



City Research Online

City, University of London Institutional Repository

Citation: Hotaling, J., Jarvstad, A., Donkin, C. & Ben, N. (2019). How to Change the Weight of Rare Events in Decisions from Experience. *Psychological Science*, 30(12), pp. 1767-1779. doi: 10.1177/0956797619884324

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/22647/>

Link to published version: <https://doi.org/10.1177/0956797619884324>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

How to Change the Weight of Rare Events in Decisions from Experience

Jared M. Hotaling^{1*}, Andreas Jarvstad^{2,3*}, Chris Donkin¹, Ben R. Newell¹

¹University of New South Wales

²University of Oxford

³City, University of London

*Equal contribution

Please send all correspondence to:

Jared Hotaling

School of Psychology

University of New South Wales

Sydney 2052, Australia

jaredhotaling@gmail.com

Abstract

When making risky choices, two kinds of information are crucial: outcome values and outcome probabilities. Here, we demonstrate that the juncture at which value and probability information is provided has a fundamental effect on choice. Across four experiments involving 489 participants, we compare two decision making scenarios: one where value information is revealed during sampling (Standard), and one where value information is revealed after sampling (Value-Ignorance). On average, participants made riskier choices when value information was provided after sampling. Moreover, parameter estimates from a hierarchical Bayesian implementation of cumulative prospect theory suggested that participants overweighted rare events when value information was absent during sampling, but showed no overweighting in the Standard condition. This suggests that the impact of rare events on choice relies crucially on the timing of probability and value integration. We provide paths towards mechanistic explanations of our results based on frameworks which assume different underlying cognitive architectures.

Keywords: decisions from experience, description-experience gap, decision making, sampling, cumulative prospect theory; hierarchical Bayesian modeling.

When making risky choices, two kinds of information are crucial: outcome values and outcome probabilities (Bernoulli, 1954; von Neumann & Morgenstern, 1947). While some decisions involve choosing between options with described values and probabilities (investing in stocks based on earnings projections), others require learning about options through experience (finding a favorite restaurant in a new city). There is considerable evidence that these types of decisions differ systematically: *the description-experience gap* (see Wulff, Mergenthaler-Canseco, & Hertwig, 2018 for a recent review).

Decisions from description typically result in choices that imply overweighting of rare events (Kahneman & Tversky, 1979; Rieskamp, 2008), whereas decisions from experience do not (Camilleri & Newell, 2011; Hertwig, Barron, Weber, & Erev, 2004; Lejarraga & Gonzalez, 2011; Rakow & Newell, 2010; Yechiam & Busemeyer, 2006). Much has been made of this key difference in behavior over the past decade (Hertwig, Hogarth, & Lejarraga, 2018), yet its causes remain elusive (e.g. Glöckner, Hilbig, Henninger, & Fiedler, 2016; Wulff et al., 2018). Part of the problem may stem from a focus on the ‘gap’ – which is not, in itself, behaviour to-be-explained – rather than an attempt to predict and explain the behaviors that exist along the continuum between decisions from description and experience (e.g. Camilleri & Newell, 2013; Rakow & Newell, 2010). Here we focus on an aspect of information provision in risky choice that has received little attention in the literature.

We propose that an important aspect of the decision making process lies in the juncture at which outcome values and probabilities are integrated (Jarvstad, Hahn, Rushton, & Warren, 2013, p. 16275). We explore this integration by comparing choices under two sampling-based decisions from experience conditions (Figure 1A,B). Participants directed a ‘robot’ to sample balls from boxes. In the Standard condition (Standard, Figure 1A), the reward magnitudes were presented at the time of sampling, as in standard decision from experience tasks (e.g. Hertwig et al., 2004). In the Value-Ignorance condition (Value-Ignorance, Figure 1A), participants learned about outcome probabilities during the sampling phase, but not reward magnitude, which was revealed at the choice stage. The

participants' goal was to learn which of two boxes they would prefer to draw from 'for real' at the choice stage. Crucially, participants in both conditions learned about identical outcome distributions with deterministic reward magnitudes – the sole difference was when value information was available (c.f. Hadar & Fox, 2009, and see Discussion).

In the Value-Ignorance condition, participants can, by sampling each box, gradually learn the probability of drawing a blue ball (Hogarth & Einhorn, 1992; Sutton & Barto, 1998). At choice, they might recall the probability and integrate it with value information. Thus, participants might represent two pieces of information for each box: the probability of drawing a rewarded ball and its reward. An identical learning strategy can also be applied in the Standard condition. However, unlike in the Value-Ignorance condition, a simpler approach is possible. Participants might represent only the average outcome of each box. This value-updating strategy does not require that participants separately learn the probability of drawing a blue ball, and similar mechanisms have been proposed to account for decisions from experience (Hertwig, Barron, Weber, & Erev, 2006).

If people use a single mechanism in both the Standard and Value-Ignorance conditions, we should observe similar choices across conditions. However, if manipulating the juncture at which participants receive value information leads them to use different mechanisms, we should observe divergent choices across conditions. As is typical, we first examine the proportion of choices in favor of the riskier alternative in each condition. However, we want to understand the aspects of the decision mechanisms that produce different risk preferences, and so we use cumulative prospect theory (CPT; Tversky & Kahneman, 1992) as a measurement model to estimate the effect of our manipulations on people's treatment of probability and value information. This model has been used to measure risk preferences in a number of domains (e.g. Glöckner et al., 2016).




Figure 1. Robot-sampling task A) Main design. Each trial was composed of two phases: a sampling phase (top row), and a choice phase (bottom row). For a given sample, participants could observe either a blue ball, or a red ball. Red balls were worth \$0 and blue balls were worth some reward. In the Standard condition, the value of the blue ball was revealed during sampling (left column). In the Value-Ignorance condition, the value of the blue ball was not revealed until the choice stage. Thus, during sampling under value-ignorance, the probability of drawing a blue ball could be learned, but not its value. B) Example of a sampling sequence. Once a box was selected for sampling (having been clicked), an animation showed the box shaking (to 'mix' the balls), then a 'robot arm' reached down and grabbed a ball, lifted it up to reveal it and then dropped it back down again (illustrating sampling with replacement). Participants were required to sample each box a set number of times but were free to sample in any order.

