Evolutionary Foundation for Heterogeneity in Risk Aversion

Heller, Yuval and NEHAMA, Ilan

Bar-Ilan university

13 October 2021

Evolutionary Foundation for Heterogeneity in Risk Aversion

Yuval Heller* Ilan Nehama†

2021-10-13

Abstract

We examine evolutionary basis for risk aversion with respect to aggregate risk. We study populations in which agents face choices between aggregate risk and idiosyncratic 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 aggregate risk and obtaining its linear and harmonic mean for sure, respectively. Moreover, an 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 Robson & Samuelson, 2011, for a survey). Various papers have shown that idiosyncratic

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

*Department of Economics, Bar Ilan University. Email: yuval.heller@biu.ac.il.

†Department of Economics, Bar Ilan University. Email: ilan.nehama@mail.huji.ac.il

risk (independent across individuals) induces a higher long-run growth rate (henceforth, growth rate) than aggregate risk and as a result, natural selection should induce agents to be more risk averse when facing aggregate risk.

This general result has been presented in three main forms in the literature. The first representation highlights the fact that the optimal long-run growth rate can be achieved by maximizing a logarithmic utility with respect to aggregate risk and a linear utility with respect to idiosyncratic risk (see, e.g., Lewontin & Cohen, 1969; Robson, 1996). The second representation shows that the optimal long-run growth rate can be achieved by agents who maximize the expected relative fitness (namely, the ratio between the agent's number of offspring and the total number of offspring in her generation; see, e.g., McNamara, 1995; Grafen, 1999; Curry, 2001; Orr, 2007). Finally, the bet-hedging approach interprets this result as the population should achieve a trade-off between maximizing the expected mean number of offspring and minimizing the variance of the mean number of offspring in the population (see, e.g., Cohen, 1966; Cooper & Kaplan, 1982; Bergstrom, 1997).[1]

In this paper, we introduce a new representation of the above general result that fits situations in which agents face discrete choices between alternatives with aggregate risk (which we refer to as risky alternatives) and alternatives without aggregate (which we refer to as safe alternatives). 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. 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 near-optimal growth rate can be achieved by a simple distribution of vNM (von Neumann–Morgenstern) preferences, according to which all agents have constant relative risk aversion, and the risk coefficient is uniformly distributed between zero and two. This new representation circumvents some of the difficulties we see in applying the existing representations to the prehistoric evolution of risk preferences (as discussed in Section 5).

[1]See also Robson & Samuelson (2009) and Netzer (2009) who study the evolution of risk attitude and its impact on time preferences, Robatto & Szentes (2017) who study choices that influences fertility rate in continuous time, Robson & Samuelson (2019) who explore age-structured populations, and Heller & Robson (2021) who analyze heritable risk, which is correlated between an agent and her offspring. The approach has been applied to analyze the prevalence of overconfidence in Heller (2014) and of the equity premium in Robson & Orr (2021).

Highlights of the Model We consider a continuum population with asexual reproduction. Each agent lives a single generation, in which she faces a choice between two lotteries over the number of offspring: a *safe alternative* for which the a-priori distribution of the number of offspring is known, and a *risky alternative* $\boldsymbol{y}$ with aggregate risk (i.e., after the agents make their choices, $\boldsymbol{y}$ is instantiated to be one safe alternative out of a given set). Nature induces the population with a distribution of risk preferences with regard to aggregate risk, according to which each of the agents chooses between the risky alternative and the safe alternative (and we assume that all agents are risk neutral with regard to idiosyncratic risk).

Main Results If Nature were limited to endow all agents with the same preference, then it would be optimal that all agents evaluate any risky alternative $\boldsymbol{y}$ as having a certainty equivalent of its geometric mean. However, heterogeneous populations can induce 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.

Proposition 1 characterizes the optimal share of agents that chooses the risky alternative. That is, the share that maximizes the long-run growth rate. In particular, it shows that all 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 $\boldsymbol{y}$, $HM(\boldsymbol{y})$, satisfies $HM(\boldsymbol{y}) \geqslant \mu$. Moreover, we characterize the optimal preference distribution (Proposition 2): (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 numeric analysis shows that a nearly-optimal growth rate can be achieved by a simple distribution of vNM utilities with constant relative risk aversion, 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 attitude 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 numeric 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 has only idiosyncratic risk (henceforth, the *safe alternative*) and its expected value is μ.

