American Psychological Association

Psychological Review

Manuscript version of

The Development of Risk Behaviors and Their Cultural Transmission

Alejandro Perez Velilla, Bret Beheim, Paul E. Smaldino

Funded by:

- John Templeton Foundation
- Max Planck Institute for Evolutionary Anthropology
- National Science Foundation

The development of risk behaviors and their cultural transmission

Alejandro Pérez Velilla[1,*], Bret Beheim[2], Paul E. Smaldino[1,3]
[1] University of California, Merced
[2] Max Planck Institute for Evolutionary Anthropology
[3] Santa Fe Institute
[*] Corresponding author: Alejandro Pérez Velilla, aperezvelilla@ucmerced.edu

Abstract

We use cultural evolutionary models to examine how individual experiences and culturally-inherited information jointly shape risk-taking behavior under environmental uncertainty. We find that learning processes not only generate considerable variation in risk beliefs and behaviors, but also that conservative learning strategies – emphasizing the preservation of generational knowledge – excel in high-risk settings, promoting risk avoidance and long-term survival but limiting growth when conditions improve. In contrast, exploratory learning strategies – leveraging juvenile exploration and peer influence – foster risk-tolerant behaviors that thrive in affluent, low-risk settings where wealth buffers and social safety nets reduce the costs of miscalculations. Introducing economic stratification to the model reveals how wealth disparities and inter-class interactions reinforce these patterns, exacerbating differences in learning strategies and risk-taking behaviors within populations, and perpetuating socioeconomic inequalities through the cultural inertia of excessive risk avoidance. By uniting developmental, social, and evolutionary perspectives, our framework provides a novel lens on the cultural evolution of risk-taking behavior and its broader societal implications.

Keywords: uncertainty; social learning; cultural evolution; poverty traps; parochialism; decision making

1 Introduction

Nearly everyone must make repeated risky decisions throughout their productive lives—from starting a business to pursuing higher education, from proposing marriage to investing in the stock market. Each of these decisions might pay off, but a tactical misstep might lead to losses that can be hard or even impossible to recover from. Risking too much can lead to catastrophe, but too much caution can lead to middling performance. If one knows all the risks involved, one can—in theory, at least—optimize one's bets, but the knowledge needed to successfully navigate risky decisions can itself be risky to obtain. Social learning becomes of prime importance when learning how to navigate the risky environments we experience throughout life; learning from others allows us to reduce error and inherit the hard-won wisdom of others without incurring the costs of trial and error, which can include overly risky decisions that lead to ruin. It is therefore of interest not only how people should deal with risk in decision making, but how they should *learn* about the risks involved in their decisions.

Acquiring optimal risk behavior—that which yields the best distribution of outcomes for a given environment—might be straightforward in an environment that is fully predictable or where all risks are accurately known. But real-world conditions are inherently uncertain in the Knightian sense (Knight, 1921), and in risky settings with the potential for catastrophic loss caution often becomes essential. The challenge of optimizing risk behaviors is therefore difficult, if not impossible, for an individual to solve without help. However, research on risk-sensitive behavior and on social learning strategies has, for the most part, developed independently. This is unfortunate given the importance of social learning in helping organisms navigate environmental uncertainty (Boyd and Richerson, 1985; Laland, 2004; Kendal et al., 2018), which has been proposed as one of the drivers of humans' extreme reliance on culture (Boyd and Richerson, 1985; Boyd et al., 2011), as well as more generally explaining the extent of human psychological and behavioral variation across the globe (Chudek and Henrich, 2011; Gelfand et al., 2024).

Understanding the origins of risk behavior is central to many fields, from economics and psychology to anthropology and evolutionary biology, where risk-relevant behavior is often viewed from a life-history lens (Nettle and Frankenhuis, 2020). Evolutionarily-oriented economic models posit that a degree of risk aversion is adaptive, reflecting evolved traits for navigating uncertain environments where rational decision-making is key to survival, particularly in cases where risk is aggregate (correlated) (Robson, 1996; Levy, 2010; Zhang et al., 2014; Robson and Samuelson, 2019; Robson and Orr, 2021; Robson and Samuelson, 2022; Alger, 2023), which is consistent with work in theoretical evolutionary biology that highlights how aversion to risk (in the form of bet-hedging strategies) can evolve as a response to within- and between-generation environmental change (Starrfelt and Kokko, 2012; Schreiber, 2015; Mallpress et al., 2015) or to asymmetric fitness functions when the outcomes of behaviors under selective pressures are themselves stochastic Vercken et al. (2012); Haaland et al. (2019). Psychological research, on the other hand, tends to emphasize how risk behavior is related to cognitive and emotional factors in risk perception in specific domains Weber and Hsee (1999), with developmental perspectives viewing childhood and adolescence as critical periods for refining risk-related behaviors in response to adversity (Gopnik, 2020; Frankenhuis and Gopnik, 2023; Xu et al., 2023). Despite the existence of these rich perspectives, they remain largely disconnected. A cultural evolutionary viewpoint—bridging economic, cognitive, social, and evolutionary sciences—offers a framework for integrating these insights and generating new testable hypotheses.

Here, we develop a model that clarifies how risk perception and behavior within a decision-making domain emerge from the interplay of evolved predispositions for navigating risk-reward trade-offs, individual life experiences, and cultural transmission. We assume that the decision-making domains that drive risk behavior exhibit an asymmetric structure: recovering from losses is more difficult than achieving gains, especially when when one has less to begin with and when subsistence is at risk (Cieslik and D'Aoust, 2018). Therefore, domains that have a higher impact on subsistence and cultural influence incentivize more cautious risk-taking, especially when buffers are low. At the same time, we answer calls to incorporate culture into models of risk behavior (Weber and Hsee, 1999; Bowles, 1998) by showing how, in response to pressures to learn socially, distributions of risk behavior can change both within and across generations. We also recover several observed empirical patterns in risk behavior, such as its adversity-mediated developmental trajectory (Brooks-Gunn and Duncan, 1997; Paulsen et al., 2012; Amir et al., 2020; Frankenhuis and Gopnik, 2023; Xu et al., 2023; Leonard and Sommerville, 2024), the presence of extreme risk avoidance and extreme risk taking in economically-deprived populations (Banerjee, 2004; Haushofer and Fehr, 2014; de Courson et al., 2025), its differential cultural heritability (Dohmen et al., 2011; Shore, 2011; Dohmen et al., 2012; Zumbuehl et al., 2013; Arrondel, 2013; Necker and Voskort, 2014; Sepahvand and Shahbazian, 2017; Wolff, 2020; Sepahvand and Shahbazian, 2021; Hong et al., 2024) and the way that acute population shocks and the incentivization of conservative learning strategies can keep risk avoidance suboptimally high (Moya, 2018; Kim and Lee, 2014; Augsburger and Elbert, 2017; Nunn, 2022).

1.1 A theoretical framework for understanding risk

How do individuals discern whether their environment rewards investments or punishes them with ruinous losses? Trial-and-error learning is possible, but costly—especially when information is noisy or slow to accumulate. Childhood may act as a protective period that permits exploration and peer learning with minimal immediate risks (Gopnik, 2020; Frankenhuis and Gopnik, 2023; Xu et al., 2023) at the cost of omitting important information about those risks, although greater adversity can also lead to "shorter" childhoods, in which children are expected to significantly contribute to productive and caregiving activities in riskier settings Kramer (2005, 2021); Lew-Levy et al. (2022). As individuals mature, additional learning occurs through direct experience and from observing others' successes, failures and catastrophes. When environmental risk changes little between generations, vertical or oblique transmission of established practices from elders can be a highly effective means of stabilizing near-optimal risk behaviors (Boyd and Richerson, 1985; Reyes-García et al., 2016; Deffner and McElreath, 2022; Turner et al., 2023), avoiding optimistic (too risk tolerant) or pessimistic (too risk avoidant) behaviors. However, environmental conditions often do change and can widely differ from context to context, providing a large diversity of backgrounds for developmental trajectories.

For example, ruin boundaries might not be the same everywhere. While wealth, as a social phenomenon, is present in subsistence economies (Borgerhoff Mulder et al., 2009) and even in some non-human systems (Strauss and Shizuka, 2022), what can be considered an absorbing boundary or point of "ruin" is likely to vary significantly from context to context (both within and across generations), as the currencies that constitute wealth itself can be very different from place to place and from time to time. In our model, we treat ruin as a complete loss of cultural influence on prospective learners from future generations. This is because financial ruin often goes hand in hand with loss of status and social ostracism in contemporary large-scale societies (Mills and Zavaleta, 2015), and social capital itself can be seen as a form of wealth that can be completely lost, especially in subsistence economies where it feeds into the structure of material exchange networks (Borgerhoff Mulder et al., 2009; Jones and Ready, 2022). For example, among the Pokot of Northwestern Kenya, cohesive and durable livestock exchange networks are frequently likened to a form of banking and insurance (see Bollig (1998), p. 142), while classic accounts of Melanesian "Big Men" consistently emphasize not only how social wealth is garnered through savvy reciprocal exchanges, but also how those who do not participate in these social games can quickly lose influence relative to those who do (see Strathern (1971), p. 10, "Moka as a Game"). Loss of influence stemming from loss of wealth is also apparent in modern-day mixed economies: in Kangiqsujuaq, a Canadian Inuit community, foraging for country food and the sharing of harvests (activities that can exhibit a significant economic entry barrier) often confer status and sociopolitical influence for those wealthy enough to afford these costly activities (Ready and Power, 2018). In all of these contexts, the cultural influence of individuals directly stems from their investment on these social currencies, and failing to invest correctly can lead to total irrelevance.

Environmental conditions also vary widely within and across populations, and not all individuals within an environment face identical conditions. Baseline wealth, institutional stability, and social safety nets can strongly influence the consequences of risk-taking. Risk behaviors are not fixed; they develop over the lifespan and evolve across generations, shaped by changing conditions, cultural norms, and individual adversity (Amir and Jordan, 2017; Amir et al., 2018, 2020; Xu et al., 2023). For example, younger individuals often show higher risk tolerance, which then declines with age (Paulsen et al., 2012; Dohmen et al., 2017; Leonard and Sommerville, 2024). Early-life adversity can accelerate the onset of risk avoidance (Gopnik, 2020; Frankenhuis and Gopnik, 2023; Xu et al., 2023), and shocking or traumatic experiences can create long-lasting shifts toward heightened caution (Moya, 2018; Kim and Lee, 2014; Augsburger and Elbert, 2017). Moreover, while individual variation in risk avoidance has been largely ignored in models of optimal risk-taking, empirical work suggests that different socioeconomic strata can exhibit complex risk behavior distributions with distinct variance signatures, especially in deprived scenarios where high risk avoidance can coexist with disproportionate risk taking (Banerjee, 2004; Haushofer and Fehr, 2014; de Courson et al., 2025).

Environmental conditions can also change with time, such that strategies that previously worked well are no longer effective. Such changes can undermine the value of conservative learning strategies that maintain the steady transmission of elders' acquired information (Deffner and McElreath, 2022), helping explain the findings that the transmission of risk-related attitudes and behavior from parents and other elders might strongly depend on local cultural, social and economic conditions (Dohmen et al., 2011, 2012). This is exemplified studies indicating stronger intergenerational transmission in non-Western countries like Burkina Faso (Sepahvand and Shahbazian, 2017; Wolff, 2020; Sepahvand and Shahbazian, 2021), China and Korea (Hong et al., 2024) than in Western countries such as Germany (Zumbuehl et al., 2013; Necker and Voskort, 2014), France (Arrondel, 2013) and the United States (Shore, 2011), as well as observations of higher reliance on inter-generational transmission by Bangladeshi immigrant populations in the United Kingdom (Mesoudi et al., 2016).

Taking these lines of research into account, we propose that risk behaviors emerge from a combination of evolutionary and developmental processes, as coarse evolved predispositions for avoiding aggregate risk interact with the finer-grained plasticity of learning from experiential and cultural inputs. Four key factors shape this process: 1) a baseline level of risk aversion, shaped by evolutionary pressures to avoid aggregate risk during the human evolutionary trajectory; 2) low-risk environmental sampling during early developmental periods, reinforced by horizontal learning from peer interactions (exploratory learning); 3) vertical and oblique cultural transmission of risk-related knowledge from previous generations (conservative learning); and 4) lifetime adjustments in response to social information that indicates that one's current risk behaviors are overly optimistic (sensitivity to peer ruin). As different learning strategies have differential success in growing payoffs while avoiding ruin, natural selection leads to the transmission of successful strategies, which themselves give shape to the cultural selection that acts on the lifelong acquisition of risk beliefs and behaviors. Our emphasis in this paper therefore deviates from the more common concern with identifying optimal risk behaviors; instead, we focus on the optimal strategies for *learning about risk*. The tremendous capacity of humans to adapt to a wide range of local physical and social environments is a testament to the power of our cultural evolution, and highlights that many of our behavioral strategies, including those for managing risk, are learned (Rogers, 1988; Boyd et al., 2011). While the model itself is agnostic to whether learning strategies are inherited culturally or genetically, we see them as cultural traits passed down by individuals who avoided ruin during their lifetimes, a form of phenotypic natural selection.

We study a simple investment dynamic in which agents must gamble to increase their wealth in a single domain and face a risk-reward trade-off: the more risk they take, the higher their potential return on investment but also the greater their risk of ruin. To this we add developmental stages in which agents use individual and peer learning as ways to acquire information and estimate their environmental conditions, building risk perceptions that lead to risk behaviors attuned to their present conditions, as well as to quickly adjust those perceptions in response to witnessing the ruin of other individuals. We then add a layer of cultural inheritance, in which agents can use information from previous generations to further hone their perception and risk behavior.

Our model reveals how the cultural evolution of different strategies for learning about risk—ranging from conservative to exploratory—can yield distinct cultural equilibria in risk behavior across domains, environments and between socioeconomic strata. Conservative learning strategies, characterized by the strong influence of information inherited from elders, thrive in high-risk settings but can delay adaptation when conditions improve. Exploratory learning strategies, which rely more on immediate environmental feedback and peer influence, are better suited to improving conditions, but may lead to ruin if misapplied in risky environments. We find that conservative learning strategies typically outperform exploratory ones, but perform worse under low-risk, high-wealth conditions where the insulation provided by ample wealth buffers encourages boldness and diversification in information sources. We also consider the differential consequences of when environmental change is correlated across individuals (what we call *aggregate uncertainty*) versus when such change is uncorrelated across individuals (*idiosyncratic uncertainty*), showing that conservative social learning is robust to both forms of change under high risk, but as risk decreases idiosyncratic uncertainty disincentivizes most forms of social learning, except for sensitivity to peer ruin.

We also examine the effects of socioeconomic stratification by allowing group-specific differences in environmental risk and wealth buffers. In such scenarios, wealthy groups can afford more risk-taking and may adopt payoff bias as a parochial strategy, reinforcing their advantage and cultural distinctiveness. In contrast, disadvantaged groups, facing harsher conditions, must resort to parochial learning (which reduces the size of their pool of learning models) to avoid ill-fitting information gleaned from wealthier out-groups. This leads to entrenched socioeconomic disparities and a heightened cultural inertia of risk behaviors.

Our framework provides a coherent mechanistic explanation for a wide range of observed phenomena in human risk behavior, linking economic conditions, cultural processes, life history, and social learning strategies. The framework further predicts relationships between wealth, learning strategies, and risk behaviors that can guide subsequent empirical research. We offer a unified mechanistic explanation for juvenile optimism, the increase of pessimism throughout the lifetime, the transmission of acquired pessimism to new generations, and the relationship of these processes to wealth, environmental adversity, and structural inequalities. In the following sections, we detail our model, present the results of our analyses, and discuss the implications for understanding the cultural evolution of responses to risk.

2 The model

2.1 Optimal gambling dynamics under cliff-edge effects and risk of ruin

We begin by describing a simple model of optimal gambling, upon which we will then introduce risk of ruin, development, learning, population structure, and cultural evolutionary dynamics.

Wealth growth dynamics

We start with a single-asset risky investment dynamic (Kelly, 1956), representing investment decisions in a single domain by agents who can opt into a multiplicative wealth growth scheme (i.e. the stock market). An agent starting with an initial wealth V_0 chooses a stake $s \in [0, 1]$, representing the fraction of their wealth they will bet, and then proceeds to throw a biased coin every time period (for a total of T periods), so that each attempt results in a success with probability u and a failure with probability $1 - u$. An agent's chosen stake thus describes their confidence that their gamble in this domain will pay off and the fraction of their wealth they are willing to forego in case of failure.

