Cognitive Psychology 159 (2025) 101735

Contents lists available at ScienceDirect

Cognitive Psychology

journal homepage: www.elsevier.com/locate/cogpsych

Decision making under extinction risk ☆,☆☆

Maximilian Maier [a], *, Adam J.L. Harris [a], David Kellen [b], Henrik Singmann [a]

[a] Department of Experimental Psychology, University College London, United Kingdom
[b] Department of Psychology, Syracuse University, United States of America

ARTICLE INFO

Dataset link: Decision Making Under Extinction Risk (Original data)

Keywords:
Decision making
Risky choice
Extreme outcomes
Extinction
Mixture modelling
Computational modelling

ABSTRACT

In everyday life, people routinely make decisions that involve irredeemable risks such as death (e.g., while driving). Even though these decisions under extinction risk are common, practically important, and have different properties compared to the types of decisions typically studied by decision scientists, they have received little research attention. The present work advances the formal understanding of decision making under extinction risk by introducing a novel experimental paradigm, the Extinction Gambling Task (EGT). We derive optimal strategies for three different types of extinction and near-extinction events, and compare them to participants' choices in three experiments. Leveraging computational modelling to describe strategies at the individual level, we document strengths and shortcomings in participants' decisions under extinction risk. Specifically, we find that, while participants are relatively good in terms of the qualitative strategies they employ, their decisions are nevertheless affected by loss chasing, scope insensitivity, and opportunity cost neglect. We hope that by formalising decisions under extinction risk and providing a task to study them, this work will facilitate future research on an important topic that has been largely ignored.

Throughout their daily lives, people make decisions involving risks that could change their lives as they know them: a potential jaywalker decides whether to cross the street in front of an incoming car, a driver decides whether they have enough space to make an overtaking manoeuvre, and a boxer decides whether to risk a career-ending injury in a high-stakes bout. At a collective level, people need to decide how much to invest in managing small probability high-stakes risks, such as extreme climate change or pandemics.

The common denominator across all these examples is the risk of irredeemable extreme outcomes. In the current paper, hopefully without sounding too ominous, we refer to such examples as *decisions under extinction risk*. Extant literature, recent history, and everyday anecdotes ('*that* driver') suggest that people are bad at making decisions about these types of risks (Elga et al., 2024; Wiener, 2016). However, in any real-world decision under extinction risk, a large set of psychological factors are at play simultaneously and it is difficult to disentangle the relative influence of different factors. For instance, inadequate social distancing

☆ This work was supported by NSF CAREER Award, United States (ID2145308) attributed to DK, Grant 100014_179121 from the Swiss National Science Foundation (SNSF) to HS, a PhD scholarship provided by Longview Philanthropy, United Kingdom to MM, and a grant by the Forethought Foundation for Global Priorities Research, United Kingdom to MM.

☆☆ We thank Rani Moran for helpful feedback on a previous draft and suggesting the name Reset condition. We are also grateful to Dirk Wulff for helpful feedback on the initial task design and theoretical ideas. We thank Maarten Speekenbrink, Duarte Goncalves, and the members of the UCL Centre for Behavioural Data Science for their helpful suggestions and ideas for improvement. We also thank Jack Stecher for pointing us to a number of relevant choice paradigms. Finally, we thank Vanessa Cheung for helpful comments and suggestions on a previous version of this manuscript.

Experiment 1 of this manuscript has been previously included in the non-archival *Proceedings of the Cognitive Science Society* (Maier et al., 2024). All other experiments and the computational modelling results were added in this version.

* Corresponding author.
E-mail address: maximilianmaier0401@gmail.com (M. Maier).

https://doi.org/10.1016/j.cogpsych.2025.101735
Received 16 January 2025; Received in revised form 13 May 2025; Accepted 13 May 2025
Available online 24 June 2025
0010-0285/© 2025 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

during the Covid-19 pandemic could have been the result of self-interested decision making (especially with regards to younger people for whom COVID-19 was on average less severe), scope insensitivity (i.e., people do not take the large number of people affected from a pandemic outbreak into account), or underestimating the probability and cost of getting severely ill or even dying from the disease, to name just a few of many plausible factors.

Anyone tackling the challenges posed by extinction events will find little relief in the pages of mainstream decision making research. To see this, consider decision making research's 'fruit-fly', the binary-choice lottery paradigm (e.g., Kahneman & Tversky, 1979; Rieskamp, 2008). In this paradigm, participants are presented with one lottery problem at a time, for example choosing between a sure win of £50 versus a 50% chance of winning £110. These lottery problems are independent in the fullest sense of the term: at no point do outcomes of these decisions affect the possibility of future choices or their outcomes. By contrast, the outcomes of decisions under extinction risk might reduce the options available in future decisions—or prevent any future decision making at all.

Of course, the binary lottery paradigm is only one among the many different options available to researchers. Paradigms such as the Balloon Analogue Risk Task (Lejuez et al., 2003, 2002; Pleskac, 2008; Pleskac et al., 2008; Wallsten et al., 2005), or the Bomb Risk Elicitation Task (BRET, Crosetto & Filippin, 2013), mitigate some of the limitations of the binary lottery paradigm. However, these tasks were not designed to study extinction risks as discussed in this article. They therefore come with several limitations that make them unsuitable to study decisions under extinction risk, such as dependencies between choice and event probabilities, impossibility of learning from experience, or the probability of the extinction event being much higher than realistic in real-world settings (for a detailed review of potentially relevant tasks and their limitations, see Appendix A).

Given that existing paradigms do not allow studying extinction risks, what predictions can be derived from the extant decision making literature? Prospect theory would suggest that low probability events are overweighted and people are loss-averse (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992). We might therefore expect people to be risk-averse for decisions involving low-probability, high-impact options (and thus avoid jaywalking). However, this overweighting of low probabilities does not necessarily extend to contexts in which people learn about the probabilities from experience, a discrepancy known as the description-experience gap (Hertwig et al., 2004; Kellen et al., 2016; Lejarraga & Hertwig, 2021; Wulff et al., 2018). For many extinction risks, people often lack access to descriptive information and instead rely on personal experience (or the experiences of those they know). Further complicating this picture, other studies involving experienced options point to people overweighting extreme events (i.e., events with extremely good or bad utilities; see Lieder et al., 2015, 2018; Ludvig et al., 2014; Madan et al., 2014; Sunstein & Zeckhauser, 2011). As extinction events clearly have extreme (dis)utility, this literature would predict their overweighting.

In addition to the probability of the extinction event, another factor that would likely influence people's decisions is the perceived badness of extinction. Several well-documented psychological effects imply an underestimation of the value lost by extinction. For instance, scope insensitivity (i.e., not valuing a good in proportion with its scope or size; Desvousges et al. 1993) and narrow bracketing (i.e., viewing one choice in isolation rather than as part of a wider set of choices; Read et al. 2000) may cause people to underestimate the cost of being precluded from a larger number of choices or experiences in the future. Likewise, opportunity cost neglect (Frederick et al., 2009) may cause people to underestimate the opportunity cost of not being able to do what they could have done had they avoided an extinction event.

These factors would influence decision making even if people had accurate, unbiased information about the probability of the extinction event. These factors are likely amplified by biases in the available samples, in particular survivorship or observer selection effects (Ćirković et al., 2010; Wilson, 2023): conditional on thinking about the probability of the extinction event, the decision maker has not been subjected to one themselves, and will therefore have experienced zero instances of it (e.g., someone who is reflecting on the risk of jaywalking will have zero personal experience with deadly accidents, as otherwise they would not be in a position to ponder this question). The impossibility of sampling extinction events from memory will likely lead to underestimation of extinction risk.

In summary, several considerations point toward risk aversion, such as the overweighting of small probabilities from descriptions or of extreme events, while others suggest risk-seeking tendencies, including the description–experience gap, underestimation of extinction's utility loss, and observer selection effects. Given the conflicting predictions that can be drawn from past research and the practical importance of the topic, it is surprising that there is a "notable scarcity of research on human behaviour in the context of low-probability high-impact events" (Sundh, 2024, p.7). One of the main reasons for this absence may be the lack of a behavioural paradigm to study these types of risks, especially one that can be connected with normative benchmarks.

The present work contributes to the understanding of how people make choices under extinction risk. After providing the necessary clarification on what is meant by an *extinction event*, and offering two distinct definitions, we introduce an experimental paradigm suitable for its study—the *Extinction Gambling Task* (EGT). The EGT, like most decision-making paradigms, isolates key features of decisions under extinction-risk within a lab-based framework involving financial gains and losses (Frey et al., 2017; Kahneman & Tversky, 1979). This approach not only enables a controlled first exploration of such decisions but also supports the derivation of optimal choice strategies for different variants of extinction events, which can serve as performance benchmarks.

After introducing the EGT, we report three experiments demonstrating its utility in examining people's choices under extinction risk. Experiment 1 shows that participants are sensitive to differences between two different kinds of extinction events in ways that qualitatively align with optimal strategies. Experiment 2 introduces a new type of catastrophic non-extinction event and shows that participants treat this type of event differently from the extinction events, again consistent with optimal predictions. Finally, Experiments 3a and 3b demonstrate how the EGT can be used to study factors that might plausibly influence people's decision making under extinction risk.

1. Towards an Extinction Gambling Task (EGT)

It is typical for scientific investigations to call for regimentations of ordinary-language concepts through the development of technical definitions (e.g., going from *warmth* to *temperature*). The case of human decision making is no different. What is perhaps distinct is the co-existence of multiple possible definitions. Take the case of 'risk', which can be defined as tracking the probability of the worst-possible outcome, or the variance among possible outcomes with known probabilities, to name just two possible definitions (see Levy, 2015, Chap. 1; see also Konovalova & Pachur, 2021). The same issue applies to 'extinction events', in part due to the wide range of domains to which the term could be applied (e.g., health, finance, climate). In the present work, we consider cases of repeated choice. Participants are required to repeat the same decision (between an option containing the possibility of extinction and one without this possibility) a maximum of 100 times. Within these scenarios, there are two possible types of extinction events. The first we define as *Keep*, which establishes extinction as an end to the possibility of gains while keeping any gains made thus far. As a real-world example, consider a boxer who risks career-ending injuries every time they take on a high quality opponent for a high value payday. As devastating as career-ending injuries are, they do not undo the benefits gained from such bouts in the past. Our second definition, *Lose*, is much more drastic, referring to a complete wipeout of gains—past, present, and future. As an example, consider the case of a bank whose investment losses are greater than all of the profits generated in their entire history, which results in them going out of business. These definitions of extinction risk also apply to extinction risks involving life, depending on how one views death. While Alice and Bob might agree that death is an end, Alice might consider that this, nevertheless, does not erase all the good experiences accumulated throughout life (Keep). Bob, however, might consider that those experiences are also forgotten and therefore erased (Lose). Alice's view is reminiscent of philosophical positions that underscore the experiences we have collected as fundamental to the meaning of life (Eagleton, 2008; Mitsis, 2020; Seneca et al., 2004), whereas Bob's view would be reflected in positions such as existential nihilism, which emphasise the futility of human endeavours in light of death (Kuhn, 1992; Sartre, 1972).

These two definitions of extinction mirror debates in the literature on the value of a statistical life and mortality-risk valuation (Kniesner & Viscusi, 2019). Some researchers have argued that the value of statistical life should vary with age – an idea sometimes morbidly labelled the "senior death discount" – on the grounds that older individuals have fewer remaining life years (Harris, 1985; Krupnick, 2007; Lockwood, 1988; Williams, 1997). However, empirical work on whether people's preferences align with this perspective shows inconsistent findings: stated preference studies show mixed results (a review by Krupnick (2007) found that 14 of 26 studies supported a "senior discount", while the remaining 12 found no effect or even a "senior premium"); in contrast, revealed preference studies suggest an inverse U-shaped pattern, where the lives of both younger adults and seniors are valued lower than the lives of middle-aged adults, likely reflecting confounds such as income (Aldy & Viscusi, 2007; Kniesner & Viscusi, 2019). In policy applications, using a senior discount is highly contentious and usually a uniform value of life is applied independent of age (Kniesner & Viscusi, 2019; Krupnick, 2007). Retaining both the Lose definition of extinction (which reflects roughly similar costs of extinction independent of age) and the Keep definition (which reflects reducing costs with increasing age) makes our task applicable to views that involve a senior discount and views that assume age-independent valuations of life.

Both definitions of extinction can be implemented in the context of a repeated risky-choice task. Consider a lottery problem with a 'risky' and a 'safe' option:

$$\texttt{Risky} = \begin{pmatrix} r_1 & r_2 & r_E \\ p_1 & p_2 & p_E \end{pmatrix} \qquad \texttt{Safe} = \begin{pmatrix} s_1 & s_2 \\ q_1 & q_2 \end{pmatrix}$$

Each lottery is comprised of mutually exclusive outcomes (e.g., r_1) that occur with a given probability (e.g., p_1). Note that $p_1 + p_2 + p_E = q_1 + q_2 = 1$. The outcome r_E corresponds to the extinction event. The consequences of this event will depend on the definition adopted.

To give a more concrete example:

$$\texttt{Risky} = \begin{pmatrix} £10 & £0 & \texttt{Extinct} \\ .475 & .475 & .05 \end{pmatrix} \qquad \texttt{Safe} = \begin{pmatrix} £1 & £0 \\ .50 & .50 \end{pmatrix}$$

If a decision maker only encounters a single lottery problem once, then there is no difference between the potential £0 outcome from the Safe option and the Extinct outcome from the Risky option. For the two outcomes to become distinguishable, it is necessary for the decision maker to engage with multiple lottery problems and to accumulate gains. In the present implementation of the EGT, they will encounter the same lottery problem multiple times. Knowledge of how many times a lottery problem will be encountered allows the decision maker to better evaluate the exposure to risk vis-à-vis the opportunity for maximising gains. This evaluation depends on the operating definition of extinction. In the next section, we outline the optimal strategies for both the Keep and Lose definitions of extinction.

1.1. Optimal solution for the keep definition

The Keep definition of extinction captures the idea of an end to the accumulation of gains, without any loss of past gains. In the EGT, where a lottery problem is encountered multiple times, this definition implies that the expected payoff of a strategy strongly depends on the position in the sequence of trials where the risky choices are played. Early in the experiment, extinction carries a higher (opportunity) cost than it does later in the experiment, as one misses out on more opportunities to earn money. By contrast, in the final trial, the opportunity cost is zero and the cost of drawing the extinction option is equal to the £0 outcome. Therefore, it would make sense to play the safe lottery in the first trials and the risky lottery towards the end.

Specifically, we show that the expected value is maximised by following a strategy with a single switch point, before which participants should choose the safe option and after which they should choose the risky option. The expected value ($\mathbb{E}$) for the Keep condition, assuming that a participant follows this optimal order of choices (that is, they first make N_{safe} safe choices and afterwards N_{risky} risky choices for a fixed total of $N_{\text{safe}} + N_{\text{risky}} = N_{\text{total}}$ choices), is given by

$$\begin{aligned}\mathbb{E}(N_{\text{risky}}) = &\underbrace{\bar{s} \times N_{\text{safe}}}_{\substack{\text{Expected value of}\\ \text{safe-choice trials}}}\\ &+ \underbrace{p_E \times \sum_{i=0}^{N_{\text{risky}}-1} (1-p_E)^i \times i \times \bar{r}}_{\substack{\text{Expected value of}\\ \text{the risky-choice trials}\\ \text{if one goes extinct at some point}}}\\ &+ \underbrace{(1-p_E)^{N_{\text{risky}}} \times N_{\text{risky}} \times \bar{r}}_{\substack{\text{Expected value of}\\ \text{the risky-choice trials}\\ \text{if one does not go extinct}}},\end{aligned} \tag{1}$$

where $\bar{s}$ denotes the expected value of choosing the safe lottery, and $\bar{r}$ denotes the expected value of choosing the risky lottery – extinction excluded – and p_E denotes the probability of going extinct when playing the risky lottery. If the safe choices are made first, then the expected value from the safe-choice trials is the number of safe-choice trials multiplied by the expected value per trial (first line of Eq. (1)). The second line denotes the probability of going extinct in the risky-choice trial multiplied by the payoff of this type of extinction (i.e., the expected value from the risky-choice trials over all extinction outcomes). Finally, the third line denotes the probability of surviving the entire experiment multiplied by the payoff of the risky-choices in this case (in the main text we derive all optimal solutions on the basis of expected value; the qualitative patterns of the optimal solutions are very similar if one instead uses expected utility, although with somewhat fewer risky choices, see Appendix B).

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

$$\underset{N_{\text{risky}}}{\operatorname{argmax}}\ \mathbb{E}(N_{\text{risky}}). \tag{2}$$

But what if participants do not follow the optimal ordering? If the risky-choice trials are not played strictly after the initial block of N_{safe} safe-choice trials, then the expected value for the latter (first term of Eq. (1)) changes. For example, consider the case where we replace two safe choices from this block with risky choices, namely the jth and kth trials, with $j < k < N_{\text{safe}}$. Then

$$\begin{aligned}\bar{s} \times N_{\text{safe}} \xrightarrow{\text{replaced with}} &\underbrace{\bar{s} \times (j-1)}_{\text{first } j-1 \text{ safe choices}}\\ &+ \underbrace{(\bar{s} \times (k-j) + \bar{r}) \times (1-p_E)}_{\text{risky choice at } j \text{ and safe choices between } j+1 \text{ and } k-1}\\ &+ \underbrace{(\bar{s} \times (N_{\text{safe}} - k - 1) + \bar{r}) \times (1-p_E)^2}_{\substack{\text{risky choice at } k\\ \text{and safe choices between } k+1 \text{ and } N_{\text{safe}}}}.\end{aligned} \tag{3}$$

Comparing Eqs. (1) and (3), we can see why the expected value is reduced whenever one deviates from the optimal order of playing all safe lotteries first and then all risky lotteries. The expected value for some of the safe choices is multiplied by the probability of survival (a number smaller than one, here $(1-p_E)$ and $(1-p_E)^2$).

Finally, consider the case that the decision maker is given some initial endowment. Then, the expected value for both choice options would simply change by a constant, namely the value of that endowment. This illustrates a more general fact: the optimal strategy for the Keep definition does not depend on initial earnings, as the decision maker can keep them regardless of their choice in the next trial.

1.2. Optimal solution for the lose definition

While it is relatively obvious that one should play risky in the final trial under the Keep definition, it is not so clear in the Lose definition, where one might lose all earnings gained on all previous trials. Indeed, under the Lose definition, the optimal solution is *dynamic* and has to account for participants' luck (i.e., their winnings from previous trials). For example, if a participant chooses the risky option in the first three trials, they may receive the maximum payoff three times, or they could receive zero payoff three times. In the first case, they have more to lose in subsequent risky choices than in the second case, and therefore they should play less risky in the following trials.

