



UNIVERSITY OF LEEDS

This is a repository copy of *Simultaneous Underweighting and Overestimation of Rare Events: Unpacking a Paradox*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/142931/>

Version: Accepted Version

Article:

Szollosi, A, Liang, G, Konstantinidis, E orcid.org/0000-0002-4782-0749 et al. (2 more authors) (2019) *Simultaneous Underweighting and Overestimation of Rare Events: Unpacking a Paradox*. *Journal of Experimental Psychology: General*, 148 (12). pp. 2207-2217. ISSN 0096-3445

<https://doi.org/10.1037/xge0000603>

© American Psychological Association, 2019. This paper is not the copy of record and may not exactly replicate the authoritative document published in the APA journal. Please do not copy or cite without author's permission. The final article is available, upon publication, at <http://www.doi.org/10.1037/xge0000603>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

© 2019, American Psychological Association. This paper is not the copy of record and may not exactly replicate the final, authoritative version of the article. Please do not copy or cite without authors' permission. The final article will be available, upon publication, via its DOI: 10.1037/xge0000603

Simultaneous Underweighting and Overestimation of Rare Events: Unpacking a Paradox

Aba Szollosi¹, Garston Liang¹, Emmanouil Konstantinidis^{1,2}, Chris Donkin¹, Ben R. Newell¹

¹School of Psychology, University of New South Wales, Sydney, Australia

²Centre for Decision Research, University of Leeds, Leeds, England

Author Note

This work was supported by the Australian Research Council (DP140101145; DP160101186). Portions of the data were published in the Proceedings of the 39th Annual Meeting of the Cognitive Science Society (Liang, Konstantinidis, Szollosi, Donkin, & Newell, 2017), and were presented at the 58th Annual Meeting of the Psychonomic Society and at the 39th Annual Conference of the Society for Judgment and Decision Making. All data and the analysis script are available in the Online Supplemental Materials and at the following link: <https://osf.io/3c8tf/> All authors contributed to the conceptualization and the design of the experiments. GL and AS collected the data. AS, CD, and EK analyzed the data. AS wrote the paper. BRN provided critical revisions. All authors approved the final version of the article.

Corresponding Author: Aba Szollosi, School of Psychology, University of New South Wales, Sydney, NSW 2052, Australia, e-mail: aba.szollosi@gmail.com

Abstract

We investigated previous findings suggesting a paradoxical inconsistency of people's beliefs and choices: When making decisions under uncertainty, people seem to both overestimate the probability of rare events in their judgments and underweight the probability of the same rare events in their choices. In our re-examination, we found that people's beliefs *are consistent* with their decisions, but they do not necessarily *correspond* with the environment. Both overestimation and underweighting of the rare event seemed to result from (most, but not all) participants' mistaken belief that they can infer and exploit sequential patterns in a static environment. In addition, we found that such inaccurate representations can be improved through incentives. Finally, detailed analysis suggested a mixture of individual level response patterns, which can give rise to an erroneous interpretation of group-level patterns. Our results offer an explanation for why beliefs and decisions can appear contradictory and present challenges to some current models of decisions under uncertainty.

Keywords: risky choice, learning under uncertainty, judgment, underweighting, overestimation

Psychology has long been concerned with uncovering the sources of mistakes in people's decisions. Normative models of choice suggest that if people want to maximize their gains, their beliefs about the environment need to correspond with the actual environment and their decisions need to be consistent with these beliefs (Hammond, 2000). Thus, making mistakes can have two major sources: errors in the correspondence of people's representations of the environment with the actual environment, and errors in the consistency with which they use these representations to inform their decisions. Inconsistencies between people's beliefs and choices have received much attention in risky and uncertain decision making research, such as the non-linear probability and value weighting functions of Prospect Theory (Kahneman & Tversky, 1979). Although such models still assume some level of consistency (Fox & Tversky, 1998), two lines of research suggest that in situations with risky rare events, beliefs and behavior may completely diverge.

On the one hand, people seem to underweight rare events when they make choices under uncertain conditions (Hertwig, Barron, Weber, & Erev, 2004; Rakow & Newell, 2010; Regenwetter & Robinson, 2017; Wulff, Mergenthaler-Canseco, & Hertwig, 2018). Work on experience-based decision-making, in which people do not know the outcome probabilities a-priori but can learn them through repeated experience, shows that people consistently choose as if they underweight rare events (Camilleri & Newell, 2011; Hertwig et al., 2004; Lejarraga & Gonzalez, 2011; Newell & Rakow, 2007; Yechiam & Busemeyer, 2006). In contrast, another line of research observes overestimation in people's probability judgments of rare events (Erev, Wallsten, & Budescu, 1994; Fischhoff et al., 2000; Viscusi & Hakes, 2008).

Most interestingly perhaps, Barron and Yechiam (2009) found this dissociation within the same experiment. Participants in their study appeared to simultaneously underweight the same rare event in their choice but overestimate its occurrence in their probability judgments. Specifically, participants repeatedly chose between a safe option that always provided the same outcome (-3 points, $p = 1$) and a risky option that either provided a common (0 points, p

= .85) or a rare outcome (-20 points, $p = .15$). In a static setting such as this, since the outcomes are generated by random processes, people should learn the long-run outcome probabilities and then exploit the option with the largest expected pay-off. In this case, the pay-offs of both options were the same (-3 points), thus there was no normative reason to prefer either option over the other. Intriguingly, although participants overwhelmingly preferred the risky over the safe option in their choices, suggesting *underweighting* of the rare event, when they were asked about the probability of the rare event, their estimates averaged around .25, showing *overestimation* of the same rare event. These results suggest a dissociation between choices and probability judgments, because if .25 was the actual probability of rare-event occurrence, participants should have preferred the safe option over the (subjectively worse) risky option (as the subjective expected utilities would be -3 points and -5 points respectively¹).

But do people really behave inconsistently with their beliefs or can such a pattern also arise as a result of representations that do not correspond with the environment? Decades of research on the well-known Gambler's and Hot-hand fallacies have demonstrated that people are sensitive to spurious patterns in sequential choice tasks (e.g., Hahn & Warren, 2009). More generally, people were found to assume that the environment is dynamic even in static settings (Navarro, Newell, & Schulze, 2016). That is, they seem to believe that predictable, rather than random processes generate outcomes. Plonsky, Tedorescu, and Erev (2015) observed choice patterns consistent with assuming a dynamic environment in experience-based gambles, such as the one used in the study of Barron and Yechiam (2009). Specifically, they showed that people's choices follow a "wavy recency" pattern: after the observation of a rare outcome, people's choices reflect that they expect this event not to occur again for a

¹ Based on the aggregate probability estimate for the rare outcome, the expected utility (U) of the risky option is the sum of the possible outcome values (here -20 and 0) multiplied by their respective subjective probability estimates (here .25 and .75), that is, $U(risky) = -20 * .25 + 0 * .75 = -5$. We use this calculation to help describe the data, but we do not mean to imply that people calculate expected utilities this way to make decisions. See Supplementary Materials (section #5) for more detail.

short initial period, followed by a gradual increase in this expectation. Thus, the observation of a rare outcome elicits a gradually decreasing negative recency effect.

