

The high availability of extreme events serves resource-rational decision-making

Falk Lieder (falk.lieder@berkeley.edu)

Helen Wills Neuroscience Institute, University of California at Berkeley, CA, USA

Ming Hsu (mhsu@haas.berkeley.edu)

Haas School of Business, University of California at Berkeley, CA, USA

Thomas L. Griffiths (tom_griffiths@berkeley.edu)

Department of Psychology, University of California at Berkeley, CA, USA

Abstract

Extreme events come to mind very easily and people overestimate their probability and overweight them in decision-making. In this paper we show that rational use of limited cognitive resources can generate these 'availability biases.' We hypothesize that availability helps people to quickly make good decisions in very risky situations. Our analysis shows that agents who decide by simulating a finite number of possible outcomes (sampling) should over-sample outcomes with extreme utility. We derive a cognitive strategy with connections to fast-and-frugal heuristics, and we experimentally confirm its prediction that an event's extremity increases the factor by which people overestimate its frequency. Our model also explains three context effects in decision-making under risk: framing effects, the Allais paradox, and preference reversals.

Keywords: availability; Bayesian; bounded rationality; cognitive biases; heuristics; judgment and decision-making

Introduction

People overestimate the probability of extreme events such as terrorism (Sunstein & Zeckhauser, 2011) and other threats (Lichtenstein, Slovic, Fischhoff, Layman, & Combs, 1978; Rothman, Klein, & Weinstein, 1996) and overreact accordingly (Sunstein & Zeckhauser, 2011). This phenomenon has been explained by the *availability bias* (Tversky & Kahneman, 1973) according to which people overestimate the probability of events that come to mind very easily.

The availability bias appears irrational, but here we argue that it reflects the rational use of finite time (*resource-rationality*; Lieder, Griffiths, & Goodman, 2013, Griffiths, Lieder, & Goodman, in revision). In brief, we hypothesize that the high availability of extreme events helps decision-makers to allocate their finite time towards considering the most important consequences their decision might have. We model the strategy that might determine which events come to mind first and how they influence judgment and decision-making. Our model explains not only why people overestimate the probability of extreme events, but it also explains three context effects in decision-making under risk.

The plan of this paper is as follows: The first section introduces the theoretical and empirical background. The second section derives a rational model of decision-making in high-risk situations. The third section presents an experiment testing the model's predictions for frequency judgment, and the fourth section applies the model to explain context-effects in decision-making under risk. The final section discusses our results and their implications.

Sampling as a Decision Strategy

To evaluate a potential action a , decision makers should integrate the probabilities $P(o|A = a)$ of possible outcomes o with their utilities $u(o)$ into the action's *expected utility* $\mathbb{E}_{p(o|A=a)}[u(o)]$ (Von Neumann & Morgenstern, 1944). In the real-world—unlike in the laboratory—each action has infinitely many possible outcomes. As a consequence, the expected utility of action a becomes an integral:

$$\mathbb{E}_{p(o|A=a)}[u(o)] = \int p(o|a) \cdot u(o) do. \quad (1)$$

In the general case, this integral is intractable to compute, but it can be approximated by *sampling methods* (Hammersley & Handscomb, 1964). Mental simulations of a decision's potential consequences can be thought of as samples. In fact, there is neural evidence (Fiser, Berkes, Orbán, & Lengyel, 2010) and behavioral evidence (Vul, Goodman, Griffiths, & Tenenbaum, 2014; Denison, Bonawitz, Gopnik, & Griffiths, 2013; Griffiths & Tenenbaum, 2006) that the brain handles uncertainty by sampling. For instance, people's predictions of an uncertain quantity X given partial information y are roughly distributed according to the posterior distribution $p(X|y)$ as if they were sampled from it (Griffiths & Tenenbaum, 2006; Vul et al., 2014). These results suggest that people often use only one or very few samples, and this is what rational agents with finite computational resources (*bounded rational agents*) should do (Vul et al., 2014). The evidence for sampling in human cognition raises the question which sampling algorithm(s) are implemented in the brain.