When a gamble is successful, the player increases their wealth by a factor $1 + s$, otherwise their wealth decreases by a factor $1 - s$. This gives us

$$V_{t+1} = V_t \left[(1+s)x + (1-s)(1-x)\right] \tag{1}$$

where V_t is the wealth of the agent at time t and $x \sim \text{Bernoulli}(u)$. Agents seek to maximize the long-term growth rate of their wealth over many time periods T. An agent's expected wealth at the end of its lifetime can be written as

$$V_T(s) = V_0 \prod_{t=1}^{T} \exp\left[g_t(s)\right] = V_0 \exp\left[\sum_{t=1}^{T} g_t(s)\right] \tag{2}$$

with $g_t(s) = \log(1+s)$ if successful or $g_t(s) = \log(1-s)$ otherwise.

Written this way, we can see that the growth rate of V_T depends on a sum of random variables, given by $\log \frac{V_T}{V_0} = \sum_{t=1}^{T} g_t(s)$. To find the average change in wealth in one time period, we divide this sum by T. As T grows large, by the law of large numbers, this average log growth rate of V_T converges to the long-run log geometric mean of the outcomes, given by

$$\begin{aligned} G(s) &= \lim_{T\to\infty} \frac{\log \frac{V_T}{V_0}}{T} = \lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} g_t(s) \\ &= u \log(1+s) + (1-u)\log(1-s) \end{aligned} \tag{3}$$

so that the long-run growth can be approximated by

$$V_T(s) \approx V_0 \exp\Big[T \cdot G(s)\Big] \tag{4}$$

for large enough T. In other words, finding an optimal stake s^* that maximizes the expected long-term growth rate of wealth merely requires maximizing $G(s)$. This leads to an optimal stake that can be calculated explicitly. Popularly known as the Kelly criterion or Kelly bet (Poundstone, 2010) after its initial derivation by Kelly (1956), for this simple model it takes the following form:

$$s^* = \begin{cases} 2u - 1 : u \geq \frac{1}{2} \\ 0 \qquad : \text{otherwise} \end{cases} \tag{5}$$

This means that a decision-maker who knows the rate of success u can maximize their growth rate by following this formula, and should bet nothing if their *environmental edge* (the extent to which an environment provides an advantage to the gambler, which we denote $\epsilon = u - \frac{1}{2}$) is equal to or less than 0. Indeed, it is easy to show that an equal number of successes and failures yields an overall loss in the agent's wealth: note that the payoff of an agent with n successes and n failures is $V_0(1+s)^n(1-s)^n = V_0[(1+s)(1-s)]^n = V_0(1-s^2)^n$, which is necessarily less than or equal to V_0. Losses therefore carry more weight than gains in this dynamic, so it is important to be playing with a favorable edge. Equivalently, the is a positive correlation between the possible rewards from investment and the risk of losses: a *risk-reward tradeoff.* The leftmost plot of Figure 1 shows how this asymmetry translates into an asymmetry in the average growth rate as a function of stake: a lower-than-optimal stake leads to less of a penalty than a higher-than-optimal stake, a problem structure commonly refereed to as the "cliff-edge" effect. This is also reflected in the expression for s^*: a success rate of $\frac{1}{2}$ or less yields an optimal Kelly bet of 0.

Since the above process involves optimizing an exponential growth rate, optimal decision-making in this simple game happens when individuals' psychologies are represented by a logarithmic utility function (Bernoulli, 1954; Levy, 2010, 2024) and when they have perfect information on their odds of success; that is, when the expected utility of a gamble is evaluated as the expectation of the logarithms of its possible outcomes, a notion which aligns with the concept of geometric mean fitness in evolutionary biology (Gillespie, 1974), and which has been shown to be tied to decision problems in the presence of aggregate risk (Robson, 1996; Robson and Samuelson, 2019).

We assume individuals start life *as if* they possessed logarithmic utility functions as a baseline level of risk aversion. This represents an evolutionarily-optimized adaptive risk aversion that always leads to choosing the optimal stake in this dynamic, given perfect information about the environmental edge and given that there are no absorbing boundaries in place. This decision is meant to capture our broader assumption of the existence of a baseline degree of adaptive risk aversion that matches the unconstrained wealth growth dynamic (no absorbing boundary, full information), and should not be confused with a claim about the universality of the logarithmic utility function for decision-making. In principle, choosing any other value for the fixed curvature of the utility function would lead to a shifted but qualitatively similar version of the results we explore here and, as we show below, allowing for imperfect information about environmental edges and non-zero absorbing boundaries add uncertainty that requires agents to adjust this initial degree of risk-taking, which we model through changes in risk perception while keeping the curvature of the utility function fixed (Meder et al., 2021).

Absorbing boundaries: meeting a terrible fate

Gambling carries with it the risk of ruin. Wealth, for example, can plummet to a level from which it becomes impossible to recover (Smith et al., 2005; Rodems and Pfeffer, 2021; Miller et al., 2021). This "point of no return" is also known as an *absorbing boundary*, which we denote V_B. If wealth ever falls under this value during the T time periods of an agent's lifetime, we say that the agent is *ruined*: they cannot continue gambling and their wealth is set to zero. A decision-making domain with a high absorbing boundary thus incentivizes caution in betting strategies. When $V_B = 0$, the absorbing boundary can only be reached when an agent chooses to stake it all ($s = 1$), such that the first disfavorable outcome leads to total loss. $V_B = 0$ is also the condition under which Equation 5 actually represents the optimal stake. However, individuals are often unable to continue making risky gambles even if they have not technically lost everything, because a baseline amount of resources is required to support basic needs. We represent these scenarios by considering cases where $V_B > 0$.

Unless stated otherwise, we fix $V_B = 1$ for the remainder of the paper. Our minimum starting wealth will thus be $V_0 = 1$. When agents attempt to optimize their resource growth rates, the positive absorbing boundary therefore leads to optimal stakes below those indicated by Equation 5. The placement of the absorbing barrier at $V_B = 1$ is arbitrary; however, the important thing for agents' wealth is not the absolute value of the boundary, but rather their relative distance from it. Adding non-zero absorbing boundaries to the model introduces the need for numerical simulation in order to characterize important aspects of the system's behavior, such as the payoff distribution. However, we can use tools from the theory of martingales in order to find a Brownian approximation of the risk of ruin p_{ruin}, which is given by the following (see Corollary 2.4 in Asmussen and Albrecher (2010), p. 26):

$$p_{\text{ruin}}(s) \approx \left(\frac{V_B}{V_0}\right)^{\frac{2\mu}{\sigma^2}} \tag{6}$$

where

$$\mu = G(s) = u \log(1 + s) + (1 - u) \log(1 - s) \tag{7}$$

and

$$\sigma^2 = u\Big[\log(1+s) - \mu\Big]^2 + (1-u)\Big[\log(1-s) - \mu\Big]^2 \tag{8}$$

Equation 6 implies that it is the ratio of V_B to V_0, and not individual magnitudes of these quantities, that has an influence over the probability of ruin. Thus, we define the *wealth buffer* as

$$\aleph = 1 - \frac{V_B}{V_0} \tag{9}$$

which measures the extent to which initial wealth V_0 exceeds the absorbing boundary. When $\aleph = 0$, then $V_0 = V_B$ and the focal agent starts out ruined, while a focal agent with $\aleph = 1$ has an effectively infinite wealth buffer since this implies $V_0 = +\infty$ as long as $V_B > 0$. In this case, wealth growth rates are undefined, and so we restrict our use of $\aleph = 1$ to the case of zero absorbing barriers $V_B = 0$. This way, the wealth buffer $\aleph$, ranging from 0 to 1, fully parameterizes the effects of absorbing barriers. It represents a safety net that gives agents a cushion during difficult times (e.g., inherited wealth and/or stable institutions providing a social safety net).

To understand $\aleph$ intuitively, consider that for a fixed stake s, the wealth buffer $\aleph$ is the proportion of initial wealth that must be lost in an uninterrupted string of gambles in order for an agent to become ruined. If we represent the number of losses at ruin as B, we get this equation relating V_0 and V_B:

$$V_B = (1-s)^B V_0 \iff \frac{V_B}{V_0} = (1-s)^B \tag{10}$$

The ratio $\frac{V_B}{V_0}$ is therefore equal to the proportion of initial wealth an agent betting s and starting at V_0 has after an uninterrupted string of losses that eventually leads to ruin. This means the wealth buffer can also be written as

$$\aleph = 1 - (1-s)^B \tag{11}$$

We can also approximate the optimal stake (the Kelly bet) for cases involving absorbing boundaries by numerically solving the following expression:

$$s^*_{\text{ruin}} = \arg\max_{\text{s}} \bar{V}(s; \epsilon, \aleph) \tag{12}$$

where

$$\bar{V}(s; \epsilon, \aleph) = \exp([1 - p_{\text{ruin}}(s; \epsilon, \aleph)] \cdot G(s)), \tag{13}$$

which converges to Equation 5 in the limit of $\aleph \to 1$. This is an example of optimization with respect to a "safety-first expected utility" model, which assumes agents make decisions based on both the expected utility of gambles as well as their risk of hitting an absorbing boundary (Levy and Levy, 2009). Figure 1 illustrates the average growth rate as a function of stake for several values of the environmental edge (ϵ) and the wealth buffer ($\aleph$). Optimal stake decreases with both environmental edge and wealth buffer. It is clear that when absorbing boundaries are non-zero and wealth buffers are finite, adopting risk behavior (stakes) that are optimized for boundary-less scenarios will likely lead to ruin due to overbetting.

Given that the wealth growth rates resulting from particular stakes depend on both the favorability of the environment to gambles (the environmental edge) and individual wealth buffers, optimizing risk behaviors requires the acquisition of quite a lot of information. Subsequent analyses will consider

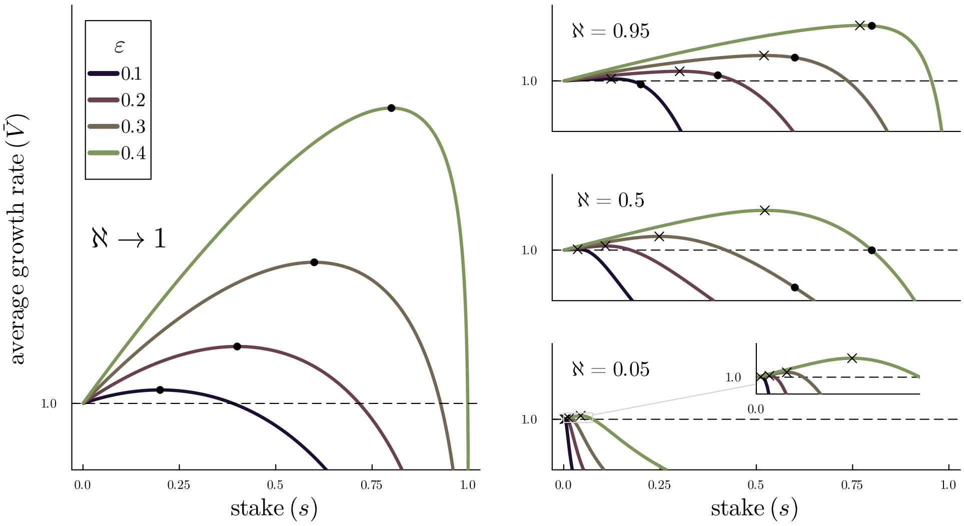

Figure 1: Average geometric wealth growth rates $\bar{V}$, as a function of stake (s). The three plots on the right show the average geometric growth rate of wealth, incorporating the effects of absorbing barriers. The stakes given by Equation 5 are shown as black dots, while the growth-maximizing stakes given by Equation 12 are indicated by an X. Without absorbing boundaries (right plot), growth rate is maximized according to Equation 5

what happens when agents face non-zero absorbing boundaries and have no initial knowledge of the environmental edge or the extent of their wealth buffer. Hence, they have no direct knowledge of the risk of ruin or of the optimal stakes they should bet under their particular conditions. Instead, they must try to infer these through learning via direct experience as well as through social transmission. In particular, there are a number of key social cues that an adaptive risk-psychology can plausibly attend do: the success rates in a risk-free developmental learning period, the rate of ruin among peers, and the beliefs of successful members of the previous generation.

Learning strategies and wealth growth

We assume that agents must infer their environmental edge from individual experience and social learning. When the environmental edge is positive but close to 0, environments are very uncertain: gambles are expected to fail nearly as often as they succeed. The variance of outcomes is higher, and a "lucky streak" can lead to a highly biased estimate of the environmental edge. A higher edge means less uncertainty and, consequently, easier learning.

To incorporate learning throughout the lifetime, we assume agents update a belief distribution over the environmental success probability u, given by $\Gamma_i \sim \text{Beta}(a_i, b_i)$, which always starts as uniformly distributed, $\Gamma_i^0 \sim \text{Beta}(1, 1)$. In other words, we assume agents start their lives with maximal uncertainty about environmental states, and while we keep environmental edges in the favorable range ($\epsilon = u - 0.5 > 0$) and agent's estimate of the edge might become negative through unlucky sampling or pessimistic cultural influence, reflecting the development of pessimistic beliefs (see below). The Beta distribution permits iterated learning using straightforward Bayesian updating. As the agents incorporate information from different sources, they update their belief distribution by adding positive values to their Beta's parameters, where additions to a represent encountering information about successes and additions to b represent encountering information about failures. The precise way that information is incorporated into the posterior updates will depend on the learning strategies, described below.

We rely on numerical simulation throughout the remainder of the paper in order to explore the effects of learning on the stakes that agents choose. When an agent i uses a learning strategy L_i, we summarize their overall success by computing the average growth rate of their wealth achieved with that strategy. This is given by

$$V_T(L_i|\aleph, \epsilon) = \exp\left[\frac{1}{T+1}\sum_{t=0}^{T} g_t(L_i|\aleph, \epsilon)\right] \tag{14}$$

where the notation indicates that the agent's growth rate depends on the learning strategy they choose as well as their wealth buffer and the uncertainty in their environment. Because learning is a stochastic process, two agents employing the same learning strategy will not necessarily end up using the same stake after a finite number of gambles. We thus simulate N agents for every learning strategy implemented, and take the average geometric growth rate as an aggregated measure of the learning strategy's success:

$$\bar{V} = \frac{1}{N}\sum_{i=1}^{N} \bar{V}_T(L_i|\aleph, \epsilon) \tag{15}$$

An average growth rate of less than 1 indicates that agents using a particular learning strategy are, on average, experiencing losses and wealth shrinkage, while an average growth rate greater than 1 means that on average, agents' wealth is increasing. Below, we consider how learning throughout the lifespan can lead to the adoption of optimistic or pessimistic risk behaviors.

2.2 The development and cultural transmission of responses to risk and uncertainty

Here, we establish the full dynamic model for the development and cultural transmission of risk responses to consider the long-term effects of different learning strategies over generational time. This model involves the initial acquisition of risk perception during a juvenile period involving risk-free exploration of the environment (*individual learning*), information-sharing between generational peers coupled with socio-environmental feedback from observations of peer ruin during adulthood (*horizontal social learning*), and information inheritance from previous-cohort elders (*vertical and oblique social learning*). We use this model to also consider the effects of environmental change, socioeconomic stratification, and social learning biases. In the remainder of this section, we describe this model in detail. In the following section, we present the results of our investigative simulations.

Juvenile learning phase: acquiring risk behaviors using intra-generational learning

Consider a population of N individuals in a young generational cohort who all experience an initial wealth buffer $\aleph$, such that they have starting wealth $V_0 = \frac{1}{1-\aleph}$. Perhaps they are the children in a population of migrants who have entered mysterious new lands full of opportunity and, potentially, danger. For this first generation, let us assume that information from elders is unavailable or ignored. The young cohort must therefore interact with the environment to learn its ways. Subsequent generations will be able to learn from the knowledge acquired by their elders.

In their juvenile phase, agents can sample the environment without staking anything consequential. This risk-free sampling is meant to represent individual learning in early development: juveniles get to "play" in risk-free settings and use that information to perform in settings where stakes actually matter. Such low-risk play during childhood is common in humans and other complex social animals (Smaldino et al., 2019; Burghardt et al., 2024). The objective during this period is to estimate the environmental edge, ϵ. We divide the *juvenile learning phase* into two key processes: individual learning and peer influence.

