UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Introducing the Extinction Gambling Task

Permalink

https://escholarship.org/uc/item/4qk2v6s2

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors

Maier, Maximilian
Harris, Adam
Kellen, David
et al.

Publication Date

2024

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at https://creativecommons.org/licenses/by/4.0/

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

Introducing the Extinction Gambling Task

Maximilian Maier
(maximilianmaier0401@gmail.com)
UCL, Department of Experimental Psychology
26 Bedford Way, London, WC1H 0AP, UK

Adam Harris (adam.harris@ucl.ac.uk)
UCL, Department of Experimental Psychology
26 Bedford Way, London, WC1H 0AP, UK

David Kellen (davekellen@gmail.com)
Syracuse University, Department of Psychology
430 Huntington Hall, 13240 Syracuse NY, USA

Henrik Singmann (singmann@gmail.com)
UCL, Department of Experimental Psychology
26 Bedford Way, London, WC1H 0AP, UK

Abstract

Decisions about extinction risks are ubiquitous in everyday life and for our continued existence as a species. We introduce a new risky-choice task that can be used to study this topic: The Extinction Gambling Task. Here, we investigate two versions of this task: a *Keep* variant, where participants cannot accumulate any more earnings after the extinction event, and a *Lose* variant, where extinction also wipes out all previous earnings. We derive optimal solutions for both variants and compare them to behavioural data. Our findings suggest that people understand the difference between the two variants and their behaviour is qualitatively in line with the optimal solution. Further, we find evidence for risk-aversion in the Keep condition but not in the Lose condition. We hope that this task can facilitate further research on this vital topic.

Keywords: risky choice; extinction risks; decision making; repeated choice

Introduction

People often need to make choices that trade-off a benefit against a small risk of extinction (i.e., death or even human extinction). For instance, on the individual level, people may decide whether to speed or jaywalk when they are late for an appointment, which will increase the chance of being on time but comes with a small chance of death in an accident. On a collective level, people need to decide how much to invest to manage small probability high-stakes events, such as asteroid impacts or risks from extreme climate change.

Properties of Extinction Events

The defining property differentiating extinction events from other catastrophic events is that, after an extinction event, no further events can be experienced, thus precluding the possibility of any kind of utility gains. In some situations (for example, when working towards a goal that can either be reached or not), extinction may even wipe out all of the utility gains accumulated so far. This property makes it uniquely tricky to reason about extinction events because people need to take into account both their current situation and – more difficultly – the expected value of what could be achieved in the future that is at stake (e.g., the opportunity cost of going extinct). That being said, these difficulties do not appear to prevent people from often choosing risky options that include the risk of extinction, such as speeding. In other words, they see a benefit in going with the risky option.

Given these considerations, we argue that investigating people's choices when extinction risks are involved calls for a task with the following characteristics:

1. Be a multi-trial experiment, so that the extinction event can influence earnings accumulated so far, as well as earnings that would have incurred after the extinction event.
2. Allow participants to make decisions between a safe and a risky choice.
3. The risky choice should have higher payoff per trial but also include the possibility of extinction.

In this paper, we introduce a novel risky-choice task, the *Extinction Gambling Task*, that satisfies these desiderata. In this task, participants face a sequence of choices between a safe and a risky lottery, where the risky option results in a higher per-trial payoff but incurs a small chance of extinction. One attractive feature of this task is that one can derive the optimal decision strategy, which can be directly contrasted with people's choice behaviour.

How Might People Reason About Extinction?

One focus in the extant literature on risky choices is the difference between the objective probabilities and payoffs of presented options and their subjective counterparts, subjective probabilities and utilities. Both of these components likely also play a role in people's reasoning about extinction events. For example, standard models of decision-making under risk, such as prospect theory (Kahneman, 1979; Tversky & Kahneman, 1992), assume that people overestimate small probabilities. This would imply that people are likely risk-averse for (by definition rare) extinction events. On the other hand, the convex shape of prospect theory's value function for losses implies an underweighting of more extreme events which suggests people are likely risk-seeking when presented with options that include possibility of an extinction event. By risk-seeking and risk-averse preferences, we mean that, when presented with two (non-dominated) options that differ in terms of the variance of their outcomes, they prefer the one with largest or smallest variance, respectively (this definition will become clearer later on).

