

Elston, Thomas W.; Leuthold, Hartmut; Mackenzie, Ian G.; Mittelstädt, Victor

Article — Published Version

Extreme Outcomes Accentuate Overweighting of Low Probabilities and Underweighting of High Probabilities in Experience-Based Choice

Journal of Behavioral Decision Making

Provided in Cooperation with:

John Wiley & Sons

Suggested Citation: Elston, Thomas W.; Leuthold, Hartmut; Mackenzie, Ian G.; Mittelstädt, Victor (2024) : Extreme Outcomes Accentuate Overweighting of Low Probabilities and Underweighting of High Probabilities in Experience-Based Choice, Journal of Behavioral Decision Making, ISSN 1099-0771, Wiley, Hoboken, NJ, Vol. 37, Iss. 2, <https://doi.org/10.1002/bdm.2380>

This Version is available at:

<https://hdl.handle.net/10419/290182>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.



If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<http://creativecommons.org/licenses/by-nc/4.0/>

RESEARCH ARTICLE OPEN ACCESS

Extreme Outcomes Accentuate Overweighting of Low Probabilities and Underweighting of High Probabilities in Experience-Based Choice

Thomas W. Elston¹  | Hartmut Leuthold² | Ian G. Mackenzie² | Victor Mittelstädt² 

¹Helen Wills Neuroscience Institute, University of California at Berkeley, Berkeley, California, USA | ²Department of Psychology, University of Tübingen, Tübingen, Germany

Correspondence: Thomas W. Elston (telston@nurhopsi.org) | Victor Mittelstädt (victor.mittelstaedt@uni-tuebingen.de)

Received: 5 August 2023 | **Revised:** 20 February 2024 | **Accepted:** 21 February 2024

Funding: Data collection was supported by internal funding from the Department of Psychology at the University of Tübingen. This research was funded as part of the Excellence Strategy of the German Federal and State Governments.

Keywords: decision-making | prospect theory | reinforcement learning | subjective value | uncertainty

ABSTRACT

Subjective inferences of probability are critical for decisions involving uncertainty. Recent studies have suggested that extreme outcomes bias beliefs about the value of uncertain options toward the best/worst outcome possible when learning the odds through experience, leading to increased preferences for uncertain options over equivalent sure bets when there is the prospect of gain and, conversely, aversion to uncertain options when there is the prospect of loss. However, prior studies regarding the influence of extreme outcomes on decisions involving uncertainty have only done so using 50/50 gambles, and it was unclear whether extreme outcomes biased probability perception more broadly. Across three pre-registered experiments, we found that when people made decisions between equivalent certain and uncertain options, they particularly preferred uncertain options at low probabilities (20%) when there was the prospect of gain and avoided them when there was the prospect of loss, with these preferences being reduced or even reversed at medium (50%) and high (80%) probabilities. We also found that uncertainty preferences were influenced by outcome extremity and the relative certainty associated with safe reference options. We conclude that extreme outcomes accentuate the overweighting of low probabilities and the underweighting of high probabilities in experience-based choice.

1 | Introduction

Economic decision-making depends on subjective inferences of choice option values. Often, these inferences correspond to the objective payoffs of choice options; however, decisions involving uncertain outcomes can lead to distortions in subjective valuation. Many studies examined these distortions in scenarios where the outcomes and likelihoods were explicitly described (e.g., Kahneman and Tversky 1979). For example, people might be asked explicitly to choose between the prospects of either a guaranteed 50€ or a 50/50 chance at 100€. When faced with these description-based choices, participants generally preferred the gamble when

mitigating losses and avoided it when seeking gains, suggesting that people treat the prospect of loss as more consequential than the prospect of gain. However, decisions outside of the laboratory are often based on experience of the consequences of those decisions. People frequent certain cafes but not others, decide whether or not to get vaccinated, or whether to risk crossing the street when the light is still red—all largely based on their own experiences. Experience-based choices can lead to radically different behavior than description-based choices (Camilleri and Newell 2011, 2013; Elston, Mackenzie, and Mittelstädt 2021; Hertwig and Erev 2009; Hertwig et al. 2004; Ludvig and Spetch 2011; Wulff, Mergenthaler-Canseco, and Hertwig 2018).

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial](https://creativecommons.org/licenses/by-nc/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2024 The Authors. *Journal of Behavioral Decision Making* published by John Wiley & Sons Ltd.

A common finding from studies of experience-based decision-making is that people tend to prefer uncertain options when seeking gains and avoid them when mitigating losses (Konstantinidis, Taylor, and Newell 2018; Ludvig, Madan, and Spetch 2013; Ludvig and Spetch 2011). For example, Ludvig, Madan, and Spetch (2013) found that when people decided between a 50/50 gamble for +40/−40 points or a guaranteed +20/−20 points in win/loss contexts, respectively, participants typically preferred the uncertain option (+40) when there was the prospect of gain, whereas they preferred the safe option (−20) when faced with the prospect of loss—precisely the opposite of what occurs in description-based choices. Ludvig, Madan, and Spetch (2013) explained that these effects in terms of an “extreme outcome rule” (EOR), which posits that participants tend to bias their beliefs about an uncertain option’s value toward the best/worst outcome it could yield, thereby leading to preferences for uncertain gains and aversion to uncertain losses.

Subsequent research has shown that extreme outcomes are more salient in perception and memory than more moderate outcomes (Konstantinidis, Taylor, and Newell 2018; Ludvig et al. 2018; Madan, Ludvig, and Spetch 2014, 2017). This suggests that the relative extremity associated with uncertain outcomes relative to safe ones leads to biased subjective valuation—a notion characterized through a neurophysiologically supported risk-sensitive reinforcement learning model (Niv et al. 2012). However, prior studies examining uncertainty preferences across gain and loss contexts have typically examined uncertainty preferences only with 50/50 gambles (e.g., 50/50 chance of 40 or 0). Decisions involving uncertainty outside of the laboratory often involve risk levels beyond 50%, and it remains unclear whether the EOR generalizes to other probabilities. Thus, the overall goal of our experiment was to assess how extreme outcomes influence uncertainty preferences across a wider range of probabilities than previously investigated using a task design similar to those previously used to examine the role of extreme outcomes.

A simple generalization of the EOR predicts that participants will prefer the more uncertain options across all probabilities when there is the prospect of gain and, conversely, avoid them when there is the prospect of loss (see Figure 2). This perspective suggests that extreme outcomes modify uncertainty preferences uniformly across all probabilities. Prior work by Elston, Mackenzie, and Mittelstädt (2021) suggests that such a generalization is plausible. The critical condition in their experiment involved choices between equiprobable gambles where the choice options differed in their associated uncertainty but were matched in terms of their payoff probabilities and amounts. Specifically, they used partially filled bars for

the “more certain” options where, for example, 20% filledness indicated a 20% payoff likelihood; emojis were used as “more uncertain” options because no aspect of the stimulus described its payoff likelihood. Thus, the probabilities associated with the bars could be directly inferred from the stimulus itself, reducing their predictive uncertainty, whereas the probabilities associated with emojis could only be inferred through learning and were therefore relatively more uncertain. Elston et al.’s participants preferred the more certain, bar-based options across all probabilities (20%, 50%, and 80%) in the loss domain and, conversely, preferred the more uncertain emoji options across all probabilities in the gain domain.

While Elston, Mackenzie, and Mittelstädt’s (2021) study appears consistent with the notion that people overweight outcomes associated with the more uncertain options, it did not directly assess the roles of outcomes themselves. This is because the outcomes were the same for all stimuli (either +1 or 0 in the gain domain; either −1 or 0 in the loss domain). Because the actual outcome amounts were matched, Elston et al.’s study cannot address whether extreme outcomes influence uncertainty preferences uniformly or disjointly across all probabilities.

In the present study, we addressed this question through three pre-registered experiments where the critical condition involved decisions between equivalent safe and uncertain options across three probabilities (20%, 50%, and 80%) with the goals of either maximizing gains or minimizing losses. Uncertainty levels were defined as the probability of the more extreme event (e.g., an option with an 80% chance of +40 or else +20 would be defined as an 80% option; see Tables 1, 2, and 3); equivalent choices were defined as those where the options had equal expected value (e.g., 50/50 chance of +40 (or nothing) vs. 100% chance of +20) (see Figure 1). Thus, in contrast to Elston, Mackenzie, and Mittelstädt (2021), the uncertain options always had more extreme outcomes relative to the safe options. If extreme outcomes uniformly inflate subjective value estimates uniformly across all probabilities, participants will generally prefer the more uncertain options across all probabilities when there is the prospect of gain and, conversely, avoid them when there is the prospect of loss. In other words, as the EOR seems to solely refer to the valuation of outcomes, a simple generalization of this rule across probabilities would imply that different probabilities do not modulate choice behavior.

