Title

People Missallocate Risk-Reduction Investments in Decisions Under Extinction Risk

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 48(0)

Authors

Maier, Maximilian
Wu, Leah Ruo Tong
Singmann, Henrik

Publication Date

2026

People Missallocate Risk-Reduction Investments in Decisions Under Extinction Risk

**Maximilian Maier (max.maier@wbs.ac.uk)[1,2], Leah Wu (leah.wu.22@ucl.ac.uk)[2] &
Henrik Singmann (singmann@gmail.com)[2]**
[1]Behavioural Sciences Group, Warwick Business School, University of Warwick
[2]Department of Experimental Psychology, University College London

Abstract

Many real-world decisions involve small probabilities of irrecoverable loss, where failure eliminates the possibility of making future decisions. Such decisions under extinction risk differ fundamentally from standard risky choice. We study how people allocate resources to reduce extinction risks. Participants play 100 trials with variable extinction probabilities of 2.5%, 5%, and 10%. If a participant draws the extinction outcome, they lose all accumulated bonus payment and are precluded from further earnings. Participants can reduce extinction risk by investing earnings: for every £0.02 invested, the extinction probability in the current trial is halved. We use dynamic programming to derive the optimal investment decision for each choice and compare it with participants' investments. Our results show that participants overinvest in reducing small risks, underinvest in reducing larger risks, and fail to adjust their investment appropriately across trials, leading them to go extinct at a higher rate compared to the optimal strategy.

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

Introduction

Not all risks we encounter in everyday life are of the same nature; the consequences of things going wrong can be relatively small (e.g., if we snooze the alarm once too often in the morning, we risk being late to work) or larger (e.g., if we cross the road without paying attention, we risk being hit by a car). In some situations, decisions even incur the (typically small) possibility of irrecoverable losses: At an individual level, severe accidents can end one's life, and on the collective level, large-scale events like extreme climate change can lead to global extinction. *Decisions under extinction risk* (Maier et al., 2025) describe decisions that involve a risk of ending the possibility of further decisions.

Decisions under extinction risk are not only practically important, but their unique potential outcome (preventing all future decisions) also implies different theoretical properties compared to decisions typically studied by cognitive scientists. Taking extinction risk seriously requires one to consider the full set of future choices that could still be made, and the optimal strategy depends strongly on the time horizon (Maier et al., 2025). In reality, people often exhibit narrow bracketing (i.e., viewing each choice in isolation, Read et al., 2000) and scope insensitivity (i.e., not scaling one's concern with the number of individuals affected, Kahneman et al., 1999). Further, decisions under extinction risk introduce an extremely biased learning signal: Conditional on one's existence, one cannot have experienced an extinction event before, and therefore the experience of those alive will always underestimate the risk (Ćirković et al., 2010; Wilson, 2023). These properties make it difficult to generalize from extant paradigms in decision sciences to decisions under extinction risk (see Maier et al., 2025, Appendix A, for a comparison with existing paradigms).

Crucially, extinction risks are not uncommon. In everyday life, individuals face multiple small extinction risks that compete for limited resources—such as the risk of fatal accidents when driving, fires when cooking, or undetected disease when foregoing medical check-ups. While each of these risks can be reduced through investment (e.g., safer cars, induction stoves, private screenings), it is neither feasible nor optimal to eliminate all risks entirely. Instead, decision makers must allocate finite resources across several extinction risks that may differ in probability. Understanding how people solve this allocation problem is, therefore, an important question for cognitive science. However, so far there is almost no work on this question. To address this gap, the present paper derives optimal strategies for allocating resources to reduce time-varying extinction risks and compares these normative benchmarks to participants' actual investment behaviour.

Prior Work

One set of prior work has investigated decisions under extinction risks or catastrophic risk in static decision-making environments. The key feature of this work is that while decisions entail the risk of extremely negative outcomes, decisions are investigated in isolation or decision makers specify a policy in advance without being able to adjust their decisions based on the experienced outcomes in previous choices. For example, Perfors and Van Dam (2018) asked participants to make choices between two lotteries where either one or both options entailed a small chance of an extremely large loss, either in a single shot setting or specifying a policy for 2,000 choices in advance. Results showed that if only one lottery entailed the possibility of a large loss, participants were risk averse (i.e., avoided this lottery even when it had positive expected value). But if both lotteries entailed the possibility of a large loss, they were risk seeking (i.e., chose the risky option even if it had negative expected value).