Interestingly, both risk-averse and risk-seeking preferences were observed in a study investigating risk-taking for *black swan events* (Perfors & Van Dam, 2018) – extremely rare events that wipe out current earnings but do not prevent participants from accumulating further earnings (i.e., very bad

In L. K. Samuelson, S. L. Frank, M. Toneva, A. Mackey, & E. Hazeltine (Eds.), *Proceedings of the 46th Annual Conference of the Cognitive Science Society.* ©2024 The Author(s). This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY).

events that are not extinction events). In this study, participants were asked to choose between two described lotteries, one more risky and one less risky. In the condition in which only the more risky lottery included a black-swan event, participants were found to be risk-averse – they generally preferred the less risky lottery even when it wasn't beneficial. In contrast, in the condition in which both lotteries included a black-swan event, participants were generally found to be risk-seeking – they now preferred the more risky lottery even when it wasn't beneficial. Whilst these results are instructive for our study, they also have important limitations. In addition to black-swan events differing from extinction events, participants did not experience the outcomes of any of their choices as they needed to specify the whole sequence of choices in advance (i.e., there was no feedback or learning).

Description Versus Experience. Further insights into how people understand the probabilities of rare events come from research into decisions from experience. When participants learn about small probabilities from experience (i.e., they decide between two options several times and observe the outcomes), they choose *as though* they underestimate rather than overestimate the probability of rare events (Hertwig et al., 2004; Lejarraga and Hertwig, 2021; Wulff et al., 2018; but see Kellen et al., 2016). Even in cases where participants also have full descriptive information about the options, it is often found that experience influences choices more than the probabilistic description of the risk (Hertwig & Wulff, 2022). This phenomenon is found not only in lab studies but also in real-world decision-making: In a review of the literature on natural hazard risk perception, Wachinger et al. (2013) found that personal experience is a strong factor in risk perception, while "the likelihood of a disaster is barely taken into account when making judgments about perceived risk" (Wachinger et al., 2013, p. 1051). These effects are likely exacerbated in the case of extinction events, as these events imply an insidious feedback regimen: conditional on thinking about the probability of the extinction event, one is alive and, therefore, has zero personal experience with it.

Overweighting Extreme Events. Another line of research shows that people overweight events that are extreme (e.g., in terms of severity) in both decisions from description and experience (Ludvig et al., 2014; Madan et al., 2014; Sunstein & Zeckhauser, 2011). This phenomenon is possibly a consequence of the availability heuristic (Tversky & Kahneman, 1973). Lieder et al. (2015, 2018) showed that overweighting of extreme events is in line with formulations of optimal information processing under limited cognitive resources and is, therefore, likely adaptive for many real-life decisions. As extinction events are, by definition, extreme events, the overweighting implied by this research suggests that people should hold risk-averse preferences.

Reasoning About the (Life-)Time Lost by Extinction. A key feature of extinction events is that their badness depends on what is foregone by going extinct. For example, on the individual level, most people would consider the death of a child worse than the death of an elderly person, and one reason for this is that the child loses out on much more potential lifetime due to their death; on the collective level, human extinction is often considered uniquely bad because of its impact on all future humans that could not come into existence if humanity goes extinct (Parfit, 1984, p. 453, Schubert et al., 2019).

When looking at the decision-making literature on how people might reason about possible future earnings, there are a number of different accounts that all point to the same general conclusion. People have problems adequately accounting for possible future earnings, which should lead to risk-seeking in the present. Firstly, people tend to exhibit opportunity cost neglect (Frederick et al., 2009): they fail to consider what they could have earned had they made other choices. Secondly, research on choice bracketing (Read et al., 2000; Tversky & Kahneman, 1981) shows that people have a prevalence for narrow bracketing (i.e., focusing mostly on the current choice rather than contextualising it within a wider set of relevant choices). When adopting a narrow bracketing and focusing only on the current choice, the risky option almost always has a higher payoff, and the disadvantages only become clear when adopting a wider bracketing and considering its effect on (preventing one from making) future choices. Consequently, choice bracketing would almost always imply that it is better to choose the risky option. Finally, research on scope insensitivity has shown that people do not value a good in proportion to its scope or size (Kahneman et al., 1999; Maier et al., 2023). In the context of extinction events, scope insensitivity would suggest that people are not sufficiently sensitive to the amount of time that is lost by extinction, especially when this amount is large (Coleman et al., 2023).