An alternative possibility is that extreme outcomes modify uncertainty preferences in conjunction with probability weighting, whereby people behave as if a given option is more or less likely to yield a payoff or loss than it objectively does. Many studies

TABLE 1 | Payoff scheme and stimuli for Experiment 1.

Gain domain		Loss domain	
Stimulus	Payoff(s)	Stimulus	Payoff(s)
Emoji 1	20% +40; 80% +20	Emoji 4	20% −40; 80% −20
Emoji 2	50% +40; 50% +20	Emoji 5	50% −40; 50% −20
Emoji 3	80% +40; 20% +20	Emoji 6	80% −40; 20% −20
Green 20%-filled bar	100% +24	Red 20%-filled bar	100% −24
Green 50%-filled bar	100% +30	Red 50%-filled bar	100% −30
Green 80%-filled bar	100% +36	Red 80%-filled bar	100% −36

TABLE 2 | Payoff scheme and stimuli for Experiment 2.

Gain context		Loss context	
Stimulus	Payoff(s)	Stimulus	Payoff(s)
Emoji 1	20% +40; 80% +0	Emoji 4	20% −40; 80% −0
Emoji 2	50% +40; 50% +0	Emoji 5	50% −40; 50% −0
Emoji 3	80% +40; 20% +0	Emoji 6	80% −40; 20% −0
Green 20%-filled bar	100% +8	Red 20%-filled bar	100% −8
Green 50%-filled bar	100% +20	Red 50%-filled bar	100% −20
Green 80%-filled bar	100% +32	Red 80%-filled bar	100% −32

TABLE 3 | Payoff scheme and stimuli for Experiment 3.

Gain context		Loss context	
Stimulus	Payoff(s)	Stimulus	Payoff(s)
Emoji 1	20% +40; 80% +20	Emoji 7	20% −40; 80% −20
Emoji 2	50% +40; 50% +20	Emoji 8	50% −40; 50% −20
Emoji 3	80% +40; 20% +20	Emoji 9	80% −40; 20% −20
Emoji 4	100% +24	Emoji 10	100% −24
Emoji 5	100% +30	Emoji 11	100% −30
Emoji 6	100% +36	Emoji 12	100% −36

have found that people perceive low and high probabilities nonlinearly. For example, studies of choices between written, explicitly described safe, and risky options suggest that people overweight low probabilities and underweight high ones—consistent with Prospect Theory (see Figure 2) (Kahneman and Tversky 1979; Ruggeri et al. 2020). In contrast, other studies suggest that when people infer an uncertain option's payoff likelihood through sampling (and thus experiencing) its outcomes, people appear to underweight low probabilities and overweight high ones (Fox and Hadar 2006; Hertwig et al. 2004; Hertwig and Erev 2009; Kellen, Pachur, and Hertwig 2016; Wulff, Mergenthaler-Canseco, and Hertwig 2018).

Figure 2 illustrates the behavior patterns predicted by a simple generalization of the EOR, sampling-based decisions from experience studies, as well as Prospect Theory in the context of our task environment. We note that although we designed our task explicitly to study the EOR, it remains unclear whether extreme outcomes uniformly influence uncertainty preferences or whether extreme outcomes might modulate probability weighting. Thus, while our aim was not to delineate between these three simplified patterns directly, our task design did not preclude the possibility that these factors could interact. Thus, it still seems useful to consider these perspectives if the results are not entirely compatible with the EOR.

2 | General Methods and Materials

2.1 | Transparency and Openness

All experiments were pre-registered on the Open Science Framework (OSF). Raw data associated with each experiment is also available via the OSF. The JavaScript code used to control the experiments is available on Github, as are the Python scripts used to analyze the data.

- OSF project page with all pre-registrations and raw data: https://osf.io/7qbp8/?view_only=57351bf1248245ffbadb2db2d34ec4f9
- Experiment code (JavaScript): <https://github.com/igmmgi/jsPsychExperiments/tree/master/Experiments/RiskyProbability>
- Analysis code (Python): https://github.com/t-elston/safe_risky

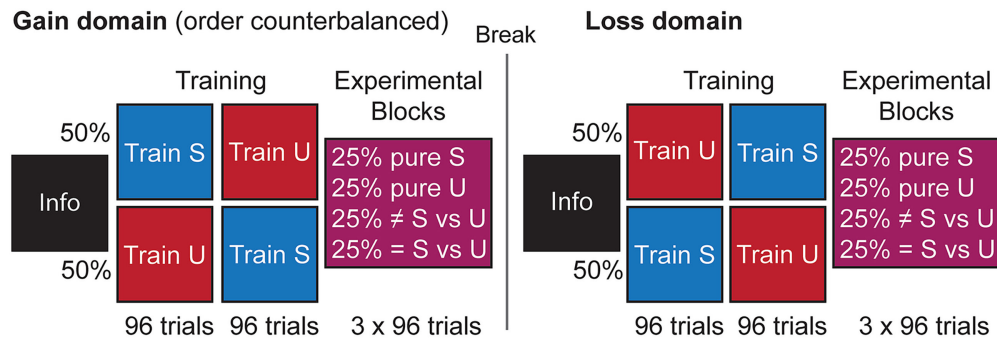
2.2 | Participants and Ethics Standards

All participants in this and the following experiments gave informed consent prior to testing. Furthermore, all experiments were in accordance with the ethical standards of the institutional (University of Tübingen) and German national research committees and with the 1964 Declaration of Helsinki (and its later amendments) as well as with the EU and German national data security regulations (GDPR and DSGVO). The exact number of participants in each experiment are described in Section 3.

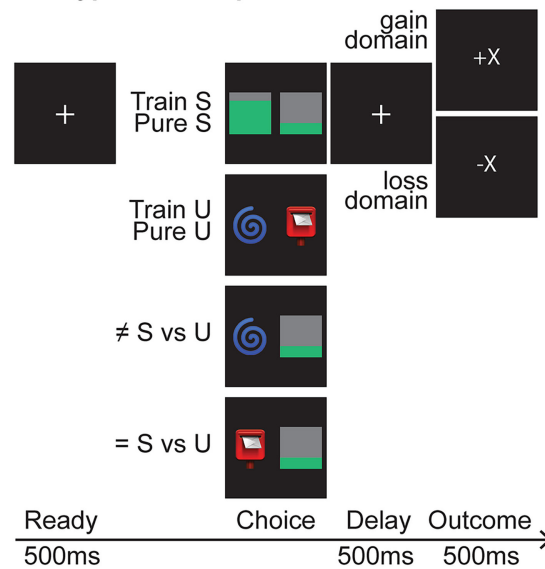
2.3 | Online Behavioral Testing Environment and Stimuli

Experiments were conducted online using the JavaScript library jsPsych (de Leeuw 2015). A centrally positioned “+” on a light gray background served as a fixation point. The uncertain options were represented as six emojis (three unique options for each domain—gain and loss) that were randomly selected from a pool of 34 emojis. All candidate emojis were simple in color and form and did not contain faces or face-adjacent information (Elston, Mackenzie, and Mittelstädt 2021). The probabilities of 20%, 50%, and 80% were randomly assigned to the emojis for each subject. Across experiments, we changed

A - Experimental Design and Session Flow



B - Trial Types and Sequences



C - Example Equivaluable Choices




FIGURE 1 | Session organization and trial structure. (A) Experimental design and events of a single experimental session. “S” refers to trials with “safe” options and “U” refers to trials with uncertain options. (B) Trial types and sequences of screens (+/−X refers to the gain or loss of X points, depending on the selected stimulus). Green-filled bars are used here to depict the safe options although either red or green fill (depending on the gain/loss domain) were experimentally used. (C) Examples of equivaluable (=SvsU) choices in the gain context of Experiment 2 with the corresponding payoff contingencies. The screen side of stimuli were counterbalanced across trials and no text was present on the screen during decisions.