The second alternative bears aggregate risk (henceforth, the *risky alternative*). That is, after the agents make their choices, $\boldsymbol{y}$ is instantiated to be a safe alternative out of a given finite set (e.g., due to an environmental dependence among the lotteries taken by the individuals in a given generation, where the environment for each generation is an i.i.d. draw from a known distribution). Specifically, the mean number of offspring of agents who choose the risky alternative is a random variable $\widetilde{\boldsymbol{y}}$ with a finite support. Henceforth, we identify $\widetilde{\boldsymbol{y}}$ with $\boldsymbol{y}$ and denote its distribution by $supp\left(\boldsymbol{y}\right) = \{y_1, ..., y_n\} \in \mathbb{R}_+$ and $\Pr\left[\boldsymbol{y} = y_i\right] = p_i$.

We justify our choice to model an alternative $\boldsymbol{y}$ as a distribution over the mean number of offspring of agents who choose $\boldsymbol{y}$ on two grounds. First, due to an exact large law of large numbers for continuum populations, the mean number of offspring of agents who choose a safe alternative equals to its expectation, and hence the mean number of offspring is a sufficient statistic for analyzing the long-run growth rate. Second, from an agent's perspective, this choice is equivalent to assuming risk neutrality with respect to idiosyncratic risk, which is a plausible and common assumption (see, e.g., Robson, 1996). We further discuss this assumption in Section 5.

Let $\mathscr{Y}$ denote the set of risky alternatives (i.e., distributions over non-negative numbers). A risky alternative $\boldsymbol{y} \in \mathscr{Y}$ is *non-degenerate* if $\left|supp\left(\boldsymbol{y}\right)\right| > 1$. We identify a constant number of offspring μ with the degenerate distribution that yields μ with probability 1.

Growth rate Let $\boldsymbol{w}\left(t\right)$ denote the size of the population in time t. Let $gr_{\alpha}\left(\boldsymbol{y}, \mu\right)$ denote the long-run growth rate (henceforth, *growth rate*) of a population in which in each generation a share of α of the population chooses the risky option $\boldsymbol{y}$ and the remaining agents obtain the safe option μ. It is well-known (see, e.g., Robson, 1996) that the growth rate $gr_\alpha(\boldsymbol{y}, \mu)$ is equal to the geometric mean of $\alpha\boldsymbol{y} + (1-\alpha)\mu$:

$$gr_\alpha(\boldsymbol{y}, \mu) \equiv \lim_{T\to\infty} \sqrt[T]{\frac{\boldsymbol{w}(T)}{\boldsymbol{w}(1)}} = GM(\alpha\boldsymbol{y} + (1-\alpha)\mu) = \prod_{i\leqslant n} (\alpha y_i + (1-\alpha)\mu)^{p_i}. \quad (1)$$

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 at generation t, i.e., $\boldsymbol{z}(t)$ is a sequence of i.i.d. variables which are distributed like the random variable $\alpha\boldsymbol{y} + (1-\alpha)\mu$. Hence, the size of the population size at time T equals

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

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

$$\alpha^\star(\boldsymbol{y}, \mu) = \operatorname*{argmax}_{\alpha\in[0,1]} (gr_\alpha(\boldsymbol{y}, \mu)). \quad (2)$$

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 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 over 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 $y \sim \mu$.

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.,

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

Utility A common way to represent the preference of Agent a is using a *utility function*, $U_a \colon \mathscr{Y} \to \mathbb{R}_+$, s.t. Agent a strictly prefers an alternative $\boldsymbol{y} \in \mathscr{Y}$ over 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 which 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 non-degenerate 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 non-degenerate 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

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

Let $HM(\boldsymbol{y})$ and $GM(\boldsymbol{y})$ denote the harmonic and geometric means of $\boldsymbol{y}$, respectively, i.e.,

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

It is well-known that:

1. $HM(\boldsymbol{y}) \leqslant GM(\boldsymbol{y}) \leqslant \mathbb{E}[\boldsymbol{y}]$ with a strict inequality whenever $\boldsymbol{y}$ is non-degenerate.
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 (as was originally observed by Lewontin & Cohen, 1969). 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 their extent of 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 choosing the risky alternative. For example, if agents face a choice between a safe alternative yielding $\mu = 1$ 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 $gr_1(\boldsymbol{y}, \boldsymbol{\mu}) = GM(\boldsymbol{y}) = 4^{0.5} \cdot 0.25^{0.5} = 1$). By contrast, a heterogeneous population in which agents differ in their extent of risk aversion such that half the population choose $\boldsymbol{y}$ and the others choose the safe alternative induces a substantially higher growth rate of

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

Our first result characterizes the optimal share of agents that choose the risky alternative.[2]

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

[2]Similar results to Proposition 1 in related setups have been presented in the literature (see, e.g., the relative fitness condition in Brennan & Lo, 2012). For completeness we present a short proof.

1. $\alpha^{\star}(\boldsymbol{y},\mu) = \operatorname{argmax}_{\alpha\in[0,1]}\left(gr_{\alpha}(\boldsymbol{y},\mu)\right)$ *is unique.*
2. $\alpha^{\star}(\boldsymbol{y},\mu) = 0$ *iff* $\mathbb{E}[\boldsymbol{y}] \leqslant \mu$, *and*
   $\alpha^{\star}(\boldsymbol{y},\mu) = 1$ *iff* $HM(\boldsymbol{y}) \geqslant \mu$.
3. *If* $\mu \in (HM(\boldsymbol{y}), \mathbb{E}[\boldsymbol{y}])$, *then* $\alpha^{\star}(\boldsymbol{y},\mu) \in (0,1)$ *is the unique solution of the following equation:*

$$\mathbb{E}\left[\left(\frac{\mu}{\boldsymbol{y}-\mu}+x\right)^{-1}\middle|\, \boldsymbol{y}\neq\mu\right] = 0. \tag{5}$$

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

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

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

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

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., $\mu \leqslant HM(\boldsymbol{y})$.
- If $\mu \in (HM(\boldsymbol{y}), \mathbb{E}[\boldsymbol{y}])$, then $\alpha^{\star}(\boldsymbol{y},\mu) \in (0,1)$ is the unique solution of

$$\mathbb{E}\left[\left(\frac{\mu}{\boldsymbol{y}-\mu}+x\right)^{-1}\middle|\, \boldsymbol{y}\neq\mu\right] = 0. \qquad \square$$

Given a 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 y, the induced distribution $\Phi_{\boldsymbol{y}}$ has a support that starts in $\boldsymbol{y}$'s harmonic mean and ends in $\boldsymbol{y}$'s arithmetic mean, and fully characterizes these induced distributions.

Proposition 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*

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

In particular, for any (non-degenerate) risky alternative $\boldsymbol{y} \in \mathscr{Y}$,

- *The support of $\Phi_{\boldsymbol{y}}$ is $[HM(\boldsymbol{y}), \mathbb{E}[\boldsymbol{y}]]$.*
- *For any $\lambda \in (0,1)$ the λ-median of $\Phi_{\boldsymbol{y}}$ is the unique solution of*

$$\mathbb{E}\left[\left(\frac{\boldsymbol{y}}{\boldsymbol{y}-x} - \lambda\right)^{-1} \middle| \boldsymbol{y} \neq x\right] = 0.$$

Note that by Proposition 1, $(1 - \alpha^{\star}(\boldsymbol{y}, x))$ is indeed a CDF for any (non-degenerate) risky alternative $\boldsymbol{y}$. $\alpha^{\star}(\boldsymbol{y}, x)$ equals one when $x \leqslant HM(\boldsymbol{y})$, to zero when $x \geqslant \mathbb{E}[\boldsymbol{y}]$, and to the solution of $\mathbb{E}\left[\left(\frac{x}{\boldsymbol{y}-x} + \alpha\right)^{-1} \middle| \boldsymbol{y} \neq x\right] = 0$, otherwise. The function $x \mapsto \alpha^{\star}(\boldsymbol{y}, x)$ is continuous and strictly downward monotone in $x \in (HM(\boldsymbol{y}), \mathbb{E}[\boldsymbol{y}])$ since the function $(x, \alpha) \mapsto \mathbb{E}\left[\left(\frac{x}{\boldsymbol{y}-x} + \alpha\right)^{-1} \middle| \boldsymbol{y} \neq x\right]$ is continuous, with a bounded domain, and strictly downward monotone in x. I.e., $1 - \alpha^{\star}(\boldsymbol{y}, x)$ equals zero for $x \leqslant HM(\boldsymbol{y})$, 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 $\boldsymbol{y}$ over a safe alternative μ iff her certainty equivalent value of $\boldsymbol{y}$ is higher than μ, which holds for a proportion $1 - \mathrm{CDF}_{\Phi_{\boldsymbol{y}}}(\mu)$ 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 Proposition 1, for any risky alternative $y \in \mathscr{Y}$ the support of $\Phi_{\boldsymbol{y}}$ is $[HM(\boldsymbol{y}), \mathbb{E}[\boldsymbol{y}]]$.