Such a dynamic representation of the environment may provide an explanation for Barron and Yechiam's (2009) findings. For example, if people assumed that the sequences of common and rare outcomes are predictable, it would be consistent with their beliefs to pick the risky option on the trials that they expect the common outcome to occur and to pick the safe option when they expect the rare outcome to occur. The apparent underweighting could therefore be the result of simply predicting the common outcome more often than the rare outcome. On the other hand, people's beliefs in a dynamic environment might increase the variance in the probability estimates, enhancing the regression effect that "draws" the more extreme values on the response scale (such as $p = .15$) towards less extreme values ($p = .50$, the middle of the response scale), and thus resulting in overestimation of the rare outcome probability (Erev et al., 1994; Fiedler & Unkelbach, 2014).

In the current study, we conducted three experiments that had two related aims. First, we were interested in whether the underweighting-overestimation pattern observed by Barron and Yechiam (2009) can arise from people's mistaken belief that the environment is dynamic rather than static. Although this pattern has been interpreted as an inconsistency of beliefs and choice, and therefore challenging belief-based accounts of decision-making (e.g., Fox & Tversky, 1998), we aimed to provide an empirical test of whether it is instead a result of a lack of correspondence between beliefs and the environment (Hammond, 2000).

Our second aim was to investigate whether inaccurate representations of the environment can be improved. Therefore, we manipulated the incentives for participants' probability estimation performance in an experience-based decision task, attempting to make people attend more to the relevant aspects of the environment (Camerer, Hogarth, Budescu, & Eckel, 1999). In three between-subjects conditions, participants were either i) not incentivized on their probability estimation but incentivized on their choices; ii) they were

incentivized on both their estimation and their choices; or iii) they were incentivized on their probability estimation but they did not have to make choices. We expected that these manipulations would respectively increase the effort that people allocate to building representations of the probabilities. We also expected that this increase in effort would in turn increase the accuracy of participants' estimates by decreasing overestimation².

The task used by Barron and Yechiam (2009) presents a unique opportunity to investigate these questions in detail. We adopted this task, because unlike the vast majority of research on risky and uncertain choice, it invites participants to make repeated choices while simultaneously estimating the probability of an outcome. As such, we have a more direct window into how participants attempt to use probability judgments to guide their choices. Typically such information about probability estimates is inferred from choices, implied by attempts to model those choices, or only elicited following many trials of experience (e.g., Ungemach, Chater, & Stewart, 2009). The elicitation of both choice and probability estimates simultaneously provides a vehicle for assessing how people's representations of rare event probabilities build up over time and what their role is in determining or influencing choices. This level of examination allows our work to contribute to broader debates concerning the role of propositional or prior beliefs in learning (e.g., Gershman, Blei, & Niv, 2010; Mitchell, De Houwer, & Lovibond, 2009; Shanks, 2010) as well as research on belief-updating in static and dynamic environments (e.g., Ashby, Konstantinidis, & Yechiam, 2017; Gallistel, Krishan, Liu, Miller, & Latham, 2014; Navarro et al., 2016; Plonsky et al., 2015).

Methods

Participants

² We expected the incentives to improve people's performance but note that incentives can sometimes have detrimental effects on people's performance (see Arkes, Dawes, & Christensen, 1986). We thank an anonymous referee for bringing this literature to our attention.

We recruited 40 participants for Experiment 1 (22 female, mean age = 18.93, $SD = 1.87$), 132 for Experiment 2 (100 female, mean age = 19.52, $SD = 4.59$), and 240 for Experiment 3 (167 female, mean age = 22.11, $SD = 4.54$) through the subject pool of the University of New South Wales. For all experiments, we determined a minimum sample size of 40 per experimental condition based on previous research in similar domains. We obtained informed consent from all participants and compensated them either with course credit or a flat payment of AUD 7.50 (in addition to incentive-based payment – see below for details). The UNSW School of Psychology Ethics Committee (Ref. #2909) approved all experimental work.

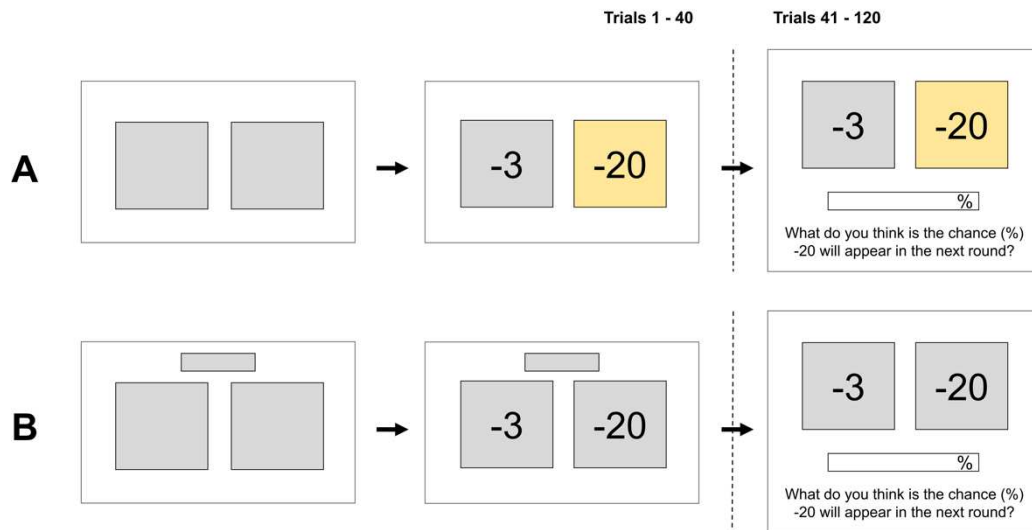


Figure 1. A. Illustration of the Choice- and dual-incentivized conditions. On trials 1-40, participants made repeated choices and received feedback on their choices (button turned yellow and displayed outcome following selection) and the option they did not choose (button remained grey, but outcome was displayed). In addition to this, at the end of each trial on trials 41-120, they gave probability estimates on the rare event occurrence. **B.** Participants in the probability-incentivized condition clicked one button (positioned above the two options) which revealed the two outcomes. Participants in this group did not make choices, but they gave probability estimates on trials 41-120.

Procedure

Experiment 1. In a computerized task, participants made repeated uncertain choices for 120 trials. They were presented with two buttons on every trial, representing a risky and a safe option (Figure 1A). Participants started the experiment with 1000 points. Choosing the risky option resulted in a loss of 20 points with a probability of .15 or no loss with a probability of .85, whereas choosing the safe option resulted in a loss of 3 points on every trial. Participants received full feedback on each trial: they were informed both about the outcome of the chosen option and about the outcome of the non-chosen option. A constant reminder of their total points was also visible to them throughout the experiment. At the beginning of the experiment, participants were informed that their starting amount reflected AUD 5.00, and that their final payment will be converted back from their remaining points. In the cases where they were paid for their participation, they received this payment in addition to their show-up fee.

To allow participants to learn about the outcomes associated with each option, they exclusively made choices for the first 40 trials. Following these trials, participants were also asked to estimate the probability of the rare event before each choice (“What do you think is the chance (%) -20 will appear on Button [A/B] in the next round?” Note that, as the position of the risky option was counterbalanced, participants were asked about either Button A or Button B). They did not receive feedback on the accuracy of their estimates. The sequence of outcomes on the risky option was pseudo-randomized for each participant; it was randomly selected from a number of pre-generated sequences, in order to make sure that all participants experienced exactly the same number of rare events. Details of the pre-generated sequences are available in the Supplementary Materials (section #2).