Across four experiments, we find that behavior at the choice stage differs systematically across conditions. Specifically, participants made a higher proportion of riskier choices in the Value-Ignorance condition than the Standard condition. We unpack this result with a hierarchical Bayesian implementation of CPT, which shows that participants overweighted small probabilities in Value-Ignorance, but used neutral weighting in Standard. Moreover, the observed neutral weighting of outcomes in the Standard condition was disrupted when rarer events were perceptually highlighted during sampling (Experiments 3 and 4), with this manipulation causing these events to be overweighted

like in the Value-Ignorance condition¹. We argue that these results provide clear evidence for the operation of different decision mechanisms in the Standard and Value-Ignorance conditions. In the Discussion we outline contrasting mechanistic explanations based on frameworks assuming different cognitive architectures.

Method

Ethics

Ethical approval for all experiments was obtained through the institutional review boards of the School of Psychology at the University of New South Wales (UNSW).

Participants

All participants were UNSW students and received course credit plus a monetary bonus (\$0 - \$21) based on a randomly selected trial. Eighty² participated in Experiment 1, 83 (45 female, age 17-42, $M = 19.44$, $SD = 2.95$) participated in Experiment 2, 149 (99 female, age 18-53, $M = 22.93$, $SD = 4.63$) participated in Experiment 3, and 177 (106 female, age 18-58, $M = 20.49$, $SD = 3.92$) participated in Experiment 4. Within each experiment, participants were randomly distributed across conditions.

Procedure

After giving informed consent, participants were placed in a computer booth where they read the following instructions:

"In this task you will draw balls from pairs of virtual boxes. In each box, there are 100 balls, some of which are blue and some of which are red. Blue balls are associated with reward and red balls are not (reward for a red ball = \$0)."

¹ The lowest probability outcomes ranged from approximately .1 to .4 across problems (see Method), however, for simplicity we define 'rare' events as those that occur with a probability of less than .5.

² Demographic information was not collected in Experiment 1.

Participants first completed a practice trial to familiarize them with the task (Figure 1) before beginning the experimental trials. Participants were instructed that each trial involved a new pair of boxes and that they would have to learn anew the values and proportions of balls within each box. To emphasize that boxes were different across trials, each box had a unique color. On each trial, participants were required to sample the entire sample set for both alternatives before making their choice. Participants were able to sample freely (e.g., alternating between boxes, sampling exhaustively from one then the other, etc.), with sampling disabled once the entire set from each box had been seen. Participants were instructed that one of their choices would be used to draw a ball for a bonus payment at the end of the experiment³.

Importantly, the samples that participants observed matched the true underlying probabilities of each outcome, thus mitigating other factors that may give rise to illusory ‘gaps’ (e.g., biased sampling and reliance on small samples; Hau, Pleskac, Kiefer, & Hertwig, 2008; Hertwig & Pleskac, 2010; Rakow, Demes, & Newell, 2008).

Materials and Design

Experiments 1 and 2. The values of boxes (monetary gambles) were determined as follows. Each choice alternative was defined by a reward value, v (range \$1 to \$21), and a probability of reward, π (range .083 to 1). With these values we created a sample set for each alternative representing the proportion of red and blue balls. The size of the sample set ranged from 10 to 12 and the frequency of rewards was determined by π .

Red balls were always worth \$0. The value of blue balls was fixed within each box, but varied across boxes and trials (\$1 to \$21). For example, the value of a blue ball may be \$16 in the left box and

³ In Experiment 1 we collected responses to described versions of the gambles from a separate group of participants. In Experiment 2 all participants completed a description phase following the sampling phase in which they responded to the same set of gambles. We do not consider those data further in this report. No description-data were collected in Experiments 3 and 4. See Supplementary Materials for further information.

\$2 in the right box. In the Standard condition (Figure 1A), each sampled ball was labeled with the outcome value. In the Value-Ignorance condition, sampled balls were not labeled with values, though the instruction indicated that red balls were worth \$0 and the blue balls were worth some reward. Participants could therefore learn the relative proportions of balls in each box, but not their values, with value revealed in the choice phase (Value-Ignorance Figure 1A).

Choice pairs were constructed with the goal of exposing participants to a range of problems. For example, problems could involve zero, one, or two risky options (i.e. $\pi < .5$), and equal or unequal expected values. To better understand the task, consider an example trial involving a riskier option on the left offering a 10% chance of winning 16 points, and a safer option on the right offering an 80% chance of winning 2 points. While sampling from the riskier box participants would observe one blue ball and nine red balls. From the safer box they would sample eight blue balls and two red balls. In Experiment 1 each participant received the same underlying decision problems, and in Experiment 2 five of the decision problems were the same across participants and six of the problems were randomly generated. Experiment 2 also involved minor graphical changes, such a larger font size. In all other aspects, Experiment 1 and 2 were identical (see Supplementary Materials for decision problem details).

Experiments 3 and 4. In Experiment 3 we introduced a new manipulation motivated by our finding of the differential influence of rare events in Experiments 1 and 2 (see Results). Specifically, we sought to examine whether the relatively lower weighting placed on rare events in the Standard, compared to the Value-Ignorance condition, could be increased by emphasizing the rare outcome to participants. To do this, we introduced a Salience condition in which, during sampling, some balls were highlighted. When a highlighted ball was drawn, a tone played and the ball flashed for approximately 700ms before being returned to the box as usual.

The highlighting occurred whenever participants sampled the rare event for the riskier alternative. This resulted in two types of problems. For Type 1 (*best-outcome salient*) problems, salience

highlighted a rare reward, and was expected to increase the likelihood of choosing the risky option. For Type 2 (*worst-outcome salient*) problems, salience highlighted an outcome of \$0, and could be expected to decrease the likelihood of choosing the risky option. The No-Salience conditions replicated the conditions in Experiments 1 and 2, and no highlighting occurred. Each participant completed the same set of fourteen Type 1 trials and six Type 2 trials in a random order. Experiment 4 was a preregistered replication of Experiment 3 (details can be found at <https://osf.io/ry75j>). See the Supplementary Materials for the specific gambles used in Experiments 3 and 4.

