Human decision making in black swan situations

Amy Perfors (amy.perfors@unimelb.edu.au)
University of Melbourne

Nicholas T. Van Dam (nicholas.vandam@unimelb.edu.au)
University of Melbourne

Abstract

Real-world decisions often involve "black swan" choices with extremely low probability chances of catastrophic loss, like riding a motorcycle or going on a dangerous trip. These have several characteristics that make them especially difficult from the perspective of decision theory. How do people assign utilities to losses like "go bankrupt" or "die"? Do people have the representational resolution to encode differences between extremely tiny probabilities? We address these questions in two experiments in which people make decisions involving very low probabilities (as low as 1 in 10,000) of losing all of their points (and monetary bonus). Our results indicate that people mostly appear not to encode differences between tiny probabilities and are indifferent to the magnitude of losses. These factors lead to a startling qualitative shift in behaviour between scenarios with the same expected value and very similar absolute risk levels: people are risk averse when only one option is a black swan but become strongly risk seeking when both are.

Keywords: decision-making; probabilistic reasoning; risk; loss

Introduction

Should you ride a motorcycle? Go sky-diving? Go on a backpacking vacation overseas? Life abounds with decisions that combine the reasonably high odds of moderate gain (e.g., having fun) with extremely low odds of catastrophic loss (e.g., major injury or death). Known colloquially as "black swan" decisions, these situations pose multiple challenges. Yet much remains unknown about how people approach such decisions. This paper presents preliminary work aimed at that gap.

Black swan decisions are difficult for a variety of reasons. By definition, the event probabilities involved are extremely small; as a result, most people will not be faced with negative outcomes directly. For instance, consider that around 1 in 10,000 motorcyclists experience a fatality. This statistic means that most people will never personally know somebody who dies in a motorcycle accident; their lived experience puts the probability of death-by-motorcycle at zero. Conversely, for those few who *do* know someone who died, their fatality estimate would be much larger than the true rate. Decisions about black swan events must therefore be based on descriptions rather than direct experience.

The reason this matters is that a wealth of research suggests that people behave very differently depending on how they learn about the event probabilities involved. Consistent with Prospect Theory (Kahneman & Tversky, 1979), people generally overweight small probabilities when they are presented as descriptions but underweight them when they are learned through experience (Hertwig, Barron, Weber, & Erev, 2004). This may be due to effects of memory, statistical sampling, rational use of limited cognitive resources, or estimation error; regardless, the phenomenon is extremely robust (Lichtenstein, Slovic, Fischhoff, Layman, & Combs, 1979; Hertwig & Erev, 2009). These effects may be magnified by the fact that in addition to small probabilities, people also overweight extreme events (Madan, Ludvig, & Spetch, 2014; Lieder, Griffiths, & Hsu, 2018). Taken together, these factors predict that people should be averse to selecting options with a black-swan component, even if the expected value is higher than other options and the probabilities involved are very low.

However, another concern – representational resolution – yields a more muddled prediction. This issue revolves around the precision with which people encode small probabilities. It appears both that people generally represent numerical information logarithmically (Dehaene, 2003) and often do so in a fuzzy, gist-like fashion in which only qualitative distinctions between quantities are encoded (Reyna & Brainerd, 1991). If people do tend to treat all very small probabilities the same, this implies that the differences between very-low-probability options may be ignored more than (logarithmically equivalent) differences between higher-probability options. Indeed, if people encode sufficiently small probabilities as zero, that would mean that they might even be dismissed entirely.

Yet another complexity of black swan decisions is that it is not immediately obvious how to translate massive losses into utilities, since our normal trick of applying numbers is not necessarily appropriate. How should you represent the utility of death or major injury? Even apparently more straightforward outcomes like "lose all your money" are not necessarily straightforward, because it's not obvious that their utilities correspond to the numerical quantity of money involved. We know that in general people do not treat verbally- and numerically-presented information the same (Zwick, Wallsten, & Budescu, 1993) and that framing has a robust and pervasive effect on decision-making choices (Kühberger, 1998). Given these factors, it might be surprising if people *did* encode events like losing all their money in a straightforward numerical manner. But if they don't, how do they do it?

