What is an extreme outcome in risky choice?

Joel Holwerda (j.holwerda@unsw.edu.au)
School of Psychology
University of New South Wales
NSW, 2052, Australia

Ben R. Newell (ben.newell@unsw.edu.au)
School of Psychology
University of New South Wales
NSW, 2052, Australia

Abstract

Numerous experiments have suggested that extreme outcomes are disproportionately influential when we make decisions involving risk, but there is less consensus on what it actually means to be extreme. Existing accounts broadly fall into two categories: those that suggest that the best and worst outcomes are uniquely influential and those that suggest that outcomes become more influential with increasing deviation from the centre of the distribution. We conducted two experiments that aimed to tease apart these explanations. Although there was some evidence that the distance from the centre influences memory, neither account was able to fully explain the choices made by participants. This finding has implications for the viability of these explanations as well as for the generalisability of the effect and the interpretation of the method used to assess memory.

Keywords: decision-making; extreme-outcome effect; peak-end rule; memory; risky choice

Introduction

When answering a question like "how was your weekend?", our responses are usually full of highs and lows. We'll talk about the perfect weather at the beach or a disastrous attempt at baking a cake, while omitting the majority of the experience that was merely "fine". These highs and lows have a disproportionate influence on our behaviour and importantly this depends on more than just aspects of the experience in isolation (Schmidt, 1991). A trip to the beach that was remembered one weekend might have been almost completely forgotten the weekend that you went skydiving from 14,000 feet—what matters is its relative position within its context.

The influence of these highs and lows extends far beyond the way you recount your weekend and there is growing evidence that they play a role in the way people make decisions involving risk. This was initially demonstrated in a series of experiments that required participants to make choices between multiple pairs of options and to learn about the outcomes associated with each option through experience (Ludvig, Madan, & Spetch, 2014). In the simplest version of this task, participants made interspersed choices between a low-value pair and a high-value pair that each consisted of a safe and risky option that had equal expected value. This design ensured that the risky low-value option sometimes resulted in the worst outcome and the risky high-value option sometimes resulted in the best. Although the safe and risky options within each pair had equal expected value, participants displayed greater willing to choose the risky option when deciding between the high-value pair where it allowed the possibility of acquiring the best outcome than in the low-value pair where it might have resulted in the worst outcome.

Based on these experiments, Ludvig et al. (2014) devised an *extreme-outcome rule* that the outcomes with the highest or lowest values in a given context are more influential than they should be relative to the intermediate outcomes. They also provided evidence that these extreme outcomes are more likely to come to mind and that their frequency is overestimated (Madan, Ludvig, & Spetch, 2014). This paints a picture in which the best and worst outcomes are uniquely advantaged in memory and that this leads to a unique influence over our decisions, but this is not the only interpretation that is compatible with the evidence. All of these experiments have involved pairs of options that were symmetrically positioned so that the best and worst outcomes were equidistant from the centre, and therefore, identical results would have been acquired if outcomes, instead, become more influential the further they deviate from the centre of the value distribution. In other words, it is possible that instead of unique extreme outcomes, what matters is degrees of extremity.

The idea that extreme outcomes are unique was heavily influenced by the peak-end rule, which was devised to explain a curious phenomenon regarding the way people remember events: they seem to neglect the duration of experiences and instead evaluate events based on a small number of salient aspects, including the best or worst moments (for a review, see Fredrickson, 2000). The extreme-outcome rule appears, at least superficially, to be little more than a generalisation of this rule to include decisions involving risk, but there are also non-trivial differences between phenomena that they seek to explain. The peak-end rule was based upon experiences that were continuous, affect-laden, and for which there was no easily applicable aggregation strategy, such as averaging. Contrast this with the discrete numerical outcomes used in the tasks conducted by Ludvig et al. (2014) and the connection between the two might seem somewhat more equivocal—or at least not entirely self-evident.

