

Over-representation of Extreme Events in Decision-Making: A Rational Metacognitive Account

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

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

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

¹Department of Electrical & Computer Engineering, McGill University

²School of Computer Science, McGill University

³Department of Psychology, McGill University

Abstract

The Availability bias, manifested in the over-representation of extreme eventualities, is a well-known cognitive bias, and is generally taken as evidence of human irrationality. In this work, we present the first rational, metacognitive account of the Availability bias, formally articulated at Marr’s algorithmic level of analysis. Concretely, we present a normative, metacognitive model of how a cognitive system should over-represent extreme eventualities, depending on the amount of time available for decision-making. Our Sample-based Expected Utility model also accounts for two well-known framing effects in human decision-making under risk—the fourfold pattern of risk preferences in outcome probability (Tversky & Kahneman, 1992) and in outcome magnitude (Markovitz, 1952)—thereby providing the first metacognitively-rational basis for the aforementioned effects. Empirical evidence confirms an important prediction of our model. Surprisingly, our model is strikingly robust with respect to its focal parameter. We discuss the implications of our work for studies on human decision-making, and conclude by presenting a counter-intuitive prediction of our model, which, if confirmed, would have intriguing implications for human decision-making under risk. To our knowledge, our model is the first metacognitive, resource-rational process model of cognitive biases in decision-making. Notably, our work also contributes to the fields of artificial intelligence and computational statistics, by presenting a previously unknown proposal distribution, with firm rational grounds, broadly applicable to the influential sub-field of importance sampling Monte Carlo methods.

Keywords: Availability bias; Decision-making under uncertainty and risk; Metacognitively rational models; Fourfold pattern of risk preferences

1 Introduction

Which one comes to your mind more easily? The most horrible car crash of your life, or the event of driving home safely on any given day? Among the great many cognitive biases documented in the literature, the Availability bias (Tversky & Kahneman, 1972) is a notable one: People overestimate the probability of events that easily come to mind. A number of notable effects can be explained by this cognitive bias: people’s overestimation of the frequency of extreme events like an earthquake (Lichtenstein, Slovic, Fischhoff, Layman, & Combs, 1978) and people’s overreaction to threats like terrorism (Lichtenstein et al., 1978; Rothman, Klein, & Weinstein, 1996, Sunstein & Zeckhauser, 2011). Neurobiological work shows that the strength of a memory is modulated by the salience of its positive or negative valance (Cruciani et al., 2011), thereby providing a possible explanation of the bias.

Recently, Lieder, Griffiths, and Hsu (2014, 2017) proposed a boundedly-optimal, rational process model of the Availability bias which can explain a wide range of findings in the

human-decision making literature. Drawing on the *importance sampling paradigm*, their account aimed to minimize the mean squared error (MSE) of an expected utility estimator, as a well-established and normatively-justified measure of quality of an estimator (Poor, 2013). Since the variance of the estimator is the asymptotically-dominant term in the MSE (i.e., for large sample size, variance becomes an accurate proxy for MSE), Lieder et al. (2014, 2017) suggest that people adopt the following importance distribution (as the importance distribution minimizing the variance):

$$q(o) \propto p(o)|u(o) - \mathbb{E}_p[u(o)]|, \quad (1)$$

for mental simulations of events. In (1), o denotes an arbitrary event, p the objective probability of event o , $u(o)$ the utility of event o , q the probability distribution one adopts for their mental simulations (i.e., the subjective probability of event o), and, finally, $\mathbb{E}_p[\cdot]$ the expectation with respect to p .

Note that the expression in (1) does not depend on the number of samples drawn before deciding (denoted by s). In that light, Lieder et al.’s (2014, 2017) account implies that *time availability*, i.e., the amount of time a decision-maker has at their disposal, should have no implications on what importance distribution q one adopts. While a cognitively-rational agent is ignorant about adapting their importance distribution q based on time availability, a *metacognitively-rational* agent would plausibly make such considerations in their choice of q . That is, the metacognitively-rational agent chooses, among all q ’s, the one which is normatively-justified based on time availability considerations—this essentially makes it a strategy selection task guided by time availability. In agreement with this view, a large body of psychological work on decision-making suggests that (1) people evoke different strategies for decision making under time pressure vs. no time pressure, and (2) people adapt their strategies in accord with time availability (see e.g. Svenson & Maule, 1993; Svenson, 1993).