These considerations are not just idle academic worries; they are the fundamental component of the colourfully-named "Tyranny of Catastrophic Risk", which is a major concern for policy makers and economists who need to decide how to allocate resources toward black-swan scenarios like nuclear war (Buchholz & Schymura, 2012). There is no obviously correct solution. Representing the utility of black swan events as negative infinity means that one should devote infinite resources to ensuring that they do not happen. Conversely, if the probabilities of these events happening are small enough then any numerical values would effectively be ignored entirely.

A final interesting aspect of black swan problems is that the cumulative effect of many such decisions may be very different from any single one. If I have only 1:1000 chance of dying every time I get on my motorcycle, on any given day it might be sensible to go for a ride. Yet if I ride three times a day for decades, the cumulative odds of having a significantly shortened lifespan are inordinately high. How do people handle this discrepancy? Even if they appear not to notice very small probabilities in single events, do they take them into consideration when reasoning about long sequences?

It is perhaps because of these various and highly intermingled considerations that (to our knowledge) true black swan events are relatively understudied, at least in lab-based, controlled settings. Rare events are an intense focus of research, but "rare" is typically defined as less than 20% or so, with the vast majority of cases at 5% or higher. The smallest probabilities in laboratory tasks that we have found hover around 1 in 200 (e.g., Yechiam, Rakow, & Newell, 2015). That is rare, but it may not yet be in the realm where the issues above become a major concern, nor were those issues been the focus of that work. Similarly, when it comes to utilities, most lab tasks involve losses or gains corresponding to a fraction of people's points rather than all of them (or worse).

Given the complexity and tangledness of these issues, our work here is of necessity preliminary. Our goal was to capture the most important qualitative characteristics of black swan decisions (that they involve both a very low probability and an outcome with very extreme negative utility). Within that framework, we address three main questions. First, do people appear to track numerical utilities when they are framed in qualitative terms like "lose everything"? Second, do people represent the differences between extremely small probabilities? If they do, do they overweight the probabilities or underweight them? And third, do these behaviours change if people must make repeated decisions rather than a single one?

Experiment 1

We address these questions by putting people in a typical decision-making task with several modifications. First, we present people with situations in which at least one of the options is a black swan, with a very low probability of extreme loss (i.e., losing all of their points). We vary the number of points each person is given and frame the black swan outcomes as "lose everything" rather than referring to the specific number of points (although the specific points are always visible). If people make the same choices regardless of how many points they have to lose, this suggests that they are not making decisions based on the precise numbers involved.

Second, in order to determine whether people are ignoring these very small probabilities or are simply encoding them with poor resolution, we compare behaviour in two conditions with equivalent expected values but different numbers of black swan options. If people are ignoring small probabilities entirely, then they should prefer an outcome with high expected gain and a black-swan loss over an outcome with low expected gain but no loss. If, conversely, they are encoding small probabilities at low resolution, then people should focus on gains only and ignore the difference in magnitude between black-swan losses when both options have them.

Third, in order to determine whether people are aware of the implications of repetition when evaluating black-swan events, we ask people to choose a policy for playing 2000 games involving the same options they made one choice earlier.

Method

Participants A total of $N = 530$ US participants were run on Amazon Mechanical Turk; of these, 11 were excluded for failing an attention check, leaving 519 in the full dataset.[1] Ages ranged from 19 to 74 (mean: 36.5) and 255 (49.1%) were female. People were guaranteed a payment of $1.25USD for the 5-10 minute experiment regardless of their performance, with the possibility of a bonus if their choices resulted in additional gains.

Design This study had a 2x5 between-subjects design in which people were randomly allocated to one of two risk conditions (ONE RISK, $N = 266$ or TWO RISK, $N = 253$) and one of five starting point levels (3000, 4000, 5000, 6000, and 7000). People were told they would receive a bonus at the end of the study based on final points, and were asked to choose between a More Risky and Less Risky option in two ways: first, by making a single (ONE-SHOT) choice, and second – before receiving feedback on their first choice – by developing a POLICY according to which the computer would simulate 2000 games for them.