Two additional between-subjects conditions of this experiment (in which participants received descriptive summaries about the outcome probabilities) are not reported here, as they are not relevant to our current investigation (for more information about these conditions see Liang, Konstantinidis, Szollosi, Donkin, & Newell, 2017). We also collected prospective

estimates at the end of the experiment, which are not reported here but are available in the supplementary data files. The condition reported here served as a basis to develop Experiment 2 and 3.

Experiment 2 and 3. The procedure for Experiment 2 and 3 was similar to that of Experiment 1, but in these two experiments, participants were randomly assigned to one of three between-subjects conditions (Figure 1). We manipulated the incentive structure between these conditions the following way. In the *choice-incentivized* condition, participants were only paid based on their choices (same as in Experiment 1). In the *dual-incentivized* condition, they were paid based on both their choices and their probability judgments: when the probability estimation task started, in addition to the instruction the choice-incentivized group received, they were also informed that, on each trial, each percentage point deviation from the correct probability will reduce their final payoff by 0.025 cents. In the *probability-incentivized* condition, participants' final payoff was solely based on their judgments: they were penalized by 0.05 cents for each percentage point deviation from the objective probability on each trial, and in contrast with the other two conditions, instead of actively choosing between the options, they were only presented with the outcome of the two options and were asked to pay attention to these outcomes. They received the same instructions for the estimation task as the choice-incentivized group, and neither of the two groups received feedback on their probability estimation performance. Additionally, at the end of the experiment, participants were asked to (retrospectively) estimate the number of rare events they had seen during the experiment ("Thinking about your overall experiences: How many times do you think you saw the -20 outcome?"). Similar to Experiment 1, other measures were collected at the end of the experiment, which are not reported here but are available in the supplementary data files.

The possible outcomes and their probabilities were the same in Experiments 1 and 2. In Experiment 3, there was an additional manipulation that was factorially crossed with the

three conditions from Experiment 2: a loss and a gain domain. In the loss domain, the safe option resulted in a loss of 13 points, whereas the risky option either resulted in a loss of 10 points with a probability of .85 or a loss of 30 points with a probability of .15. In the gain domain, the safe option returned 27 points, while the risky option returned 30 with a probability of .85 or 10 with a probability of .15. To ensure equal expected value between the respective conditions, in the loss domain participants started with 4800 points and in the gain domain they started with 0 points (therefore, the expected final points for both groups were 3240 points).

Analysis and estimation

The data were analyzed in R (R Core Team, 2017). For group level analyses, we used linear mixed-effects models with the lme4 (Bates, Mächler, Bolker, & Walker, 2015) and the lmerTest (Kuznetsova, Brockhoff, & Christensen, 2016) packages. Results and detailed specification of the models are available in the Supplementary Materials (section #3).

Results

We describe the results of the three experiments jointly. First, we report group level patterns of simultaneous overestimation and underweighting (Barron & Yechiam, 2009) and the wavy recency effect (Plonsky et al., 2015). We also report the effect of incentives on both judgments and choices at the group level. Since aggregating data can be misleading regarding individual behavior (Estes, 1956; Regenwetter & Robinson, 2017), for example, if individuals employ different response strategies, we then turn to individual level data to explore the possibility for such differences and find evidence for a mixture of response types. We also explore the effect of the incentives on these different individual level patterns. Finally, we draw connections between the group level and individual level effects.

Group level results

Simultaneous overestimation-underweighting pattern. Figure 2 displays participants' risky choice proportions and average probability estimates in blocks of 10 trials for each experimental condition. At the group level, we replicated the pattern of simultaneous underweighting and overestimation of the rare event when participants made both choices and estimates in all experiments: participants' risky choice rates were reliably above .50, and their probability estimates were above .15. The only exception to this overestimation pattern was in the probability-incentivized condition, where the average estimates were closer to the objective long-run probability.




Figure 2. Participants' risky choice rates (PR) and trial-by-trial probability estimates in blocks of 10 trials for the respective experimental conditions. The y-axis represents both risky choice rates (solid lines; 0 – safe choice; 1 – risky choice) and average probability estimates (dashed lines). Error-bars represent standard errors of the mean. The figure demonstrates simultaneous overestimation and underweighting of the rare event on a group level in the choice- and dual-incentivized conditions; that is, people's average probability estimates of the rare event reflected a lower implied expected value for the risky option, yet people chose that option more frequently than chance (.50; solid grey line). Participants' mean estimates were closest to the objective value (.15; dashed grey line) in the probability-incentivized condition.

Wavy recency effect. Aggregating the data based on the distance from the most recent rare outcome trial replicated the wavy recency effect for both probability estimates and choices (Plonsky et al., 2015). Figure 3 shows this pattern (collapsed between Experiments 2 and 3). The pattern is typified by an initial negative recency pattern following the observation of a rare outcome (trials 1-5 since rare outcome), where participants expectation of the rare event seemed relatively low, reflected by higher risky choice rates and lower probability estimates for the rare event. This negative recency gradually decreased with every observation of a common outcome (trials 6-10 since rare outcome), where the expectation of the rare event increased, reflected again in both lower risky choice rates and higher probability estimates for the rare event. It is important to note that although, the elicitation of probability estimates could have produced the expectation that the probabilities are changing, we observe the same wavy recency pattern in participants' choices as has been reported in previous studies in which trial-by-trial probability estimates were not elicited (Plonsky et al., 2015). This suggests that such behavior is not caused by our trial-by-trial elicitation of probability.

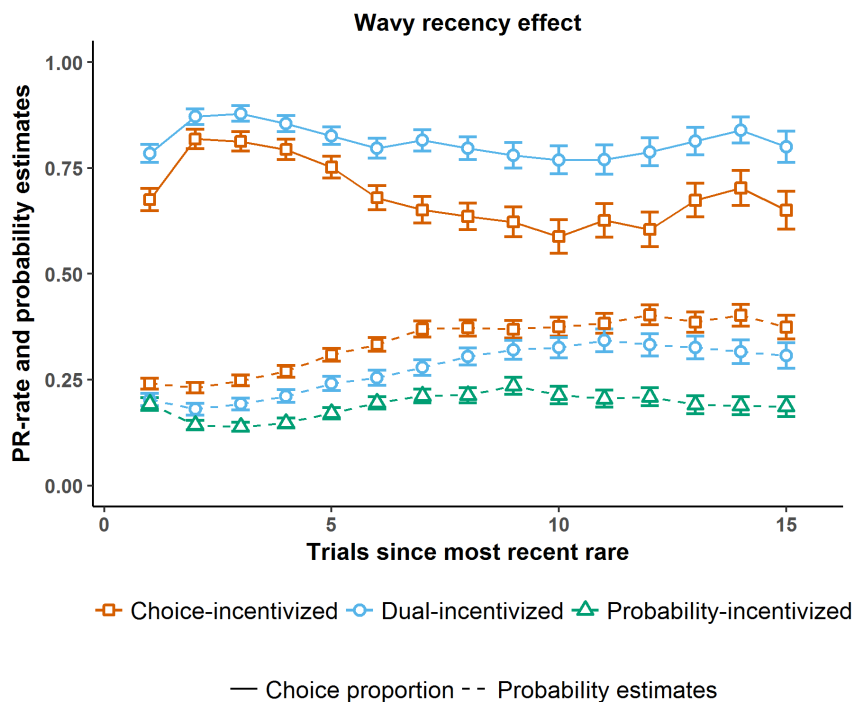


Figure 3. Wavy recency effect in risky choice rates and trial-by-trial probability estimates in all experiments. Choice rates and estimates are averaged on the trials following the observation of a rare event. Error-bars represent standard error of the mean. The wavy recency effect is characterized by an initial (5 trials) lower expectation of rare-event occurrence, followed by an increase in this expectation. The pattern is reflected in both the risky choice rates and the trial-by-trial probability estimates. Estimates given on trials >15 are not displayed.