Importance sampling is a popular sampling algorithm in computer science and statistics (Hammersley & Handscomb, 1964; Geweke, 1989), and it has been shown to have connections to both neural network (Shi & Griffiths, 2009) and psychological process models (Shi, Griffiths, Feldman, & Sanborn, 2010). *Self-normalizing* importance sampling estimates the expected value of a function by the weighted average of the function's values at points drawn from a distribution q :

$$X_1, \dots, X_s \sim q, \quad w_j = \frac{p(x_j)}{q(x_j)} \quad (2)$$

$$\mathbb{E}_p[f(X)] \approx \hat{E}_{q,s}^{\text{IS}} = \frac{1}{\sum_{j=1}^s w_j} \cdot \sum_{j=1}^s w_j \cdot f(x_j), \quad (3)$$

where w_i is the *weight* of the i^{th} sample. With finitely many samples, this estimate is generally biased. Following Zabaras

(2010), we approximate its bias and variance by

$$\text{Bias}[\hat{E}_{q,s}^{\text{IS}}] \approx \frac{1}{s} \cdot \int \frac{p(x)^2}{q(x)} \cdot (\mathbb{E}_p[f(x)] - f(x)) dx \quad (4)$$

$$\text{Var}[\hat{E}_{q,s}^{\text{IS}}] \approx \frac{1}{s} \cdot \int \frac{p(x)^2}{q(x)} \cdot (f(x) - \mathbb{E}_p[X])^2 dx \quad (5)$$

Importance sampling can be used to approximate the expected utility $\mathbb{E}_{p(O|A=a)}[u(O)]$ of taking action a and to estimate the optimal decision $a^* = \arg \max_a \mathbb{E}_{p(O|A=a)}[u(O)]$:

$$\hat{a}^* = \arg \max_a \hat{U}_{q,s}^{\text{IS}}(a), \quad \hat{U}_{q,s}^{\text{IS}}(a) \approx \mathbb{E}_{p(O|a)}[u(O)] \quad (6)$$

$$\hat{U}_{q,s}^{\text{IS}}(a) = \frac{1}{\sum_{j=1}^s w_j} \sum_{j=1}^s w_j \cdot u(o_j), \quad o_1, \dots, o_s \sim q. \quad (7)$$

Note that importance sampling is a family of algorithms: each importance distribution q yields a different estimator, and two estimators may recommend opposite decisions. We thus consider which distribution q yields the best decisions.

Which distribution should we sample from?

Intuitively, the rational way to make a decision is to simulate consequences o according to one's best knowledge of the probability p that they will occur, i.e. $q = p$:

$$\hat{U}_s^{\text{RS}}(a) = \frac{1}{s} \sum_{i=1}^s u(o_i), \quad o_1, \dots, o_s \sim p(O). \quad (8)$$

This importance sampling algorithm, which we call *representative sampling*, computes unbiased utility estimates. Yet—surprisingly—representative sampling is insufficient for making good decision with very few samples. Consider, for instance, the choice to play Russian roulette with the popular six-round NGant M1895 revolver. Playing the game will most likely, i.e. with probability $p_1 = \frac{5}{6}$, reward you with a thrill and a gain in status ($u(o_1) = 1$) but kill you otherwise ($p_2 = \frac{1}{6}$, $u(o_2) = -10^9$). Ensuring that representative sampling declines a game of Russian roulette at least 99.99% of the time, would require 51 samples—potentially a very time-consuming computation. Many real-life decisions involve high risks and are even more challenging, because the probability of disaster is orders of magnitude smaller than $\frac{1}{6}$ but may or may not be large enough to warrant caution. Examples include risky driving, the stock market, and air travel. For some of these choices (e.g. texting while driving) there may be a one in a million chance of disaster while all other outcomes have negligible utilities:

$$u(o_d) = -10^9, p(o_d) = 10^{-6}, \quad \forall i \neq d : |u(o_i)| \leq 1 \quad (9)$$

If people decided based on n representative samples, they would completely ignore the potential disaster with probability $1 - (1 - 10^{-6})^n$. Thus to have at least a 50% chance of taking the potential disaster into account the agent would have to generate almost 700 million samples. This is clearly infeasible; thus one would almost always take this risk even

though the expected utility is about -1000 . In conclusion, under high risk, representative sampling is insufficient for resource-bounded decision-making.

What is the problem with representative sampling and how can it be overcome? Representative sampling fails when it neglects crucial eventualities. Neglecting some eventualities is necessary, but particular eventualities are more important than others. Intuitively, the importance of potential outcome o_i is determined by $|p(o_i) \cdot u(o_i)|$ because neglecting o_i amounts to dropping the addend $p(o_i) \cdot u(o_i)$ from the expected-utility integral (Equation 1). Thus, intuitively, the problem of representative sampling can be overcome by considering outcomes whose importance ($|p(o_i) \cdot u(o_i)|$) is high and ignoring those whose importance is low.