In this work, we present the first normative, metacognitive model of how an agent should over-represent extreme eventualities, depending on the amount of time available at their disposal for decision making. Concretely, our work serves as a rational, meta-level model for the work by Lieder et al. (2017, 2014). More specifically, the importance distribution suggested by Lieder et al. (2017, 2014) naturally follows from our metacognitive account, when s is large. In contrast to Lieder et al. (2017, 2014), our meta-level account also

specifies how a decision-maker should rationally choose their importance distribution when they can only afford to collect very few samples (i.e. when making decision under extremely high time pressure).¹ Importantly, recent work has provided mounting evidence suggesting that people often use very few samples in probabilistic judgments and reasoning under uncertainty (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), elevating the importance of developing process models specifically directed at few number of samples.

We show that our model can account for two well-known framing effects in human decision-making under risk: the fourfold pattern of risk preferences in outcome probability (Tversky & Kahneman, 1992) and in outcome magnitude (Markovitz, 1952). Despite being often taken as strong evidence for human irrationality, we provide the first metacognitively-rational basis for these effects. Furthermore, empirical evidence, confirms an important prediction of our model: over-representation of extreme events regardless of their frequencies. Our model also makes a counterintuitive (normative) prediction, which, if confirmed, would have surprising implications for human decision-making under risk.

2 Proposed Model

In this section, we formally present our metacognitively-rational model for the Availability bias (Tversky & Kahneman, 1973). According to the expected utility theory (von Neumann & Morgenstern, 1944), an agent chooses an action a , with the highest expected utility

$$\mathbb{E}[u(o)] = \int p(o|a)u(o)do, \quad (2)$$

where $p(o|a)$ denotes the distribution over outcomes o resulting from taking action a , $u(o)$ the subjective utility associated to outcome o , and $\mathbb{E}[\cdot]$ the expectation operation. Since the computation of (2) is generally intractable, we assume that the agent estimates (2) using *sampling methods* (Hammersley & Handscomb, 1964). Substantial neural and behavioral evidence supports this hypothesis (see e.g. Fiser, Berkes, Orbán, & Lengyel, 2010; Vul, Goodman, Griffiths, & Tenenbaum, 2014; Denison, Bonawitz, Gopnik, & Griffiths, 2013; Griffiths & Tenenbaum, 2006). Concretely, following Lieder et al. (2014, 2017), we assume that the agent estimates (2) using (self-normalized) importance sampling (Hammersley & Handscomb, 1964; Geweke, 1989), which is shown to be related to both neural networks (Shi & Griffiths, 2009) and cognitive process models (Shi, Griffiths, Feldman, & Sanborn, 2010):

$$\hat{E} = \frac{1}{\sum_{j=1}^s w_j} \sum_{i=1}^s w_i u(o_i), \quad \forall i: o_i \sim q, w_i = \frac{p(o_i)}{q(o_i)}. \quad (3)$$

¹The optimality of Lieder et al.'s (2017, 2014) model hinges on the number of samples s being large. When s is small, Lieder et al.'s (2017, 2014) model is no longer optimal. Our model, however, remains rational for both small and large values of s .

In Eq. (3), s denotes the total number of mental simulations performed by the agent, o_i the i^{th} mentally simulated outcome, $u(o_i)$ the utility of o_i , p the objective probability of event o_i , q the probability distribution the agent adopts for their mental simulations (i.e., the subjective probability of event o_i), and, \hat{E} the (normalized) importance sampling estimator of $\mathbb{E}[u(o)]$ given in (2).

