Evolutionary Foundation for Heterogeneity in Risk Aversion

Yuval Heller and Ilan Nehama

Abstract

We examine the evolutionary basis for risk aversion with respect to aggregate risk. We study populations in which agents face choices between alternatives with different levels of aggregate risk. We show that the choices that maximize the long-run growth rate are induced by a heterogeneous population in which the least and most risk-averse agents are indifferent between facing an aggregate risk and obtaining its linear and harmonic mean for sure, respectively. Moreover, approximately optimal behavior can be induced by a simple distribution according to which all agents have constant relative risk aversion, and the coefficient of relative risk aversion is uniformly distributed between zero and two.
Keywords: Evolution of preferences, risk interdependence, long-run growth rate.
JEL Classification: D81.

1 Introduction

Our understanding of risk attitudes can be sharpened by considering their evolutionary basis in situations in which agents face choices that affect their number of offspring (see 28, for a survey). Lewontin and Cohen [17] shows that idiosyncratic risk (independent across individuals) induces a higher expected long-run growth rate (henceforth, abbreviated as *growth rate*) than aggregate risk (correlated across individuals) with the same marginal distribution (i.e., for each individual the distribution of offspring induced is the same). Specifically, the growth rate induced by an idiosyncratic-risk lottery is equal to its arithmetic mean, while the growth rate induced by an aggregate-risk lottery is equal to its geometric mean.

In his seminal paper, Robson [24] shows that when a population faces a choice between several aggregate-risk lotteries, the optimal growth rate can be achieved by choosing the mixture of these lotteries that maximizes the expected logarithm of the number of offspring.[1,2] For example, consider a choice between an alternative that yields 2 offspring

Appeared in *Journal of Economic Theory*, Volume 208, 2023. [15]

[1] In this paper, we focus on characterizing the type that maximizes the growth rate (and do not explicitly model the underlying dynamics). This focus is motivated by the following sketched argument (which is an adaptation of Fisher's Fundamental Theorem; see 30, Section 3.6 for a textbook exposition). Consider a large population in which various groups of agents have different heritable types, where the type determines how the agent chooses between lotteries over the number of offspring (importantly, the number of the agent's offspring is independent of the types' frequencies in the population). Occasionally, new types are introduced into the population following a genetic mutation. Observe that the population share of agents of the type that induces the highest growth rate grows until, in the long run, almost all agents are of this type (see, e.g., 28, for a detailed argument of why natural selection induces agents to have types that maximize the growth rate).

[2] See also Robson and Samuelson [27] and Netzer [20] who study the evolution of risk attitudes and their impact on time preferences, Heller [14] who argues that the evolution of risk attitudes induces overconfidence, Robatto and Szentes [23] who study choices that influence fertility rate in continuous time, Robson and Samuelson [29] who explore age-structured populations, Netzer et al. [21] who argue that constrained optimal perception affects people's risk attitudes and induces probability weighting, Robson and Orr [26] who study the relation between aggregate risk and the equity premium, and Heller and Robson [16] who analyze heritable risk, which is correlated between an agent and her offspring.

to each agent for sure (the *safe* alternative) or a *risky* alternative yielding either 1 offspring or 5 offspring to all agents who choose it, with equal probabilities. One can show that the expected logarithm of the number of offspring is maximized (and the optimal growth rate is achieved) by having $1/3$ of the agents choosing the safe alternative, and the remaining $2/3$ of the agents choosing the risky alternative. This mixture yields an average number of offspring in the population that is either 4 or $4/3$ in each generation. This heterogeneity in the choices of the agents can be interpreted as *bet-hedging* the expected number of offspring in the population (4, 5, 1). The existing literature typically does not specify how the optimal bet-hedging mixtures are implemented.[3]

Note that each agent faces a non-convex choice between alternatives, and thus she cannot hedge her own personal risk (for example, there is no choice that yields her either 4 or $4/3$ offspring in the above example). One way to implement the optimal level of bet-hedging is to induce each agent to randomly choose her alternative according to the probability that maximizes the population's growth rate. Observe that this behavior is complicated in the sense that it is not induced by a von Neumann–Morgenstern (vNM) utility. For example, in the above scenario she strictly prefers choosing a lottery to either of the two alternatives. Hence, when facing a choice between two alternatives, she needs to evaluate her utility also from various lotteries over them.

In this paper, we introduce a new mechanism for implementing the optimal growth rate. We show that the optimal growth rate is induced by a heterogeneous population of utility maximizers in which agents have different levels of risk aversion with respect to the aggregate risk (and all agents are risk neutral with respect to idiosyncratic risk). In this population, the most risk-averse agent is indifferent between obtaining a risky lottery $\boldsymbol{y}$ and obtaining the harmonic mean of $\boldsymbol{y}$ for sure, while the least risk-averse agent is risk neutral, that is, indifferent between obtaining the risky lottery $\boldsymbol{y}$ and obtaining the arithmetic mean of $\boldsymbol{y}$ for sure. Moreover, we show that a nearly-optimal growth rate can be achieved by a simple distribution of vNM utilities, according to which all agents have constant relative risk aversion, and the risk coefficient is uniformly distributed between zero and two.

Highlights of the Model We consider a continuum population with asexual reproduction. Each agent lives for a single generation, during which she faces a choice between two lotteries over the number of offspring: a *safe* (degenerate) *alternative* that gives μ offspring for sure, and a *risky alternative* $\boldsymbol{y}$ with aggregate risk.[4] Nature induces a distribution of risk preferences with regard to aggregate risk, according to which each agent in the population (deterministically) chooses between the risky alternative and the safe alternative.

Main Results If nature were limited to endowing all agents with the same preference, then it would be optimal for all agents to evaluate any risky alternative $\boldsymbol{y}$ as having a certainty equivalent of its geometric mean. However, heterogeneous populations can induce a substantially higher growth rate, because heterogeneity in risk aversion allows the population to hedge the aggregate risk by enabling scenarios in which only a portion of the population (the less risk-averse agents) chooses the risky alternative.

Theorem 1 characterizes the optimal share of agents that choose the risky alternative, that is, the share that maximizes the long-run growth rate. In particular, it shows that all

[3] McNamara [19] shows that the optimal bet-hedging can be implemented when the agents maximize their relative fitness (namely, the ratio between an agent's number of offspring and the total number of offspring in her generation (see also 10, 6, 22). However, this maximization requires the agents to know the aggregate behavior in the population, which was, arguably, implausible in most of our evolutionary past.

[4] A restriction of the present analysis is that it assumes that in any generation there exists a single risky alternative (as discussed in Section 5). We leave for future research the important question of how to extend our analysis to a more general setup with multiple (non-degenerate) risky alternatives.

agents should choose the safe option iff $\mathbb{E}[\boldsymbol{y}] \leqslant \mu$, and all agents should choose the risky option iff the harmonic mean of $\boldsymbol{y}$, $\mathbb{HM}[\boldsymbol{y}]$, satisfies $\mathbb{HM}[\boldsymbol{y}] \geqslant \mu$. Moreover, we characterize the optimal preference distribution (Theorem 2) as follows: (1) the least risk-averse agent in the population is risk neutral, (2) the most risk-averse agent in the population has harmonic utility, i.e., she evaluates any risky alternative $\boldsymbol{y}$ as having a certainty equivalent of its harmonic mean, and (3) all agents in the population should be risk averse, but less risk averse than the harmonic utility.