Formally, the agent's goal is to maximize the expected utility of a decision made from only s samples. The utility foregone by choosing a sub-optimal action can be upper-bounded by the error in a rational agent's utility estimate. Therefore the agent should minimize the expected squared error of its utility estimate, which is the sum of its squared bias and variance, i.e. $\mathbb{E}[(\hat{U}_{q,s}^{\text{IS}} - \mathbb{E}[U])^2] = \text{Bias}[\hat{U}_{q,s}^{\text{IS}}]^2 + \text{Var}[\hat{U}_{q,s}^{\text{IS}}]$. As the number of samples s increases, the estimate's squared bias decays much faster ($O(s^{-2})$) than its variance ($O(s^{-1})$); see Equations 4-5. Therefore, as the number of samples s increases, minimizing the estimator's variance becomes a good approximation to minimizing its expected squared error.

According to variational calculus the variance (Equation 5) of the utility estimate in Equation 7 is minimized by

$$q^{\text{var}}(o) \propto p(o) \cdot |u(o) - \mathbb{E}_p[U]|. \quad (10)$$

Interestingly, the optimal sampling distribution overrepresents outcomes with large absolute utility. Thus, biased sampling can lead to better decisions. Unfortunately, importance sampling with q^{var} is intractable, because it presupposes the expected utility $\mathbb{E}_p[U]$ that importance sampling is supposed to approximate. A priori the expected utility of a prospect whose outcomes may be good or bad is equally likely to be positive or negative. The same is true for choosing action a over action b or vice versa. Therefore, the most principled guess an agent can make for the expected utility $\mathbb{E}_p[U]$ in Equation 10—before computing it—is 0. Thus when the expected utility is not too far from zero, then the importance distribution q^{var} is efficiently approximated by

$$\tilde{q}(o) \propto p(o) \cdot |u(o)|. \quad (11)$$

This confirms our intuition and leads to an importance sampling scheme that we call *utility-weighted sampling*:

$$\hat{U}_{\tilde{q},s}^{\text{IS}} = \frac{1}{\sum_{j=1}^s 1/|u(o_j)|} \sum_{j=1}^s \text{sgn}(u(o_j)), \quad o_j \sim \tilde{q}, \quad (12)$$

where $\text{sgn}(x)$ is $+1$ for positive x and -1 for negative x . Utility-weighted sampling is a simple and psychologically plausible strategy for decision-making under uncertainty: It generates examples of possible outcomes by an appropriately

biased simulation and tallies how many of them are positive. If more than half of the examples for the utility of choosing action a over action b are positive, then the agent will prefer action a ; if more than half of them are negative, then the agent will prefer action b ; otherwise it will be indifferent.

Utility-weighted sampling succeeds where representative sampling failed. In Russian roulette utility-weighted sampling requires only 1 rather than 51 samples to recommend the correct decision at least 99.99% of the time, because the first sample is almost always the most important potential outcome, i.e. death. Likewise one single utility-weighted sample suffices to consider the potential disaster (Equation 9) at least 99.85% of the time, whereas even 700 million representative samples would miss the disaster almost half of the time. Thus, utility-weighted sampling would allow people to avoid both disasters even under extreme time pressure. The hypothesis that people decide by utility-weighted sampling makes two predictions that we test in the next two sections:

1. People overestimate an event’s probability more, the more extreme the event is.
2. Extreme potential outcomes are over-weighted in decision-making, and extremity is context-dependent.

Overestimation of extreme events’ frequencies

We hypothesize that the mind re-uses utility-weighted sampling (Equation 12) to estimate event frequencies, because evolutionary fitness depended on good decisions rather than accurate statements of probability. We therefore predict that people’s estimate \hat{p}_k of the probability $p_k = p(o_k)$ is

$$\hat{p}_k = \frac{\sum_{i=1}^s w_i \cdot \mathbb{1}(o_i = o_k)}{\sum_{i=1}^s w_i}, \quad w_i = \frac{1}{|u(o_i)|}, \quad o_i \sim \tilde{q}. \quad (13)$$

Since the importance density $\tilde{q}(o) \propto p(o) \cdot |u(o)|$ over-represents events proportionally to their extremity $|u|$, i.e. $\frac{\tilde{q}(o)}{p(o)} \propto |u(o)|$, we predict that people’s relative over-estimation $\frac{\hat{p}_k}{p_k}$ is a monotonically increasing function of the event’s extremity $|u_k|$. Formally, the bias (Equation 4) of utility-weighted probability estimation (Equation 13) implies that $\frac{\hat{p}_k - p_k}{p_k} = \frac{1}{s} \left(-\frac{1}{|u_i|} + C \right)$, for some constant C . We tested this prediction with a simple experiment.