The mean-squared error (MSE) of the estimator in (3), as a standard normative measure of the quality of an estimator (Poor, 2013), can be decomposed as follows: $\mathbb{E}[(\hat{E} - \mathbb{E}[u(o)])^2] = (\text{Bias}[\hat{E}])^2 + \text{Var}[\hat{E}]$, where The bias $\text{Bias}[\hat{E}]$ and variance $\text{Var}[\hat{E}]$ of the estimator \hat{E} can be approximated by (Zabaras, 2010):

$$\text{Bias}[\hat{E}] \approx \frac{1}{s} \int \frac{p(o)^2}{q(o)} (\mathbb{E}_p[u(o)] - u(o)) do, \quad (4)$$

$$\text{Var}[\hat{E}] \approx \frac{1}{s} \int \frac{p(o)^2}{q(o)} (\mathbb{E}_p[u(o)] - u(o))^2 do. \quad (5)$$

Under mild technical conditions, it can be shown that the rational importance distribution for minimizing the MSE of the estimator \hat{E} is given by:

$$q_{meta}^* \propto p(o) |u(o) - \mathbb{E}_p[u(o)]| \sqrt{\frac{1 + |u(o) - \mathbb{E}_p[u(o)]| \sqrt{s}}{|u(o) - \mathbb{E}_p[u(o)]| \sqrt{s}}}, \quad (6)$$

where p denotes the objective probability of event o , and $\mathbb{E}_p[\cdot]$ the expectation with respect to distribution p .² We refer to q_{meta}^* given in (6) as the metacognitively-rational importance distribution the agent should adopt for mental simulation of events for decision-making under uncertainty.

Comparing expressions (1) and (6) reveals that the multiplicative factor $\sqrt{\frac{1 + |u(o) - \mathbb{E}_p[u(o)]| \sqrt{s}}{|u(o) - \mathbb{E}_p[u(o)]| \sqrt{s}}}$, which we term *metacognitive rationality factor* (MCRF), differentiates Lieder et al.'s (2014, 2017) cognitively-rational model (see Eq. (1)) from our metacognitively-rational model. In the remainder of this work, we show that MCRF plays a crucial role in accounting for two important framing effects in decision-making under risk. It is crucial to note that q_{meta}^* takes into account the amount of time available for making a decision (i.e., time availability), as evidenced by expression (6) explicitly depending on the number of mental simulations s performed by the agent.

Following Lieder et al. (2017, 2014), and for ease of exposition, we assume $\mathbb{E}_p[u(o)] = 0$ hereinafter.

3 Extreme Eventualities are Over-represented Regardless of their Frequency

A simple investigation of the metacognitively-rational importance distribution q_{meta}^* given in (6) yields an important prediction of our model: Extreme eventualities should be over-represented in decision-making, regardless of how rare or frequent they are. Importantly, this effect is already empirically

²For the derivation of the expression given in (6), the reader is referred to <https://arxiv.org/abs/1801.09848>




Figure 1: A metacognitively-rational agent should over-represent extreme events precisely according to Lieder et al.’s (2014, 2017) cognitively-rational model, evidenced by the curves converging to 1 as $u(o) \rightarrow +\infty$. Importantly, however, a metacognitively-rational agent should also over-represent mundane events significantly more than what a merely cognitively-rational model prescribes, evidenced by the curves overshooting at the neighborhood of $u(o) = 0$.

confirmed (Lieder et al., 2017). Note that this coverage is to be expected as our proposed model subsumes the model outlined in Lieder et al. (2014, 2017).

Importantly, a detailed analysis of MCRF reveals that a metacognitively-rational agent should over-represent extreme events precisely according to Lieder et al.’s (2014, 2017) cognitively-rational model, however, it should also over-represent mundane events significantly more than what the cognitively-rational model by Lieder et al. prescribes. These findings are depicted in Fig. 1.