Statistical analyses

Because each experiment used nearly identical methods, procedures, and designs, and because we expected small effect sizes given our well-controlled design (see Camilleri & Newell, 2011; Glöckner et al., 2016; Wulff et al., 2018), and to make use of the large amount of data collected across experiments – both in terms of number of participants and different choice options – we report and analyze them together. In designing each experiment, we selected sample sizes such that pooling across experiments with hierarchical Bayesian analyses would yield a pooled sample size sufficient to demonstrate reliable evidence for the effects we observe, and one considerably larger than in previous decisions from experience studies (see Wulff et al., 2018). We pool in two different ways: pooling across Experiment 1-4 for Standard vs Value-Ignorance contrasts, and pooling across Experiment 3-4 to examine the effect of the salience manipulation (absent in Experiments 1-2).

Type 1 (Best-Outcome Salient) and Type 2 (Worst-Outcome Salient) gambles. In pooling these data, we treated Type 2 gambles differently. When considering the effect of value information (Figures 2, 4 and 5), we use all data from Experiments 1-4 because we expected the effect of value information to be the same for both Type 1 and Type 2 gambles. However, our salience manipulation carried the risk of introducing a demand characteristic whereby participants were encouraged to choose the riskier option, regardless of which outcome was highlighted. Type 2 problems therefore served as a

manipulation check because salience highlighted an outcome of \$0, rather than a rare reward. Since we expected the effect of salience to be different for Type 1 and Type 2 gambles, we analyzed only Type 1 problems when considering the effect of salience (Figures 3 and 6). In the Supplementary Materials, we show that the salience manipulation had no effect on choices for Type 2 problems.

Proportion of risky choices. We model the proportion of risky choices made by participant i in condition j as coming from a Binomial process with rate parameter θ_{ij} . We can think of our data across the four experiments as contributing different amounts of information to four conditions in a 2 (sampling) x 2 (salience) design. Because there was no salience manipulation in Experiments 1 and 2, these data only contribute to the No-Salience conditions. Experiments 3 and 4, on the other hand, contribute data to all four conditions. Finally, we assume that the rate parameters for each participant come from a group-level Normal distribution, $\theta_{ij} \sim N(\mu^{\theta_j}, \sigma^{\theta})$, truncated to be between 0 and 1, with a group-level mean, μ^{θ_j} , for each condition, and a variability parameter, σ^{θ} , that is shared across conditions for parsimony. We used Uniform(0,1) priors for each group-level mean, μ^{θ_j} , and the prior for the group-level precision was set so that $1/\sigma_{\theta_j}^2 \sim \text{Gamma}(.001, .001)$.⁴

We also examine the difference between the (population-level) proportion of risky choices in the Standard and Value-Ignorance conditions. We calculate the posterior distribution of the differences between μ^{θ} in the two sampling conditions, $\Delta\mu^{\theta}$. To calculate a Bayesian equivalent of a frequentist p-value, we evaluate the empirical cumulative density function of $\Delta\mu^{\theta}$ values at 0. This Bayesian p-value tells us how likely it is that the difference between μ^{θ} in the two sampling conditions is below zero. Extremely small or large p-values are thus associated with a high probability that participants behaved differently in Standard and Value-Ignorance conditions.

⁴ Note that the population mean parameters, μ^{θ} , in this hierarchical model is not equivalent to the average proportions typically used to summarize behavior. We use the hierarchical Bayesian approach here, because it outperforms the standard approach in most contexts (Gelman et al., 2013), but we present the standard measures in Tables S1-S3 in the Supplementary Materials.

Model-Based Analysis of Risk Preferences. We used a hierarchical Bayesian implementation of prospect theory (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992) to estimate participants' risk preferences. This analysis allows us to interpret our results in terms of psychologically interesting latent parameters, while naturally accounting for the variability observed across gambles, which is ignored in many traditional analyses. We can think of this as extending our pooled data analysis to include a psychological model of the θ parameters used earlier, accounting for participants' risk preferences with four latent parameters.

Prospect Theory was parameterized following Stott (2006): a power utility function $v(x) = x^\alpha$ capturing preferences for money, and a one-parameter probability weighting function $w(p) = e^{(-(-\ln p)^\gamma)}$ (Prelec, 1998) capturing preferences for probabilities. Choice options were modeled in terms of whether they appeared on the left or right of the screen, with the differences in prospect values between the left and right options given by: $\Delta_{prospect} = v(x)_{left}w(p)_{left} - v(x)_{right}w(p)_{right}$. We mapped this difference onto probability of choosing the left option using a two-parameter logistic function, $p_{left} = \frac{1}{1 + e^{-1 \kappa [\beta + \Delta_{prospect}]}}$. Higher values of sensitivity parameter, κ , indicate better discrimination between prospects (i.e., more deterministic choice). Bias parameter, β , captures the extent to which one side of the screen was favored irrespective of prospect value.

In a first analysis, we fit the model to the Standard and Value-Ignorance conditions, pooling data across Experiments 1-4. Participants' risk preferences, described by the four parameters in the model, were modeled as coming from four separate population-level distributions: $\alpha_i \sim e^{N(\mu^\alpha, \sigma^\alpha)}$, $\gamma_i \sim e^{N(\mu^\gamma, \sigma^\gamma)}$, $\kappa_i \sim e^{N(\mu^\kappa, \sigma^\kappa)}$, $\beta_i \sim N(\mu^\beta, \sigma^\beta)$. For α , γ , and κ , participant-level samples were drawn from exponentiated normal distributions. Priors for the population-level means were: $\mu^\alpha \sim U(-10, 10)$, $\mu^\gamma \sim U(-10, 10)$, $\mu^\kappa \sim U(-2, 10)$, $\mu^\beta \sim U(-10, 10)$, which are sufficiently broad to capture extreme distortions of money and probability and a wide range of logistic functions. Population-level parameters in log-space accounts for the fact that α and γ have qualitatively different meanings between 0 and 1,

and between 1 and $+\infty$. The priors for α and γ were symmetric around 0 to assign equal prior probability to under/overweighting of small probabilities and increasing/diminishing marginal utility respectively. The prior for k was not symmetric as discriminability increases monotonically with k . The priors for all population-level standard deviations were identical across parameters: e.g., $\sigma^\alpha \sim U(1 \times 10^{-5}, 6)$. Sampling was truncated by the respective prior range. For plotting purposes, we report best-fit population-level parameters in standard parameter space.