1. **Individual learning**: Agents draw τ samples from the environment, representing independent and consequence-free experiences of would-be success or failure. That is, in each of τ time periods during their juvenile stage, agents throw a biased coin, where the probability of success is $u = \epsilon + \frac{1}{2}$. The value of τ represents the length of risk-free juvenile periods. Each agent i adds its observed number of successes w_i during its juvenile period to the a parameters of their belief distribution, and the number of failures $l_i = \tau - w_i$ to the b parameter. This results in a beliefs given by the random variable
$$\Gamma_i^I = \Gamma_i^{0'} \sim \text{Beta}(1 + w_i, 1 + l_i) \tag{16}$$
at the end of the juvenile learning period. This distribution is the agent's initial estimate of the environmental risk, u.
2. **Peer influence**: After individual learning is finished, each agent observes the behavior of n other agents chosen at random from within their cohort (henceforth their *peer learning set*), from which they further update their estimates of u. Such learning from peers is often called "horizontal" transmission in the cultural evolution literature (Cavalli-Sforza and Feldman, 1981). We assume agents incorporate information from their peer learning sets by considering the sum of their peers' environmental samples:
$$\Gamma_i^P = \Gamma_i^{I'} \sim \text{Beta}(a_i^P, b_i^P) \tag{17}$$
$$a_i^P = 1 + w_i + \beta_i \sum_{j \in P} w_j \tag{18}$$
$$b_i^P = 1 + l_i + \beta_i \sum_{j \in P} l_j \tag{19}$$
j indexes individuals in i's peer learning set P, and $\beta_i \in [0, 1]$ is the *weight of peer influence*, which is the extent to which an individual's belief distribution changes after sampling peers at the end of juvenile learning. This process thus reinforces the information encountered during risk-free individual learning and, in the case where there is no idiosyncratic uncertainty (that is, when all agents experience fully correlated environmental edges) it can also be seen as a way to expand the sample size of juvenile environmental exploration. Because we are interested in the evolution of risk-learning strategies, we allow β_i to evolve in our analyses.

After the juvenile learning phase, agents apply their risk behaviors to gambles that have consequences, during which time they can continue to update their risk behaviors. To simplify notation, we will refer to the beliefs about environmental risk at gambling period t as the random variable $\Gamma_{i,t}$ and the parameters of the (Beta) belief distribution it follows as $a_{i,t}$ and $b_{i,t}$. This means that in the absence of information from elder cohorts, the beliefs at the end of the juvenile period, Γ_i^P, are initially used to face the perils of adulthood, so that $\Gamma_{i,0} = \Gamma_i^P$.

Gambling phase: adjusting risk behaviors in response to observed ruin

At the end of their exploratory juvenile period, individuals venture into the world and use their beliefs about risk to make meaningful gambles. They play one gamble for each of T periods, all while keeping an eye on their peer learning set. Gambles are played by sampling the beliefs $\Gamma_{i,t}$ every time period t, obtaining an estimate of u (which we call $\hat{u}_{i,t}$), and passing this estimate through Equation 5 to determine their stake.

The multi-dimensionality of wealth as a social phenomenon and its dependence on variable currencies, each with their own context-dependent rules and dynamics, thus makes it difficult, if not downright impossible, for natural selection to effectively optimize for stable, hard-wired risk strategies. Similarly, children cannot typically infer optimal adult behavior as a direct result of their early ruin-free explorations. As societies change, wealth buffers can change too. Thus, learners can acquire knowledge about how buffered they really are only through direct experience as adults, through observing ruin in adult peers, or through information inherited from previous generations (which may be flawed if the individual's wealth buffer differs substantially from that of their elders).

We first analyze the case in which naive learners can learn only from the environment and peers. We then add the possibility of culturally inherited risk beliefs. We denote the stake used in period t as

$$\hat{s}_{i,t} = s^*(\hat{u}_{i,t}) \tag{20}$$

Hence, the observable risk behaviors of individuals with less concentrated belief distributions are assumed to correspond to greater variability in their chosen stakes from one time period to the next, the product of greater uncertainty about the true value of the environmental edge. Once the stake has been wagered, each agent performs an independent Bernoulli trial U_t with success probability u (that is, they draw a 1 with probability u and a 0 with probability $1 - u$) and their wealth is updated according to their success or failure. In the log-wealth scale, this is expressed as the following stochastic recursion:

$$g_{i,t+1} = \begin{cases} g_{i,t} + \log(1 + \hat{s}_i) & \text{when } U_t = 1 \\ g_{i,t} + \log(1 - \hat{s}_i) & \text{when } U_t = 0 \end{cases} \tag{21}$$

We allow agents to keep learning from their trials, such that they continue to update their beliefs $\Gamma_{i,t}$ in the same way as they did during juvenile individual learning, iteratively adding successes and failures throughout their lifetime of T periods, so that

$$\Gamma_{i,t+1} | U_t \sim \text{Beta}\Big(a_{i,t} + U_t,\ b_{i,t} + (1 - U_t)\Big) \tag{22}$$

If at any point in the T periods of the gambling process an agent's wealth falls to or below the absorbing boundary, the agent becomes ruined, their growth rate is set to 0 (g_i is set to $-\infty$), and they make no further gambles. Overly optimistic betting (i.e., overly large stakes) will often lead to ruin. In this case, continued peer learning can help agents avoid their peers' mistakes. In particular, if an agent observes a chosen cohort peer go to ruin, they may adjust their risk behaviors to be less optimistic.

Sensitivity to peer ruin: fear is the risk-killer

At the end of childhood, individuals are likely to have accurately estimated the environmental edge through a combination of individual and social learning. However, the childhood phase does not provide any information about the individual's wealth buffer, and agents following the Kelly criterion implicitly assume that there is no such risk. If their adult lives are heavily buffered by large wealth reserves, agents applying heuristic rules like the Kelly criterion will produce nearly optimal betting behavior. Conversely, if wealth buffers are small, the Kelly heuristic leads to over-betting even with extremely precise estimates of environmental risk. As adults, however, learners can observe ruin events among peers, and so plausibly benefit by incorporating information about how close they are to an absorbing boundary.

The fact that juvenile learning can easily lead to catastrophic overconfidence calls for a mechanism that will aid in pessimistic adjustment throughout the lifetime. Here we add to the model the possibility for peer learning in adulthood, such that agent i reduces their stakes when they observe peers go to ruin as a function of how similar their belief are to i's beliefs (which they take as a cue that their own risk behaviors are overly optimistic). We assign agents a *peer ruin sensitivity* $\delta \in [0, 1]$ (which we refer to as "sensitivity" for conciseness), which is the extent of their response to observing ruin among peers. We remind the reader that $a_{i,t}$ and $b_{i,t}$ represent the accumulated positive and negative beliefs (respectively) of agent i due to all information influences up to (adult) time period t. Let k be the count of peer ruin events i observed in time period t after gambling, and $b_{i,t} = b_{(i,t,k)}$ are the negative beliefs after observing these k peer ruin events at time period t, such that $b_{(i,t,0)}$ are the negative beliefs after gambling in period t but before observing any peer ruin event. Pessimistic adjustment thus behaves according to the following equation:

$$b_{i,t} = b_{(i,t,k)} = (1 + \delta_i)^k \cdot b_{(i,t,0)} \tag{23}$$

where i indexes the focal agent and sensitivity δ controls the strength of the agent's pessimistic response. As $\delta_i \to 1$, fully sensitive agents double their present failure count when they observe a ruin event. When $\delta \to 0$, the focal agent is not affected by observed ruin events. $\delta \in (0, 1)$ interpolates the cases between these two extremes, allowing for a full continuum of sensitivity to peer ruin and consequently for different strengths of social trauma response.

Figure 2 shows an example of how this learning and adjustment process looks like throughout the lifetime of three brothers: Juan, Gabriel and Roberto. All three end childhood optimistic, with beliefs shaped by their lack of elder influence and reliance on exploration and peer observation. If it wasn't for peer influence from their siblings, Gabriel would have entered adulthood fully pessimistic and Juan even more optimistic. As the brothers enter adulthood, Roberto experiences early ruin, prompting Juan to adjust his behavior toward greater pessimism. Gabriel, however, remains insufficiently sensitive and eventually faces ruin as well, prompting the more sensitive Juan to adjust once again to a degree of risk that secures his accumulated wealth in the last quarter of his life. Juan's mix of luck and adaptability enables him to reach the end of his life with a favorable outcome and having acquired adaptive pessimism, which he survives to pass on to the next generation (even though he started adulthood as the most dangerously optimistic brother). Gabriel, on the other hand, is doomed by the peer influence of his optimistic siblings during childhood and by his lack of sensitivity in adulthood. This emphasizes the dual importance of pessimism and sensitivity in navigating risk across the lifetime, with Roberto's and Gabriel's failures serving as a cautionary legacy for Juan's future offspring. A population-level view is shown in Figure 3, which plots the mean stakes of agents simulated coming out of juvenile learning and at the end of their lifetimes for a fixed value of sensitivity ($\delta = 0.1$). It shows how, as wealth buffers and environmental edges decrease, agents start their lives with more overly-optimistic risk behavior, and thus require higher degrees of pessimistic adjustment to get closer to the optimal stake. It also shows how high adversity can lead to over-adjustment, generating extreme levels of pessimism that aid in safety but minimize any future growth (see i.e. middle row, with $\aleph = 0.5$).

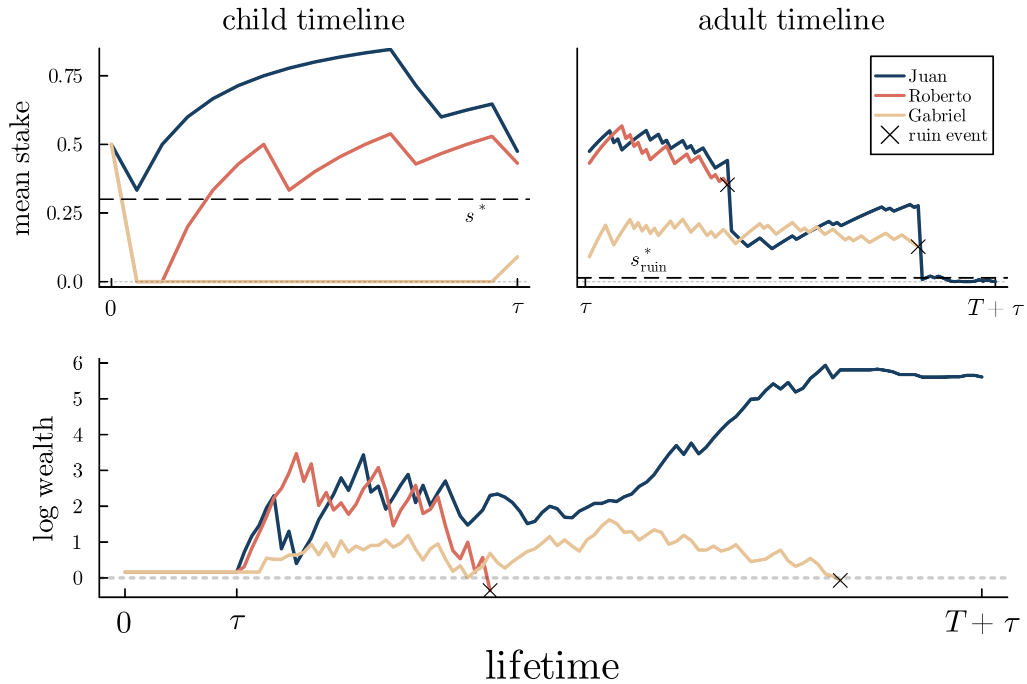

Figure 2: **A tale of three brothers.** The trajectories of three brothers, Juan, Roberto, and Gabriel, illustrate the interplay between optimism, sensitivity, and life outcomes under risk. 3 agents were simulated going through the juvenile learning phase and gambling phase with $\tau = 15$, $\epsilon = 0.15$, $\aleph = 0.15$ and $\beta = 0.25$ and sensitivities sampled from a uniform distribution between 0 and 1. They then experience $T = 100$ rounds of gambling, adjusting their beliefs accordingly. The ruin-free kelly stake from Equation 5 is plotted as a horizontal dashed line on the upper left plot, while the optimal stake given by Equation 12 is plotted on the top-right plot, and ruin events are marked with an X.

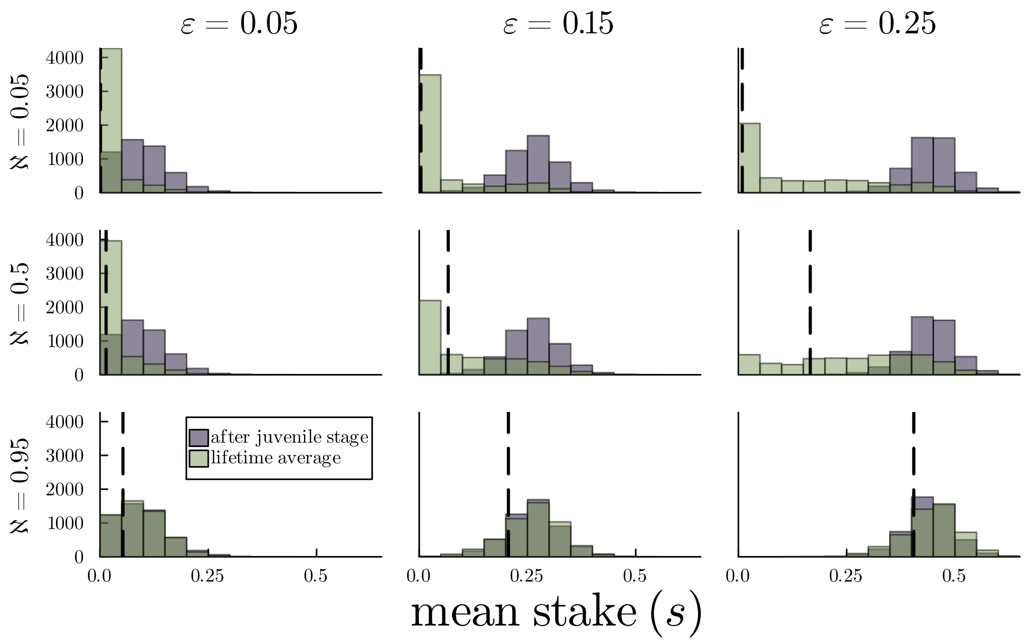

Figure 3: Distributions of mean stakes ($\bar{s}$) at the end of juvenile learning stages (gray) and at the end of gambling (green) when agents have a social trauma response given by their sensitivity to peer ruin. High risk of ruin scenarios (low edge, low wealth buffer) lead to a higher divergence between the distributions, indicating higher optimism at the end of juvenile learning and, as a consequence, a greater degree of acquired pessimism during the lifetime. $N = 10^4$ agents were simulated going through the juvenile learning phase with $\tau = 10$ and $\beta = 0.15$. They then experience $T = 500$ rounds of gambling while sampling their belief distribution and adjusting their beliefs according to a sensitivity of $\delta = 0.1$. The optimal stake given by Equation 12 is plotted as a vertical dashed line for each case.

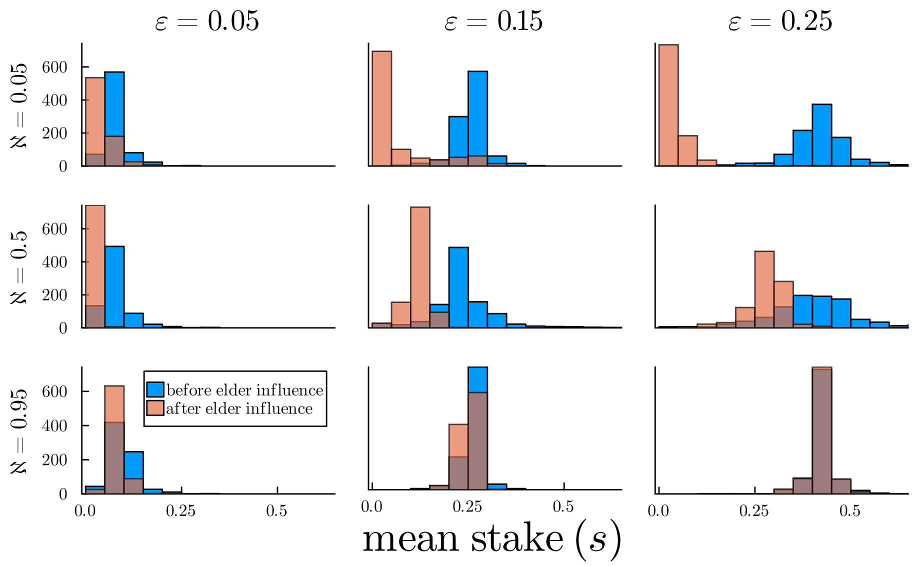

Figure 4: Distributions of mean stakes ($\bar{s}$) at the end of juvenile learning stages, before elder influence (blue) and after inheriting elder beliefs (orange). Populations of $N = 750$ agents were simulated with $T = 500$, $\tau = 10$ and $m = 10$, allowing learning strategies to evolve.

Cultural inheritance: using information from those who came before

We have so far ignored the value of information from previous generations and the possibility of culturally-inherited risk behavior through the inheritance of beliefs. Yet there are good reasons to expect such capacities to be adaptive. Elders from the previous generation have already experienced the adult environment, and have updated their risk preferences accordingly. When offspring are likely to face the same environmental uncertainty with the same wealth buffer as their parents, attending to elder risk preferences at the end of childhood could plausibly yield adaptive benefits. This is especially true for estimating one's wealth buffer, which is impossible for children learning only from peers and from their direct experiences. Let us now consider the scenario where agents have access to information from elders.