Next, we formally show that when the number of samples s is sufficiently large, our proposed metacognitively-rational importance distribution q_{meta}^* converges to the cognitively-rational importance distribution of Lieder et al. (2014, 2017) given in (1).³

³More accurately, in formal terms, q_{meta}^* converges to (1) almost surely, except at $u(o) = 0$. Notice that despite the unboundedness of MCRF at $u(o) = 0$ (see Fig. 1), q_{meta}^* remains bounded at

Proposition 1. When the number of mental simulations s is large, q_{meta}^* converges to the importance distribution given in (1). Formally, assuming $u(o) - \mathbb{E}_p[u(o)] \neq 0$,

$$\lim_{s \rightarrow +\infty} q_{meta}^* = \frac{1}{Z} p(o) |u(o) - \mathbb{E}_p[u(o)]|, \quad (7)$$

where Z is a normalizing constant (aka partition function).⁴

Proposition 1 formally establishes that our metacognitively-rational model of Availability bias serves as a rational, meta-level model for the work by Lieder et al. (2017, 2014), with our model converging to Lieder et al.’s when the number of samples s is large. Note that, since Lieder et al.’s importance distribution was specifically derived under the assumption that s is large, the result presented in Proposition 1 is intuitive, and, importantly, attests to the claim that our metacognitively-rational model subsumes Lieder et al.’s cognitively-rational model, with the rationality of our model holding for both small and large s ’s while that of Lieder et al.’s only for large s ’s.

4 Framing Effect in Decision-Making

Past work has documented that people’s risk preferences are inconsistent and context-dependent (see e.g., Tversky & Kahneman, 1992; Markovitz, 1952). For example, in choosing between a safe gamble (low payoff with high probability) and a risky gamble (high payoff with low probability), risk preferences change depending on the probabilities of the gambles (Tversky & Kahneman, 1992), the amount offered (Markovitz, 1952), and whether those gambles are framed as a gain or loss (Tversky & Kahneman, 1992).

In what follows, we show that our metacognitively-rational model can account for two well-known framing effects in human decision-making under risk: the fourfold pattern of risk preferences in outcome probability (Tversky & Kahneman, 1992) and in outcome magnitude (Markovitz, 1952). Thus,

$u(o) = 0$.

⁴For a formal proof of Proposition 1, the reader is referred to <https://arxiv.org/abs/1801.09848>.

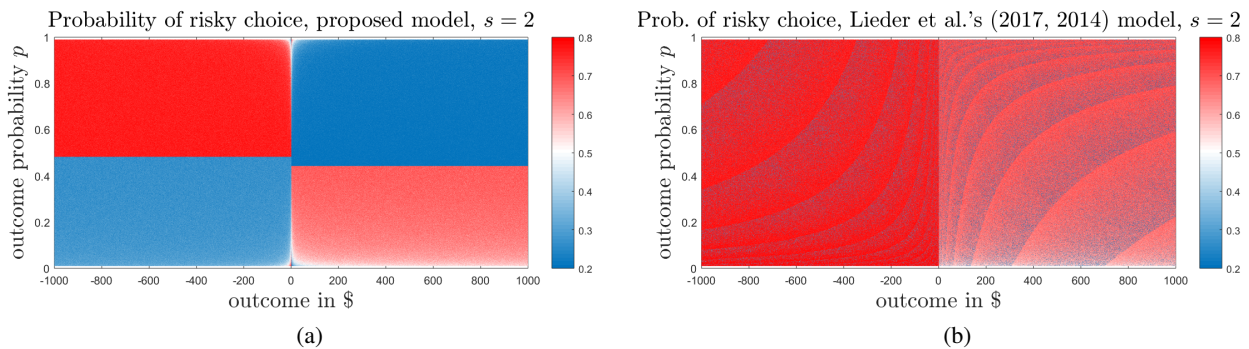


Figure 2: Accounting for the fourfold pattern of risk preferences in outcome probability (Tversky & Kahneman, 1992), with few samples ($s = 2$) and the utility function given in (8) based on prospect theory (Tversky & Kahneman, 1992). (a) Our metacognitively-rational model can account for the fourfold pattern of risk preferences in outcome probability, with $s = 2$ and the utility function given in (8). (b) The prediction of Lieder et al.’s (2014, 2017) cognitively-rational model for the probability of choosing the risky choice, with $s = 2$ and the utility function given in (8).

our model establishes the first metacognitively-rational basis for those effects.