Deriving the dynamic optimal solution for this scenario 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.

Let $\mathcal{Z}$ denote current earnings. We can estimate the expected value for the total game following a risky choice with N *remaining trials* as

$$\begin{aligned}\mathbb{E}_{\text{risky}}(N,\mathcal{Z}) &= p_1 \times \mathbb{E}(N-1,\mathcal{Z}+r_1)\\ &\quad + p_2 \times \mathbb{E}(N-1,\mathcal{Z}+r_2),\end{aligned} \tag{4}$$

with

$$\mathbb{E}_{\text{risky}}(0,\mathcal{Z}) = \mathcal{Z}. \tag{5}$$

Analogously, we can estimate the expected value of the total game following a safe choice,

$$\begin{aligned}\mathbb{E}_{\text{safe}}(N,\mathcal{Z}) &= q_1 \times \mathbb{E}(N-1,\mathcal{Z}+s_1)\\ &\quad + q_2 \times \mathbb{E}(N-1,\mathcal{Z}+s_2).\end{aligned} \tag{6}$$

The optimal choice in a given trial is then risky if $\mathbb{E}_{\text{risky}}(N,\mathcal{Z}) > \mathbb{E}_{\text{safe}}(N,\mathcal{Z})$ and otherwise it is safe. Finally, the expected value of the whole game, when following the optimal dynamic strategy, is determined by the maximum of $\mathbb{E}_{\text{risky}}(N,\mathcal{Z})$ and $\mathbb{E}_{\text{safe}}(N,\mathcal{Z})$, given that an optimal agent would simply play whichever gamble has higher expected value,

$$\mathbb{E}(N,\mathcal{Z}) = \max(\mathbb{E}_{\text{risky}}(N,\mathcal{Z}), \mathbb{E}_{\text{safe}}(N,\mathcal{Z})). \tag{7}$$

Unlike the Keep variant, the Lose variant does not imply a single optimal switch point. Instead, there are different possible strategies that would approximate the optimal solution well. Following the optimal policy closely would imply a gradual decrease of the number of risky choices; however, strategies that start safe or risky and then switch to risky or safe, or a strategy playing mostly safe throughout the experiment would lead to similar expected value, as long as the total number of risky choices is similar. The reason being that, ultimately, the most important factor that determines the payoff in the Lose variant is the total number of risky choices played and, due to the large number of trials, this number is affected relatively little by luck. This is also reflected by the fact that a non-dynamic optimal strategy (i.e., a version of the optimal strategy were all choices are specified in advance and which does not allow taking the current endowment into account) leads to payoffs that are relatively close to payoffs from the dynamic solution (see Appendix C).

2. Experiment 1: Empirical choice patterns for the lose and keep conditions with 100 trials

In the previous sections, we (1) argued that the extant literature makes conflicting predictions regarding whether participants would be risk-seeking or risk-averse when deciding whether to engage in actions that risk extinction, and (2) derived optimal solutions for the EGT. We now report an implementation of the EGT that operationalises the Lose and Keep definitions of extinction and compare participants' choices against the optimal strategy.

2.1. Method

2.1.1. Participants & power

Participants were paid £1.50 for a 10 min study. On top of that base payment, participants in the Lose condition earned an average bonus of £0.50, interquartile range (IQR) [£0, £0.77], and in the Keep condition £1.01, IQR [£0.39, £1.39]. The study was approved by the ethics chair of UCL's Department of Experimental Psychology (EP/2021/005). All experiments in this paper were hosted online on a JATOS (Lange et al., 2015) server, and participants for all experiments were recruited via Prolific, which has been shown to have high data quality compared to other crowdsourcing vendors (Douglas et al., 2023; Eyal et al., 2021; Stagnaro et al., 2024). Initially, 196 participants signed up for the study. We excluded participants who failed one or more of four different comprehension and attention checks: two participants indicated the wrong number of trials, four participants did not indicate correctly that the piggy bank in the top right corner of the screen showed their current bonus earnings, four participants indicated that they did not remember the possible payoffs in the experiment, and 33 participants indicated that they did not understand the nature of the extinction event. Some participants failed multiple of these 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. A simulation-based power analysis showed close to 100% power with this design to detect a condition difference of 1 on the logit scale, equivalent to a ≈17% difference[1] in the proportion of risky choices between conditions (see Appendix C for details on the power analysis, which further demonstrates the robustness of our analytical approach to the selection effects introduced by the extinction event).

2.1.2. Design and materials

The experiment was programmed in lab.js (Henninger et al., 2021) and exported as a JATOS file for hosting. The full experiment is shared on the OSF (https://osf.io/qhkw6/). Participants played a sequence of 20 practice trials, followed by 100 incentivised

[1] In estimated average riskiness, the main quantity of interest introduced in more detail below. For comparison, the difference in estimated average riskiness was 35% in Experiment 1.

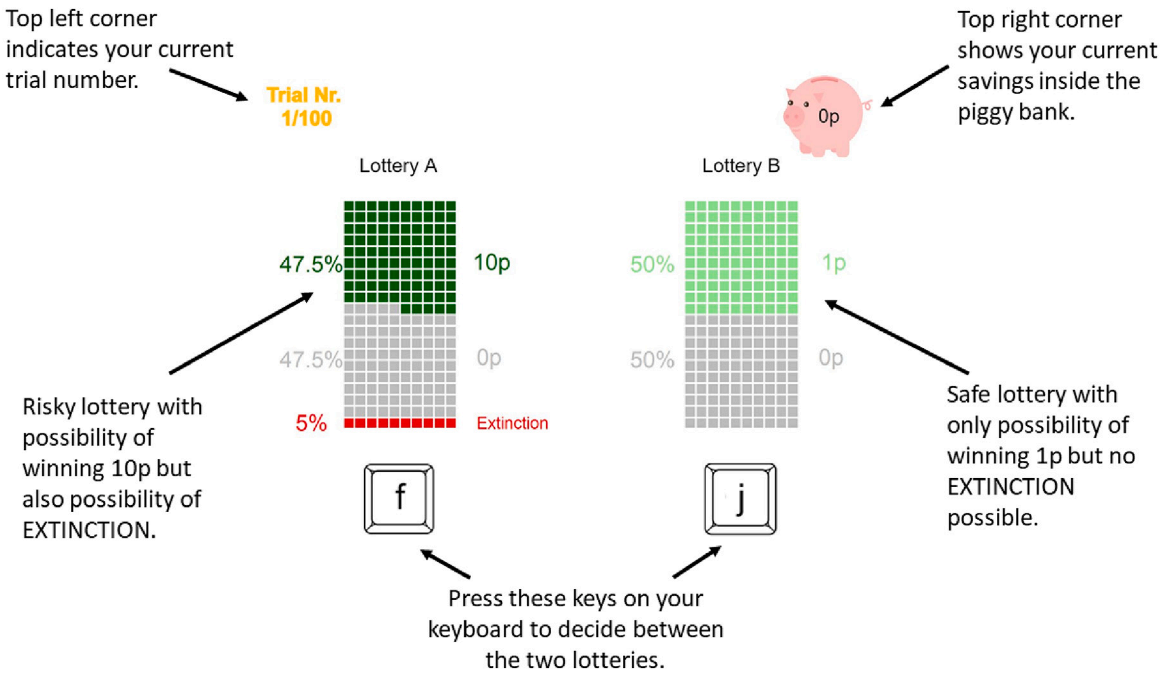

Fig. 1. The experimental setup as shown to participants on the instruction screen.

trials. Throughout the experiment, participants could see how many trials they had left and how much money they had currently accumulated. On each trial, participants had to choose between two lotteries: 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 (see Fig. 1 for an example trial). The task parameters were chosen by simulating from the optimal strategy and selecting values where optimality implies playing risky a few times but less than half of the time in both conditions (Lose ≈ 10%, Keep 44%). These parameters reflect that, although selecting the risky option can occasionally be advantageous, regularly taking extinction-level risks – especially in the Lose condition – is generally unfavourable (the values can easily be adjusted in future studies, and we modify the payoffs in Experiment 3). In Experiment 1, the risky lottery was always displayed on the left, and the safe lottery on the right (Experiment 3b replicated the findings of Experiment 1, where lottery position was counterbalanced between participants).

Our two experimental conditions represented the two extinction definitions outlined in the Introduction. In the Keep condition, participants could keep what they had earned if they experienced the extinction event but could not earn additional money in the following trials. In the Lose condition, participants' entire earnings would be wiped out when experiencing the extinction event and they also could not earn anything in future trials. After encountering an extinction event, participants still needed to play the remaining trials (to avoid any incentives for early extinction); however, the piggy bank in the top right corner turned into a mushroom cloud to indicate extinction (we did not analyse choices made after the extinction event was experienced). At the end of the experiment, participants received their earnings as a bonus payment.

2.1.3. Procedure

Participants started the experiment on a welcome page with a high-level description of the experiment and the following two warnings: 'We advise against participation if you have or had problems with gambling (i.e., gambling addiction) at any point in your life', and 'At one point during the experiment, you may view a picture of a mushroom cloud. If you are not comfortable with this, please do not continue'. After agreeing to the consent forms, participants received detailed task instructions. These instructions included the probabilities and outcomes of all the gambles, what would happen if they drew the extinction event, and an annotated illustration of the lottery screen (see Fig. 1).

After reading all instructions, participants made decisions in 20 practice trials. Some participants drew an extinction event during these practice trials. In this and the following experiments only a minority of participants experienced the extinction event during practice (Experiment 1: 43 out of 157, Experiment 2: 63 out of 173, Experiment 3a 46 out of 147 participants, Experiment 3b: 116 out of 305 participants), and we found no difference in the proportion of risky choices based on whether or not participants drew the extinction event during the practice trials (Experiment 1: $\chi^2(1) = 0.31, p = .575$; Experiment 2: $\chi^2(1) = 0.78, p = .376$; Experiment 3a: $\chi^2(1) = 0.96, p = .328$; Experiment 3b: $\chi^2(1) = 0.84, p = .361$).

After the practice trials, participants answered two attention checks: 'What happens if you choose the risky lottery and you draw the extinction option?' (correct answer for the Keep condition: 'I will keep all my past bonus money, but I cannot make more bonus money in the next trials', correct answer for the Lose condition: 'I will lose all my bonus money, and I cannot get any more

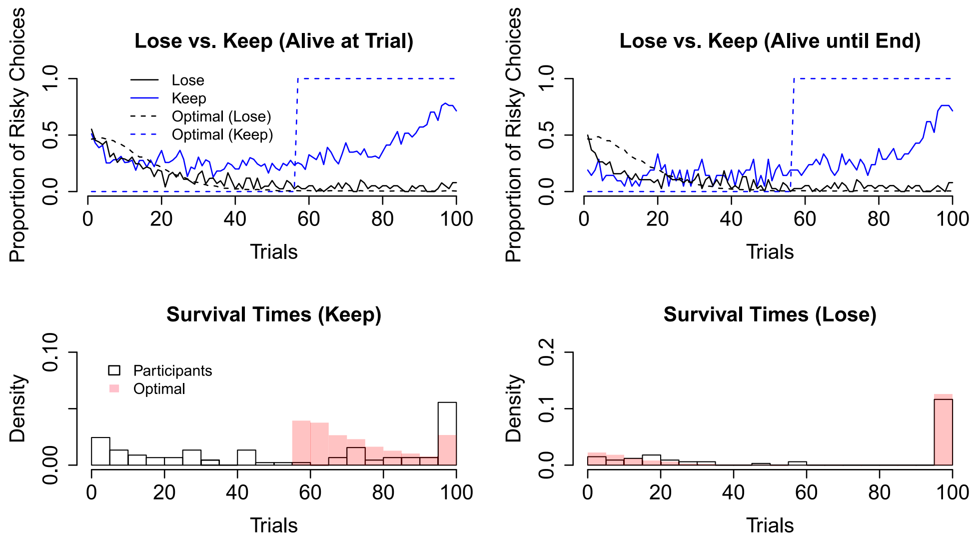

Fig. 2. Risky choices and survival times in Experiment 1. The optimal solution for the Lose condition is based on simulating 5000 participants who follow the dynamic solution (Eqs. (4)–(7)). The decision is executed based on a softmax decision function, as otherwise tiny differences in expected value in the first choices would lead to deterministic switching between the safe and risky options. We use a very low temperature (0.02) so that the solution is very similar to a deterministic decision rule. The same qualitative pattern arises for a range of temperature values. For a version of this figure that shows model predictions from the fitted mixed effects model along with 95% confidence bands see https://osf.io/af8sr. See also Fig. E1 for a visualisation of the optimal choices in the Lose condition as a function of trial number and current endowment.

bonus payments for future trials'.), 'What does the piggy bank on the top right indicate'? (correct answer: 'Your current bonus'). Participants continued to the main part of the experiment independent of their answers to these questions; however, those who failed the attention checks were excluded from data analysis.

In the main task, participants chose between the lotteries 100 times. After all trials were completed, we asked participants two further attention checks before they could finish the experiment: 'How many trials were there in the main part of the experiment?' (Correct answer: '100'), 'What were the possible bonus payments per trial in this study?' (Correct answer: 'Extinction, 0, +1, +10'; the other incorrect choice options also included 'extinction' to avoid any uncertainties about whether 'extinction' counts as a bonus payment or not).

2.2. Results

Results of Experiment 1 are shown in Fig. 2. The optimal solution is represented in the top row of Fig. 2 as a dashed line. In the Keep condition (blue), the optimal strategy is to switch from choosing the safe lottery to choosing the risky lottery after trial 56. In the Lose condition (black), the aggregate optimal strategy is to start with more risky choices and then reduce the proportion of risky choices throughout. The top row of Fig. 2 suggests that participants' choices (solid lines) differed between the two extinction conditions in a manner qualitatively consistent with the difference in the optimal solutions. Specifically, risky choices increased across trials in the Keep condition, but decreased across trials in the Lose condition. The top row of Fig. 2 also suggests that choices were quantitatively closer to the optimal solution in the Lose condition than in the Keep condition.

The bottom row of Fig. 2 shows the survival times of participants (white bars, where 100 trials indicates that the participant survived until the end of the experiment), as well as the survival times expected for participants following the optimal strategy (red bars). This shows that more participants went extinct in the Keep compared to the Lose condition. Again, this indicates sensitivity to the core task characteristics, as extinction is more costly in the Lose condition. Further, in the Keep condition, some participants went extinct earlier than implied by the optimal strategy, whereas in the Lose condition, the survival times tracked the optimal expectations remarkably well (see also Fig. E1 for a visualisation of the optimal choice for the Lose condition as a function of trial number and money, and Fig. G1 for a by-participant summary of the data).

To statistically test people's sensitivity to the different extinction conditions, we applied a logistic mixed effects model which predicts choosing the risky option as a function of trial number (with both a linear and a quadratic effect, to account for a potential U-shape), and condition (Keep vs. Lose), as well as an interaction between trial number and condition. The optimal strategy implies a higher proportion of risky choices overall in the Keep condition. We perform this comparison across conditions using the *estimated average riskiness* (i.e., the predicted riskiness based on the model's fixed and random effects assuming each participant responded to all trials, which is then averaged per condition; this measure takes potential non-linearities into account and addresses the selective dropout of more risky participants). For more details on our approach, including a simulation study to verify its accuracy, see Appendix C as well as the online supplementary materials. We found a significant difference in the estimated average riskiness between conditions, $z = 6.01$, $p < .001$, with the estimated proportion of risky choices being 52%, 95% CI [43%, 62%] in the Keep condition, as opposed to 17%, 95% CI [10%, 24%] in the Lose condition. For comparison, if all participants had followed the optimal strategy, these proportions would be 44% and 10%.

Further, the optimal strategy implies increasing risky choices during the task for the Keep condition and decreasing risky choices for the Lose condition. In line with this, we found a significant interaction between trial number (linear effect) and condition, $\chi^2(1) = 54.66, p < .001$, with a positive marginal linear slope in the Keep condition $b = 3.68$, 95% CI $[1.29, 6.08]$, and a negative marginal linear slope in the Lose condition $b = -6.41$, 95% CI $[-8.02, -4.79]$. Overall, these results suggest that participants were sensitive to the differences between conditions in a way that was (qualitatively) consistent with the optimal strategy.[2]

2.3. Computational modelling of individual level strategies

The aggregate pattern of results in the previous section likely arose from a mixture of different strategies on the level of individual participants. For example, it is plausible that the steady increase in the probability of choosing risky in the Keep condition (top-row of Fig. 2) is actually the result of different participants switching from choosing safe to choosing risky at different time points. Specifically, in the EGT participants can: (1) play the risky option with a constant probability throughout; (2) gradually increase or decrease the probability of playing the safe or risky option; or (3) switch between playing safe and playing risky. These strategies could be played in either a deterministic or probabilistic way (e.g., switching from always safe to always risky vs. switching from playing safe 90% of the time to playing risky 90% of the time).

Modelling the individual-level strategies confers several benefits. First, some of these strategies are qualitatively closer to the optimal solution than others. For example, in the Keep condition, the optimal solution would imply a sudden switch from safe to risky (rather than a constant increase or a reduction), whereas switching from risky to safe, or playing only risky, would result in a lower expected payoff. Therefore, modelling individual-level strategies helps us understand to what extent participants recognised the optimal strategy and were only impeded by their ability to estimate the correct parameters (e.g., some participants may have recognised that a single switch strategy was optimal, but did not know where the optimal switch point was), or did not even qualitatively recognise the optimal strategy. Second, the optimal strategies differ between conditions in systematic ways. For example, in the Keep condition, it is best to switch from playing safe to playing risky, whereas this strategy is less strongly implied in the Lose condition. Therefore, modelling individual level strategies allows us to better understand differences between conditions and sensitivity to the type of extinction event. Third, we can use parameters from computational models to get a better understanding of specific elements of participants' behaviour. For example, comparing the model estimated switch point to the optimal switch point gives some evidence of whether participants are risk-averse or risk-seeking. In the next sections, we first outline the model specification and then illustrate the three benefits discussed here through application of the model to the experimental data.

2.4. Model specification

To describe the above strategies, we implemented a hierarchical (dependent) mixture model with three components: a *safe state*, in which participants would play mostly the safe choice; a *risky state*, in which participants would play mostly the risky choice; and a *gradual state*, which describes a gradual increase or decrease of the probability of playing risky throughout the experiment (implemented with logistic regression). The model allows participants to switch between the safe and risky states during the experiment to accommodate the different types of strategies outlined above (e.g., switching from safe to risky, or switching from risky to safe). However, participants cannot switch away from or into the gradual state. Further, we constrained the states so that the probability of playing risky when in the risky state is always > 0.8, and the probability of playing risky when in the safe state is always < 0.2 (for a detailed model specification, including information on prior distributions, see Appendix D).