Effect of incentives. The incentives affected people's judgments in line with our expectations: the mean probability estimates differed from each other in all conditions (all $t_s < -2.45$, all $p_s < .05$); the estimates in the choice-incentivized condition deviated most from the objective value (.15), while they were closer in the dual-incentivized condition, and the closest in the probability-incentivized condition. Interestingly, this decrease in overestimation was coupled with an increased rate of choosing the risky option ($z = 1.37$, $p = .17$, in Experiment 2; $z_s > 2.59$, $p_s < .05$ in Experiment 3). This can be seen in Figure 2 with the separation between the solid and dashed blue lines (dual-incentivized) being larger than the separation between the solid and dashed orange lines (choice-incentivized). The effect of the

incentives was observable on both estimates and choices even when the data were aggregated conditional on the recent rare outcome (Figure 3).

Individual level results

Since aggregate level data can be misleading regarding individual behavior, we also investigated individual level patterns. First, we explored individual differences in the trial-by-trial probability judgments as it provided rich data on participants' beliefs about the environment. Next, to assess the consistency of participants' beliefs with their choices, we looked at the relationship of estimates and choices. Finally, to compare trial-by-trial behavior with participants' retrospective knowledge, we report their retrospective estimates about the frequency of rare events.

Mixture of individual response patterns. Our visual inspection of the data suggested three typical temporal patterns in the probability estimates (Figure 4A-C). The data were aggregated in a similar way to Figure 3: responses on the trials were grouped based on their distance from the most recent rare event. Representative cases of these patterns are displayed in Figure 4A-C. The pattern we observed in Figure 4A was the normatively correct pattern: the participant's estimates were not affected by their distance from the most recent rare outcome. The participant seemed to assume a static environment and attempted to learn the long-run probability of the rare outcome. The patterns observed in Figure 4B and Figure 4C were more consistent with the wavy recency effect: the participants' initial lower probability estimates gradually transitioned into higher estimates. In addition to this, in the pattern in Figure 4C we observe an eventual decrease in the estimates. Taken together, these results indicate that participants differed in their responding for the probability estimation task.

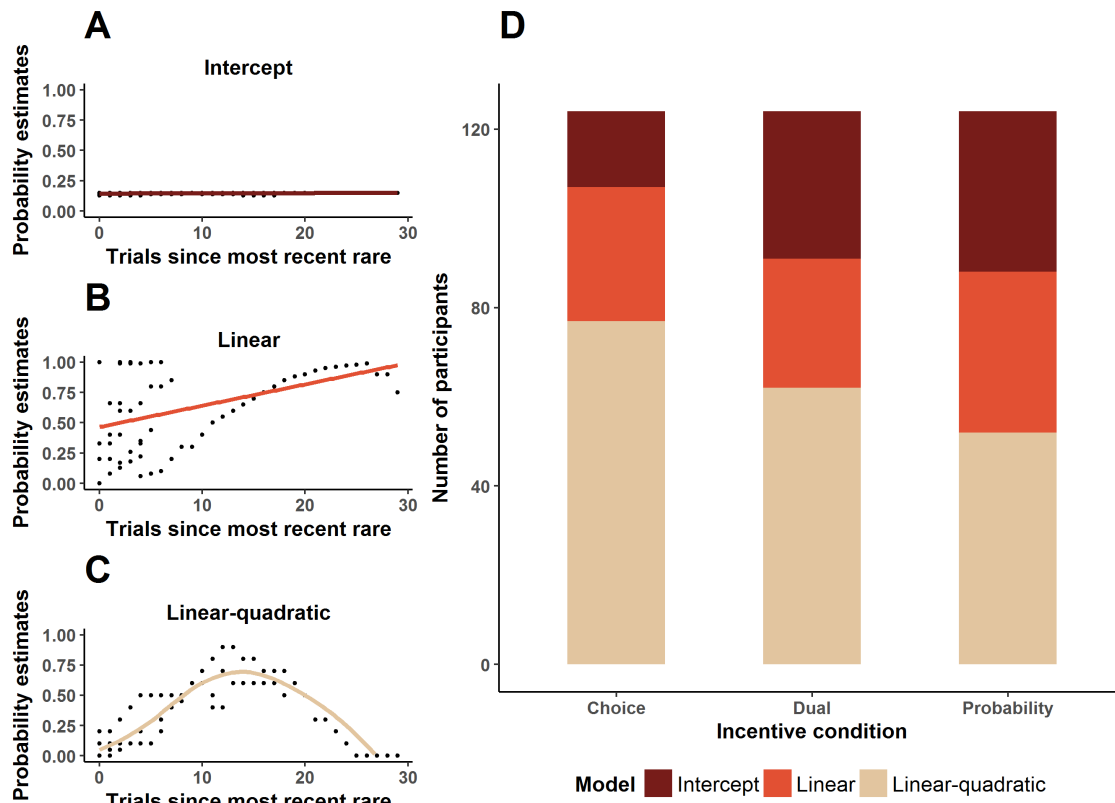


Figure 4. A-C. Typical probability estimation patterns; estimates are contingent on the number of trials since the most recent rare outcome. Similar individual probability estimation plots for each participant are available in the Supplementary Materials (Figures S1, S2, and S3). **A.** No pattern-seeking behavior: the estimates seem to be independent of the occurrence of the rare event. Solid line represents linear fit. **B.** Linear wavy-recency pattern: after the occurrence of the rare event, initial estimates indicate that the rare event is less likely to happen, but over time the expectation of the occurrence of this event increases. Solid line represents linear fit. **C.** Linear-Quadratic wavy-recency pattern: similar to the linear pattern, but the initial increase in rare-event expectation is followed by a decrease. Solid line represents loess fit. **D.** Proportion of participants whose probability estimates were best fitted by intercept, linear, or linear-quadratic models in the respective experimental conditions collapsed for Experiments 2 and 3 based on AIC indices.

Effect of incentives on individual response patterns. To determine how often each of these three patterns occurred in our experiments, we fitted simple linear models to each participant's responses (for participants in Experiments 2 and 3). We implemented the pattern described in Figure 4A as an intercept-only model; the pattern in Figure 4B as a linear model (in which the number of trials since the rare event linearly predicted the probability

estimates); and the pattern in Figure 4C as a linear-quadratic model (in which the number of trials since the rare event predicted the probability estimates with a linear and quadratic term). We then selected the best fitting model for each individual based on AIC model fit indices (we obtained similar results when we used different categorization methods, such as assessing model fit with a BIC criterion, or asking human-raters to categorize the patterns; see Supplementary Materials, section #4).