The optimal distribution of preferences is quite complicated. By contrast, our numerical analysis shows that a nearly-optimal growth rate can be achieved by a simple distribution of vNM utilities with constant relative risk aversion (CRRA), where the relative risk coefficient is uniformly distributed between zero (risk neutrality) and two (harmonic utility). The predictions of our model fit reasonably well with the empirical works on the distribution of risk attitudes in the population, as discussed in Section 5.

Structure The paper is structured as follows. In Section 2 we describe the model. Section 3 presents the analytic results, which are supplemented by a numerical analysis in Section 4. We conclude with a discussion in Section 5.

2 Model

Consider a continuum population with an initial mass one. Reproduction is asexual. Time is discrete, indexed by $t \in \mathbb{N}$. Each agent lives a single time period (which is interpreted as a generation). In each time period, each agent in the population faces a choice between two alternatives, where each alternative is a lottery over the number of offspring. The first alternative (henceforth, the *safe alternative*) yields all agents who choose it with the same number of offspring, μ. The second alternative bears aggregate risk (henceforth, the *risky alternative*). That is, all agents who choose the risky alternative have the same number of offspring, but this number is a random variable $\boldsymbol{y}$ with finite support $supp(\boldsymbol{y}) = \{y_1, ..., y_n\} \in \mathbb{R}_+$ and distribution $\Pr[\boldsymbol{y} = y_i] = p_i$.

Let $\mathscr{Y}$ denote the set of risky alternatives (i.e., distributions over nonnegative numbers). A risky alternative $\boldsymbol{y} \in \mathscr{Y}$ is *degenerate* if $|supp(\boldsymbol{y})| = 1$, and we identify it with the respective constant that it yields.

Remark 1. Our model can capture a more general setup that also includes idiosyncratic risk (which is independent between different agents). It is well known (see, e.g., 24) that the long-run growth rate is the same for all idiosyncratic random variables with the same expectation.[5] In this extended setup, one should interpret the safe lottery as one that gives an expected number of offspring μ, where the number of offspring of each agent choosing the safe alternative might be the result of an arbitrary idiosyncratic lottery with expectation μ. Similarly, the risky alternative should be interpreted as a random variable $\boldsymbol{y}$ over the expected number of offspring (i.e., for each of the realizations $\boldsymbol{y} = y_i$ of the risky alternative, the number of offspring of each agent choosing the risky alternative might be any idiosyncratic random variable with expectation y_i). An additional assumption in this extended model is that all agents are endowed with risk neutrality with respect to idiosyncratic risk (i.e., they care only about the expected number of offspring).

Growth rate Let $\boldsymbol{w}(t)$ denote the size of the population in time t. Let $gr_\alpha(\boldsymbol{y}, \mu)$ denote the long-run growth rate (henceforth, *growth rate*) of a population in which in each

[5]This is implied by an exact law of large numbers for continuum populations. We refer the interested reader to Duffie and Sun [7] (and the citations therein) for details on how the exact law of large numbers is formalized in a related setup.

generation a share α of the population chooses the risky alternative $\boldsymbol{y}$ and the remaining agents obtain the safe option μ. It is well known (see, e.g., 24) that the growth rate $gr_{\alpha}(\boldsymbol{y},\mu)$ equals the geometric mean of $\alpha\boldsymbol{y}+(1-\alpha)\mu$:

$$\begin{aligned} gr_{\alpha}(\boldsymbol{y},\mu) & \equiv \lim_{T\to\infty} \sqrt[T]{\frac{\boldsymbol{w}(T)}{\boldsymbol{w}(1)}} = \mathbb{GM}\left[\alpha\boldsymbol{y}+(1-\alpha)\mu\right] = \\ & = \prod_{i\leqslant n}\left(\alpha y_i+(1-\alpha)\mu\right)^{p_i} = e^{\sum_{i\leqslant n} ln\left(p_i\cdot(\alpha y_i+(1-\alpha)\mu)\right)}. \end{aligned} \tag{1}$$

The intuition for Equation (1) is as follows. Let $\boldsymbol{z}(t)=\frac{\boldsymbol{w}(t+1)}{\boldsymbol{w}(t)}$ be the mean number of offspring in generation t; i.e., $\boldsymbol{z}(t)$ is a sequence of i.i.d. variables that are distributed like the random variable $\alpha\boldsymbol{y}+(1-\alpha)\mu$. Hence, the size of the population at time T equals

$$\begin{aligned} & \frac{\boldsymbol{w}(T)}{\boldsymbol{w}(1)} = \prod_{t<T}\frac{\boldsymbol{w}(t+1)}{\boldsymbol{w}(t)} = e^{\left(\sum_{t\leqslant T}\ln(\boldsymbol{z}(t))\right)} \\ & \Rightarrow \lim_{T\to\infty}\sqrt[T]{\frac{\boldsymbol{w}(T)}{\boldsymbol{w}(1)}} = \lim_{T\to\infty} e^{\left(\frac{1}{T}\sum\limits_{t<T}\ln(\boldsymbol{z}(t))\right)} \stackrel{(\star)}{=} e^{\mathbb{E}[\ln(\alpha\boldsymbol{y}+(1-\alpha)\mu)]} = \prod_{i\leqslant n}\left(\alpha y_i+(1-\alpha)\mu\right)^{p_i}, \end{aligned}$$

where the equality marked by $(\star)$ is implied by the law of large numbers.

Let $\alpha^{\star}\in[0,1]$ be the share of agents who choose the risky alternative that maximizes the long-run growth rate:

$$\alpha^{\star}(\boldsymbol{y},\mu) = \underset{\alpha\in[0,1]}{\operatorname{argmax}}\left(gr_{\alpha}(\boldsymbol{y},\mu)\right). \tag{2}$$

We show in Theorem 1 that $\alpha^{\star}(\boldsymbol{y},\mu)$ is unique.

Preferences Each agent is endowed with a preference over the lotteries, i.e., a linear order $\succcurlyeq$ over the set $\mathscr{Y}$ (and we use the notation $\sim$ for indifference). That is, Agent a chooses the risky alternative iff $\boldsymbol{y}\succ_a\mu$ (the tie-breaking rule that is applied when $\boldsymbol{y}\sim_a\mu$ has no impact on our results since it holds for a share of measure zero of the population). A preference $\succcurlyeq$ is *regular* if it satisfies the following two mild assumptions: (1) monotonicity of the safe alternatives: $\mu<\mu'$ implies that $\mu\prec\mu'$, and (2) any risky alternative has a certainty equivalent: for any $y\in\mathscr{Y}$, there exists a safe alternative μ such that $\boldsymbol{y}\sim\mu$.

Let $\mathscr{U}$ denote the set of regular preferences. Observe that any regular preference $\succcurlyeq$ can be represented by a *certainty equivalent function* $CE_{\succcurlyeq}:\mathscr{Y}\to\mathbb{R}_+$, which evaluates each risky alternative in terms of the equivalent safe alternative (i.e., $CE_{\succcurlyeq}(y)=\mu$ iff $\boldsymbol{y}\sim\mu$).

We assume that nature endows the population with a distribution Φ of regular preferences, and that each agent uses her preference to choose an alternative. A distribution of regular preferences Φ induces a choice function $\alpha_{\Phi}:\mathscr{Y}\times\mathbb{R}_+\to[0,1]$, which describes the share of agents who choose the risky option for any pair of alternatives.