To evaluate the effect of the salience manipulation on risk preferences, we pooled data from Experiments 3 and 4, and estimated population level α and γ for each of the four conditions (Standard & Salience; Standard & No-Salience; Value-Ignorance & Salience; Value-Ignorance & No-Salience). Based on the Standard/Value-Ignorance analysis, for which we observed no condition-specific effects on κ and β (see Results), these parameters were modeled as coming from the same population-level distributions. In all other aspects, this model was identical to the model above.

Results

The left panel of Figure 2 shows our analysis of choice proportion data pooled across Experiments 1 to 4. According to the figure, participants who did not know the values associated with each outcome during sampling (Value-Ignorance condition) made more risky choices than participants for whom the value information was present during sampling. The right panel of Figure 2 plots the posterior distribution of the difference in the proportion of risky choices between the Standard and Value-Ignorance conditions. Here we see that the effect of value information on risky choice was reliable, with the posterior distribution of the difference in proportion risky choices between the Standard and Value-Ignorance conditions sitting almost entirely above zero ($p = .009$).

Figure 3 pools data from Experiments 3 and 4 to show the effects of value information and salience. In the top panel we see that in both the Salience and No-Salience conditions participant made

more risky choices if they did not know the values associated with outcomes during sampling. The bottom panel of Figure 3 shows that in the No Saliency condition, the effect of value information on risky choice was relatively large and reliable when there was no saliency manipulation, with the posterior distribution of the difference in proportion risky choices between the Standard and Value-Ignorance conditions again sitting almost entirely above zero ($p = .004$). When the saliency manipulation was present the difference between Standard and Value-Ignorance conditions was less robust (due to the saliency manipulation increasing risky choices in the Standard conditions), though there was some evidence of its reliability with the posterior distribution of the difference in proportion risky choices between Standard and Value-Ignorance conditions having most of its density above zero ($p = .08$).




Figure 2. Left: The median and central 95% of the posterior distributions for the population-level mean of the proportion of risky choices, θ , in the Standard and Value-Ignorance conditions. Right: Posterior distribution of the difference between population-level θ parameters in the Standard and Value-Ignorance conditions. Positive values on the horizontal axis reflect more risky choices when value information was absent during sampling (Value-Ignorance).