2.5. Computational modelling results

The model showed good convergence (all $\hat{R} < 1.01$, Vehtari et al., 2021) and recovered the patterns in the data well (see Figs. F1 and F2 for posterior predictives). The computational modelling allowed us to identify six types of individual strategies, as visualised in Fig. 3:

1. **Gradual:** Participants who continuously increased or decreased their probability of playing risky during the experiment.
2. **Constant Risky:** Participants who mostly played the risky option, and their probability of choosing risky did not change during the experiment.
3. **Constant Safe:** Participants who mostly played the safe option, and their probability of choosing safe did not change during the experiment.
4. **Risky → Safe:** Participants who switched from playing mostly risky to playing mostly safe at a certain trial number.
5. **Safe → Risky:** Participants who switched from playing safe to playing mostly risky at a certain trial number.
6. **Many Switches:** Participants who switched between playing mostly safe and playing mostly risky multiple times during the experiment.

[2] There was also a significant interaction between trial number (quadratic effect) and condition, $\chi^2(1) = 10.48, p = 0.001$ (Keep: $b = 11.64$, 95% CI $[4.70, 18.59]$; Lose: $b = -0.365$, 95% CI $[-5.70, 4.97]$).

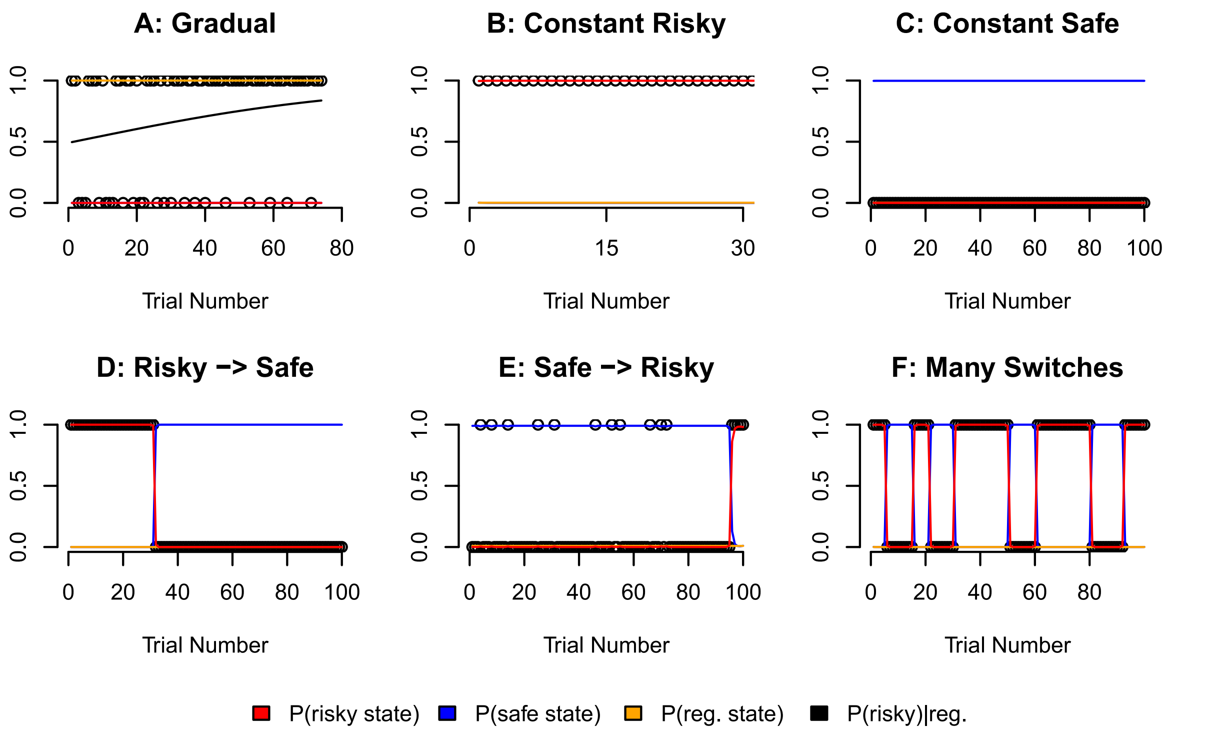

Fig. 3. Six different strategies identified by the (Dependent) mixture model. Each panel illustrates one participant who is representative of one of the six strategies. The y-axis denotes the probability of a given state shown in a coloured line (red, blue, and orange lines), the probability of choosing risky given the regression state (black line), and the actual choices as circles (1 denoting a risky choice and 0 denoting a safe choice). When one strategy has a posterior probability of almost 100%, the lines for the other two strategies will overlay at the bottom of the plot (i.e., only one is visible). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

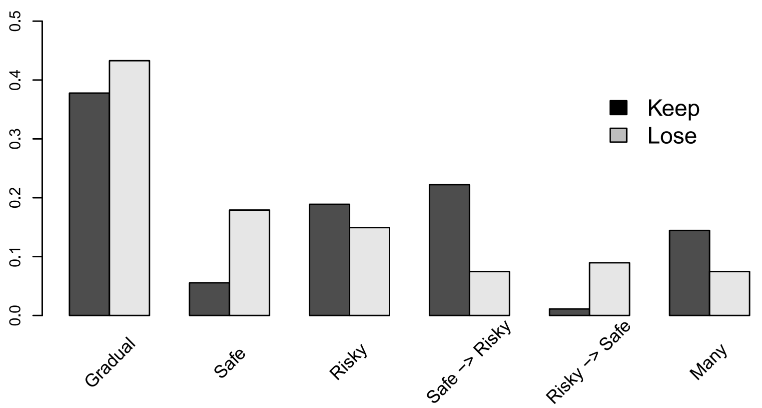

Fig. 4. Proportion of participants allocated to each of the six strategies in Experiment 1.

Fig. 4 shows that the proportion of strategies employed by different participants varied between conditions, $\chi^2 = 18.24, p = .002$.[3] This again indicates sensitivity to the core task characteristics and qualitative understanding of the optimal strategies in several ways. First, participants in the Keep condition switched from safe to risky more often than in the Lose condition, as implied by the optimal strategy. Second, participants in the Lose condition played the safe strategy considerably more often than participants in the Keep condition. This is in line with the optimal strategy: extinction is more costly in the Lose condition than in the Keep condition; consequently, in the Lose condition, playing fewer risky choices is optimal.

In both conditions, the majority of participants were assigned to the gradual strategy (Fig. 4), which was implemented with logistic regression. Based on the optimal strategies, we would expect the regression to have a positive slope in the Keep condition

[3] We use a simulation-based chi-square test as the chi-square approximation may be incorrect given the number of observations per cell. Therefore, no degrees of freedom are provided.

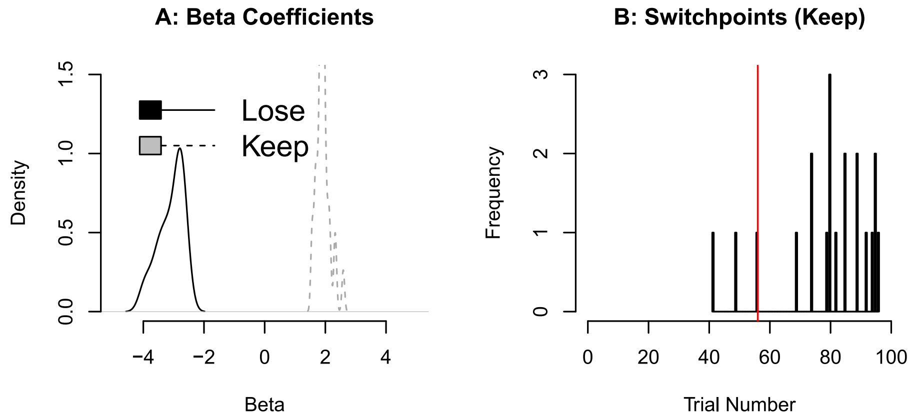

Fig. 5. Parameter estimates from the computational models in Experiment 1. Panel A shows beta coefficients for both conditions and Panel B shows switch points for the Keep condition. Each panel shows the distribution of posterior mean estimates across participants. Beta coefficients are only plotted for those participants who were assigned to the gradual strategy and switch points are only shown for those participants who were assigned to the Safe → Risky strategy.

and a negative slope in the Lose condition.[4] In line with this, all regression slopes were negative in the Lose condition, and positive in the Keep condition (Fig. 5, Panel A).

Finally, for the Keep condition, we compared the switch point estimated by the model to the optimal switch point to assess whether those participants who understood that the optimal strategy involved a single switch from safe to risky tended to switch too late or too early (Fig. 5, Panel B). This analysis indicated that most people switched later than the optimal switch point, $t(19) = 6.74$, $p < .001$, consistent with risk-aversion.

2.6. Discussion

Experiment 1 compared participants' behaviour to the optimal strategies in the Keep and Lose conditions of the EGT. We found that participants' aggregate behaviour was qualitatively in line with the optimal strategies, approximating the optimal solution in the Lose condition particularly well. Computational modelling of individual differences revealed a variety of different strategies employed by individual participants. Even though only a small proportion of participants followed the optimal strategy exactly (i.e., single switch or gradual), most participants used a strategy that at least qualitatively resembled the optimal approach (e.g., constant increase of p(risky) in the Keep condition).

3. Experiment 2: Introducing a reset condition

In Experiment 2, we introduce a new experimental condition—*Reset*. In this condition, participants lose all their earnings when the extinction event occurs, but they can continue playing and accumulate new gains and potentially draw the extinction event again. In practical terms, Reset is a mirror version of the Keep condition implemented in Experiment 1 (in which participants can keep all the money they have earned, but cannot make new gains after extinction).

The motivation for this new condition is twofold. First, it can be mapped onto real-world risk scenarios, where one loses everything but can recover (e.g., if one's house gets destroyed in a natural disaster). Second, it allows us to test the influence of opportunity cost neglect. Opportunity costs are foregone benefits that one would have incurred if one had chosen an alternative option. A key feature of extinction events is that they have high opportunity cost: all the benefits one could have incurred after the extinction event are now lost. Research shows that people tend to exhibit opportunity cost neglect; that is, they insufficiently take into account opportunity costs as compared to 'direct' costs (Frederick et al., 2009; Persson & Tinghög, 2020; Spiller, 2019).

Comparing the choice patterns in the Keep and Reset conditions of the EGT allows us to make inferences about how well participants can reason about opportunity costs versus 'direct' costs. In the Keep condition, a participant's endowment is not at stake, but their opportunity to keep playing in later trials is at stake (i.e., extinction leads to opportunity cost). Conversely, in the Reset condition, a participant's endowment is at stake, but their opportunity to continue playing in later trials is not at stake (i.e., extinction leads to 'direct' costs). In particular, we can compare the choices at the beginning of the task in the Keep condition (where extinction would lead to high opportunity cost but low direct cost) versus at the end of the task in the Reset condition (where extinction would lead to a high 'direct' cost but low opportunity cost). If participants display opportunity cost neglect, we would expect them to play too risky at the beginning of the task in the Keep condition, but we would *not* expect them to play too

[4] Although the optimal strategy predicts a negative slope for the Lose condition, this prediction is not as strict as the positive slope predicted in the Keep condition. As discussed above, the expected value may be very similar for other trajectories in the Lose condition as long as the total number of risky choices is similar to the one implied by the optimal strategy.

risky at the end of the task in the Reset condition. Consequently, we might expect more people in the Reset condition to follow the optimal single switch strategy than in the Keep condition (where some people would start too risky).

Overall, Experiment 2 has one main hypothesis and one research question.

1. **Sensitivity to Extinction Condition**: In the Keep condition, participants will increase the proportion of risky choices as they proceed through the experiment, whereas in the Reset condition, they will reduce the proportion of risky choices (preregistered).[5]
2. **Opportunity Cost**: Are participants in the Reset condition more in line with the optimal policy than participants in the Keep condition? (not preregistered)[6]

3.1. Optimal strategy for the reset condition

Because participants can lose all earnings accumulated so far in the Reset condition, the optimal strategy depends on how lucky participants have been and how many earnings they have won (as in the Lose condition). The optimal strategy for the Reset condition therefore also follows a dynamic solution. It is the same as in the Lose condition, with the addition that the optimal solution now also factors in the possibility of 'resets' (i.e., an event where one loses one's endowment but can continue playing). Specifically, the expected value of a risky choice with N remaining trials corresponds to

$$\begin{aligned}\mathbb{E}_{\text{risky}}(N,\mathcal{Z}) = {} & p_1 \times \mathbb{E}(N-1,\mathcal{Z}+r_1) \\ & + p_2 \times \mathbb{E}(N-1,\mathcal{Z}+r_2) \\ & + p_3 \times \mathbb{E}(N-1,0).\end{aligned} \tag{8}$$

Apart from this modification, the optimal solution follows the Lose condition—Eq. (6) determines the expected value of a safe choice, the value when zero trials are left equals the endowment (Eq. (5)), and the choice between safe and risky is determined by whichever of the two options has higher expected value (Eq. (7)).

Even though the Reset condition can conceptually be viewed as a mirror version of the Keep condition (as we argue in the introduction of this experiment), the optimal solution in the Reset condition implies a larger proportion of risky choices than in the Keep condition for the same maximum trial number, probabilities, and outcomes (67.6% vs. 44%). The reason for this is that risky trials in the Reset condition can be played both before and after an extinction event is experienced. Consider a participant in the Reset condition who starts risky and is reset to zero after 40 trials. This participant should then continue playing risky until the endowment becomes so large that it has a higher expected value to play safe rather than risky, which is for much longer than 4 more trials (44 trials is the optimum in the Keep condition).[7]

3.2. Participants

Participants were paid £1.20 for the 8 min study. On top of that base payment, participants in the Reset condition earned an average bonus of £1.22 IQR [£0.79, £1.57] and in the Keep condition £1.03 IQR [£0.40, £1.35]. The study was approved by the ethics chair for UCL's Department of Experimental Psychology (EP/2021/005). Initially, 197 participants completed the study. One participant indicated the wrong number of trials, six participants did not indicate correctly that the piggy bank in the top right corner of the screen showed their current bonus earnings, six participants did not remember the possible payoffs in the experiment correctly, and 20 participants indicated that they did not understand the nature of the extinction event (some participants failed multiple of these checks). After these exclusions, we obtained a final sample of 173 participants (77 in the Reset condition and 96 in the Keep condition). The mean age of participants was 40.60 years (SD = 12.15). Eighty-eight participants were female and 85 male.

3.3. Materials and procedure

The materials and procedure were the same as in Experiment 1, except that we replaced the Lose condition with the Reset condition. In addition, the position of the risky and safe choices on the screen was counterbalanced between participants.

3.4. Results

As in Experiment 1, the differences in participants' choices across conditions were qualitatively in line with the predictions of the optimal solutions (see Fig. 6). During the experiment, the proportion of risky choices increased in the Keep condition and decreased in the Reset condition.

---

[5] The preregistration is available at https://osf.io/8uas6. Note that in the preregistration, we use a previous terminology and refer to the Keep condition as the Keep-stop condition and the Reset condition as the Lose-continue condition.

[6] We initially preregistered that the main test for the role of opportunity cost neglect is whether there is a significant difference in the number of risky choices between both conditions. However, as this test was based on a static optimal policy for the Reset condition, this is no longer the most informative test (see section 'Optimal Strategy for the Reset Condition'). Following the preregistration with a test comparing the average riskiness between conditions, would come to the same conclusions.

[7] This difference in optimal strategy is also why the effect of opportunity cost neglect cannot be directly tested by comparing the proportion of risky choices, but instead requires computational modelling of the individual level strategies.

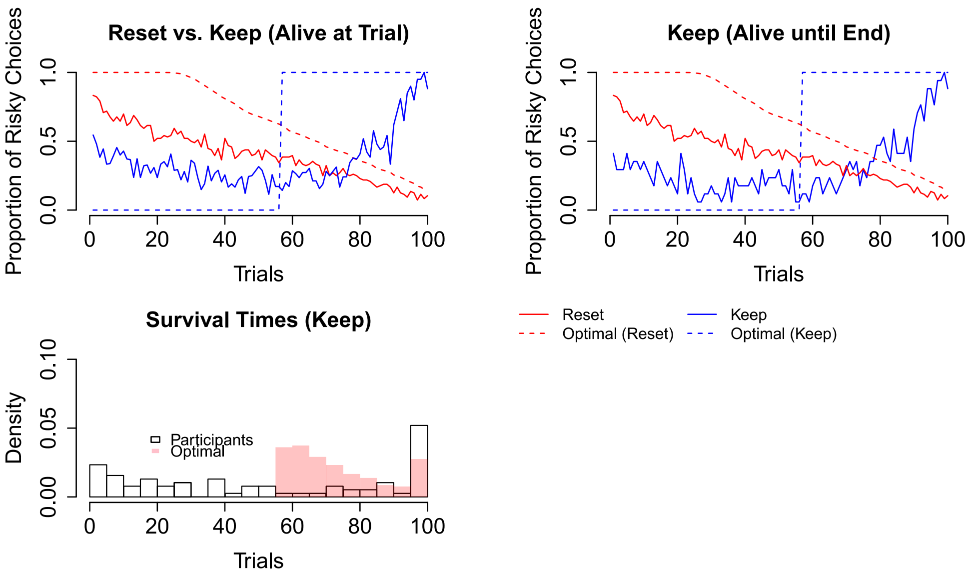

Fig. 6. Risky choices and survival times in Experiment 2. The optimal solution for the Reset condition is based on simulating 5000 participants, which follow the dynamic solution (Eqs. (5)–(8)). The decision is executed based on a softmax decision function, as otherwise tiny differences in expected value in the first choices would lead to deterministic switching between 0 and 1. We use a very low temperature (0.02) so that the solution is very similar to a deterministic decision rule. The same qualitative pattern arises for a range of temperature values. In the Reset condition, participants cannot drop out; therefore, the 'Alive at Trial' and 'Alive until End' lines are identical in this condition. For a version of this figure that shows model predictions from the fitted mixed effects model along with 95% confidence bands see https://osf.io/mrvg7. Also see Fig. E2 for a visualisation of the optimal strategy in the Reset condition as a function of current endowment and trial number.

We again use a mixed effects model with the same fixed effects specification as in Experiment 1 (effects of condition, linear and quadratic effect of trial number, as well as the interactions of condition with trial number).[8] Consistent with the predictions of the optimal strategies, we found a significant interaction between trial number (linear effect) and condition, $\chi^2(1) = 1372.28, p < .001$, with a positive marginal linear slope in the Keep condition, $b = 2.44$, 95% CI [2.11, 2.76], and a negative marginal linear slope in the Reset condition, $b = -4.46$, 95% CI [−4.68, −4.24]. We further found a just below threshold significant difference between the conditions in estimated average riskiness, $z = 1.98, p = .048$ (when using the maximal model, which showed convergence issues, this difference was no longer significant, $z = 1.29, p = .197$), with the estimated proportion of risky choices being 40%, 95% CI [35%, 46%] in the Reset condition and 49%, 95% CI [42%, 56%] in the Keep condition.