Imagine that our first generation of agents, which relied solely on individual and peer learning, has given way to the next generation. The old generation retires from productive activity, but they stick around long enough to give advice to the new generation. For simplicity, we assume individuals in this next generation inherit their parents' initial wealth buffer and undergo the same juvenile learning period described above. We then add the possibility of a new learning stage, *elder influence*, which we implement after juvenile learning but before the gambling stage. While peer influence is about the social reinforcement of immediate within-cohort juvenile experiences, elder influence represents the conservative force of traditionalism.

We implement a blending inheritance process in which the focal agent observes m randomly selected non-ruined agents from the previous generation (henceforth the *elder learning set* H_i), averaging their belief distribution parameters and blending these averages into her own belief distribution (with the strength of blending controlled by the weight of elder influence). Elder-influenced agents thus effectively incorporate a "conventional" previous-generation beliefs, in the sense that they inherit the population average with some sampling error. Note that, if $m = 1$, this is equivalent to the sort of unbiased learning rule used in many cultural evolutionary models. We assume that agents ignore elders who have gone to ruin, as their information may be deemed low-quality—this can be viewed as a sort of cultural viability selection. The elder learning set H_i is thus formed by sampling m elders from the population and "dropping" those with an average growth rate of 0. If, by chance, an agent ends up sampling only ruined agents, then no elder learning set is formed and no elder influence occurs.

At the end of their juvenile period, agents begin adulthood with estimates of the environmental edge drawn from two sources: intra-generational exploratory learning (individual and peer learning), and inter-generational conservative learning (elder influence). From generation 2 onwards, agents at the end of their juvenile phase begin adulthood with beliefs $\Gamma_{i,0}$ that weight these two influences:

$$\Gamma_{i,0} = \Gamma_i^{P'} \sim \text{Beta}(a_{i,0}, b_{(i,0,0)}) \tag{24}$$

where

$$a_{i,0} = a_i^P + \alpha_i \frac{1}{|H_i|} \sum_{h \in H_i} a_{h,T} \tag{25}$$

and

$$b_{(i,0,0)} = b_i^P + \alpha_i \frac{1}{|H_i|} \sum_{h \in H_i} b_{h,T} \tag{26}$$

where h indexes individuals in i's elder learning set H_i, $|H_i|$ is its cardinality (the number of elements in the set), and $\alpha \in [0, 1]$ is the *weight of elder influence*, which determines how much agents rely on intra- versus inter-generational learning. Both $a_{h,T}$ and $b_{h,T}$ represent elder h's accumulated lifetime environmental impressions. We refer to learning strategies that heavily rely on elder influence (α close to 1) as more *conservative*, as they conserve the knowledge accumulated by previous generations, and strategies that are less reliant on this information as more *exploratory.* A conservative learning strategy essentially conserves generational information ("tradition"), which can be appreciated by comparing Figure 4 with Figure 3.

The dangerous allure of payoff bias

Another form of demographic filtering can be implemented by introducing an explicit payoff bias to cultural learning. For a payoff-biased social learner, the higher the wealth of an observed elder, the higher the probability of adopting their risk behavior. While payoff-biased social learning can outperform unbiased learning if payoff is reliably associated with the adoption of adaptive behaviors (McElreath et al., 2013), payoff-biased learning can also be problematic for learning risk behaviors in uncertain environments. This is because those agents with the highest payoffs will tend to be those that made large gambles and got lucky, even though most individuals gambling with similarly optimistic stakes will experience ruin. By focusing on these rare successes, agents using payoff-biased social learning may adopt excessively optimistic risk behavior, exposing themselves to significant risks of ruin (Van Tilburg and Mahadevan, 2020; Baldini, 2012). On the other hand, in a socioeconomically stratified society, success can act as a *de facto* social marker for the wealthy classes, directing social learning to targets of similar wealth (Smaldino and Pérez Velilla, 2025).

We implemented payoff bias in elder influence by assuming that learners selectively learn from the agent in their elder learning set with the highest average growth rate, ignoring the rest. Symbolically, we denote by $\mathcal{L}_i \in \{\text{UB, PB}\}$ the *learning bias* of agent i, with UB being shorthand for *unbiased blending* and PB for *payoff bias.*

Environmental change, group structure and parochialism

In our baseline model, we have assumed that populations face static environments, such that environmental edges do not change across time and wealth buffers are homogeneous within populations. However, environmental change is a reality that most, if not all living populations contend with. In order to explore environmental change, we introduce two environmental uncertainty parameters: the probability of correlated change that affects the entire population, λ, which we call *aggregate uncertainty*, and the probability of individual environmental change, representing individual variation in environmental conditions, which we call *idiosyncratic uncertainty*, μ. At the beginning of every generation, a new environmental edge for the whole population is drawn from a Uniform$(0, 0.5)$ distribution with probability λ, otherwise it remains the same as in the previous generation. Each agent then samples a new local environmental edge from a Uniform$(0, 0.5)$ distribution with probability μ.

Many populations also exhibit class stratification and other forms of structural inequalities. Individuals may learn from elders within or across class lines, though they can also be likely to exhibit a preference for their own social class (Smaldino and Pérez Velilla, 2025). To model this preference, we add another behavior to the learning strategy: parochial social learning (also called similarity-biased social learning). We model class stratification by creating two classes of agents, each of which experiences a class-specific environmental edge and initial wealth buffer. Agents are assigned an observable class tag $\mathrm{C}_i \in \{c_0, c_1\}$, which are represented in the population with frequencies f and $1 - f$, respectively. We fix $f = 0.5$ for simplicity. We use c_0 to refer to the *advantaged group* (which has a higher $\aleph$ that is kept constant) and c_1 to the *disadvantaged group* (which has a lower $\aleph$ that is allowed to vary). Parochialism is then implemented by a binary variable $\mathcal{P}_i \in \{0, 1\}$, so that if $\mathcal{P}_i = 1$, an agent is parochial and drops any sampled outgroup members from its learning set. Note that this implies an endogenous cost to parochialism, as dropping learning models from learning sets leads to higher noise in the inherited acquired social information or can even pose the risk of empty learning sets (Smaldino and Pérez Velilla, 2025).

Cultural evolution and its consequences

The model simulates cultural evolution through an evolutionary algorithm in which learning strategies evolve based on fitness outcomes across generations. For each generation $\mathcal{G} \in 1, \ldots, \mathcal{G}_{\max}$, N agents are spawned. If $\mathcal{G} = 0$, then every agent i samples a learning strategies vector $\mathbf{L}_i = (\alpha_i, \beta_i, \delta_i, \mathcal{L}_i, \mathcal{P}_i)$ uniformly from their respective ranges or sets and is assigned a class tag C_i. If $\mathcal{G} > 0$, agents inherit their learning strategy vectors from previous-generation agents (see below) and sample a class tag without replacement from a previous-generation individual, such that the initial class tag distribution is kept constant throughout the run.

Generation $\mathcal{G}$ agents then go through the stages of juvenile learning and adult gambling described above, updating beliefs and payoffs accordingly. For every time period t in the adult gambling phase, non-ruined agents are scheduled at random to make a bet by taking a sample of their belief distribution and immediately update their beliefs depending on the outcome. After all agents have done their betting within the time period, non-ruined agents check for newly-ruined peers in their learning sets and do pessimistic adjustment according to their sensitivity to peer ruin. Only then does the population move to the next time period of the adult gambling stage.

At the end of the lifetime ($t = T$), the fitness of generation $\mathcal{G}$ agents is calculated:

$$\mathcal{F}_i = \exp\left(\frac{g_{i,T} - g_{i,0}}{T}\right)$$

and normalized into weights:

$$W_i = \frac{\mathcal{F}_i}{\sum_{j=1}^{N} \mathcal{F}_j}$$

for ever agent i. N generation $\mathcal{G}+1$ agents are then initialized, sampling their learning strategies (each component independently, using fitness weights) and class tags (uniformly, without replacement) from the now-inactive generation $\mathcal{G}$ agents. For all continuous learning strategy components $L_{i,j}$ in agent i's strategy vector, the agent independently samples a new continuous learning strategy uniformly from the $[0, 1]$ interval to replace the one inherited from the previous generation with probability m_d. For discrete learning strategies, a random learning strategy from the feasible set replaces the inherited one with probability m_d.

After the new generation goes through juvenile learning and updates their beliefs using elder information, generation $\mathcal{G}$ agents are retired.

Numerical results throughout the paper are obtained using scripts written in the Julia programming language (Bezanson et al., 2012), with individual-based simulations assisted by the Agents.jl package (Datseris et al., 2022). All simulation code and data, as well as a Pluto notebook (JE and CT, 2021) that can generate the plots in the paper can be found at `https://github.com/datadreamscorp/PessimisticLearning`. For all results, 25 randomly-seeded simulation runs were performed per parameter combination. Simulations were run for 2500 generations in order to obtain evolutionary equilibria. All fixed model parameters, along with their default values used in our simulations, are listed in Table 1.

Table 1: Default fixed parameter settings for evolutionary dynamics.

Population parameters	**Description**	**Default**
V_B	Absorbing boundary	1
N	Population size	750
T	Duration of gambling phase (time periods)	500
τ	Number of risk-free juvenile samples	10
n	Number of peers sampled during learning	10
m	Number of elders sampled (generation 2+)	10
f	Relative frequency of class c_0	1
Evolutionary parameters		
$\mathcal{G}_{\max}$	Number of generations	2500
m_c	Mutation rate for continuous strategies	0.001
σ_c	Standard deviation of Gaussian mutation noise	0.01
m_d	Mutation rate for discrete strategies	0.01

3 Analyses and results

To aid in understanding, we present the results of our analyses in a stepwise manner. We begin by considering populations with fixed wealth buffer and environmental edge in unchanging environments, for which we evaluate the types of learning strategies that evolve. We then consider time-varying environments, characterized by both aggregate (correlated) and idiosyncratic (uncorrelated) environmental change. Next, we consider stratified societies in which social classes can vary

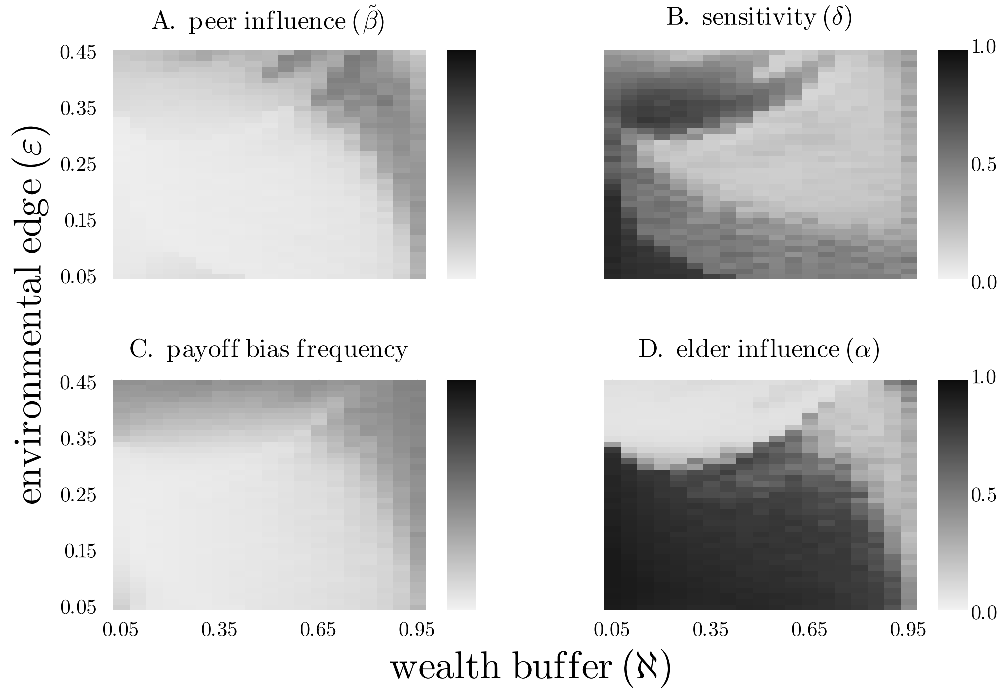

Figure 5: **Optimal learning strategies under static environments.** Median learning strategies after 2500 generations. Wealth buffers (ℵ) are varied from 0.05 to 0.95 in steps of 0.05, while environmental edges (ϵ) are varied from 0.55 to 0.95 in steps of 0.01.

in their wealth buffers, and examine the how learning strategies diverge, including in the parochial use of class markers. Finally, we turn our focus from social learning to the resulting patterns of risk behavior the model produces, including how deprivation leads to the coexistence of risk avoidance and extreme risk tolerance, how risk tolerance declines over the lifespan, and how environmental change can lead to the emergence of generational "poverty traps" among disadvantaged populations and social classes.

3.1 The evolution of learning strategies in static environments

We first consider the case where environments do not change, either globally or locally, and so $\lambda, \mu = 0$. Instead of reporting the raw peer influence weight β, we report the weighted peer influence $\hat{\beta} = (1 - \alpha)\beta$, in order to focus results not only on the intensity of peer influence but also of its overall contribution to beliefs.

As wealth buffers (ℵ) increase from low to high, elder influence (α) generally remains high across the range of wealth buffer, only declining significantly when environmental edges are high (Figure 5, D). Under low-to-moderate wealth buffers, relying on the accumulated knowledge of previous generations allows populations to avoid repeated missteps, stabilizing risk behaviors and preventing ruin. Under high wealth buffers, the value of elder influence diminishes, as younger generations can more easily discover successful strategies on their own without risking ruin.

When the environmental edge is low and wealth buffers are limited, populations are not helped by either peer influence ($\hat{\beta}$) or payoff-biased social learning (Figure 5, plots A and C), because peer influence leads to the "overfitting" of beliefs to juvenile environments (or, equivalently in the case of static environments, longer juvenile periods) and payoff bias often involves emulating previous-generation individuals who were merely lucky rather than genuinely well-adapted. By contrast, when wealth buffers and environmental edges are large, the cost of being wrong decreases, juvenile environments resemble adult environments more, and individuals are luckier throughout life. Under these more forgiving conditions, both peer influence and payoff bias become viable in low intensities/frequencies. In other words, strong wealth buffers and environmental edges transform peer influence and payoff bias from reckless strategies into lower-stakes gambles. It is important to note that these are the same conditions in which elder influence decreases, so the effects of payoff bias are modest and the effects of peer influence are stronger. Also, payoff bias is always a high-risk choice for learning about risk behavior, and it is fittingly never adopted by a majority of the population in wealth-homogenous societies (Figure 5, C).

At low wealth buffers but relatively high environmental edges, there is a sharp transition from high reliance on elder influence to almost none, with peer influence also remaining low, but sensitivity (δ) experiencing a precipitous comeback. In other words, agents rely on individual exploration while keeping a high sensitivity to peer ruin (Figure 5, B). When the high environmental edge favors gambling but individuals start close to the point of ruin, it is favorable to start life optimistic, while the high risk of ruin is attenuated by a high sensitivity.

While payoff bias evolves to remain relatively low under risky conditions, it increases slightly under the deprivation of very low wealth buffers and environmental edges ($\aleph, \epsilon << 1$), where high elder influence also makes its effects considerable among payoff-biased individuals. This is due to the fact that in these extremely impoverished conditions, some lucky individuals who adopted payoff bias and proceeded to use overly-optimistic high stakes did not experience ruin and thus achieved significantly higher growth rates than the median surviving individual, who played it safe. These reckless individuals in turn have a disproportionate, albeit still limited, cultural evolutionary contribution to the learning strategies of the next generation. While most of the contribution to the next generation's learning strategies will still come from individuals that played it safe and used non-payoff biased learning, the high payoffs of the very few surviving reckless individuals allow them to leave a cultural imprint on the following cohort. In these scenarios, most of the population will exhibit high risk avoidance, but there is a stable fat tail of payoff-biased, highly risk-tolerant agents (see Figure 8, B). In other words, in deprived populations, individuals who take extreme risks can perpetuate a maladaptive cycle of catastrophic risk tolerance, because a few of them get lucky and exhibit payoffs that are much higher than those of the median cautious individual, even if most individuals who take these risks end up ruined.

3.2 The evolution of learning strategies under environmental change

Real world environments change, and it may be difficult or impossible to know when or how often changes will occur. Here, we repeat the analyses of the previous subsection while allowing for environmental change. Based on well-known differential effects on social behavior when risks are correlated vs. uncorrelated among members of a population, we consider the separate contributions of aggregate uncertainty ($\lambda > 0$) and idiosyncratic uncertainty($\mu > 0$), as defined in the previous section. The results of these analyses are presented in Figure 6.