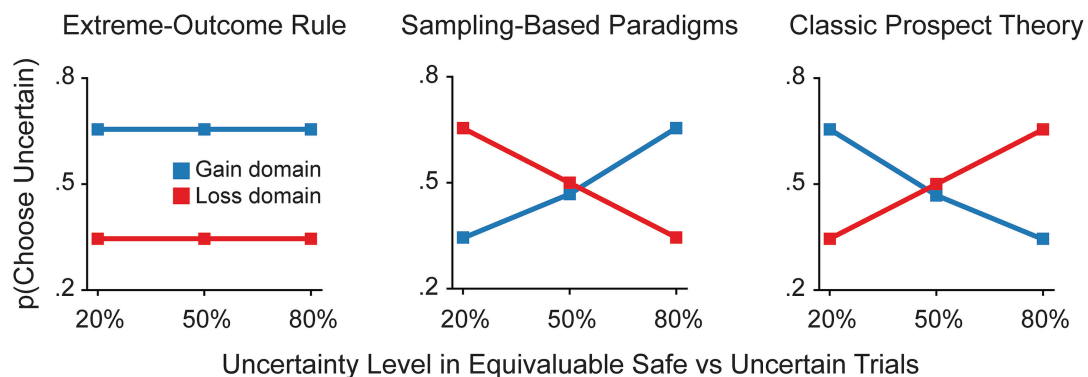


FIGURE 2 | Choice patterns consistent with the extreme outcome rule, early studies of sampling-based experiential learning, and classic prospect theory.

the format of the safe, reference options. In Experiments 1 and 2, we used partially filled bar stimuli where the fill percentage was related to the sure-bet number of points gained or lost

(depending on the gain/loss context). Each of the three safe options had the same long-run expected values as an equivaluable (probability × amount) uncertain option. We used two

different fill colors to differentiate across contexts: green fill was used to denote gains, whereas red fill was used to denote losses. In Experiment 3, the safe options were randomly selected emojis.

2.4 | Session Organization and Trial Types

Each participant was tested in 10 blocks of 96 trials (see Figure 1B). Half of the participants experienced the gain context in the first 5 blocks (and the loss context in the remaining blocks). The remaining half experienced the reverse order. Within a context, the first two blocks were training blocks consisting entirely of either “safe” or “uncertain” options. The purpose of these training blocks was to ensure that participants adequately sampled the appropriate outcome distribution for each stimulus and learned the stimulus values. Before starting the first training block of the gain context, participants were told to select the best option to gain as many points as possible. Before starting the loss context, participants were told to retain as many points as possible (i.e., to mitigate their losses). Participants were also told that the top 10% of participants finishing with the most points at the end of the session would win a 10€ voucher. The order of safe and uncertain training blocks was counterbalanced across contexts and subjects. A self-paced pause was introduced between each block. To ensure participants did not accrue negative points throughout the experiment, all participants were given 20,000 points prior to the start of the first block of the first context.

In the training blocks, participants made choices exclusively between either safe or uncertain options. In the main experimental blocks, four equally occurring trial types were presented: (i) “pure safe (Pure S)” trials where participants decided between two “safe” options (identical to training blocks); (ii) “pure uncertain (Pure U)” trials where participants decided between two “uncertain” options (identical to training blocks); (iii) “unequal safe versus uncertain (\neq SvsU)” trials where participants decided between both “safe” and “uncertain” options but one was objectively better; (iv) “equal safe versus uncertain ($=$ SvsU)” trials where participants choose between a “safe” and “uncertain” option with identical expected values. All stimulus combinations were equally counterbalanced in terms of frequency and spatial location (i.e., all stimuli were presented an equal number of times on the left and right sides of the computer screen).

2.5 | Structure and Events of Individual Trials

As schematized in Figure 1B, trials began with the presentation of a central “+” for 500ms. Participants were then presented with two different stimuli of either the same or different type (either “safe” or “uncertain” stimuli) on the left/right side of the computer screen. Participants responded with their left and right index fingers by pressing the “Q” (left) and “P” (right) keys on QWERTZ keyboards. The stimuli remained on the screen until participants responded; there was no response deadline on choice reaction times (RTs). After each trial, participants received feedback about the points their choice accumulated.

2.6 | Exclusion Criterion and Data Preparation

As pre-registered, we excluded subjects that performed worse than 60% across trials where decisions were between the same type (safe vs. safe and risky vs. risky) across both the training and experimental blocks. We report the exact number of excluded subjects in Section 3. Furthermore, as per our pre-registrations, we excluded individual trials with outlier choice RTs. Based on our previous study (Elston, Mackenzie, and Mittelstädt 2021) and a manual inspection of the data, we defined outlier RTs as those less than 200 ms and greater than 1500 ms. These excluded individual trials were rare and never comprised more than 3% of an individual’s responses.

2.7 | Data Analysis

Analyses were pre-registered on the Open Science Framework and conducted via custom scripts written in Python. The code for this experiment is available on our Open Science Framework project page and Github (links are provided in Section 2.1). We used the Pingouin (Vallat 2018) and statsmodels (Seabold and Perktold 2010) statistics modules in Python. All statistical tests were two-sided, and we used $p < 0.05$ as our significance threshold. In the figures, all markers denote the mean \pm standard error of the mean (SEM) unless stated otherwise.

3 | Results

3.1 | Experiment 1: High Expected Values and Descriptive Safe Options

In our first experiment, our primary question was whether participants’ preference for uncertain gain versus loss options was independent of the probability of the outcome. Ludvig et al. (2013) found that participants were more likely to overweight the odds associated with uncertain options when the expected values of all choice options were relatively higher and every trial led to a concrete gain or loss (no “zero”/nothing outcomes). Our first experiment capitalized on these prior results in order to examine whether uncertainty preferences were biased at low and high probabilities. In addition to using larger expected values (i.e., larger potential payoffs/losses) and no “nothing” outcomes, we also used descriptive bar stimuli as the safe options to heighten the contrast between the safe, “sure-bet,” and uncertain options. The exact payoff scheme and stimuli used in Experiment 1 are presented in Table 1.

3.1.1 | Participants

Forty-eight University of Tübingen undergraduate psychology students participated in this experiment ($M \pm SEM$ age: 21 ± 0.62 years old; 38 female). Three participants were excluded according to our pre-registered inclusion criterion.

3.1.2 | Task Validation: Assessing Sampling Error and Task Performance

Given the potentially confounding role that a sampling error can play when interpreting studies of experience-based choice

(Barron and Ursino 2013; Camilleri and Newell 2011; Fox and Hadar 2006; Hau et al. 2008; Hertwig et al. 2004), we first verified that participants actually experienced the outcome likelihoods associated with the uncertain options in the training blocks. For example, we tested how often selecting a 20% uncertain option in the gain context led to a large reward. As illustrated in Figure 3A (top), participants adequately experienced underlying probability associated with each uncertain option (all $p > 0.05$ when t -testing experienced probability against objective probability for a given option). This result indicates that participants adequately sampled and experienced the probabilities associated with the uncertain options and subsequent results cannot be explained as an artifact of sampling error.

Next, we assessed whether participants adequately learned the option values and thus chose the best option when a best option was available. We defined the “best” option as the one with the higher expected value (probability \times amount). As shown in Figure 3B–D (top), across all choice conditions in the training and experimental blocks, participants tended to select the better option when one was available (all $ts(44) > 4.1$ all $ps < 0.001$, all Cohen's $ds > 0.61$, t -tests against 50%). These results indicate that the participants adequately learned the option values and were able to generalize and adequately compare their inferred values across the safe and uncertain stimulus types (i.e., in \neq SvsU trials).

3.1.3 | Core Results: Choice Patterns During Equivaluable ($=$ SvsU) Trials

Our experiment's critical condition involved decisions between equivaluable safe and uncertain options—options that had the same expected value (probability \times amount). We used three uncertainty levels (20%, 50%, and 80%), and to isolate the effect of probability weighting, all uncertain options paid/lost the same number of points, depending on whether a trial was in the gain or loss domain. The uncertain option always had a more extreme outcome than the safe option, and the only difference in preferences across the probabilities could be due to how the probabilities themselves were weighted. This allowed us to directly measure participant probability weightings through how often the uncertain option was selected versus an equivaluable safe option.

To that end, we conducted repeated measures analyses of variance (rmANOVAs) with the within-subject factors of uncertainty level (20%, 50%, and 80%) and context (win and loss). The rmANOVA revealed significant main effects of probability ($F(2, 88) = 7.28$, $p = 0.001$, $\eta_p^2 = 0.14$) and context ($F(1, 44) = 47.24$, $p < 0.001$, $\eta_p^2 = 0.52$). These main effects indicate that participants were generally more likely to select uncertain options for larger probabilities and in the gain context than in the loss context.

