Journal of Behavioral Decision Making, *J. Behav. Dec. Making*, **27**: 146–156 (2014)
Published online 11 July 2013 in Wiley Online Library (wileyonlinelibrary.com) **DOI:** 10.1002/bdm.1792

Extreme Outcomes Sway Risky Decisions from Experience

ELLIOT A. LUDVIG[1,2*,†], CHRISTOPHER R. MADAN[3,4,†] and MARCIA L. SPETCH[3]

[1]*Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA*

[2]*Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ, USA*

[3]*Department of Psychology, University of Alberta, Edmonton, AB, Canada*

[4]*Department of Systems Neuroscience, University Medical Center Hamburg-Eppendorf, Hamburg, Germany*

ABSTRACT

Whether buying stocks or playing the slots, people making real-world risky decisions often rely on their experiences with the risks and rewards. These decisions, however, do not occur in isolation but are embedded in a rich context of other decisions, outcomes, and experiences. In this paper, we systematically evaluate how the local context of other rewarding outcomes alters risk preferences. Through a series of four experiments on decisions from experience, we provide evidence for an extreme-outcome rule, whereby people overweight the most extreme outcomes (highest and lowest) in a given context. As a result, people should be more risk seeking for gains than losses, even with equally likely outcomes. Across the experiments, the decision context was varied so that the same outcomes served as the high extreme, low extreme, or neither. As predicted, people were more risk seeking for relative gains, but only when the risky option potentially led to the high-extreme outcome. Similarly, people were more risk averse for relative losses, but only when the risky option potentially led to the low-extreme outcome. We conclude that in risky decisions from experience, the biggest wins and the biggest losses seem to matter more than they should.

KEY WORDS risky choice; decisions from experience; extreme outcomes; peak-end rule; decision context; framing effects

Many behavioral economic studies on risky decisions present people with scenarios in which the outcomes and their probabilities are explicitly described (e.g., Kahneman & Tversky, 1979). For example, people might be explicitly asked whether they would prefer a guaranteed \$20 or a 50/50 chance at \$40. When faced with these risky decisions from description, people are usually risk averse for gains and risk seeking for losses—a pattern of risk preference known as the *reflection effect*. In life, however, people often make economic decisions based on their past experience with the consequences of those decisions. People frequent certain stores but not others, decide which products to buy, and risk whether or not to pay for the parking meter for a short dash into the store—all based in part on their own experiences. These decisions from experience can sometimes lead to markedly different behavior than those from description (e.g., Barron & Erev, 2003; Camilleri & Newell, 2011; Hertwig & Erev, 2009; Hertwig et al., 2004; Ludvig & Spetch, 2011; Weber et al., 2004).

The most prominent difference between description and experience occurs when one outcome is relatively rare. People overweight rare outcomes in description, whereas they underweight rare events in experience—often known as the description–experience (DE) gap. In the absence of rare events, there are usually no systematic differences between the described and experienced cases (e.g., Erev et al., 2010). In a recent study, however, we found exactly such a DE gap. Using 50/50 outcomes (no rare events), we found a clear reversal of the reflection effect in experience, but not in description (Ludvig & Spetch, 2011). People were more risk seeking for gains than losses in experience but, conversely, were more risk seeking for losses than gains in description. In those experiments, unlike many studies that examine the DE gap, we used a within-subject design (but see Camilleri & Newell, 2009). This design intermingled decisions from experience and descriptions as well as those between gains and losses, suggesting that perhaps the decision context is crucial for determining the pattern of risky choice (Ludvig & Spetch, 2011). In this paper, we focus solely on decisions from experience and present a series of experiments that systematically evaluates how the decision context influences this experience-based risky choice.

In perception, context effects abound: A surface can appear different colors because of the surrounding colors (e.g., Lotto & Purves, 2000), and the apparent length of a line can depend on the direction of the arrowheads (Müller-Lyer, 1889). Similarly, in choice, the local context in which a decision is made can greatly influence that decision (Simonson, 1989; Simonson & Tversky, 1992). For example, when choosing between two items, people's preference between the two options can be altered by introducing a third, seemingly irrelevant, option. If the third option is similar, but slightly worse than one of the original two options, this change will often lead to an increase in preference for the similar option (Heath & Chatterjee, 1995; Huber et al., 1982). This preference bump occurs even though nothing has changed about the original two options except for the surrounding context of other options. Risky decisions can also be influenced by how the decision is framed (Tversky & Kahneman, 1981) or the alternative options available (Erev, Glozman & Hertwig, 2008; Stewart et al., 2003). Thus, as with psychophysical judgments (e.g., Helson, 1947; Thomas & Jones, 1962), risky decisions can be context dependent and altered by the comparison set experienced.

*Correspondence to: Elliot Ludvig, Princeton Neuroscience Institute, Princeton University, NJ 08540, USA. E-mail: eludvig@princeton.edu
†These authors contributed equally to this paper.

In the present studies, we examined how the decision context in which repeated choices are made affects risky decisions from experience. Specifically, we were interested in whether the edges of the experienced distribution (the biggest wins and losses) might be overweighted in decisions from experience. A similar effect, termed the peak-end rule, is observed in retrospective judgments of affective experiences. The post-hoc valuation depends primarily on the point of maximum intensity (the peak) and on the end (Fredrickson, 2000). In a classic illustration of this rule, post-operative pain judgments of patients undergoing a colonoscopy were strongly correlated with the peak intensity of pain (as judged in real time) and with the pain intensity at the end, but not with the duration of the procedure (Kahneman et al., 1993; Redelmeier & Kahneman, 1996; Stone et al., 2005). Delayed judgments relating to positive experiences, such as vacations, are similarly influenced by peak and end intensities (Mitchell et al., 1997).

Inspired by the peak portion of the peak-end rule, we hypothesized that the extreme outcomes in a decision context may be overweighted in decisions from experience, as they are in delayed judgments. Following this *extreme-outcome rule*, when a risky option occasionally leads to the best possible gain in a given context, that extreme gain should be overweighted in the valuation of that risky option. As a result, when pitted against another option of similar expected value, but without the possibility of an extreme gain, that option would be chosen more frequently. Similarly, when a risky option occasionally leads to the worst possible outcome in a context, people should overweight that extreme loss. When pitted against another option of similar expected value, but without the possibility of an extreme loss, that option should be chosen less frequently. Thus, the extreme-outcome rule predicts that people will become more risk seeking for gains than for losses in decisions from experience, but only when the risky choices include the most extreme outcomes in the decision context (i.e., the biggest gain or loss). This prediction about the effects of rewarded experience on subsequent decisions from experiences runs opposite to the usual reflection effect in decisions from description (e.g., Kahneman & Tversky, 1979).

Some evidence for this potential extreme-outcome rule comes from studies with non-human animals, which can *only* rely on experience for learning about outcomes. Many of these studies have also reported risk seeking for gains (e.g., Hayden et al., 2008; Heilbronner & Hayden, 2013; Kacelnik & Bateson, 1996; McCoy & Platt, 2005; O'Neill & Schultz, 2010), and some evidence suggests that this risk seeking for gains may be driven by extreme outcomes in a context. For example, the risk-seeking behavior of rhesus macaques was shown to be sensitive to the magnitude of the *jackpot* or largest reward on a two-outcome choice. When the jackpot was reduced, risk seeking declined; when the jackpot was increased, risk seeking increased. An identical manipulation of the smaller, non-extreme reward had no influence on risk preference (Hayden et al., 2008).