In environments with high wealth buffers, elder influence decreases due to elder knowledge becoming obsolete when the aggregate environment changes. However, we find that elder influence remains at high levels when wealth buffers are low, regardless of whether or how environments change (Figure 6, row B). This occurs because it is more important to avoid ruin during difficult periods than it is to grow during favorable times, and high elder influence allows for the uninterrupted

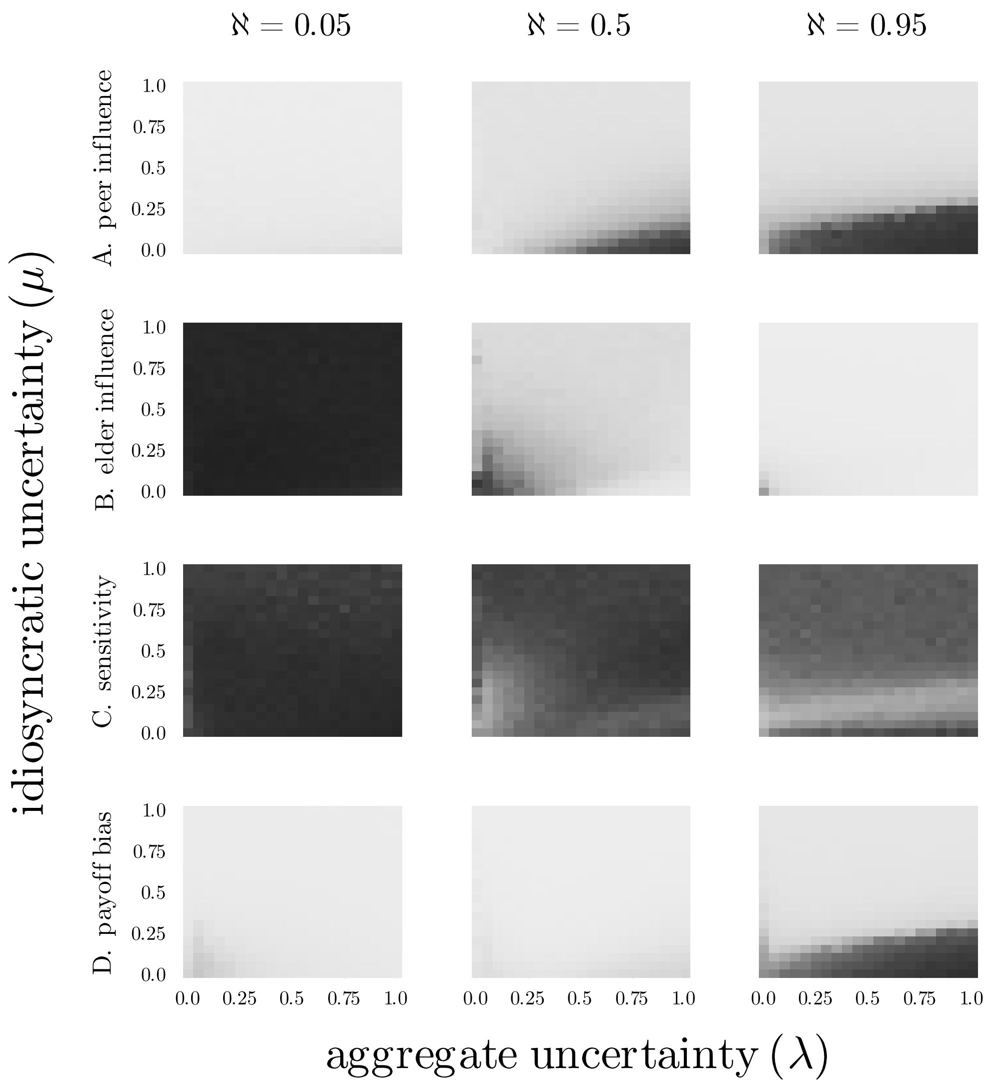

Figure 6: **The effects of environmental change on learning strategies.** The evolution of optimal learning strategies under aggregate (λ) and idiosyncratic (μ) uncertainty for low, intermediate and high values of $\aleph$. Median values for learning strategies are plotted as both μ and λ are varied from 0 to 1 in steps of 0.05. Initial environmental edges are sampled uniformly at random from the interval $(0.0, 0.5]$. The color scale goes from 0 (white) to 1 (black) for all strategies, in the same fashion as Figure 5.

transmission of risk avoidance throughout generations. As wealth buffers increase, elder influence can only remain high in an increasingly smaller region of environmental stability. This is consistent with previous theoretical results of Deffner and McElreath (2022), who studied an age-structured evolutionary model of social learning without risk, and found that learning from elders becomes favored as environments increase in temporal stability. Their results can thus be seen as equivalent to ours in the limit of no risk ($\aleph \to 1$).

Consistent with our results from static environments, low wealth buffers select against peer influence and payoff-biased learning. Peer influence can increase as wealth buffers and aggregate uncertainty increase, as peer influence aids in the rapid adaptation to new conditions. This occurs as long as environmental change is sufficiently correlated among individuals, otherwise the value of peer information diminishes. Payoff bias is only favored when wealth buffers are very high and increases with aggregate uncertainty (and again, its effects are muted by the low amounts of elder influence that evolve in these settings). At high enough idiosyncratic uncertainty, peer influence and payoff bias disappear, consistent with classic results showing that the evolution of social learning is not favored under idiosyncratic environmental change, such as that induced by high migration (Boyd and Richerson, 1985) (Figure 6, rows A and D).

Sensitivity to peer ruin remains high under most conditions, even when other forms of social learning are selected against (Figure 6, row C). While its strength is higher at lower wealth buffers, being sensitive to the losses of peers remains useful across the wealth buffer spectrum as a form of pessimistic social learning in the face of risk. At low wealth buffers, high elder influence accompanies high sensitivity, while at intermediate wealth buffers the conditions that favor high elder influence allow sensitivity to peer ruin to relax to lower levels. With large wealth buffers, sensitivity decreases the most when idiosyncratic uncertainty is present but mild enough so that only a minority of individuals deviate from aggregate environments. This is because reacting too intensely to the harsh conditions that only a minority might be experiencing experiences can hurt growth. However, as idiosyncratic uncertainty increases, sensitivity to peer ruin returns to high values, while all other forms of social learning decrease to minimal values. Under high idiosyncratic uncertainty, pessimistic social learning strategies like sensitivity to peer ruin can dominate by virtue of their general ability to promote safely risk avoidant behavior.

3.3 The evolution of parochial learning strategies under class inequality

To explore the effects of structural socioeconomic heterogeneity, we introduced class structure into the population. We operationalize this by dividing the population into two classes that differ in their wealth buffers, reflecting systematic inequalities in resource access and stability. We refer to the class with the higher wealth buffer as *advantaged*, and to the class with the lower wealth buffer as *disadvantaged*, reflecting their relative positions. This introduces another form of environmental heterogeneity, in the sense that individuals from different classes may be able to incur very different amounts risk, even if they face similar environmental edges.

We allow for the evolution of *parochialism*, whereby agents preferentially learn from their own class and ignore information from out-group members. Figure 7 shows the evolutionary outcomes for payoff-biased and parochial social learning, as well as the other traits related to learning, among each of the two classes in a stratified population. We limit our analysis to two class groups of equal size ($f = 0.5$). The red heatmaps on the left represent the optimal learning strategies of the advantaged group (c_0), while the blue heatmaps on the right represent those of the disadvantaged group (c_1). We explore learning differences between groups under both aggregate and idiosyncratic uncertainty, and plot results for a case of extreme inequality ($\aleph_{c_0} = 0.95$, $\aleph_{c_1} = 0.05$) to create stark contrast between conditions faced by each class. For these analyses, we expand population size to $N = 1000$

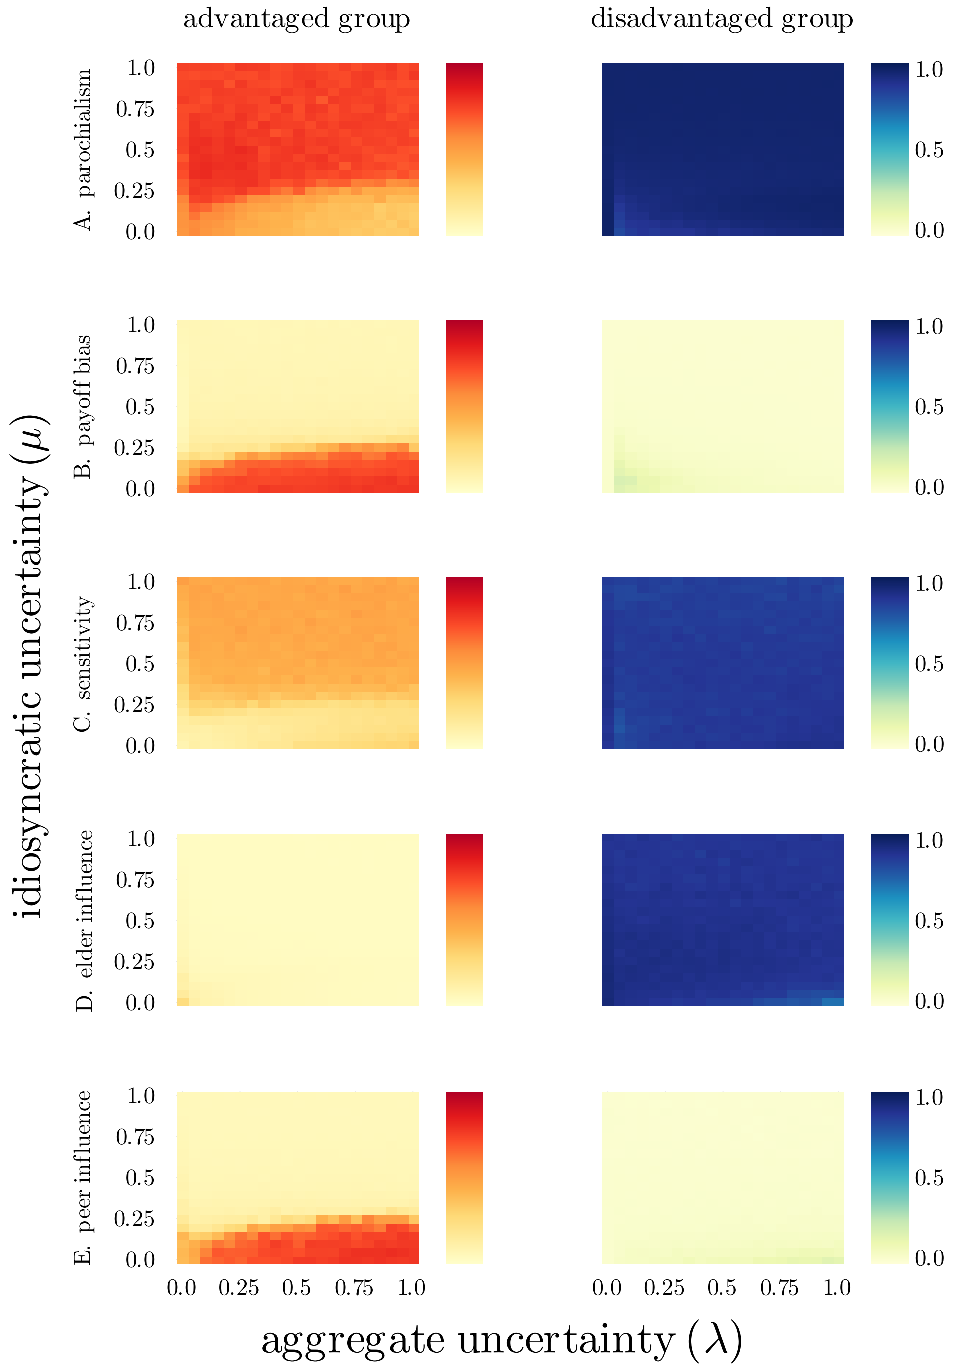

Figure 7: **Effects of class inequality and parochialism on learning strategies and risk behaviors.** Two classes, c_0 (red) and c_1 (blue), face distinct wealth buffers $\aleph_{c_0} = 0.95$ and $\aleph_{c_1} = 0.05$. Populations are initialized with $f = 0.5$. Each panel shows median evolved social learning strategy values for combinations of aggregate (λ) and idiosyncratic (μ) uncertainty values, which are varied from 0 to 1 in steps of 0.05. Initial environmental edges are sampled uniformly at random from the interval $(0, 0.5]$.

and use $n = m = 20$. In the SI we show that lowering the degree of inequality (by increasing the wealth buffer of the disadvantaged group) attenuates these patterns in a predictable manner.

Overall, the results for the evolution of peer influence, elder influence and sensitivity are coherent with our previous results for non-stratified societies. Disadvantaged groups are highly sensitive, rely primarily on elder influence, and eschew payoff bias and peer influence. In advantaged groups, we observe a severely reduced reliance on elder influence and find that sensitivity decreases while peer influence and payoff bias increase with aggregate uncertainty, but these effects are undermined at high levels of idiosyncratic uncertainty Figure 7, rows C, D and E).

Previous work has shown that reliance on parochial cues limits the integration of out-group information, which can be adaptive if behavioral strategies that adaptive for one group are suboptimal for another (Smaldino and Pérez Velilla, 2025). We recover this pattern in the current results, where the disadvantaged group evolves a high frequency of parochial social learning in order to avoid acquiring the beliefs and thus the riskier behavior of the advantaged group (Figure 7, row A, right). When idiosyncratic uncertainty is low, the advantaged group evolves only moderate levels of parochial social learning. This is due to two factors. First, there is an asymmetry regarding optimal stake under the risk-reward trade-off: overshooting it is much worse than undershooting it, so advantaged individuals who learn from disadvantaged individuals face less severe consequences than do disadvantaged individuals who learn from members of the advantaged class. Second, because they can risk more and thereby experience higher growth rates, individuals from advantaged groups will enjoy higher payoffs, which allows payoff bias to act as a *de facto* parochial learning strategy for the advantaged group. As idiosyncratic uncertainty increases, however, there is a sudden shift to much higher parochialism among the advantaged group (Figure 7, row A, left). Because both peer influence and payoff bias become obsolete as idiosyncratic uncertainty increases, sensitivity to peer ruin becomes the main strategy for narrowing down belief distributions and stabilizing stakes during the lifetime. This prompts a sudden increase in parochialism accompanied by a modest increase in sensitivity, as it becomes important for advantaged individuals to only react to the ruin of their fellow wealthy group members. Idiosyncratic uncertainty thus incentivizes members of an advantaged group to rely primarily on individual exploration and to be moderately sensitive only to the ruin of their own group members (Figure 7, rows A and C, left).

3.4 The effects of evolved learning strategies on risk behavior

So far we have focused on how risk and uncertainty affect the evolution of optimal learning strategies under different environmental and socioeconomic conditions. We now turn our attention towards the patterns of risk behavior that these learning strategies produce. We will show how our model can generate and explain the shape of risk behavior distributions under deprivation, the decrease of risk tolerance during the lifetime, and the emergence of "poverty traps" as cultural evolutionary mismatch under histories of high risk that incentivize high traditionalism.