A different type of static extinction risk was investigated by Elga et al. (2025). In their study, participants had to repeatedly

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

allocate a set amount across three options to minimize the risk of any of them occurring in a given trial. Each option was associated with a specific, often large extinction probability (up to 99%), and any investment reduced the risk linearly. One feature of the optimal policy that people failed to follow in their specific setup was to allocate more funds to the options with largest extinction probabilities to reduce overall extinction risk. However, even though participants made choices across multiple rounds, each round was independent, and participants never risked going extinct, in the sense of not being able to make choices in later rounds.

A dynamic choice environment for investigating decisions under extinction risk was introduced by Maier et al. (2025). In their *Extinction Gambling Task*, participants repeatedly choose between two lotteries: one safe option with a small reward, and one risky option with a larger reward but a small (i.e., 5%) chance of extinction. The dynamic feature of the task was that participants made choices sequentially across up to 100 trials and received feedback after each choice. If participants chose the risky option and drew the extinction outcome in any trial, they were barred from making any further choices. Thus, participants needed to balance their desire for obtaining a large reward with the possibility of going extinct when choosing the risky option. The dynamic choice environment is arguably more representative of decisions under extinction risk outside the lab than static environments, as it involves factors such as observer selection effects that distinguish decisions under extinction risk from other types of risky choices. Results from Maier et al. (2025) showed that, while participants understood the requirements of the optimal policy relatively well and acted qualitatively in line with the corresponding optimal policies, their decisions were also affected by loss chasing (Gainsbury et al., 2014; Lister et al., 2016) and opportunity cost neglect (Frederick et al., 2009).

The Present Research

The aim of the present work is to use the dynamic features of the Extinction Gambling Task (Maier et al., 2025) for studying investment decisions to decrease the extinction probability (Elga et al., 2025). One shortcoming of the previous work with the Extinction Gambling Task is that it only considered binary choices between two options, one risky and one safe. While this setup matches some situations in which one of the options entails an extinction risk (e.g., driving somewhere versus staying at home), it does not capture other situations where people can invest limited resources to reduce risk (e.g., spending more on a car with better safety features).

In our *Investment Extinction Gambling Task*, each trial presents participants with a single lottery with two outcomes: one small monetary gain (5p) and a small chance of extinction. We adopted the 'lose' definition of extinction (see Maier et al., 2025), meaning that if participants draw the extinction outcome they lose all bonus payments earned so far and are also barred from earning further bonuses. The initial extinction probability differs across trials, randomly selected from 2.5%, 5%, or 10%. Participants choose how much of their endowment to invest to reduce the extinction probability. They can invest up to 10p in steps of 2p, with each 2p halving the extinction probability (e.g., if the initial extinction probability is 10%, 2p would reduce the probability to 5% and 4p would reduce the probability to 2.5%). Thus, in contrast to Elga et al. (2025), investment has a non-linear effect on extinction probability. As in Maier et al. (2025), participants must play a total of 100 trials without going extinct to obtain their bonus. Participants received an initial endowment of 50p, allowing them to invest money to reduce the extinction probability in the first trial(s).

We compare participants' decisions to the optimal policy. As described in detail in the next section, the optimal policy differs based on the initial extinction probability, participants' current endowment (bonus payment), and the remaining number of trials. We investigate how each of these three factors affected participants' investment decisions and what determines how strongly participants deviate from the optimal policy.

Optimal Investment Decisions

General Optimal Strategy. To derive the optimal investment for different extinction probabilities, we use Bellman Equations (Dixit, 1990). Bellman Equations use recursive backward induction to derive the expected value of different investments given the remaining number of trials and the current endowment. The basic idea is to break down the sequential decision problem by calculating the optimal policy at each step, assuming that future decisions will also be optimal (conditioning on the outcomes of the current decision). We can estimate the expected value of choosing an investment i, given a current endowment ϵ and N remaining trials:

$$\mathrm{EV}_i(N, \epsilon) = p(s \mid i) \sum_{x \in \mathcal{X}} p(x \mid s)\, \mathrm{EV}(N-1, \epsilon + r - i, \mathcal{X}), \tag{1}$$

where $p(\mathrm{s} \mid i)$ denotes the probability of survival under investment i, r denotes the payoffs after survival, and $\mathcal{X}$ is the set of possible extinction probabilities. Equation 1 denotes that the expected value of an investment i equals the chance of survival (given i) times the expected value of playing the optimal policy across the remaining $N-1$ trials with the new endowment $\epsilon + r - i$.