The Extinction Gambling Task

The previous section provided an overview of a number of considerations relevant to thinking about extinction related risky decisions. From this vantage point, it seems very difficult to predict, a priori, whether people will be risk-seeking or risk-averse when faced with an extinction related risky decision. To address this issue, we designed the Extinction Gambling Task, where participants repeatedly choose between a safe lottery and a risky lottery that includes a possibility of extinction. Specifically, participants were asked to choose 100 times between the following two lotteries:

- Risky Lottery:
  - 47.5% chance of winning £0,
  - 47.5% chance of winning £0.10,
  - 5% chance of extinction.
- Safe Lottery:
  - 50% chance of winning £0,
  - 50% chance of winning £0.01.

Based on these options, we calculated the expected value *per trial* of a safe choice as

$$\bar{r}_{\text{t, safe}} = 0.5 \times £0 + 0.5 \times £0.01 = £0.005, \tag{1}$$

and the expected value *per trial* of a risky choice, *assuming one does not go extinct* as

$$\bar{r}_{\text{t, risky}} = 0.5 \times £0 + 0.5 \times £0.1 = £0.05, \tag{2}$$

In this kind of repeated choice scenario with one extinction event, the key question is whether extinction merely means not being able to accumulate additional earnings in future trials or whether it means losing all earnings (both future and past). Therefore, we provide two alternative versions of the Extinction Gambling Task with associated optimal solutions, the *Keep* condition, in which participants can keep all earnings accumulated until they experience the extinction event, and the *Lose* condition, in which drawing the extinction event wipes out all earnings.

Keep Condition. In the Keep condition, participants are precluded from earning more from future trials when drawing the extinction event; however, they keep all earnings accumulated so far. Intuitively, this implies that the expected payoff of a strategy strongly depends on the position in the sequence of trials where the risky choices are played: It is rational to play safe gambles early on in the experiment because, in the beginning, the opportunity cost of going extinct is higher. In contrast, in the last trial the opportunity cost is zero and the cost of drawing the extinction option is equal to the £0 outcome. Therefore, the risky choice should be much more attractive at the end of the experiment. This further implies that it is always better to play all the safe gambles first and afterwards switch to playing the risky gambles. The expected value (EV) for the Keep condition assuming that a participant follows this optimal strategy – they first play all safe choices and afterwards all N_{risky} risky choices – is given by

$$\begin{aligned}\text{EV}(N_{\text{risky}}) = & \underbrace{\bar{r}_{\text{t, safe}} \times (N_{\text{total}} - N_{\text{risky}})}_{\substack{\text{Expected value of safe} \\ \text{trials}}} \\ & \underbrace{+ (1 - p(s)) \times \sum_{i=0}^{N_{\text{risky}}-1} (p(s)^i \times i \times \bar{r}_{\text{t, risky}})}_{\substack{\text{Expected value} \\ \text{from the risky trials} \\ \text{if one goes extinct}}} \\ & \times \underbrace{p(s)^{N_{\text{risky}}} \times N_{\text{risky}} \times \bar{r}_{\text{t, risky}}}_{\substack{\text{Expected value} \\ \text{from the risky trials} \\ \text{if one does not go extinct}}},\end{aligned} \tag{3}$$

where $\bar{r}_{\text{t, safe}}$ denotes the expected value of choosing the safe lottery, $\bar{r}_{\text{t, risky}}$ denotes the expected value of choosing the risky lottery (assuming the player does not go extinct), N_{total} denotes the total number of trials of the experiment, and $p(s)$ denotes the probability of surviving (i.e., $1 - \text{p(extinction)}$) when playing the risky gamble.

If the safe trials are played first, then the expected value from the safe trials is the number of safe trials times the expected value per trial (first line of Equation 3). The second line denotes the probability of going extinct in the risky trial multiplied by the payoff of this type of extinction, in other words, the expected value from the risky trials over all extinction outcomes. Finally, the third line denotes the probability of surviving the entire experiment multiplied by the payoff of the risky scenarios in this case.