This is particularly pertinent given that there are several alternate theories in which the influence of outcomes is proportional to their distance from the centre of the distribution of experienced values. One notable account suggests that

biasing memory retrieval towards the edges of the distribution offers a rational strategy for allocating limited resources (Lieder, Griffiths, & Hsu, 2018). Another suggests that interference between items leads to better memory for items that are towards the edges of a distribution because they possess fewer neighbours (Brown, Neath, & Chater, 2007). This model has successfully explained results spanning a large number of domains—although there is some evidence that it might not be driving the effect in risky choice (Ludvig, Madan, McMillan, Xu, & Spetch, 2018). Yet another account suggests that increased memory for extreme outcomes is caused by greater prediction error for outcomes further from the average experienced outcome (Rouhani, Norman, & Niv, 2018).

There are numerous accounts of the disproportionate influence of extreme outcome and the peak-end rule is clearly not the only theory relevant to risky choice. The aim of this paper was to tease apart the explanations that consider extreme outcomes as unique and those that characterise extremity as increasing with distance from the centre.

Experiment 1

The aim of the first experiment was to examine whether preferences regarding the best and worst outcomes are influenced by their distance from the centre of the distribution of experienced outcomes (see the design section for more details). Experiment 1a and 1b were identical except that in the latter, the context options were presented twice as often and the outcomes associated with them were pushed closer towards the best or worst outcome in order to strengthen the manipulation. The difference between conditions in their average outcome was doubled from around 20 points in Experiment 1a to 40 points in Experiment 1b—equivalent to half the range of the experienced outcomes.

Methods

Participants A total of 210 undergraduate psychology students from UNSW Sydney participated in Experiment 1 (80 in Experiment 1a and 130 in Experiment 1b). The average age was 19.6 years (SD = 3.5) and 146 participants were female. In addition to receiving course credit, they were able to earn a small amount of money depending on their performance in the task (M = AU$5.10, SD = AU$2.74).

Design The decision task required participants to make a series of choices between pairs of coloured squares representing options (see Figure 1). Similarly to the experiments conducted by Ludvig et al. (2014), a total of four options were presented in an interspersed fashion—two pairs, each consisting of a risky option and a safe option that always resulted in the expected value of the associated risky option (see Table 1). Participants were not given information about the distribution of outcomes associated with each option but instead were able to learn about the options by clicking on them and receiving feedback on the number of points won. No feedback was given for the option that was not chosen.

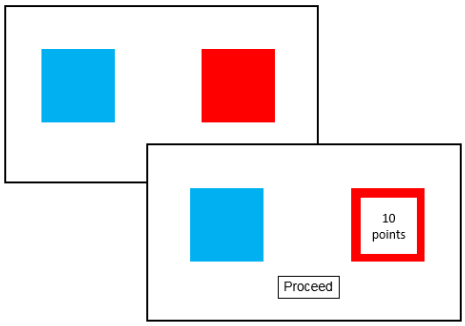

Figure 1: *Example of a standard trial in Experiment 1.*

Table 1: The number of points associated with each option in Experiment 1. Each of the outcomes associated with a risky option occurred with equal probability. The greater proportion of Context trials in Experiment 1b contributed to the increased difference between the averages of the conditions.

		Low average		High average	
		Safe	**Risky**	**Safe**	**Risky**
Experiment 1a	**Target**	50	10/90	50	10/90
	Context	30	20/40	70	60/80
Experiment 1b	**Target**	50	10/90	50	10/90
	Context	20	15/25	80	75/85

One of the pairs (target options) was identical across the two conditions. The risky target option resulted in either the best outcome (90 points) or the worst outcome (10 points) with equal probability. Choices between this option and the equivalent safe option were used to determine whether the best and worst outcomes exerted a different degree of influence depending on their distance from the centre. A second pair (context options) was included solely to manipulate the distribution of points experienced during the task so that the centre was closer to either the best or worst outcome.

In the Low-average condition, the outcomes associated with the context options were adjacent to the worst outcome (e.g., 15, 20, and 25 points) so that the worst outcome was much closer to the average than the best outcome. Likewise, in the High-average condition, these outcomes were adjacent to the best outcome (e.g., 75, 80, and 85 points) so that the best outcome was much closer to the average than the worst outcome. If the distance between an outcome and the centre of the distribution determines whether it is disproportionately influential, participants should have chosen the risky target option more often in the Low-average condition where the best outcome is an outlier than in the High-average condition where the worst outcome is an outlier.