4.1 Fourfold Pattern of Risk Preferences in Outcome Probability

Framing outcomes as losses rather than gains can reverse people’s risk preferences (Tversky & Kahneman, 1992): In the domain of gains people prefer a lottery (o dollars with probability p) to its expected value (risk seeking) when $p < 0.5$, but when $p > 0.5$ they prefer the expected value (risk-aversion). To the contrary, in the domain of losses people are risk averse for $p < 0.5$ but risk seeking for $p > 0.5$. This phenomenon is known as the fourfold pattern of risk preferences in probability outcome. Next we show that our metacognitively-rational model can simulate this effect. Following the prescriptions of the *prospect theory* (Tversky & Kahneman, 1992), as did Lieder et al. (2014) postulate, we assume that the agent’s utility function can be modeled by:

$$u(o) = \begin{cases} o^{0.85} & \text{if } o \geq 0, \\ -|o|^{0.95} & \text{if } o < 0. \end{cases} \quad (8)$$

Normatively, people should make their choice depending on whether the expected value of the utility difference $\Delta u(o)$ is negative or positive:

$$\Delta u(o) = \begin{cases} u(o) - u(p \times o) & \text{with probability } p, \\ -u(p \times o) & \text{with probability } 1 - p. \end{cases} \quad (9)$$

Fig. 2(a) shows that our metacognitively-rational model can account for the fourfold pattern of risk preferences in outcome probability (Tversky & Kahneman, 1992), with the utility function given in (8) based on prospect theory (Tversky & Kahneman, 1992) and very few samples ($s = 2$). This result is fully consistent with past work suggesting that people often use very few samples in probabilistic inference and reasoning under uncertainty (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).

Fig. 2(b) shows the prediction of Lieder et al.’s (2014, 2017) cognitively rational model for the probability of choosing the risky option, with $s = 2$ and the utility function given in (8) based on prospect theory (Tversky & Kahneman, 1992). Lieder et al.’s cognitively-rational model seems unable to account for the *probability of risky choice* suggested by Tversky and Kahneman (1992) using a suggested utility function of prospect theory given in (8); our simulations suggest that this apparent failure also holds for other values of s . However, Lieder et al.’s (2014, 2017) cognitively-rational model can partially account for this effect (see Fig. 3) based on the expected value of the importance sampling estimator given in (3), $\mathbb{E}[\hat{E}]$, replicating the finding reported in Lieder et al.’s (2014) Fig 3.⁵

⁵Importantly, the evaluation of $\mathbb{E}[\hat{E}]$ hinges on the assumption of averaging over multiple runs of mental simulations.




Figure 3: Expected value of the importance sampling estimator given in (3), $\mathbb{E}[\hat{E}]$, with $s = 2$ and the utility function given in (8), showing that Lieder et al.’s (2014, 2017) cognitively-rational model can partially account for the fourfold pattern of risk preferences in outcome probability. This replicates the finding reported in Lieder et al.’s (2014) Fig 3. However, Lieder et al.’s (2017, 2014) model appears to be unable to account for the *probability of risky choice* suggested by Tversky and Kahneman (1992), using the utility function given in (8) based on prospect theory; cf. Fig. 2(b).

In their recent work, Lieder et al. (2017) showed that their cognitively-rational model can better account for the fourfold pattern of risk preferences in outcome probability, provided that the utility function is noisy (*efficient neural coding*, Summerfield and Tsetsos, 2015); see Fig. 4 in Lieder et al. (2017).⁶ The result reported in Fig. 2(a) strongly suggests that this effect can be accounted for by a purely metacognitively-rational model together with a utility function fully consistent with prospect theory (Tversky & Kahneman, 1992), without necessarily having to invoke a noisy utility function (see Lieder et al., 2017, Appendix C).

4.2 Fourfold Pattern of Risk Preferences in Outcome Magnitude

Past work in behavioral economics has documented another curious inconsistency in human decision making under risk: the fourfold pattern of risk preferences in outcome magnitude (Markowitz, 1952; Hershey & Schoemaker, 1980; Scholten & Read, 2014). Concretely, in choosing between a sure thing and a low-probability risky gamble people demonstrate the following behavioral pattern: In moderate-to-large outcomes, people are risk-averse for gains and risk-seeking for losses. This pattern reverses when outcomes are small, with people being risk-seeking for gains and risk-averse for losses. For example, people would rather choose a sure 1 million dollar option rather than a (low-probability) risky gamble yielding \$10 million dollars with probability 0.1 and nothing otherwise (Hershey & Schoemaker, 1980). When framed in the context of losses, people prefer a risky gamble yielding \$10 million dollar loss with probability 0.1 and nothing otherwise, over a sure loss of \$1 million dollars (Markowitz, 1952).