Moreover, for any $\lambda \in (0,1)$ the λ-median of $\Phi_{\boldsymbol{y}}$, m_{λ}, satisfies

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

so

$$\mathbb{E}\left[\left(\frac{m_\lambda}{\boldsymbol{y}-m_\lambda}+(1-\lambda)\right)^{-1}\middle|\,\boldsymbol{y}\neq m_\lambda\right]=0,$$

and m_λ is a solution to

$$\mathbb{E}\left[\left(\frac{\boldsymbol{y}}{\boldsymbol{y}-x}-\lambda\right)^{-1}\middle|\,\boldsymbol{y}\neq x\right]=0.$$

Last, since

$$\begin{aligned}\mathbb{E}\left[\left(\tfrac{\boldsymbol{y}}{\boldsymbol{y}-x}-\lambda\right)^{-1}\middle|\,\boldsymbol{y}\neq x\right] &= \left(\Pr\left[\boldsymbol{y}\neq x\right]\right)^{-1}\cdot\left(\mathbb{E}\left[\tfrac{\boldsymbol{y}-x}{\boldsymbol{y}-\lambda(\boldsymbol{y}-x)}\right]-\mathbb{E}\left[\mathbb{1}_{\boldsymbol{y}=\boldsymbol{x}}\cdot\tfrac{\boldsymbol{y}-x}{\boldsymbol{y}-\lambda(\boldsymbol{y}-x)}\right]\right)\\ &= \left(\Pr\left[\boldsymbol{y}\neq x\right]\right)^{-1}\cdot\mathbb{E}\left[\tfrac{\boldsymbol{y}-x}{\boldsymbol{y}-\lambda(\boldsymbol{y}-x)}\right],\end{aligned}$$

$\mathbb{E}\left[\left(\tfrac{\boldsymbol{y}}{\boldsymbol{y}-x}-\lambda\right)^{-1}\middle|\,\boldsymbol{y}\neq x\right]$ equals zero iff $\mathbb{E}\left[\tfrac{\boldsymbol{y}-x}{\boldsymbol{y}-\lambda(\boldsymbol{y}-x)}\right]$ equals zero. $\mathbb{E}\left[\tfrac{\boldsymbol{y}-x}{\boldsymbol{y}-\lambda(\boldsymbol{y}-x)}\right]$ is strictly downward monotone in x,

$$\frac{d}{dx}\mathbb{E}\left[\frac{\boldsymbol{y}-x}{\boldsymbol{y}-\lambda\left(\boldsymbol{y}-x\right)}\right]=-\mathbb{E}\left[\frac{\boldsymbol{y}}{\left(\boldsymbol{y}-\lambda\left(\boldsymbol{y}-x\right)\right)^2}\right]<0,$$

and hence there is a unique solution to $\mathbb{E}\left[\left(\tfrac{\boldsymbol{y}}{\boldsymbol{y}-x}-\lambda\right)^{-1}\middle|\,\boldsymbol{y}\neq x\right]=0.$ □

Corollary 1. *By Proposition 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:*

- $HM\left(\boldsymbol{y}\right)$ *if* $a=0$,
- $\mathbb{E}\left[\boldsymbol{y}\right]$ *if* $a=1$, *and*
- *the unique solution to*

$$\mathbb{E}\left[\left(\frac{\boldsymbol{y}}{\boldsymbol{y}-x}-a\right)^{-1}\middle|\,\boldsymbol{y}\neq x\right]=0,\ \textit{otherwise.}$$