The recursive function stops at zero trials, at which point the expected value is equal to the current endowment,

$$\mathrm{EV}(0, \epsilon, \mathcal{X}) = \epsilon. \tag{2}$$

The optimal decision can then be calculated as the investment which has the highest expected value,

$$\mathrm{EV}(N, \epsilon) = \max_{i \in \mathcal{I}} \mathrm{EV}_i(N, \epsilon, \mathcal{X}). \tag{3}$$

Optimal Strategies for the Present Task. We implemented investment decisions under extinction risk in a task with 100 trials and an initial endowment of ϵ = 50p. For every trial in which participants avoided extinction, they would earn an

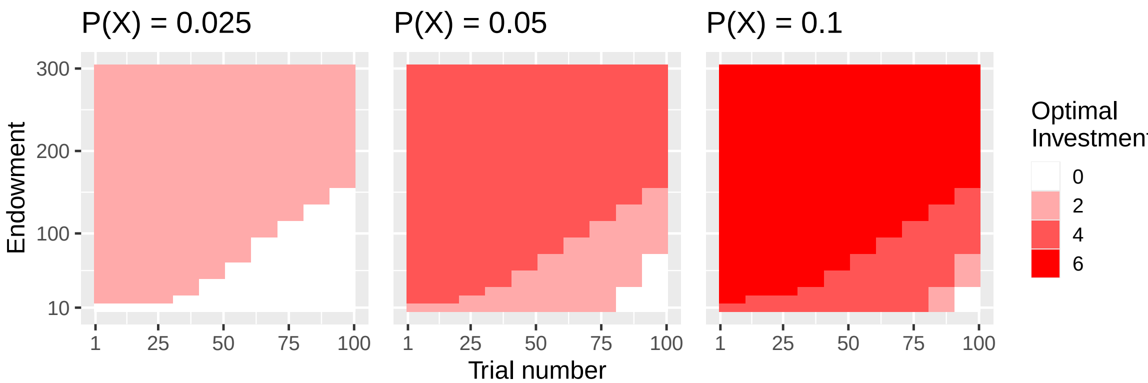

Figure 1: Optimal policy as a function of initial extinction probability $P(X)$ (in separate panels), where $X = 0.025, 0.05, 0.10$, the current endowment (on the y-axis), and trial number (on the x-axis). The colours indicate the optimal investment (in pence) for the payoffs and investment costs in the present experiment.

additional payoff of $r = 5$p. The set of possible extinction probabilities was $\mathcal{X} = 0.025, 0.05, 0.10$ and each probability was equally likely (i.e., randomly chosen for each trial).

Because recursive programming requires all possible combinations of future choices to be evaluated, it quickly reaches a combinatorial explosion for very fine-grained decisions. We, therefore, restricted the possible investments to values from 0p to 10p in increments of 2p. For every 2p invested, the probability of extinction in the current trial would be halved.

Figure 1 shows the optimal investment decision across trials and different probabilities in our setup. First, unsurprisingly, the figure indicates that the more endowment the player has, the more they should invest to avoid extinction (as the cost of extinction is higher with a higher endowment). However, for smaller endowments, the difference in endowment required to reach a larger investment is smaller than for larger endowments, as indicated by the relatively large area taken up by endowments of 2p (for $P(X) = 0.025$), 4p (for $P(X) = 0.025$), and 6p (for $P(X) = 0.1$). Further, the earlier in the task (i.e., lower trial number), the more a player should invest to avoid extinction, given a constant endowment. This is explained by the fact that the opportunity cost of going extinct is higher with more trials remaining (i.e., one foregoes more potential future +5p earnings upon extinction when more trials are left).