A distribution of regular preferences $\Phi^{\star}$ is *optimal* if for any $\boldsymbol{y}\in\mathscr{Y}$ and $\mu\in\mathbb{R}_+$ it maximizes the growth rate, i.e.,

$$\alpha_{\Phi^{\star}}(\boldsymbol{y},\mu) = \alpha^{\star}(\boldsymbol{y},\mu) = \underset{\alpha\in[0,1]}{\operatorname{argmax}}\left(gr_{\alpha}(\boldsymbol{y},\mu)\right). \tag{3}$$

Dynamic interpretation Our interpretation of the model is that in each generation, all agents face a choice between the same pair of a safe alternative and a risky alternative. The choice changes over time, such that in different generations there might be different alternatives to choose from. The optimal distribution of preferences is required to maximize the expected long-run growth rate, and it is not hard to see that this is equivalent to maximizing the growth rate $gr_{\alpha}(\boldsymbol{y},\mu)$ for any combination of a risky alternative $\boldsymbol{y}$ and a safe alternative μ (which appears with non-zero frequency).

Further, observe that the optimal distribution is required to maximize the growth rate among all possible choice profiles (including choice functions that are not consistent with having preferences). Thus, our results show that the restriction of agents to regular preferences over the set of risky alternatives does not decrease the maximal feasible growth rate (see 25, 28 for arguments in favor of endowing agents with utilities).

Utility A common way to represent the preference of Agent a is to use a *utility function*, $U_a\colon \mathscr{Y} \to \mathbb{R}_+$, such that Agent a strictly prefers an alternative $\boldsymbol{y} \in \mathscr{Y}$ to another alternative $\boldsymbol{y'} \in \mathscr{Y}$ iff $U_a(\boldsymbol{y}) > U_a(\boldsymbol{y'})$. We note that for a given regular preference $\succcurlyeq$, its certainty equivalent function $CE_{\succcurlyeq}$ is in particular a utility function that represents $\succcurlyeq$.

A preference $\succcurlyeq$ is a *vNM (von Neumann–Morgenstern) preference* if it has an expected utility representation, that is, if there exists a Bernoulli utility function $u\colon \mathbb{R}_+ \to \mathbb{R}$ such that $\succcurlyeq$ is represented by the utility function $\mathbb{E}[u(\boldsymbol{y})] = \sum_i p_i \cdot u(y_i)$ for any $\boldsymbol{y} \in \mathscr{Y}$.

Risk aversion and constant relative risk aversion (CRRA) preferences A preference $\succcurlyeq$ is *risk averse* (resp., risk neutral) if $CE_{\succcurlyeq}(\boldsymbol{y}) < \mathbb{E}[\boldsymbol{y}]$ (resp., $CE_{\succcurlyeq}(\boldsymbol{y}) = \mathbb{E}[\boldsymbol{y}]$) for any nondegenerate risky alternative $\boldsymbol{y} \in \mathscr{Y}$. A preference $\succcurlyeq$ is *more risk averse* than a preference $\succcurlyeq'$ if $\mathrm{CE}_{\succcurlyeq}(\boldsymbol{y}) < \mathrm{CE}_{\succcurlyeq'}(\boldsymbol{y})$ for any nondegenerate risky alternative $\boldsymbol{y} \in \mathscr{Y}$.

For any $\rho \geqslant 0$, let CRRA_ρ denote the constant relative risk aversion preference with relative risk coefficient ρ, i.e., the expected utility preference defined by

$$\widehat{u}_\rho(y_i) = \begin{cases} \frac{y_i^{1-\rho}-1}{1-\rho} & \rho \neq 1 \\ \ln(y_i) & \rho = 1 \end{cases}. \tag{4}$$

Let $\mathbb{HM}[\boldsymbol{y}]$ and $\mathbb{GM}[\boldsymbol{y}]$ denote the harmonic and geometric means of $\boldsymbol{y}$, respectively,

$$\mathbb{HM}[\boldsymbol{y}] = \left(\mathbb{E}\left[\boldsymbol{y}^{-1}\right]\right)^{-1} = \frac{1}{\sum_i p_i/y_i}, \quad \mathbb{GM}[\boldsymbol{y}] = \prod_i y_i^{p_i}.$$

It is well known that:

1. $\mathbb{HM}[\boldsymbol{y}] \leqslant \mathbb{GM}[\boldsymbol{y}] \leqslant \mathbb{E}[\boldsymbol{y}]$ with strict inequality whenever $\boldsymbol{y}$ is nondegenerate.
2. Under the utilities CRRA_0, CRRA_1, and CRRA_2, the certainty equivalent values of any risky alternative $\boldsymbol{y} \in \mathscr{Y}$ are its arithmetic, geometric, and harmonic means, respectively. Hence, we also refer to CRRA_0 as the *linear utility*, to CRRA_1 as the *logarithmic utility*, and to CRRA_2 as the *harmonic utility*.

Last, a distribution Φ of regular preferences is *monotone* if its support is a chain with respect to the strong risk aversion order; i.e., for any two preferences $\succcurlyeq$ and $\succcurlyeq'$ in the support of Φ, $\succcurlyeq$ is either strictly more risk averse or strictly less risk averse than $\succcurlyeq'$.

3 Results

Observe that if nature is limited to a homogeneous population in which all agents have the same risk preference, then the maximal long-run growth rate is attained by the logarithmic utility CRRA_1, according to which the certainty equivalent of a risky alternative is its geometric mean (see, e.g., 17). This is an immediate corollary of Equation (1) and the definition of CRRA preferences.

Fact 1. For any $\boldsymbol{y} \in \mathscr{Y}$ and $\mu \in \mathbb{R}_+$, $gr_1(\boldsymbol{y}, \mu) \geqslant gr_0(\boldsymbol{y}, \mu) \Leftrightarrow \widehat{U}_1(\boldsymbol{y}) \geqslant \widehat{U}_1(\mu)$.

Our analysis is motivated by the fact that heterogeneous populations in which agents differ in the extent of their risk aversion can induce substantially higher growth rate, because the heterogeneity allows the population to hedge the aggregate risk by having only the more risk-averse agents choose the risky alternative. For example, if agents face a choice between a safe alternative μ yielding one offspring to each agent or a risky alternative $\boldsymbol{y}$ yielding either

4 offspring or 0.25 offspring with equal probabilities, then any homogeneous population in which all agents share the same risk preference (with a deterministic tie-breaking rule) yields a growth rate of 1 (because both $gr_0(\boldsymbol{y},\mu)=\mu=1$ and $gr_1(\boldsymbol{y},\mu)=\mathbb{GM}[\boldsymbol{y}]=4^{0.5}\cdot 0.25^{0.5}=1$). By contrast, a heterogeneous population in which agents differ in the extent of their risk aversion such that half the population choose the risky alternative $\boldsymbol{y}$ and the others choose the safe alternative μ induces a substantially higher growth rate of

$$gr_{0.5}(\boldsymbol{y},\mu)=\mathbb{GM}[0.5\cdot\boldsymbol{y}+0.5\cdot 1]=2.5^{0.5}\cdot 0.625^{0.5}=1.25.$$

Our first result characterizes the optimal share of agents that choose the risky alternative.[6]

Theorem 1. Fix $\boldsymbol{y}\in\mathscr{Y}$ and $\mu\in\mathbb{R}_+$. Then:

1. $\alpha^\star(\boldsymbol{y},\mu)=\operatorname{argmax}_{\alpha\in[0,1]}(gr_\alpha(\boldsymbol{y},\mu))$ is unique.
2. $\alpha^\star(\boldsymbol{y},\mu)=0$ iff $\mathbb{E}[\boldsymbol{y}]\leqslant\mu$, and
   $\alpha^\star(\boldsymbol{y},\mu)=1$ iff $\mathbb{HM}[\boldsymbol{y}]\geqslant\mu$.
3. If $\mu\in(\mathbb{HM}[\boldsymbol{y}],\mathbb{E}[\boldsymbol{y}])$, then $\alpha^\star(\boldsymbol{y},\mu)\in(0,1)$ is the unique solution to the following equation:
$$\mathbb{HM}[1+x\cdot(\boldsymbol{y}/\mu-1)]=1. \tag{5}$$

Proof. The long-run growth rate when a share α of the population chooses $\boldsymbol{y}$ is (see Equation (1))

$$gr_\alpha(\boldsymbol{y},\mu)=\mathbb{GM}[\alpha\cdot\boldsymbol{y}+(1-\alpha)\cdot\mu]=e^{\mathbb{E}[\ln(\alpha\cdot\boldsymbol{y}+(1-\alpha)\cdot\mu)]}.$$

Hence, $gr_\alpha(\boldsymbol{y},\mu)$ is maximized iff $\ln(gr_\alpha(\boldsymbol{y},\mu))=\mathbb{E}[\ln(\alpha\cdot\boldsymbol{y}+(1-\alpha)\cdot\mu)]$ is maximized, and since

$$\begin{aligned}\tfrac{d}{d\alpha}\ln(gr_\alpha(\boldsymbol{y},\mu)) &= \mathbb{E}\left[\tfrac{\boldsymbol{y}-\mu}{\alpha\cdot\boldsymbol{y}+(1-\alpha)\cdot\mu}\right]\\ \tfrac{d^2}{d^2\alpha}\ln(gr_\alpha(\boldsymbol{y},\mu)) &= -\mathbb{E}\left[\left(\tfrac{\boldsymbol{y}-\mu}{\alpha\cdot\boldsymbol{y}+(1-\alpha)\cdot\mu}\right)^2\right]<0,\end{aligned}$$

there is exactly one maximizer for $gr_\alpha(\boldsymbol{y},\mu)$ in $[0,1]$, and the following three statements hold:

- $\alpha^\star(\boldsymbol{y},\mu)=0$ iff $\frac{d}{d\alpha}\ln(gr_\alpha(\boldsymbol{y},\mu))|_{\alpha=0}=\mathbb{E}[\boldsymbol{y}/\mu]-1\leqslant 0$, i.e., $\mathbb{E}[\boldsymbol{y}]\leqslant\mu$.
- $\alpha^\star(\boldsymbol{y},\mu)=1$ iff $\frac{d}{d\alpha}\ln(gr_\alpha(\boldsymbol{y},\mu))|_{\alpha=1}=1-\mathbb{E}[\mu/\boldsymbol{y}]\geqslant 0$, i.e., $\mu\leqslant\mathbb{HM}[\boldsymbol{y}]$.
- If $\mu\in(\mathbb{HM}[\boldsymbol{y}],\mathbb{E}[\boldsymbol{y}])$, then $\alpha^\star(\boldsymbol{y},\mu)\in(0,1)$ is the unique solution to

$$\mathbb{E}\left[\frac{\boldsymbol{y}-\mu}{x\cdot\boldsymbol{y}+(1-x)\cdot\mu}\right]=0.$$

Noting that for $x\neq 0$,

$$\mathbb{E}\left[\tfrac{\boldsymbol{y}-\mu}{\mu+x\cdot(\boldsymbol{y}-\mu)}\right]=\tfrac{1}{x}\cdot\left(1-\mathbb{E}\left[\tfrac{\mu}{\mu+x\cdot(\boldsymbol{y}-\mu)}\right]\right)=\tfrac{1}{x}\cdot\left(1-\left(\mathbb{HM}[1+x\cdot(\boldsymbol{y}/\mu-1)]\right)^{-1}\right),$$

we get that $\alpha^\star(\boldsymbol{y},\mu)\in(0,1)$ is the unique solution to

$$\mathbb{HM}[1+x\cdot(\boldsymbol{y}/\mu-1)]=1.\quad\square$$

Given a distribution Φ of regular preferences and a risky alternative $\boldsymbol{y}\in\mathscr{Y}$, we define $\Phi_{\boldsymbol{y}}$ to be the distribution of certainty equivalent values of $\boldsymbol{y}$ in the population.Our next result characterizes the optimal distribution of risk preferences. Specifically, it shows that for any risky alternative $\boldsymbol{y}$, **(1)** the support of $\Phi_{\boldsymbol{y}}$ is the range between $\boldsymbol{y}$'s harmonic mean and $\boldsymbol{y}$'s arithmetic mean, and **(2)** the λ-median of $\Phi_{\boldsymbol{y}}$ is the unique solution to a simple equation.

[6] Similar results to Theorem 1 have been presented in related setups (see, e.g., the relative fitness condition in 2). For completeness we present a short proof.

Theorem 2. Let Φ be a distribution of regular preferences. Then, Φ is optimal iff for any risky alternative $\boldsymbol{y} \in \mathscr{Y}$, the cumulative density function (CDF) of $\Phi_{\boldsymbol{y}}$ is

$$\mathrm{CDF}_{\Phi_{\boldsymbol{y}}}(x) = 1 - \alpha^{\star}(\boldsymbol{y}, x).$$

In particular, for any (nondegenerate) risky alternative $\boldsymbol{y} \in \mathscr{Y}$,
1. The support of $\Phi_{\boldsymbol{y}}$ is $[\mathbb{HM}[\boldsymbol{y}], \mathbb{E}[\boldsymbol{y}]]$.
2. For any $\lambda \in (0,1)$ the λ-median of $\Phi_{\boldsymbol{y}}$ is the unique solution to

$$\mathbb{HM}[\lambda + (1-\lambda) \cdot {}^{\boldsymbol{y}}\!/_{x}] = 1.$$

Note that by Theorem 1, $(1 - \alpha^{\star}(\boldsymbol{y}, x))$ is indeed a CDF for any (nondegenerate) risky alternative $\boldsymbol{y}$. $\alpha^{\star}(\boldsymbol{y}, x)$ equals one when $x \leqslant \mathbb{HM}[\boldsymbol{y}]$, equals zero when $x \geqslant \mathbb{E}[\boldsymbol{y}]$, and equals the solution to $\mathbb{HM}[1 + \alpha \cdot ({}^{\boldsymbol{y}}\!/_{x} - 1)] = 1$ otherwise. The function $x \mapsto \alpha^{\star}(\boldsymbol{y}, x)$ is continuous and strictly downward monotone in $x \in (\mathbb{HM}[\boldsymbol{y}], \mathbb{E}[\boldsymbol{y}])$ since the function $(x, \alpha) \mapsto \mathbb{HM}[1 + \alpha \cdot ({}^{\boldsymbol{y}}\!/_{x} - 1)]$ is continuous, with a bounded domain, and strictly downward monotone in x. That is, $1 - \alpha^{\star}(\boldsymbol{y}, x)$ equals zero for $x \leqslant \mathbb{HM}[\boldsymbol{y}]$, equals one for $x \geqslant \mathbb{E}[\boldsymbol{y}]$, and is continuous and strictly upward monotone in between.