In this paper, we present four experiments that test the extreme-outcome rule by systematically manipulating the decision context. Figure 1 illustrates the basic task (cf. Ludvig & Spetch, 2011). On most trials, people decided between two doors (Figure 1A). Picking one door always led to a fixed outcome, and picking the other (risky) door led with a 50/50 chance to more or less than the fixed outcome. For example, in Experiment 1, the fixed-gain door always led to +20, whereas the risky-gain door led to a 50/50 chance of 0 or +40. Conversely, the fixed-loss door always led to −20, whereas the risky-loss door led to a 50/50 chance of −40 or 0. In this case, the extreme outcomes were +40 and −40. By the proposed extreme-outcome rule, these highest-valued and lowest-valued outcomes in the decision context would be overweighted, leading to more risk seeking for gains, yet more risk aversion for losses.

A partial-feedback procedure was used in which participants only saw the outcome for the chosen option, but not the foregone option (see Camilleri & Newell, 2011; Hertwig & Erev, 2009). Outcomes for the forgone option were not included out of concern that people might confuse which option led to which outcome. There were three types of trials. On decision trials (Figure 1A), which always involved a choice between two loss doors or between two gain doors, the objective expected value of the fixed and risky door was equal. Interspersed catch (Figure 1B) and single-choice (Figure 1C) trials ensured that people indeed learned and experienced the correct contingencies. This partial-feedback procedure was designed to highlight the relationship between the option chosen and the outcome received.

The series of four experiments each examined a different facet of the extreme-outcome rule. Table 1 details the exact contingencies in each experiment. Experiment 1 examined decisions from experience with gains and losses intermingled in a single decision context, without concurrent decisions from description (cf. Ludvig & Spetch, 2011). Experiment 2 directly tested the alternative hypothesis that zero values were underweighted rather than extreme values being overweighted. Experiments 3 tested whether all larger magnitude options were overweighted or only extreme ones, by adding into the decision context additional doors that potentially led to more extreme outcomes. Experiment 4 evaluated a novel prediction of the extreme-outcome rule that only relative extremes matter, independent of whether they are absolute gains or losses. In all cases, based on the extreme-outcome rule, we predict more risk seeking for risky options that potentially led to the high extreme in the experiment and, conversely, less risk seeking for risky options that potentially led to the low extreme.

EXPERIMENT 1: MIXED GAINS AND LOSSES

Experiment 1 evaluated risk preferences in a decision context with intermixed gain and loss problems. There were four possible doors that each led to a different outcome (see Table 1): a fixed-gain door (+20), a risky-gain door (50/50 chance of either 0 or +40), a fixed-loss door (−20), and a risky-loss door (50/50 chance of either −40 or 0). People repeatedly made choices between pairs of these doors. In this decision context, +40 and −40 were the two extremes. According to the extreme-outcome rule, people should overweight these two extremes in the decision process, leading to

DOI: 10.1002/bdm

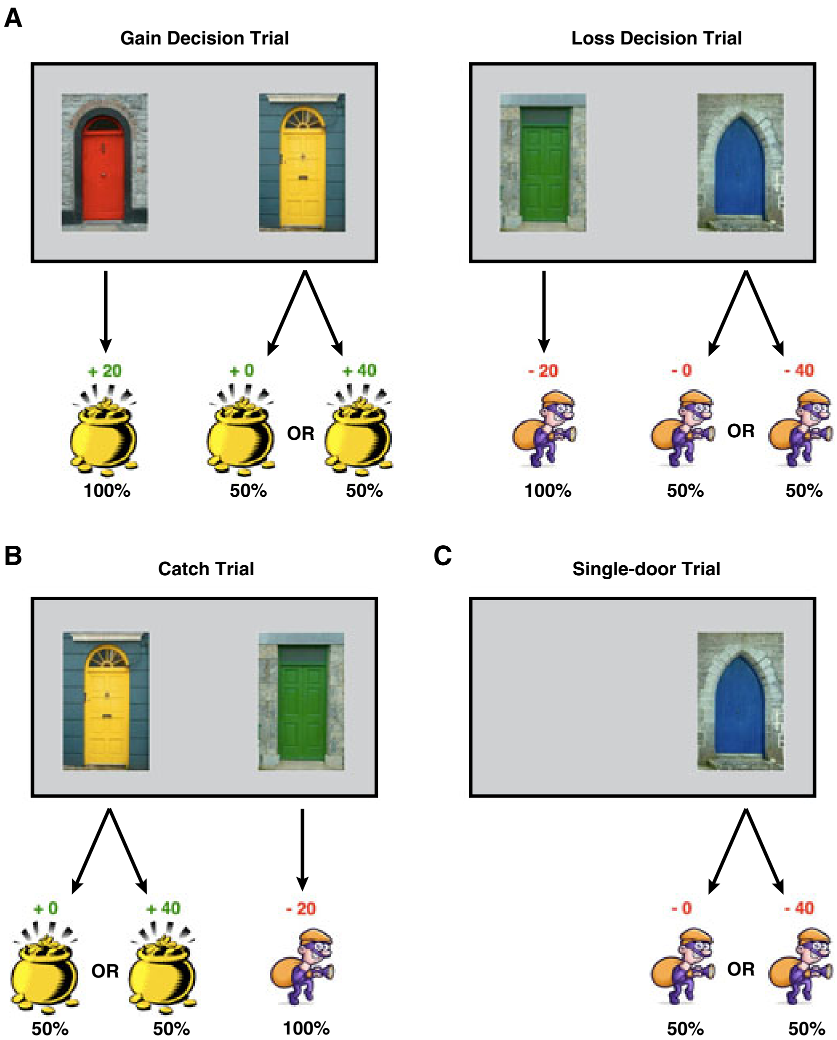

Figure 1. Schematic of the general method used. Specific values correspond to Experiment 1. (A) *Decision trials* involved choices between two gain doors or two loss doors. One door always led to a gain (or loss) of a fixed number of points, and the other door led equiprobably to one of two possible outcomes. Choices were followed by feedback about the amount gained or lost. (B) *Catch trials* involved choices between one gain door and one loss door and ensured that participants paid attention to their decisions. (C) *Single-door trials* only presented one door and ensured that participants occasionally experienced the planned reward contingencies

Table 1. Details of all experimental manipulations and summary of main results

Experiment	Manipulation	Decision type	Fixed outcome	Risky outcomes (50/50)	Degree of risk seeking
1	Mixed problems	Gain	+20	0/**+40**	Gains > losses
		Loss	−20	**−40**/0	
2	No zeroes	Gain	+25	5/**+45**	Gains > losses
		Loss	−25	**−45**/−5	
3	Magnitude	X gain	+40	0/**+80**	X: gains > losses
		X loss	−40	**−80**/0	
		NX gain	+20	0/+40	NX: no difference
		NX loss	−20	−40/0	
4G	All gains	HX	+60	+40/**+80**	HX > BX > LX
		BX	+40	**0**/**+80**	
		LX	+20	**0**/+40	
4L	All losses	HX	−20	−40/**0**	HX > BX > LX
		BX	−40	**−80**/**0**	
		LX	−80	**−80**/−40	

Bold font indicates extreme outcome in an experiment. X, extreme outcome; NX, no extreme; HX, high extreme; LX, low extreme; BX, both extremes.

more risk seeking in the gain case and more risk aversion in the loss case. This experiment also evaluated more generally whether a reverse-reflection effect (greater risk seeking for gains than losses) would occur in decisions from experience without concurrent decisions from description (cf. Ludvig & Spetch, 2011).

DOI: 10.1002/bdm

Methods

Participants

Twenty-nine introductory psychology students at the University of Alberta participated for course credit (15 females; $M_{\text{age}} = 18.9$ years, $SD = 1.7$). Each participant gave written informed consent. The study was approved by a university ethics board.