As Figure 1 shows, options were constructed so that in both the ONE RISK and TWO RISK conditions, the preferences based on expected utility (EU) theory followed the same pattern: if one has fewer than 5000 points, the More Risky choice is more EU-optimal, but this flips when one already has more than 5000 points. This occurs because the More Risky option is a black swan whose outcome (losing all points) is worse the more points one has. Simulating a sequence of 2000 games reveals a slightly different effect in which the More Risky option resulting in approximately 3000 points less at the end of the sequence than the Less Risky one (7000 to 10,000 or so).[2] Our question is whether people act correspondingly differently when asked a POLICY query.

The ONE RISK and TWO RISK conditions primarily differed in the nature of their Less Risky option, as shown in Figure 1. In the ONE RISK condition the Less Risky option might perhaps better be called riskless, because it involves

[1] 30 were pilot data used only to check for bugs in the code and not analysed beforehand, and thus included in the full sample.

[2] The reason simulations are different than the EU results in Figure 1 is that the simulation takes into account the fact that points vary as a function of prior success. The simulations spend more time at less than 5000 points thanks to occasionally losing everything, which is the region where the Less Risky option is preferred; the Less Risky option therefore does better overall. Technically, the EU-optimal strategy of switching at 5000 points performs slightly better than Less Risky alone; however, this amounted to a difference of about 50 points (around 0.5% of the total).

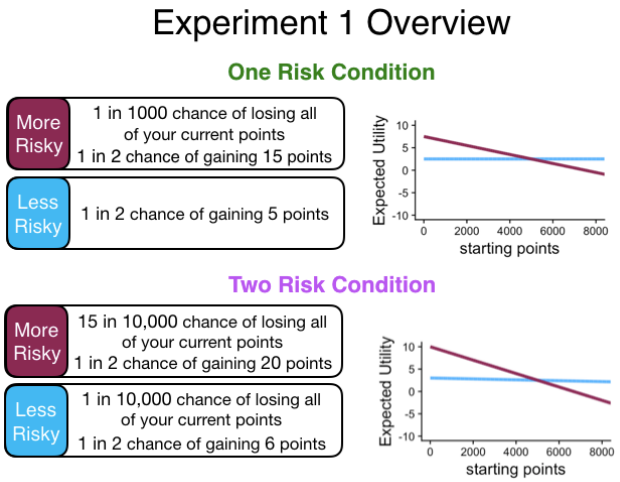

Figure 1: **Choices and expected utilities for Experiment 1**. Participants were randomly assigned to one of two risk conditions, each of which offered one More Risky option and one Less Risky option. In the ONE RISK condition, the Less Risky option was actually riskless, while in the TWO RISK condition, there was a 1 in 10,000 chance of losing all current points. Values were chosen so that the predictions assigned by expected utility theory followed the same qualitative pattern in both conditions: if people already had 5000 or fewer points, the More Risky option was preferred by EU theory, after which the Less Risky had higher expected utility. The *y* axes show the expected utility of a single game. Participants were randomly allocated different numbers of starting points (3000 to 7000) in order to evaluate whether people changed their strategies as EU theory would predict.

only a 50% chance of small (5 point) gain. The TWO RISK condition was designed as a comparison, because in it the Less Risky option is a black swan as well, albeit one that is extremely unlikely both in absolute terms (1 in 10,000) and relative to the More Risky option (fifteen times less probable). To the extent that people's behaviour in black swan situations is affected by the presence of risk at all (rather than the magnitude), we predicted that the ONE RISK and TWO RISK conditions would show different patterns of behaviour.

Procedure After filling out demographic information, participants took two short personality surveys[3] which were included as exploratory analysis to determine whether certain personality features were correlated with choices. For space reasons we defer analysis of this data to future work.