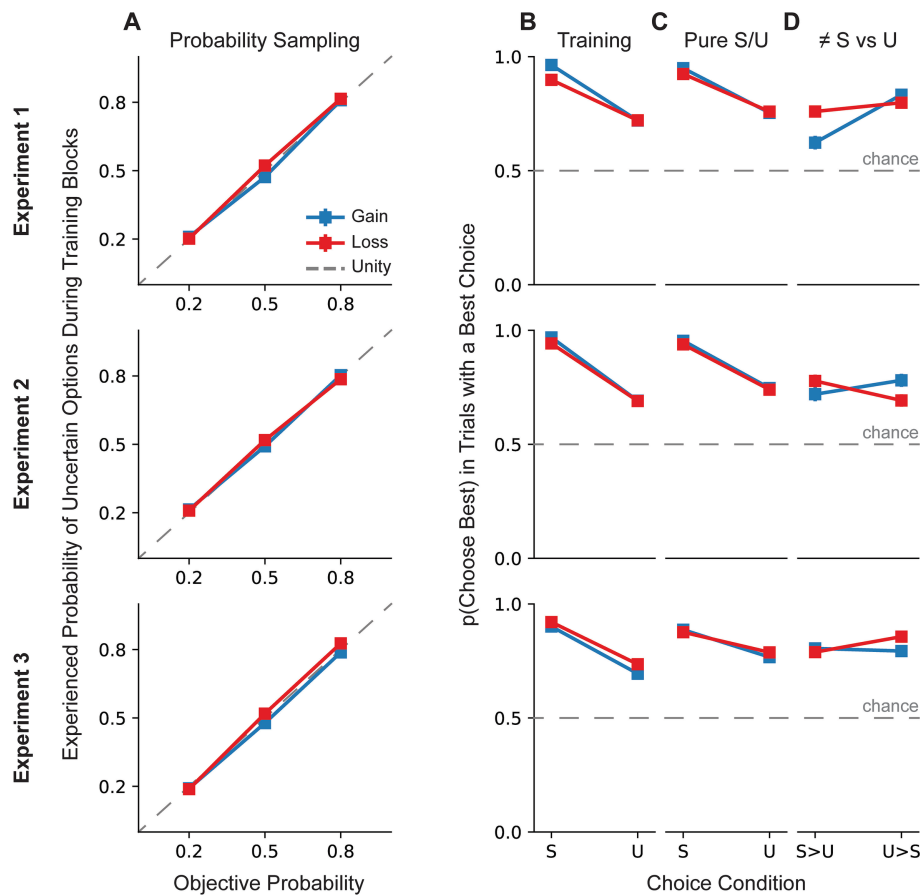


FIGURE 3 | Sampling behavior and task performance. (A) Ruling out sampling error. Participants experienced the outcome odds associated with the uncertain stimuli in the training blocks that preceded the main, experimental blocks. (B) Choice patterns during the initial training blocks and. (C) The “pure safe/uncertain” trials of the experimental blocks. In both cases, these were trials in which participants decided exclusively between either safe (“S” on x-axis) or uncertain (“U” on x-axis) options. (D) Choice patterns during trials in which participants decided between both a safe and an uncertain option and one was better. “S>U” denotes trials where the safe option was better than the uncertain one and “U>S” denotes trials where the uncertain option was better than the safe one. Each row in the figure corresponds to one experiment. Markers denote \pm SEM.

Critically, the probability \times context interaction was also significant ($F(2, 88) = 11.28, p < 0.001, \eta_p^2 = 0.20$). To better understand the nature of this interaction effect, we first computed the difference between p (choose uncertain) during the gain and loss domain; positive values indicate a preference for the uncertain option in the gain domain and, conversely, uncertainty-aversion in the loss domain. Negative values indicate a preference for the uncertain option in the loss domain and, conversely, uncertainty-aversion in the gain domain. We regressed these contextual differences in uncertainty preference against the uncertainty level associated with each equivalent condition. The resulting beta weight was significantly negative ($\beta = -0.46, p < 0.001$; see Figure 3D), indicating that participants especially preferred the uncertain option in the gain domain (and, conversely, avoided the uncertain option in the loss domain) at the lowest probabilities and tended toward neutrality at the middle and high probabilities. These results suggest that low probabilities were especially overweighted, leading participants to prefer low-probability options in the gain domain and to avoid them in the loss domain.

3.2 | Experiment 2: Low Expected Values and Descriptive Safe Options

In Experiment 1, participants strongly overweighted low probabilities, preferring/avoiding uncertain options in the gain/loss domains. However, the size of this effect, although still present, diminished at intermediate and high probabilities. Thus, this pattern was not directly predicted by a simple generalization of the EOR. The goal of Experiment 2 was to investigate the extent to which this pattern would replicate when using different outcomes.

Ludvig et al. (2013) previously showed such an effect at the 50% probability level and found that the size of the across-domain differences appeared to depend on the payoff scheme. Specifically, they found that a nonzero payoff scheme, as used in Experiment 1, potentiated uncertainty preferences. Next, we pre-registered and conducted Experiment 2 to examine whether using a payoff scheme with lower expected values and outcomes that included “0” would diminish probability overweighting (as predicted by Ludvig et al.’s observations) and whether such an effect would be uniform across low, intermediate, and high probabilities. See Table 2 for complete details of the payoff scheme. To make Experiment 2 directly comparable with Experiment 1, we continued to use descriptive bar stimuli as the safe options.

3.2.1 | Participants

A new sample of 48 University of Tübingen undergraduate psychology students participated in this experiment ($M \pm SEM$ age: 22 ± 0.66 years old; 32 female). Three participants were excluded according to our pre-registered inclusion criterion.

3.2.2 | Task Validation: Assessing Sampling Error and Task Performance

As in Experiment 1, participants in Experiment 2 adequately sampled the uncertain options during the training blocks (all $p > 0.05$ when t -testing experienced probability against objective probability for a given option; see Figure 3A, middle) selected the best option across all conditions when a best option was

available (all one-sample t -tests with $ts(44) > 7.17, ps < 0.001$, and Cohen’s $ds > 1.06$; see Figure 3B–D, middle). These results indicate that participants learned the values associated with the choice options and that subsequent results cannot be explained as an artifact of sampling error.

3.2.3 | Core Results: Choice Patterns During Equivalent (=SvsU) Trials

Figure 4B shows participant choice patterns during the equivalent (=SvsU) trials as a function of uncertainty level and context. Critically, the interaction was again significant ($F(2, 88) = 8.07, p < 0.001, \eta_p^2 = 0.16$; $ps > 0.06$ for the main effects). To characterize this interaction, we again assessed the difference between uncertainty preferences across contexts and regressed the differences against the probability associated with equivalent conditions. As in Experiment 1, we found a significant negative beta weight ($\beta = -0.51, p < 0.001$; see Figure 4E). The negative slope indicates that participants were most likely to prefer the uncertain option for the 20% chance of gain (and, conversely, avoided the uncertain option with the 20% chance of loss). These results again suggest that low probabilities were overweighted and that probabilities associated with common losses, including at the 50% level, were perceptually underweighted. Although the pattern of probability weighting we observed is not what would be expected from a simple generalization of the EOR, when comparing across Experiments 1 and 2, we note that probability weights appeared to be globally diminished when expected values were lower, consistent with the spirit of the EOR. We directly statistically compare these experiments in a later subsection.

3.3 | Experiment 3: High Expected Values and Nondescriptive Safe Options

The previous two experiments showed that participants strongly preferred uncertain gains over losses at low probabilities, and this preference diminished or even reversed at medium and high probabilities. Thus, extreme outcomes did not uniformly influence uncertainty preferences across probabilities, as predicted by our simple generalization of the EOR. Before discussing these results in more detail, we conducted another experiment to determine whether these results were solely due to differences in stimulus material between uncertain and certain options. Specifically, it was unclear whether our use of descriptive bar stimuli in Experiments 1 and 2 as the safe options somehow influenced how people weigh probabilities. We had initially used the descriptive reference stimuli as a way to accentuate the difference between the safe and uncertain options, but this could have impacted how people perceived the uncertain options relative to the safe ones. For example, our prior study indicated that people perceive descriptive stimuli as more certain than emojis (Elston, Mackenzie, and Mittelstädt 2021) when making probabilistic choices. Thus, we pre-registered and conducted a third experiment in which all stimuli, both safe and uncertain, were arbitrary images with no organizing principle, to investigate whether the pattern found in the previous experiments will replicate. Because we, like Ludvig, Madan, and Spetch (2013), observed the clearest effects with higher expected values and no “nothing” outcomes, we again opted to use that payoff scheme (see Table 3).