Procedure

Participants played a computer-based task in which they were told to try and earn as many points as possible. As illustrated in Figure 1, on most trials, participants were presented with pictures of two doors, and they indicated their choice by clicking on one of those doors. Choices were immediately followed by feedback for 1.2 s, which showed the number of points won or lost along with a cartoon graphic. Feedback was only given for the chosen door, as in a partial-feedback procedure (e.g., Camilleri & Newell, 2011; Hertwig & Erev, 2009). The total accumulated points were continuously displayed at the bottom of the screen. An inter-trial interval of 1 to 2 s separated each trial.

Sessions were each organized into five blocks of 48 trials, separated by a brief break. Each block included a mixture of three trial types. *Decision trials* involved choices between two gain doors or two loss doors (Figure 1A). For both gains and losses, the *fixed* door always led to the same outcome, and the *risky* door led equiprobably (i.e., with a 50/50 chance) to one of two outcomes: one smaller and one larger than the fixed outcome. The objective expected value on these decision trials was always equal for the fixed and risky doors. Across the risky option in all four experiments, the experienced likelihood of receiving either outcome never deviated significantly from .5. On s*ingle-door trials*, there was only one door presented, which had to be clicked on to continue (Figure 1C). These trials ensured that all doors were sometimes selected and that participants occasionally experienced all reward contingencies. *Catch trials* presented two doors with substantially different objective expected values—typically a choice between a gain door and a loss door (Figure 1B). These trials provided the opportunity to gain points over the session and ensured that participants paid attention to their decisions. Participants that chose the gain door on fewer than 60% of these catch trials, corresponding to behavior not distinguishable from chance responding at $p < .05$ with 80 catch trials, were excluded from further analyses. In Experiment 1, data from one participant were excluded for poor performance on catch trials.

As detailed in Table 1, there were four different-colored doors in Experiment 1: a fixed gain (100%: +20 points), risky gain (50%: 0 and 50%: +40), fixed loss (100%: −20), or risky loss (50%: 0 and 50%: −40). The 48 trials in each block were divided among the three trial types as follows: 24 decision trials between the two gains and the two losses (12 of each), 16 catch trials, and 8 single-door trials. Each trial was equally incentivized with people earning (or losing) points on all 240 trials in the experiment. There were no separate sampling and choice phases, mirroring a repeated-choice design, rather than a sampling design (e.g., Erev et al., 2010; Hertwig & Erev, 2009). Trial order was randomized, but the total number of trials (240) and their distribution was constant across all participants. The door color associated with the fixed or risky gain or loss was counterbalanced across participants.

The numbers of different trial types were chosen to ensure that each door appeared equally often on both sides of the screen to prevent any side biases. Both the decision trials and the single-door trials in each block were equally divided between gain and loss trials and thus had a total expected value of 0. Only the catch trials provided an opportunity to earn a net gain of points. Participants were encouraged to maximize their number of points, and good performance on the catch trials provided independent evidence that participants were adequately incentivized by the points.

Data analysis

Risk preference was calculated as the probability of choosing the risky door. For each experiment, risk preference was averaged over the final three blocks and compared using *t*-tests or ANOVAs as appropriate. The final three blocks were selected as the primary dependent measure because that afforded participants sufficient opportunity to learn the outcomes associated with each option, while providing a long enough sample to get a reliable measure of their risk preference over time. Linear-trend analyses across all blocks were conducted to look for learning effects. All tests were repeated measures. Given the a priori hypotheses, all *t*-tests were one-tailed. The Greenhouse–Geisser correction for non-sphericity was applied where appropriate. Inferential statistics were calculated using SPSS (IBM Inc., Armonk, NY) and MATLAB (The Mathworks Inc., Natick, MA).

Results and discussion

The left bars in Figure 2A depict the average risk preference (proportion of risky choices) over the final three blocks, split by gains and losses. Participants were more risk seeking for gains than losses, displaying a reversal of the usual reflection effect [$t(27) = 1.82$, $p < .05$, $d = .51$]. In this experiment, +40 and −40 were the extreme outcomes in the decision context, and thus, we predicted that these outcomes would be overweighted in the decision process. The observed difference in risk seeking for gains and losses was consistent with this prediction. Figure 2B plots these risk preferences by experimental block and shows an interaction between valence (gain or loss) and block [$F(4, 71) = 3.63$, $p < .05$, $\eta_p^2 = .12$]. There was an increase in risk aversion for losses across the experiment [linear effect of Block: $F(1, 27) = 8.76$, $p < .01$, $\eta_p^2 = .25$], but no change in risk preference for gains [linear effect of Block: $p > .1$, $\eta_p^2 = .009$].

Across the population, risk preference hovered around .5, raising the possibility that people have a tendency toward equal allocation of their responses across the two options. To evaluate this possibility, Figure 2C plots the risk preferences for each individual for both gains and losses. There is clearly a wide spread of risk preferences in the population (from almost never to almost always). Independent of this

DOI: 10.1002/bdm

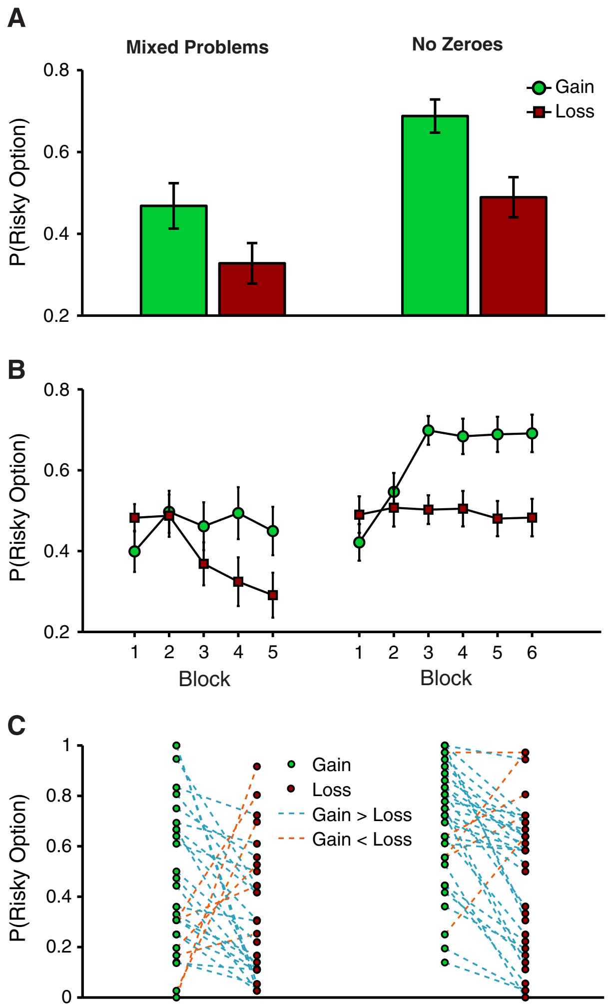

Figure 2. Results of Experiments 1 and 2. (A) Mean risk preference (± SEM) for gain and loss doors averaged over the last three blocks. (B) Mean risk preference (± SEM) as a function of block for all blocks. (C) Mean risk preference for individuals over the last three blocks. Dashed lines connect data from the same individual. From left to right, the plots show results for mixed problems (Experiment 1; $N = 28$) and no zeroes (Experiment 2; $N = 34$). As predicted, in both cases, risk preferences are higher for gains than losses, consistent with the extreme-outcome rule

individual variability in overall risk preference, most people show the same pattern: more risk seeking for gains than losses [22 out of 28 participants = 79%; $p < .001$ by a binomial test].