Proof. Let Φ be a distribution of regular preferences. An agent prefers $\boldsymbol{y}$ to a safe alternative μ iff her certainty equivalent value of $\boldsymbol{y}$ is higher than μ, which holds for a share $1 - \mathrm{CDF}_{\Phi_{\boldsymbol{y}}}(\mu)$ of the population, i.e.,

$$\alpha_{\Phi}(\boldsymbol{y}, \mu) = 1 - \mathrm{CDF}_{\Phi_{\boldsymbol{y}}}(\mu).$$

Hence, Φ is optimal iff for any $\boldsymbol{y} \in \mathscr{Y}$ and $\mu \in \mathbb{R}_{+}$,

$$\alpha_{\Phi}(\boldsymbol{y}, \mu) = \alpha^{\star}(\boldsymbol{y}, \mu), \text{ i.e., } \mathrm{CDF}_{\Phi_{\boldsymbol{y}}}(\mu) = 1 - \alpha^{\star}(\boldsymbol{y}, \mu).$$

In particular, by Theorem 1, for any risky alternative $y \in \mathscr{Y}$, the support of $\Phi_{\boldsymbol{y}}$ is $[\mathbb{HM}[\boldsymbol{y}], \mathbb{E}[\boldsymbol{y}]]$. Moreover, for any $\lambda \in (0,1)$, the λ-median of $\Phi_{\boldsymbol{y}}$, m_{λ}, satisfies

$$\lambda = \mathrm{CDF}_{\Phi_{\boldsymbol{y}}}(m_{\lambda}) = 1 - \alpha^{\star}(\boldsymbol{y}, m_{\lambda}).$$

Therefore, $\mathbb{HM}\left[1 + (1-\lambda) \cdot \left(\frac{\boldsymbol{y}}{m_{\lambda}} - 1\right)\right] = 1$ and m_{λ} is a solution to $\mathbb{HM}[\lambda + (1-\lambda) \cdot {}^{\boldsymbol{y}}\!/_{x}] = 1$.

Lastly, $\mathbb{HM}[\lambda + (1-\lambda) \cdot {}^{\boldsymbol{y}}\!/_{x}]$ is strictly downward monotone in x, and hence there is a unique solution to $\mathbb{HM}[\lambda + (1-\lambda) \cdot {}^{\boldsymbol{y}}\!/_{x}] = 1$. □

Corollary 1. By Theorem 2, the following is the unique monotone optimal distribution of regular preferences $\Phi^{\star}$. We index the agents by $[0,1]$ and define the preference of Agent $a \in [0,1]$ by defining her certainty equivalent value for any risky alternative $\boldsymbol{y} \in \mathscr{Y}$ to be:
- $\mathbb{HM}[\boldsymbol{y}]$ if $a = 0$,
- $\mathbb{E}[\boldsymbol{y}]$ if $a = 1$, and
- the unique solution to $\mathbb{HM}[a + (1-a) \cdot {}^{\boldsymbol{y}}\!/_{x}] = 1$ otherwise.

Remark 2. We note that $\mathrm{CE}_a(\boldsymbol{y})$ is continuous in the parameter a. It is easy to see that $\mathbb{HM}[\boldsymbol{y}]$ is the limit solution when $a \to 0$ for $\mathbb{HM}[a + (1-a) \cdot {}^{\boldsymbol{y}}\!/_{x}] = 1$.

For $a \to 1^{-}$, noticing that $(\mathbb{HM}[a + (1-a) \cdot {}^{\boldsymbol{y}}\!/_{x}])^{-1} =$

$$= 1 - (1-a) \cdot \left(\frac{1}{x} \cdot \mathbb{E}\left[\frac{\boldsymbol{y}}{1+(1-a)\cdot\left(\frac{\boldsymbol{y}}{x}-1\right)}\right] - \mathbb{E}\left[\frac{1}{1+(1-a)\cdot\left(\frac{\boldsymbol{y}}{x}-1\right)}\right]\right),$$

we get that $\mathrm{CE}_a(\boldsymbol{y})$ satisfies

$$\mathrm{CE}_a(\boldsymbol{y}) = \mathbb{E}\left[\frac{\boldsymbol{y}}{1 + (1-a) \cdot \left(\frac{\boldsymbol{y}}{\mathrm{CE}_a(\boldsymbol{y})} - 1\right)}\right] \cdot \left(\mathbb{E}\left[\frac{1}{1 + (1-a) \cdot \left(\frac{\boldsymbol{y}}{\mathrm{CE}_a(\boldsymbol{y})} - 1\right)}\right]\right)^{-1},$$

and hence $\lim_{a \to 1^{-}} \mathrm{CE}_a(\boldsymbol{y}) = \mathbb{E}[\boldsymbol{y}]$.

We note that the behavior of the least and the most risk-averse agents in $\Phi^\star$ is simple and intuitive. The choices of the least risk-averse agent, Agent 1, are consistent with CRRA_0 (risk neutrality), and the choices of the most risk-averse agent, Agent 0, are consistent with CRRA_2. By contrast, for any $a \in (0,1)$ and $\boldsymbol{y} \in \mathscr{Y}$, the choices of Agent a are consistent with $\text{CRRA}_{\rho_a(\boldsymbol{y})}$ for some $\rho_a(\boldsymbol{y}) \in (0,2)$; the dependency of $\rho_a(\boldsymbol{y})$ on $\boldsymbol{y}$ makes the representation of the preferences of these agents more cumbersome and in particular, as we show in Appendix A, they do not have an expected utility representation.

4 Numerical Analysis

Section 4.1 presents a Monte Carlo simulation that we use to evaluate what percentage of the theoretically optimal growth rate is induced by various distributions of utilities. The numerical results (Section 4.2) show that simple distributions of preferences, in which all agents have CRRA utilities and the relative risk coefficient is distributed between zero (risk neutrality) and two (harmonic utility), achieve (on average) 99.86% of the optimal long-run growth rate. In Section 4.3, we interpret these result as suggesting that simple distributions of CRRA utilities could be the result of the evolutionary process.

4.1 Setup and Simulation

The code of the simulation is detailed in the online supplementary material.

Distributions of utilities We compare 15 distributions of utilities:

1. Five homogeneous populations in which all agents have the same utility:
   (a) *Extreme risk loving*: All agents always choose the risky alternative (as long as $\Pr[\boldsymbol{y} > \mu] \neq 0$).
   (b) *Extreme risk aversion*: All agents always choose the safe alternative (as long as $\Pr[\boldsymbol{y} < \mu] \neq 0$).
   (c) *Risk neutrality (*CRRA_0*)*: All agents evaluate risky choices by their arithmetic mean.
   (d) *Logarithmic utility (*CRRA_1*)*: All agents evaluate risky choices by their geometric mean.
   (e) *Harmonic utility (*CRRA_2*)*: All agents evaluate risky choices by their harmonic mean.
2. Two classes of heterogeneous populations with monotone distributions. In each class, each agent is endowed with a value $\beta \in [0,1]$ (each class includes 5 distributions of β as detailed below). All classes have the property characterized by Corollary 1, namely, that the most and the least risk-averse agents (corresponding to $\beta = 1$ and $\beta = 0$, respectively) evaluate a risky alternative $\boldsymbol{y}$ as having a certainty equivalent value of the harmonic mean and arithmetic mean of $\boldsymbol{y}$, respectively.
   The behavior of the agent endowed with value $\beta \in [0,1]$ in each class is as follows:
   (a) *Heterogeneous constant relative risk aversion: populations*: All agents have $\text{CRRA}_{2\beta}$ preferences where β's distribution is detailed below.
   (b) *Heterogeneous weighted-average populations*: All agents evaluate risky alternatives as a weighted average of their harmonic mean and arithmetic means: $\text{CE}_\beta(\boldsymbol{y}) = \beta \cdot \mathbb{HM}[\boldsymbol{y}] + (1-\beta) \cdot \mathbb{E}[\boldsymbol{y}]$, where β's distribution is detailed as follows.[7]