Method

Participants

We recruited 100 UK participants using Prolific. We first excluded three participants due to technical issues,[1] and one incomplete submission, leaving us with 96 participants. From this sample, we further excluded one participant who failed an attention check asking about the maximum number of trials, and 16 participants who failed an attention check asking what the three different pre-investment extinction probabilities in the task were. After exclusions we retained 79 participants (mean age 43.01, SD = 12.79; 28 female and 51 male). The study was approved by the ethics chair of UCL's Department of Experimental Psychology (EP/2021/005).

Design, Materials, & Procedure

Participants signed up for a 10-minute study with the opportunity to earn a bonus in addition to their base payment. After giving consent, they were provided with detailed task instructions. Participants were instructed about the total number of trials (100), the starting endowment (50p), the different initial extinction probabilities (2.5%, 5%, and 10%), and how investing money influenced the probability (every 2p invested halved the initial probability). Further, participants were provided with a practice slider for an example probability of 10% to familiarise themselves with how the investment affects the extinction probability, and a screenshot of the experimental task (see Figure 2).

To continue, participants needed to answer three forced comprehension checks correctly: 1) "What happens if you choose the risky lottery and you draw the extinction option?" (correct answer: "I will lose all my bonus money, and I cannot get any more bonus payments for future trials."), 2) "What does the piggy bank on the top right indicate?" (correct answer: "Your current bonus."), 3) "How does the investment slider impact the extinction risks if you choose to invest?" (correct answer: "For every 2p, the extinction risk is halved.").

Participants then played 20 practice trials[2] before proceeding to the main task involving a single session of 100 trials. When going extinct, participants still needed to complete the remaining trials, but we did not analyse the choices made

[1] One duplicate submission and two that seemed to have skipped or repeated a trial due to issues with internet connection

[2] Whether or not participants experienced the extinction event in the practice trials had no effect on their choices in the main task, neither as a main effect nor in interaction with initial extinction probability (using a linear mixed effect model; $p > .14$).

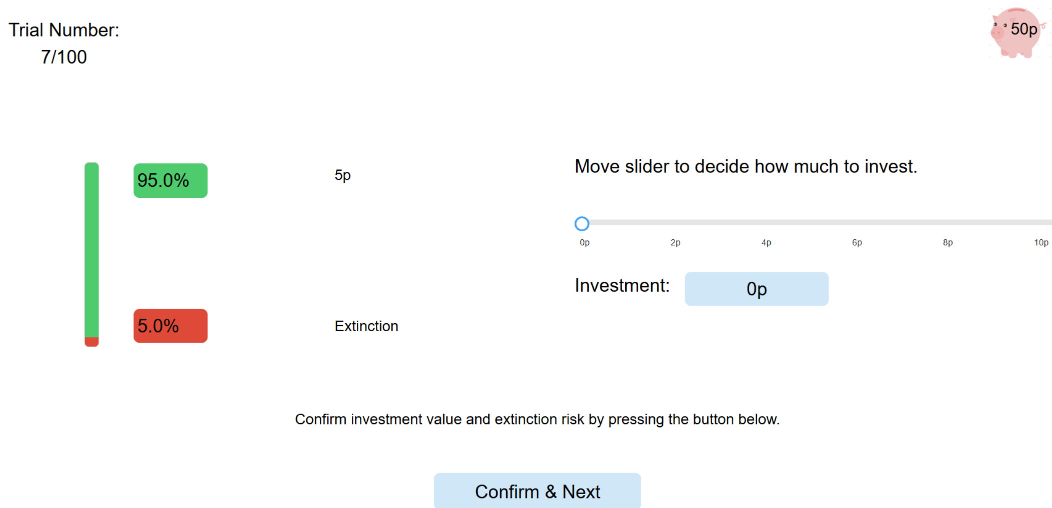

Figure 2: Experimental setup as shown to participants. Participants can move the slider to decide how much to invest and the probability will update dynamically. The top left shows the current trial and total number of trials, and the top right shows the current endowment (including the initial 50p endowment). Participants played a single session of 100 trials with probabilities varying across trials (2.5%, 5%, and 10%). We only analyse choices made before an extinction event occurred.

after extinction. Each participant could only experience the extinction event once during the main part of the experiment.