The optimal number of risky choices can then be obtained by finding the maximum EV across all possible N_{risky},

$$\underset{N_{\text{risky}}}{\text{argmax}}\ \text{EV}(N_{\text{risky}}). \tag{4}$$

But what if participants do not follow the optimal ordering? If the risky trials are not played strictly after the safe trials, the expected value from the safe trials (first line of Equation 3) changes. For example, if the jth and kth trials are risky, we could calculate the expected value via

$$\begin{aligned}\text{EV}(\text{j}_{\text{risky}} \,\&\, \text{k}_{\text{risky}}) = & \underbrace{(\bar{r}_{\text{t, safe}} \times j - 1)}_{\substack{\text{In the first j-1 trials} \\ \text{you can't be extinct}}} \\ & \underbrace{+ (\bar{r}_{\text{t, safe}} \times ((k-1) - (j+1)) + \bar{r}_{\text{t, risky}}) \times p(s)}_{\substack{\text{between j+1 and k-1} \\ \text{the probability of existence is p(s)}}} \\ & \underbrace{+ (\bar{r}_{\text{t, safe}} \times (N - (k+1)) + \bar{r}_{\text{t, risky}}) \times p(s)^2}_{\substack{\text{between k+1 and N} \\ \text{the probability of existence is p(s) squared}}}\end{aligned} \tag{5}$$

In general, all safe choices played after risky choices would be multiplied with $p(s)^{N_{\text{risky so far}}}$. Finally, if the participant has already earned some endowment, the expected value would simply change by a constant of that endowment (as an already earned endowment cannot be lost in the Keep condition). This implies that the optimal strategy does *not* depend on the endowment or the earnings. Further, because the earnings so far do not influence the choice, optimality in the Keep condition implies always switching with the same distance from the last trial independent of the total number of trials (keeping the per-trial payoffs and probabilities constant).

Lose Condition. In the Lose condition, all past and future earnings are wiped out if the extinction event occurs. Because of this feature, playing all safe choices first and then all risky choices is not necessarily optimal anymore. Instead, the optimal solution is *dynamic* and takes participants' luck into account. For example, when starting with three risky choices, a participant could receive the maximum payoff three times, or receive zero three times. If they receive the maximum payoff three times, they have more to lose in subsequent risky choices than in the second case, which implies playing less risky in the following trials.

Obtaining the dynamic optimal solution requires the application of the *Bellman equation*, a standard method in economics (e.g., Dixit, 1990). In particular, this method uses

backward induction by using the relationship between the value function at one trial and the value function in the next trial. We can estimate the expected value of a risky choice with N remaining trials as

$$\begin{aligned} \mathrm{EV}_{\mathrm{risky}}(N,\varepsilon) =& p(r_1^{\mathrm{risky}}) \times \mathrm{EV}(N-1,\varepsilon + r_1^{\mathrm{risky}}) \\ &+ p(r_2^{\mathrm{risky}}) \times \mathrm{EV}(N-1,\varepsilon + r_2^{\mathrm{risky}}) \end{aligned} \tag{6}$$

and

$$\mathrm{EV}(0,\varepsilon) = \varepsilon, \tag{7}$$

where $p(r_1^{\mathrm{risky}})$ and $p(r_2^{\mathrm{risky}})$ denote the probabilities of the two possible payoffs of the risky option if one does not go extinct (in our experimental setup 10p and 0p with 47.5% chance each), and ε denotes the current earnings. Equation 7 simply describes that when there are no choices left the participants receives the earnings ε.

Analogously, we can estimate the expected value of a safe choice at any given time using

$$\begin{aligned} \mathrm{EV}_{\mathrm{safe}}(N,\varepsilon) =& p(r_1^{\mathrm{safe}}) \times \mathrm{EV}(N-1,\varepsilon + r_2^{\mathrm{safe}}) \\ &+ p(r_2^{\mathrm{safe}}) \times \mathrm{EV}(N-1,\varepsilon + r_2^{\mathrm{safe}}), \end{aligned} \tag{8}$$

where $p(r_1^{\mathrm{safe}})$ and $p(r_2^{\mathrm{safe}})$ denote the probabilities of the two possible payoffs of the safe option (in our experimental setup 0 and 1p with 50% chance each)

The choice in a given trial is then risky if $\mathrm{EV}_{\mathrm{risky}}(N,\varepsilon) > \mathrm{EV}_{\mathrm{safe}}(N,\varepsilon)$ and otherwise safe.

Finally, the expected value of the gamble, given that the participant follows the optimal dynamic strategy, is determined by

$$\mathrm{EV(N},\varepsilon) = \max(\mathrm{EV}_{\mathrm{risky}}(N,\varepsilon), \mathrm{EV}_{\mathrm{safe}}(N,\varepsilon)) \tag{9}$$

Experiment

In the previous sections we (1) argued that the extant literature makes conflicting predictions regarding whether participants would be risk-seeking or risk-averse for extinction events and (2) derived optimal solutions for the Extinction Gambling Task. We next test how risk-seeking participants are, and compare their choices to the optimal strategies.