Overall, the results indicated that most participants engaged in some kind of pattern seeking behavior: 51% were best described by a linear-quadratic model, with a further 26% by a linear model, and only the remaining 23% by the no-pattern-seeking intercept only model. The proportion of individuals in these categories differed by experimental conditions, $\chi^2(4, N = 372) = 13.16, p = .011$ (Figure 4D). Participants appeared to be less pattern seeking in the probability-incentivized condition compared to the choice-incentivized condition irrespective of the categorization method (for the dual-incentivized condition, the results varied more substantially across different methods for categorization; see Supplementary Materials, section #4).

Consistency of judgment and choice. We assessed individual level pattern seeking in participants' decisions by investigating the association between their probability estimates and choices. Logistic regressions predicting choices from the preceding probability estimates for each participant (pooling data from Experiments 2 and 3) indicated that risky choices were negatively associated with their preceding probability estimate for 83% of the participants (see section #5 in the Supplementary Materials for more details). That is, most participants chose the risky option less often when they estimated that the rare event was more likely to occur. These results suggest that participants' choices are roughly consistent with their probability ratings (see section #5 Supplementary Materials for more details on the level of consistency by a normative standard). This correlation also indicates that

participants' choices were sensitive to the distance from the most recent rare outcome not just at the aggregate but at the individual level as well.

Retrospective estimates. Participants were also asked to recall the number of rare events that they encountered during the experiment. In contrast with their trial-by-trial probability estimates (but in line with previous research; e.g., Camilleri & Newell, 2009; Hau et al., 2008; Ungemach, Chater, & Stewart, 2009), their responses were close to the objective rare event frequency (aggregated across Experiment 2 and 3: Choice-incentivized: $M = 15.52$, $SD = 8.69$; Dual-incentivized: $M = 15.28$, $SD = 7.71$; Probability-incentivized: $M = 14.57$, $SD = 13.47$; each participant was presented with 18 rare events). This result suggests that participants were able to keep track of the objective frequencies; this knowledge was just not reflected in some participants' trial-by-trial estimates, potentially resulting from their assumption that the environment had a dynamic structure.

Summary of results: Relating the individual to the group level

The high level of heterogeneity at the individual level suggests that, although the group level effects observed in previous studies are robust, they consist of a mixture of different individual behaviors. In the following section, we discuss how the individual level effects could give rise to the group level phenomena.

Simultaneous overestimation-underweighting pattern. At the group level, we observed the pattern of simultaneous overestimation and underweighting of the rare event. On the one hand, participants' probability estimates of the rare outcome were, on average, higher than the objective probability (Figure 2, dashed lines). Some participants' assumption of a dynamic environment can explain this pattern. We observed that participants' probability estimates changed systematically as a function of the number of trials since the last rare event at the aggregate level (wavy-recency analysis; Figure 3). This influence was also apparent for most participants at the individual level (response pattern analysis; Figure 4), suggesting they

believed the environment to be dynamic. Believing the environment to be dynamic led to the overestimation of the probability of the rare event, because the increased variance in probability estimates introduced a regression to the mean effect that pushed the average estimates above .15 (Fiedler & Unkelbach, 2014). That is, because individual estimates were variable, and there are more values above .15 than below, the averages were “drawn” towards the midpoint of the scale ($p = .50$), resulting in an overestimation of the rare event. This regression pattern can be seen in Figure 5, which plots aggregated trial-by-trial probability estimates as a function of participants’ response-type categorization (i.e., intercept, linear, linear-quadratic). The increased variance in the average rating of the groups whose estimates implied belief in a dynamic environment (i.e., linear and linear-quadratic; Figure 4B-C) is readily apparent in the figure.

On the other hand, participants showed a preference for choosing the risky option, appearing to underweight the rare event, at the group level (Figure 2, solid lines). Similar to the overestimation pattern, we argue that most participants’ assumption of a dynamic environment can explain this pattern. First, we observed that the distance from the most recent rare outcome affected participants’ choices similarly to their probability judgments at the aggregate level (see Wavy-recency analysis, Figure 3), indicating that the alternation of participants risky and safe choices was not random but influenced by the dynamics of the environment. Moreover, participants alternated their choices *based on* their beliefs about the dynamics of the environment, suggested by the finding that choices and judgments were strongly correlated at the individual level (see the analysis under the *Consistency of judgment and choice* subsection). This link between beliefs on choices and their causal direction was further demonstrated by the experimental manipulation of participants’ judgments that resulted in changes in their choice behavior (see the effect of incentive manipulation; Figure 2).

Taken together, participants appeared to select the option that they thought would provide the best outcome on the current trial based on their beliefs about the sequential structure of the environment. Trials on which participants did not expect the rare outcome (-20 in Experiment 1 and 2; -30 and 30 in Experiment 3) to occur, it was more beneficial to choose the risky option: the common outcome of the risky option was better than the safe outcome (0 vs -3 in Experiment 1 and 2; -10 vs -13 and 30 vs 27 in Experiment 3). Since such participants predicted the rare outcome less often (since these outcomes indeed occurred less often), on most trials they would have picked the risky option – resulting in a preference for the risky option at the group level, and therefore in the apparent underweighting pattern. These findings demonstrate that inaccurate representation of the environment could account for the simultaneous overestimation-underweighting of the rare event.

Wavy recency effect. We observed a wavy recency effect at the group level (Figure 3). Although many participants engaged in pattern-seeking behavior, not all of them did. Specifically, a significant proportion of participants (correctly) seemed to assume that the environment is static and tried to learn the long-run probability of the rare event (Figure 4A). The responses of these participants were overwhelmed in the group-level wavy recency effect (Figure 3) by other participants, who seemed to assume that the environment was dynamic, and tried to predict the exact occurrence of the rare outcome (Figure 4B-C).

Effect of the incentives. We found that the accuracy of participants' estimates increased in the dual- and in the probability-incentive conditions at the group level, evidenced by the average responses being closer to the objective probability ($p = .15$). This may have been the result of the increased number of participants whose responses were best described by an intercept model (Figure 4D). These participants responded normatively and, unlike the estimates of participants whose responses were best described by a linear or a linear-quadratic model, most of their estimates were centered on the objective probability (Figure 5). Exploratory analyses suggested that the number of these participants increased in the dual-

and probability-incentive groups, which could explain why these groups' estimates became more accurate at the aggregate level.

We also found that risky choice rates of the dual-incentive group were higher than that of the choice-incentivized group. This could be resulting from the fact that the average estimate of many participants in the “intercept” group was below .15 (about 52% of the sample had estimates below .15; see Figure 5), thus they seemed to underestimate the rare event probability. At this probability, the risky option *is* the subjectively preferable alternative. An increased proportion of such participants (as can be seen in Figure 4D) can explain the increase in risky choice rates. To summarize, changes in the proportion of what kind of environment people inferred may explain the effect of the incentives.