3.5. Computational modelling results

We applied the same computational model as in Experiment 1 to investigate participants' decision strategies on an individual level. The model again showed good convergence (all $\hat{R} < 1.03$). In the Keep condition, the optimal strategy still involves switching to playing risky after trial 56. In the Reset condition, the optimal strategy involves switching in the other direction (Risky → Safe); however, because the solution is dynamic, this switch does not always occur at the same trial. Instead, the timing of the optimal switch point varies between participants. If participants do not take opportunity cost into account, we would expect those in the Keep condition to start more risky than implied by the optimal strategy, and consequently fewer participants in this condition would follow the optimal strategy.

Participants' strategies differed considerably between conditions (see Fig. 7; $\chi^2 = 57.17, p < .001$). In line with the optimal solutions, participants were more likely to switch from risky to safe in the Reset condition, and more likely to switch from safe to risky in the Keep condition. In line with our hypothesis that participants find it easier to reason about 'direct' costs than opportunity costs, more participants followed the optimal strategy (switching from risky to safe) in the Reset condition than in the Keep condition (switching from safe to risky), $\chi^2(1) = 4.80, p = .028$.

A comparison of participants' switch points to optimal switch points suggests risk aversion in both conditions, as in Experiment 1. As shown in Fig. 8, those in the Reset condition switched from choosing risky to safe earlier than the optimal solution suggested, $t(38.85) = 3.81, p < .001$, whereas those in the Keep condition descriptively switched from safe to risky later than suggested, though

[8] The model also had by-participant random intercepts and initially by-participant random slopes for both the linear and the quadratic effect of trial number. As the maximal model, showed convergence issues (degenerate Hessian with 4 negative eigenvalues) and candidate models with random slopes for only the linear effect of trial number also did not converge without issues, we report results from a reduced model with only by-participant random intercepts. The maximal model showed the same pattern of statistical significant effects as the reduced model, apart from the main effect of condition as indicated in the main text. In addition to the effects reported in the main text, we also found a significant interaction effect between condition and trial number (quadratic effect), $\chi^2(1) = 157.39, p < .001$; Keep: $b = 9.87$, 95% CI [8.77, 10.97]); Reset: $b = -0.72$, 95% CI [−1.48, 0.04]).

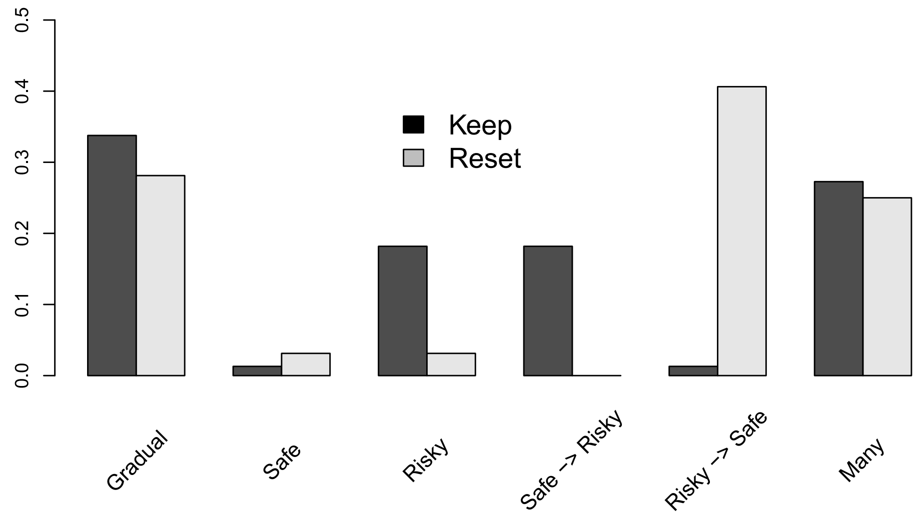

Fig. 7. Proportion of participants allocated to each of the six strategies in Experiment 2. More Participants Switch from Safe to Risky in the Reset Condition than from Risky to Safe in the Keep Condition.

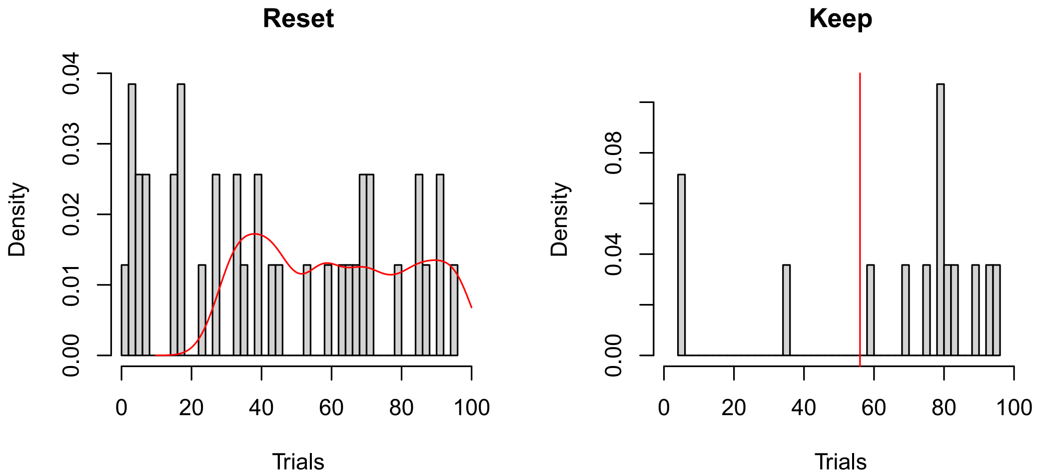

Fig. 8. Participants' switch points and optimal switch points for Experiment 2. Histogram of switch points for those participants that followed a single switch strategy. Red denotes the distribution of optimal switch points in Reset and the single optimal switch point in Keep. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

this effect is not statistically significant, $t(13) = 1.32, p = .208$ (this test is based on only 14 participants, three of whom switched very early). As in Experiment 1, the slopes of those participants assigned to the gradual strategy showed the expected direction (mostly negative in the Reset condition and mostly positive in the Keep condition, though there is more variability in the slopes in the reset condition, see Fig. G2).

3.6. Discussion

In Experiment 2, we introduced the Reset condition, in which participants lose everything if they draw the extinction event, but can continue to play and acquire future gains. As in Experiment 1, participants demonstrated sensitivity to the core task characteristics and we found some alignment between individual-level strategies and the optimal solution. Furthermore, we also found that more participants recognised the optimal single switch strategy in the Reset than in the Keep condition. This suggests that it is possibly easier for participants to reason about ‘real’ losses (i.e., their endowment, as represented by money in the piggy bank) than opportunity costs (i.e., not being able to continue playing and earning money).

4. Experiment 3: What psychological factors shape decisions under extinction risk?

In the first two experiments, we introduced three different definitions of extinction or catastrophic risks (Keep, Lose, & Reset) and tested the influence of opportunity cost neglect via comparison between the Keep and Reset conditions. Experiment 3 aims to further demonstrate the EGT's potential by varying features of the task *within* extinction conditions to test the influence of a variety of psychological factors on decision making under extinction risk. In particular, we explore the influence of loss chasing (Gainsbury et al., 2014; Lister et al., 2016), opportunity cost neglect (e.g., Frederick et al., 2009), and scope insensitivity (e.g., Desvousges et al., 1993) on decision making under extinction risk. Our aim is to demonstrate the utility of the EGT for studying the influence of such phenomena. Consequently, the following does not provide a definitive account of these phenomena, but rather a first demonstration of how they can be investigated using the EGT.

4.1. Experiment 3a: Endowment and losses

In Experiment 3a, we extended the Lose condition by introducing small losses and a starting endowment (Lose + Endowment condition). This new condition is identical to the Lose condition from Experiment 1, with the addition that participants now start with a small endowment of 50p and decide between a safe lottery that has a 50% chance of winning 1p and a 50% chance of *losing* 1p; and a risky lottery that has a 47.5% chance of winning 10p, a 47.5% chance of *losing* 1p, and a 5% chance of extinction. In other words, the 0p outcome from before is now replaced with a −1p outcome.

The Lose + Endowment condition has both practical and theoretical motivations. From a practical perspective, in the real world people usually do not start from zero but with something to lose. In addition, there is usually a possibility of small losses even under the safe option. For instance, someone who decides *not* to jaywalk may risk being late for a meeting. This would likely result in more risky choices because people often play more risky to recover from previous losses (a phenomenon termed 'loss chasing' in the literature on casino gambling; Gainsbury et al. 2014, Lister et al. 2016), especially when these losses have not yet been realised (i.e., paper losses or losses that have not yet been 'cashed out' in the eyes of the participant; Imas 2016). Losing money in the piggy bank in our task would likely constitute an unrealised loss as this money has not yet been awarded to the player.

This condition additionally allows us to probe in more detail why participants' proportion of risky choices declined during the task and, consequently, approximated the optimal solution in the Lose condition remarkably well in Experiment 1 (Fig. 4). One possible explanation for this decline is that participants play less riskily the higher their endowment (i.e., the more money is displayed in the piggy bank), similar to the idea of opportunity cost neglect investigated in Experiment 2. If this was indeed the case, we would expect fewer risky choices at the start of the Lose + Endowment condition, where there is a higher starting endowment, than in the normal Lose condition (note that this difference would not be affected by loss chasing, as participants have not yet experienced any losses at this point). The proportion of risky choices should, however, decline more slowly in the Lose + Endowment condition, as participants' endowment will not increase as quickly as in the normal Lose condition due to the −1p outcome.

Therefore, this study has two main research questions:

1. Is overall risk-seeking increased in the Lose + Endowment condition compared to the normal Lose condition? (not preregistered).
2. Does the proportion of risky choices decline faster in the normal Lose condition than in the Lose + Endowment condition? (preregistered)[9]

4.1.1. Participants

The study was approved by the ethics chair of UCL's Department of Experimental Psychology (EP/2021/005). Participants were paid £1.20 to participate in an 8 min study. On top of that base payment, participants in the normal Lose condition earned an average bonus of £0.44 IQR [£0.00, £0.60] and in the Lose + Endowment condition £0.41 IQR [£0.00, £0.63]. Initially, 210 participants signed up for the study. Three participants failed the attention check asking them about the total number of trials, nine failed the check asking about the possible outcomes of the lotteries, 25 did not know what the extinction event indicated, and six did not know what the piggy bank indicated. We also excluded three participants that seemed to have signed up twice, due to a technical problem. After excluding participants who failed these attention checks, we were left with a final sample of 147 (76 in the normal Lose condition and 71 in the Lose + Endowment condition). The average age was 41.21 (SD = 14.73), 72 participants were female and 75 were male.

4.1.2. Materials and procedure

The normal Lose condition was the same as in Experiment 1. For the Lose + Endowment condition, we updated the instructions and lottery screens to reflect the possibility of small losses, and the piggy bank on the first page of the experiment started with a value of 50. As in Experiment 2, we counterbalanced the position of the options on the screen.

4.1.3. Results

Fig. 9 (top row) suggests that the overall proportion of risky choices was higher in the Lose + Endowment condition than in the original Lose condition. Furthermore, as in Experiment 1, the proportion of people who survived until the end was similar to the optimal solution in the normal Lose condition; however, it was lower than the optimal solution in the Lose + Endowment condition (Fig. 9, bottom row), indicating greater risk-seeking in the Lose + Endowment condition.

[9] The preregistration is available at https://osf.io/8dntk. If the results of this test were to be statistically significant, we would have followed up by testing the difference at the first trial to delineate this from the influence of loss chasing.

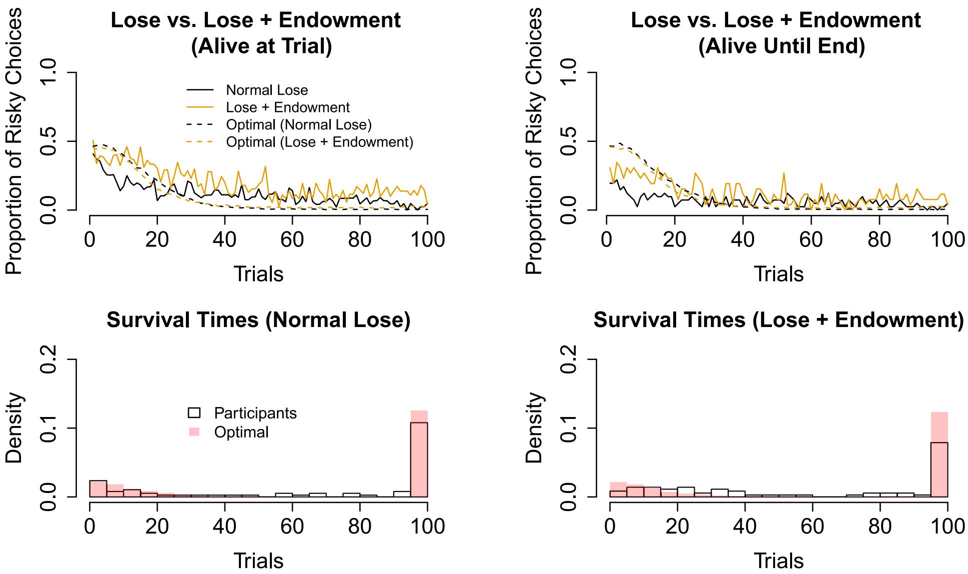

Fig. 9. Risky choices and survival times for Experiment 3a. The optimal solution is based on simulating 5000 participants who follow the dynamic solution (Eqs. (4)–(7)). The decision is executed based on a softmax decision function, as otherwise tiny differences in expected value in the first choices would lead to deterministic switching between 0 and 1. We use a very low temperature (0.02) so that the solution is very similar to a deterministic decision rule. The same qualitative pattern arises for a range of temperature values. For a version of this figure that shows model predictions from the fitted mixed effects model along with 95% confidence bands, see https://osf.io/g6ck3.

As in the previous experiments, we used a logistic mixed effects model with effects of condition, trial number (linear & quadratic effect), and their interaction. This indicated a significant difference in estimated average riskiness, $z = 2.74, p = .006$, with the estimated proportion of risky choices being 17%, 95% CI [11%, 23%] in the normal Lose condition and 28%, 95% CI [18%, 38%] in the Lose + Endowment condition (for comparison, the optimal would be 10% in both conditions). Further, we found no evidence for an interaction between trial number (linear effect) and condition, $\chi^2(1) = 2.20, p = 0.138$, with a negative marginal slope in both the normal Lose condition, $b = -3.30$, 95% CI [−4.78, −1.81], and the Lose + Endowment condition, $b = -1.88$, 95% CI [−3.61, −0.14].[10] These results are consistent with RQ1 (more risky choices after the introduction of Losses) but not RQ2 (faster decrease in risky choices in the normal Lose condition compared to the Lose + Endowment condition).

4.1.4. Computational modelling results

The model again showed good convergence (all $\hat{R} < 1.01$). Since the optimal strategy for both conditions was virtually the same in this experiment, we did not expect any major differences in terms of participants' assignment to the different strategies. In line with this, we found no evidence for differences in strategies between conditions, $\chi^2 = 3.50, p = .643$ (see Appendix G for a visualisation of the computational modelling results for this experiment).

4.1.5. Discussion

Overall, Experiment 3a provides evidence for risk-seeking after (paper) losses, as participants' responses were more risk-seeking in the condition including small (paper) losses (Lose + Endowment) than in the condition without such losses (Lose). Further, we did not find evidence that participants' risky choices are mainly determined by the amount of money in the endowment rather than the expected value of the total game. Instead, risky choices decreased during the experiment at a similar rate for both the Lose + Endowment and normal Lose conditions. This suggests that factors other than the endowment (e.g., exploration in the beginning of the task) cause the reduction in risky choices during the experiment.

4.2. Experiment 3b: Varying the maximum number of trials

In Experiment 3b, we turn to another important feature of our task: unlike most other gambling tasks, the maximum number of trials has considerable impact on the optimal strategies in the EGT. Intuitively, as the maximum number of trials decreases, less can be lost by drawing the extinction outcome, and consequently participants should choose the risky lottery more often. A similar line of reasoning also applies to real-world decisions under extinction risk: when more is at stake in terms of investment or future time lost from going extinct, it becomes more important to avoid extinction risk. However, research on choice bracketing (Read et al.,

[10] We also did not find a significant interaction between trial (quadratic effect) and condition ($\chi^2(1) = 0.04, p = .839$) with a positive slope estimate in both conditions (Lose: $b = 1.69$, 95% CI [−1.71, 5.09], Lose + Endowment: $b = 2.16$, 95% CI [−1.18, 5.51]).

2000) and scope insensitivity (Kahneman et al., 1999) suggests that people may not be sufficiently sensitive to the value that could be lost in case of extinction.

Choice bracketing refers to whether people, when faced with a sequence or set of choices, are more likely to integrate across them ('broad bracketing') or to consider each of the choices in isolation ('narrow bracketing'). People tend to view individual choices in isolation rather than integrating across the set; that is, they exhibit narrow bracketing (Bland, 2019; Rabin & Weizsäcker, 2009; Read et al., 2000). If narrow bracketing also pertains to the EGT, it would imply insensitivity to the total number of choices. Scope insensitivity describes a similar phenomenon whereby people do not value a good in proportion to its scope or size (Kahneman et al., 1999). It has been shown in a variety of domains, such as contingent valuation judgment (Baron & Greene, 1996; Desvousges et al., 1993; Kahneman & Knetsch, 1992), charitable giving (Maier et al., 2023; Västfjäll & Slovic, 2020), or in the reaction to mass suffering and genocides (Cameron & Payne, 2011; Dickert et al., 2015; Slovic & Västfjäll, 2010). Like narrow bracketing, scope insensitivity would suggest that people are not sufficiently sensitive to the amount of time that is lost by extinction (or the number of future choices), and would therefore likely lead to deviations from optimal decision making under extinction risk.

In Experiment 3b, we investigated how people adjust their choices as the maximum trial number changes. We decided on the maximum trial number in the different conditions based on three considerations. First, we chose maximum trial numbers that imply a sufficient difference in optimal strategies to maximise the chance of finding an effect of maximum trial number if it is present. Second, we tried to avoid extreme cases where the optimal solution implies always or never choosing the risky option. Third, we wanted to retain a maximum trial number of 100 in one of the two conditions to allow a direct replication of Experiment 1.