Method

Participants. Participants were paid $1.50 to participate in a 10 minute study. The study was approved by the departmental ethics committee of UCL Experimental Psychology (EP/2021/005). Initially, 196 participants signed up for the study. We excluded participants based on four different comprehension/attention checks. After these exclusions, we obtained a final sample of 157 participants (55 female, 102 male; mean age = 38). 90 participants were in the Keep condition and 67 participants in the Lose condition.

Design, Materials, and Procedure. Participants played a sequence of 20 practice trials followed by 100 incentivised trials. Figure 1 visualizes the two lotteries participants faced in each trial: a safe lottery with a 50% chance of not winning anything and a 50% chance of winning 1p; and a risky lottery with a 47.5% chance of not winning anything, a 47.5% chance of winning 10p, and a 5% chance of extinction. Participants were informed about the probabilities of the different events, the associated payoffs, and the total trial number.

We implemented two different extinction conditions between subjects. In the Keep condition, participants could keep what they had earned so far but could not earn additional endowment in following trials. In the Lose condition participants' entire earnings would be wiped out and they could not earn any more money in future trials. After encountering an extinction event, participants still needed to play the remaining trials; however, the piggy bank in the top right turned into a mushroom cloud and they could not accumulate any more earnings. At the end of the experiment, participants were paid their earnings as bonus payment on top of their regular payment.

Results

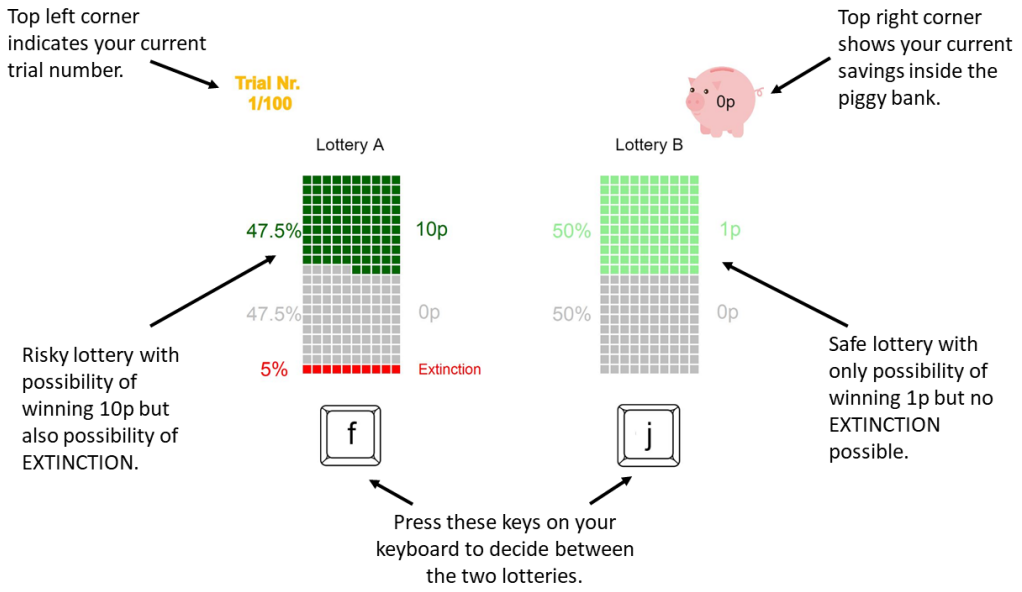
The top row of Figure 2 shows the optimal solution (red and orange), as well as participants' choices as a function of trial number. To account for the fact that changes in risky choices may be driven by selective attrition, we plot both the results for all participants who were still alive when a certain trial was played (black line) and the results only for those participants who survived the whole experiment (blue line). In the Keep condition (left panel), the optimal strategy is to switch from choosing the safe lottery to choosing the risky lottery after trial 56. In line with this optimal strategy, participants' risky choices increase toward the end of the task; however, participants' behaviour is still far from optimal, as evidenced by the distance to the red line. For the Lose condition (right panel), the optimal strategy is to start more riskily and then reduce the proportion of risky choices throughout; in this condition, participants' choices are remarkably similar to the optimal strategy.