After completing the surveys, people were told that they had been allocated n points to begin with (with n varying between 3000 and 7000 by random allocation, as described above) and would be able to "cash out" any points they had at the end of the experiment as a bonus at an exchange range of 10 cents per 1000. To make sure they understood, the instructions explained what that would mean if the experiment were to end now; for instance, people who were allocated 5000 points would have read that "if the HIT were ending now you would be paid $1.75 (the $1.25 guaranteed rate plus 50 cents bonus)." The instructions further explained that people would be asked to make choices between different possibilities with different payoffs and losses. At the end the computer would simulate games following the choices they made, with the bonus depending on the outcome of the simulations.

Participants were not allowed to proceed further without correctly answering a series of questions (on a separate page) designed to ensure that they understood the task. These included a question about the task in the study, how many points they had been allocated, what the exchange rate was, and how their total payment was calculated.

After answering those questions, people were then shown the ONE-SHOT question on a new page which stated "You currently have n points. Please choose between the following two options:" and presented them with the More Risky and Less Risky options, called A and B (the mapping of option to A/B was random for each person). Upon making a choice, people were then shown a new page with the POLICY question, the top of which stated "You started off with n points. Before we tell you the results of the choice you just made, we have another question for you." The full instructions were:

> Here options A and B are the same as before (shown below). This time, though, instead of having the computer simulate one play, we're going to have it simulate 2000 plays. At each play it will make the choice you tell it to, and the results of that will be added or subtracted from your current point total. So for instance, if you had an option of 100% chance of gaining 1 point, then simulating 2000 plays would mean you end up earning 2000 points.

People were then given four options: Play A all 2000 times; Play B all 2000 times; Play A while I have more than x points, otherwise play B; and Play B while I have more than x points, otherwise play A. The latter two options had a blank text box in the place of the x, within which people could enter any number they wished. After making their choice, people went to another page where they were informed about the outcome of the simulations and their final bonus. Bonuses ranged from $0.00 to $2.58, with a mean of $0.82.

Results

Figure 2 shows what proportion of people chose each option in both risk conditions and for both types of queries. When people were asked to select between the options one time in the ONE-SHOT query, people in the ONE RISK condition preferred the Less Risky option slightly more than the More Risky one (although this result was unconvincing, $\chi^2(1) = 3.85, p = 0.0498$). However, this preference flipped dramatically in the TWO RISK condition, with many more people preferring the More Risky option ($\chi^2(1) = 35.8, p < 0.0001$).

The difference between risk conditions is significant ($\chi^2(1) = 31.3, p < 0.0001$). It suggests a mild degree of risk aversion when given a riskless option in the ONE RISK condition, coupled with a strong degree of risk *seeking* in the TWO

[3]The 12-item Intolerance of Uncertainty scale and a 20-item subset of the Big Five focusing on neuroticism and volatility.

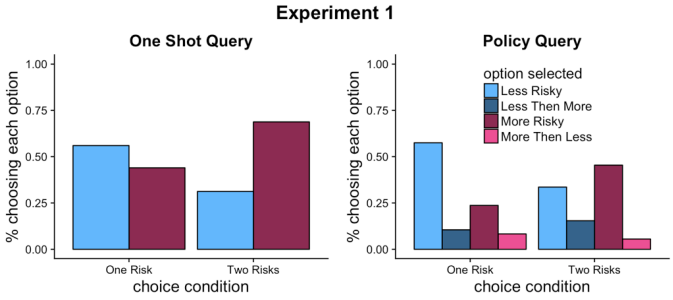

Figure 2: **Options chosen for each query type in Experiment 1**. *Left panel*: In the ONE-SHOT query in which people were asked to make one choice between the two options, they slightly favoured the Less Risky option overall in the ONE RISK condition. This preference flipped in the TWO RISK condition, despite the fact that both conditions were equivalent in terms of EU. Indeed, people were strongly risk-seeking even though the Less Risky option was nearly riskless, with only 1 in 10,000 chance of loss. *Right panel*: People showed the same qualitative pattern when asked to create a POLICY for playing 2000 games. Relatively few people chose a policy that involved switching between the More Risky and Less Risky, options, despite the fact that this is optimal. Of the remainder, as before, people were risk averse in the ONE RISK situation and risk seeking in the TWO RISK condition.