Based on these considerations, we chose 60 and 100 as our maximum trial numbers for comparison. As the maximum remaining trials should mostly have an effect in the Lose and Keep conditions, where earnings from a large number of future trials can be foregone upon extinction, this experiment focuses on those two conditions. This resulted in the following four groups, with the optimal number of risky choices indicated in brackets:

1. 60 maximum trials—Lose condition (optimal strategy ≈14/60 risky choices)
2. 100 maximum trials—Lose condition (optimal strategy ≈10/100 risky choices)
3. 60 maximum trials—Keep condition (optimal strategy 44/60 risky choices)
4. 100 maximum trials—Keep condition (optimal strategy 44/100 risky choices)

Even though the difference between 60 and 100 trials might appear small, the implied difference in the proportion of risky choices is substantial. In the Lose condition, participants should choose the risky option more than twice as often for 60 compared to 100 maximum trials (23% vs. 10% of the time), and in the Keep condition, a bit less than twice as often (73% vs. 44% of the time).

In addition to testing the effect of maximum trial number, Experiment 3b also aimed to replicate the sensitivity to the core task characteristics found in the previous experiments (in particular, Experiment 1, which compared the same conditions but only for 100 maximum trials).

Overall, Experiment 3b tested one set of predictions and one research question:

1. **Sensitivity to Extinction Condition:** We expected the same effects as observed in Experiment 1. The proportion of risky choices in the Keep conditions will be higher than in the Lose conditions. Additionally, we expect the same interaction between Extinction Condition and trial number: Participants in the Keep conditions will increase the proportion of risky choices across trials, whereas participants in the Lose conditions will not. (preregistered, see: https://osf.io/wesxa)
2. **Sensitivity to Maximum Number of Trials:** Are participants playing more risky when the maximum number of trials is lower? (not preregistered)[11]

4.2.1. Participants

Participants were paid £1.20 to participate in an 8 min study. On top of that base payment, participants in the Lose condition earned an average bonus of £0.52 IQR [£0.00, £082] and in the Keep condition £0.77 IQR [£0.22, £1.19]. The study was approved by the ethics chair of UCL's Department of Experimental Psychology (EP/2021/005). Initially, 393 participants completed the study. We excluded participants based on four different comprehension/attention checks. Four participants indicated the wrong number of trials, 14 participants did not indicate correctly that the piggy bank in the top right corner showed their current bonus earnings, 18 participants did not remember the possible payoffs in the experiment correctly, and 65 participants indicated that they did not understand the nature of the extinction event. Some participants failed multiple of these checks. After these exclusions, we obtained a final sample of 305 participants. The mean age of participants was 40.35 years (SD = 13.03). 158 participants were female and 147 male. 142 participants were in the Lose condition (85 in the 60-trial condition and 57 in the 100-trial condition) and 163 participants in the Keep condition (96 in the 60-trial condition and 67 in the 100-trial condition).

4.2.2. Design, materials and procedure

The method was the same as in Experiment 1 aside from the addition of the maximum trial number manipulation (60 vs. 100), and the counterbalancing of the position of the two lotteries (as in Experiments 2 & 3a). This resulted in a 2 × 2 (Extinction Condition × Maximum Trials) between-participants design.

[11] For this analysis, we deviate from the preregistration with an analysis that maps better onto our theoretical question. We report the preregistered analysis in a footnote in the Results section.

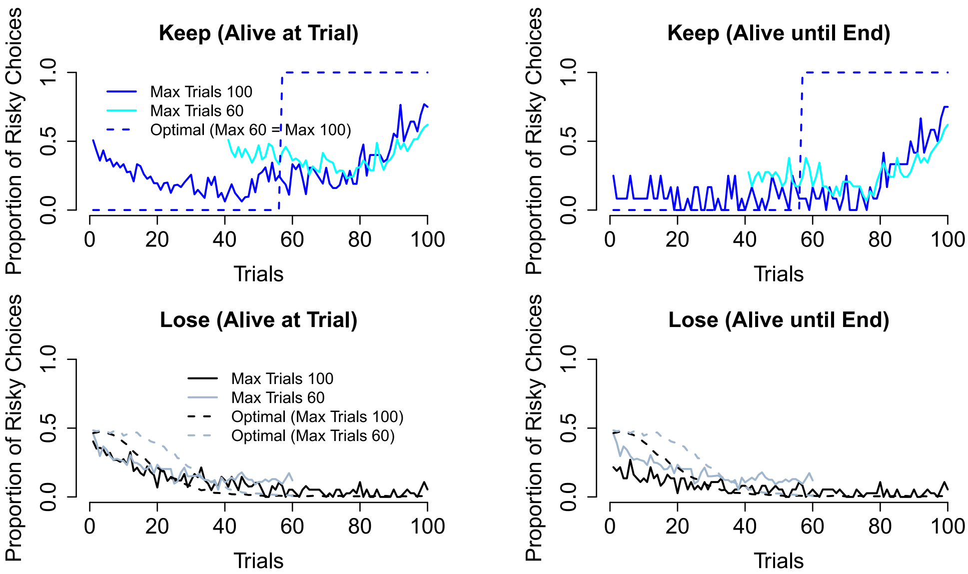

Fig. 10. Risky choices for each maximum trial number and extinction condition for Experiment 3b. In the Keep condition, the 60-trial condition is overlaid on the 100-trial condition in such a way that they both show the same end point. Therefore, the lines for the optimal solutions exactly overlay for 60 vs. 100 Maximum Trials. The optimal solution for the Lose condition is based on simulating 5000 participants, which follow the dynamic solution (Eqs. (4)–(7)). 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) so that the solution is very similar to a deterministic decision rule. The same qualitative pattern arises for a range of temperature values. For a version of this figure that shows model predictions from the fitted mixed effects model along with 95% confidence bands see https://osf.io/pz7hg.

4.2.3. Results

As in Experiment 1, participants' choices differed between the Keep and Lose conditions. In line with the optimal strategies, participants made more risky choices at the end of the experiment in the Keep condition, whereas they made more risky choices at the beginning of the experiment in the Lose condition (Fig. 10). Regarding the influence of Maximum Trials, we found that in the Keep conditions, participants increased their risky choices a similar number of trials from the end of the experiment, regardless of whether they would see a maximum of 60 or 100 trials (Fig. 10, top panels, in which the 60-trial condition has been plotted according to the distance from the final trial). In the Lose conditions, participants played a similar proportion of risky choices in the first 60 trials, again regardless of whether they would see a maximum of 60 or 100 trials (Fig. 10, bottom panels).

We first tested participants' sensitivity to the Extinction Condition using a logistic mixed effects model with effects of trial number (linear & quadratic effect), Extinction Condition, Maximum Trials, their two-way interactions, and their three-way interaction. We standardised the trial number by subtracting the midpoint within each condition.[12] As predicted, and as in Experiment 1, we found a significant difference in the estimated average riskiness, $z = 9.04, p < .001$, and a significant interaction between Extinction Condition and trial number (linear effect), $\chi^2(1) = 472.92, p < .001$, with a positive marginal slope in the Keep conditions, $b = 1.74$, 95% CI [1.47, 2.00], and a negative marginal slope in the Lose conditions, $b = -2.17$, 95% CI [−2.46, −1.89].[13]

The visual pattern from Fig. 10 suggests that in the Keep conditions, participants played the same proportion of risky choices with the same distance from the *end* of the experiment, whereas in the Lose conditions, participants played the same proportion of risky choices at the same distance from the *beginning* of the experiment. We therefore estimated two additional mixed effects models, one for each Extinction Condition, that compares the estimated average riskiness of the conditions with 60 maximum trials to the estimated average riskiness in the first 60 trials (Lose) and the last 60 trials (Keep) of the conditions with 100 maximum trials. Consistent with the visual impression from Fig. 10, this indicated no evidence for an effect of Maximum Trial Number in the Keep conditions, $z = 0.65, p = .518$, nor in the Lose conditions, $z = 1.51, p = .131$.[14]

[12] Specifically, for the 60-trial conditions: $trialNr_{stan} = (trialNr - 30.5)$ and the 100-trial conditions: $trialNr_{stan} = (trialNr - 50.5)$ (Note that the mean of trials ranging from 1 to 100 is 50.5 not 50).

[13] We further found a significant interaction between trial number (quadratic term) and condition $\chi^2(1) = 23.24$; Keep: $b = 6.25$, 95% CI [5.37, 7.13]; Lose: $b = 2.75$, 95% CI [1.78, 3.72]. These results were based on a random intercepts only model. The maximal model, which showed convergence issues (Hessians with negative Eigenvalues) lead to the same conclusions. Further, in this experiment, we preregistered only a linear effect of trial number. Removing the quadratic effect of trial number lead to the same conclusions.

[14] Testing for an effect of maximum trial number using the first GLMM (i.e., the one used to test for the main effect of Extinction Condition), shows evidence for an effect of Maximum Trial Number in the Lose condition, $z = 2.34, p = .020$, but not in the Keep condition, $z = 0.15, p = 0.883$. The reason is that this analysis does not compare matching trials but rather compares the estimated average riskiness across 60 trials in the 60-trial condition to the estimated average riskiness across 100 trials in the 100-trial condition.

4.2.4. Computational modelling results

The models again showed good convergence (all $\hat{R} < 1.01$). As in Experiment 1, we found a significant difference in strategies employed between Extinction Conditions, $\chi^2 = 33.31, p < .001$. More participants followed the risky strategy and the safe to risky switch strategy in the Keep than in the Lose conditions. We found no evidence for an impact of the Maximum Trial Number on the employed strategy (Lose: $\chi^2 = 7.39, p = .191$, Keep: $\chi^2 = 3.18, p = .682$; see Appendix G for visualisations of the strategies as a function of Extinction Condition and Maximum Trial Number).

4.2.5. Discussion

Experiment 3b replicated the findings of Experiment 1, where participants responded to the differences between the extinction conditions in ways that are qualitatively in line with the optimal strategies (i.e., they increased risky choices towards the end of the experiment in the Keep but not in the Lose conditions). Regarding the impact of the maximum trial number, our results suggest that in the Lose conditions, participants' choices depended on the number of trials already played (independent of the maximum trial number), while in the Keep conditions, their choices depended on the number of trials remaining (again independent of the maximum trial number). These results are consistent with the hypothesis that choice bracketing and scope insensitivity may lead people to be insensitive to the maximum trial number.

However, in addition to this insensitivity, participants' behaviour also showed a surprisingly adaptive element in the Keep conditions, which diminishes deviation from optimality: they switched at a constant distance from the last trial (rather than, for instance, after a constant proportion of trials). Therefore, insensitivity to trial number only leads to a deviation from optimality in the Lose conditions, where the number of risky choices changes as the number of trials reduces. Participants in this condition played somewhat too safely when the maximum trial number is reduced (since the optimal strategy implies more risky choices for fewer trials, red dotted line in Fig. 10). In the Keep conditions, insensitivity to trial number leads to well-calibrated behaviour, because people switch on average with the same distance from the end, a behaviour that is aligned with the optimal strategy (black line in Fig. 10).

5. General discussion

This paper aimed to advance our understanding of decision making in the face of irredeemable losses—namely, extinction risks. We proposed three definitions of extinction or catastrophic risk for study in a decision making task: (1) Lose, where all endowment is wiped out upon extinction and the participant cannot earn any payoffs for trials after the extinction event; (2) Keep, where the endowment is retained upon extinction, but the participant cannot earn any payoffs for trials after the extinction event; and (3) Reset, where the endowment is wiped out, but the participant can continue playing. We derived optimal strategies for these three definitions and operationalised them in a new experimental paradigm, the Extinction Gambling Task (EGT).

Leveraging the EGT and the computational models developed for it, we were able to obtain a nuanced understanding of people's choices in the task, which highlights both strengths and weaknesses in human decision making under extinction risk. On the positive side, the strategies employed by participants indicated a general qualitative understanding of the differences between extinction conditions. However, our studies also indicate several deficits: consistent with opportunity cost neglect (Frederick et al., 2009), we found that people are more in line with the optimal strategy and less risk-seeking when their endowment is at stake than when the opportunity to keep playing is at stake (Experiment 2); in line with research on loss chasing (Gainsbury et al., 2014; Lister et al., 2016), we found that introducing losses leads to excessive risk-taking (Experiment 3a); in line with research on scope insensitivity (Desvousges et al., 1993) and choice bracketing (Read et al., 2000), we found that participants in the Lose condition are not appropriately sensitive to the maximum trial number (Experiment 3b).

A pertinent question concerns how the results from the 'small world' (Savage, 1972, pp. 82–91) of our task relate to decision making about extinction risk in the real world. There are several real-world examples suggesting that people might be dealing with extinction risks rather poorly (e.g., Wiener, 2016). For instance, despite the recent pandemic and the possibility of another more deadly pandemic during our lifetimes (Marani et al., 2021), governments' investment in pandemic preparedness is relatively small compared to other areas of spending (Clark et al., 2022; Michael & Mark, 2024). Some of our findings are consistent with poor decision making about extinction risks. Specifically, people's choices in our task are affected by irrelevant factors (e.g., the introduction of losses and endowment in Experiment 3a), and sometimes participants are not sufficiently sensitive to other factors that should affect decision making (e.g., the total number of trials in the Lose condition of Experiment 3b). Some of the phenomena that we document have also been found in more realistic (albeit less controlled) settings (e.g., Coleman et al., 2023; Slovic & Weber, 2013). These observations suggest that the task may indeed be appropriate for isolating psychological factors that affect real-world decision making.

However, we do not only find limitations in human decision making. Participants were relatively good in terms of the qualitative strategies that they employed and they were sensitive to the differences between extinction conditions. To the extent that people make reasonably good decisions in our task but bad decisions in a corresponding real-world scenario, this would suggest shortcomings in real-world decisions other than an inability to understand the nature of the risk. In the next paragraphs, we outline several modifications of our task that may further the understanding of decision making under extinction risk,

5.1. Where next? The future of the extinction gambling task

The current experiments, and the current version of the EGT, provide a starting point for research into decision making under extinction risk. The task offers a flexible framework that can be extended to address a wide range of theoretical and applied questions in both individual and collective decision making. Here, we outline some directions for future research using the EGT.

5.1.1. Answering theoretical and practical questions in different risk-taking domains

Within the EGT, decisions under extinction risk are investigated via monetary gambles. This has a variety of advantages, in particular the ability to derive optimal strategies and compare participants' behaviour to them. However, our results should therefore be interpreted cautiously as speaking most directly to decisions under extinction risk within this domain. The observed choice patterns may be substantially different when examining different domains, types of gambles, or increasing the stakes (Blais & Weber, 2006; Camerer & Hogarth, 1999; Hertwig & Ortmann, 2001; Weber et al., 2002). The task could therefore be modified to investigate these differences. For instance, one could imagine a version with a realistic cover story related to pandemic response, where not locking down an area would correspond to the risky option, which can result in a large-scale outbreak or a global pandemic (corresponding to an extinction event), but in the more likely event that an outbreak does not happen, may lead to more economic activity and fewer people affected by control measures. Such modifications would allow bridging the gap between the current financial gambles and decision making under extinction risk in specific real-world domains. In general, decisions under extinction risk can be made increasingly more realistic by giving up internal validity for increased external validity: first, one could introduce realistic cover stories, while retaining the information about monetary outcomes (so it is still possible to calculate optimal strategies); second, one could create a purely hypothetical task with a realistic cover story but no monetary element; third, one could investigate risky-choices in the real world and see whether age-related changes reflect patterns of optimal strategies in our task.

Future work might additionally explore changes of the task structure, for instance by introducing multiplicative rather than additive payoffs (e.g., one's wealth grows by a constant percentage every trial if one does not go extinct), variable risk over trials, situations where an extinction event is possible for both lotteries, or allowing players to invest some money in order to reduce the risk. This creates further links to research in fields adjacent to psychology, particularly optimal foraging theory (Stephens & Krebs, 1986) and economics, and allows modelling behaviours such as saving for retirement (Pensions, 2024; Vanguard, 2021), saving in preparation for potential unexpected costs (Wang-Ly & Newell, 2024), and investment portfolio choice. Indeed, we can already see links between the present findings and commonly observed effects in these domains. For example, in investment portfolio choice, the finding that wealthier individuals typically invest in more risky assets (Guiso & Paiella, 2008; Peress, 2004) is reminiscent of the increase in riskiness during the task in the Keep condition (though important differences remain in terms of the implied optimal strategies and the observed behaviour, for instance the Keep condition has an optimal single switch strategy, which is not applicable in portfolio choice).

Further, the EGT could be modified to test to what extent people follow through on their planned behaviour and how this affects decision making about extinction risks. In the experiments reported here, participants engaged with the EGT dynamically (i.e., they experienced the outcomes of their choices in each trial before making the decision for the next trial). We argue that this dynamic implementation mimics the way people make most decisions under extinction risk in the real world. Consequently, the EGT should be prone to some of the same biases that operate in real-world decision making under extinction risk, in particular, survivorship or observer selection effects (e.g., when alive, one has never experienced a fatal event). Previous work has investigated decision making about extreme risks in static setups (i.e., all choices are specified in advance; Perfors and Van Dam 2018) or a combination of static and dynamic setups (Crosetto & Filippin, 2013). Moving forward, it would be interesting to leverage the dynamic aspect of the EGT to study discrepancies between planned and on-the-spot actions, also known as dynamic inconsistencies (e.g., Barkan & Busemeyer, 1999; Cubitt et al., 1998; Hotaling & Kellen, 2022).

Finally, all elements of the task were known to participants in our experiments (maximum trial number, payoffs and probabilities, and current endowment), as this information is required to calculate the optimal strategies. Future work could investigate what happens when information about these factors is reduced. Regarding the endowment, removing the display may reduce participants' ability to take their accumulated earnings into account, which could increase risk-taking, whereas not knowing the probabilities and outcomes could reduce risk-taking, similarly to ambiguity aversion (Ellsberg, 1961). The effect of not knowing the maximum trial number is more difficult to predict and would likely depend on people's prior beliefs.

5.1.2. Studying collective and social decisions

Many of the extinction risks that people are exposed to are collective rather than individual risks. For instance, when taking measures to mitigate climate change, if one country does not reduce its emissions but other countries do, this country could "free ride" and enjoy the benefits of mitigated climate change without the cost of reducing its own emissions. Public goods tasks that study cooperation in the context of preserving resources for the next generation (Hauser et al., 2014; Jacquet et al., 2013) find that, absent any interventions to improve cooperation, defection is usually too high to preserve the public good. In order to investigate collective decisions under extinction risk, we are currently working on adapting the task into a public goods game, where players play in groups of five. In this type of task, choosing the risky option constitutes a form of defection as the player who plays risky accumulates all the payoffs from the risky decision themselves, but if they draw the extinction option, the whole group will go extinct. By comparing risk-taking in the individual-level task to risk-taking in the collective task, we can quantify to what extent non-optimal decision making about extinction risk is driven by cooperation problems. Additionally, a collective task will allow us to test interventions to improve collective decision making, such as median voting (Hauser et al., 2014) or increased communication between participants.