Deprivation leads to a population-level pattern of both extreme risk avoidance and extreme risk tolerance

Under conditions of severe deprivation, where wealth buffers are minimal ($\aleph = 0.05$), populations evolve towards a distribution of risk behaviors that is highly skewed (Figure 8A). Most individuals converge on extremely low stakes, avoiding nearly all risk, as even a small gamble could lead to ruin. Yet, a minority of individuals—those who learn from individuals with incidental early successes or anomalous circumstances—adopt disproportionately high stakes, demonstrating extreme risk tolerance. This asymmetry emerges because under deprivation, conservative strategies dominate but

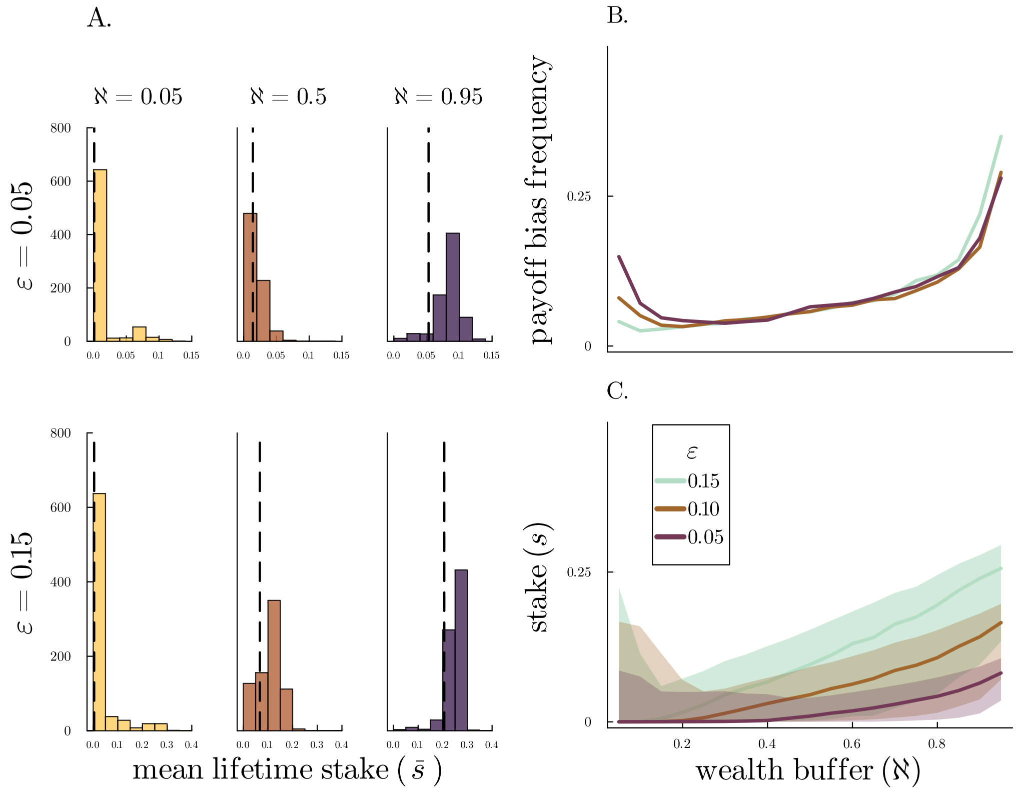

Figure 8: **Deprivation leads to skewed risk-attitude distributions.** The left column show distributions of mean lifetime stakes (s) at $T = 2500$ for different combinations of wealth buffers and environmental edges. Vertical dashed lines indicate the approximate Kelly stake given by Equation 12. On the bottom right, median lifetime stakes (solid lines) are plotted with a shaded region representing the area bounded by the 5% and 95% quantiles. Top right plots a non-monotonic pattern exhibited by payoff bias that at low wealth buffers drives the expansion of a long tail of risk takers. Plotted for static environments ($\lambda, \mu = 0$).

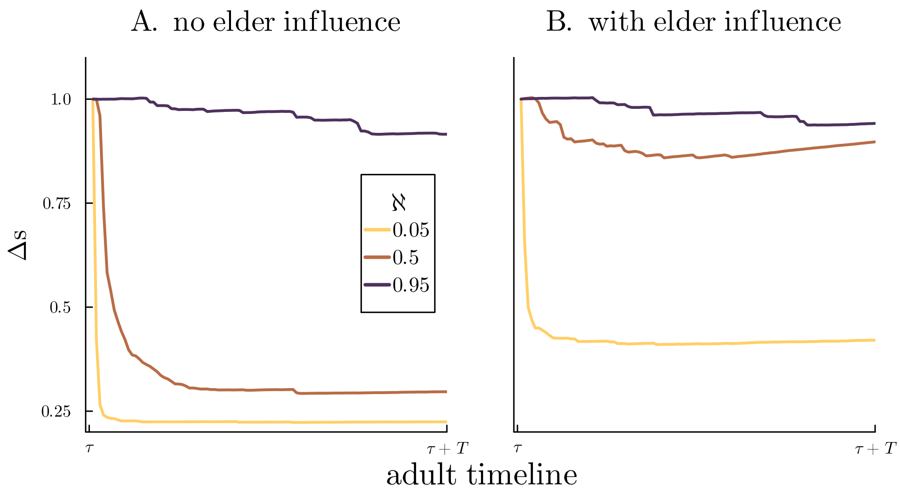

Figure 9: **Average proportional change in stake throughout adulthood.** Change is relative to initial levels at the end of juvenile periods under varying wealth buffers, with 1.0 indicating no change. Elder influence was either turned off ($\alpha = 0$; left) or allowed to evolve (right). Each line corresponds to a different wealth buffer ($\aleph$) and shows how the average stake changes as agents progress through their adult lifetime ($T = 500$ periods). Stake trajectories were averaged across $N = 750$ individuals per condition, and environmental edges were fixed at $\epsilon = 0.15$.

occasional outliers can still achieve and maintain high-risk behaviors if they survive early hazards, and their high growth rates can lead to cultural evolutionary contributions to the next generation when they occur in an environment dominated by small-growth risk-avoiders. The acquisition of risk tolerant behavior under low wealth buffers is afforded by payoff-biased learning, so it is the payoff-biased individuals that compose this long tail of catastrophic optimism.

While most of the individuals who acquire highly risky stakes through payoff bias go to ruin, those who do not earn payoffs that eclipse those of the median individual. This leads to disproportionate cultural evolutionary contributions to the next generation, which generates a significant fraction of payoff-biased individuals in the next generation who go on to repeat the maladaptive cycle of risk tolerance acquisition. The result is a stable population distribution composed primarily of extremely cautious risk-avoiders with a conspicuous tail of payoff-biased optimists.

Risk tolerance declines throughout the lifetime

Figure 9 illustrates how agents adjust their stakes over the course of their adult lifetimes for three different wealth buffer values. Each trajectory shows the average proportional change in stake relative to the starting level (Δs), providing insight into how agents dynamically modulate their risk-taking behavior as they gain experience. Stakes show a decreasing tendency throughout the lifetime, in a way that is mediated by both wealth buffers and the availability of elder information. Higher wealth buffers and availability of elder information attenuate the decrease in stakes, which has straightforward mechanistic explanation: under high wealth buffering, there is less risk of ruin overall and juvenile environments resemble adult environments more, making pessimistic adjustment less necessary. When elder information is present, individuals start adulthood already pessimistically-adjusted, which reduces the need for further adjustment.

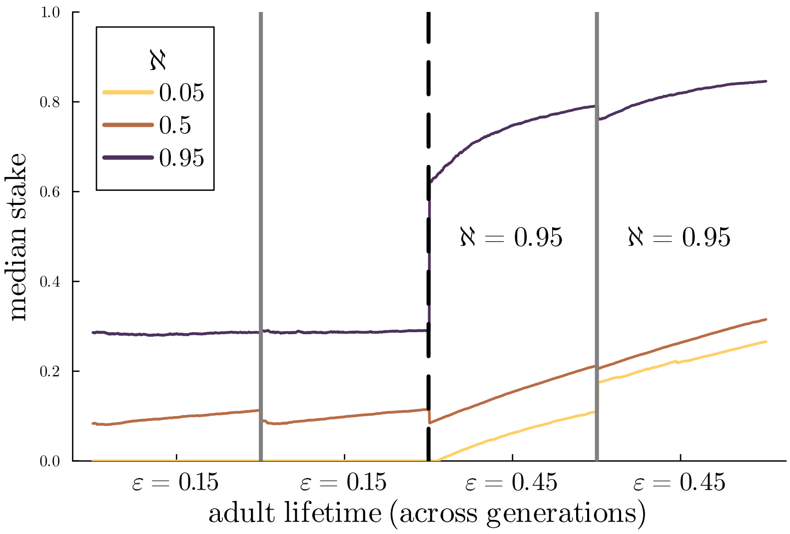

Figure 10: **Median stakes before and after environmental changes.** Median stakes across the lifetime in three populations that differ in wealth buffer, all with environmental edge $\epsilon = 0.15$. Vertical grey lines indicate the end of a generation and the start of a new one, while the vertical dashed line indicates an induced change in which all populations suddenly experience equally advantaged conditions ($\aleph = 0.95, \quad \epsilon = 0.45$). Runs were done by allowing populations to converge for 2500 generations with $\lambda = 0$ and $\mu = 0.1$ before fixing environmental edges to the values indicated in the figure.

Generational trauma and poverty traps as cultural evolutionary mismatch under histories of high risk

We have seen how high risk, particularly that which comes from having low wealth buffers, leaves little margin for error in terms of the stakes agents can adopt and still avoid ruin. In such environments, agents tend to develop highly conservative strategies that are effective in preventing catastrophic losses but simultaneously restrict opportunities for wealth accumulation and growth. This behavior, deeply ingrained through generations of navigating harsh conditions, creates a cultural legacy that can persist even when the environment becomes less perilous.

Figure 10 illustrates these dynamics. We simulated three different populations with varying wealth buffers and all with a fairly small environmental edge, and allowed each to evolve social learning strategies for risk behaviors attuned to their particular conditions. We then induced an environmental change that placed each population in an identical, highly favorable environment ($\aleph = 0.95, \quad \epsilon = 0.45$). Prior to the induced environmental change (marked by the vertical dashed line), populations with lower wealth buffers consistently display lower median stakes over their lifetimes compared to those with higher buffers. This disparity is not merely a consequence of immediate resource limitations but reflects the inherited risk-averse strategies that evolve under chronic high-risk conditions. After the shift to equalized conditions, conservative behaviors can persist in the previously-disadvantaged groups, as the cultural imprint of past risk continues to shape decision-making even after the environment has improved.

This persistence highlights a critical aspect of what we might term *cultural evolutionary mismatch* (cf. Nunn, 2022) and its relationship to what is usually referred to as *generational trauma*. The decision-making frameworks that once maximized survival under high-risk conditions lead to lagging adaptability when the potential rewards of taking higher risks increase. The slow recalibration of these cultural strategies may keep populations in a state of underinvestment and missed opportunities for generations. However, we stress that this mismatch is not a “trap” in the self-reinforcing sense that keeps a population in a suboptimal equilibrium, but rather an evolutionary lag in which the persistence of highly conservative learning strategies (including parochial learning in stratified populations) after population shocks effectively slows down the speed of adaptive cultural change. With enough time, populations are, *ceteris paribus*, expected to reach an adaptive equilibrium state.

4 Discussion

We developed a cultural evolutionary model to explore how risk behaviors in particular domains emerge from the interplay of risk-reward trade-offs, individual life experiences, cultural inheritance, and environmental conditions. Our framework highlights the importance of wealth buffers, environmental edges, and social learning strategies—encompassing juvenile exploration, peer influence, elder influence, payoff bias, and parochialism—in shaping the development and transmission of risk behaviors. By simulating how learning strategies evolve over multiple generations, we have generated insights into how cultural processes can generate variation in domain-specific risk behaviors within and between populations, as well as across socioeconomic strata.

4.1 Explanations and testable predictions

By studying this multi-level process of developmental and cultural adaptation, we obtain a unified explanation for observed differences in risk behaviors across developmental stages, socioeconomic strata, environmental conditions, and cultural contexts. In decision-making domains that involve risk-reward trade-offs and incomplete information, individuals tend to develop pessimistic risk beliefs at lower wealth buffers (high risk) which are obtained through high reliance on information from previous generations and observations of peer ruin, while optimistic risk beliefs are developed at high wealth buffers (low risk) which are facilitated by more exploratory learning strategies like individual and peer learning. This logic translates to within population dynamics in socioeconomically stratified societies, where cultures of caution, conservative learning patterns and risk avoidance, will tend to be promoted among at-risk subpopulations, while cultures of exploratory learning and risk tolerance will generally promoted among the wealthy and buffered. In cultural evolutionary terms, higher wealth buffering weakens selection for any single source of information, allowing individuals to be more “playful” with the way they learn about the world. And when environments change from generation to generation, these exploratory forms of learning can even provide an adaptive advantage.

Sensitivity to peer ruin, whereby individuals adjust their own risk behaviors when they observe the losses of others, evolves as a psychological mechanism for balancing risk-taking and caution by increasing pessimism throughout the lifetime, but can also leave agents prone to extreme risk avoidance, especially in high risk settings. Since a high sensitivity to peer ruin is favored in very adverse environments, individuals are able to adjust their beliefs and attitudes quickly, but also run the risk of responding too strongly to observed ruin, leaving an overly-pessimistic imprint that can be passed down to new generations. When environmental conditions change between generations, sensitivity to peer ruin remains adaptive across the risk spectrum. The success of riskier learning strategies like peer influence and payoff-biased learning depends on how useful social information is: they see support in environments that change in an aggregate manner (correlated across individuals) but they are undermined by idiosyncratic environmental changes (uncorrelated across individuals).

This leads to a life trajectory in which individuals in both deprived and wealthy conditions tend to be optimistic at the start of adulthood and become more pessimistic throughout their lifetimes. In deprived settings, juvenile optimistic risk behaviors can be adjusted through both strong elder influence and by sensitivity to peer ruin when observing peers go to ruin. In wealthy settings, individuals face much less pressure to become more risk avoidant, but still do so to some degree by reacting to observed peer ruin.

Under conditions of severe deprivation, populations can develop a highly-skewed risk culture, in which most individuals become highly conservative in their learning and pessimistic in their risk behaviors, while a small minority adopt a maladaptive payoff bias that fosters extreme optimism (i.e., high risk tolerance), driven by rare outlier successes that disproportionately influence cultural norms. This mirrors empirical findings of high proportions of risk-taking in marginalized groups despite widespread caution (Banerjee, 2004; de Courson et al., 2025), and is a candidate explanation for why extreme high-risk behaviors, such as criminal behavior, appear disproportionately (although still rarely) in materially-deprived populations. In these settings, while most people who learn from high-payoff individuals (and thus take large risks) face almost certain ruin, a very small number "luck out" and experience significant wealth growth. This growth can be massive relative to most non-ruined individuals in deprived settings, and provides a dangerous allure for future generations of learners, who observe a highly-skewed wealth distribution among their elders. As a result, while most individuals in the next generation choose the route of caution, a substantive minority can be influenced to adopt payoff-biased social learning, leading to recklessness.

Finally, when structural inequality is introduced in the form of group differences in wealth buffering, learning strategies diverge along class lines. Wealthy, better-buffered groups who adopt more exploratory learning (leading to greater risk tolerance), will also face incentives to adopt payoff bias in higher numbers as a parochial strategy, reinforcing their advantage. On the other hand, disadvantaged groups become even more reliant on conservative strategies and evolve more parochial learning cultures (in which information from outgroup individuals is actively disregarded). This divergence in learning strategies can drive the maintenance, and even further exacerbation, of inequality in both risk behaviors and wealth, leading to "poverty traps" that can persist for generations.

The main testable predictions of the model can be summarized as follows:

1. Individuals start adulthood with more optimistic risk attitudes (higher risk tolerance), which tend to become more pessimistic (risk-avoidant) over the lifespan. This shift results from accumulating personal experience and observing peer failures (sensitivity to peer ruin).
2. Individuals exposed to higher early-life adversity will exhibit:
   - Shorter periods of exploratory learning,
   - Greater reliance on inter-generational (conservative) learning, and
   - Greater sensitivity to peer ruin.
3. Individuals vary in how much they adjust their risk preferences after observing peer failures, and this "sensitivity to peer ruin" is predictive of long-term conservatism in risk-taking.
4. Exposure to repeated peer or community ruin leads to over-updating toward pessimistic beliefs, even when current conditions no longer warrant caution. This over-updating creates long-lasting conservatism (generational trauma) that can be passed culturally.
5. Greater wealth (or perceived buffering) increases openness to exploratory learning and risk tolerance; low wealth promotes conservative learning.
6. In environments of high inequality and high adversity, a small minority of individuals will be disproportionately influenced by highly successful (but rare) models, leading to excessive risk-taking.
7. Individuals from disadvantaged backgrounds will exhibit more parochial learning preferences (preference for in-group role models) due to the poor fit of out-group information.

8. Individuals who experience ruin or observe many ruin events become more resistant to updating even under improved conditions.

Conservative vs. exploratory learning in risky environments

Coming from a wealthy family and/or having access to robust societal safety nets allows individuals to take more risk, and thus to achieve higher lifetime growth, while keeping the risk of ruin low. Our analysis shows that when environments pose significant risks of ruin (due to low wealth buffers and/or uncertain payoffs), conservative learning strategies—characterized by high reliance on elder influence and low payoff bias—are generally favored. These conservative strategies curb overly-optimistic risk-taking and help prevent catastrophic losses. They transmit accumulated knowledge about environmental hazards across generations, producing cultural inertia that can stabilize risk-averse attitudes. Such conservatism is adaptive in the sense that it prioritizes survival and avoids ruin, especially when a single miscalculation can have irreversible consequences.

In particularly deprived conditions, individuals tend to start adulthood with optimistic risk behaviors that quickly crash after they witness their (similarly optimistic) peers experience ruin, leading to the development of high risk avoidance soon after exposure to real-life risks begins. However, groups under extreme deprivation also produce a small but stable proportion of extreme risk-takers, of which some (even fewer) lucky ones will enjoy great relative success, and who proceed to have significant cultural contributions to the next generation. These results mirror the empirical results of de Courson et al. (2025) and provide a mechanistic explanation for them: in poverty, those individuals who take extreme risks tend to fail spectacularly, but a few of them will attain payoffs that are so high relative to other individuals in the population that they manage to culturally influence a small but significant proportion of the next generation. In accordance with the suggestions of de Courson et al., this pattern can explain why extremely risky behaviors, such as property crime, are disproportionally observed in impoverished populations even though most individuals in this category are likely to be extremely risk averse (de Courson et al., 2025), as well as why these behaviors tend to occur more often in impoverished juveniles (Whitbeck and Simons, 1993).