[7] For any $\beta \neq 0, 1$, this preference cannot be represented using expected utility. Consider the lottery $\boldsymbol{L} = \begin{cases} 6 & 1/2 \\ 2 & 1/2 \end{cases}$, and the two degenerate lotteries $\boldsymbol{M} = 4 - \beta$ and $\boldsymbol{N} = 4$. Then,

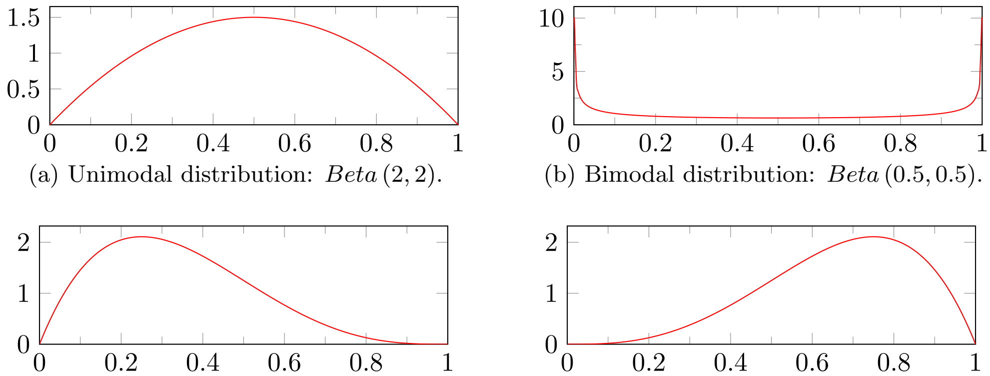

(a) Unimodal distribution: $Beta\,(2,2)$.
(b) Bimodal distribution: $Beta\,(0.5, 0.5)$.
(c) Positively skewed dist.: $Beta\,(2,4)$.
(d) Negatively skewed dist.: $Beta\,(4,2)$.

Figure 1: Probability Density Function (PDF) of $Beta\,(\alpha, \beta)$.

We use five beta distributions for $\beta \in [0,1]$ for the two classes (as demonstrated in Figure 1):

(a) Uniform distribution: $\beta \sim Beta\,(1,1)$.
(b) Unimodal distribution: $\beta \sim Beta\,(2,2)$.
(c) Bimodal distribution: $\beta \sim Beta\,(0.5, 0.5)$.
(d) Positively skewed distribution: $\beta \sim Beta\,(2,4)$.
(e) Negatively skewed distribution: $\beta \sim Beta\,(4,2)$.

Description of the simulation The simulation evaluates the performance of each distribution of utilities over $10.7M$ choices between a safe alternative μ and a binary risky alternative $\boldsymbol{y}$ yielding either a low realization ℓ or a high realization h. We run the following scenario for the distribution of the risky alternative and the safe alternative in each generation (the alternatives in different generations are independent of each other):

- In each generation, the two alternatives are defined by three independent uniform random numbers $p, q, r \in [0,1]$, where:[8]
  - p is the probability of the risky alternative yielding its high realization h: $p = \Pr[\boldsymbol{y} = h]$.
  - q is the ratio between the low realization of the risky alternative and the value of the safe alternative: $q = \ell/\mu$.
  - r is the ratio between the value of the safe alternative and the high realization of the risky alternative: $r = \mu/h$.

Without loss of generality, we normalize the value of the safe alternative to be $\mu = 1$. In each simulation run we calculate the theoretically optimal growth rate $gr_{\alpha^\star}(\boldsymbol{y}, \mu)$, and then evaluate the percentage of this optimal growth rate achieved by each of the 15 distributions of utilities. Finally, we calculate the geometric mean of this percentage for each distribution over all the simulation runs, which evaluates the relative performance of each distribution (in terms of its long-run growth rate) in a setup in which the risky and safe alternatives

- $\mathrm{CE}_\beta(\boldsymbol{L}) = \mathrm{CE}_\beta(\boldsymbol{M}) = 4 - \beta$.
- $\mathrm{CE}_\beta\left(1/2\boldsymbol{L} + 1/2\boldsymbol{N}\right) = 4 - \frac{4\beta}{7}$
- $\mathrm{CE}_\beta\left(1/2\boldsymbol{M} + 1/2\boldsymbol{N}\right) = \frac{1}{2\cdot(8-\beta)} \cdot \left(64 - 16\beta + \beta^2 - \beta^3\right)$

and hence Agent β is indifferent between $\boldsymbol{L}$ and $\boldsymbol{M}$ but not between $1/2\boldsymbol{L} + 1/2\boldsymbol{N}$ and $1/2\boldsymbol{M} + 1/2\boldsymbol{N}$ (in violation of the *independence axiom* of vNM) and, in particular, the preference of Agent β cannot be represented by expected utility.

[8] Equivalently, p, ℓ, and h are independent given μ and are sampled as follows: $p \in_{\mathbb{U}} [0,1]$, $\ell \in_{\mathbb{U}} [0,\mu]$, and $h \in [\mu, \infty)$ with the inverse-uniform distribution with parameters $\langle 0, 1\rangle$ $\left(F(x) = 1 - \mu/x\ ;\ f(x) = \mu/x^2\right)$.

change from one generation to the next.

4.2 Numerical Results

The results are summarized in Table 1. The optimal growth rate in our setup is 1.428 (which is calculated as the geometric mean of the growth rate achieved in each generation). We evaluate the performance of each distribution of preferences according to the decline in the relative growth rate, i.e., according to the percentage of the optimal growth rate that is lost under this distribution of preferences. The best homogeneous population is the one in which all agents have logarithmic utility, and it achieves a loss of 2.1 relative to the optimal growth rate. Heterogeneous CRRA populations reduce this loss substantially to less than 1 (which is better than what can be achieved by the heterogeneous weighted-average populations). Moreover, heterogeneous CRRA populations in which β is distributed uniformly (or according to the unimodal distribution) reduce this loss further to 0.15%.

Table 1: Summary of results of simulation runs ($10.7M$ generations).

Class	Distribution of β	Empirical mean of α	Long-run growth rate	Relative growth rate loss
		$\mathbb{E}[\alpha]$	$\mathbb{GM}[gr_\alpha(\boldsymbol{y},\mu)]$	$1-\frac{\mathbb{GM}[gr_\alpha(\boldsymbol{y},\mu)]}{\mathbb{GM}[gr_{\alpha^\star}(\boldsymbol{y},\mu)]}$
Optimal	(Corollary 1)	0.500	1.4251	0.00%
Homogeneous populations	Extreme risk loving	1.0	0.9949	30.2%
	Extreme risk aversion	0.0	1.0000	29.8%
	Risk neutrality	0.644	1.3152	7.7%
	Logarithmic utility (CRRA_1)	0.499	1.3949	2.1%
	Harmonic utility (CRRA_2)	0.357	1.3172	7.6%
Heterogeneous CRRA populations	Uniform	0.500	1.4232	0.14%
	Unimodal	0.500	1.4231	0.15%
	Bimodal	0.500	1.4211	0.29%
	Positively skewed	0.551	1.4118	0.93%
	Negatively skewed	0.448	1.4117	0.94%
Heterogeneous weighted-average populations	Uniform	0.530	1.4054	1.4%
	Unimodal	0.535	1.3983	1.9%
	Bimodal	0.524	1.4103	1.0%
	Positively skewed	0.577	1.3757	3.5%
	Negatively skewed	0.491	1.4054	1.4%

Robustness check To check the robustness of our results, we tested other parameter distributions (30 additional distributions in total) as follows:

- By taking the probability, $\Pr[\boldsymbol{y}=h]$, and the two ratios, ${}^{\mathbb{GM}[\boldsymbol{y}]}/_{\mu}$ and ${}^{\mu}/_{\mathbb{E}[\boldsymbol{y}]}$, to be three i.i.d. uniformly distributed random numbers in $[0,1]$.
- By conditioning the two distributions on the event $[\mathbb{GM}[\boldsymbol{y}] \leqslant \mu \leqslant \mathbb{E}[\boldsymbol{y}]]$ and the events $\left[\frac{i-1}{k} \leqslant \frac{\mu-\mathbb{GM}[\boldsymbol{y}]}{\mathbb{E}[\boldsymbol{y}]-\mathbb{GM}[\boldsymbol{y}]} \leqslant \frac{i}{k}\right]$ for $k=2,\ldots,5$ and $i=1,\ldots,k$.