These results extend our previous results (Ludvig & Spetch, 2011), which also showed greater risk seeking for gains than losses in decisions from experience, but that design also included interspersed decisions from description. Here, we found more risk seeking for gains than losses in experience-based choice, even in the absence of described choices.

EXPERIMENT 2: NO ZEROES

Experiment 1 provided evidence for an extreme-outcome rule, whereby the largest and smallest outcomes in a decision context are overweighted in choice. In that experiment, however, the potential non-extreme outcomes for the risky options were always 0 (see also Ludvig & Spetch, 2011). Zero, however, is neither an absolute gain nor an absolute loss and may instead be treated as a special number (e.g., Shampanier et al., 2007). Thus, a possible alternative hypothesis to the overweighting of extreme outcomes is that, instead, the 0 outcomes may be underweighted. In addition, because the 0 outcome potentially followed the risky option in both the gain and the loss case in Experiment 1, it occurred twice as frequently as the extreme outcomes. As a result, perhaps this increased frequency of the 0 outcome led to a reduced weighting of those 0 outcomes relative to the extreme (+40 or −40) outcomes. This second alternative hypothesis is unlikely, however, as it goes against the wealth of existing evidence that rare events are underweighted in experience-based choice (e.g., Hertwig & Erev, 2009; Hertwig et al., 2004; Weber et al., 2004).

This experiment aimed to rule out both these alternative hypotheses by shifting the absolute values of the different outcomes by 5 points. Thus, as indicated in Table 1, people were presented with four possible doors: a fixed-gain door (+25), a risky-gain door (50/50 chance of either +5 or +45), a fixed-loss door (−25), and a risky-loss door (50/50 chance of either −45 or −5). People repeatedly made choices between pairs of these doors. The extreme outcomes are now +45 and −45. By the extreme-outcome rule, we expect more risk seeking for gains and more risk aversion for losses. The two alternative hypotheses, in contrast, would predict that this reversal of the reflection effect should disappear when the more frequent 0 outcome is removed from the mix.

Methods

Participants

Thirty-four students drawn from the same subject pool as Experiment 1 participated (15 females; $M_{age} = 19.1$ years, $SD = 2.0$).

Procedure

The procedure was identical to Experiment 1, except that the absolute values of the potential outcomes that followed the four doors (Table 1) were all shifted by 5 points and participants were run for an additional (6th) block of 48 trials. The outcomes for the four doors were a fixed gain (100%: +25), risky gain (50%: +5 and 50%: +45), fixed loss (100%: −25), or risky loss (50%: −5 and 50%: −45). No participants in this experiment were excluded for poor performance on the catch trials.

Results and discussion

The right bars in Figure 2A plot the average risk preference on gain and loss trials over the final three blocks in Experiment 2. Participants were once again more risk seeking for gains than losses—a clear reversal of the reflection effect [$t(33) = 4.46$, $p < .001$, $d = .77$]. Figure 2B clearly shows how risk preference changed across the six blocks [main effect of Block: $F(3, 99) = 10.15$, $p < .001$, $\eta_p^2 = .24$; Block × Valence interaction: $F(3, 101) = 5.47$, $p < .001$, $\eta_p^2 = .14$]. Across the blocks, risk seeking for gains increased [linear

DOI: 10.1002/bdm

effect of Block: $F(1, 37) = 16.93$, $p > .001$, $\eta_p^2 = .34$], but there was little change in risk preference for losses [linear effect of Block: $p > .1$, $\eta_p^2 = .002$]. Figure 2C plots the individual differences, where a significant majority of people are more risk seeking for gains than losses [29 out of 34 participants = 85%; $p < .001$ by a binomial test]. The higher level of risk seeking for gains than losses in the absence of zero values rules out the two alternative hypotheses. The pattern of risk preferences observed in Experiments 1 and 2 were not due to underweighting the zero outcomes. The results, however, offer further support for the extreme-outcome rule, as they are congruent with an overweighting of the two extremes in the decision context, which were +45 and −45 in this experiment.

EXPERIMENT 3: OUTCOME MAGNITUDE

One clear prediction of the extreme-outcome rule is that a given outcome will only be overweighted when it is either the largest and smallest outcomes in a decision context. In Experiment 3, we directly tested this prediction by adding new doors with higher-magnitude outcomes into a decision context with the same gain and loss problems as in Experiment 1. In this case, as detailed in Table 1, participants encountered eight possible doors. The four non-extreme (NX) doors led to the same outcomes as in Experiment 1: a guaranteed gain/loss of 20 or a 50/50 chance of a gain/loss of 40. The other four extreme (X) doors led to exactly double those outcomes: a guaranteed gain/loss of 40 or a 50/50 chance of a gain/loss of 80. The highest and lowest outcomes were +80 and −80 respectively, and the extreme-outcome rule predicts more risk seeking for gains than losses for decisions involving those extreme outcomes. For the NX doors, even though the outcomes were identical to those in Experiment 1, the decision context has changed so that +40 and −40 are no longer extreme outcomes. As a result, the extreme-outcome rule predicts that the reverse-reflection effect observed in Experiment 1 (see Figure 2A; Ludvig & Spetch, 2011) should not be present for these NX doors.

Methods

Participants

Thirty-nine students drawn from the same subject pool as Experiments 1–2 participated (16 females; $M_{age} = 18.9$ years, $SD = 2.6$).

Procedure

The procedure followed the same protocol as Experiments 1 and 2 with the following minor changes because there were now eight total doors. The outcomes for the four non-extreme (NX) doors were the same as in Experiment 1: a fixed gain (100%: +20 points), risky gain (50%: 0 and 50%: +40), fixed loss (100%: −20), or risky loss (50%: 0 and 50%: −40). The outcomes for the extreme (X) doors were twice the magnitude: a fixed high gain (100%: +40), risky high gain (50%: 0 and 50%: +80), fixed high loss (100%: −40), or risky high loss (50%: 0 and 50%: −80). Trials were presented in five blocks of 72 trials that each contained 32 decision trials between fixed and risky gains or losses of the same magnitude level (eight of each), 24 catch trials, and 16 single-door trials. Three types of catch trials were used: choices between the extreme gain and loss doors, choices between the non-extreme gain and loss doors, and choices between two gain or two loss doors of different expected values (e.g., +20 vs. 0/+80). Only the choices between a gain door and a loss door were used to exclude participants. Two participants were excluded.