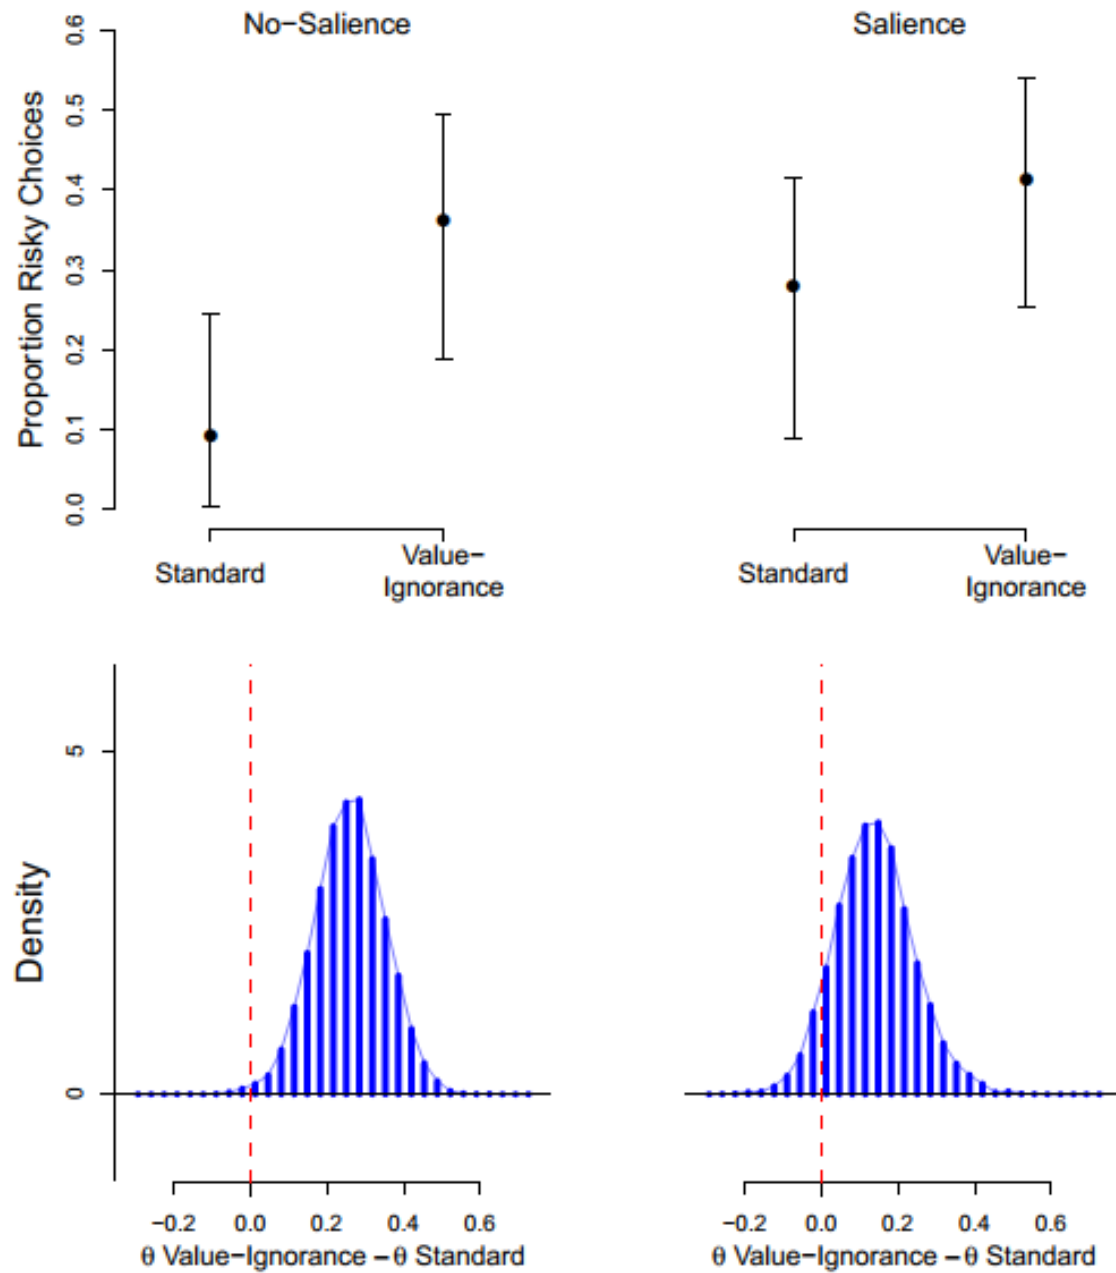


Figure 3. Top: The median and central 95% of the posterior distributions for the population-level mean of the proportion of risky choices, θ , in the Standard and Value-Ignorance conditions for the No-Saliency (left) and Saliency (right) conditions in Experiments 3 and 4. Bottom: Posterior distribution of the difference between population-level θ parameters in the Standard and Value-Ignorance conditions for the No-Saliency (left) and Saliency (right) conditions. Positive values on the horizontal axis reflect more risky choices when value information was absent during sampling (Value-Ignorance). Saliency refers to the perceptual highlighting of a drawn ball during sampling.

The previous analysis shows that participants made riskier choices when they had to wait until the choice phase to integrate probability and value (Value-Ignorance), compared to when they could integrate this information at the sampling phase (Standard). To see whether this pattern holds at the level of individual gambles, Figure 4 explores gamble-wise choices from Experiments 3 and 4 (which used a wide range of reward probabilities and had substantial data for each gamble). Specifically, Figure 4 plots the proportion of risky choices as a function of the probability of drawing a blue ball from the riskier box (i.e. the box with the lower probability of reward).

As can be seen, when the riskier option was unlikely ($p < .5$), participants in the Value-Ignorance condition made riskier choices than participants in the Standard condition. This result matches the overall effect on risky choice observed above. However, for gambles in which the riskier option was likely ($p > .5$), we find either no appreciable difference between conditions, or a trend towards *less* risky choices in the Value-Ignorance condition. This is the pattern one would expect if the Value-Ignorance condition induced Prospect Theory-like weighting of probability at the time of choice (overweighting of small probabilities, underweighting of large probabilities).




Figure 4. Gamble-wise proportion of risky choices as a function of the probability of drawing a blue ball from the riskier option. Each data point comes from a unique gamble in one of the two conditions. The plot uses data from Experiments 3 and 4, for which there was a wide range of reward probabilities and approximately 80 observations per gamble. The vectors indicate the direction of change in the Value-Ignorance condition relative to the Standard condition. X-axis values have been jittered to improve readability (avoiding overlap).

The previous analyses show that having to integrate probability and value at the time of choice induced riskier choices relative to the Standard condition, and that this average effect was mainly due to a strong preference for choosing the risky option when its probability was low (Figure 4). What these previous analyses do not show, is which of the two critical pieces of information were affected: value or probability. To quantify this, we used Prospect Theory as a measurement model, and estimated its parameters using hierarchical Bayesian methods.

Figure 5 shows the result of fitting Prospect Theory to the data, with the mean posterior utility (Panel A) and probability (Panel B) weighting functions for the Standard (red symbols) and Value-Ignorance (black symbols) condition respectively. Participants in both conditions ($\mu_{Standard}^{\alpha} =$

.71, $\mu_{Value-Ignorance}^{\alpha} = .59$) showed decreasing sensitivity to increasing monetary values (diminishing marginal utility) as expected (Figure 5A). Comparing across the two conditions, there was a small difference with those in the Standard condition showing more linear preferences ($p = .032$, Figure 5E).




Figure 5. Risk preferences for the Standard (red lines) and Value-Ignorance (black lines) conditions. Panel A: mean posterior utility functions mapping money onto subjective utility. Panel B: mean posterior probability weighting functions mapping probability onto subjective decision-weights. Panel C: posterior mean and standard deviation of the sensitivity parameter of the logistic function. Panel D: posterior mean and standard deviation of the side bias (positive = left-side bias). Panel E-H: posterior distributions of differences in population-level parameters between the Standard and the Value-Ignorance conditions. X-axis were scaled to facilitate comparison of effect sizes. For x-axis scaled to fit the distributions see Figure S1 in Supp. Mat.

In terms of probability weighting (Figure 5B), the Standard condition resulted in near-linear weighting and Value-Ignorance in overweighting of small probabilities ($\mu_{Standard}^Y = 1.02$, $\mu_{Value-Ignorance}^Y = .69$). This between-condition difference in probability weighting was highly reliable ($p = 1 \times 10^{-4}$, Figure 5F). This pattern corresponds well to the standard D-E gap, in that we observed near-neutral weighting for the standard experience-based condition (i.e. experience-like weighting, e.g., Camilleri & Newell, 2011), and overweighting of small probabilities when value information was absent at the time of sampling (i.e. description-like weighting, e.g., Kahneman & Tversky, 1979). There was very little evidence for differences between conditions in the parameters of the logistic choice function (Figure 5C-D, 4G-H).

In Experiments 3 and 4, we additionally manipulated the salience of the riskier option. We found that perceptually highlighting rare events during sampling increased the proportion of risky choices in the Standard condition but had little effect on sampling under value-ignorance (Figure 3). The corresponding prospect theory analyses are shown in Figure 6⁵.