Methods

We recruited 100 participants on Amazon Mechanical Turk. Each participant estimated how many of 1000 randomly selected American adults had experienced each of 39 events in 2013 and judged the events’ valence (good or bad) and extremity (0: neutral – 100: extreme). The 39 events comprised 30 stressful life events from Hobson et al. (1998), four lethal events (suicide, homicide, lethal accidents, and dying from disease/old age), three more mundane events (going to the movies, headache, and food-poisoning), and two attention-checks. Overestimation was measured by the ratio of a participant’s estimate over the event’s frequency according to of-




Figure 1: Judged frequency and extremity by event type.

ficial statistics.¹ The complete experiment can be inspected online.² Out of 100 participants 17 failed the attention check (wrong answer on attention-check questions, or mean judged extremity of lethal events less than 75%) and were excluded.

Results and Discussion

Consistent with our theory, participants overestimated the frequencies of all lethal events (all $p < 0.0005$) but none of the mundane events (all $p > 0.048 > \alpha_{\text{Sidak}} = 0.0014$, Sidak-correction for multiple comparisons). Participants also overestimated 23 of the 30 stressful life events. Figure 1 shows that across event types overestimation gradually increase with extremity: Absolute overestimation ($\hat{p} - p$) was significantly larger for stressful life events than for mundane events ($p < 0.01$) and even larger for lethal events ($p = 0.02$). While our participants’ probability estimates were very accurate for mundane events, their (implicit) estimate of the annual death rate was 25-times higher than its true value. For stressful life events overestimation and judged extremity are both intermediate. This pattern answers the open question whether availability or *regressed frequency* causes the overestimation of extreme events (Hertwig, Pachur, & Kurzenhäuser, 2005): if estimates were merely regressive towards a mean frequency, as proposed by Gigerenzer (2008a), then people should underestimate frequent mundane events, but our participants did not. Figure 2 shows that the average relative overestimation of the 37 events increases with their judged extremity (Spearman’s rank correlation $\rho = 0.53, p < 0.001$). Furthermore, we observed the same effect at the level of individual judgments (Spearman’s rank correlation $\rho = 0.28, p < 10^{-15}$). On average, the extremest event’s frequency (murder, 98% extreme) was overestimated by a factor of 972, whereas the frequency of the least extreme event (headache, 20% extreme) was overestimated by merely 6% (n.s.). Consistent with this result, Rothman et al. (1996) found that prevalence is more heavily overestimated for suicide than for divorce.

In conclusion, the experiment confirmed our theory’s prediction that an event’s extremity increase the relative overestimation of its frequency. However, additional experiments are

¹Hobson & Delunas (2001), www.cdc.gov/nchs/fastats/deaths.htm, www.mpa.org/resources/3037b7a4-58a2-4109-8012-58fca3abdflb.pdf, www.cdc.gov/foodborneburden/, Rasmussen, Jensen, Schroll, & Olesen (1991).

²<http://sites.google.com/site/falkliedez/freq.estimation.html>




Figure 2: Overestimation $\frac{\hat{p}-p}{p}$ increases with perceived extremity ($|u|$). A cross represents one event’s average ratings.

required to disentangle the effects of extremity and low probability, because these two factors were anti-correlated. Future work will formally test our model against competing theories and investigate whether the effect of an event’s extremity is mediated by how many instances of this event people imagine or retrieve from memory (cf. Hertwig et al., 2005).