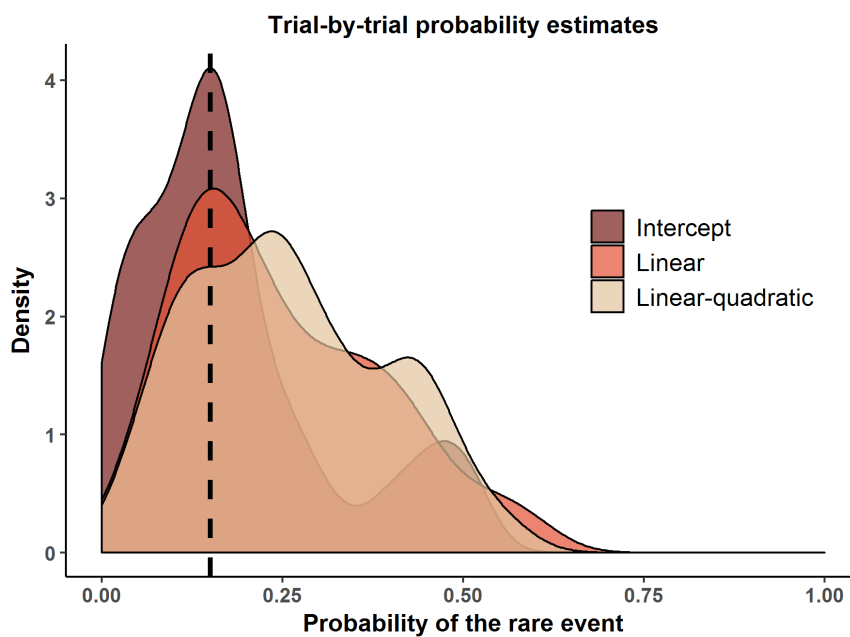


Figure 5. Smoothed histogram of trial-by-trial probability estimates of the rare event, averaged over each participant, broken down by the best fitting model of their probability estimates (intercept, linear, or linear-quadratic) according to AIC indices. The figure demonstrates the increased variance in the estimates of the linear and the linear-quadratic groups. The dashed line depicts the objective probability (.15).

Discussion

We investigated whether people's inaccurate representations of the environment could account for what appeared to be a simultaneous overestimation of the probability of rare events in people's judgments and an underweighting of the same probability in their choices (Barron & Yechiam, 2009). Our results suggest that (most, but not all) participants' misconception that the environment was dynamic rather than static gave rise to the group level dissociation of choices and judgments. We observed that both participants' probability estimates and their choices varied conjointly as a function of the number of trials since the last rare event, indicating that they believed the environment to be dynamic. We argue that this belief led to the *overestimation* of the probability of the rare event, because the increased variance in probability estimates introduced a regression to the mean effect that pushed the average estimates above .15. Conversely, based on what they believed were the dynamics of the environment, participants appeared to choose the option that they thought would provide the best outcome on the current trial – and since the risky option provided the best outcome on most trials, they chose that option more often resulting in the *underweighting* pattern. We also investigated whether such inaccurate representation of the environment can be improved and found that incentivizing people to allocate more effort to the probability estimation task increased the accuracy of their representation. Taken together these results offer an explanation for why beliefs and decisions can appear contradictory and present challenges to some current models of decisions under uncertainty.

Unpacking the Paradox

The main cause of overestimation was the incorrect mapping between participants' belief about the probabilities and the objective probabilities. Most participants seemed to believe that there were environmental regularities that they could exploit (Navarro et al., 2016): we observed the effect of a gradually decreasing negative recency on beliefs and

behaviors in the form of a wavy recency effect (Plonsky et al., 2015). Instead of giving (an increasingly accurate) long-run probability estimate on every trial, most participants appeared to offer a local prediction (“What will happen on the next trial?”) based on the assumption that the sequences of common and rare outcomes are predictable. As a result, the variance of the probability estimates increased, which amplified the regression effect (Erev et al., 1994; Fiedler & Unkelbach, 2014) that caused the apparent overestimation.

Intriguingly, the retrospective estimates were centered on the objective frequency of rare events. This finding is consistent with research suggesting that online and memory-based judgments are affected by different factors (Hastie & Park, 1986); in this case, the retrospective frequency estimates were presumably not subject to participants’ beliefs about the environmental structure. Although participants remembered the long-run frequency of the rare events correctly, they were unsuccessful in identifying the process that generated the samples they observed (Fiedler, 2012). Instead, they seemed to rely on their belief that the sequences of common and rare events are predictable. Limitations of short term memory can cause such “biased” randomness perception: people’s actual experiences with the environment can make them (understandably) believe that some sequences are more probable than others (Hahn & Warren, 2009).

Inferring sequential patterns can explain the overestimation of the rare event, but how can this be reconciled with the apparent underweighting in choices? Again, mistaken beliefs that the environment is dynamic rather than static and attempts to exploit spurious sequential patterns may provide an answer (Schulze & Newell, 2015; Speekenbrink & Konstantinidis, 2015). In dynamic environments, it is beneficial for people to predict a pattern of outcomes and choose the highest paying option on the trial level. In contrast with how people should respond in static environments (exploiting the option with the highest pay-off), alternating between the options based on one’s prediction can be a superior strategy. Since participants observed the (higher pay-off) common outcome more often than the rare outcome, they

(reasonably) predicted more common outcomes than rare outcomes, and therefore chose the risky option more often, leading to the apparent underweighting of the rare outcomes at the aggregate level.

Importantly, our results suggest that incorrect beliefs about the environment can be improved. We found an increased correspondence between probability estimates and the environment in the dual- compared to the choice-incentivized condition, and this correspondence was even larger in the probability-incentivized condition. This increased accuracy resulted from a decrease in the number of individuals that assumed a dynamic environment. Although reduced pattern-seeking under increased effort may seem counterintuitive in that it is a seemingly more cognitively taxing strategy, pattern-seeking behavior in sequential environments appears to be the default strategy that participants follow (Gaissmaier & Schooler, 2008; Schulze & Newell, 2016). Indeed, few real-life environments are static, so it seems reasonable that people's default assumption in a new and unknown environment is that it is dynamic (Navarro et al., 2016). We speculate that our manipulations allowed (some) participants to allocate more effort to identifying the correct (static) response strategy. Additionally, in the probability-incentivized condition, the removal of the choice task may have eliminated the affective costs associated with realized losses, which may have affected participants' representations in the other two conditions (Mellers, Schwartz, & Ritov, 1999).

People face uncertainty and need to learn from feedback for many high-impact real life decisions and there is substantial evidence that misrepresenting the environment can lead to problematic outcomes. For example, having misconceptions about the effects of vaccines (Brewer, Chapman, Rothman, Leask, & Kempe, 2017) or being exposed to misinformation about the causes of climate change or lung cancer (Lewandowsky, Ecker, Seifert, Schwarz, & Cook, 2012) can lower the odds of engaging in preventative behaviors. With our experiments, we attempted to provide a simple model of an uncertain environment in which being misled

by spurious patterns can result in similar misconceptions. The results suggest that people can build a more accurate environmental representation when they are pushed to allocate more effort to the task, which may aid them in making better decisions.