In each trial, participants were presented with the screen shown in Figure 2. Their only decision was how much money to invest to reduce the extinction risk. The initial extinction probability was randomly chosen for each trial from the three possibilities (2.5%, 5%, and 10%). Furthermore, the starting position of the slider was randomly chosen for each trial and could either be at the leftmost option (i.e., 0p) or rightmost option (i.e., 10p). After making their decision and confirming, participants were shown the outcome of their choice (i.e., either 5p or extinction), and proceeded to the next trial. Throughout the 100 trials, the current trial number, the maximum trial number, and the participant's current bonus were always displayed on the screen. If participants drew the extinction outcome, the piggy bank with the current bonus was replaced by a mushroom cloud and 0p, but they still needed to play the remaining trials so that going extinct to save time would not be a viable strategy. We did not analyse any choices that were made after an extinction event occurred.

Analysis

Our main dependent variable is the misinvestment, which we calculated for each participant and trial by subtracting their investment from the optimal investment. The optimal investment is specific to each participant's decision and considers their current endowment, trial number, and initial (pre-investment) extinction probability on that trial (see Figure 1). One particular challenge when analyzing decisions under extinction risk are selection effects: less data is available in the later trials than in the earlier trials, and participants who played more risky are more likely to be missing in later trials. To account for these selection effects, we compared the *estimated average misinvestment*, which is the estimated misinvestment based on a linear mixed models' fixed and random effects structure, assuming each participant had responded to all trials. This measure accounts for the selective dropout of participants and also takes non-linearities into account (we fit a model with linear and quadratic trial number term). For more details on this approach, including a simulation study showing that the approach indeed accounts for selective dropout, see Maier et al. (2025, Appendix C.1).

Results

Figure 3 visualizes the main results. The first row shows participants' investment (black) compared to the optimal investment (red). This illustrates two key findings. First, on average across all trials, participants seem to invest too much relative to the optimal policy for smaller probabilities ($P(X) = 0.025$) and too little relative to the optimal policy for larger probabilities ($P(X) = 0.10$). In other words, participants are not sufficiently sensitive to differences in risk magnitude.

Second, in the last trials participants invested more than implied by the optimal strategy. This can also be seen in Figure 3B, which compares participants' endowments on each trial with that of 1,000 simulated agents following the optimal strategy. Initially, participants were accumulating more money relative to the optimal policy, indicating that they invested too little to reduce the risk. However, at $\approx$ trial 60, participants almost stopped accumulating more money entirely, whereas the optimal strategy shows a much steeper slope. Because of this strong risk aversion toward the end, participants who survived until the end ultimately end up with somewhat less earnings than agents following the optimal strategy. Im-