Results and discussion

Figure 3 shows how, as predicted, there was an interaction between decision type (X/NX) and reward valence in the final three blocks [$F(1, 36) = 4.27$, $p < .05$, $\eta_p^2 = .11$; no main effects]. Risk seeking was greater for gains than for losses in X decisions [$t(36) = 1.84$, $p < .05$, $d = .40$], but not in NX decisions [$t(36) = .80$, $p > .1$, $d = .17$]. Note how the outcomes that followed the NX doors were identical to the ones that followed the doors in Experiment 1, yet the results were reversed in the new decision context (cf. Figure 2). Across the session, Figure 3B shows how, for the X doors, participants became increasingly risk averse for losses [linear effect of Block: $F(1, 36) = 5.05$, $p < .05$, $\eta_p^2 = .12$], but the visual trend toward an increase in risk seeking for gains was not significant [$p > .1$]. For the NX doors, risk preference did not change across blocks for either gains or losses [both $ps > .1$]. The greater risk seeking for gains than losses for X decisions, but not NX decisions, provides strong

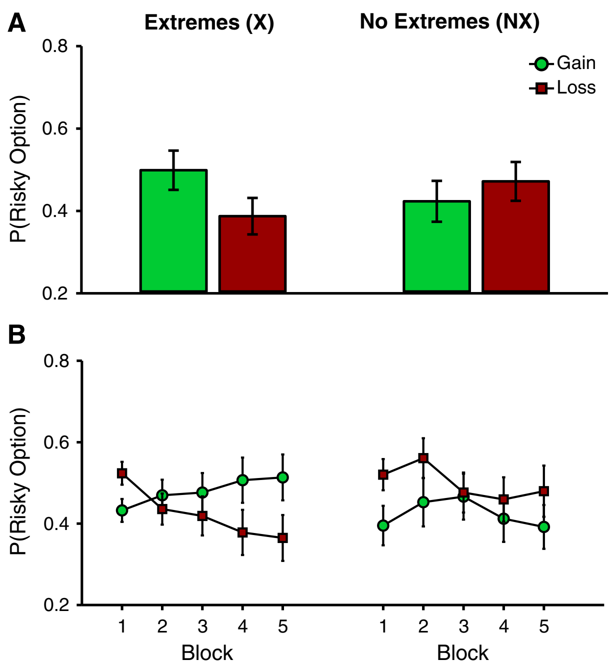

Figure 3. Results of Experiment 3. (A) Mean risk preference (± SEM) for gain and loss doors averaged over the last three blocks. (B) Mean risk preference (± SEM) as a function of block for all blocks. Decisions with extreme (X) outcomes are on the left, and decisions with no extreme (NX) outcomes are on the right ($N = 37$). As predicted by the extreme-outcome rule, risk preferences were higher for gains than losses in the X condition, but not in the NX condition

DOI: 10.1002/bdm

support for the prediction that only the most extreme outcomes in the decision context would be overweighted.

EXPERIMENT 4: FRAMING EFFECTS

In the first three experiments, the extreme values were always gains and losses of the same magnitude (40 in Experiment 1, 45 in Experiment 2, and 80 in Experiment 3). A further prediction of the extreme-outcome rule is that the largest and smallest values in a decision context should be overweighted in the decision process, independent of whether they are absolute gains or losses. For example, in a decision context with all gains, the lowest gain would be an extreme and should be overweighted.

To test this prediction, Experiment 4 split gains and losses across participants to test whether relative extremes are sufficient to elicit changes in risk sensitivity (Table 1). In the *All-Gain* group (Experiment 4G), the largest possible gain (+80) was the high-extreme (HX) outcome, and the smallest possible gain (0) was the low-extreme (LX) outcome. Conversely, in the *All-Loss* group (Experiment 4L), the smallest possible loss (0) was the HX outcome, and the largest possible loss (−80) was the LX outcome. Following the extreme-outcome rule, in both groups, there should be more risk seeking for the doors with potential HX outcomes and more risk aversion for the doors with potential LX outcomes.

In addition, we attempted to evaluate the relative weightings of the high and low extremes in the decision process. To do so, we also included risky options that potentially led to both extremes (BX). Given the loss aversion that characterizes many decisions (e.g., Kahneman & Tversky, 1979; Tom et al., 2007; Tversky & Kahneman, 1992), we might expect that the LX outcome (the relative loss) would be more heavily weighted, leading to risk aversion in both the gain and loss groups. In experience-based choice, however, significant loss aversion is often not observed (e.g., Erev, Ert, & Yechiam, 2008; Yechiam & Hochman, 2013), suggesting that the two extremes might be more evenly weighted, leading to an intermediate level of risk preference.

Methods

Participants

A total of 79 students from the same subject pool participated (58 females; $M_{age} = 19.5$ years, $SD = 2.3$; $N = 39$ and 40 for the All-Gain and All-Loss groups, respectively).

Procedure

The basic procedure was the same as in Experiments 1–3 (Figure 1) except that participants in the All-Gain group (Experiment 4G) experienced only gain doors and participants in the All-Loss group (Experiment 4L) experienced only loss doors. For both groups, there were six different-colored doors. As detailed in Table 1, for the All-Gain group, LX decisions were between a fixed gain (100%: +20 points) or a risky gain (50%: 0 and 50%: +40), the HX decisions were between a fixed high gain (100%: +60) or a risky high gain (50%: +40 or 50%: +80), and the BX decisions were between a fixed intermediate gain (100%: +40) and a risky gain that included both extremes (50%: 0 and 50%: +80). For the All-Loss group, the doors led to an LX fixed loss (100%: −60), LX risky loss (50%: −40 and 50%: −80), BX fixed loss (100%: −40), BX risky loss (50%: 0 and 50%: −80), HX fixed loss (100%: −20), and HX risky loss (50%: 0 and 50%: −40). Thus, for gains, the high and low extremes were +80 and 0, and for losses, the high and low extremes were 0 and −80. In both groups, we expected more risk seeking for HX decisions than for LX decisions, and we expected that risk seeking for the BX decisions would be closer to the LX decisions or fall between the other two decision types.

For both groups, trials were presented in five blocks of 60 trials that were divided among the three trial types: 36 decision trials between fixed and risky gains or losses of the same value level (12 of each), 12 catch trials, and 12 single-door trials. Two types of catch trials were used. Easy catch trials consisted of a choice between an HX door and an LX door (e.g., −20 vs. −40/−80). Subtle catch trials consisted of a choice between a BX door and either an HX or LX door (e.g., −40 vs. −20/−40). Subtle catch trials were included in the design to match the number of presentations of each door, but performance on these trials was not used to exclude participants. A total of 18 participants (10 in the All-Loss group and 8 in All-Gain group) were excluded for poor performance on the easy catch trials.

To ensure that both groups ended the experiment with a similar number of points, the All-Loss group started with approximately twice the number with which they would end the experiment (24 000 points). The All-Gain group started with zero points as in Experiments 1–3. The two groups did not differ in the number of points remaining/earned at the end of the experiment [$t(77) = .28$, $p > .1$].

Results and discussion

Consistent with the extreme-outcome rule, Figure 4A shows that people were more risk seeking for HX doors than for LX doors in the final three blocks for both groups [$F(2,\ 108) = 16.68$, $p < .001$, $\eta_p^2 = .22$]. Across the session, there was an increase in risk aversion for the LX decisions [$F(1,\ 60) = 11.00$, $p < .01$, $\eta_p^2 = .16$], but there was no change in risk preference for HX and BX decisions [both ps $> .1$, η_p^2s $< .05$]. Note how, in the All-Gain group, the LX decision was between a guaranteed +20 and a 50/50 chance at +40 (Table 1). This decision was identical to the gain decision in Experiment 1, yet, in this decision context, people were now a lot more risk averse [cf. Figure 2; $t(59) = 2.40$, $p < .05$, $d = .63$]. Similarly, in the All-Loss group, the HX decision was between a guaranteed −20 and a 50/50 chance at −40. This decision was identical to the loss decision in Experiment 1, yet, in this decision context, people were now a lot more risk seeking [cf. Figure 2; $t(55) = 3.88$, $p < .001$, $d = 1.05$]. This comparison across groups and experiments provides strong evidence for the extreme-outcome rule.