In addition to trials that involved choices between either the target options or context options, there were also catch trials that involved a choice between one target option and

one context option. Finally, there were a number of single-option trials designed to prevent participants from completely avoiding an option that was initially unfavourable. Participants completed five blocks that each consisted of 48 trials in Experiment 1a (12 target and context trials, 16 catch trials, and eight single-option trials) or 60 trials in Experiment 1b (12 target trials, 24 context trials, 16 catch trials, and eight single-option trials).

Procedure Upon entering the laboratory, participants were told that the experiment involved a decision task in which they would have the opportunity to earn real money based on the choices they made. At the start of the task, detailed instructions were presented on the computer screen including that one decision would be randomly selected following completion of the experiment and that they would be paid $1 for every ten points earned on that choice. Participants then completed the decision task followed by two measures of memory: participants were asked to report the outcome that came to mind first and estimate the frequency of each outcome when presented with the stimulus associated with each outcome in random order.

Analysis All posterior distributions were determined by Hamiltonian Monte Carlo using the brms package in R (Bürkner, 2017). Weakly regularising priors were used for each parameter: for the logistic regression models, Gaussian distributions with a mean of 0 and standard deviation of 1 were used for the intercept and slope parameters and a half-Cauchy distribution with a location of 0 and scale of 1 was used for the standard deviation parameter in the hierarchical model of choices. For the percentage estimates, a Gaussian distribution with a mean of 0 and standard deviation of 25 was used for the intercept and slope parameters. A half-Cauchy distribution with a location of 0 and scale of 25 was used for the standard deviation parameter. All parameters had an effective sample size greater than 10000 and an $\hat{R} < 1.01$ indicating adequate chain convergence.

Results and discussion

The choices that participants made were not consistent with the distance from the average conceptualisation of extremity, which predicted that they would be more willing to choose the risky option when the best outcome rather than the worst outcome was an outlier (see Figure 2). To assess this hypothesis, we used Bayesian hierarchical logistic regression predicting the option chosen when presented with the target pair, with experiment and condition (Low-average/High-average) as fixed predictor variables and varying intercepts for each participant. The mean posterior log-odds for condition were 0.32 (95% CI = -0.13, 0.78) providing some evidence for a small difference in the opposite direction than was predicted by the distance from the average hypothesis. There was weaker evidence of an interaction between experiment and condition: the mean posterior log-odds were -0.15 (95% CI = -0.73, 0.43). If anything, this would suggest that the difference between conditions—in the wrong direction—was greater in Experiment 1b even though the manipulation was designed to be stronger.

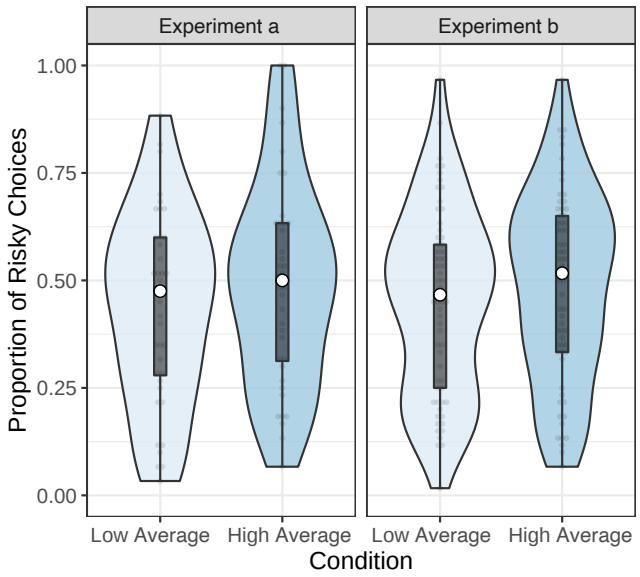

Figure 2: *The proportion of choices for the target risky option (10/90 points) when it was paired with the safe option that had equal expected value (50 points) in Experiment 1.*