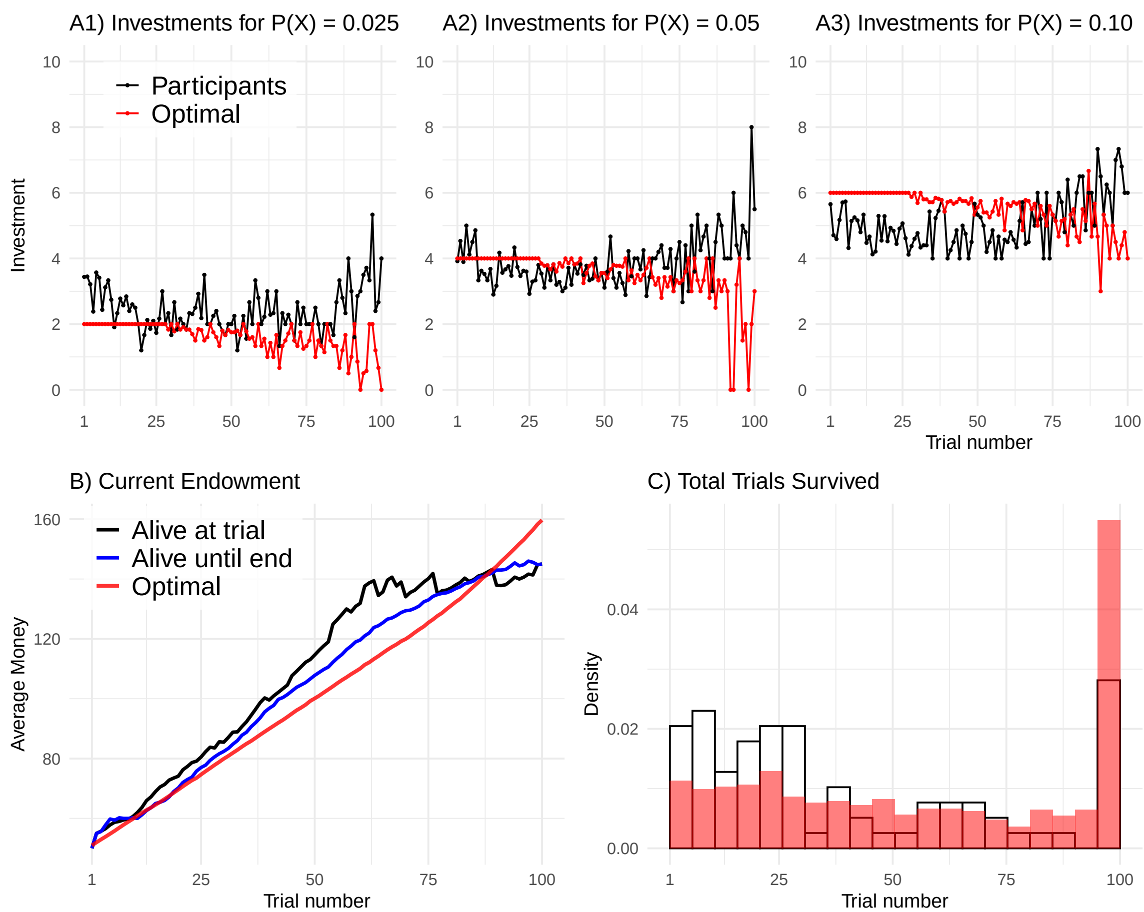

Figure 3: Panels A1-A3 show participants' average investment for the current trial number and initial extinction probability relative to the optimal investment, which is calculated using the endowments of those participants who made a decision for this trial and extinction probability. Panel B shows the current endowment of all participants that are not extinct at each trial, relative to the endowments of 1,000 simulated participants following the optimal strategy. Panel C shows participants' survival times (white) relative to the survival times of 1,000 participants following the optimal strategy (red).

portantly, this change in slope holds for both participants alive at a given trial (black line) and participants alive until the end (blue line), albeit to a weaker extent. Therefore, this pattern is not only due to earlier dropout of riskier participants.

Finally, Figure 3C compares participants' survival times to that of 1,000 simulated optimal agents. In line with the misallocation of investments over time (i.e., investing too little early and too much later) and the misallocation across the different probabilities (i.e., investing too much in smaller probabilities and too little in larger probabilities), more participants go extinct relative to the simulated optimal strategy, especially in the first 20 trials of the experiment. Overall, only 11.39% of participants (9/79) survived until the end of the task, whereas 27.2% of simulated agents following the optimal policy survived until the end. This also results in a substantially lower average payout for participants (13.75p, 95%-CI [3.34p, 24.15p]) relative to the optimal policy (43.61p).[3]

To further corroborate these patterns we estimated a linear mixed effects model predicting the difference between participants' investment and the optimal investment (i.e., the misinvestment) for each choice using trial number, trial number squared, initial extinction probability (2.5%, 5%, or 10%), the slider start point (0p or 10p; as a factor), and all interactions as fixed effects.[4]

Consistent with Figure 3, this model shows a main effect of initial extinction probability, $F(2, 39.75) = 36.24$, $p < .001$, indicating that participants' misinvestments differed with ini-

[3]The difference to the values in Figure 3B is that the present numbers are calculated using 0 earnings for all participants who drew the extinction event rather than being conditional on survival.