Conversely, when wealth buffers are large and environmental conditions are more forgiving, agents can afford more exploratory and diversified learning strategies. Low-risk juvenile exploration and peer influence become more impactful, and the costs of mistakes diminish. Under these conditions, risk-taking becomes a viable strategy for achieving greater wealth growth, and agents exhibit a higher stability in their risk behaviors throughout their lifetime (Leonard and Sommerville, 2024). Elder influence remains useful but is likely to decline in importance as newer generations learn to navigate the environment effectively on their own. If environments change frequently, elder influence may become maladaptive at high wealth buffers while remaining useful for those at low wealth buffers due to how wealth constrains exploration. Moreover, payoff bias emerges as a "luxury of the rich," enabling agents in affluent conditions to imitate rare high performers, thereby amplifying success stories and fostering greater tolerance for risk, even when not widely adaptive. Elon Musk famously goes "all in" on every hand of Texas Hold 'Em poker (Isaacson, 2023). This strategy, described in glowing terms by economic elites and elite aspirants, works because Musk, due to his obscene wealth, can always buy more chips; emulating his strategy would be catastrophically foolish for most people.

It is important to clarify that our model predicts that risk tolerance increases with *perceived wealth buffering* rather than with absolute income or national wealth per se. This distinction is important when interpreting cross-national comparisons of risk preferences. Large-scale studies such as Rieger et al. (2015), Falk et al. (2018), and Vieider et al. (2019a) suggest that individuals from wealthier countries tend to be more risk averse, but these patterns are difficult to interpret due to issues of confounding. Both the Rieger et al. and Vieider et al. datasets are based on university students, who in poorer countries are more likely to be drawn from relatively privileged strata. The Falk et al. study, which uses more representative samples, finds no robust relationship between income and risk attitudes once broader heterogeneity is considered. In fact, within-country analyses do support a positive relationship between income and risk tolerance. For instance, Vieider et al. (2019b) find a negative correlation between risk aversion and income in Vietnam, while Vieider et al. (2018) report that in Ethiopia, higher income, job stability, and other local economic security indicators (proxies for the wealth buffer, in the paper treated as "income proxies") are associated with greater willingness to take risks—consistent with our core prediction that increased buffering enables exploratory behavior. Apparent contradictions between macro-level and micro-level patterns can also stem from how local environments, institutions, and life histories shape subjective perceptions of security and buffering (Weber and Hsee, 1998; Di Falco and Vieider, 2022). For example, households in agrarian economies may report low income while enjoying robust informal safety nets, whereas urban households in wealthier nations may face greater precarity due to weaker social ties despite higher income at the time of the study. Studies focusing solely on income—without accounting for these contextual factors—can therefore produce misleading conclusions. A more accurate approach to identifying country-level differences in risk behavior would be to compare the strength of the *within-population correlation* between wealth buffer proxies and risk preferences across countries. Measures of consumption smoothing capacity in response to shocks may serve as better proxies for effective wealth buffering than income alone or aggregate indicators such as GDP.

While the current model is not gender-specific, it offers a useful lens for interpreting observed gender differences in risk attitudes. If women, on average, perceive themselves to be less buffered against adverse outcomes—due to structural inequalities, life-history constraints, or differences in access to resources and social capital—then our model predicts they will favor more conservative learning strategies and develop greater risk aversion. This interpretation is consistent with extensive empirical findings showing lower risk tolerance among women (Eckel and Grossman, 2008; Borghans et al., 2009) and suggests that such patterns may emerge not from intrinsic preferences, but from adaptive responses to different developmental and environmental contexts (Schubert et al., 1999). Extending the model to incorporate gender-specific trajectories represents a potential avenue for future work.

Our results are also compatible with a growing empirical literature on the inter-generational transmission of risk behaviors and the role of socialization and investment in their formation (Dohmen et al., 2012; Zumbuehl et al., 2013). Since even a small degree of elder influence is adaptive in a wide variety of wealth buffer scenarios, our model implies that significant inter-generational transmission of risk behaviors is to be expected across a wide variety of human populations. But we also provide hypotheses regarding the strength of elder influence as wealth buffers change under different forms of environmental uncertainty, which can serve as explanations for observations of stronger elder influence in populations with lower wealth buffering (Sepahvand and Shahbazian, 2017; Wolff, 2020; Sepahvand and Shahbazian, 2021) relative to populations with higher wealth buffering (Dohmen et al., 2011; Shore, 2011; Arrondel, 2013; Necker and Voskort, 2014), or in recent immigrant groups in wealthy nations relative to the broader population (Mesoudi et al., 2016), providing possible directions for cross-cultural empirical testing.

Early adversity, generational trauma and the importance of (socially) learning from failure

Our model clarifies how cultural inheritance interacts with individual-level development and environmental adversity to produce changing risk behaviors over the lifespan (Amir and Jordan, 2017; Amir et al., 2018, 2020). Juveniles start off optimistic, benefiting from low-risk exploration and peer learning that shields them from the immediate consequences of over-betting. As they transition into adulthood, observing peers who face ruin leads to increased sensitivity and more cautious, pessimistic risk behavior—especially in harsher environments. This pattern is coherent with the idea that early adversity can accelerate the shift from exploratory to exploitative behavior (Frankenhuis and Gopnik, 2023; Gopnik, 2020) and aligns with empirical findings that childhood hardship can result in greater risk aversion later in life (Leonard and Sommerville, 2024; Malmendier and Nagel, 2011).

Our model shows that pessimistic adjustment can be adaptive at a population level, even if at an individual level it can lead to cases of over-adjustment and extreme risk avoidance. Sensitivity to peer ruin can be adaptive at the population level, allowing individuals to pull back overly-optimistic risk behaviors, but it can also lead to some agents becoming "traumatized" (extremely risk averse) due to unfortunate combinations of high sensitivity and witnessing too much ruin. Deprived populations and individuals in highly-adverse environments (such as warzones and areas experiencing ecological catastrophe) will be more prone to experiencing cases of trauma, as individuals in these conditions need to be highly sensitive to adjust their risk behaviors as quickly as possible. This makes sensitivity to peer ruin a possible evolutionary mechanism behind culturally-mediated psychological trauma responses, such as post-traumatic stress disorder and its social components (Kim and Lee, 2014; Huh et al., 2016; Ruderman et al., 2016; Zefferman and Matthew, 2021; Bayer and Shtudiner, 2024), akin to a sociocultural version of the "hot stove effect" (Denrell and March, 2001). Its robustness across simulated conditions of stability and environmental change in our model is due to the cliff-edged nature of the problem of choosing a risky stake combined with the stochasticity of possible outcomes from investment, effects that have been leveraged in evolutionary psychiatry to explain the prevalence of mood disorders such as depression (Nesse, 2015). Bet-hedging mechanisms that aid individuals to avoid falling off an "evolutionary cliff-edge" can also lead to disproportionately risk-averse responses, since avoiding the steeper side of the cliff-edge becomes a priority (Haaland et al., 2019). This comes at a sharp contrast with payoff bias, which is favored under a much narrower range of conditions. Under risk, learning from failure is better than learning from success.

The interplay of learning strategies and risk also helps explain cross-cultural variation in learning strategies (Mesoudi et al., 2016; Kline, 2015), as well as why certain groups seem to become trapped in risk-averse equilibria, as the pessimistic risk behavior acquired throughout the lifetime is passed on to successive generations. One or more generations experiencing a period of considerable adversity may lead to the rapid evolution of "poverty traps" through the cultural-evolutionary mismatch generated by the transmission of the acquired generational trauma (Nunn, 2022) borne from population-level shocks (Moya, 2018; Kim and Lee, 2014; Augsburger and Elbert, 2017) or long generational histories of high-risk conditions. When deprived of substantial buffers, populations risk frequent ruin whenever they try to act in exploratory ways, and repeated exposure to failures can entrench conservative, pessimistic beliefs, making it difficult to seize opportunities for growth even if conditions improve (McPeak and Barrett, 2001; Dercon and Christiaensen, 2011; Barrientos, 2013; Senadjki et al., 2017; Visser et al., 2020). By contrast, prosperous conditions that are stable for several generations permit the cultural evolution of flexible, exploratory strategies, allowing risk behaviors to adjust more readily to improving circumstances, which might feed into other benefit-inducing processes, including innovation.

The rise of socioeconomic inequality, parochialism, and class-based risk cultures

Introducing class inequality into our model shows how structural differences in wealth buffers and risk landscapes generate divergent learning strategies and risk cultures. Disadvantaged classes facing harsher conditions and more frequent ruin tend to evolve more reliance on information from elders, greater sensitivity to peer ruin, and strong parochialism as bulwarks against catastrophic optimism. Wealthy classes, operating in safer environments, can tolerate greater risk-taking and integrate more diverse sources of information, including using payoff-biased strategies to avoid learning from the poor when inequality becomes extreme. Over time, this divergence can produce stark differences in risk cultures between classes. While disadvantaged classes remain trapped in a conservative, low-stake equilibrium, the affluent classes can gamble with proportionally greater stakes, using more exploratory learning to adapt flexibly. These findings align with the notion that socioeconomic disparities can perpetuate through cultural channels, as wealthier groups accumulate and transmit behaviors that lead to continued growth while poorer groups remain stuck in risk-averse equilibria that limit upward mobility.

Previous work has shown that the Pareto distribution of wealth usually observed empirically is better explained by differential luck than by differential investment ability (Levy, 2003). Our results bolster the theoretical foundations for this view. Adding to this, our model shows that it is possible for these emergent inequalities to become exacerbated by the entrenchment of conservative learning strategies in disadvantaged groups, leading to an exploitation-first mindset that resists innovation. Advantaged groups who are well-buffered can be more exploratory in their learning, which may grant tangible advantages when it comes to the adoption of cultural innovations and thereby compound between-group inequalities. In other words, those who are initially lucky have not only the potential to obtain higher wealth, but also tend to reap the generational benefits of cultural exploration. The downtrodden, on the other hand, must be extremely cautious about their choices and the behaviors they adopt from others, limiting their exploratory options and locking them into tried-and-true solutions. In this latter group, those who put themselves at a significant risk have a small chance to grow their capital, which they exchange for a great chance of becoming ruined.

Connections to existing cultural and psychological frameworks

Our approach provides a plausible way to unify several lines of discussion about cultural variation across the social sciences—including cultural conservatism (Morin, 2022), tightness-looseness (Gelfand et al., 2011), and WEIRDness (Henrich et al., 2010; Mesoudi et al., 2016)—under a broad theoretical umbrella. Cultural differences that have been described as distinct phenomena could be reinterpreted as manifestations of the same underlying adaptive dynamics responding to risk and scarcity. In high-risk, resource-scarce contexts, conservative and parochial learning strategies can lead to the emergence of tight cultural norms as adaptive responses to mitigate ruin. In low-risk, resource-rich environments, exploratory and payoff-biased strategies that lead to loose norms and greater individualism may thrive because the consequences of failure are less dire. Moreover, these differences can occur within sectors of a single population due to the effects of class stratification.

Our framework also suggests that changes in social conditions that effectively alter wealth buffers can precipitate cultural shifts in risk tolerance through the incentives they induce on learning strategies, although the presence of large inequalities can slow down these shifts. If risk behaviors are indeed the product of the learning processes we have explored here, then interventions aiming to reduce risk and inequality will have to factor in the cultural inertia of evolved conservative learning strategies, especially for groups that have lived in conditions of deprivation and inequality for several generations. This cultural inertia of conservative learning can look like a "poverty trap" if one expects rapid changes in risk behaviors after the implementation of an intervention (Barrett and McPeak, 2006; Barrett and Carter, 2013; Barrett et al., 2016).

Conclusion

Our model highlights the importance of cultural evolutionary dynamics in shaping risk behaviors across diverse social and economic landscapes. By showing how individual development, cultural transmission, and environmental conditions interact to produce stable—but malleable—risk strategies, we provide a general and extendable framework for understanding human risk-taking behavior and its cultural variation over both developmental and generational time. In doing so, we offer new insights into how cultural inertia, social inequality, and environmental uncertainty can perpetuate or transform the risk behaviors that influence economic decisions, resource distribution, and social cohesion.

Author note

All simulation code, analysis scripts, and materials necessary to reproduce the results reported in this manuscript are openly available at `https://github.com/datadreamscorp/PessimisticLearning`. No empirical data were collected for this study.

An earlier version of this article was posted as a preprint in Pérez Velilla et al. (2025). The ideas and results reported here have not been published elsewhere.

Correspondence concerning this article should be addressed to Alejandro Pérez Velilla, Department of Cognitive and Information Sciences, University of California, Merced.
Email: aperezvelilla@ucmerced.edu.

References

Alger, I. (2023). Evolutionarily stable preferences. *Philosophical Transactions of the Royal Society B*, 378(1876):20210505.

Amir, D. and Jordan, M. R. (2017). The behavioral constellation of deprivation may be best understood as risk management. *Behavioral and Brain Sciences*.

Amir, D., Jordan, M. R., McAuliffe, K., Valeggia, C. R., Sugiyama, L. S., Bribiescas, R. G., Snodgrass, J. J., and Dunham, Y. (2020). The developmental origins of risk and time preferences across diverse societies. *Journal of Experimental Psychology: General*, 149(4):650.

Amir, D., Jordan, M. R., and Rand, D. G. (2018). An uncertainty management perspective on long-run impacts of adversity: The influence of childhood socioeconomic status on risk, time, and social preferences. *Journal of Experimental Social Psychology*, 79:217–226.

Arrondel, L. (2013). Are "daddy's boys" just as rich as daddy? the transmission of values between generations. *The Journal of Economic Inequality*, 11:439–471.

Asmussen, S. and Albrecher, H. (2010). *Ruin probabilities*, volume 14. World scientific.

Augsburger, M. and Elbert, T. (2017). When do traumatic experiences alter risk-taking behavior? a machine learning analysis of reports from refugees. *PLoS ONE*, 12(5):e0177617.

Baldini, R. (2012). Success-biased social learning: cultural and evolutionary dynamics. *Theoretical Population Biology*, 82(3):222–228.

Banerjee, A. (2004). The two poverties. *Insurance against poverty*, pages 59–75.

Barrett, C. B. and Carter, M. R. (2013). The economics of poverty traps and persistent poverty: Empirical and policy implications. *The Journal of Development Studies*, 49(7):976–990.

Barrett, C. B., Garg, T., and McBride, L. (2016). Well-being dynamics and poverty traps. *Annual Review of Resource Economics*, 8(1):303–327.

Barrett, C. B. and McPeak, J. G. (2006). Poverty traps and safety nets. *Poverty, Inequality and Development: Essays in Honor of Erik Thorbecke*, 1:131–154.

Barrientos, A. (2013). Does vulnerability create poverty traps? *Chronic Poverty: Concepts, Causes and Policy*, pages 85–111.

Bayer, Y. M. and Shtudiner, Z. (2024). Sirens of stress: Financial risk, time preferences, and post-traumatic stress disorder: Evidence from the israel-hamas conflict. *Journal of health psychology*, 29(13):1489–1502.

Bernoulli, D. (1954). Exposition of a new theory on the measurement. *Econometrica*, 22(1):23–36.

Bezanson, J., Karpinski, S., Shah, V. B., and Edelman, A. (2012). Julia: A fast dynamic language for technical computing. *arXiv preprint arXiv:1209.5145*.

Bollig, M. (1998). *Kinship, Networks and Exchange*, chapter 7, pages 137–157. Cambridge University Press.

Borgerhoff Mulder, M., Bowles, S., Hertz, T., Bell, A., Beise, J., Clark, G., Fazzio, I., Gurven, M., Hill, K., Hooper, P. L., et al. (2009). Intergenerational wealth transmission and the dynamics of inequality in small-scale societies. *Science*, 326(5953):682–688.

Borghans, L., Heckman, J. J., Golsteyn, B. H., and Meijers, H. (2009). Gender differences in risk aversion and ambiguity aversion. *Journal of the European Economic Association*, 7(2-3):649–658.

Bowles, S. (1998). Endogenous preferences: The cultural consequences of markets and other economic institutions. *Journal of Economic Literature*, 36(1):75–111.

Boyd, R. and Richerson, P. J. (1985). *Culture and the evolutionary process*. University of Chicago press.

Boyd, R., Richerson, P. J., and Henrich, J. (2011). The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences*, 108(supplement_2):10918–10925.

Brooks-Gunn, J. and Duncan, G. J. (1997). The effects of poverty on children. *The future of children*, pages 55–71.