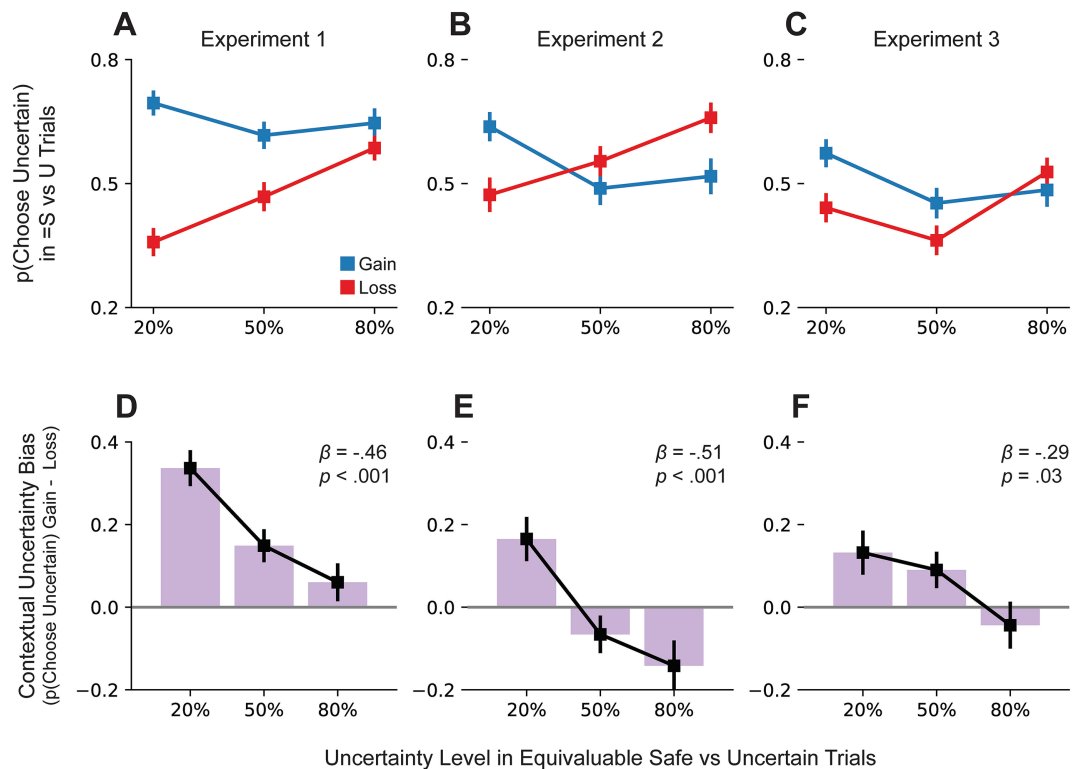


FIGURE 4 | Probability weightings revealed through choice patterns in =SvsU trials. Each column represents each experiment. (A–C) Equivaluable trials in which participants decided between a safe and an uncertain option with equal expected value. The x-tick values indicate the uncertainty level—for example, 20% indicates a choice between a uncertain option with a 20% hit rate and the safe option with the same expected value. See main text for statistical details. (D–F) Contextual differences in uncertainty preferences grouped by probability. The y-axis is the gain-loss context difference in p (choose uncertain) and the x-axis is the uncertainty level in the =SvsU trials. These are simply the differences between the red and blue curves in panels (A)–(C) at each uncertainty level. Markers denote mean \pm SEM.

3.3.1 | Participants

Sixty University of Tübingen undergraduate psychology students participated in this experiment ($M \pm \text{SEM}$ age: 21 ± 0.56 years old; 38 female). Six participants were excluded according to our pre-registered criterion.

3.3.2 | Task Validation: Assessing Sampling Error and Task Performance

As in Experiment 1 and 2, participants in Experiment 3 adequately sampled the uncertain options during the training blocks (all $p > 0.05$ when t -testing experienced probability against objective probability for a given option; see Figure 3A, bottom) selected the best option across all conditions when a best option was available (all one-sample t -tests with $ts(53) > 9.94$, $ps < 0.001$, and Cohen's $ds > 1.35$; see Figure 3B–D, bottom). These results indicate that participants learned the values associated with the choice options and that subsequent results cannot be explained as an artifact of sampling error.

3.3.3 | Core Results: Choice Patterns During Equivaluable (=SvsU) Trials

As in Experiments 1 and 2, we used rmANOVAs to assess choice behavior. The main effect of probability was significant, indicating that participants generally selected uncertain options more often at low and high probabilities than at medium probabilities

($F(2, 106) = 5.96$, $p = 0.004$, $\eta_p^2 = 0.10$). The main effect of context was marginally significant, reflecting that participants were more likely to select uncertain options in the gain context than in the loss context ($F(1, 53) = 3.77$, $p = 0.057$, $\eta_p^2 = 0.07$).

Critically, the interaction was again significant ($F(2, 106) = 3.23$, $p = 0.04$, $\eta_p^2 = 0.06$) (see Figure 4C). We again characterized the nature of the interaction effect by regressing contextual differences in uncertainty preferences against probability and again detected a significant, negative beta weight ($\beta = -0.29$, $p = 0.03$, see Figure 4F). The rmANOVA interaction result and the negative beta weight were critical as they indicated that our findings of overweighting of low probabilities and underweighting of high probabilities was primarily driven by the differential uncertainty associated with safe and uncertain options and cannot be explained due to differences in context, payoff scheme, or information format.

3.4 | Assessing the Roles of Probability, Information Format, and Expected Value on Probability Weighting Across the Three Experiments

We consistently observed that participants preferred uncertain gains more than uncertain losses when probabilities were low (20%) and often preferred uncertain losses over uncertain gains at high probabilities (80%). However, as seen in Figure 4, choice behavior during equivaluable trials also appeared to be

influenced by the expected values of the choice options (via the payoff schemes) and the information format of the safe option (whether it was descriptive or not). Next, we explored how these factors influenced probability weighting across experiments via a general linear regression model (GLM). The dependent variable was the gain-minus-loss difference of p (choose uncertain) in the =SvsU trials (see Figure 4D–F). In addition to probability, our other predictors were expected value and information format. We coded the expected value regressor as +1 for the experiments with higher expected values (Experiments 1 and 3) and –1 for Experiment 2, which had a lower expected value for each choice option. Similarly, we coded the information format regressor as +1 for Experiments 1 and 2, where the safe option was descriptive; the regressor was coded as –1 for Experiment 3, where the safe option was an arbitrary symbol. Coding the variables this way eased interpretation as positive values indicated a positive contribution to uncertainty preference. We included all interactions in the model.

As in our earlier analysis, we detected a significant, negative effect of probability ($\beta = -0.31$, $p < 0.001$), indicating that the same latent value function of over/underweighting low/high probabilities was present across all three experiments. Critically, the GLM revealed significant main effects of expected value ($\beta = 0.09$, $p < 0.001$) and information format ($\beta = 0.11$, $p < 0.001$). The expected value main effect indicates that increased expected values (as in Experiments 1 and 3) generally led to greater probability overweighting (as indicated by a preference for uncertain options in the gain domain and uncertainty-aversion in the loss domain). Similarly, the positively signed information format main effect indicates that participants preferred (and thus presumably overweighted) uncertain options when the safe option descriptively conveyed its payoff value.

We also observed a significant interaction of probability and information format ($\beta = -0.17$, $p < 0.001$). We note that this interaction was negatively signed, indicating that information format accentuated probability weighting (compare Experiments 1 and 3 in Figure 4 where the expected value was the same but information formats differed). It is also noteworthy that expected value did not interact with probability ($\beta = -0.06$, $p = 0.24$). No other terms in the GLM were significant (expected value \times information format: $\beta = -0.009$, $p = 0.76$; probability \times expected value \times information format: $\beta = 0.08$, $p = 0.1$). Together, these results suggest that expected value shifts the probability weighting curve up and down without modifying its slope whereas the information format of the safe option appears to modify the sharpness of the curve.