Prospect theory (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992), in its original form, cannot account

⁶Specifically, Lieder et al. (2017) adopt the noisy utility function $u(o) = \frac{o}{o_{max} - o_{min}} + \epsilon$, where ϵ is an additive Gaussian noise, i.e., $\epsilon \sim N(0, \sigma^2)$.




Figure 4: Simulating both the fourfold pattern of risk preferences in outcome probability and in outcome magnitude, with few samples ($s = 2$) and the normalized logarithmic utility function in (10) with $\alpha = 0.032$ and $\beta = 0.0031$. **(a)** Our model can simultaneously account for both effects: *(Top-Right)* Moving from left to right along the x -axis within the boxed region clearly shows the risk preference reversal from risk-seeking to risk-aversion (in losses), back to risk-seeking and finally to risk-aversion (in gains). *(Bottom-Right)* The fourfold pattern of risk preferences in outcome probability. **(b)** Lieder et al.'s (2014, 2017) cognitively-rational model prediction under the same setting as **(a)**.

for the fourfold of risk preferences in outcome magnitude (Scholten & Read, 2014). However, Scholten and Read (2014) show that, armed with a particular choice of utility function, prospect theory can accommodate this effect. Concretely, they show that prospect theory can best account for this effect by adopting the *normalized logarithmic* utility function (Rachlin 1992; Scholten & Read, 2010; Kirby, 2011; Kontek, 2011):

$$u_{nlog}(o) = \begin{cases} \frac{1}{\alpha} \log(1 + \alpha \times o) & \text{if } o \geq 0, \\ -\frac{\lambda}{\beta} \log(1 - \beta \times o) & \text{if } o < 0, \end{cases} \quad (10)$$

where $\alpha, \beta \in \mathbb{R}^{>0}$, $\lambda \geq 1$ are free parameters.

Using empirical data, Scholten and Read (2014) found the maximum-likelihood estimates of α and β to be 0.032 and 0.0031, respectively (see Scholten and Read, 2014, Table 4). Adopting the normalized logarithmic utility function in (10) with $\lambda = 1$, $\alpha = 0.032$, $\beta = 0.0031$, we show that our Sample-based Expected Utility model can *simultaneously* account for both the fourfold pattern of risk preferences in outcome probability (Tversky & Kahneman, 1992) and in outcome magnitude (Markovitz, 1952); see Fig. 4(a). However, Lieder et al.'s (2014, 2017) cognitively-rational model appears to be unable to account for these effects under the same setting. Again, our simulations suggest this apparent failure holds for other values of s . These findings indicate that the fourfold pattern of risk preferences in outcome magnitude could stem from the optimization of a boundedly-rational agent's decision strategy at the metacognitive level, as suggested by (6).

5 Sensitivity Analysis

As discussed earlier, a metacognitively-rational agent optimizes their decision strategy (in our case, their importance distribution for mental simulations) according to time availability. This requires the agent to have a good estimate of the number of samples s they will likely draw within the given time frame, using which they can appropriately select their importance distribution q_{meta}^* . However, a crucial question

immediately presents itself: What happens if the agent is inaccurate at approximating the number of samples they get to draw before making their decision? After all, it seems plausible to assume that the agent would only have a rough estimate of the parameter s . Thus, it is very likely that there would be a mismatch between the number of samples the agent thinks they can draw, and the actual number of samples they ultimately draw. Our model nicely allows for a quantitative investigation of the effects of such a mismatch. The parameter s in (6) indicates the the number of samples the agent *thinks* they can draw, whereas the parameter s in (3) reflects the the number of samples the agent *draws* before making a decision. It is worth noting that the cognitively-rational model by Lieder et al. (2014, 2017) does not permit the investigation of the possible mismatch alluded to above, as the parameter s does not feature in Lieder et al.'s importance distribution (Eq. (1)).