In stark contrast, however, with the results from the decision task, the responses to both memory measures suggested that outcomes further from the centre of the distribution were disproportionately influential. To assess responses to the first-to-mind task (see Figure 3), we used Bayesian logistic regression predicting whether participants reported the extreme or non-extreme outcome as coming to mind first when presented the coloured square associated with the risky target option, with experiment and condition as fixed predictor variables. Responses that were not a valid experienced outcome for the presented option were excluded from this analysis. Starting with a model that included only the intercept, the mean posterior log-odds for the intercept were -0.77 (95% CI = -1.07, -0.48) indicating that, on average, the outcome further from the centre of the distribution was considerably more likely to be reported than the closer outcome. For the model including all parameters, the mean posterior log-odds for condition (Low-average/High-average) were 0.76 (95% CI = -0.02, 1.55) providing some evidence that this influence of the extreme outcome was stronger in the Low-average than the High-average condition. The mean posterior log-odds for experiment were -0.49 (95% CI = -1.27, 0.29) providing some evidence that the effect of extreme outcomes was stronger in Experiment 1b.

To assess the results of the percentage estimation task (see Figure 4), we used Bayesian linear regression predicting the difference between the percentage estimates given for the extreme and non-extreme outcomes when presented with the

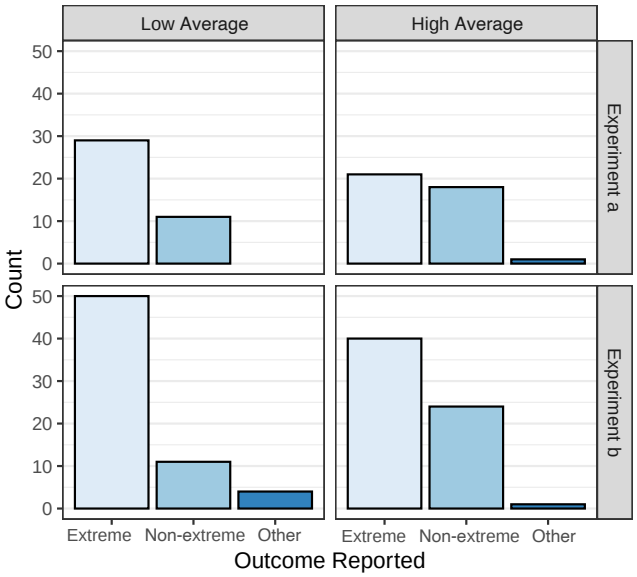

Figure 3: *The number of participants that responded with each outcome as coming to mind first when presented with the risky target option in Experiment 1.*

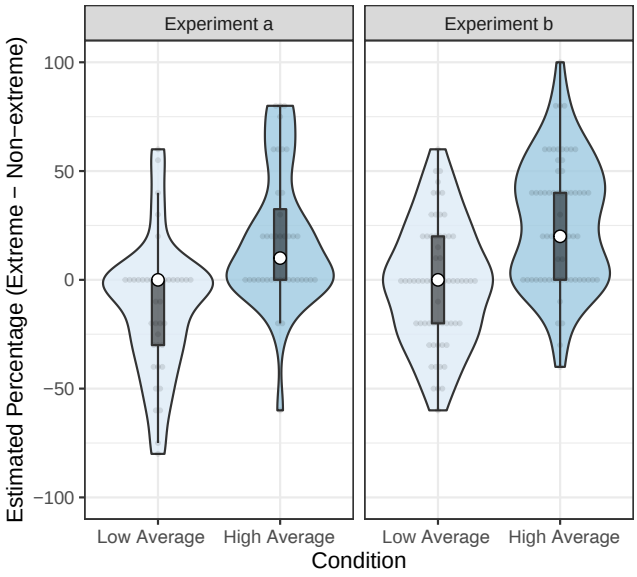

Figure 4: *The difference between percentage estimates for the extreme and non-extreme outcomes when presented with the risky target option in Experiment 1*

risky target option, with experiment and condition as fixed predictor variables. Again, starting with the model that included only the intercept, the mean posterior estimate for the intercept was 8.64 (95% CI = 4.14, 13.14) suggesting that, on average, the outcome further from the centre was reported as occurring more frequently. In the model that included all parameters, however, there was evidence that this was largely driven by a bias towards extreme outcomes in the High-value condition: the mean posterior estimate for condition was 28.84 (95% CI = 16.77, 40.98). There was also some evidence that the effect of extremity was stronger in Experiment 1a: the mean posterior estimate for experiment was 9.34 (95% CI = -1.60, 20.30)