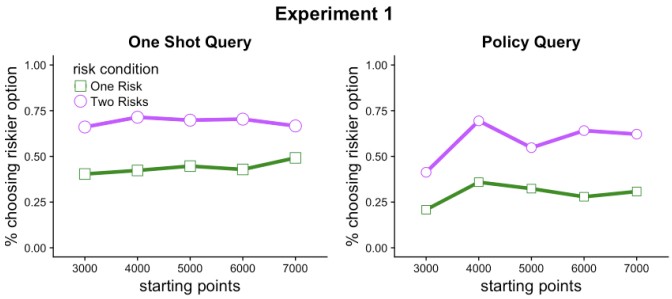

Figure 3: **Options chosen as a function of points in Experiment 1**. For both the ONE-SHOT and POLICY queries, people's choices did not differ based on their current number of points. Regardless of how much they had, people chose the riskier option in the TWO RISK condition and the less risky option in the ONE RISK condition. This behaviour contradicts the EU-optimal choice of switching from More Risky to Less Risky at the 5000-point mark.

RISK condition. This preference flip is consistent with the possibility that people are not precisely encoding the magnitude of the black swan probabilities but are noting the presence or absence of risk. They thus prefer the riskless option in the ONE RISK condition but treat the black swans in the TWO RISK case as qualitatively identical, resulting in a preference for the More Risky option because it has higher relative gains.

The qualitative effects are similar when we look at how people answered the POLICY question in which they were asked to choose a strategy for simulating 2000 games. Although expected utility theory favours a strategy that changes depending on one's current points, relatively few people chose that (fewer than 20% in both risk conditions). When we consider the 80% of people who opted to stay with one choice for all 2000 games, we see a similar pattern as before: preference for the Less Risky option in the ONE RISK condition ($\chi^2(1) = 37.5, p < 0.0001$) and slightly for the More Risky option in the TWO RISK condition, ($\chi^2(1) = 4.5, p = 0.034$). Risk aversion in general seemed to be greater when asked to make a POLICY decision, but not enough to change this qualitative difference between the ONE RISK and TWO RISK condition, which itself was significant ($\chi^2(1) = 32.9, p < 0.0001$).

To what extent are people's choices affected by how many points they currently have and thus how much they have to lose in the event of catastrophic loss? Interestingly, as Figure 3 illustrates, regardless of what query they were being asked, people showed little if any change in their decisions based on how many points they currently had.[4]

Taken together, these results suggest several things. First, the fact that behaviour did not change regardless of the starting points suggests either that the "lose everything" framing made people not pay attention to the numerical values involved, or that the difference between 3000 and 7000 points was not sufficiently large to affect behaviour. Given that this corresponds to a bonus that is almost a third the size of the actual payment and larger than has had an effect in other studies, this latter possibility seems unlikely but we cannot rule it out. Second, the fact that people preferred the Less Risky option in the ONE RISK case suggests that people are not encoding even very low (1:1000) odds as zero. However, the fact that they prefer the More Risky option in the TWO RISK case suggests that they *also* may not be encoding even a fifteen-fold difference in black-swan probabilities, instead treating them equivalently and making decisions based on the gains. Third, the fact that people are qualitatively similar but quantitatively somewhat distinct for the ONE-SHOT and POLICY queries suggests that they are attuned to the difference in reasoning about single vs multiple instances, albeit only slightly.

Experiment 2 is a follow-up study with two main goals. The secondary one is to determine whether these three findings replicate. The primary one is to determine whether the failure to encode differences occurs even when the probabilities are ten times larger: still black swan, but much less extreme. This results in EU calculations that now strongly favour the Less Risky option. If people still choose the More Risky option under these circumstances, this is a strong indication that they are truly indifferent to these distinctions in the absolute magnitudes of small probabilities, even when they have a profound impact on expected value.