Panels A and C of Figure 6 show the mean posterior functions for the Standard condition, and Panels B and D show the functions for the Value-Ignorance condition. The best-fit utility functions (top-row, Figure 6A-B) show that there were trends toward a lower sensitivity to value (more non-linear utility functions) in the Salience conditions (full lines) compared to the No-Salience conditions (dashed lines) for both the Standard ($\mu_{Standard,No-Salience}^\alpha = .75$, $\mu_{Standard,Salience}^\alpha = .58$) and Value-Ignorance ($\mu_{Value-Ignorance,No-Salience}^\alpha = .73$, $\mu_{Value-Ignorance,Salience}^\alpha = .61$) conditions (Standard $p = .057$, Value-Ignorance $p = .196$ respectively).

For probability weighting, the salience manipulation affected the Standard condition differently than the Value-Ignorance condition. For the latter, the manipulation had essentially no effect. In other

⁵ These analyses were performed on Type 1 gambles only. Analyzing both Type 1 and 2 gambles produced near-identical results (see Figure S2, Supp Mat.).

words, when outcome values were absent during sampling, highlighting a rare event did not induce a change in how learned probabilities were treated at the time of choice ($\mu_{Value-Ignorance, No-Salience}^{\gamma} = .84$, $\mu_{Value-Ignorance, Salience}^{\gamma} = .78$, $p = .333$, Figure 6D). For the Standard condition, however, highlighting the rewards for the riskier option had a dramatic impact on how the gambles were evaluated, resulting in an overweighting of low probability outcomes ($\mu_{Standard, No-Salience}^{\gamma} = 1.06$, $\mu_{Standard, Salience}^{\gamma} = .75$, $p = .005$, Figure 6C).

The population-level estimates of κ (sensitivity to differences between prospects) were comparable to those from the model used to contrast the Standard and the Value-Ignorance conditions ($\mu^{\kappa} = 2.2$, $\sigma^{\kappa} = 1.08$), and there was a similar trend towards a left-option bias ($\mu^{\beta} = .015$, $\sigma^{\beta} = .013$).

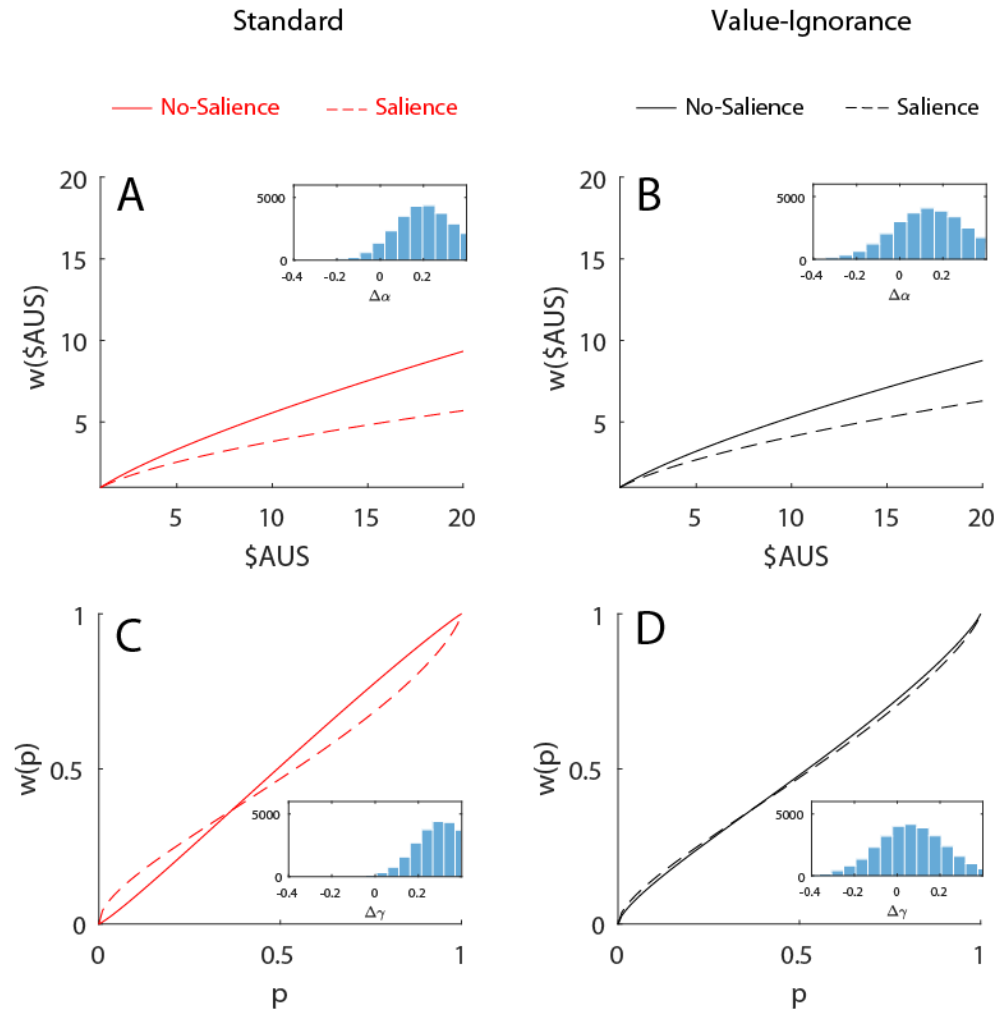


Figure 6. Value and probability weighting functions for the salience manipulation. Red lines represent risk preferences when value information was available at the time of sampling (Standard). Black lines represent risk preferences when value information was revealed at the time of choice (Value-Ignorance). The functions show best-fit population-level utility functions (Panel A,B) and probability weighting functions (Panel C,D), for the No-Salience (solid line) and the Salience (dashed line) conditions. Histograms show the posterior distribution of population-level differences in parameters (i.e., No-Salience - Salience).

General Discussion

Our findings show that subtle changes to an experience-based decision task can substantially affect both choice behavior and risk preferences. By manipulating whether outcome values were present while participants learned the probabilities of receiving a reward, we produced effects

reminiscent of the much-publicized description-experience gap. Across four experiments we observed riskier choices, on average, when outcome values were absent during sampling, compared to when they were present. This average increased preference for risk was driven by situations in which there was a small probability of the best outcome (see Figure 4). These patterns suggest that having to integrate value and probability information at the time of choice leads people to treat rare events differently.