When interpreted together, the results of the decision and memory tasks present us with a puzzle. If anything, participants showed a slight preference for the risky target option when the worst rather than best outcome was an outlier while simultaneously showing increased memory for outcomes further from the centre of the distribution. Most explanations of the influence of extreme outcomes in risky choice suggest that increased memory for extremes causes them to disproportionately influence choices. This was clearly not the case in this experiment. It is possible that different mechanisms drive responses in the choice and memory tasks—something overlooked due to the homogeneity in previous experimental designs in which the best and worst outcomes were always equidistant for the centre. Or perhaps, something about the asymmetrical distributions of outcomes used in the current experiment led to a preference towards the risky option in the High-average condition that obscured the effect of extremity.

Experiment 2

To address this potential for a difference in designs being responsible for the pattern of results found in the first experiment, the second experiment aimed to manipulate distance from the centre in a task that more closely resembled the experiments conducted by Ludvig et al. (2014) in which there were always Low- and High-value pairs of options that were equidistant from the centre. To achieve this, we manipulated the variance of the outcomes associated with the risky target options whilst ensuring that they were not the best or worst outcomes experienced in the experiment. This also made it possible to directly examine whether participants would be more willing to choose the High-value risky option, even though it never resulted in the best outcome, than the Low-value risky option, which never resulted in the worst outcome. This finding would cast doubt upon the idea that the best and worst outcomes are uniquely influential.

Methods

Participants A total of 130 undergraduate psychology students from UNSW Sydney participated in Experiment 2. The average age was 19.7 years (SD = 4.0) and 78 participants were female. In addition to receiving course credit, they were able to earn a small amount of money depending on their performance in the task (M = AU$4.00, SD = AU$0.00).

Design Similarly to the first experiment, participants made a series of choices between pairs of safe and risky options (see Table 2). Both conditions included a Low-value and a High-value target pair of options. In the Low-variance condition, the better and worse outcomes associated with each of the target risky options differed by only ten points—they

were both a similar distance from the centre of the distribution (e.g., 20 and 30 points). In the High-variance condition, these outcomes were three times as far apart, with one situated close to the centre and the other close to the edge (e.g., 10 and 40 points). Therefore, if distance from the centre underlies the disproportionate influence of extreme outcomes, the higher risk preferences for the High-value than Low-value pair should be stronger in the High-variance condition because the difference in distance between the better and worse outcomes in each pair is greater.

A third (extreme) pair of options was included that resulted in the best and worst outcomes (0 and 100 points) to surround the intermediate outcomes. This pair was not used in the analysis but instead was included to ensure that the range of outcomes was constant and to prevent the outcomes associated with the target options from being the best or worst outcomes experienced in the experiment. Participants completed five blocks, each consisting of 72 trials (12 Low-value target, High-value target, and extreme trials, 24 catch trials, and 12 single-option trials).

Table 2: The number of points associated with each option in Experiment 2. Each of the outcomes associated with a risky option occurred with equal probability.

	Low variance		**High variance**	
	Safe	**Risky**	**Safe**	**Risky**
Low-value	25	20/30	25	10/40
High-value	75	70/80	75	60/90
Extreme	50	0/100	50	0/100

Procedure The procedure was identical to Experiment 1 except that a point tally was displayed on the screen and participants were informed that they would be given $1 for every 5000 points earned during the task. This change was implemented to remove a difference between our study and the previous experiments conducted by Ludvig et al. (2014).

Analysis The same priors were used as in Experiment 1 with a few additions. Firstly, an LKJ distribution with a η of 4 was used for the correlation between the intercept and slope parameters in the hierarchical model of choices (Lewandowski, Kurowicka, & Joe, 2009). Secondly, a half-Cauchy distribution with a location of 0 and scale of 1 and a half-Cauchy distribution with a location of 0 and scale of 25 were used for the group-level standard deviation parameters for the hierarchical models of first-to-mind and percentage estimation responses, respectively. All parameters had an effective sample size greater than 10000 and an $\hat{R} < 1.01$ indicating adequate chain convergence.