The bottom row of Figure 2 shows the survival times of participants and the optimal solution (100 trials for participants that survived until the end of the experiment). We can see that more participants go extinct in the Keep compared to the Lose condition. Again, this shows sensitivity to the core task characteristics, as extinction is much worse in the Lose condition (people lose everything).

To statistically test people's sensitivity to the different task conditions, we fit a logistic mixed-effects model on participants' choices of the risky option with by-participant random effects. As fixed effects we included trial number (standardised to a range between -1 and 1 by substracting 50 and dividing by 50), trial number^2 (to account for the U shape seen in the Keep condition), and condition (Keep vs. Lose), as well as an interaction between trial number, trial number^2, and condition.[1] The optimal strategy would show increased

[1]The model with random slopes for trial number^2 did not converge successfully. We therefore used a random intercepts only

Figure 1: Explanation of the Experimental Setup as Shown to Participants on the Instruction Screen



risky choices as trial number increases in the Keep condition, and the opposite in the Lose condition. Further, more risky choices overall are expected in the Keep condition. In line with the optimal strategy, we see a significant difference between conditions, $z = 5.54, p < .001$, with the estimated proportion of risky choices at trial 50 being 10% [95% CI: 6%, 17%] in the Lose condition and 47% [95% CI: 35%, 60%], in the Keep condition. Furthermore, we find a significant interaction between trial number and condition, $z = 21.55, p < .001$, with the marginal slope being negative in the Lose condition $b = -1.29$, 95% CI $[-1.51, -1.06]$, and positive in the Keep condition $b = 1.86$, 95% CI $[1.69, 2.03]$. Overall, these results suggest that participants were sensitive to the differences between conditions in a way that was (qualitatively) consistent with the optimal strategy.

Are participants risk-seeking or risk-averse in comparison to the optimal strategy? To answer this question, we can compare the proportion of people who survived until the end of the task in our study to the proportion of participants who would have survived till the end if participants followed the optimal strategy.[2] If participants followed the optimal strategy, we would expect 10% to survive until the end of the task in the Keep condition and 67% in the Lose condition. A significantly higher proportion of our participants survived until the end of the task in the Keep condition (23%; 95%CI: 15%, 34%; $\chi^2[1] = 14.55, p < .001$). In the Lose condition, the proportion of survivors did not differ significantly from that predicted by the optimal solution (57%; 95% CI: 44%, 69%; $\chi^2[1] = 2.59, p = 0.11$). Overall, these findings suggest risk-aversion in the Keep condition and calibrated risk-taking in the Lose condition.

Concluding Comments

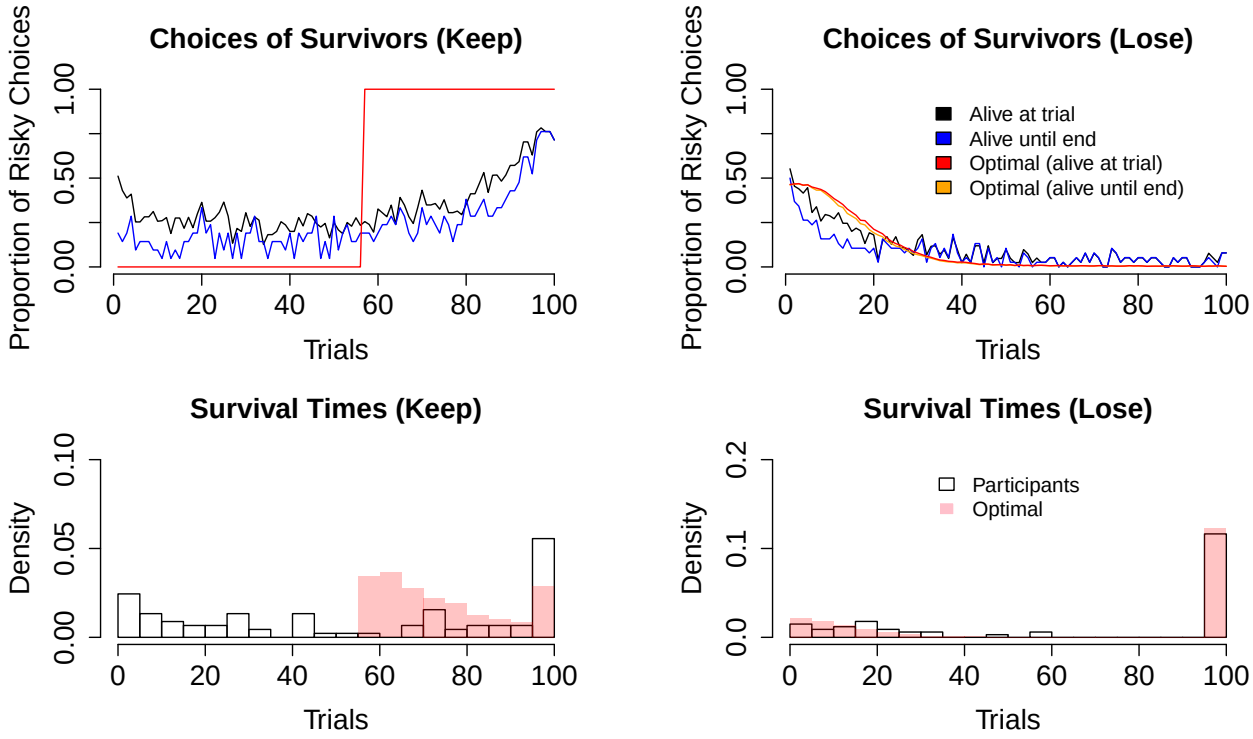
In this paper, we introduced the Extinction Gambling Task to study decision-making with extinction risks. We derived optimal solutions for two different types of extinction events: the Keep case, where participants can keep what they already earned but need to stop playing when drawing the extinction event, and the Lose case, where participants lose everything. We further compared participants' behaviour to the optimal strategies. We find that participants' behaviour is qualitatively in line with the optimal strategies, and we see evidence for risk-aversion in the Keep, but not in the Lose condition. We note that this pattern is consistent with well-known reports on gain-loss asymmetries, such as the reflection effect popularized by Tversky and Kahneman (1974) (assuming that people perceive losing the money in the piggybank as more of a loss than losing the opportunity to keep playing). Future work will explore the connection between these two phenomena and theoretical concepts such as 'loss aversion' and 'opportunity cost neglect'.