The logic behind this conclusion is based on the observation that the two versions of our task can be solved using different methods. In the Value-Ignorance condition, participants must somehow integrate probability information – which they learned during sampling – with value information presented on the choice screen. On the other hand, participants in the Standard condition can avoid this direct integration by learning the expected reward produced by each option during sampling. The behavioral differences we observe across conditions indicate that our manipulation did affect participants' decision processes, though in subtle ways depending on the exact combinations of probabilities and rewards (see Figure 4). To better examine this interpretation, we used a hierarchical Bayesian CPT model to analyze the risk preferences underlying participants' choices. Here too we found compelling evidence for a difference between conditions, most notably in people's probability weighting⁶. Those in the Value-Ignorance condition showed overweighting of rare events, while those in the Standard condition did not. These results emphasize the importance of mapping and measuring behavior on the continuum between decisions from description and experience (Camilleri & Newell, 2013; de Palma, et al., 2014; Rakow & Newell, 2010).

Toward mechanistic explanations

Our results raise important questions regarding psychological mechanism. Why does the weighting of rare events increase when outcome values are presented only at choice? How do

⁶ We found some evidence that utility functions were affected by our manipulations, but we refrain from discussing these until more robust evidence is found.

manipulations of outcome salience influence choice and risk preference? Several possibilities exist, but we focus on two potential frameworks of explanation.

Reinforcement-learning. Our results could be cast in the form of Reinforcement-learning (RL). Imagine an RL agent making choices in the Standard condition. After sampling a ball from Box A, the value for Box A is updated according to $v_A^{t+1} = v_A^t + \delta * (o_A^t - v_A^t)$, where v_A^t is the estimated value at sample t, o_A^t is the value of the sampled ball, and δ is a learning rate parameter between 0 and 1. After completing sampling, the agent compares the estimated value of each option and chooses accordingly. Contrast this mechanism with that of another agent in the Value-Ignorance condition. This agent learns the probability of receiving a reward from Box A according to $p_A^{t+1} = p_A^t + \delta * (c_A^t - p_A^t)$, where p_A^t is the estimated probability of a reward at sample t and c_A^t is the outcome of the sample (1 if a reward, otherwise 0). On the choice screen, the value of a reward from Box A is revealed and the RL agent integrates this value with the learned probability and would thus – unlike the box-averaging agent – explicitly represent both value and probability. These different methods of combining value and probability – value-updating in Standard, and direct integration in Value-Ignorance – may give rise to different behavior. For instance, if the explicit representation of value and probability triggers cognitive mechanisms responsible for CPT-like preferences (as seen in decisions from description), the models can produce the behavior we see in our data.

The salience effect could also be explained within an RL framework: modeled as a multiplicative gain on the impact of highlighted outcomes, where highlighting ‘boosts’ value signals leading to an overestimation of the riskier option. This could also explain the absence of a salience effect for Type 2 (worst-outcome salient) problems (see Section 2 in Supplementary Materials) because a multiplicative gain would have no impact on a \$0 outcome. Future studies could test this by manipulating the value of the lower outcome.

Memory encoding and retrieval. The allocation of attentional resources provides another mechanistic account. In the Value-Ignorance condition, participants must retrieve information about the frequency of outcomes from memory in order to make a choice. Our results suggest that when participants retrieve this information, they place additional attention (and higher decision weight) on the rare outcome. This is also supported by the observation that prompting decision makers to mentally “repack” experienced events (sampled in the absence of outcome knowledge) leads to choices more consistent with decisions from description (Fox et al, 2013, unpublished manuscript cited in de Palma et al., 2014).

A related explanation suggests that the mere presentation of outcome values at the choice phase in the Value-Ignorance condition might lead to more equal emphasis being placed on those outcomes than is warranted. This argument, similar to that proposed by Erev, Glozman, and Hertwig (2008) invokes the notion of a *propositional* (symbolic) *representation* (e.g., blue ball = \$2 vs. blue ball = \$16) leading to greater attentional allocation to the rare outcome in memory. In contrast, when outcomes and frequencies are learned simultaneously, Erev et al. (2008) suggest an *analogical representation* is formed (e.g., 0, 0, 0, 0, 16, 0, 0, 16, 0) from which the “frequency of the option’s events can be read off directly” (Hertwig, 2016, p. 258). This ‘read-off’ of frequencies from memory would lead to the near linear weighting we observe in the Standard condition.

Allocation of attentional resources can also explain the impact of outcome salience on choice and risk preference. In the Standard condition, highlighting rare outcomes caused overweighting of rare events, and brought risk preferences more in line with those in Value-Ignorance. This could be explained – in keeping with the mere presentation proposal of Erev et al. (2008) – by an analogical representation in which rare outcomes are promoted in the mental sample (e.g., 0, 0, 0, 0, **16**, 0, 0, **16**, 0) during encoding, leading to greater attention and higher decision weighting at the time of choice. The absence of a salience effect under Value-Ignorance suggests that events must be tied to outcome values to

receive this promotion. That the effect would also disappear for Type 2 problems (see Section 2 in Supplementary Materials) likewise indicates the importance of outcome values. Perhaps participants view receiving \$0 as a ‘non-event’, rendering it immune to attentional distortions.

Consistency with existing literature and a way forward. Although we observed consistent effects across all experiments (see Supplementary Materials), our results highlight the importance of statistical power, experimental control, and good measurement models when studying choice behavior. We analyze data from 489 participants choosing between many different gambles, and use a hierarchical Bayesian CPT model to understand how behavior varied across trials, conditions, and participants (see also Glöckner et al., 2016). Our findings show that the parameters of CPT support stronger, more psychologically grounded inferences, compared to the parameters from our statistical models.

Methodological considerations may also shed light on discrepancies between our results and those reported by Hadar and Fox (2009). As part of a larger study focused on the under-sampling of rare events, they contrasted conditions in which reward information was revealed at the time of choice with more standard experience-based conditions, but found no differences. Perhaps one reason we observed differences was because, unlike Hadar and Fox, we ensured that the sample of outcomes participants observed were representative of the properties of the gambles, thus eliminating noise arising from sampling error. Another important factor is statistical power. With our much larger sample size, we observed consistent but modest effects that may have gone undetected in Hadar and Fox (2009).