Another way in which the EGT could be modified to account for social interactions is by incorporating competition between different players, which would be similar to some existing dice games, and has been shown to affect risky choices in previous work

(Schulze et al., 2015).[15] A competitive setup would introduce additional game-theoretic elements as the optimal strategy in this type of game also depends on what the other players would choose. For example, if the other players are very risky and likely to go extinct, one can win easily by being safe; on the other hand, if the other players follow a safer strategy, it becomes optimal to play more risky to increase variance and have a chance of achieving a score that is higher than each of the other players. These competitive dynamics may play out in real-world decisions, such as when multiple companies compete to develop a potentially dangerous technology. The competitive setup would allow us to test to what extent risk-taking is increased compared to the individual level game and allow testing interventions to reduce risk-taking (e.g., communication and contracts between the players).

5.1.3. Can the task be used to measure individual differences?

The EGT as implemented in this paper was not designed for measuring individual differences and we caution against directly using the task as individual difference measure, as the current one-shot setup with a single block and one extinction event likely has low reliability. Reliability could, however, be increased by chaining the current task multiple times in a row. The current duration of the main part of the task is around 6 min, so chaining the task would be feasible in terms of duration and participant engagement. However, creating a multi-block version would distort the nature of the extinction event, as it would not wipe out all earnings, but only the earnings for this block. It would also allow participants to learn from their experiences of extinction events. Both of these properties are in conflict with the definition of extinction events we outlined in the introduction. Future research investigating the potential differences between repeated and one-shot designs is therefore necessary before versions of the task that are suitable for measuring individual differences can be employed.

These extensions of the EGT go hand in hand with extensions to the modelling of the choice data on an individual level. The (dependent) mixture model employed here can be considered a descriptive or measurement model that allows us to classify the strategies based on the choice data, but the parameters of the model do not have a straightforward psychological interpretation. In future work, it would be interesting to consider process models, for instance, reinforcement learning models ((Watkins, 1989; Watkins & Dayan, 1992), especially so for modifications of the task without known probabilities and/or payoffs), or risky choice models (Palma et al., 2008; Krueger et al., 2024; Tversky & Kahneman, 1992). However, when applying risky choice models, additional questions arise, such as whether risk aversion should operate on a per-trial basis or over the whole set of choices in the task. To address these questions, one could compare responses and parameter estimates between the EGT and decisions in a risky-choice task not involving extinction risk.

5.2. Closing remarks

Until now, little psychological research has investigated decisions under extinction risk. This paper introduced the Extinction Gambling Task (EGT) as the first paradigm designed to study these decisions. While our findings offer valuable psychological insights, the primary aim of this work was to lay a robust foundation for future research in this area. Given the applied relevance of decisions under extinction risk, it is important to consider the usefulness of a laboratory task like the EGT. Some researchers have argued that much of decision-making research has had very limited real-world impact, owing to its emphasis on task designs tailored to testing formal models in artificial settings rather than examining truly "consequential decisions" (Weiss & Shanteau, 2021). While the study of consequential decisions is indeed vital, we contend that meaningful laboratory tasks can serve this goal, provided they capture features that are genuinely relevant to real-world decision contexts (for relevant discussions, see Fiedler, 2018; Garcia-Marques & Ferreira, 2018). The EGT has been developed primarily with this goal in mind. The computational modelling conducted here was designed to understand the nuances in the observed choice data, rather than for testing different formal models. This enables researchers to observe individual choice strategies in scenarios with meaningful features that isolate key properties of decisions under extinction risk, such as repeated choice opportunities, wealth accumulation, and the possibility of irredeemable outcomes. Given how these features map onto real-world scenarios, we are optimistic about the EGT's potential for facilitating scientific insight. We hope that our work and the development of the EGT will help to bring much-needed research attention to a significant domain that has been largely ignored.

CRediT authorship contribution statement

Maximilian Maier: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Adam J.L. Harris:** Writing – review & editing, Validation, Supervision, Conceptualization. **David Kellen:** Writing – review & editing, Validation, Conceptualization. **Henrik Singmann:** Writing – review & editing, Supervision, Software, Methodology, Conceptualization.

Appendix A. Existing paradigms and their limitations

Among the most popular paradigms (for a recent overview, see Pedroni et al., 2017), the Balloon Analogue Risk Task (BART) stands out as including something resembling an extinction event (Lejuez et al., 2003, 2002; Pleskac, 2008; Pleskac et al., 2008;

[15] For instance, in "Pass the Pigs", two players need to roll a set of two-pigs and are scored depending on which side the pig falls on. For a mathematical treatment, see https://www.youtube.com/watch?v=ULhRLGzoXQ0.

Wallsten et al., 2005). In the BART, participants' earnings increase each time they pump a balloon. Each additional pump, however, risks bursting the balloon and earning nothing. While the bursting of the balloon might be viewed as an extinction event, the BART does not capture how people reason about the types of extinction risks described above. This is because the losses are limited to that specific balloon trial, leaving the earnings from past – and future – balloons unaffected. Also, the BART imposes a strong relationship between choice and event probabilities: the more you pump, the greater the chances that the balloon will pop at each subsequent pump. However, in many real-life examples, the frequency of committing a risky act does not necessarily increase the risk of extinction with each subsequent action. That is, ceteris paribus, your 20th jaywalk is not more risky than your 1st.

The Bomb Risk Elicitation Task (BRET, Crosetto & Filippin, 2013) is another paradigm worth highlighting. In this task, participants decide how many boxes to collect out of 100. One of these boxes contains a bomb. Participants receive more earnings with each box they collect, but if they happen to collect the box containing the bomb, their earnings are reduced to zero. The BRET can be administered in a static and a dynamic version. In the static version, people specify the number of boxes to collect in advance. In the dynamic version, one box is collected every second until the participant presses the stop button. While the BRET has been shown to be a useful tool to study risky choices in general, it was not designed to apply to extinction risks and differs from decisions involving extinction risks in at least two respects. First, exactly one extinction outcome is among the set of options with certainty. This, again, results in a dependency between the number of boxes selected and the probability of finding the bomb, similar to the BART. Second, in the BRET, participants either choose the number of boxes they wish to collect in advance, or observe boxes being collected sequentially (without making an active decision on each instance) until they decide to hit the stop button. Both of these versions are quite different from real-life decisions about extinction risks, such as jaywalking, which involve repeated choices between different options.

Because of the growing interest in human action in the face of extreme risks (for theoretical papers see e.g., Slovic & Weber, 2013; Sundh, 2024; Weber, 2006), it is unsurprising that the need for experimental research has been recognised: Perfors and Van Dam (2018) proposed a new choice paradigm in which people make repeated choices under the risk of extreme *black swan* events that completely wipe out their accumulated earnings. Specifically, participants first executed a single choice and then (without feedback from their first choice) specified a policy according to which they would choose 2000 times in a row. The riskier lottery included a small-probability event that wiped out all previous earnings. However, these black swan events cannot be considered to be extinction events as defined in this paper (where losses are unrecoverable), given that participants were still able to accumulate future gains. Also, the initial commitment to a specific number of risky choices prevents any influence of experience and the investigation of how people frame and adjust their choices across trials (e.g., Bland, 2019; Rabin & Weizsäcker, 2009; Read et al., 2000). Further, Elga et al. (2024) discuss resource allocation in the face of extinction risks; however, their results mostly apply to scenarios where the risk is very high, and the number of repeated choices is low, which is different from the setup proposed in our paper and extinction risks in the real-world.

Appendix B. Expected value vs. Expected utility optimal strategies for the three conditions

See Fig. B1.

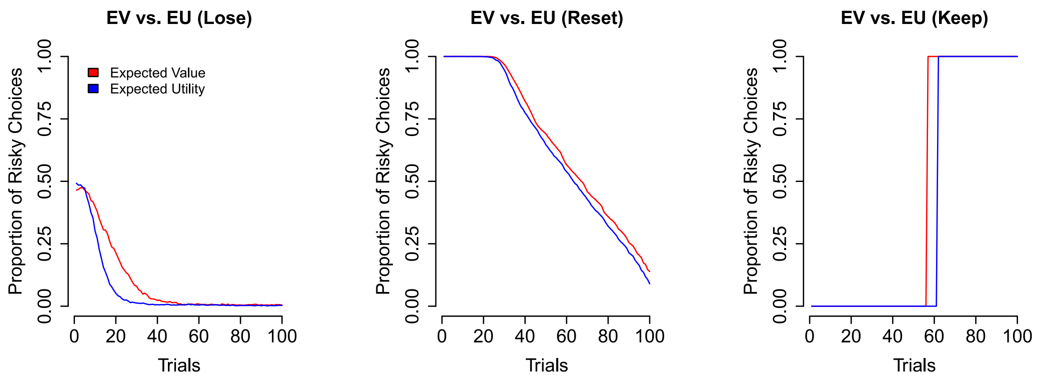

Fig. B1. Optimal strategies according to expected value vs. expected utility. Payoffs, probabilities and trial number are the same as in Experiments 1 and 2. We used an exponential function to model the diminishing returns with an exponent of 0.66 based on Kellen et al. (2016).

Appendix C. A priori (non-dynamic) optimal strategy for the lose condition

What if people do not take the current endowment into account to dynamically adjust the optimal strategy but instead commit to a single number of risky choices in advance, which they play? We can model this type of *a priori* expected value of different numbers of risky choices in the complete extinction case using Eq. (9),

$$\mathrm{EV}_{\text{a priori}}(N_{\text{risky}}) = (\underbrace{\bar{r}_{\text{t, safe}} \times (N_{\text{total}} - N_{\text{risky}})}_{\text{Expected value of safe trials if survival}} + \underbrace{\bar{r}_{\text{t, risky}} \times N_{\text{risky}}}_{\text{Expected value of risky trials if survival}} + \underbrace{\mathcal{Z}}_{\text{Endowment}}) \times \underbrace{p(s)^{N_{\text{risky}}}}_{\text{Total probability of survival}}. \tag{9}$$

$\bar{r}_{\text{t, safe}}$ denotes the expected value of choosing the safe lottery. $\bar{r}_{\text{t, risky}}$ denotes the expected value of chosing the risky lottery (assuming the player does not go extinct). N_{total} denotes the total number of trials of the experiment. N_{risky} denotes the number of risky trials being played. $p(s)$ denotes the probability of surviving (i.e., $1 -$ p(extinction)) when playing the risky gamble. Finally, $\mathcal{Z}$ denotes the endowment. When starting in the first trial with zero endowment this term is simply set to zero.

Based on this expected value the optimal a priori number of risky choices is simply

$$\underset{N_{\text{risky}}}{\operatorname{argmax}}\ \mathrm{EV}_{\text{a priori}}(N_{\text{risky}}). \tag{10}$$

Notably, the difference in terms of expected payoff between the dynamic and the a priori strategy is relatively small. For the payoffs and probabilities used in Experiment 1, the expected payoff following the dynamic strategy would be 57.87p, whereas the expected payoff following the a priori strategy would be 57.05p.

C.1. Estimating false positive rate and power

To estimate the statistical power of our design and to verify that our analysis would indeed accurately capture the data generating fixed effects between conditions in the presence of selection effects, we conducted a simulation-based power analysis for mixed effects models. While the R script shared on the OSF page can easily be adapted for a variety of settings, this appendix focuses on the setting of an interaction between trial number and extinction condition as well as a main effect of extinction condition (similar to Experiment 1). We simulated from a trial number slope, which starts from a probability of selecting the risky option of .5 in the first trial, declining to .05 during the experiment in Condition A, while it increases from 0.05 to 0.5 in Condition B. Further, we added a main effect of condition, whereby condition B is shifted by 1 on the logit scale (in comparison, the condition effect found in Experiment 1 was 1.12). We simulated from models without this main effect to assess the false positive rate and from models with this main effect to assess the statistical power. The random intercepts and slopes of trial number were 1 on the logit scale. We assume a sample size of 80 per condition. We also compared the mixed effects model approach to simply t-testing the difference in the proportion of risky choices between conditions. We repeated each condition 1000 times.

Table C1 summarises the results. The first two columns show the performance of both models when there is no selection. This indicates an almost nominal error rate, calibrated estimates and good power to detect an effect. The second two columns show what happens if we introduce selection effects (i.e., a chance of dropping out when choosing the risky option). The performance of the mixed effects models is similar, with slightly lower power due to the reduced number of trials. We can also see a slightly above nominal error rate, likely due to non-convergence of some models. In the manuscript, we always compare to a reduced, converged random effects structure in this case. However, merely t-testing the difference in proportions between conditions dramatically increases the error rate as the selective dropout distorts the simple proportion estimates (unlike the mixed model, which can account for this through modelling the trial number effect) and actually has a higher chance of claiming an effect when there is none than when one is present (due to the different starting points form the fixed slope and the condition effects cancelling each other out under selection).

Table C1
Power analysis for mixed model and t-testing proportions.

Selection	No		Yes	
True condition difference (logit)	0	1	0	1
Power mixed model	3.9%	>99.9%	6.4%	>99.9%
Power difference between proportions	4.1%	>99.9%	>99.9%	73%
Estimate mixed model (probability)	0.00*	0.16*	0.01*	0.18*
Estimate difference between proportions	0.00*	0.16*	0.22*	−.10*

Note. Estimates indicated with * are on linear probability scale rather than logodds scale.

Appendix D. Full model specification and implementation

D.1. Full model specification

Our model had the three following response models that describe the probability of choosing risky ($p(\mathrm{r})$) either in terms of a constant probability or as a function of the current trial number (t):

$$P(\mathrm{r}|\text{safe st.}) = \text{logit}(\alpha_{s,i})^{-1} \times 0.2 \quad \text{(Safe State)} \tag{11}$$

$$P(\mathrm{r}|\text{reg. st.}) = \text{logit}(\alpha_{reg,i} + \beta_{reg} \times N_{\text{trial}})^{-1} \quad \text{(Logistic Regression State)} \tag{12}$$

$$P(\mathrm{r}|\text{risky st.}) = \text{logit}(\alpha_{r,i})^{-1} \times 0.2 + 0.8 \quad \text{(Risky State)} \tag{13}$$

N_{trial} denotes the trial number. We used a hierarchical implementation for α_s, α_{reg}, α_r, and β_{reg} with random intercepts and random slopes. We did not directly model the covariation between intercepts and slopes, as this would increase the number of parameters that need to be estimated, considerably, impeding model convergence.

We then introduced a hidden Markov model component that allowed transitions between the safe and the risky states to model the switch points. However, as described above, we did not allow switches to and away from the regression state. The transition matrix below shows the structure of the model with the rows and columns being ordered as (1) safe, (2) regression, (3) risky and ω denoting the probability of staying in the safe or risky state once in it:

$$\Gamma_{i,j} = \begin{pmatrix} \omega & 0 & (1-\omega) \\ 0 & 1 & 0 \\ (1-\omega) & 0 & \omega \end{pmatrix} \tag{14}$$

Finally, the starting probability for each state is denoted by the uniform three-dimensional simplex ρ with a separate ρ estimated for each participant (i.e., no hierarchical structure for ρ).

D.1.1. Prior distributions

For the grand mean and standard deviation on all α and β we always use a normal and half-normal distribution:

$$\mu \sim \text{normal}(0,1) \tag{15}$$

$$\sigma \sim \text{normal}(0,1)_+ \tag{16}$$

For the transition probabilities simplex $[\omega, 1-\omega]$ we use

$$[\omega, 1-\omega] \sim \text{Dirichlet}(30,1), \tag{17}$$

which is equivalent to

$$\omega \sim \text{beta}(30,1), \tag{18}$$

and reflects the assumption that participants who are in a certain state (i.e., play mostly risky or mostly safe) are more likely to play risky or safe again in the next trial than to change the state. Finally, the initial state probability has the uniform prior

$$\rho \sim \text{Dirichlet}(1,1,1) \tag{19}$$

D.1.2. Implementation

We implemented the model in `cmdstan` (version, 2.34.1 cite) and fitted it using `cmdstanr` (Version 0.6.1, cite). We ran each modelith 1000–3000 warmup iterations and 2000–6000 sampling iterations (for Experiments 1 and 3 we ran 3000 iterations, of which 1000 were warmup iterations using adapt delta = 0.95. For Experiment 2 we run 9000 iterations with 3000 warmup iterations using adapt delta = 0.98). For Experiment 4, we run 9000 iterations, 3000 warmup with adapt delta .95 in the Lose condition and adapt delta .98 in the Keep condition). All $\hat{R}$ were smaller 1.03.[16]

[16] For Experiment 1 all Rhats were smaller 1.01, for Experiment 2 all Rhats were smaller 1.03, for Experiment 3a all Rhats were smaller 1.01, for Experiment 3b all Rhats were smaller than 1.02.

Appendix E. Heatmaps of optimal choice as a function of endowment and trial number

See Figs. E1 and E2.