One remarkable result of our study is the similarity between participants' choices and the optimal solution in the lose case (Figure 2, right panels). This is particularly surprising given that the optimal solution involves solving a recursive function that is impossible to evaluate without a computer and even the authors of this paper did not initially suspect that the optimal solution would involve starting risky and continuously becoming more safe. Future research may investigate which heuristic people use to achieve this similarity. It is plausible that, rather than variations around a unique

model. The results are qualitatively the same and quantitatively very similar for the (non-converged) random slopes model.

[2]We consider this metric rather than other metrics, such as the proportion of risky choices as those are affected by selective attrition (i.e., more risky participants may drop out earlier, which would lead to overestimation of the proportion of risky choices).

Figure 2: Distribution of Survival Times and Choices of Survivors



Note. Alive at trial includes all participants that were not extinct at the current trial. Alive until the end includes all participants that survived until the end of the experiment. For the keep condition, those two lines are identical for the optimal solution; we, therefore, plot only the red line. The optimal solution for the Lose condition is based on simulating 10000 participants which follow the dynamic solution (Equations 6-9). The decision is executed based on a softmax decision function, as otherwise tiny differences in expected value in the first choices would lead to deterministic switching between 0 and 1. We use a very low temperature (0.02) to keep the solution as close as possible to the deterministic case. Note that the qualitative pattern arises for a range of temperature values.

strategy, different participants employ qualitatively different strategies' (e.g., Hey & Knoll, 2007; Kellen et al., 2017; Luce, 2010). We are, therefore, currently using mixture modelling approaches to better understand interindividual variability and the specific strategies that participants are using (cf. Lee & Courey, 2021). It would further be interesting to see how the resulting classifications relate with participants' ability to plan their choices ahead (e.g., Hey and Knoll, 2007; for a review, see Hotaling and Kellen, 2022). Finally, future research may introduce uncertainty by exploring generalizations of the task to settings where the maximum number of trials or the probabilities associated with the different rewards are unknown.

We conclude that the Extinction Gambling Task opens up many avenues for interesting theoretical and applied research. We hope that the new task will facilitate research into decision-making about extinction risks, a topic that is ubiquitous both in everyday life and for our continued existence as a species.

Data and Materials Availability

Data, materials, and analysis code are available at https://osf.io/cy6an/.

Acknowledgments

This work was supported by NSF CAREER Award (ID 2145308) attributed to David Kellen, Grant 100014_179121 from the Swiss National Science Foundation (SNSF) to Henrik Singmann, a PhD scholarship provided by Longview Philanthropy to MM, and a grant by the Forethought Foundation for Global Priorities Research to MM.