Burghardt, G. M., Pellis, S. M., Schank, J. C., Smaldino, P. E., Vanderschuren, L. J., and Palagi, E. (2024). Animal play and evolution: Seven timely research questions about enigmatic phenomena. *Neuroscience & Biobehavioral Reviews*, 160:105617.

Cavalli-Sforza, L. L. and Feldman, M. W. (1981). *Cultural transmission and evolution: A quantitative approach*. Number 16. Princeton University Press.

Chudek, M. and Henrich, J. (2011). Culture–gene coevolution, norm-psychology and the emergence of human prosociality. *Trends in Cognitive Sciences*, 15(5):218–226.

Cieslik, K. and D'Aoust, O. (2018). Risky business? rural entrepreneurship in subsistence markets: evidence from burundi. *The European Journal of Development Research*, 30:693–717.

Datseris, G., Vahdati, A. R., and DuBois, T. C. (2022). Agents. jl: a performant and feature-full agent-based modeling software of minimal code complexity. *Simulation*, page 00375497211068820.

de Courson, B., Frankenhuis, W., and Nettle, D. (2025). Poverty is associated with both risk avoidance and risk taking: empirical evidence for the desperation threshold model from the uk and france. *Proc. R. Soc. B.*

Deffner, D. and McElreath, R. (2022). When does selection favor learning from the old? social learning in age-structured populations. *PloS ONE*, 17(4):e0267204.

Denrell, J. and March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science*, 12(5):523–538.

Dercon, S. and Christiaensen, L. (2011). Consumption risk, technology adoption and poverty traps: Evidence from ethiopia. *Journal of Development Economics*, 96(2):159–173.

Di Falco, S. and Vieider, F. M. (2022). Environmental adaptation of risk preferences. *The Economic Journal*, 132(648):2737–2766.

Dohmen, T., Falk, A., Golsteyn, B. H., Huffman, D., and Sunde, U. (2017). Risk attitudes across the life course. *Economic Journal*, (605):F95–F116.

Dohmen, T., Falk, A., Huffman, D., and Sunde, U. (2012). The intergenerational transmission of risk and trust attitudes. *The Review of Economic Studies*, 79(2):645–677.

Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J., and Wagner, G. G. (2011). Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the european economic association*, 9(3):522–550.

Eckel, C. C. and Grossman, P. J. (2008). Men, women and risk aversion: Experimental evidence. *Handbook of experimental economics results*, 1:1061–1073.

Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D., and Sunde, U. (2018). Global evidence on economic preferences. *The quarterly journal of economics*, 133(4):1645–1692.

Frankenhuis, W. E. and Gopnik, A. (2023). Early adversity and the development of explore–exploit tradeoffs. *Trends in Cognitive Sciences.*

Gelfand, M. J., Gavrilets, S., and Nunn, N. (2024). Norm dynamics: Interdisciplinary perspectives on social norm emergence, persistence, and change. *Annual Review of Psychology*, 75(1):341–378.

Gelfand, M. J., Raver, J. L., Nishii, L., Leslie, L. M., Lun, J., Lim, B. C., Duan, L., Almaliach, A., Ang, S., Arnadottir, J., et al. (2011). Differences between tight and loose cultures: A 33-nation study. *science*, 332(6033):1100–1104.

Gillespie, J. H. (1974). Natural selection for within-generation variance in offspring number. *Genetics*, 76(3):601–606.

Gopnik, A. (2020). Childhood as a solution to explore–exploit tensions. *Philosophical Transactions of the Royal Society B*, 375(1803):20190502.

Haaland, T. R., Wright, J., Tufto, J., and Ratikainen, I. I. (2019). Short-term insurance versus long-term bet-hedging strategies as adaptations to variable environments. *Evolution*, 73(2):145–157.

Haushofer, J. and Fehr, E. (2014). On the psychology of poverty. *Science*, 344(6186):862–867.

Henrich, J., Heine, S. J., and Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and brain sciences*, 33(2-3):61–83.

Hong, B., Kim, K., and Su, Y. (2024). The intergenerational transmission of risk preferences: Evidence from field experiments in china and korea. *Journal of Family and Economic Issues*, 45(1):151–173.

Huh, H. J., Baek, K., Kwon, J.-H., Jeong, J., and Chae, J.-H. (2016). Impact of childhood trauma and cognitive emotion regulation strategies on risk-aversive and loss-aversive patterns of decision-making in patients with depression. *Cognitive neuropsychiatry*, 21(6):447–461.

Isaacson, W. (2023). *Elon Musk*. Simon and Schuster.

JE, O. and CT, T. (2021). Reactive, reproducible, collaborative: computational notebooks evolve. *Nature*, 593.

Jones, J. H. and Ready, E. (2022). Subsistence risk-management networks. *SocArXiv. November*, 10.

Kelly, J. L. (1956). A new interpretation of information rate. *The Bell System Technical Journal*, 35(4):917–926.

Kendal, R. L., Boogert, N. J., Rendell, L., Laland, K. N., Webster, M., and Jones, P. L. (2018). Social learning strategies: Bridge-building between fields. *Trends in cognitive sciences*, 22(7):651–665.

Kim, Y.-I. and Lee, J. (2014). The long-run impact of a traumatic experience on risk aversion. *Journal of Economic Behavior & Organization*, 108:174–186.

Kline, M. A. (2015). How to learn about teaching: An evolutionary framework for the study of teaching behavior in humans and other animals. *Behavioral and Brain sciences*, 38:e31.

Knight, F. H. (1921). *Risk, uncertainty and profit*. Hart, Schaffner and Marx.

Kramer, K. L. (2005). Children's help and the pace of reproduction: cooperative breeding in humans. *Evolutionary Anthropology*, 14(6):224–237.

Kramer, K. L. (2021). Childhood teaching and learning among savanna pumé hunter-gatherers: Mismatch between foraging and postindustrial societies. *Human Nature*, 32(1):87–114.

Laland, K. N. (2004). Social learning strategies. *Animal Learning & Behavior*, 32(1):4–14.

Leonard, J. A. and Sommerville, J. A. (2024). A unified account of why optimism declines in childhood. *Nature Reviews Psychology*, pages 1–14.

Levy, H. and Levy, M. (2009). The safety first expected utility model: Experimental evidence and economic implications. *Journal of Banking & Finance*, 33(8):1494–1506.

Levy, M. (2003). Are rich people smarter? *Journal of Economic theory*, 110(1):42–64.

Levy, M. (2010). Evolution of risk aversion: The'having descendants forever'approach. *Available at SSRN 1688265*.

Levy, M. (2024). Relative risk aversion must be close to 1. *Annals of Operations Research*, pages 1–9.

Lew-Levy, S., Reckin, R., Kissler, S. M., Pretelli, I., Boyette, A. H., Crittenden, A. N., Hagen, R. V., Haas, R., Kramer, K. L., Koster, J., et al. (2022). Socioecology shapes child and adolescent time allocation in twelve hunter-gatherer and mixed-subsistence forager societies. *Scientific Reports*, 12(1):8054.

Mallpress, D. E., Fawcett, T. W., Houston, A. I., and McNamara, J. M. (2015). Risk attitudes in a changing environment: An evolutionary model of the fourfold pattern of risk preferences. *Psychological Review*, 122(2):364.

Malmendier, U. and Nagel, S. (2011). Depression babies: Do macroeconomic experiences affect risk taking? *The quarterly journal of economics*, 126(1):373–416.

McElreath, R., Wallin, A., and Fasolo, B. (2013). The evolutionary rationality of social learning. In Hertwig, R., Hoffrage, U., and ABC Research Group, editors, *Simple heuristics in a social world*, pages 381–408. Oxford University Press.

McPeak, J. G. and Barrett, C. B. (2001). Differential risk exposure and stochastic poverty traps among east african pastoralists. *American Journal of Agricultural Economics*, 83(3):674–679.

Meder, D., Rabe, F., Morville, T., Madsen, K. H., Koudahl, M. T., Dolan, R. J., Siebner, H. R., and Hulme, O. J. (2021). Ergodicity-breaking reveals time optimal decision making in humans. *PLoS Computational Biology*, 17(9):e1009217.

Mesoudi, A., Magid, K., and Hussain, D. (2016). How do people become weird? migration reveals the cultural transmission mechanisms underlying variation in psychological processes. *PloS one*, 11(1):e0147162.

Miller, P., Podvysotska, T., Betancur, L., and Votruba-Drzal, E. (2021). Wealth and child development: differences in associations by family income and developmental stage. *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 7(3):154–174.

Mills, C. and Zavaleta, D. (2015). Shame, humiliation and social isolation: Missing dimensions of poverty and suffering analysis. *World suffering and quality of life*, pages 251–266.

Morin, O. (2022). Cultural conservatism. *Journal of Cognition and Culture*, 22(5):406–420.

Moya, A. (2018). Violence, psychological trauma, and risk attitudes: Evidence from victims of violence in colombia. *Journal of Development Economics*, 131:15–27.

Necker, S. and Voskort, A. (2014). Intergenerational transmission of risk attitudes–a revealed preference approach. *European Economic Review*, 65:66–89.

Nesse, R. M. (2015). Evolutionary psychology and mental health. *The handbook of evolutionary psychology*, pages 903–927.

Nettle, D. and Frankenhuis, W. E. (2020). Life-history theory in psychology and evolutionary biology: one research programme or two? *Philosophical Transactions of the Royal Society B*, 375(1803):20190490.

Nunn, N. (2022). On the dynamics of human behavior: the past, present, and future of culture, conflict, and cooperation. In *AEA Papers and Proceedings*, volume 112, pages 15–37. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Paulsen, D. J., Platt, M. L., Huettel, S. A., and Brannon, E. M. (2012). From risk-seeking to risk-averse: the development of economic risk preference from childhood to adulthood. *Frontiers in Psychology*, 3:313.

Pérez Velilla, A., Beheim, B., and Smaldino, P. E. (2025). The development of risk attitudes and their cultural transmission. *SoArXiv*.

Poundstone, W. (2010). *Fortune's formula: The untold story of the scientific betting system that beat the casinos and Wall Street*. Hill and Wang.

Ready, E. and Power, E. A. (2018). Why wage earners hunt: food sharing, social structure, and influence in an arctic mixed economy. *Current Anthropology*, 59(1):74–97.

Reyes-García, V., Gallois, S., and Demps, K. (2016). A multistage learning model for cultural transmission: Evidence from three indigenous societies. *Social learning and innovation in contemporary hunter-gatherers: Evolutionary and ethnographic perspectives*, pages 47–60.

Rieger, M. O., Wang, M., and Hens, T. (2015). Risk preferences around the world. *Management Science*, 61(3):637–648.

Robson, A. and Samuelson, L. (2022). The evolution of risk attitudes with fertility thresholds. *Journal of Economic Theory*, 205:105552.

Robson, A. J. (1996). A biological basis for expected and non-expected utility. *Journal of economic theory*, 68(2):397–424.

Robson, A. J. and Orr, H. A. (2021). Evolved attitudes to risk and the demand for equity. *Proceedings of the National Academy of Sciences*, 118(26):e2015569118.

Robson, A. J. and Samuelson, L. (2019). Evolved attitudes to idiosyncratic and aggregate risk in age-structured populations. *Journal of Economic Theory*, 181:44–81.

Rodems, R. and Pfeffer, F. T. (2021). Avoiding material hardship: The buffer function of wealth. *Journal of European Social Policy*, 31(5):517–532.

Rogers, A. R. (1988). Does biology constrain culture? *American Anthropologist*, 90(4):819–831.

Ruderman, L., Ehrlich, D. B., Roy, A., Pietrzak, R. H., Harpaz-Rotem, I., and Levy, I. (2016). Posttraumatic stress symptoms and aversion to ambiguous losses in combat veterans. *Depression and anxiety*, 33(7):606–613.

Schreiber, S. J. (2015). Unifying within-and between-generation bet-hedging theories: an ode to jh gillespie. *The American Naturalist*, 186(6):792–796.

Schubert, R., Brown, M., Gysler, M., and Brachinger, H. W. (1999). Financial decision-making: are women really more risk averse? *American economic review*, 89(2):381–385.

Senadjki, A., Mohd, S., Bahari, Z., and Hamat, A. F. C. (2017). Assets, risks and vulnerability to poverty traps: A study of northern region of malaysia. *The Journal of Asian Finance, Economics and Business*, 4(4):5–15.

Sepahvand, M. H. and Shahbazian, R. (2017). Intergenerational transmission of risk attitudes: The role of gender, parents and grandparents in burkina faso. Technical report, Working Paper.

Sepahvand, M. H. and Shahbazian, R. (2021). Intergenerational transmission of risk attitudes in burkina faso. *Empirical economics*, 61(1):503–527.

Shore, S. H. (2011). The intergenerational transmission of income volatility: Is riskiness inherited? *Journal of Business & Economic Statistics*, 29(3):372–381.

Smaldino, P. E., Palagi, E., Burghardt, G. M., and Pellis, S. M. (2019). The evolution of two types of play. *Behavioral Ecology*, 30(5):1388–1397.

Smaldino, P. E. and Pérez Velilla, A. (2025). The evolution of similarity-biased social learning. *Evolutionary Human Sciences*, 7:e4.

Smith, D. M., Langa, K. M., Kabeto, M. U., and Ubel, P. A. (2005). Health, wealth, and happiness: Financial resources buffer subjective well-being after the onset of a disability. *Psychological science*, 16(9):663–666.

Starrfelt, J. and Kokko, H. (2012). Bet-hedging—a triple trade-off between means, variances and correlations. *Biological Reviews*, 87(3):742–755.

Strathern, A. (1971). *The rope of Moka: Big-men and ceremonial exchange in Mount Hagen New Guinea*. Number 4. CUP Archive.

Strauss, E. D. and Shizuka, D. (2022). The ecology of wealth inequality in animal societies. *Proceedings of the Royal Society B*, 289(1974):20220500.

Turner, M. A., Moya, C., Smaldino, P. E., and Jones, J. H. (2023). The form of uncertainty affects selection for social learning. *Evolutionary Human Sciences*, 5:e20.

Van Tilburg, W. A. and Mahadevan, N. (2020). When imitating successful others fails: Accidentally successful exemplars inspire risky decisions and can hamper performance. *Quarterly Journal of Experimental Psychology*, 73(6):941–956.

Vercken, E., Wellenreuther, M., Svensson, E. I., and Mauroy, B. (2012). Don't fall off the adaptation cliff: when asymmetrical fitness selects for suboptimal traits. *PLoS One*, 7(4):e34889.

Vieider, F. M., Beyene, A., Bluffstone, R., Dissanayake, S., Gebreegziabher, Z., Martinsson, P., and Mekonnen, A. (2018). Measuring risk preferences in rural ethiopia. *Economic Development and Cultural Change*, 66(3):417–446.

Vieider, F. M. et al. (2019a). All over the map: A worldwide comparison of risk preferences. *Quantitative Economics*, 10(1):185–215.

Vieider, F. M., Martinsson, P., Nam, P. K., and Truong, N. (2019b). Risk preferences and development revisited. *Theory and Decision*, 86(1):1–21.

Visser, M., Jumare, H., and Brick, K. (2020). Risk preferences and poverty traps in the uptake of credit and insurance amongst small-scale farmers in south africa. *Journal of Economic Behavior & Organization*, 180:826–836.

Weber, E. U. and Hsee, C. (1998). Cross-cultural differences in risk perception, but cross-cultural similarities in attitudes towards perceived risk. *Management science*, 44(9):1205–1217.

Weber, E. U. and Hsee, C. K. (1999). Models and mosaics: Investigating cross-cultural differences in risk perception and risk preference. *Psychonomic Bulletin & Review*, 6(4):611–617.

Whitbeck, L. B. and Simons, R. L. (1993). A comparison of adaptive strategies and patterns of victimization among homeless adolescents and adults.

Wolff, F.-C. (2020). The intergenerational transmission of risk attitudes: Evidence from burkina faso. *Review of Economics of the Household*, 18(1):181–206.

Xu, Y., Harms, M. B., Green, C. S., Wilson, R. C., and Pollak, S. D. (2023). Childhood unpredictability and the development of exploration. *Proceedings of the National Academy of Sciences*, 120(49):e2303869120.

Zefferman, M. R. and Matthew, S. (2021). Combat stress in a small-scale society suggests divergent evolutionary roots for posttraumatic stress disorder symptoms. *Proceedings of the National Academy of Sciences*, 118(15):e2020430118.

Zhang, R., Brennan, T. J., and Lo, A. W. (2014). The origin of risk aversion. *Proceedings of the National Academy of Sciences*, 111(50):17777–17782.

Zumbuehl, M., Dohmen, T., and Pfann, G. A. (2013). Parental investment and the intergenerational transmission of economic preferences and attitudes. Technical report, IZA Discussion Papers.