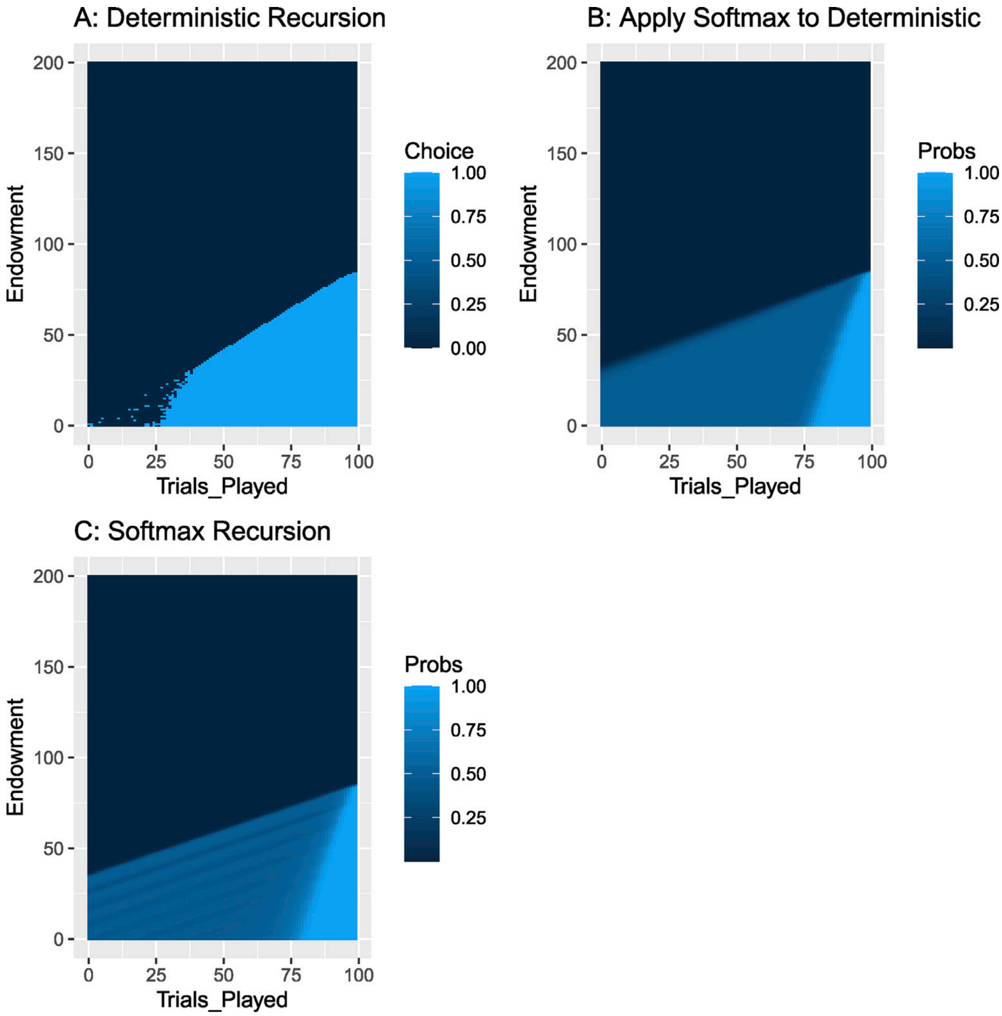

Fig. E1. Heatmaps visualising the optimal choice as a function of trial number and money for Experiment 1 - Lose Condition. Panel A shows a surprising pattern, where there is no clear boundary between playing safe and risky choices. Panels B and C help us understand the reason for this by using a softmax transformation to derive choice probabilities based on the expected value of the two options with a very low temperature of 0.02. Panel B applies a softmax transformation to the deterministic recursive function, while Panel C uses a recursive softmax function directly. Comparing both of these panels to Panel A suggests that in a large area of the parameter space, there is almost no difference between playing safe and playing risky in terms of the dynamic solution. So the pattern of Panel A is the result of only tiny differences in EV between the safe and risky option that can flip based on how sequences of choices starting from a certain tile make it easier to reach the safe vs. risky state.

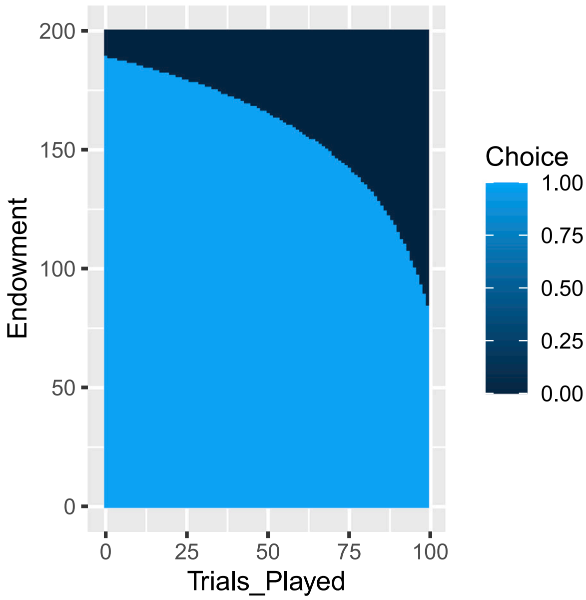

Fig. E2. Heatmap visualising the optimal choice as a function of trial number and money for Experiment 2 - reset condition.

Appendix F. Comparison of posterior predictions and data

See Figs. F1 and F2.

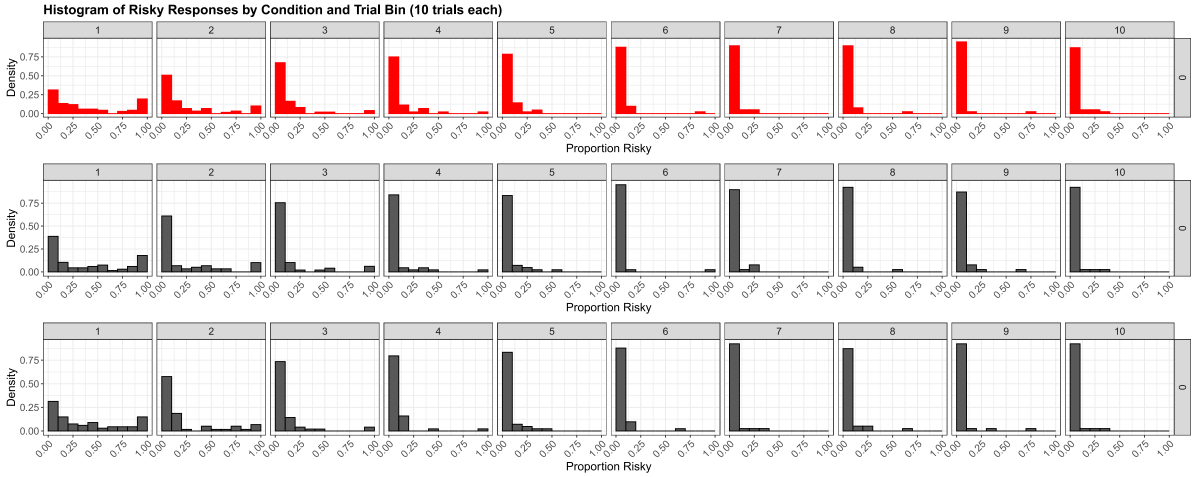

Fig. F1. Model predictions vs. data in the lose condition (Exp. 1). Histograms show the distribution of risky choices within ten trials. For example, panel "1" shows the distribution of average probabilities to make a risky choice across participants for trials 1 to 10. The top row indicates the data and the rows 2 and 3 below each show one random draw from the posterior predictive distribution.

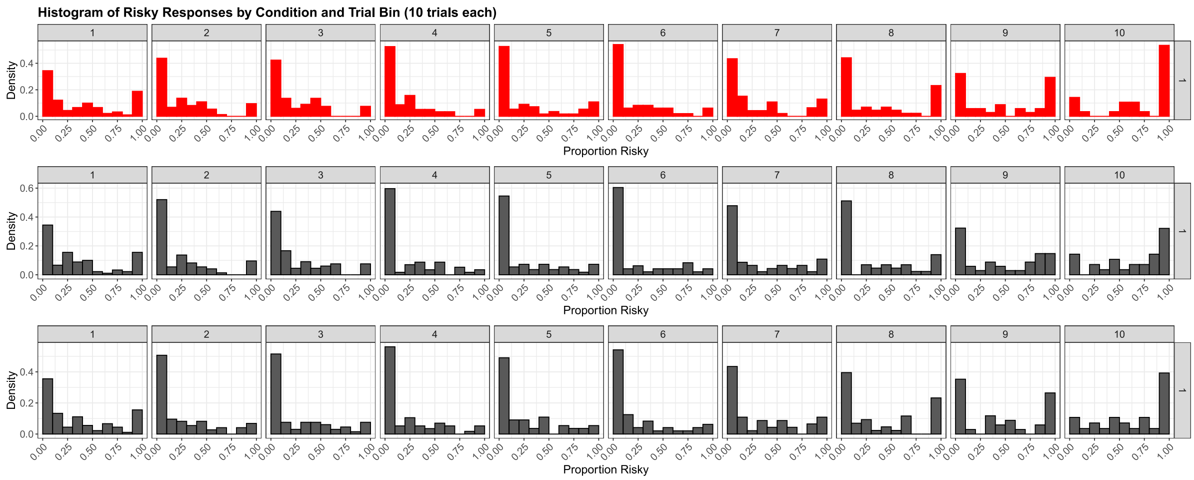

Fig. F2. Model predictions vs. data in the keep condition (Exp. 1). Histograms show the distribution of risky choices within ten trials. For example, panel "1" shows the distribution of average probabilities to make a risky choice across participants for trials 1 to 10. The top row indicates the data and the rows 2 and 3 below each show one random draw from the posterior predictive distribution.

Appendix G. Additional results

G.1. Visualisation of additional variables for Experiment 1

See Fig. G1.

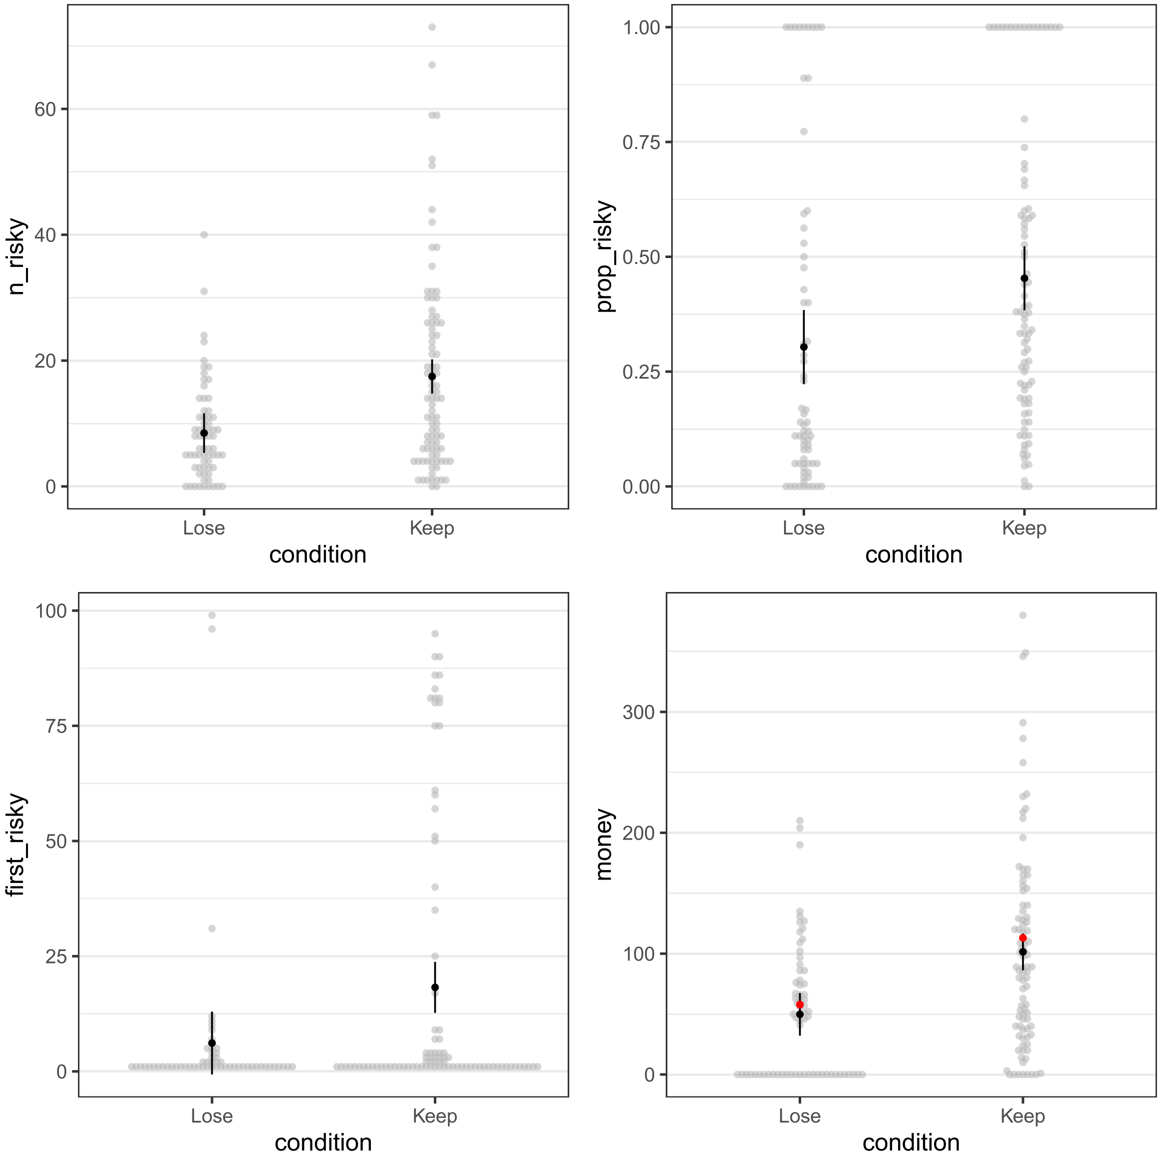

Fig. G1. Number of risky choices, proportion of risky choices, trial number of first risky choice, and money earned between conditions for Experiment 1. The red dots in the bottom right panel indicate how much money an agent following the optimal strategy would make on average. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

G.2. Beta coefficients for Experiment 2

See Fig. G2.

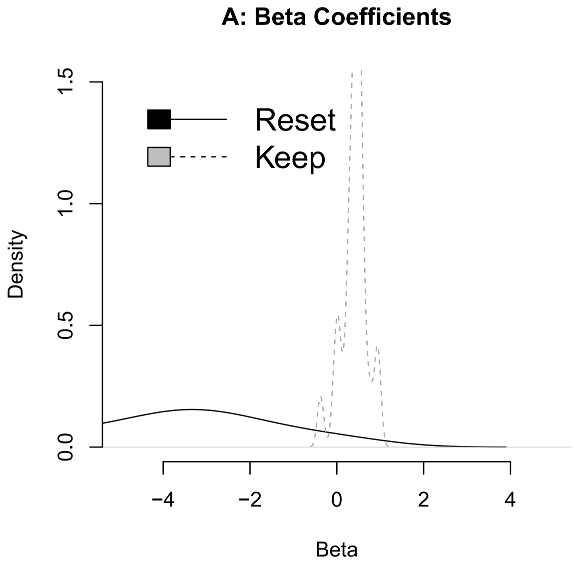

Fig. G2. Beta coefficients of those participants assigned to the gradual strategy.

G.3. Computational modelling results for Experiment 3

See Fig. G3.

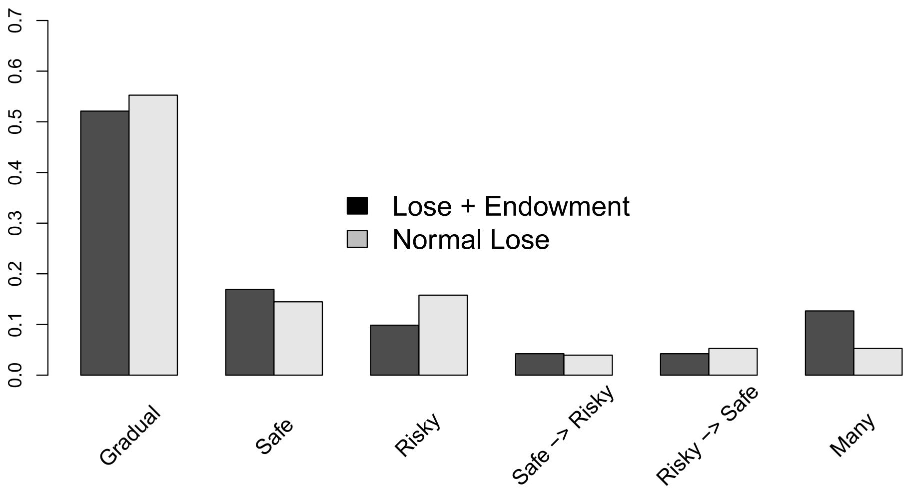

Fig. G3. Introducing endowment and losses does not affect the strategies participants employ.

G.4. Computational modelling results for Experiment 4

See Figs. G4–G6.

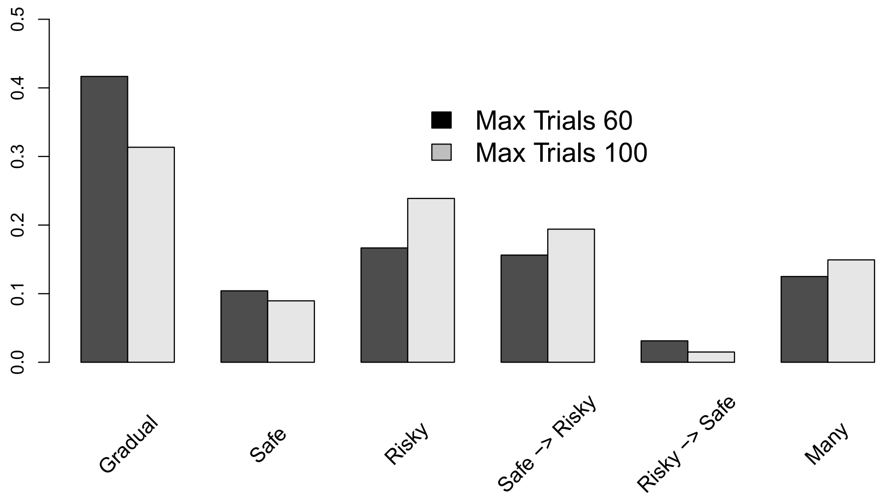

Fig. G4. Proportion of strategies for the two different maximum scopes in the keep condition.

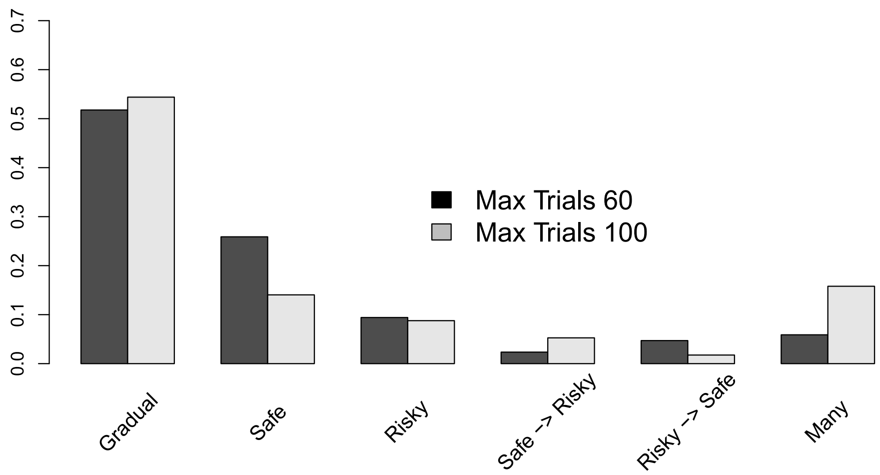

Fig. G5. Proportion of strategies for the two different maximum scopes in the lose condition.