[4]We further included by-participant random intercepts and by-participant random slopes for all fixed effects (but no correlations to ensure model convergence). Because the maximal model showed a singular fit, we report the results of a model in which three random slopes were removed: the two three-way interactions and the two-way interactions of trial number squared. The results of this reduced model are almost identical to the maximal model.

tial extinction probability. To further understand this pattern, we calculated the estimated average misinvestment for each initial extinction probability, which all differed significantly from each other, ps < .001. For small initial extinction probabilities (i.e., 2.5%), participants invested on average 0.72p, 95%-CI [0.18p, 1.26p], more than was optimal. For medium initial extinction probabilities (i.e., 5%), participants' misinvestments did not differ significantly from the optimal policy, (0.07p, 95%-CI [-0.43p, 0.56p]). For large initial extinction probabilities (i.e., 10%), participants invested too little compared to the optimal policy, with an average misinvestment of -0.73p, 95%-CI [-1.27p, -0.18p].

The linear mixed model also indicated a significant quadratic effect of trial number on misinvestments, $F(1, 38.44) = 13.45$, $p < .001$, but no linear effect, $F(1, 28.62) = 1.46$, $p = .238$. This indicates that participants' misinvestments exhibited a U-shaped pattern (see Figure 3A), in line with the pattern in Figure 3B. The effect of trial number did not differ across initial extinction probability or slider start point (for all interactions, ps > .170). To further understand the pattern of misinvestment across trials, we calculated the estimated average misinvestment for the trials split into four equally sized blocks (i.e., block 1 = trials 1–25, block 2 = trials 26–50, etc.). This analysis showed that average misinvestment was negative for the first three blocks, but only significantly so for block 2 (block 1: -0.16, 95%-CI [-0.54, 0.23]; block 2: -0.50, 95%-CI [-0.94, -0.05]; block 3: -0.15, 95%-CI [-0.69, 0.39]). In block 4, the misinvestment was not significantly positive, 0.88, 95%-CI [-0.03, 1.80].

Finally, we also found a significant effect of slider start point, $F(1, 62.32) = 5.88$, $p = .018$. This effect indicated that when the slider started at 0 participants estimated average misinvestments, were somewhat lower than when the slider started at 10, (-0.22, 95%-CI [-0.77, 0.32] vs. 0.26, 95%-CI [-0.28, 0.80]). However, slider start position did not interact with any other variable, ps > .170.

Discussion

Our goal was to understand how people allocate resources over time when facing extinction risks of varying magnitudes. We investigated people's investment decisions in a lottery with a small potential payout of 5p and an initial extinction probability of either 2.5%, 5%, or 10% across 100 trials. Each invested 2p halved the initial extinction probability up to a maximum investment of 10p. People's choices diverged from the optimal policy in two ways. First, while participants understood that a larger initial extinction probability required a larger investment, they did not adapt their investment decisions sufficiently to the different initial risk levels: they overinvested for small probabilities and underinvested for large probabilities. Second, while the optimal policy prescribes a relatively flat investment pattern across trials with a reduction of investment if only a few trials are left, participants showed a U-shaped pattern, investing too little in the middle of the task. Thus, participants qualitatively failed to match the pattern prescribed by the optimal policy for how investments should change across trials (Figure 3B). Combining these two effects led to the pattern in Figure 3C: participants went extinct earlier and at a higher rate relative to the optimal policy.

Our result that people invest too much to reduce small risks and too little to reduce larger risks is consistent with previous work. Elga et al. (2025) found the same result in a static choice task where people needed to allocate resources across three large risks in each trial. Further, in a static task that required participants to specify an optimal policy in advance, Perfors and Van Dam (2018) found that people did not distinguish sufficiently between smaller and larger risks if any risk was present, leading them to become risk seeking when facing decisions between two extinction risks. The convergent evidence across different experimental paradigms suggests that this is an important factor that may lead to suboptimal decisions under extinction risk in the real world. A crucial direction for future research is therefore to identify interventions that improve people's sensitivity to risk magnitude and their ability to appropriately scale investments. The present study already incorporated a visual display of risk magnitude to improve people's sensitivity to different magnitudes of extinction risk. Future work could examine alternative formats (e.g., Gigerenzer & Hoffrage, 1995; Hammitt & Graham, 1999).