Context effects in decision-making

While our goal was to explain why people overestimate the probability of extreme events, utility-weighted sampling also predicts that people over-weight certain outcomes in economic decisions. To simulate such decisions we model the utility of winning or losing money by the following function based on *prospect theory* (Tversky & Kahneman, 1992):

$$u(o) = \begin{cases} o^{0.85} & , \text{ if } o \geq 0 \\ -|o|^{0.95} & , \text{ if } o < 0 \end{cases} \quad (14)$$

When choosing between receiving a high payoff with low probability (*risky gamble*) or a low payoff with high probability (*safe gamble*), people often prefer the risky gamble in one context but the safe gamble in another. This suggests that people’s *risk preferences* are inconsistent and context-dependent (Tversky & Kahneman, 1992). Utility-weighted sampling can explain several such inconsistencies: (i) *framing effects*, (ii) *the Allais paradox*, and (iii) *preference reversals*. We now consider these in turn.

Framing effects on risk attitudes

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 seeking for $p < 0.5$ but risk averse for $p > 0.5$. This phenomenon is known as the *fourfold pattern of risk preferences*.

Whether or not people should choose the gamble depends on the expected value of the utility difference ΔU :

$$\Delta U = \begin{cases} u(o) - u(p \cdot o) & \text{with probability } p \\ -u(p \cdot o) & \text{with probability } 1 - p \end{cases} \quad (15)$$




Figure 3: Risk preferences ($\mathbb{E}[\Delta\hat{U}]$) predicted by utility-weighted sampling with $s = 2$.

According to our theory, people simulate ΔU by sampling from $\tilde{q}(\Delta u) \propto |\Delta u| \cdot p(\Delta u)$ and estimate $\mathbb{E}[\Delta U]$ by $\Delta\hat{U}_{\tilde{q},s}^{IS}$ according to Equation 12. When the bias of this estimate ($\text{Bias}(\Delta\hat{U}_{\tilde{q},s}^{IS})$) is positive, the agent will fancy the gamble (risk-seeking), but when it is negative, the agent will be risk-averse. The importance distribution \tilde{q} is proportional to $|\Delta u|$. Therefore the agent will overweight the gain/loss o of the lottery if $p(o)$ is small, because then $|u(o) - u(p \cdot o)| > |u(p \cdot o)|$. Conversely, the agent will underweight outcome o if $p(o)$ is large, because then $|u(o) - u(p \cdot o)| < |u(p \cdot o)|$. This renders the agent’s bias positive (risk-seeking) for improbable gains and probable losses but negative (risk-aversion) for probable gains and improbable losses (see Figure 3). Thus our model predicts the fourfold pattern of risk preferences which is used to explain how a person who is so risk-averse that he buys insurance can also be so risk-seeking that he plays the lottery.

Next we show by simulation that a bounded rational agent might indeed make both choices. First, we simulated the decision whether or not to play the Powerball lottery.³ The jackpot is at least \$40 million, but the odds of winning it are less than 1:175 million. In brief, people pay \$2 to play a gamble whose expected value is only \$1. We found that utility-weighted sampling does, in expectation, fancy the lottery when the number of samples is less than 4 ($\mathbb{E}[\hat{U}_{\tilde{q},s}^{IS}] \geq 59.8$ for $s \leq 3$). If the agent chose based on one sample (cf. Vul et al., in press) and was given the choice twice a week, it would buy about seven lottery tickets per year. This value is a lower bound, because it ignores debt and other factors. Nevertheless, it demonstrates that playing the lottery is compatible with resource-rational decision-making. Second, we simulated the decision whether or not to buy insurance. The magnitude of an insured loss l can be modeled by the Pareto distribution (e.g., $p(l) = \alpha \cdot x_{\min}^{\alpha} \cdot x^{-\alpha-1}$, $x_{\min} < x < 10^9$ with $\alpha = x_{\min} = 1$). We found that the bias of utility-weighted sampling (Equation 4) would make people overestimate the value of insurance against such a loss by $\frac{244}{s}\%$, where s is the number of samples. This resolves the apparent paradox of being willing to buy both lottery tickets and insurance.

³ www.calottery.com/play/draw-games/powerball

Table 1: Allais Gambles

	(o_1, p_1)	(o_2, p_2)	(o_3, p_3)
$L_1(z) :$	$(z, 0.66)$	$(2500, 0.33)$	$(0, 0.01)$
$L_2(z) :$	$(z, 0.66)$	$(2400, 0.34)$	

Table 2: Utility-weighted sampling explains the Allais paradox.

