Springer Science & Business Media.
- Poulsen, A., & Sonntag, A. (2019). *Focality is intuitive-experimental evidence on the effects of time pressure in coordination games* (Tech. Rep.). School of Economics, University of East Anglia, Norwich, UK.
- Schmidt, D., Shupp, R., Walker, J. M., & Ostrom, E. (2003). Playing safe in coordination games: the roles of risk dominance, payoff dominance, and history of play. *Games and Economic Behavior*, 42(2), 281–299.
- 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.
- Stahl, D. O., & Wilson, P. W. (1994). Experimental evidence on players' models of other players. *Journal of Economic Behavior & Organization*, 25(3), 309–327.
- Straub, P. G. (1995). Risk dominance and coordination failures in static games. *Quarterly Review of Economics and Finance*, 35(4), 339–364.
- 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.
- 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.
- Weber, R., Camerer, C., Rottenstreich, Y., & Knez, M. (2001). The illusion of leadership: Misattribution of cause in coordination games. *Organization Science*, 12(5), 582–598.