Our results provide new insights into the relationships between task demands, experimental procedures, risk preferences, and decision processes. Using subtle manipulations – adjusting when value information was presented – we found notable differences in the way that value and probability (frequency) information were combined. These results offer novel empirical and modeling pathways for investigating how risk preferences develop in other domains (e.g., losses) across the entire description-

experience continuum. At present, frameworks assuming very different cognitive architectures provide plausible accounts of our results, auguring fertile avenues of future investigation to distinguishing between them.

Research disclosure statements:

All data exclusions have been reported and justified in the Method sections. All independent variables and manipulations, whether successful or failed, have been reported in the Method sections. All dependent measures that were analyzed for this article's target research question have been reported in the Method sections.

Author contributions:

AJ and BRN developed the original concept and design for Experiments 1 and 2. JMH, BRN, CD, and AJ developed Experiments 3 and 4. AJ wrote the software and JMH modified it for Experiments 3 and 4. JMH and BRN supervised data collection. AJ, JMH, CD, and BRN analyzed and interpreted the data. JMH wrote the first draft, and AJ, BRN, and CD provided critical revisions.

Open practices statement:

Methods, hypotheses, and analyses for Experiment 4 were preregistered on the Open Science Framework. All data and the analysis script are available at the following link: <https://osf.io/ry75j>.

Acknowledgements

JMH, CD, and BRN were supported by the Australian Research Council (DP160101186). AJ was supported by a British Academy Postdoctoral Fellowship (PF150005). We thank Garston Liang and Dominic Tran for help with data collection and Casimir Ludwig for valuable discussions.

References

- Bernoulli, D. (1954). Exposition of a new theory on the measurement of risk. [Specimen theoriae novae de mensura sortis]. *Econometrica*, 22, 23-26.
- Camilleri, A. R., & Newell, B. R. (2011). When and why rare events are underweighted: A direct comparison of the sampling, partial feedback, full feedback and description choice paradigms. *Psychonomic Bulletin & Review*, 18(2), 377-384.
- Camilleri, A. R., & Newell, B. R. (2013). Mind the gap? Description, experience, and the continuum of uncertainty in risky choice. In *Progress in brain research* (Vol. 202, pp. 55-71): Elsevier.
- de Palma, A., Abdellaoui, M., Attanasi, G., Ben-Akiva, M., Erev, I., Fehr-Duda, H., . . . Picard, N. (2014). Beware of black swans: Taking stock of the description–experience gap in decision under uncertainty. *Marketing Letters*, 25(3), 269-280.
- Erev, I., Glozman, I., & Hertwig, R. (2008). What impacts the impact of rare events. *Journal of Risk and Uncertainty*, 36(2), 153-177.
- Fox, C. R., Long, A., Hadar, L., & Erner, C. (2013). *Unpacking decisions from description and experience*. Unpublished manuscript, UCLA Anderson School of Management.
- Gelman, A., Stern, H. S., Carlin, J. B., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis*: Chapman and Hall/CRC.
- Glöckner, A., Hilbig, B. E., Henninger, F., & Fiedler, S. (2016). The reversed description-experience gap: Disentangling sources of presentation format effects in risky choice. *Journal of Experimental Psychology: General*, 145(4), 486-508.
- Hadar, L., & Fox, C. R. (2009). Information asymmetry in decision from description versus decision from experience. *Judgment and Decision Making*, 4(4), 317-325.
- Hau, R., Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008). The description–experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making*, 21(5), 493-518.
- Hertwig, R. (2016). Decisions from experience. In G. Keren & G. Wu (Eds.), *The Wiley Blackwell Handbook of Judgment and Decision Making* (Vol. 1, pp. 239-267). Oxford, UK: Blackwell Wiley.
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15(8), 534-539.
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2006). The role of information sampling in risky choice. In K. Fiedler & P. Juslin (Eds.), *Information sampling and adaptive cognition* (pp. 72-91). New York: Cambridge University Press.
- Hertwig, R., Hogarth, R. M., & Lejarraga, T. (2018). Experience and Description: Exploring Two Paths to Knowledge. *Current Directions in Psychological Science*, 27(2), 123-128.
- Hertwig, R., & Pleskac, T. J. (2010). Decisions from experience: Why small samples? *Cognition*, 115(2), 225-237.
- Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive psychology*, 24, 1-55.
- Jarvstad, A., Hahn, U., Rushton, S. K., & Warren, P. A. (2013). Perceptuo-motor, cognitive, and description-based decision-making seem equally good. *Proceedings of the National Academy of Sciences*, 110(40), 16271-16276.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision making under risk. *Econometrica*, 47, 263-292.
- Lejarraga, T., & Gonzalez, C. (2011). Effects of feedback and complexity on repeated decisions from description. *Organizational Behavior and Human Decision Processes*, 116(2), 286-295.
- Prelec, D. (1998). The probability weighting function. *Econometrica*, 66(3), 497-527.

- Rakow, T., Demes, K. A., & Newell, B. R. (2008). Biased samples not mode of presentation: Re-examining the apparent underweighting of rare events in experience-based choice. *Organizational Behavior and Human Decision Processes*, 106(2), 168-179.
- Rakow, T., & Newell, B. R. (2010). Degrees of uncertainty: An overview and framework for future research on experience - based choice. *Journal of Behavioral Decision Making*, 23(1), 1-14.
- Rieskamp, J. (2008). The probabilistic nature of preferential choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(6), 1446-1465.
- Stott, H. P. (2006). Cumulative prospect theory's functional menagerie. *Journal of Risk and Uncertainty*, 32(2), 101-130.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction* (Vol. 1). Cambridge: MIT press.
- 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. (1947). *Theory of games and economic behavior*. Princeton, New Jersey: Princeton University Press.
- Wulff, D. U., Mergenthaler-Canseco, M., & Hertwig, R. (2018). A meta-analytic review of two modes of learning and the description-experience gap. *Psychological Bulletin*, 144(2), 140-176.
- Yechiam, E., & Busemeyer, J. R. (2006). The effect of foregone payoffs on underweighting small probability events. *Journal of Behavioral Decision Making*, 19(1), 1-16.