	ΔU	p	\tilde{q}	\tilde{q}/p
$z = 2400:$	0	0.66	0	0
	$u(2500) - u(2400)$	0.33	0.54	1.6
	$-u(2400)$	0.01	0.46	46
	<hr/>	<hr/>	<hr/>	<hr/>
	0	$0.66 \cdot 0.67$	0	0
$z = 0 :$	$-u(2400)$	$0.67 \cdot 0.34$	0.5	2.19
	$u(2500) - u(2400)$	$0.33 \cdot 0.34$	0.01	0.08
	$u(2500)$	$0.33 \cdot 0.66$	0.49	2.26

Note: The agent’s simulation yields $\Delta U = \Delta u$ with probability $\tilde{q}(\Delta u) \propto p(\Delta u) \cdot |\Delta u|$ where p is Δu ’s objective probability.

The Allais paradox

In the two gambles $L_1(z)$ and $L_2(z)$ defined in Table 1 the chance of winning z dollars is exactly the same. Yet, when $z = 2400$ most people prefer L_2 over L_1 , but when $z = 0$ the same people prefer L_1 over L_2 . This inconsistency is known as the Allais paradox. Table 2 shows how our theory resolves this paradox: According to the importance distribution \tilde{q} (Equation 11) people overweight the event for which the utility difference between the two gambles’ outcomes (O_1 and O_2) is largest ($\Delta U = u(O_1) - u(O_2)$). Thus when $z = 2400$, the most over-weighted event is the possibility that gamble L_1 yields $o_1 = 0$ and gamble L_2 yields $o_2 = 2400$ ($\Delta U = -u(2400)$); consequently the bias is negative and the first gamble appears inferior to the second ($\mathbb{E}[\Delta \hat{U}_{\tilde{q},2}^{IS}] = -152$). But when $z = 0$, then L_1 yielding $o_1 = 2500$ and L_2 yielding $o_2 = 0$ ($\Delta U = +u(2500)$) becomes the most over-weighted event making the first gamble appear superior ($\mathbb{E}[\Delta \hat{U}_{\tilde{q},2}^{IS}] = +4.1$). This explains why people’s preferences switch from the second to the first gamble as z switches from 2400 to 0.

Preference Reversals

When people first price a risky gamble and a safe gamble with similar expected value and then choose between them, their preferences are inconsistent almost 50% of the time: most people price the risky gamble higher than the safe one, but many of them nevertheless choose the safer one (Lichtenstein & Slovic, 1971). Utility-weighted sampling over-estimates the expected utility $\mathbb{E}_{p_{\text{risky}}} [u(O_{\text{risky}})]$ of high-payoff (risky) lotteries more than for low-payoff (safe) lotteries, because $\tilde{q}(o) \propto |u(o)|$. This explains why people price the risky gamble higher than the safe one. To choose between two lotteries our model estimates the expected utility difference, i.e. $\mathbb{E}_{p_{\text{safe,risky}}} [u(O_{\text{safe}}) - u(O_{\text{risky}})]$. In this estimation problem there are 2×2 rather than just 2 possible outcomes, and the importance of positive outcomes is counterbalanced by the importance of the negative outcome. As a result, utility-weighted sampling’s bias in favor of the riskier option is weaker in choice than in pricing: so weak that it is overwritten by the risk-aversion due to concavity (flattening) of the

Table 3: Utility-weighted sampling explains preference reversals.

utility of	$\mathbb{E}[U]$	$\mathbb{E}[\hat{U}_{\tilde{q},2}^{IS}]$
safer gamble ($L_s : \$1, p = 0.8$)	0.80	1
riskier gamble ($L_r : \$2, p = 0.4$)	0.72	1.8
choosing L_s over L_r	0.079	0.075

utility function (cf. Equation 14); see Table 3. This illustrates that utility-weighted sampling weights events differently depending on the problem to be solved.

General Discussion

Our resource-rational analysis of decision-making in high-risk situations suggested that people should decide by utility-weighted sampling (Equation 12). Utility-weighted sampling explains not only how we are able to make sensible decisions under high risk but also why we overestimate the frequency of extreme events and have inconsistent risk preferences (framing effects, the Allais paradox, and preference reversals).