Unlike the first experiment, the second experiment involved some options that always resulted in better outcomes than others. Therefore, following from Ludvig et al. (2014), data from participants that chose the better option on less than 60% of choices between a Low-value and High-value option—one participant in the Low-variance condition and three in the High-variance condition—were excluded from the analysis.

Results and discussion

The choices that participants made in the second experiment were again not consistent with the distance from the centre conceptualisation of extremity but neither were they kind to the idea that the best and worst outcomes are unique (see Figure 5). These results were analysed using Bayesian hierarchical logistic regression predicting the choices that participants made when presented the target pairs of options, with condition (High-variance/Low-variance) as a fixed predictor variable and with intercepts and slopes for option type (High-value/Low-value) that varied for each participant. There was evidence that participants were more willing to choose the risky High-value option than the risky Low-value option despite the fact that none of the associated outcomes were the best or worst: the mean posterior log-odds for option type were 0.43 (95% CI = 0.09, 0.78). There was no evidence, however, suggesting that this willingness was stronger in the High-variance than the Low-variance condition, despite the distance from the centre of the more extreme outcome being much greater: the mean posterior log-odds for the interaction between option type and condition were 0.06 (95% CI = -0.42, 0.55).

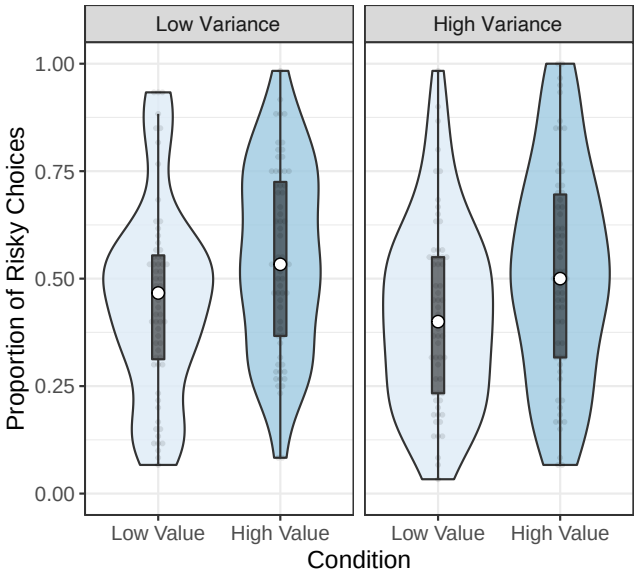

Figure 5: *The proportion of choices for the risky option on target trials (between risky and safe options that had equal expected value) in Experiment 2.*

Turning now to the memory tasks, the responses to the first-to-mind questions (see Figure 6) were analysed using Bayesian logistic regression predicting whether participants reported the extreme or non-extreme outcome as coming to mind first when presented with the risky target options, with

experiment and condition as fixed predictor variables and varying intercepts for each participant. Responses that were not a valid experienced outcome for the presented option were excluded from this analysis. Starting with the model that included only the intercept, the mean posterior log-odds for the intercept were -0.66 (95% CI -0.93, -0.39) providing evidence that outcomes further in rank from the centre were, on average, remembered better than intermediate outcomes. In the model including all parameters, there was some evidence that this effect was stronger in the High-variance condition than the Low-variance condition: the mean posterior log-odds for condition were -0.78 (95% CI = -1.52, -0.05). There was also some evidence, however, that this was largely driven by responses to the Low-value option in the High-variance condition: the mean posterior log-odds for the interaction were 0.75 (95% CI = -0.18, 1.68).