For all these distributions, we see similar qualitative results in which the heterogeneous CRRA populations outperform the other populations and the optimal growth rate is approximated by a heterogeneous CRRA population under a simple distribution of the relative risk parameter (uniform and unimodal).

4.3 Interpretation of Results

A shift from the optimal homogeneous population of agents (in which everyone has a logarithmic utility, CRRA_1) to a heterogeneous population (in which agents have different CRRA utilities) requires evolution to engineer a gene that manifest itself in different levels of risk aversion for different people in the right proportion. Our numerical result suggests that such a shift is reasonable as: (1) the added advantage of increasing the long-run growth rate by 2% per generation (when shifting from a homogeneous population to a uniform (or unimodal) distribution of CRRA utilities) is substantial when compounded over many generations, and (2) we show two different simple distributions of CRRA utilities (uniform and unimodal) that achieve similar nearly-optimal behavior, which facilitates the genetic implementation, as the gene can implement one of various nearly-optimal distributions.[9]

The numerical results further suggest that it would be difficult for the evolutionary process to improve the growth rate relative to these simple distributions of CRRA utilities. A shift to the optimal distribution of utilities (or to any other complicated mechanism that implements exactly the optimal bet-hedging mixtures) would only improve the mean growth rate by an additional 0.14%. The cognitive capacity of humans is limited by various constraints (see, e.g., 18), such as the energy consumption of the human brain (20% of the body energy) and the newborn baby's brain size, which is limited by the birth canal. Allocating more cognitive resources to one activity has a shadow cost induced by decreasing cognitive resources allocated to other activities (a decrease that will reduce the agent's growth rate). Although we do not have a formal model of these shadow costs, it seems reasonable to conjecture that a shift from a simple CRRA utility to a complicated non-vNM utility requires more cognitive resources, and that their shadow costs outweigh the small benefit of 0.14%. This suggests that when taking into account cognitive costs, it might be optimal for the evolutionary process to induce a simple distribution of CRRA utilities with relative risk aversion coefficients between 0 and 2.

5 Discussion

In what follows we discuss various aspects of our model and their implications.

Empirical predictions Our model suggests that natural selection endowed the population with (1) risk-averse preferences and (2) heterogeneity in the level of risk aversion such that the agents' certainty equivalent values for a given lottery are distributed between the lottery's harmonic mean and its expectation, and that (3) the preference distribution can be approximated by constant relative risk aversion utilities with relative risk aversion between zero and two. Our model deals with lotteries with respect to the number of offspring (fitness), but it is plausible that people apply these endowed risk attitudes when dealing with lotteries over money, which is what is typically tested in the empirical literature.[10] Chiappori and Paiella [3] rely on large panel data and show that the elasticity of the relative risk aversion index with respect to wealth is small and statistically insignificant,

[9] Additional simulation runs suggest that any moderately-unimodal beta distribution around 0.5 (i.e., any $\beta \sim Beta\,(x,x)$ for $1 \leq x \leq 2$) achieve a nearly-optimal growth rate.

[10] A nonlinear relationship between consumption and fitness in our evolutionary past might shift the optimal levels of risk aversion with respect to money. Specifically, if the expected number of offspring is a concave function of consumption, then the support of the optimal distribution of relative risk aversion with respect to consumption will be shifted to the right.

which supports our first prediction of people having constant relative risk aversion utilities. Halek and Eisenhauer [13] relies on life insurance data to estimate the distribution of the levels of relative risk aversion in the population. Their data suggests that there is substantial heterogeneity in the levels of relative risk aversion in the population, and that about 80% of the population have levels of relative risk aversion between zero and two (13, Figure 1).

Multiple risky alternatives If there are multiple sources of risky alternatives, each with its own shared risk (e.g., multiple foraging techniques, where agents using the same foraging technique have correlated risk), then we implicitly assume that agents use some decision rule to choose between the different risky sources, and the single risky alternative in our model represents a combination of these sources. For example, if there are several independent and identically distributed risky alternatives $\boldsymbol{y}^1, .., \boldsymbol{y}^n$, then it is not hard to show that it is optimal for the population to choose these alternatives with equal shares, which can be modeled by the single risky alternative $\boldsymbol{y} = \frac{\boldsymbol{y}^1+...+\boldsymbol{y}^n}{n}$. We do not analyze the general question of how to optimally diversify risk among different sources of correlated risk.

Monomorphic heterogeneous populations The population in our model is a monomorphic population rather than a polymorphic population; that is, all agents in our model have the same genotype, which manifests itself in different degrees of risk aversion in different people. In the biological literature, such phenomena in which a single genotype induces heterogeneous behavior is known as genetic expressivity (see, e.g., 11, Section 6.4), and its usage in biological evolutionary models has been popularized in Grafen [8, 9]. Hence, the relative performance of individuals does not affect the path of the evolutionary dynamics.

Random expected utility Our interpretation of the optimal distribution of preferences in the population is heterogeneity in the population; namely, some agents are more risk averse than others. We note that the optimal distribution can also be implemented by random expected utility (see, e.g., 12); namely, each agent is endowed with the optimal distribution of preferences, and in each decision problem each agent randomly applies one of these preferences.

6 Conclusion

The existing literature has shown that when agents face a choice between alternatives with various levels of aggregate risk, the optimal growth rate induces the population to choose the mixture of these lotteries that achieves the optimal level of bet-hedging. The main contribution of this paper is to present a new mechanism that allows evolution to implement the optimal level of bet-hedging. This is done by nature inducing a heterogeneous population with different levels of risk aversion, such that the most risk-averse agent is indifferent between obtaining a risky lottery and obtaining its harmonic mean for sure, while the least risk-averse agent is risk neutral. Although, the exactly-optimal distribution of risk-averse utilities is quite complicated, we show numerically that a nearly-optimal growth rate is induced by a population of expected utility maximizers with constant relative risk aversion preferences, in which the risk coefficient is distributed between zero and two according to a simple distribution (uniform or unimodal). Such a distribution might be optimal, if one takes into account the cognitive costs.

Acknowledgements

We thank the JET editor (Tilman Börgers), the anonymous referees of COMSOC & JET, and the participants at the Neuroeconomics and the Biological Basis of Economics conference at Simon-Fraser University, the 16th Meeting of the Society for Social Choice and Welfare, the 12th Annual conference of the Israeli Chapter of the Game Theory Society the Bar-Ilan Game and Economic theory seminar, and the LEG2022 conference at the IMT School for Advanced Studies (Lucca) for various helpful comments.