While overestimation was previously explained by ‘availability’ (Tversky & Kahneman, 1973), our theory specifies what exactly the availability of events should correspond to—namely their importance distribution \tilde{q} (Equation 11)—and why this is useful. Inconsistent risk preferences were previously explained by *regret* (Loomes & Sugden, 1982) or *salience* (Bordalo, Gennaioli, & Shleifer, 2012). Like these explanations our theory assumes an amplified impact of large utility differences. While *regret theory* explains this amplification by altered subjective utilities, utility-weighted sampling and *salience theory* explain the amplification by altered probability-weighting. Despite this commonality, our account offers three advances over salience theory. First, we do not only describe the effect of utility on probability-weighting, but we also model the cognitive strategy that generates it. Second, our theory reconciles this seemingly irrational effect with rational information processing. Concretely, the resource-rational basis of the salience of a utility difference $\Delta U = u(O_1) - u(O_2)$ is the relative frequency with which it should be simulated, i.e. the importance distribution $\tilde{q}(\Delta u) \propto p(\Delta u) \cdot |\Delta u|$.⁴ Third, since our explanation instantiates a more general theoretical framework—resource-rationality—it can be extended to multi-alternative decisions, decisions from experience, and many other problems.

We have proposed a psychologically plausible cognitive strategy that approximates the normative—but intractable—decision-rule of expected utility theory in two simple steps: (i) generate example outcomes by biased mental simulation, and (ii) count how many of them are positive. Interestingly, these steps resemble two *fast-and-frugal* heuristics (Gigerenzer, 2008b): Biased mental simulation (stochastically) considers the most important consequence first—like *take-the-best*—and binary choices are made by tallying if there are more positive than negative simulated outcomes—as in the *tallying heuristic*. The fact that we derived this strategy

⁴This definition satisfies two of Bordalo et al.’s (2012) three axioms of salience.

as a resource-efficient approximation to normative decision-making (resource-rational analysis) sheds light on why fast-and-frugal heuristics work and how they can be generalized to harder problems (cf. Lieder, Griffiths, & Goodman, 2013). Future research will also compare the rationality and descriptive accuracy of heuristics derived by resource-rational analysis to established heuristics, decision-by-sampling (Stewart, Chater, & Brown, 2006), and other models of risky choice (Erev et al., 2010). Our theory makes three novel predictions:

1. The *probability-weighting function* (Tversky & Kahneman, 1992) depends on the ratio of the outcomes' utilities.
2. As time pressure or *cognitive load* increase, people's risk preferences become more inconsistent.
3. The more concave a person's utility function, the less she will overestimate and over-weight extreme events.

In conclusion, utility-weighted sampling is a promising rational process model of probability judgment and decision-making. This strategy works not despite but because it is biased (cf. Gigerenzer & Brighton, 2009). Biased minds can not only make better inferences but also better decisions. However, our results highlight a tension between good inference and good decision-making: Bounded sample-based agents should over-sample extreme events even though this leads to overestimation (bad inference), and people appear to do the same. In more general terms, the human mind should, and appears to, sacrifice the rationality of its beliefs (*theoretical rationality*) for the rationality of its actions (*practical rationality*, Harman, 2013), because limited computational resources necessitate tradeoffs. Concretely, our analysis suggested that the availability bias is a manifestation of resource-rational decision-making. This conclusion supports the emerging view that cognitive biases are a window on resource-rational computation rather than a sign of irrationality (Lieder, Griffiths, & Goodman, 2013; Lieder, Goodman, & Griffiths, 2013). Being biased can be resource-rational.

Acknowledgments. This work was supported by grant number N00014-13-1-0341 from the Office of Naval Research and grant number RO1 MH098023 from the National Institutes of Health.