[4]Since so few people chose the Less Then More or More Then Less options, for simplicity we restrict ourselves in this analysis only to the approximately 80% of people who chose Less Risky or More Risky alone. However, including everyone results in qualitatively identical statistical findings. For ONE-SHOT queries: ONE RISK: $\chi^2(4) = 2.44, p = 0.6546$; TWO RISK:$\chi^2(4) = 1.46, p = 0.8337$. For POLICY queries: ONE RISK: $\chi^2(4) = 1.84, p = 0.7649$; TWO RISK:$\chi^2(4) = 1.04, p = 0.9031$.

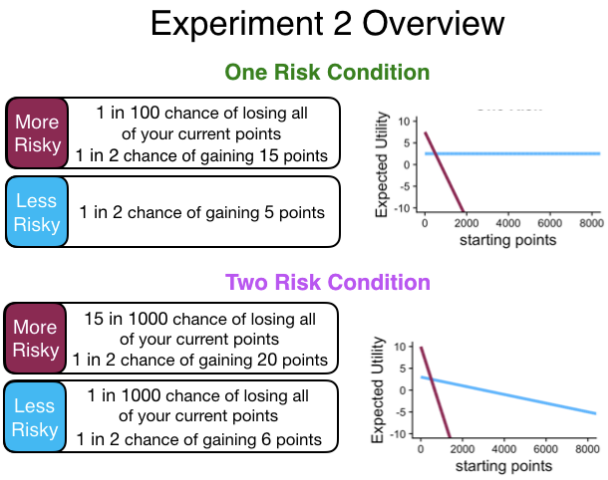

Figure 4: **Choices and expected utilities for Experiment 2**. As in the previous experiment, people were assigned to one of two risk conditions, each of which offered one More Risky option and one Less Risky option. This time EU theory strongly favours the Less Risky option at all point levels in the experiment, thus predicting that participants should always choose the Less Risky option.

Experiment 2

Method

Participants A total of $N = 500$ US participants were run on Amazon Mechanical Turk; of these, 15 were excluded for failing an attention check, leaving 485 in the full dataset. Ages ranged from 18 to 88 (mean: 36.7) and 231 (47.6%) were female. Payment was identical to Experiment 1 and none of the people had participated in Experiment 1.

Design and Procedure This experiment was exactly the same as Experiment 1, except that the probabilities and payoffs were different, as described in Figure 4. This time the probability of catastrophic loss (losing all points) was higher, although still quite infrequent in absolute terms. As a result, EU theory strongly favours the Less Risky option over the More Risky option for all of point levels endowed in the experiment (3000 to 7000 in increments of 1000, as before).

Because the utilities of the More Risky option fall off much more dramatically in the TWO RISK condition than they did in Experiment 1, the difference in expected outcomes after repeated games is more extreme as well. Simulations of a 2000-game sequence reveal that in the ONE RISK condition, the Less Risky option averages over 9000 points more than the More Risky option; in the TWO RISK condition this advantage still occurs but is reduced to about 2600.[5] Thus, although expected utility favours the Less Risky option regardless, it is favoured *more* when only one choice is a black swan. Because the probability of loss was higher in Experiment 2, the bonuses were smaller: from \$0.00 to \$1.32 with a mean of \$0.47.

[5] As before, technically the More Then Less policy (switching at 500 points) was most EU-optimal, but since switching occurred so early the difference between it and Less Risky only was minimal.

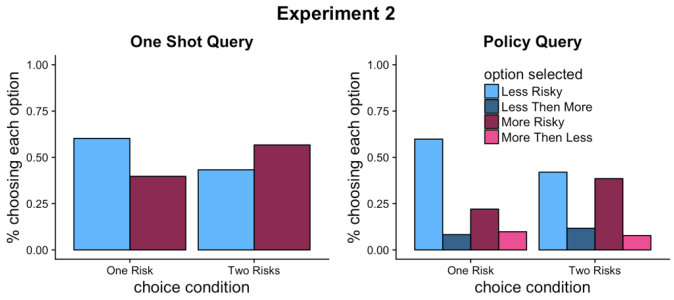

Figure 5: **Options chosen for each query type in Experiment 2**. As before, people asked both ONE-SHOT and POLICY queries were risk averse in the ONE RISK condition and risk seeking in the TWO RISK condition. Although the degree of risk seeking was somewhat lower than in Experiment 1, any at all strongly contradicts the predictions from EU theory, which very strongly favours the Less Risky option across the board in Experiment 2.