We gratefully acknowledge the financial support of the European Research Council (#677057), US–Israel Bi-national Science Foundation (#2020022), and Israel Science Foundation (#2566/20, #2443/19, and #1626/18).

References

[1] Ted Bergstrom. Storage for good times and bad: Of rats and men. UC Santa Barbara: Department of Economics. Retrieved from https://escholarship.org/uc/item/9w50s0zs, 1997.

[2] Thomas J. Brennan and Andrew W. Lo. An evolutionary model of bounded rationality and intelligence. *PLOS ONE*, 7(11):1–8, 11 2012.

[3] Pierre-André Chiappori and Monica Paiella. Relative risk aversion is constant: Evidence from panel data. *Journal of the European Economic Association*, 9(6):1021–1052, 2011.

[4] Dan Cohen. Optimizing reproduction in a randomly varying environment. *Journal of Theoretical Biology*, 12(1):119–129, 1966.

[5] William S Cooper and Robert H Kaplan. Adaptive coin-flipping: A decision-theoretic examination of natural selection for random individual variation. *Journal of Theoretical Biology*, 94(1):135–151, 1982.

[6] Philip A Curry. Decision making under uncertainty and the evolution of interdependent preferences. *Journal of Economic Theory*, 98(2):357–369, 2001.

[7] Darrell Duffie and Yeneng Sun. The exact law of large numbers for independent random matching. *Journal of Economic Theory*, 147(3):1105–1139, 2012.

[8] Alan Grafen. Biological signals as handicaps. *Journal of theoretical biology*, 144(4): 517–546, 1990.

[9] Alan Grafen. Sexual selection unhandicapped by the fisher process. *Journal of theoretical biology*, 144(4):473–516, 1990.

[10] Alan Grafen. Formal darwinism, the individual-as-maximizing-agent analogy and bet-hedging. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 266 (1421):799–803, 1999.

[11] Anthony J.F. Griffiths, Susan R. Wessler, Sean B. Carrol, and John Doebley. *An Introduction to Genetic Analysis*. W. H. Freeman, eleventh edition, 2015.

[12] Faruk Gul and Wolfgang Pesendorfer. Random expected utility. *Econometrica*, 74(1): 121–146, 2006.

[13] Martin Halek and Joseph G Eisenhauer. Demography of risk aversion. *The Journal of Risk and Insurance*, 68(1):1–24, 2001.

[14] Yuval Heller. Overconfidence and diversification. *American Economic Journal: Microeconomics*, 6(1):134–153, 2014.

[15] Yuval Heller and Ilan Nehama. Evolutionary foundation for heterogeneity in risk aversion. *Journal of Economic Theory*, 208, 2023.

[16] Yuval Heller and Arthur Robson. Evolution, heritable risk, and skewness loving. *Theoretical Economics*, 16(2):403–424, 2021.

[17] Richard C. Lewontin and Daniel Cohen. On population growth in a randomly varying environment. *Proceedings of the National Academy of Sciences*, 62(4):1056–1060, 1969.

[18] Falk Lieder and Thomas L Griffiths. Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and brain sciences*, 43, 2020.

[19] John M. McNamara. Implicit frequency dependence and kin selection in fluctuating environments. *Evolutionary Ecology*, 9(2):185–203, 1995.

[20] Nick Netzer. Evolution of time preferences and attitudes toward risk. *American Economic Review*, 99(3):937–955, 2009.

[21] Nick Netzer, Arthur Robson, Jakub Steiner, and Pavel Kocourek. *Endogenous risk attitudes*. Centre for Economic Policy Research, 2021.

[22] H. Allen Orr. Absolute fitness, relative fitness, and utility. *Evolution: International Journal of Organic Evolution*, 61(12):2997–3000, 2007.

[23] Roberto Robatto and Balázs Szentes. On the biological foundation of risk preferences. *Journal of Economic Theory*, 172:410–422, 2017.

[24] Arthur J Robson. A biological basis for expected and non-expected utility. *Journal of Economic Theory*, 68(2):397–424, 1996.

[25] Arthur J Robson. The biological basis of economic behavior. *Journal of economic literature*, 39(1):11–33, 2001.

[26] Arthur J. Robson and H. Allen Orr. Evolved attitudes to risk and the demand for equity. *Proceedings of the National Academy of Sciences*, 118(26), 2021.

[27] Arthur J. Robson and Larry Samuelson. The evolution of time preference with aggregate uncertainty. *American Economic Review*, 99(5):1925–1953, 2009.

[28] Arthur J. Robson and Larry Samuelson. The evolutionary foundations of preferences. In Jess Benhabib, Alberto Bisin, and Matthew Jackson, editors, *The Handbook of Social Economics*, volume 1, pages 221–310. Elsevier, 2011.

[29] Arthur J. Robson and Larry Samuelson. Evolved attitudes to idiosyncratic and aggregate risk in age-structured populations. *Journal of Economic Theory*, 181:44–81, 2019.

[30] Jörgen W Weibull. *Evolutionary game theory*. 1997.

Yuval Heller
Department of Economics, Bar-Ilan University
Israel
Email: yuval.heller@biu.ac.il

Ilan Nehama
Department of Exact Sciences, Haifa University
Israel
Email: inehama@sci.haifa.ac.il or Email: ilan.nehama@mail.huji.ac.il

A Optimal Monotone Distribution Does Not Have an Expected Utility Representation

Consider the following five lotteries:

- $\boldsymbol{L}$ = The degenerate (safe) lottery 3
- $\boldsymbol{M} = \begin{cases} 1.5 & 3/4 \\ 20 & 1/4 \end{cases}$
- $\boldsymbol{N} = \begin{cases} 10 & 1/2 \\ 15 & 1/2 \end{cases}$
- $\boldsymbol{X} = 1/2\boldsymbol{L} + 1/2\boldsymbol{N} = \begin{cases} 3 & 1/2 \\ 10 & 1/4 \\ 15 & 1/4 \end{cases}$
- $\boldsymbol{Y} = 1/2\boldsymbol{M} + 1/2\boldsymbol{N} = \begin{cases} 1.5 & 3/8 \\ 10 & 1/4 \\ 15 & 1/4 \\ 20 & 1/8 \end{cases}$

Here, the median agent prefers $\boldsymbol{L}$ to $\boldsymbol{M}$ and prefers $\boldsymbol{Y} = 1/2\boldsymbol{M} + 1/2\boldsymbol{N}$ to $\boldsymbol{X} = 1/2\boldsymbol{L} + 1/2\boldsymbol{N}$. But this is a violation of the *independence axiom* of vNM, and in particular, the preference of the median agent cannot be represented by expected utility.[11]

- The median agent's certainty equivalent value for $\boldsymbol{M}$ is ≈ 2.54, and hence she prefers $\boldsymbol{L}$ to $\boldsymbol{M}$.
- Her certainty equivalent value for $\boldsymbol{X}$ is ≈ 6.04, and hence she prefers the safe option 6.1 to $\boldsymbol{X}$.
- Her certainty equivalent value for $\boldsymbol{Y}$ is ≈ 6.19, and hence she prefers $\boldsymbol{Y}$ to the safe option 6.1, and $\boldsymbol{Y}$ to $\boldsymbol{X}$.

[11] By Theorem 2, the certainty equivalent value of the median agent for a lottery $\boldsymbol{y} \in \mathscr{Y}$ is the unique solution to $\mathbb{HM}\left[1 + \boldsymbol{y}/x\right] = 2$.