For the BX decisions, risk preferences were intermediate to the HX and LX decisions in both All-Gain and All-Loss

DOI: 10.1002/bdm

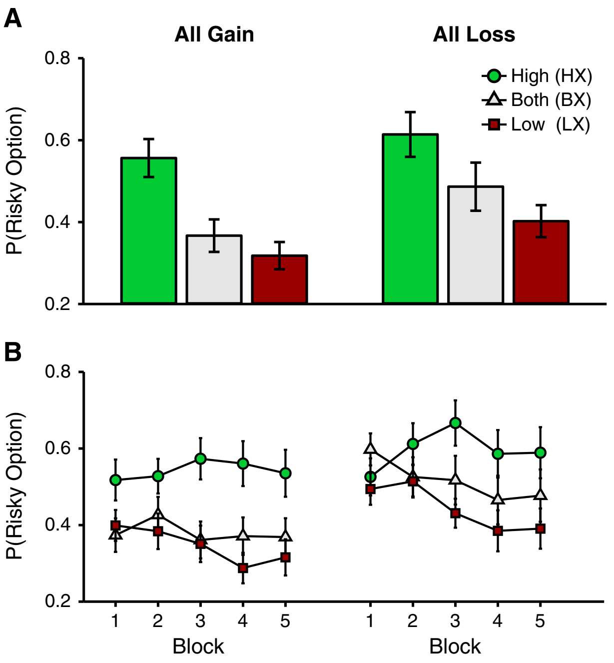

Figure 4. Results of Experiment 4. (A) Mean risk preference (± SEM) for each door averaged over the last three blocks. (B) Mean risk preference (± SEM) as a function of block for all blocks. From left to right, the plots show results for gain doors that led to high-extreme (HX), low-extreme (LX), or both extreme (BX) outcomes (Experiment 4G; $N = 32$), and loss doors that had HX, LX, or BX outcomes (Experiment 4L; $N = 29$). In both cases, risk preferences were highest for HX doors and lowest for LX doors, as predicted by the extreme-outcome rule

groups [linear effect of decision type: $F(1, 60) = 32.46$, $p < .001$, $\eta_p^2 = .35$] but were slightly risk averse and closer to the risk preference for the LX decision (particularly in the All-Gain group). This result suggests that the low extreme is weighted slightly more heavily than the high extreme.

GENERAL DISCUSSION

Across all four experiments, when decisions from experience contained the most extreme outcomes in the decision context, people were more risk seeking for relative gains than for relative losses. This result is opposite to the usual reflection effect observed with decisions from description (Kahneman & Tversky, 1979) but accords with some recent results with decisions from experience (e.g., Ludvig & Spetch, 2011; Tsetsos et al., 2012) as well as the risk seeking often observed in non-human animals (Hayden et al., 2008; O'Neill & Schultz, 2010). In accord with the extreme-outcome rule, this reversed reflection effect only occurred when the risky option potentially led to an extreme outcome. Thus, risky choice behavior was dependent not only on the potential outcomes of the current decision but also on the decision context in which it occurred (cf. Table 1).

This overall pattern can be clearly seen in Figure 5, which plots the difference between risk preferences for relative gains and losses in each experiment. In all four experiments, the risky option that led to the best possible outcome in that experiment produced more risk-seeking behavior, whereas

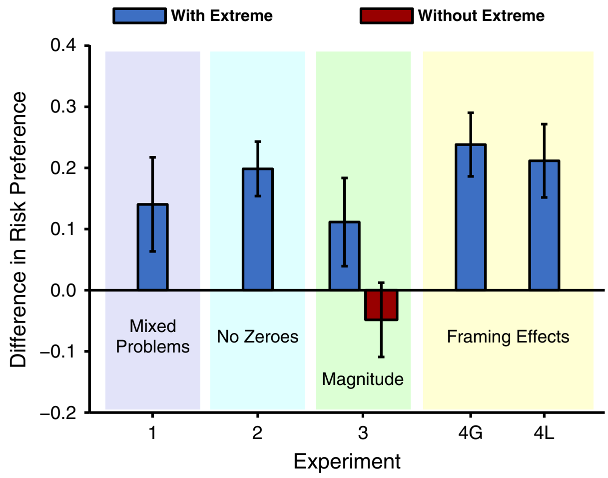

Figure 5. Mean difference (± SEM) in risk preference across the last three blocks in each experiment. For Experiments 1 to 3, the difference is calculated between the risk preferences for gains and losses. For Experiment 4, the difference is calculated between the risk preferences for high-extreme and low-extreme decisions. In all cases with extreme outcomes, the difference in risk preference was significantly above zero, indicating greater risk seeking for relative gains than relative losses

the risky option that led to the worst possible outcome in that experiment led to more risk-averse behavior. This reversal held even when all the experienced outcomes were absolute gains or losses (Experiment 4). The pattern was not apparent when the risky option never led to an extreme outcome (Experiment 3). In short, people chased the potential big win but avoided the potential big loss.

One key factor in generating this pattern of behavior is the intermingling of multiple problems in the same decision context. As a result, the outcomes that follow a risky option can be a high extreme, low extreme, or neither, pending the decision context, influencing risk preference (Figure 5). In other experiments on decisions from experience, participants occasionally encounter multiple problems sequentially (e.g., Erev et al., 2010), but they are not intermixed with one another. In those cases, even with 50/50 outcomes, there is no oversensitivity to the extreme outcomes, and the usual reflection effect is observed (e.g., Ert & Yechiam, 2010). In these other tasks with sequential presentation and 50/50 outcomes on the risky option, both outcomes become extremes—most similar to the BX option in Experiment 4 here. Indeed, isolating out that BX option, we also find the usual reflection effect, with more risk seeking for losses than gains (compare the gray bars in Figure 4).

An important question is thus what exactly constitutes the decision context. In all our experiments, there was clearly only one decision context. All problems were intermingled within each experimental block and throughout the whole experimental session. In contrast, the other studies that looked at 50/50 outcomes in decisions from experience, but found the usual reflection effect, presented problems sequentially within a session (e.g., Erev et al., 2010; Ert & Yechiam, 2010). This difference suggests that the relevant unit for the decision context is smaller than a full session. At the opposite end, we only presented one or two options at a time, yet the outcomes of other, non-presented, options influenced risk preference.

DOI: 10.1002/bdm

Thus, the decision context is clearly larger than the options immediately available (cf. Heath & Chatterjee, 1995; Huber et al., 1982; Simonson, 1989; Simonson & Tversky, 1992). More generally, we think that the decision context is the comparison set of all the other options and outcomes that are considered when making a decision—similar to the decision environment in the decision-by-sampling framework (Stewart et al., 2006). One possibility is that these other options and outcomes are linked to the context of the current decision through stimulus–stimulus associations as has been supposed in some models of animal learning (Miller & Matzel, 1988; Stout & Miller, 2007). In that case, options and outcomes that have previously co-occurred with or immediately preceded or followed either of the options under consideration would fall into the decision context.

Although opposite in direction, the context effects observed here for experience-based choices complement the previous evidence for framing effects in decisions from description (Tversky & Kahneman, 1981). The two groups in Experiment 4 highlight this relativity: zero outcomes served as both the high-extreme value for the All-Loss group, leading to more risk seeking, and the low-extreme value for the All-Gain group, leading to more risk aversion. The low and high extremes may not be weighted equally, however. In Experiment 4, the risky door for BX decisions, which led to both the best and the worst possible outcomes, produced moderate risk aversion. This finding suggests that the worst outcomes (relative losses) were weighted more heavily than the best outcomes (relative gains), reminiscent of loss aversion.