Results

As Figure 5 shows, people in Experiment 2 made essentially the same choices as did participants in Experiment 1, despite the fact that this time the Less Risky option has much higher relative expected utility across the board. As before, when people were asked just to select between the options one time in the ONE-SHOT query, people in the ONE RISK condition preferred the Less Risky option more than the More Risky one ($\chi^2(1) = 10.65, p = 0.0011$) while this preference was reversed in the TWO RISK condition ($\chi^2(1) = 4.16, p = 0.0414$). The difference between the risk conditions was significant ($\chi^2(1) = 13.25, p = 0.0003$).

Decisions were similar to Experiment 1 when considering responses to the POLICY question as well (right panel of Figure 5). As before, fewer than 20% of people chose a strategy that involved switching between the Less Risky and More Risky options. Of the remainder, most people favoured the Less Risky option in the ONE RISK condition (this time strongly so: $\chi^2(1) = 44.31, p < 0.0001$). Unlike in Experiment 2, people failed to show a preference for the More Risky option in the TWO RISK condition ($\chi^2(1) = 0.34, p = 0.5575$), suggesting that they were encoding the difference in black swan probabilities on at least some level. Still, given that the Less Risky option had *vastly* higher expected utility over the course of 2000 games, people's failure to prefer that option is somewhat striking. In any case, as before, the difference between the ONE RISK and TWO RISK conditions is still significant ($\chi^2(1) = 17.60, p < 0.0001$).

In another replication of Experiment 1, people in Experiment 2 also made similar choices regardless of how many points they were currently endowed with (see Figure 6). As before, no trends were significant; people showed similar risk aversion and risk seeking regardless how many points a catastrophic loss of all of them translated to.[6]

[6] ONE-SHOT: ONE RISK: $\chi^2(4) = 1.5, p = 0.8222$; TWO RISK:$\chi^2(4) = 3.01, p = 0.5565$. POLICY: ONE RISK: $\chi^2(4) = 2.39, p = 0.6639$; TWO RISK:$\chi^2(4) = 2.85, p = 0.5826$.

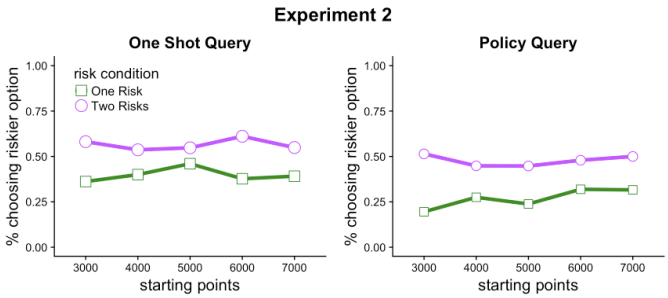

Figure 6: **Options chosen as a function of points in Experiment 2**. As in Experiment 1, for both kinds of queries, people's choices did not differ based on their current number of points. Regardless of how much they had to lose, people chose the riskier option in the TWO RISK condition and the less risky option in the ONE RISK condition. This behaviour contradicts the preference of EU theory for the Less Risky option regardless of how many points one started with.

Discussion

These two experiments reveal that people appear not to base their decisions on differences between tiny probabilities and are indifferent to the absolute magnitude of catastrophic losses. These factors lead to a qualitative shift in behaviour between scenarios with the same expected value and very similar absolute risk levels: people are risk averse when only one option is a black swan but become strongly risk seeking when both are. In some sense, our work replicates well-established findings: people in the ONE RISK condition are showing a classic certainty preference, while the TWO RISK condition demonstrates a type of isolation effect (Kahneman & Tversky, 1979). The interesting thing here is that both of these effects only emerge because of the way in which people are encoding the probabilities. Those in the TWO RISK condition are disregarding the shared black-swan component even though in a very real sense they are *not* shared: the risk is 15 times higher in one option than another. Conversely, those in the ONE RISK condition are treating 1 in 1000 odds as risky even though they are very small in absolute magnitude – much smaller than the differences being disregarded in the TWO RISK condition.