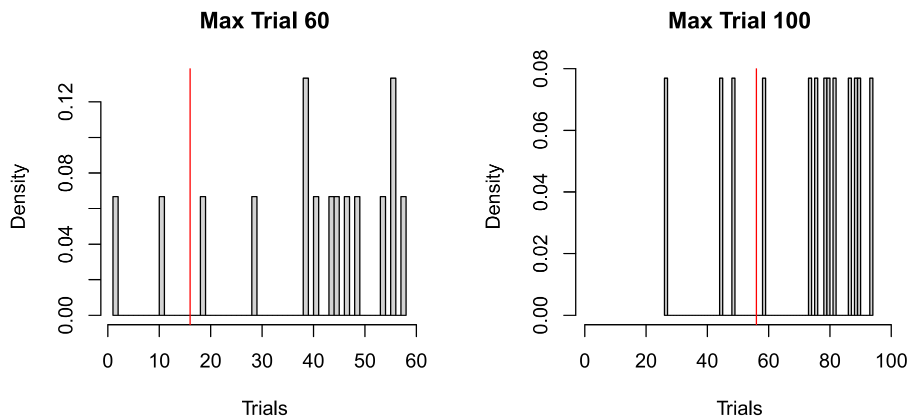

Fig. G6. Switch points in comparison to the optimal switch point in the keep condition.

Appendix H. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.cogpsych.2025.101735.

Data availability

Decision Making Under Extinction Risk (Original data) (OSF)

References

Aldy, J. E., & Viscusi, W. K. (2007). Age differences in the value of statistical life: revealed preference evidence. *Review of Environmental Economics and Policy*, *1*(2), http://dx.doi.org/10.1093/reep/rem014.

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. http://dx.doi.org/10.3758/BF03212962.

Baron, J., & Greene, J. (1996). Determinants of insensitivity to quantity in valuation of public goods: Contribution, warm glow, budget constraints, availability, and prominence. *Journal of Experimental Psychology: Applied*, *2*(2), 107. http://dx.doi.org/10.1037/1076-898X.2.2.107.

Blais, A.-R., & Weber, E. U. (2006). A domain-specific risk-taking (DOSPERT) scale for adultpopulations. *Judgment and Decision Making*, *1*(1), 33–47. http://dx.doi.org/10.1017/S1930297500000334.

Bland, J. R. (2019). How many games are we playing? An experimental analysis of choice bracketing in games. *Journal of Behavioral and Experimental Economics*, *80*, 80–91. http://dx.doi.org/10.1016/j.socec.2019.03.011.

Camerer, C. F., & Hogarth, R. M. (1999). The effects of financial incentives in experiments: A review and capital-labor-production framework. *Journal of Risk and Uncertainty*, *19*(1), 7–42. http://dx.doi.org/10.1023/A:1007850605129.

Cameron, C. D., & Payne, B. K. (2011). Escaping affect: how motivated emotion regulation creates insensitivity to mass suffering.. *Journal of Personality and Social Psychology*, *100*(1), 7–42. http://dx.doi.org/10.1023/A:1007850605129.

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

Clark, H., Cárdenas, M., Dybul, M., Kazatchkine, M., Liu, J., Miliband, D., Nordström, A., Sudan, P., Zedillo, E., Obaid, T., et al. (2022). Transforming or tinkering: the world remains unprepared for the next pandemic threat. *The Lancet*, *399*(10340), 1995–1999. http://dx.doi.org/10.1016/S0140-6736(22)00929-1.

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. http://dx.doi.org/10.31234/osf.io/qn7k5.

Crosetto, P., & Filippin, A. (2013). The "bomb" risk elicitation task. *Journal of Risk and Uncertainty*, *47*, 31–65. http://dx.doi.org/10.1007/s11166-013-9170-z.

Cubitt, R., Starmer, C., & Sugden, R. (1998). Dynamic choice and the common ratio effect: An experimental investigation. *The Economic Journal*, *108*(450), 1362–1380. http://dx.doi.org/10.1111/1468-0297.00346.

De Palma, A., Ben-Akiva, M., Brownstone, D., Holt, C., Magnac, T., McFadden, D., Moffatt, P., Picard, N., Train, K., Wakker, P., et al. (2008). Risk, uncertainty and discrete choice models. *Marketing Letters*, *19*, 269–285. http://dx.doi.org/10.1080/10407413.2013.810502.

Desvousges, W. H., Johnson, F. R., Dunford, R. W., Hudson, S. P., Wilson, K. N., & Boyle, K. J. (1993). Measuring natural resource damages with contingent valuation: tests of validity and reliability. In B. H. Baltagi, & F. Moscore (Eds.), *vol. 220*, *Contributions to economic analysis* (pp. 91–164). http://dx.doi.org/10.1016/B978-0-444-81469-2.50009-2.

Dickert, S., Västfjäll, D., Kleber, J., & Slovic, P. (2015). Scope insensitivity: The limits of intuitive valuation of human lives in public policy. *Journal of Applied Research in Memory and Cognition*, *4*(3), 248–255. http://dx.doi.org/10.1016/j.jarmac.2014.09.002.

Dixit, A. K. (1990). *Optimization in economic theory* (2nd ed.). Oxford University Press.

Douglas, B. D., Ewell, P. J., & Brauer, M. (2023). Data quality in online human-subjects research: Comparisons between mturk, prolific, CloudResearch, qualtrics, and SONA. *Plos One*, *18*(3), Article e0279720.

Eagleton, T. (2008). *The meaning of life: A very short introduction*. OUP Oxford.

Elga, A., Zhu, J.-Q., & Griffiths, T. L. (2024). People make suboptimal decisions about existential risks. URL https://www.researchgate.net/publication/384230783_People_Make_Suboptimal_Decisions_about_Existential_Risks.

Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics*, *75*(4), 643–669. http://dx.doi.org/10.2307/1884324.

Eyal, P., David, R., Andrew, G., Zak, E., & Ekaterina, D. (2021). Data quality of platforms and panels for online behavioral research. *Behavior Research Methods*, 1–20.

Fiedler, K. (2018). How to make psychology a genuine science of behavior: Comment on dolinski's thoughtful paper. *Social Psychological Bulletin*, *13*(2), 1–7, URL https://doi/10.5964/spb.v13i2.26079.

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

Frey, R., Pedroni, A., Mata, R., Rieskamp, J., & Hertwig, R. (2017). Risk preference shares the psychometric structure of major psychological traits. *Science Advances*, *3*(10), Article e1701381.

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. http://dx.doi.org/10.1016/j.psychres.2014.03.033.

Garcia-Marques, L., & Ferreira, M. B. (2018). Is observing behaviour the best way to understand behaviour? *Social Psychological Bulletin*, *13*, 1–7. http://dx.doi.org/10.5964/spb.v13i2.26076.

Guiso, L., & Paiella, M. (2008). Risk aversion, wealth, and background risk. *Journal of the European Economic Association*, *6*(6), 1109–1150. http://dx.doi.org/10.1162/JEEA.2008.6.6.1109.

Harris, J. (1985). *The Value of Life: An Introduction to Medical Ethics*. London: Routledge & Kegan Paul.

Hauser, O. P., Rand, D. G., Peysakhovich, A., & Nowak, M. A. (2014). Cooperating with the future. *Nature*, *511*, 220–223. http://dx.doi.org/10.1038/nature13530.

Henninger, F., Shevchenko, Y., Mertens, U. K., Kieslich, P. J., & Hilbig, B. E. (2021). Lab. js: A free, open, online study builder. *Behavior Research Methods*, 1–18. http://dx.doi.org/10.3758/s13428-019-01283-5.

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. http://dx.doi.org/10.1111/j.0956-7976.2004.00715.x.

Hertwig, R., & Ortmann, A. (2001). Experimental practices in economics: A methodological challenge for psychologists? *Behavioral and Brain Sciences*, *24*(3), 383–403. http://dx.doi.org/10.1017/S0140525X01004149.

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

Imas, A. (2016). The realization effect: Risk-taking after realized versus paper losses. *American Economic Review*, *106*(8), 2086–2109, URL https://www.aeaweb.org/articles?id=10.1257/aer.20140386.

Jacquet, J., Hagel, K., Hauert, C., Marotzke, J., Röhl, T., & Milinski, M. (2013). Intra-and intergenerational discounting in the climate game. *Nature Climate Change*, *3*(12), 1025–1028. http://dx.doi.org/10.1038/nclimate2024.

Kahneman, D., & Knetsch, J. L. (1992). Valuing public goods: the purchase of moral satisfaction. *Journal of Environmental Economics and Management*, *22*(1), 57–70. http://dx.doi.org/10.1016/0095-0696(92)90019-S.

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. http://dx.doi.org/10.1023/A:1007835629236.

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

Kellen, D., Pachur, T., & Hertwig, R. (2016). How (in) variant are subjective representations of described and experienced risk and rewards? *Cognition*, *157*, 126–138. http://dx.doi.org/10.1016/j.cognition.2016.08.020.

Kniesner, T. J., & Viscusi, W. K. (2019). The value of a statistical life. In *Oxford research encyclopedia of economics and finance*.

Konovalova, E., & Pachur, T. (2021). The intuitive conceptualization and perception of variance. *Cognition*, *217*, Article 104906. http://dx.doi.org/10.1016/j.cognition.2021.104906.

Krueger, P. M., Callaway, F., Gul, S., Griffiths, T. L., & Lieder, F. (2024). Identifying resource-rational heuristics for risky choice.. *Psychological Review*, http://dx.doi.org/10.1037/rev0000456.

Krupnick, A. (2007). Mortality-risk valuation and age: stated preference evidence. *Review of Environmental Economics and Policy*, http://dx.doi.org/10.1093/reep/rem016.

Kuhn, E. (1992). *Nietzsches Philosophie des europäischen Nihilismus [Nietsche's Philosophy of European Nihilism]*. De Gruyter.

Lange, K., Kühn, S., & Filevich, E. (2015). "Just another tool for online studies"(JATOS): An easy solution for setup and management of web servers supporting online studies. *PloS One*, *10*(6), Article e0130834. http://dx.doi.org/10.1371/journal.pone.0130834.

Lejarraga, T., & Hertwig, R. (2021). How experimental methods shaped views on human competence and rationality.. *Psychological Bulletin*, *147*(6), 535. http://dx.doi.org/10.1037/bul0000324.

Lejuez, C. W., Aklin, W. M., Zvolensky, M. J., & Pedulla, C. M. (2003). Evaluation of the balloon analogue risk task (BART) as a predictor of adolescent real-world risk-taking behaviours. *Journal of Adolescence*, *26*(4), 475–479. http://dx.doi.org/10.1016/s0140-1971(03)00036-8.

Lejuez, C. W., Read, J. P., Kahler, C. W., Richards, J. B., Ramsey, S. E., Stuart, G. L., Strong, D. R., & Brown, R. A. (2002). Evaluation of a behavioral measure of risk taking: the balloon analogue risk task (BART). *Journal of Experimental Psychology: Applied*, *8*(2), 75. http://dx.doi.org/10.1037//1076-898x.8.2.75.

Levy, H. (2015). *Stochastic dominance: Investment decision making under uncertainty*. Springer.

Lieder, F., Griffiths, T. L., & Hsu, M. (2015). Utility-weighted sampling in decisions from experience. In *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. http://dx.doi.org/10.1037/rev0000074.

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

Lockwood, M. (1988). Quality of life and resource allocation. *Royal Institute of Philosophy Supplements*, *23*, 33–55.

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. http://dx.doi.org/10.1002/bdm.1792.

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. http://dx.doi.org/10.3758/s13423-013-0542-9.

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*, *36*(4), Article e2335. http://dx.doi.org/10.1002/bdm.2335.

Marani, M., Katul, G. G., Pan, W. K., & Parolari, A. J. (2021). Intensity and frequency of extreme novel epidemics. *Proceedings of the National Academy of Sciences*, *118*(35), Article e2105482118. http://dx.doi.org/10.1073/pnas.2105482118.

Michael, O., & Mark, O. (2024). The world is not ready for the next pandemic. *Foreign Affairs*, URL https://www.foreignaffairs.com/united-states/world-not-ready-next-pandemic.

Mitsis, P. (2020). *Oxford handbook of Epicurus and Epicureanism*. Oxford University Press.

Pedroni, A., Frey, R., Bruhin, A., Dutilh, G., Hertwig, R., & Rieskamp, J. (2017). The risk elicitation puzzle. *Nature Human Behaviour*, *1*(11), 803–809. http://dx.doi.org/10.1038/s41562-017-0219-x.

Pensions, A. (2024). Lifecycle strategies. URL https://pensions.axa-employeebenefits.co.uk/investments/lifecycle-strategies, (Accessed 04 November 2024).

Peress, J. (2004). Wealth, information acquisition, and portfolio choice. *The Review of Financial Studies*, *17*(3), 879–914. http://dx.doi.org/10.1093/rfs/hhg056.

Perfors, A., & Van Dam, N. T. (2018). Human decision making in black swan situations.. *40*, In *Proceedings of the annual meeting of the cognitive science society*. URL https://escholarship.org/uc/item/3p76280v.

Persson, E., & Tinghög, G. (2020). Opportunity cost neglect in public policy. *Journal of Economic Behavior and Organization*, *170*, 301–312. http://dx.doi.org/10.1016/j.jebo.2019.12.012.

Pleskac, T. J. (2008). Decision making and learning while taking sequential risks. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *34*(1), 167.

Pleskac, T. J., Wallsten, T. S., Wang, P., & Lejuez, C. (2008). Development of an automatic response mode to improve the clinical utility of sequential risk-taking tasks.. *Experimental and Clinical Psychopharmacology*, *16*(6), 555. http://dx.doi.org/10.1037/a0014245.

Rabin, M., & Weizsäcker, G. (2009). Narrow bracketing and dominated choices. *American Economic Review*, *99*(4), 1508–1543. http://dx.doi.org/10.1257/aer.99.4.1508.

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). Dordrecht: Springer, http://dx.doi.org/10.1007/978-94-017-1406-8_7.

Rieskamp, J. (2008). The probabilistic nature of preferential choice. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *34*(6), 1446. http://dx.doi.org/10.1037/a0013646.

Sartre, J.-P. (1972). *Being and Nothingness*. New York: Pocket Books.

Savage, L. J. (1972). *The foundations of statistics*. Courier Corporation.

Schulze, C., van Ravenzwaaij, D., & Newell, B. R. (2015). Of matchers and maximizers: How competition shapes choice under risk and uncertainty. *Cognitive Psychology*, *78*, 78–98. http://dx.doi.org/10.1016/j.cogpsych.2015.03.002.

Seneca, L. A., et al. (2004). *On the shortness of life*. Penguin UK.

Slovic, P., & Västfjäll, D. (2010). Affect, moral intuition, and risk. *Psychological Inquiry*, *21*(4), 387–398. http://dx.doi.org/10.1080/1047840X.2010.521119.

Slovic, P., & Weber, E. U. (2013). Perception of risk posed by extreme events. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2293086.

Spiller, S. A. (2019). Opportunity cost neglect and consideration in the domain of time. *Current Opinion in Psychology*, *26*, 98–102. http://dx.doi.org/10.1016/j.copsyc.2018.10.001.

Stagnaro, M. N., Druckman, J., Berinsky, A. J., Arechar, A. A., Willer, R., & Rand, D. (2024). Representativeness versus attentiveness: A comparison across nine online survey samples. http://dx.doi.org/10.31234/osf.io/h9j2d.

Stephens, D. W., & Krebs, J. R. (1986). *Foraging theory*. Princeton University Press.

Sundh, J. (2024). Human behavior in the context of low-probability high-impact events. *Humanities and Social Sciences Communications*, *11*(1), 1–10. http://dx.doi.org/10.1057/s41599-024-03403-9.

Sunstein, C. R., & Zeckhauser, R. (2011). Overreaction to fearsome risks. *Environmental and Resource Economics*, *48*, 435–449. http://dx.doi.org/10.1007/s10640-010-9449-3.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*, 297–323. http://dx.doi.org/10.1007/BF00122574.

Vanguard (2021). Vanguard's life-cycle investing model (VLCM): A general portfolio framework. URL https://corporate.vanguard.com/content/dam/corp/research/pdf/Vanguards-Life-Cycle-Investing-Model-VLCM-A-general-portfolio-framework-US-ISGVLCM_032021_online.pdf, (Accessed 04 November 2024).

Västfjäll, D., & Slovic, P. (2020). A psychological perspective on charitable giving and monetary donations: The role of affect. In T. Zaleskiewicz, & J. Traczyk (Eds.), *Psychological perspectives on financial decision making* (pp. 331–345). Springer, http://dx.doi.org/10.1007/978-3-030-45500-2_14.

Vehtari, A., Gelman, A., Simpson, D., Carpenter, B., & Bürkner, P.-C. (2021). Rank-normalization, folding, and localization: An improved $\widehat{R}$ for assessing convergence of MCMC. *Bayesian Analysis*, *1*(1), 1–28. http://dx.doi.org/10.1214/20-BA1221.

Wallsten, T. S., Pleskac, T. J., & Lejuez, C. W. (2005). Modeling behavior in a clinically diagnostic sequential risk-taking task. *Psychological Review*, *112*(4), 862, URL https://psycnet.apa.org/doi/10.1037/0033-295X.112.4.862.

Wang-Ly, N., & Newell, B. R. (2024). Income volatility and saving decisions: Experimental evidence. *Journal of Behavioral and Experimental Finance*, *43*, Article 100941. http://dx.doi.org/10.1016/j.jbef.2024.100941.

Watkins, C. J. C. H. (1989). Learning from delayed rewards.

Watkins, C. J. C. H., & Dayan, P. (1992). Q-learning. *Machine Learning*, *8*, 279–292. http://dx.doi.org/10.1007/BF00992698.

Weber, E. U. (2006). Experience-based and description-based perceptions of long-term risk: Why global warming does not scare us (yet). *Climatic Change*, *77*(1), 103–120. http://dx.doi.org/10.1007/s10584-006-9060-3.

Weber, E. U., Blais, A.-R., & Betz, N. E. (2002). A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors. *Journal of Behavioral Decision Making*, *15*(4), 263–290. http://dx.doi.org/10.1002/bdm.414.

Weiss, D. J., & Shanteau, J. (2021). The futility of decision making research. *Studies in History and Philosophy of Science Part A*, *90*, 10–14. http://dx.doi.org/10.1016/j.shpsa.2021.08.018, URL https://www.sciencedirect.com/science/article/pii/S0039368121001345.

Wiener, J. B. (2016). The tragedy of the uncommons: On the politics of apocalypse. *Global Policy*, *7*, 67–80.

Williams, A. (1997). Intergenerational equity: an exploration of the 'fair innings' argument. *Health Economics*, *6*(2), 117–132.

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

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. http://dx.doi.org/10.1037/bul0000115.