The current study only compared participants' behaviour to a risk-neutral optimal policy. Future work should consider how to derive a policy in terms of subjective utility rather than expected value. How such a subjective utility model would be designed is directly linked to questions about how people tackle the problem set out in this task. For instance, future work could compare computational models involving narrow bracketing (considering only the current choice) and models that integrate the assessment over multiple trials with varying time horizons (Read et al., 2000). These models could, in turn, be derived from models of the underlying processes that give rise to narrow versus broad bracketing (e.g., working memory limitations, dynamic inconsistencies in planning; Barkan & Busemeyer, 1999; Hotaling & Kellen, 2022). Further, we fixed the time-horizon to 100 trials, whereas in most real-world choices the exact time-horizon is not known. An open research question is whether we still see an increase of investment toward the end with a more fuzzy time horizon or whether participants are better adapted to this type of scenario.

Individuals and societies face many risks that, if realized, can impede any possibility for making future decisions. Arguably, at the societal level, these risks are currently higher than at most times in human history, with several collective extinction risks present simultaneously (e.g., extreme climate change, Artificial Intelligence, pandemics, nuclear war). This makes understanding how people decide and how to improve decision-making in the face of these risks an important challenge for cognitive scientists in the 21st century. We hope that the present manuscript can serve as a starting point for future research investigating how individuals and societies can make better decisions under extinction risk.

Data and Materials Availability

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

Acknowledgments

We thank Jian-Qiao Zhu for helpful discussion and Vanessa Cheung for helpful comments on the manuscript.

References

Barkan, R., & Busemeyer, J. R. (1999). Changing plans: Dynamic inconsistency and the effect of experience on the reference point. *Psychonomic Bulletin & Review*, *6*(4), 547–554. https://doi.org/10.3758/BF03212962

Ćirković, M. M., Sandberg, A., & Bostrom, N. (2010). Anthropic shadow: Observation selection effects and human extinction risks. *Risk Analysis*, *30*(10), 1495–1506. https://doi.org/10.1111/j.1539-6924.2010.01460.x

Dixit, A. K. (1990, September 13). *Optimization in Economic Theory* (2nd edition). Oxford University Press.

Elga, A., Zhu, J.-Q., & Griffiths, T. L. (2025). People make suboptimal decisions about existential risks. *Cognition*, *265*, 106216. https://doi.org/10.1016/j.cognition.2025.106216

Frederick, S., Novemsky, N., Wang, J., Dhar, R., & Nowlis, S. (2009). Opportunity cost neglect. *Journal of Consumer Research*, *36*(4), 553–561. https://doi.org/10.1086/599764

Gainsbury, S. M., Suhonen, N., & Saastamoinen, J. (2014). Chasing losses in online poker and casino games: Characteristics and game play of internet gamblers at risk of disordered gambling. *Psychiatry Research*, *217*(3), 220–225. https://doi.org/10.1016/j.psychres.2014.03.033

Gigerenzer, G., & Hoffrage, U. (1995). How to improve bayesian reasoning without instruction: Frequency formats. *Psychological Review*, *102*(4), 684. https://doi.org/10.1037/0033-295X.102.4.684

Hammitt, J. K., & Graham, J. D. (1999). Willingness to pay for health protection: Inadequate sensitivity to probability? *Journal of Risk and Uncertainty*, *18*(1), 33–62. https://doi.org/10.1023/A:1007760327375

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., 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. https://doi.org/10.1023/A:1007835629236

Lister, J. J., Nower, L., & Wohl, M. J. (2016). Gambling goals predict chasing behavior during slot machine play. *Addictive behaviors*, *62*, 129–134. https://doi.org/10.1016/j.addbeh.2016.06.018

Maier, M., Harris, A. J. L., Kellen, D., & Singmann, H. (2025). Decision making under extinction risk. *Cognitive Psychology*, *159*, 101735. https://doi.org/10.1016/j.cogpsych.2025.101735

Perfors, A., & Van Dam, N. T. (2018). Human decision making in black swan situations. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *40*. https://escholarship.org/uc/item/3p76280v

Read, D., Loewenstein, G., Rabin, M., Keren, G., & Laibson, D. (2000). Choice bracketing. In B. Fischhoff & C. F. Manski (Eds.), *Elicitation of preferences* (pp. 171–202). Springer. https://doi.org/10.1007/978-94-017-1406-8_7

Wilson, P. A. (2023). The anthropic principle. In *Cosmology* (pp. 505–514). CRC Press.