One possible objection to this work is that these were not true black swans because the losses were not large in absolute dollar terms. This is true, albeit somewhat unavoidable due to the ethical problems involved in visiting *actually* catastrophic outcomes onto unwitting participants. Somewhat reassuringly, participants on AMT appear to be highly motivated by small amounts of money and produce reliable data in decision-making studies (Paolacci, Chandler, & Ipeirotis, 2010). Moreover, the fact that people adopted the same strategy regardless of their points suggests that they were probably not thinking in dollar terms at all. We do plan to follow up on this latter effect to see if people show similar insensitivity to loss magnitude when the losses are not framed as "everything" but instead highlight the actual points involved.

Another limitation is that all of this work has focused on black swan losses rather than gains (very low probabilities of amazingly awesome outcomes). In part that was because they are asymmetric situations: it is possible to "lose everything" but not "win everything" in the same way (what is *everything?* All of the money in the world? What is the positive equivalent of death?) The other reason was that it would have been far more expensive to pay people for such large gains. The difference between gains and losses in black swan situations thus remains a huge open question, and it is quite likely that there are interesting differences in how people approach these situations (Kahneman & Tversky, 1979; Sharot, 2011).

Much remains to be done, but for now we note that these findings may have interesting implications for real-world black swan problems like vaccination. If people cannot be persuaded that the risks of vaccination are actually zero, they may be encoding even extremely small risks as something to avoid as in the ONE RISK condition. A more persuasive tactic might be to emphasise that *all* decisions are risky, as in the TWO RISK condition, thus making people more likely to act on the *benefits* of vaccination. This is speculative, of course, but demonstrates the importance of understanding how people reason about these extremely unlikely but very aversive events.

Acknowledgments

This research was supported by ARC grant DP150103280.

References

Buchholz, W., & Schymura, M. (2012). Expected utility theory and the tyranny of catastrophic risks. *Ecological Economics*, *77*, 234–239.

Dehaene, S. (2003). The neural basis of the Weber–Fechner law: A logarithmic mental number line. *Trends in Cog. Sci.*, *7*(4), 145–147.

Hertwig, R., Barron, G., Weber, E., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*(8), 534–539.

Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences*, *13*(12), 517–523.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–291.

Kühberger, A. (1998). The influence of framing on risky decisions: A meta-analysis. *Org. Beh. and Human Decision Proc.*, *75*, 23–55.

Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M., & Combs, B. (1979). Judged frequency of lethal events. *Journal of Experimental Psychology: Human Learning and Memory*, *4*, 551-578.

Lieder, F., Griffiths, T., & Hsu, M. (2018). Over-representation of extreme events in decision making reflects rational use of cognitive resources. *Psychological Review*, *125*(1), 1–32.

Madan, C., Ludvig, E., & Spetch, M. (2014). Remembering the best and worst of times: Memories for extreme outcomes bias risky decisions. *Psychonomic Bulletin & Review*, *21*(3), 629–636.

Paolacci, C., Chandler, J., & Ipeirotis, P. (2010). Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making*, *5*(5), 411–419.

Reyna, V., & Brainerd, C. (1991). Fuzzy-trace theory and framing effects in choice: Gist extraction, truncation, and conversion. *Journal of Behavioural Decision Making*, *4*, 249–262.

Sharot, T. (2011). The optimism bias. *Current Biology*, *21*(23), 941-945.

Yechiam, E., Rakow, T., & Newell, B. (2015). Super-underweighting of rare events with repeated descriptive summaries. *Journal of Behavioral Decision Making*, *38*(1), 67–75.

Zwick, R., Wallsten, T., & Budescu, D. (1993). Comparing the calibration and coherence of numerical and verbal probability judgments. *Management Science*, *39*, 176–190.