References

Coleman, M., Caviola, L., Lewis, J., & Goodwin, G. (2023). *How important is the end of humanity? lay people prioritize extinction prevention but not above all other societal issues*. https://doi.org/10.31234/osf.io/qn7k5

Dixit, A. K. (1990). *Optimization in Economic Theory* (2nd edition). Oxford University Press.

Frederick, S., Novemsky, N., Wang, J., Dhar, R., & Nowlis, S. (2009). Opportunity Cost Neglect. *Journal of Consumer Research*, *36*(4), 553–561.

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

Hertwig, R., & Wulff, D. U. (2022). A description–experience framework of the psychology of risk. *Perspectives on Psychological Science*, *17*(3), 631–651.

Hey, J. D., & Knoll, J. A. (2007). How far ahead do people plan? *Economics Letters*, *96*(1), 8–13.

Hotaling, J. M., & Kellen, D. (2022). Dynamic decision making: Empirical and theoretical directions. In *Psychology of learning and motivation* (pp. 207–238, Vol. 76). Elsevier.

Kahneman, D. (1979). Prospect theory: An analysis of decisions under risk. *Econometrica*, *47*, 278.

Kahneman, D., Ritov, I., Schkade, D., Sherman, S. J., & Varian, H. R. (1999). Economic preferences or attitude expressions?: An analysis of dollar responses to public issues. *Journal of Risk and Uncertainty*, (19), 203–242.

Kellen, D., Mata, R., & Davis-Stober, C. P. (2017). Individual classification of strong risk attitudes: An application across lottery types and age groups. *Psychonomic Bulletin & Review*, *24*, 1341–1349.

Kellen, D., Pachur, T., & Hertwig, R. (2016). How (in) variant are subjective representations of described and experienced risk and rewards? *Cognition*, *157*, 126–138.

Lee, M. D., & Courey, K. A. (2021). Modeling optimal stopping in changing environments: A case study in mate selection. *Computational Brain & Behavior*, *4*, 1–17.

Lejarraga, T., & Hertwig, R. (2021). How experimental methods shaped views on human competence and rationality. *Psychological Bulletin*, *147*(6), 535.

Lieder, F., Griffiths, T. L., & Hsu, M. (2015). Utility-weighted sampling in decisions from experience. *The 2nd Multidisciplinary Conference on Reinforcement Learning and Decision Making*.

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

Luce, R. D. (2010). Interpersonal comparisons of utility for 2 of 3 types of people. *Theory and Decision*, *68*, 5–24.

Ludvig, E. A., Madan, C. R., & Spetch, M. L. (2014). Extreme outcomes sway risky decisions from experience. *Journal of Behavioral Decision Making*, *27*(2), 146–156.

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

Maier, M., Caviola, L., Schubert, S., & Harris, A. J. (2023). Investigating (sequential) unit asking: An unsuccessful quest for scope sensitivity in willingness to donate judgments. *Journal of Behavioral Decision Making*, e2335.

Parfit, D. (1984). *Reasons and persons*. OUP Oxford.

Perfors, A., & Van Dam, N. (2018). Human decision making in black swan situations. In C. Kalish, M. Rau, J. Zhu, & T. Rogers (Eds.), *Proceedings of the 40th annual conference of the cognitive science society*.

Read, D., Loewenstein, G., Rabin, M., Keren, G., & Laibson, D. (2000). Choice bracketing. *Elicitation of preferences*, 171–202.

Schubert, S., Caviola, L., & Faber, N. S. (2019). The psychology of existential Risk: Moral Judgments about Human extinction. *Scientific Reports*, *9*(1), 1–8.

Sunstein, C. R., & Zeckhauser, R. (2011). Overreaction to fearsome risks. *Environmental and Resource Economics*, *48*, 435–449.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, *5*(2), 207–232.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124–1131.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*(4481), 453–458.

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

Wachinger, G., Renn, O., Begg, C., & Kuhlicke, C. (2013). The risk perception paradox—implications for governance and communication of natural hazards. *Risk Analysis*, *33*(6), 1049–1065.

Wulff, D. U., Mergenthaler-Canseco, M., & Hertwig, R. (2018). A meta-analytic review of two modes of learning and the description-experience gap. *Psychological Bulletin*, *144*(2), 140–176.