Theoretical Implications and Future Directions

Although current models of reinforcement learning (e.g., Plonsky & Erev, 2017; Plonsky et al., 2015) may explain the apparent paradox between choices and beliefs, and predict the wavy recency effect, they cannot account for every aspect of our findings. For example, they cannot explain why participants' choices were affected by the manipulation of their beliefs. These models suggest that decisions are based on automatic, bottom-up processing, but our finding that incentivizing the estimation task changes estimates *and* their choices clearly reflects a top-down influence of participants' representations on their choices. Such models need to specify how top-down factors interact with bottom-up learning to account for such data.

An alternative way to explain these data is in terms of models that only assume top-down effects of learning, such as the propositional theory of learning (Mitchell et al., 2009; Shanks, 2010). This framework suggests that learning is an inferential process that results in propositional beliefs about the relationship between outcomes and environmental cues. Similarly in our study, participants seemed to make an inference about a rule (or process) that generated the observed relationship between the outcomes and the distance from the most recent rare outcome (for a related idea in the Bayesian cognition literature, see Gershman, Blei, & Niv, 2010; Gershman & Niv, 2012). We suspect that such inferences come via relatively simple rules (Juslin, Karlsson, & Olsson, 2008), given the capacity limitations in systems such as working memory (Lewandowsky, Oberauer, Yang, & Ecker, 2010). For example, one such rule may be a simple monotonically increasing function based on experienced spurious patterns (Hahn & Warren, 2009), such that people increase their

expectation of the rare-event occurrence with each common event that has occurred since the last rare outcome.

Another result that current reinforcement models cannot explain is that, although people's probability estimates followed a wavy recency pattern on average, we did not see this pattern for every participant (cf. Regenwetter & Robinson, 2017). Rather, we observed a large degree of individual variability in responses, even among those who assumed that there were sequential patterns in the environment to be exploited. Why do some people realize that the environment is random while others look for patterns? The effects of the incentives suggest that people are less pattern seeking when they are motivated and/or able to allocate additional effort to the estimation task (Shanks, Tunney, & McCarthy, 2002). Additionally, the consequential outcomes of the choice task may have strengthened people's expectations of patterns. The observation of two types of pattern-seeking behavior also points to the influence of expectations. One group only increased their probability estimates, while another group eventually began to reduce their estimates as the number of trials between the rare events increased. Presumably, the number of trials since a rare event in the current sequence had exceeded what they have expected based on their assumed generating process, and so began to suspect that this process might have changed (Gallistel et al., 2014).

The observation of different response types raises the additional question of whether they represent different biases within individuals, or whether they reflect features of the unique environments in which each participant found themselves. In other words, it is not clear whether these different patterns can result from a flexible overarching decision strategy that is governed by environmental regularities, or whether they consist of a collection of qualitatively different strategies that are selected based on environmental cues (Newell, 2005). The mechanism by which people select and change between these strategies, and the properties of environments that interact with them should be an interesting area for future research. A potential way to explore this question would be to conduct an experiment with a

genuinely dynamic environment and see whether a subset of participants persist with static probability estimates for the rare event, even when probability estimates are specifically incentivized. Our current results suggest that such people would be in the minority given that appropriate incentivization should promote detection of the dynamics and because the dynamic strategy appears to be more readily accessible (Gaissmaier & Schooler, 2008; Navarro et al., 2016; Schulze & Newell, 2016). However, it is conceivable, although perhaps unlikely, that the incentives induce a more static view of the environment irrespective of its actual structure.

Conclusion

Paradoxical dissociations between people's beliefs and behavior often evaporate under scrutiny (Newell & Shanks, 2014). In the current study, we found that the apparent dissociation between beliefs and choices resulted from the mismatch between beliefs and the environment. Indeed, people's behavior was consistent with their belief that they can infer and exploit sequential patterns in the environment (Plonsky et al., 2015; Schulze & Newell, 2015). Importantly, the results revealed that the accuracy of such beliefs can be improved. We also found a top-down influence of people's beliefs on their choices and a large degree of individual differences in how people form these beliefs. These findings reiterate the importance of the distinction between consistency of judgments and choices, and correspondence of judgments and the environment (Hammond, 2000).

Context

The idea for this research originated as part of an ongoing project (funded by the Australian Research Council) that aims to resolve the tension between uncertain and risky decisions. This tension is often expressed as the description-experience gap (Hertwig et al., 2004), that is the finding that people appear to overweight small probabilities in risky, but

underweight small probabilities in uncertain decisions. We thought that the simultaneous demonstration of both of these effects by Barron and Yechiam (2009) merits further investigation, especially since the design of the task provided an arguably more “real-time” window on people’s thinking about rare outcome probabilities under uncertainty than previous measures. Additionally, we found the idea of the dissociation of judgment and behavior intriguing, since this goes against the general view that people rely on reportable knowledge in their decisions (Newell & Shanks, 2014).

References

- Arkes, H. R., Dawes, R. M., & Christensen, C. (1986). Factors influencing the use of a decision rule in a probabilistic task. *Organizational behavior and human decision processes*, 37(1), 93-110.
- Ashby, N. J. S., Konstantinidis, E., & Yechiam, E. (2017). Choice in experiential learning: True preferences or experimental artifacts? *Acta Psychologica*, 174, 59–67.
<https://doi.org/10.1016/j.actpsy.2017.01.010>
- Barron, G., & Yechiam, E. (2009). The coexistence of overestimation and underweighting of rare events and the contingent recency effect. *Judgment and Decision Making*, 4(6), 447–460.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67(1).
<https://doi.org/10.18637/jss.v067.i01>
- Brewer, N. T., Chapman, G. B., Rothman, A. J., Leask, J., & Kempe, A. (2017). Increasing Vaccination: Putting Psychological Science Into Action. *Psychological Science in the Public Interest*, 18(3), 149–207. <https://doi.org/10.1177/1529100618760521>
- Camerer, C. F., Hogarth, R. M., Budescu, D. V., & Eckel, C. (1999). The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework. In B.

- Fischhoff & C. F. Manski (Eds.), *Elicitation of Preferences* (pp. 7–48). Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-94-017-1406-8_2
- Camilleri, A. R., & Newell, B. R. (2009). The role of representation in experience-based choice. *Judgment and Decision Making*, 4(7), 518–529.
- 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. <https://doi.org/10.3758/s13423-010-0040-2>
- Erev, I., Wallsten, T. S., & Budescu, D. V. (1994). Simultaneous over- and underconfidence: The role of error in judgment processes. *Psychological Review*, 101(3), 519–527. <https://doi.org/10.1037/0033-295X.101.3.519>
- Estes, W. K. (1956). The problem of inference from curves based on group data. *Psychological Bulletin*, 53(2), 134–140. <https://doi.org/10.1037/h0045156>
- Fiedler, K. (2012). Meta-Cognitive Myopia and the Dilemmas of Inductive-Statistical Inference. *Psychology of Learning and Motivation*, 57, 1–55. <https://doi.org/10.1016/B978-0-12-394293-7.00001-7>
- Fiedler, K., & Unkelbach, C. (2014). Regressive Judgment. *Current Directions in Psychological Science*, 23(5), 361–367. <https://doi.org/10.1177/0963721414546330>
- Fischhoff, B., Parker, A. M., Bruine de Bruin, W., Downs, J., Palmgren, C., Dawes, R., & Manski, C. F. (2000). Teen expectations for significant life events. *The Public Opinion Quarterly*, 64(2), 189–205.
- Fox, C. R., & Tversky, A. (1998). A Belief-Based Account of Decision Under Uncertainty. *Management Science*, 44(7), 879–895. <https://doi.org/10.1287/mnsc.44.7.879>
- Gaissmaier, W., & Schooler, L. J. (2008). The smart potential behind probability matching. *Cognition*, 109(3), 416–422. <https://doi.org/10.1016/j.cognition.2008.09.007>
- Gallistel, C. R., Krishan, M., Liu, Y., Miller, R., & Latham, P. E. (2014). The perception of