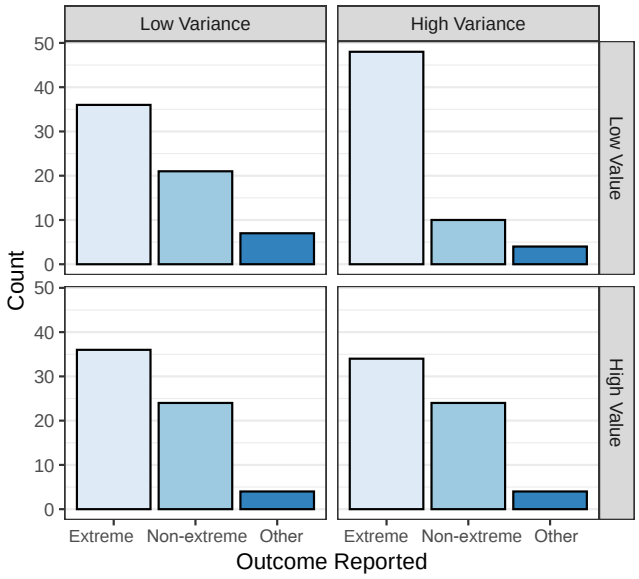

Figure 6: *The number of participants that responded with each outcome as coming to mind first when presented with the risky target options in Experiment 2.*

Responses to the percentage estimation task (see Figure 7) were analysed using Bayesian linear regression predicting the difference between the percentage estimates given for the extreme and non-extreme outcomes for the target options, with experiment and condition as fixed predictor variables and varying intercepts for each participant. Starting with the model that included only the intercept, there was evidence that, on average, the outcome ranked further from the centre was more likely to be reported as occurring with higher frequency: the mean posterior estimate for the intercept was 10.87 (95% CI = 6.02, 15.73). There was also some evidence, looking now at the model including all parameters, that the bias towards reporting the more extreme outcome was stronger for the High-variance condition: the mean posterior estimate for condition was 10.55 (95% CI = -1.57, 22.78). Similarly to the previous experiment, however, this might be largely driven by responses to the Low-value option in the High-variance condition: the mean posterior log-odds for the interaction were -12.62 (95% CI = -29.22, 3.95).

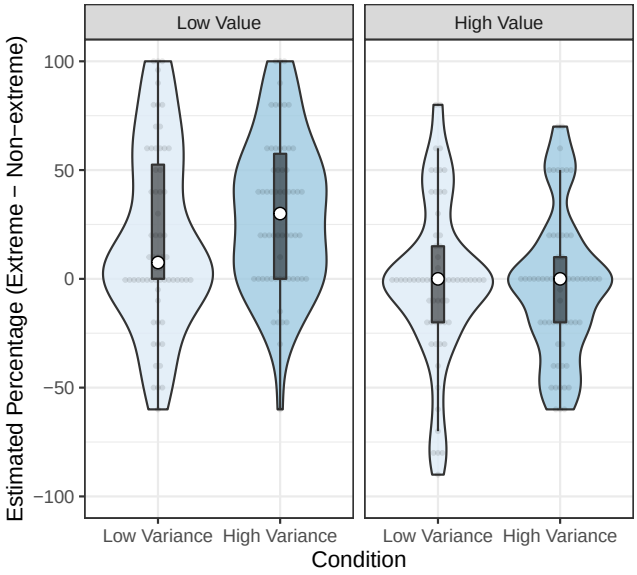

Figure 7: *The difference between percentage estimates for the extreme and non-extreme outcomes when presented with the risky target options in Experiment 2*

While there was some evidence for the disproportionate influence of extreme outcomes on responses to the memory questions, the choices that participants made in the decision task pose problems for both the distance from the centre and best and worst conceptualisations of extremity. Neither is able to fully explain the pattern of responses. Notably, higher preferences for the risky High-value option relative to the risky Low-value option were similar across the variance conditions even though the outcomes furthest from the centre that were associated with the target options in the Low-variance condition were 20 points away from the edge—for perspective, those same outcomes were 30 points from the centre of the distribution. It is difficult to reconcile this finding with the conceptualisation of extremity that the best and worst outcomes are unique, even allowing for a considerable degree of noise in encoding.

General Discussion

The experiments described here aimed to assess whether the influence of extreme outcomes should be characterised as the best and worst outcomes exerting a unique influence or as outcomes becoming more influential with increasing deviation from the centre. Contrary to our expectations, the choices that participants made were not consistent with either of these explanations: participants were more willing to choose the high-value risky option than the low-value risky option in Experiment 2 even though none of the associated outcomes were

the best or worst but there was no evidence that the distance from the centre influenced the options that they selected.