4 | General Discussion

Across three pre-registered experiments, we examined how extreme gains and losses influenced uncertainty preferences across a range of probabilities. Our aim was to test whether extreme outcomes would uniformly influence uncertainty preferences across probabilities, as predicted by our simple generalization of the EOR. Contrary to this hypothesis, we found that participants did not show a uniform preference for uncertain gains over uncertain losses across all probabilities. While we observed that

participants generally preferred uncertain options when there was the prospect of gain and avoided uncertain options when there was the prospect of loss, the largest effect was consistently observed at the 20% level. This suggests that at lower probabilities, uncertain options are perceived as especially better when seeking gains and especially worse when mitigating losses. Thus, our results suggest that extreme outcomes influence uncertainty preferences through interaction with probability weighting. Within our discussion, we elaborate on the potential theoretical implications of this unexpected empirical pattern for experience-based decision-making.

4.1 | Implications for the EOR

Our study integrates and extends previous research that investigated uncertainty preferences exclusively at the 50% level in the context of experience-based decisions involving gains and losses (Konstantinidis, Taylor, and Newell 2018; Ludvig et al. 2018; Ludvig, Madan, and Spetch 2013; Madan, Ludvig, and Spetch 2014, 2017). These studies have proposed that the saliency of extreme outcomes associated with the more uncertain options, as opposed to safe options, generally leads participants to prefer uncertain gains and to avoid uncertain losses. Indeed, when considering the average across probability levels, and at most individual probability levels in our study, we observed that participants preferred uncertain options where there was the prospect of gain and, in contrast, avoided uncertain options when there was the prospect of loss. Moreover, we observed the largest changes in uncertainty preferences in our first and third (as compared with the second) experiments, where the outcomes were more extreme. However, given that the same payoffs were used across all probabilities within each experiment, the EOR needs further refinement to account for why uncertain gains over losses were especially preferred at the 20% level, and less so—or even reversed—at other probability levels.

One possibility is that extreme outcomes associated with uncertain options accentuate underlying biases in probability weighting, whereby people treat probabilistic outcomes as more/less likely to occur than they really are. Thus, the EOR might be extended by assuming that extreme outcomes accentuate distortions in probability perception via distortions in a value-based learning process that incorporates the distribution from which outcomes are drawn. Specifically, the options with a greater outcome uncertainty yield a wider range of possible values, allowing for greater distortions in perceived value in the sense that wider distributions provide higher/lower bounds for inferred values. We theorize that such a distributional learning process interacting with an underlying probability weighting function could account for the choice patterns observed in our experiments. Below, we discuss how the specific pattern of probability weighting we observed is more consistent with a classic prospect theory rather than the prior sampling-based studies of experience-based choice.

4.2 | Relation to Experience-Based Choices in Sampling Paradigms

Our results stand in contrast to prior studies investigating experience-based choice using sampling paradigms (e.g.,

Camilleri and Newell 2011; Hertwig et al. 2004). Sampling paradigms typically involve presenting participants with two options (e.g., two colored rectangles), and usually, participants could sample as many times as they wished without consequence (gain or loss) before making a single consequential choice. After this consequential (paid) decision, a new set of options are presented, and the process begins anew. Studies using this paradigm often observe that people underweight high probabilities and overweight low probabilities, largely the opposite of what we found. Of course, our study differs in several respects from studies using sampling-like paradigms, so we can only speculate about the potential opposing effects of probability weighting. However, the most salient differences of the present study from these other studies appears to be that all of our choices were consequential and that participants experienced each choice problem hundreds of times, intermixed within blocks. Thus, these factors could conceivably play critical roles in decision-making.

The consequentialness of each decision has been shown to strongly impact choice behavior (Jessup, Bishara, and Busemeyer 2008). For example, Garcia et al. (2023) recently reported that participants who learned probabilities associated with arbitrary stimuli through consequential (rewarded) learning neglected those values when probed in blocks of trials in the absence of reward feedback. Thus, experience-based choice appears to be critically modulated by the type of ongoing feedback that led to the formation and refinement of learned outcome estimates.

Notably, some prior experience-based choice studies have implemented consequential decisions but still observed a pattern consistent with other sampling-based studies (e.g., Barron and Erev 2003). However, these studies still seemed more similar to sampling-based paradigms than our paradigm. For example, in Barron and Erev's (2003) Experiment 1, participants decided between the same options for their entire experimental session, whereas in the present study, the comparison between two options could vary trial-to-trial. Thus, one may speculate that participants pay more attention to ongoing feedback when there is some variability between the decision options.

Moreover, the structure of sampling-based studies differs markedly from our task: In sampling paradigms, participants typically encounter each pair of choice options once, sampling their outcomes in order to learn their payoff characteristics. This is important because, unlike our task, the act of sampling is in effect when participants are learning. A number of studies have shown that the effects of sampling-based studies (e.g., underweighting low probabilities) are almost entirely explained as an artifact of sampling error (Camilleri and Newell 2011; Fox and Hadar 2006; Hertwig et al. 2004). Sampling error arises due to participants not sampling the entire outcome distribution of an option prior to choosing. For instance, when sampling a rare option, 78% of Hertwig et al.'s (2004) participants never experienced the rare outcome before making their decision. From the participant's perspective, they would have learned that the uncertain option has a payoff probability of zero, which is different from probability underweighting, *per se*.

Furthermore, when Camilleri and Newell (2011) forced their participants to sample each option many times before choosing,

there was no evidence that participants underweighted low probabilities. However, Wulff, Mergenthaler-Canseco, and Hertwig's (2018) meta-analysis suggested that the underweighting of low probability experience options is only attenuated, not entirely extinguished, by mitigating sampling error. Given that Wulff et al.'s definition of probability underweighting relied on the difference between description and experience options, one may speculate that their findings, however, could also be explained by an overweighting of description-based options (and no effect on experience options, as would be consistent with Camilleri and Newell's results). In any case, the very structure of sampling paradigms where participants encounter each choice problem once appears to be a critical factor in that the results are susceptible to sampling error.

We aimed to minimize the possibility that our results stem from sampling error by including training blocks. Of course, it is still possible that the differences in choice behavior we observed—for example, participants selecting the better option more often in unequal trials—may lead to varying experiences of outcomes with options, which in turn could recursively influence subsequent choice behavior. Additionally, the incorporation of forced-choice trials (presenting only one option) could have further ensured that the outcomes of all options were more regularly experienced throughout the experiment (e.g., Konstantinidis, Taylor, and Newell 2018). Yet, even if differences in the frequency of experiencing outcomes (and probabilities) may contribute to the observed pattern, it remains unclear why this would lead to a qualitatively different pattern than that typically observed in sampling paradigms where sampling error was not controlled.

4.3 | Implications for Prospect Theory

Instead of the choice pattern associated with studies using sampling-based studies of experiential choice, our results are more consistent with earlier studies where probabilities were explicitly described to participants (Kahneman and Tversky 1979; Ruggeri et al. 2020). However, even if it seems that probability weighting as described within prospect theory largely captures the present choice behavior, additional learning-related factors beyond the scope of traditional prospect theory must also be at play because participant behavior markedly responded to manipulating whether outcomes were more or less extreme and whether the safe, reference options were more or less informative. Moreover, one limitation of prospect theory is that it is a descriptive rather than generative model—it elegantly explains behavior in a highly parameterized way but does not offer an account of the cognitive processes that undergird those parameters. In this context, it seems particularly useful to consider models relying on reinforcement learning (RL), whereby the agent learns the expected value of an option through trial and error (Sutton and Barto 2018). RL-based models have proven successful in elucidating experience-based decision-making that involve uncertainty (Niv et al. 2012; Spektor et al. 2019). Thus, RL seems a useful starting framework to consider mechanistically how biased subjective values could arise in conjunction with other parameters—for example, utility functions associated with prospect theory—to potentially explain decision-making within the present task environment. Thus, RL seems

a useful framework to consider mechanistically how biased subjective values could arise as a product of individual variability in learning processes.