- probability. *Psychological Review*, *121*(1), 96–123. <https://doi.org/10.1037/a0035232>
- Gershman, S. J., Blei, D. M., & Niv, Y. (2010). Context, learning, and extinction. *Psychological Review*, *117*(1), 197–209. <https://doi.org/10.1037/a0017808>
- Gershman, S. J., & Niv, Y. (2012). Exploring a latent cause theory of classical conditioning. *Learning & Behavior*, *40*(3), 255–268. <https://doi.org/10.3758/s13420-012-0080-8>
- Hahn, U., & Warren, P. A. (2009). Perceptions of randomness: Why three heads are better than four. *Psychological Review*, *116*(2), 454–461. <https://doi.org/10.1037/a0015241>
- Hammond, K. R. (2000). Coherence and correspondence theories in judgment and decision making. In T. Conolly, K. R. Hammond, & H. Arkes (Eds.), *Judgment and Decision Making: An Interdisciplinary Reader* (2nd ed., pp. 53–65). NY: Cambridge University Press.
- Hastie, R., & Park, B. (1986). The relationship between memory and judgment depends on whether the judgment task is memory-based or on-line. *Psychological Review*, *93*(3), 258–268. <https://doi.org/10.1037/0033-295X.93.3.258>
- 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. <https://doi.org/10.1002/bdm.598>
- 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. <https://doi.org/10.1111/j.0956-7976.2004.00715.x>
- Juslin, P., Karlsson, L., & Olsson, H. (2008). Information integration in multiple cue judgment: A division of labor hypothesis. *Cognition*, *106*(1), 259–298. <https://doi.org/10.1016/j.cognition.2007.02.003>
- Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, *47*(2), 263–292. <https://doi.org/10.2307/1914185>
- Kuznetsova, A., Brockhoff, B., & Christensen, H. B. (2016). lmerTest: Tests in Linear Mixed

Effects Models. Retrieved from <https://cran.r-project.org/package=lmerTest>

- 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. <https://doi.org/10.1016/j.obhdp.2011.05.001>
- Lewandowsky, S., Ecker, U. K. H., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and Its Correction. *Psychological Science in the Public Interest*, *13*(3), 106–131. <https://doi.org/10.1177/1529100612451018>
- Lewandowsky, S., Oberauer, K., Yang, L.-X., & Ecker, U. K. H. (2010). A working memory test battery for MATLAB. *Behavior Research Methods*, *42*(2), 571–585. <https://doi.org/10.3758/BRM.42.2.571>
- Liang, G., Konstantinidis, E., Szollosi, A., Donkin, C., & Newell, B. R. (2017). The impact of descriptions and incentives on the simultaneous underweighting and overestimation of rare events. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. J. Davelaar (Eds.), *Proceedings of the 39th Annual Meeting of the Cognitive Science Society* (pp. 736–742). Austin, TX: Cognitive Science Society.
- Mellers, B., Schwartz, A., & Ritov, I. (1999). Emotion-based choice. *Journal of Experimental Psychology: General*, *128*(3), 332–345. <http://dx.doi.org/10.1037/0096-3445.128.3.332>
- Mitchell, C. J., De Houwer, J., & Lovibond, P. F. (2009). The propositional nature of human associative learning. *Behavioral and Brain Sciences*, *32*(02), 183. <https://doi.org/10.1017/S0140525X09000855>
- Navarro, D. J., Newell, B. R., & Schulze, C. (2016). Learning and choosing in an uncertain world: An investigation of the explore–exploit dilemma in static and dynamic environments. *Cognitive Psychology*, *85*, 43–77. <https://doi.org/10.1016/j.cogpsych.2016.01.001>
- Newell, B. R. (2005). Re-visions of rationality? *Trends in Cognitive Sciences*, *9*(1), 11–15. <https://doi.org/10.1016/j.tics.2004.11.005>

- Newell, B. R., & Rakow, T. (2007). The role of experience in decisions from description. *Psychonomic Bulletin & Review*, *14*(6), 1133–1139.
<https://doi.org/10.3758/BF03193102>
- Newell, B. R., & Shanks, D. R. (2014). Unconscious influences on decision making: A critical review. *Behavioral and Brain Sciences*, *38*(01), 1–19.
<https://doi.org/10.1017/S0140525X12003214>
- Plonsky, O., & Erev, I. (2017). Learning in settings with partial feedback and the wavy recency effect of rare events. *Cognitive Psychology*, *93*, 18–43.
<https://doi.org/10.1016/j.cogpsych.2017.01.002>
- Plonsky, O., Teodorescu, K., & Erev, I. (2015). Reliance on small samples, the wavy recency effect, and similarity-based learning. *Psychological Review*, *122*(4), 621–647.
<https://doi.org/10.1037/a0039413>
- R Core Team. (2017). R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. Retrieved from <https://www.r-project.org/>
- 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. <https://doi.org/10.1002/bdm.681>
- Regenwetter, M., & Robinson, M. M. (2017). The construct–behavior gap in behavioral decision research: A challenge beyond replicability. *Psychological Review*, *124*(5), 533–550. <https://doi.org/10.1037/rev0000067>
- Schulze, C., & Newell, B. R. (2015). Compete, coordinate, and cooperate: How to exploit uncertain environments with social interaction. *Journal of Experimental Psychology: General*, *144*(5), 967–981. <https://doi.org/10.1037/xge0000096>
- Schulze, C., & Newell, B. R. (2016). Taking the easy way out? Increasing implementation effort reduces probability maximizing under cognitive load. *Memory & Cognition*,

44(5), 806–818. <https://doi.org/10.3758/s13421-016-0595-x>

Shanks, D. R. (2010). Learning: From Association to Cognition. *Annual Review of Psychology*, 61(1), 273–301. <https://doi.org/10.1146/annurev.psych.093008.100519>

Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*, 15(3), 233–250. <https://doi.org/10.1002/bdm.413>

Speekenbrink, M., & Konstantinidis, E. (2015). Uncertainty and Exploration in a Restless Bandit Problem. *Topics in Cognitive Science*, 7(2), 351–367. <https://doi.org/10.1111/tops.12145>

Ungemach, C., Chater, N., & Stewart, N. (2009). Are Probabilities Overweighted or Underweighted When Rare Outcomes Are Experienced (Rarely)? *Psychological Science*, 20(4), 473–479. <https://doi.org/10.1111/j.1467-9280.2009.02319.x>

Viscusi, W. K., & Hakes, J. K. (2008). Risk beliefs and smoking behavior. *Economic Inquiry*, 46(1), 45–59. <https://doi.org/10.1111/j.1465-7295.2007.00079.x>

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. <https://doi.org/10.1037/bul0000115>

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. <https://doi.org/10.1002/bdm.509>