There was, however, some evidence that the distance from the centre affected outcome memory following the decision task. So what are we to make of this dissociation? One possibility is that the difference between conditions was large enough to impact participants' memory but not their preferences. There is some evidence that this was not the case. Firstly, the magnitude of the effect on outcome memory in Experiment 1b was at least as large as the one that produced a difference in preferences between the low-value and high-value options in Experiment 2. Secondly, the difference between conditions in Experiment 1b was equal to roughly half the range of experienced values and the difference between conditions in Experiment 2 was contrived to be as large as possible without leading to indifference or confusion between outcomes. While it is still possible that the manipulation was insufficient, this would place severe constraints on the situations we should expect to observe the effect—perhaps constraining it to situations in which there are non-overlapping low- and high-value pairs of options.

Another possibility is that there was a confounding variable in Experiment 1 that led to the dissociation between memory and choice. The inclusion of the context variables not only changed the distance between the centre and the best and worst outcomes but also changed the rank of the equivalent safe outcome. This could explain why the choices fall in the opposite direction than was expected, but the rank of the safe option was the same in both experiment 1a and 1b and one might reasonably expect that there would be a greater effect of extremity in the latter where the difference between the average outcome of the two conditions was doubled.

A third possibility is that the dissociation between choice and memory reflects a genuine difference in the way people respond on these tasks. The memory questions required participants to reflect on the distribution of outcomes that they experienced, which might have made the outliers more salient. Admittedly, this is less parsimonious than both tasks sharing a single mechanism but it nonetheless remains a possibility because the disproportionate influence of extreme outcomes has only been observed in very specific experimental designs involving low- and high-value options. If nothing else, these results highlight the importance of not relying on a homogeneous set of experimental manipulations if we hope to understand the influence of extreme outcomes in risky choice.

References

Brown, G. D., Neath, I., & Chater, N. (2007). A Temporal Ratio Model of Memory. *Psychological Review*, *114*(3), 539–576. doi: 10.1037/0033-295X.114.3.539

Bürkner, P. C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, *80*(1). doi: 10.18637/jss.v080.i01

Fredrickson, B. L. (2000). Extracting meaning from past affective experiences: The importance of peaks, ends, and specific emotions. *Cognition and Emotion*, *14*(4), 577–606. doi: 10.1080/026999300402808

Lewandowski, D., Kurowicka, D., & Joe, H. (2009). Generating random correlation matrices based on vines and extended onion method. *Journal of Multivariate Analysis*, *100*(9), 1989–2001. doi: 10.1016/j.jmva.2009.04.008

Lieder, F., Griffiths, T. L., & Hsu, M. (2018, 1). Overrepresentation of extreme events in decision making reflects rational use of cognitive resources. *Psychological Review*, *125*(1), 1–32. doi: 10.1037/rev0000074

Ludvig, E. A., Madan, C. R., McMillan, N., Xu, Y., & Spetch, M. L. (2018, 12). Living near the edge: How extreme outcomes and their neighbors drive risky choice. *Journal of Experimental Psychology: General*, *147*(12), 1905–1918. doi: 10.1037/xge0000414

Ludvig, E. A., Madan, C. R., & Spetch, M. L. (2014, 4). Extreme Outcomes Sway Risky Decisions from Experience. *Journal of Behavioral Decision Making*, *27*(2), 146–156. doi: 10.1002/bdm.1792

Madan, C. R., Ludvig, E. A., & Spetch, M. L. (2014, 6). Remembering the best and worst of times: Memories for extreme outcomes bias risky decisions. *Psychonomic Bulletin & Review*, *21*(3), 629–636. doi: 10.3758/s13423-013-0542-9

Rouhani, N., Norman, K. A., & Niv, Y. (2018, 9). Dissociable effects of surprising rewards on learning and memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *44*(9), 1430–1443. doi: 10.1037/xlm0000518

Schmidt, S. R. (1991, 11). Can we have a distinctive theory of memory? *Memory & Cognition*, *19*(6), 523–542. doi: 10.3758/BF03197149