Intriguingly, our model demonstrates a striking insensitivity to such mismatches: Even if the the number of sample the agent thinks they can draw is unimaginably greater (to be precise, 10^8 times greater) than the the number of samples they actually draw before making their decision, the agent should still show the fourfold patterns. Figures are omitted due to lack of space.

6 General Discussion

People overestimate the probability of extreme events, and show the fourfold pattern of risk preferences in outcome probability (Tversky & Kahneman, 1992) and in outcome magnitude (Markovitz, 1952) in decision-making under risk; these effects are generally taken as evidence against human rationality. In this work, we presented the first metacognitively-rational process model which can account for the aforementioned effects, indicating that these apparent biases may not be signs of human irrationality after all, but the result of a boundedly-rational decision-maker *optimizing* their decision strategy (in our case, their importance sampling distribution for performing mental simulations) in accord with time availability. In fact, it can be shown that the

metacognitively-rational importance distribution q_{meta}^* in (6) allows the decision-maker to ensure an upper-bound on the MSE of their estimator for the expected value in (2) using *minimal* number of samples, thereby demonstrating strong signs of economy (*rational minimalist program*, Nobandegani, 2017). Furthermore, our model is unimaginably robust to inaccurate estimations of its focal parameter s , positioning it as the first rational process model we know of which scores near-perfectly in optimality, economical use of limited cognitive resources, and robustness, all at the same time. Intriguingly, this aspect of our model, together with the fact that Eq. (3) in conjunction with q_{meta}^* can be used to decide between an arbitrary number of actions a , makes our model also very appealing to the artificial intelligence community.

The metacognitively-rational process model presented in this work and Lieder et al.'s (2014, 2017) cognitively-rational process model seem to suggest that a (boundedly) rationalist approach to understanding human decision-making at the algorithmic level might be a fruitful endeavor. In fact, the influential Rescorla-Wagner model (1972) and its extension temporal-difference learning model (Sutton & Barto, 1987; Sutton & Barto, 1998) can be given solid rational grounds based on linear-Gaussian generative models and the Kalman filtering paradigm, a rational scheme in signal detection theory (Kalman, 1960).

Our model also makes a counterintuitive (normative) prediction, which, if confirmed, would have surprising implications for human decision-making under risk: In choosing between a lottery (o dollars with probability p) and its expected value ($p \cdot o$), people should qualitatively behave the same under the following two conditions: (i) making a decision based on a mere single sample (i.e., under extremely high time pressure) and (ii) making a decision based on a great many samples (i.e., after a long deliberation time). Note that, given the normative status of our model, this is exactly the behavior that a boundedly-rational agent should manifest, a finding which would be of great interest for the artificial intelligence community. If confirmed, this prediction seems to suggest an intriguing possibility for human decision-making under risk: People's performance after long deliberation time is qualitatively similar to their performance under extremely high time pressure (i.e., $s = 1$). This clearly serves as a motivation for avoiding over-thinking.

For their cognitively-rational process model, Lieder et al. (2017) proposed a neurally-plausible learning mechanism, a simple modification of which permits our metacognitively-rational to be learned in a neurally-plausible manner as well. Lieder et al. (2017) showed that their model can account for an impressively wide range of cognitive biases in decisions from experience, decisions from description, and memory recall. Future work should investigate how well our metacognitively-rational model can account for those biases. The fact that our model subsumes Lieder et al.'s (2014, 2017) model (see Proposition 1), greatly elevates the possibility of our model capturing those effects as well.

To our knowledge, our Sample-based Expected Utility model is the first *metacognitive*, resource-rational process model of cognitive biases. In that light, our work suggests perhaps those cognitive biases that do not easily lend themselves to cognitive-level accounts, might be better suited to higher-level explanations (e.g., the metacognitive-level). We hope to have shed some light on possible rational grounds of human decision-making.