In experience-based choice, decisions must be made based on the memories of past outcomes. This dependence on the past suggests that memory biases may play a role in the overweighting of extreme outcomes. In other contexts, choice is indeed influenced by the biases inherent in human memory (see Weber & Johnson, 2006). There is a well-known bias in which highly salient and emotional events are overweighted in memory tasks (e.g., Brown & Kulik, 1977; Madan et al., 2012; Madan & Spetch, 2012; Phelps & Sharot, 2008; Talarico & Rubin, 2003) and retrospective judgments, as in the peak-end rule (e.g., Fredrickson, 2000; Kahneman et al., 1993). Thus, one possibility is that the extreme outcomes are more likely to be retrieved at the time of the decision (cf. Johnson et al., 2007; Stewart et al., 2006), perhaps as a simplifying heuristic to limit the number of outcomes considered. More frequent retrieval of the extreme outcome would thus lead to more risk seeking for relative gains and more risk aversion for relative losses.

An alternate possibility is that the extreme outcomes are more salient (and thus overweighted) at the time of their occurrence, biasing the encoding of the learned values for the risky options (see Niv et al., 2012; Tsetsos et al., 2012). High extremes would increase the learned values for risky gains, causing more risk seeking, whereas low extremes would decrease the learned values for risky losses, causing more risk aversion in line with what we observe. The current results do not allow us to disambiguate these potential interpretations but suggest directions for future research.

The bulk of the literature on decisions from experience focuses on the key finding that rare events are underweighted in choice (e.g., Barron & Erev, 2003; Hertwig & Erev, 2009; Hertwig et al., 2004). Our experiments are complementary to that literature: There were no rare events, and all risky options led to two equiprobable outcomes. We found that common, extreme outcomes were overweighted in the decision process, leading to more risk seeking for relative gains than for relative losses (Figure 5; Ludvig & Spetch, 2011). The extreme-outcome rule does not explain the underweighting of rare events but does make a clear prediction: If a rare event is also an extreme, then there should be less underweighting than a parallel situation where the rare event is non-extreme. For example, take the range of outcomes from Experiment 1 (-40 to $+40$). The prediction is that there would be more underweighting of the rare event for an option that led to 95% +40 and 0 elsewise than an option that led to +40 only 5% of the time and 0 elsewise. That is, a rare, extreme outcome should be underweighted less than a rare, non-extreme outcome. Note that this prediction only holds in situations where multiple problems are intermingled so that not all outcomes are extreme in the decision context. Alternatively, it is also possible that the extreme-outcome rule would be dominated or non-applicable in situations with rare events.

Our results are not likely to be caused by a sampling bias—a particular concern with protocols that use rare events (Fox & Hadar, 2006; Hertwig & Erev, 2009; Rakow et al., 2008; Ungemach et al., 2009). All outcomes were experienced numerous times by participants, and the mean proportion of outcomes for the risky option never deviated significantly from .5 in any experiment. The risk seeking with relative gains (most notably in Experiments 2 and 4L) also rules out the "hot stove effect" (Denrell & March, 2001) that can occur from sequential sampling—whereby risk aversion emerges in experience-based choice due to the avoidance of ephemerally unlucky risky options (March, 1996; Niv et al., 2012). Furthermore, the inclusion of single-choice trials in each run assured that the planned contingencies were occasionally experienced by participants. Our results also cannot be attributed to a wealth effect (Thaler & Johnson, 1990), which would predict a consistent increase in risk seeking across the experiment, rather than the observed divergence in risk preferences between relative gains and losses.

Our task also adds some new methodological wrinkles to the study of decisions from experience (e.g., Camilleri & Newell, 2011; Erev et al., 2010; Hertwig & Erev, 2009; Hertwig et al., 2004; Ungemach et al., 2009; Weber et al., 2004). In some repeated-choice experiments, people repeatedly select (by clicking) from the same two options whose physical location is constant (for a given participant). This fixed location may introduce a "switch cost" in that it can be faster and easier for subjects to continue clicking in the same location rather than to move the mouse over to the other side. This potential switch cost may help induce some choice inertia or perseveration bias, which indeed has been observed (e.g., Erev et al., 2010).

The current task has several features that mitigate this potential bias. First, each trial was separated from the next by a short inter-trial interval ($1-2$ s) instead of

DOI: 10.1002/bdm

immediately following the previous choice. Second, the cursor was re-centered after each trial, forcing an equivalent movement to the left or the right on each choice. Third, from trial to trial, the location of the different doors were randomly counterbalanced, appearing on either side half the time—a common feature in studies with animals (e.g., Vasconcelos & Urcuioli, 2008). Thus, any perseverative side bias would appear as a random (50/50) choice and not a preference for either option. This last design feature also removes any potential confounds due to aliasing stimulus identity and stimulus location. Finally, as part of our primary manipulation, several different problems were intermingled; thus, the identity of the doors on a given trial could not be known in advance. Collectively, these design features produce a task where each choice requires active engagement with the available stimuli and their physical locations.

These features of the experimental design do indeed neutralize the perseverative bias. For example, in the problem with the 50/50 outcomes from the Technion Prediction Competition dataset (Problem 49; Erev et al., 2010), the perseveration rate was $87.0 \pm 2.4\%$ (calculated from the 20 subjects in the estimation set who encountered that problem). In contrast, in Experiment 1 here, people only selected the same side on the next trial $49.7 \pm 3.9\%$ of the time, meaning they were equally likely to perseverate or alternate. This result is not that surprising given that the location and identity of the stimuli changed from trial to trial.

As a further countermeasure to the potential disengagement of subjects, we included *catch trials* that are designed to incentivize people to attend to their choices. These catch trials provide an explicit means of ensuring that participants are, in fact, paying attention to their choices. On catch trials, the expected value of the options differs significantly. For example, in Experiment 1, some catch trials gave participants a choice between a guaranteed gain of +20 or a guaranteed loss of −20. Independent of variations in risk sensitivity, participants should choose the option with the higher expected value on these trials. By excluding participants who perform poorly (below 60%) on catch trials, we were able to ensure that the remaining participants were sufficiently incentivized. Importantly, unlike most outlier removal, this exclusion was not based on our primary dependent measure (risk preference on decision trials) but rather on a secondary measure (risk preference on the catch trials). Thus, our main results cannot be due to participants who may have ignored the stimuli and responded randomly.

In conclusion, we found that in decisions from experience, people chase the big win but avoid the big loss. The results provide evidence for an extreme-outcome rule, whereby the highest and lowest outcomes in a decision context are overweighted in choice. This potent role of extreme values in decision making has important real-world implications. For example, when gambling, people often choose between a smaller loss (the bet) and a larger win (the jackpot). Our results suggest that the overweighting of the largest wins with experience might contribute to an increased tendency to gamble.

ACKNOWLEDGEMENTS

This research was funded by grants from the Alberta Gambling Research Institute (AGRI) and the National Science and Engineering Research Council of Canada (NSERC) held by MLS. CRM was supported by NSERC, AGRI, and the Canadian Institutes of Health Research. EAL was supported by NIH Grant #P30 AG024361. We thank the Explore-Exploit Group at Princeton for insightful discussions and Ashley Rodgers for help with data collection. Door images were extracted from “Irish Doors” on fineartamerica.com with permission from Joe Bonita.

REFERENCES

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, *16*, 215–233.

Brown, R., & Kulik, J. (1977). Flashbulb memories. *Cognition*, *5*, 73–99.

Camilleri, A. R., & Newell, B. R. (2009). Within-subject preference reversals in description- and experience-based choice. *Proceedings of the Annual Conference of the Cognitive Science Society*, *31*, 449–454.

Camilleri, A. R., & Newell, B. R. (2011). When and why rare events are underweighted: A direct comparison of the sampling, partial feedback, full feedback and description choice paradigms. *Psychonomic Bulletin & Review*, *18*, 377–384.

Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science*, *12*, 523–538.

Erev, I., Ert, E., & Yechiam, E. (2008). Loss aversion, diminishing sensitivity, and the effect of experience on repeated decisions. *Journal of Behavioral Decision Making*, *21*, 575–597.