Combining aspects of prospect theory with learning-related principles may also offer an interesting way to explain why the results of the present study partly differ from our prior study (Elston, Mackenzie, and Mittelstädt 2021). As mentioned in the introduction, in our prior study, we observed that participants generally preferred the more uncertain options across all probabilities when seeking gains and avoided them when mitigating losses—consistent with a simplified extension of the EOR as shown in Figure 1. However, in this prior study, uncertain options with a certain probability (e.g., a 20% emoji) were not compared with safe options with the same expected value as in the present study, but instead with options with the same probability where a feature more explicitly described the probability (e.g., a 20% filled rectangle for a 20% option). As proposed by prospect theory, value-based decisions depend on internal reference points to classify potential outcomes as gains or losses. Thus, it may be that decision contexts involving safe options induce a different set of reference points compared with when no sure-bet option exists.

Indeed, recent experimental and computational work in humans (Hunter and Daw 2021; Hunter and Gershman 2018) found that different reference points can be induced and manipulated in the same experiment and that this can strongly influence behavior. These results suggest that subjective valuations of uncertain options may be relative to reference points and that these reference points can change over time and are likely the basis for learned outcome distributions to be formed around. Future studies might investigate the precise role that outcome uncertainty plays in the formation of reference points.

Moreover, unlike the previous study, the probability weighting observed in our current study cannot solely be attributed to differences in the information format between the safe and uncertain options. Specifically, one might argue that the bar stimuli used in the first two experiments could be perceived as descriptive options, suggesting that the observed choice patterns, when compared with emojis as experiential options, reflect a form of the description experience (DE) gap due to differing information formats. However, in Experiment 3, we used arbitrary stimuli (emojis) for both safe and uncertain options and used the same payoff scheme as in Experiment 1, but we observed a similar pattern. Therefore, our combined findings suggest that the critical factor is the level of uncertainty associated with the safe and uncertain options.

While the information format itself did not directly cause the overweighting of low probabilities, comparisons across experiments suggest that it influences probability weighting by exaggerating the overweighting or underweighting of low and high probabilities. Specifically, Experiment 1 (with descriptive safe options) elicited a broader range of probability distortions compared to Experiment 3 (featuring arbitrary emojis as safe options). Beyond the variance in information format, an alternative explanation could be that Experiment 3 imposed higher memory requirements due to its use of arbitrary emojis exclusively. This increased cognitive load could have led to more

erratic decision-making. However, as there is no indication that memory demands differentially affected choice behavior in non-equiprobable trials across experiments where there was a clearly optimal option (e.g., “pure safe” and “pure uncertain” trials), we maintain that the differences between experiments primarily reflect differences in probability weighting due to the uncertainty associated with the safe options.

4.4 | Insights From Decision Times

Finally, when theorizing about the mechanisms underlying experience-based consequential choices, considering the time required to make a decision also seems useful. Choice reaction times (RTs) are thought to reflect the duration of cognitive processes needed to make a decision, such as perception, valuation, and deliberation, with extended decision times indicating increased cognitive demands. As reported in the Appendix A, the RT data were quite consistent across experiments. First, responses to the more uncertain options (with multiple outcomes) were generally slower, in line with previous studies showing that multi-attribute options result in longer decision times (e.g., Busemeyer et al. 2019; Donkin and Brown 2018). Second, responses were generally longer during loss mitigation, perhaps reflecting risk aversion due to overweighting the prospect of loss relative to gain (cf. Kahneman and Tversky 1979). Third, and most interesting, decision times were influenced by the joint impact of probability and context of outcomes in the =SvsU equivalent trials. Specifically, we found that loss-mitigating decisions became slower as the probability of those losses increased, and conversely, gain-maximizing choices became faster as the probability of gain increased. This pattern is consistent with the notion that the outcomes of uncertain options could be internally represented as a distribution and elongated decision times could reflect additional processing required to integrate over that distribution. These results highlight the potential of considering the time needed to make a decision as a complementary measure to overt choice patterns in order to gain deeper insights into the cognitive processes underlying decisions involving uncertainty.

5 | Conclusion

The present study examined how gains and losses influenced experience-based choices between equivalent certain and uncertain extreme options across a range of probabilities. Contrary to a simple generalization of an “extreme-outcome rule,” there was no uniform preference for uncertain gains across all probabilities. Instead, extreme outcomes influenced uncertainty preferences through interaction with probability weighting. Specifically, participants particularly preferred uncertain options at low probabilities (20%) when there was the prospect of gain and avoided them when there was the prospect of loss, with these preferences being reduced or even reversed at medium (50%) and high (80%) probabilities. We also found that uncertainty preferences were influenced by outcome extremity and the relative certainty associated with safe reference options. We conclude that extreme outcomes accentuate the overweighting of low probabilities and the underweighting of high probabilities in experience-based choice.

Acknowledgments

Data collection was supported by internal funding from the Department of Psychology at the University of Tübingen. This research was funded as part of the Excellence Strategy of the German Federal and State Governments. Open Access funding enabled and organized by Projekt DEAL.

Data Availability Statement

The data that support the findings of this study are openly available in Open Science Framework at https://osf.io/7qbp8/?view_only=57351bf1248245ffbadb2db2d34ec4f9.