Acknowledgments: We would like to thank Falk Lieder for fruitful discussions. This work is supported by an operating grant to TRS from 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.
- 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.
- Cruciani, F., Berardi, A., Cabib, S., & Conversi, D. (2011). Positive and negative emotional arousal increases duration of memory traces: common and independent mechanisms. *Frontiers in behavioral neuroscience*, 5, 86.
- Denison, S., Bonawitz, E., Gopnik, A., & Griffiths, T. L. (2013). Rational variability in children's causal inferences: The sampling hypothesis. *Cognition*, 126(2), 285–300.
- Fiser, J., Berkes, P., Orbán, G., & Lengyel, M. (2010). Statistically optimal perception and learning: from behavior to neural representations. *Trends in cognitive sciences*, 14(3), 119–130.
- 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., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological science*, 17(9), 767–773.
- 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.
- Hershey, J. C., & Schoemaker, P. J. (1980). Prospect theory's reflection hypothesis: A critical examination. *Organizational Behavior and Human Performance*, 25(3), 395–418.
- 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.
- Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Journal of basic Engineering*, 82(1), 35–45.
- Kirby, K. N. (2011). An empirical assessment of the form of utility functions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37(2), 461.
- Kontek, K. (2011). On mental transformations. *Journal of Neuroscience, Psychology, and Economics*, 4(4), 235.
- 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.
- Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M., & Combs, B. (1978). Judged frequency of lethal events. *Journal of experimental psychology: Human learning and memory*, 4(6), 551.
- Lieder, F., Griffiths, T. L., & Hsu, M. (2017). Overrepresentation of extreme events in decision making reflects rational use of cognitive resources. *Psychological Review*.
- Lieder, F., Hsu, M., & Griffiths, T. L. (2014). The high availability of extreme events serves resource-rational decision-making. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 36).
- Markowitz, H. (1952). The utility of wealth. *Journal of political Economy*, 60(2), 151–158.
- Maule, A. J., & Svenson, O. (1993). Theoretical and empirical approaches to behavioral decision making and their relation to time constraints. In *Time pressure and stress in human judgment and decision making* (pp. 3–25). Springer.
- Nobandegani, A. S. (2017). *The Minimalist Mind: On Minimality in Learning, Reasoning, Action, & Imagination*. McGill University, PhD Dissertation.
- Poor, H. V. (2013). *An Introduction to Signal Detection and Estimation*. Springer Science & Business Media.
- Rachlin, H. (1992). Diminishing marginal value as delay discounting. *Journal of the Experimental Analysis of Behavior*, 57(3), 407–415.
- Rescorla, R. A., Wagner, A. R., et al. (1972). A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Classical conditioning II: Current research and theory*, 2, 64–99.
- Rothman, A. J., Klein, W. M., & Weinstein, N. D. (1996). Absolute and relative biases in estimations of personal risk. *Journal of Applied Social Psychology*, 26(14), 1213–1236.
- Scholten, M., & Read, D. (2010). The psychology of intertemporal tradeoffs. *Psychological review*, 117(3), 925.
- Scholten, M., & Read, D. (2014). Prospect theory and the forgotten fourfold pattern of risk preferences. *Journal of Risk and Uncertainty*, 48(1), 67–83.
- Shi, L., & Griffiths, T. L. (2009). Neural implementation of hierarchical bayesian inference by importance sampling. In *Advances in neural information processing systems* (pp. 1669–1677).
- Shi, L., Griffiths, T. L., Feldman, N. H., & Sanborn, A. N. (2010). Exemplar models as a mechanism for performing bayesian inference. *Psychonomic Bulletin & Review*, 17(4), 443–464.
- Sunstein, C. R., & Zeckhauser, R. (2011). Overreaction to fearsome risks. *Environmental and Resource Economics*, 48(3), 435–449.
- Sutton, R. S., & Barto, A. G. (1987). A temporal-difference model of classical conditioning. In *Proceedings of the ninth annual conference of the cognitive science society* (pp. 355–378).
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction* (Vol. 1) (No. 1). MIT press Cambridge.
- Svenson, O. (1993). *Time pressure and stress in human judgment and decision making*. Springer Science & Business Media.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4), 297–323.
- von Neumann, J., & Morgenstern, O. (1955). *The theory of games and economic behavior*. Princeton University Press.
- 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.