References

- Bordalo, P., Gennaioli, N., & Shleifer, A. (2012). Salience theory of choice under risk. *Q. J. Econ.*, *127*(3), 1243–1285.
- Denison, S., Bonawitz, E., Gopnik, A., & Griffiths, T. (2013). Rational variability in childrens causal inferences: The sampling hypothesis. *Cognition*, *126*(2), 285–300.
- Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., ... Lebiere, C. (2010). A choice prediction competition: Choices from experience and from description. *J. Behav. Decis. Making*, *23*(1), 15–47.
- Fiser, J., Berkes, P., Orbán, G., & Lengyel, M. (2010). Statistically optimal perception and learning: from behavior to neural representations. *Trends Cogn. Sci.*, *14*(3), 119–130.
- Geweke, J. (1989). Bayesian inference in econometric models using Monte Carlo integration. *Econometrica*, *57*(6), 1317–1339.
- Gigerenzer, G. (2008a). *Rationality for mortals: How people cope with uncertainty*. Oxford University Press.
- Gigerenzer, G. (2008b). Why heuristics work. *Perspect. Psychol. Sci.*, *3*(1), 20–29.
- Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Top. Cogn. Sci.*, *1*(1), 107–143.
- Griffiths, T. L., Lieder, F., & Goodman, N. D. (in review). Levels of analysis between the computational and the algorithmic. *Top. Cogn. Sci.*
- Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychol. Sci.*, *17*(9), 767–773.
- Hammersley, D. C., & Handscomb, J. M. (1964). *Monte Carlo methods*. London: Methuen & Co Ltd.
- Harman, G. (2013). Rationality. In H. LaFollette, J. Deigh, & S. Stroud (Eds.), *International Encyclopedia of Ethics*. Hoboken: Blackwell Publishing Ltd.
- Hertwig, R., Pachur, T., & Kurzenhäuser, S. (2005). Judgments of risk frequencies: tests of possible cognitive mechanisms. *J. Exp. Psychol.-Learn. Mem. Cogn.*, *31*(4), 621.
- Hobson, C. J., & Delunas, L. (2001). National norms and life-event frequencies for the revised social readjustment rating scale. *Int. J. Stress Manag.*, *8*(4), 299–314.
- Hobson, C. J., Kamen, J., Szostek, J., Nethercut, C. M., Tiedmann, J. W., & Wojnarowicz, S. (1998). Stressful life events: A revision and update of the social readjustment rating scale. *Int. J. Stress Manag.*, *5*(1), 1–23.
- Lichtenstein, S., & Slovic, P. (1971). Reversals of preference between bids and choices in gambling decisions. *J. Exp. Psychol.*, *89*(1), 46–55.
- Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M., & Combs, B. (1978). Judged frequency of lethal events. *J. Exp. Psychol. Hum. Learn. Mem.*, *4*(6), 551.
- Lieder, F., Goodman, N. D., & Griffiths, T. L. (2013). *Reverse-engineering resource-efficient algorithms* [Paper presented at NIPS-2013 Workshop Resource-Efficient ML, Lake Tahoe, USA].
- Lieder, F., Griffiths, T. L., & Goodman, N. D. (2013). Burn-in, bias, and the rationality of anchoring. In P. Bartlett, F. C. N. Pereira, L. Bottou, C. J. C. Burges, & K. Q. Weinberger (Eds.), *Adv. neural inf. process. syst.* 26. Red Hook: Curran Associates, Inc.
- Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *Econ. J.*, *92*(368), 805–824.
- Rasmussen, B. K., Jensen, R., Schroll, M., & Olesen, J. (1991). Epidemiology of headache in a general population prevalence study. *J. Clin. Epidemiol.*, *44*(11), 1147–1157.
- Rothman, A. J., Klein, W. M., & Weinstein, N. D. (1996). Absolute and relative biases in estimations of personal risk. *J. Appl. Soc. Psychol.*, *26*(14), 1213–1236.
- Shi, L., & Griffiths, T. (2009). Neural implementation of hierarchical Bayesian inference by importance sampling. In Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, & A. Culotta (Eds.), *Adv. neural inf. process. syst.* 22 (pp. 1669–1677). Red Hook: Curran Associates, Inc.
- Shi, L., Griffiths, T., Feldman, N., & Sanborn, A. (2010). Exemplar models as a mechanism for performing Bayesian inference. *Psychon. Bull. Rev.*, *17*(4), 443–464.
- Stewart, N., Chater, N., & Brown, G. D. (2006). Decision by sampling. *Cogn. Psychol.*, *53*(1), 1–26.
- Sunstein, C. R., & Zeckhauser, R. (2011). Overreaction to fearsome risks. *Environ. Resour. Econ.*, *48*(3), 435–449.
- Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cogn. Psychol.*, *5*(2), 207–232.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *J. Risk Uncertain.*, *5*(4), 297–323.
- Von Neumann, J., & Morgenstern, O. (1944). *The theory of games and economic behavior*. Princeton: Princeton university press.
- Vul, E., Goodman, N. D., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? Optimal decisions from very few samples. *Cognitive Sci.*, 1–39.
- Zabaras, N. (2010). *Importance sampling* (Tech. Rep.). Cornell University. Retrieved from <http://mpdc.mae.cornell.edu/Courses/UQ/ImportanceSampling.pdf>