References

- Barron, G., and I. Erev. 2003. "Small Feedback-Based Decisions and their Limited Correspondence to Description-Based Decisions." *Journal of Behavioral Decision Making* 16, no. 3: 215–233. <https://doi.org/10.1002/BDM.443>.
- Barron, G., and G. Ursino. 2013. "Underweighting Rare Events in Experience Based Decisions: Beyond Sample Error." *Journal of Economic Psychology* 39: 278–286. <https://doi.org/10.1016/J.JOEP.2013.09.002>.
- Busemeyer, J. R., S. Gluth, J. Rieskamp, and B. M. Turner. 2019. "Cognitive and Neural Bases of Multi-Attribute, Multi-Alternative, Value-Based Decisions." *Trends in Cognitive Sciences* 23, no. 3: 251–263. <https://doi.org/10.1016/J.TICS.2018.12.003>.
- Camilleri, A. R., and B. R. Newell. 2011. "When and Why Rare Events Are Underweighted: A Direct Comparison of the Sampling, Partial Feedback, Full Feedback and Description Choice Paradigms." *Psychonomic Bulletin and Review* 18, no. 2: 377–384. <https://doi.org/10.3758/s13423-010-0040-2>.
- Camilleri, A. R., and B. R. Newell. 2013. "Mind the Gap? Description, Experience, and the Continuum of Uncertainty in Risky Choice." *Progress in Brain Research* 202: 55–71. <https://doi.org/10.1016/B978-0-444-62604-2.00004-6>.
- de Leeuw, J. R. 2015. "jsPsych: A JavaScript Library for Creating Behavioral Experiments in a Web Browser." *Behavior Research Methods* 47: 1–12. <https://doi.org/10.3758/s13428-014-0458-y>.
- Donkin, C., and S. D. Brown. 2018. "Response Times and Decision-Making." *Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience* 1–33. <https://doi.org/10.1002/9781119170174.EPCN509>.
- Elston, T. W., I. G. Mackenzie, and V. Mittelstädt. 2021. "Outcome Uncertainty Influences Probability Perception and Risk Attitudes." *Royal Society Open Science* 8, no. 10. <https://doi.org/10.1098/RSOS.210307>.
- Fox, C. R., and L. Hadar. 2006. "'Decisions From Experience' = Sampling Error + Prospect Theory: Reconsidering Hertwig, Barron, Weber, & Erev (2004)." *Judgment and Decision Making* 1, no. 2: 159–161. <https://doi.org/10.1017/S1930297500002370>.
- Garcia, B., M. Lebreton, S. Bourgeois-Gironde, and S. Palminteri. 2023. "Experiential Values Are Underweighted in Decisions Involving Symbolic Options." *Nature Human Behaviour* 7, no. 4: 611–626. <https://doi.org/10.1038/s41562-022-01496-3>.
- Hau, R., T. J. Pleskac, J. Kiefer, and R. Hertwig. 2008. "The Description–Experience Gap in Risky Choice: The Role of Sample Size and Experienced Probabilities." *Journal of Behavioral Decision Making* 21, no. 5: 493–518. <https://doi.org/10.1002/BDM.598>.
- Hertwig, R., G. Barron, E. U. Weber, and I. Erev. 2004. "Decisions From Experience and the Effect of Rare Events in Risky Choice." *Psychological Science* 15, no. 8: 534–539. <https://doi.org/10.1111/j.0956-7976.2004.00715.x>.
- Hertwig, R., and I. Erev. 2009. "The Description–Experience Gap in Risky Choice." *Trends in Cognitive Sciences* 13, no. 12: 517–523. <https://doi.org/10.1016/j.tics.2009.09.004>.
- Hunter, L. E., and N. D. Daw. 2021. "Context-Sensitive Valuation and Learning." *Current Opinion in Behavioral Sciences* 41: 122–127. <https://doi.org/10.1016/J.COBEHA.2021.05.001>.
- Hunter, L. E., and S. J. Gershman. 2018. "Reference-Dependent Preferences Arise From Structure Learning." *BioRxiv*: 252692. <https://doi.org/10.1101/252692>.
- Jessup, R. K., A. J. Bishara, and J. R. Busemeyer. 2008. "Feedback Produces Divergence From Prospect Theory in Descriptive Choice." *Psychological Science* 19, no. 10: 1015–1022. <https://doi.org/10.1111/J.1467-9280.2008.02193.X>.
- Kahneman, D., and A. Tversky. 1979. "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica* 47, no. 2: 263. <https://doi.org/10.2307/1914185>.
- Kellen, D., T. Pachur, and R. Hertwig. 2016. "How (In)Variant Are Subjective Representations of Described and Experienced Risk and Rewards?" *Cognition* 157: 126–138. <https://doi.org/10.1016/J.COGNITION.2016.08.020>.
- Konstantinidis, E., R. T. Taylor, and B. R. Newell. 2018. "Magnitude and Incentives: Revisiting the Overweighting of Extreme Events in Risky Decisions From Experience." *Psychonomic Bulletin and Review* 25, no. 5: 1925–1933. <https://doi.org/10.3758/S13423-017-1383-8>.
- Ludvig, E. A., C. R. Madan, N. McMillan, Y. Xu, and M. L. Spetch. 2018. "Living Near the Edge: How Extreme Outcomes and Their Neighbors Drive Risky Choice." *Journal of Experimental Psychology. General* 147, no. 12: 1905–1918. <https://doi.org/10.1037/XGE0000414>.
- Ludvig, E. A., C. R. Madan, and M. L. Spetch. 2013. "Extreme Outcomes Sway Risky Decisions From Experience." *Journal of Behavioral Decision Making* 27, no. 2: 146–156. <https://doi.org/10.1002/BDM.1792>.
- Ludvig, E. A., and M. L. Spetch. 2011. "Of Black Swans and Tossed Coins: Is the Description–Experience Gap in Risky Choice Limited to Rare Events?" *PLoS ONE* 6, no. 6: e20262. <https://doi.org/10.1371/journal.pone.0020262>.
- Madan, C. R., E. A. Ludvig, and M. L. Spetch. 2014. "Remembering the Best and Worst of Times: Memories for Extreme Outcomes Bias Risky Decisions." *Psychonomic Bulletin & Review* 21, no. 3: 629–636. <https://doi.org/10.3758/S13423-013-0542-9>.
- Madan, C. R., E. A. Ludvig, and M. L. Spetch. 2017. "The Role of Memory in Distinguishing Risky Decisions From Experience and Description." *Quarterly Journal of Experimental Psychology* 70, no. 10: 2048–2059. <https://doi.org/10.1080/17470218.2016.1220608>.
- Niv, Y., J. A. Edlund, P. Dayan, and J. P. O'Doherty. 2012. "Neural Prediction Errors Reveal a Risk-Sensitive Reinforcement-Learning Process in the Human Brain." *Journal of Neuroscience* 32, no. 2: 551–562. <https://doi.org/10.1523/JNEUROSCI.5498-10.2012>.
- Ruggeri, K., S. Ali, M. L. Berge, et al. 2020. "Replicating Patterns of Prospect Theory for Decision Under Risk." *Nature Human Behaviour* 4, no. 6: 622–633. <https://doi.org/10.1038/s41562-020-0886-x>.
- Seabold, S., and J. Perktold. 2010. "Statsmodels: Econometric and Statistical Modeling With Python." 9th Python in Science Conference.
- Spektor, M. S., S. Gluth, L. Fontanesi, and J. Rieskamp. 2019. "How Similarity Between Choice Options Affects Decisions From Experience: The Accentuation-of-Differences Model." *Psychological Review* 126, no. 1: 52–88. <https://doi.org/10.1037/REV0000122>.
- Sutton, R., and A. Barto. 2018. *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA: MIT Press.

Vallat, R. 2018. "Pingouin: Statistics in Python." *Journal of Open Source Software* 3, no. 31. <https://doi.org/10.21105/joss.01026>.

Wulff, D. U., M. Mergenthaler-Canseco, and R. Hertwig. 2018. "A Meta-Analytic Review of Two Modes of Learning and the Description-Experience Gap." *Psychological Bulletin* 144, no. 2: 140–176. <https://doi.org/10.1037/bul0000115>.

Supporting Information

Additional supporting information can be found online in the Supporting Information section.

Appendix A

A.1 | Exploring Choice Reaction Times

To better understand the influences of our experimental factors on choice behavior, we also analyzed choice reaction times (RTs). Accumulator-based models of choice RT predict that multi-attribute options elicit longer RTs (Busemeyer et al. 2019; Donkin and Brown 2018), given that uncertain options, by definition, have more than one outcome and each of these outcomes could be considered an attribute of the choice option. Thus, we hypothesized that choices involving uncertain options would elicit longer RTs as compared with decisions between purely safe/certain options. Similarly, our prior results (Elston, Mackenzie, and Mittelstädt 2021) indicated that RTs are longer in the loss domain than in the gain domain. We were interested in whether such an effect would be visible in our present task design and whether RTs would additionally be modulated by expected value. We assessed the RT data with rmANOVAs with the within-subject factors of domain and uncertainty (see Figure S1) and obtained highly consistent results across all three experiments.

A.2 | RTs During Training and Pure Trials are Influenced by Uncertainty and Domain

During the training and pure-safe/uncertain blocks, participants in all three experiments responded faster when choosing between two "safe" options as compared with choosing between two uncertain options (Experiment 1: all $F_s(1, 44) > 61.53$, all $p_s < 0.001$, all $\eta_p^2 > 0.58$; Experiment 2: all $F_s(1, 44) > 65.00$, all $p_s < 0.001$, all $\eta_p^2 > 0.59$; Experiment 3: all $F_s(1, 53) > 4.09$, all $p_s < 0.05$, all $\eta_p^2 > 0.07$). We also found that participants were generally faster in the gain than the loss domain (Experiment 1: all $F_s(1, 44) > 5.15$, $p_s < 0.03$, $\eta_p^2 > 0.10$; Experiment 2: all $F_s(1, 44) > 6.54$, all $p_s < 0.01$, all $\eta_p^2 > 0.13$; Experiment 3: all $F_s(1, 53) > 8.40$, all $p_s < 0.01$, $\eta_p^2 > 0.14$).

A.3 | RTs During ≠SvsU Trials are Influenced by Domain

Consistent with the training and pure trials, RTs during ≠SvsU trials were faster in the gain domain as compared to the loss domain (Experiment 1: $F(1, 44) = 5.15$, $p = 0.03$, $\eta_p^2 = 0.10$; Experiment 2: $F(1, 44) = 6.29$, $p = 0.02$, $\eta_p^2 = 0.12$; Experiment 3: $F(1, 53) = 5.50$, $p = 0.02$, $\eta_p^2 = 0.09$). Interestingly, participants in Experiments 1 and 3, where the outcomes were more extreme, were faster when an uncertain option was the better option (Experiment 1: $F(1, 44) = 9.31$, $p < 0.01$, $\eta_p^2 = 0.17$; Experiment 3: $F(1, 53) = 9.69$, $p < 0.01$, $\eta_p^2 = 0.15$). We did not detect such an effect in Experiment 2, which used relatively less extreme outcomes.

A.4 | RTs During =SvsU Trials are Modulated by Expected Value

Across all three experiments we detected a significant domain×uncertainty level interaction (Experiment 1: $F(2, 88) = 23.09$, $p < 0.001$, $\eta_p^2 = 0.34$; Experiment 2: $F(2, 88) = 50.02$, $p < 0.001$, $\eta_p^2 = 0.53$; Experiment 3: $F(2, 106) = 54.72$, $p < 0.001$, $\eta_p^2 = 0.51$). The main effects of domain and uncertainty level were inconsistent and mostly insignificant across experiments. Taken together, these results suggest that

expected value modulates decision dynamics because participants got faster as the options got better in the gain domain and participants got slower as the options got worse in the loss domain.