Erev, I., Glozman, I., & Hertwig, R. (2008). What impacts the impact of rare events. *Journal of Risk and Uncertainty*, *36*, 153–177.

Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R. et al. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, *23*, 15–47.

Ert, E., & Yechiam, E. (2010). Consistent constructs in individuals' risk taking in decisions from experience. *Acta Psychologica*, *134*, 225–232.

Fox, C. R., & Hadar, L. (2006). “Decisions from experience” = sampling theory + prospect theory: Reconsidering Hertwig, Barron, Weber & Erev (2004). *Judgment & Decision Making*, *1*, 159–161.

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

Hayden, B. Y., Heilbronner, S. R., Nair, A. C., & Platt, M. L. (2008). Cognitive influences on risk-seeking by rhesus macaques. *Judgment & Decision Making*, *3*, 389–395.

Heath, T. B., & Chatterjee, S. (1995). Asymmetric decoy effects on lower-quality versus higher-quality brands: Meta-analytic and experimental evidence. *Journal of Consumer Research*, *22*: 268–284.

Heilbronner, S. R., & Hayden, B. Y. (2013). Contextual factors explain risk-seeking preferences in rhesus monkeys. *Frontiers in Neuroscience*, *7*, 7.

Helson, H. (1947). Adaptation-level as frame of reference for prediction of psychophysical data. *The American Journal of Psychology*, *60*, 1–29.

Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences*, *13*, 517–523.

DOI: 10.1002/bdm

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*, 534–539.

Huber, J., Payne, J. W., & Puto, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, *9*, 90–98.

Johnson, E. J., Häubl, G., & Keinan, A. (2007). Aspects of endowment: A query theory of value construction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*, 461–474.

Kacelnik, A., & Bateson, M. (1996). Risky theories: The effects of variance on foraging decisions. *American Zoologist*, *36*, 402–434.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–292.

Kahneman, D., Fredrickson, B. L., Schreiber, C. A., & Redelmeier, D. A. (1993). When more pain is preferred to less: Adding a better end. *Psychological Science*, *4*, 401–405.

Lotto, R. B., Purves, D. (2000). An empirical explanation of color contrast. *Proceedings of the National Academy Sciences, USA*, *97*, 12834–12839.

Ludvig, E. A., & Spetch, M. L. (2011). Of black swans and tossed coins: Is the description–experience gap in risky choice limited to rare events? *PLoS One*, *6*, e20262.

Madan, C. R., & Spetch, M. L. (2012). Is the enhancement of memory due to reward driven by value or salience? *Acta Psychologica*, *139*, 343–349.

Madan, C. R., Fujiwara, E., Gerson, B. C., & Caplan, J. B. (2012). High reward makes items easier to remember, but harder to bind to a new temporal context. *Frontiers in Integrative Neuroscience*, *6*, 61.

March, J. G. (1996). Learning to be risk averse. *Psychological Review*, *103*, 309–319.

McCoy, A. N., & Platt, M. L. (2005). Risk-sensitive neurons in macaque posterior cingulate cortex. *Nature Neuroscience*, *8*, 1220–1227.

Miller, R. R., & Matzel, L. D. (1988). The comparator hypothesis: A response rule for the expression of associations. In G. H. Bower (Ed.), The psychology of learning and motivation (Vol. *22*, pp. 51–92). San Diego, CA: Academic Press.

Mitchell, T. R., Thompson, L., Peterson, E., & Cronk, R. (1997). Temporal adjustments in the evaluation of events: The “rosy view”. *Journal of Experimental Social Psychology*, *33*, 421–448.

Müller-Lyer, F. C. (1889). Optische Urteilstäuschungen. *Archiv für Physiologie, Suppl*, 263–270.

Niv, Y., Edlund, J. A., Dayan, P., & O’Doherty, J. P. (2012). Neural prediction errors reveal a risk-sensitive reinforcement-learning process in the human brain. *Journal of Neuroscience*, *32*, 551–562.

O’Neill, M., & Schultz, W. (2010). Coding of reward risk by orbitofrontal neurons is mostly distinct from coding of reward value. *Neuron*, *68*, 789–800.

Phelps, E. A., & Sharot, T. (2008). How (and why) emotion enhances the subjective sense of recollection. *Current Directions in Psychological Science*, *17*, 147–152.

Rakow, T., Demes, K. A., & Newell, B. R. (2008). Biased samples not mode of presentation: Re-examining the apparent underweighting of rare events in experience-based choice. *Organizational Behavior And Human Decision Processes*, *106*, 168–179.

Redelmeier, D. A., & Kahneman, D. (1996). Patients’ memories of painful medical treatments: Real-time and retrospective evaluations of two minimally invasive procedures. *Pain*, *66*, 3-8.

Shampanier, K., Mazur, N., & Ariely, D. (2007). Zero as a special price: The true value of free products. *Marketing Science*, *26*, 742–757.

Simonson, I. (1989). Choice based on reasons: The case of attraction and compromise effects. *Journal of Consumer Research*, *16*, 158–174.

Simonson, I., & Tversky, A. (1992). Choice in context: Tradeoff contrast and extremeness aversion. *Journal of Marketing Research*, *29*, 281–295.

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, *53*, 1–26.

Stewart, N., Chater, N., Stott, H. P., & Reimers, S. (2003). Prospect relativity: How choice options influence decision under risk. *Journal of Experimental Psychology. General*, *132*, 23–46.

Stone, A. A., Schwartz, J. E., Broderick, J. E., & Shiffman, S. S. (2005). Variability of momentary pain predicts recall of weekly pain: A consequence of the peak (or salience) memory heuristic. *Personality and Social Psychology Bulletin*, *31*, 1340–1346.

Stout, S. C., & Miller, R. R. (2007). Sometimes-competing retrieval (SOCR): A formalization of the comparator hypothesis. *Psychological Review*, *114*, 759–783.

Talarico, J. M., & Rubin, D. C. (2003). Confidence, not consistency, characterizes flashbulb memories. *Psychological Science*, *14*, 455–461.

Thaler, R. H., & Johnson, E. J. (1990). Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Management Science*, *36*, 643–660.

Thomas, D. R., & Jones, C. G. (1962). Stimulus generalization as a function of the frame of reference. *Journal of Experimental Psychology*, *64*, 77–80.

Tom, S. M., Fox, C. R., Trepel, C., & Poldrack, R. A. (2007). The neural basis of loss aversion in decision making under risk. *Science*, *315*, 515–518.

Tsetsos, K., Chater, N., & Usher, M. (2012). Salience driven value integration explains decision biases and preference reversal. *Proceedings of the National Academy of Sciences USA*, *109*, 9659–9664.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*, 453–458.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*, 297–323.

Ungemach, C., Chater, N., & Stewart, N. (2009). Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)? *Psychological Science*, *20*, 473–479.

Vasconcelos, M., & Urcuioli, P. J. (2008). Deprivation level and choice in pigeons: A test of within-trial contrast. *Learning & Behavior*, *36*, 12–18.

Weber, E. U., & Johnson, E. J. (2006). Constructing preferences from memories. In S. Lichtenstein & P. Slovic (Eds.), The construction of preference (pp. 397-410). New York: Cambridge University Press.

Weber, E. U., Shafir, S., & Blais, A.-R. (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review*, *111*, 430–445.

Yechiam, E., & Hochman, G. (2013). Losses as modulators of attention: Review and analysis of the unique effects of losses over gains. *Psychological Bulletin*, *139*, 497–518.

DOI: 10.1002/bdm