We note that the behavior of the least and the most risk averse agents in $\Phi^\star$ is simple and intuitive. The actions of the least risk averse agent, Agent 1, are consistent with CRRA_0 (risk neutrality), and the actions 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 actions of Agent a are consistent with $\mathrm{CRRA}_{\rho_a(\boldsymbol{y})}$ for some $\rho_a\left(\boldsymbol{y}\right)\in(0,2)$; 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. In the next section, we demonstrate numerically that a simple distribution of preferences, in which all agents have constant relative risk aversion preferences (which are in particular, vNM preferences), and the relative risk coefficient is uniformly distributed between zero (risk neutrality) and two (harmonic utility) achieves 99.85% of the optimal long-run growth rate.

4 Numeric Analysis

In this section, we use a Monte Carlo simulation to evaluate which percentage of the theoretically optimal growth rate is induced by various simple distributions of utilities.

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 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.
3. The behavior of the β-agent in each class is as follows:
   (a) Constant relative risk aversion: All agents have $\text{CRRA}_{2\beta}$ preferences where β's distribution is detailed below.

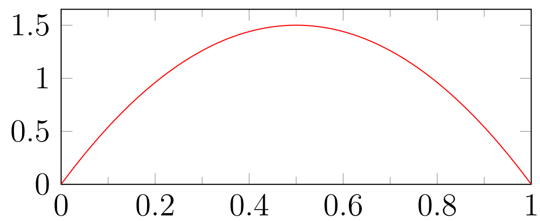

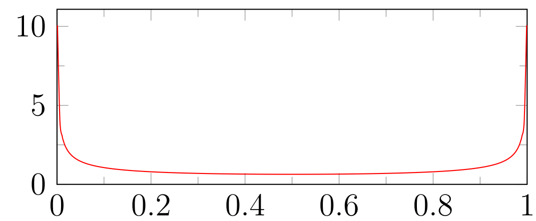

(a) Unimodal distribution: $Beta\,(2, 2)$ (b) Bimodal distribution: $Beta\,(0, 5, 0.5)$

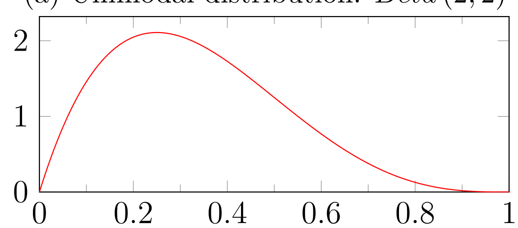

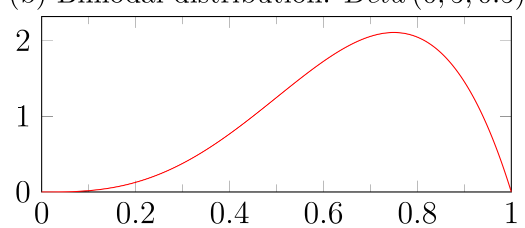

(c) Positively skewed dist.: $Beta\,(2, 4)$ (d) Negatively skewed dist.: $Beta\,(4, 2)$.

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

   (b) Weighted average of means: All agents evaluate risky alternatives as a weighted average of their harmonic mean and arithmetic means: $\mathrm{CE}_\beta\,(\boldsymbol{y}) = \beta \cdot HM\,(\boldsymbol{y}) + (1 - \beta) \cdot \mathbb{E}\,[\boldsymbol{y}]$,[3] where β's distribution is detailed below.

4. 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 option μ and a binary risky

[3] Note that for any $\beta \neq 0, 1$, this preference cannot be represented using expected utility. Consider the following three lotteries: $\boldsymbol{L} = \begin{cases} 6 & {}^1\!/_2 \\ 2 & {}^1\!/_2 \end{cases}$, and the two degenerate lotteries $\boldsymbol{M} = 4 - \beta$ and $\boldsymbol{N} = 4$.

Then,
- $\mathrm{CE}_\beta\,(\boldsymbol{L}) = \mathrm{CE}_\beta\,(\boldsymbol{M}) = 4 - \beta$.
- $\mathrm{CE}_\beta\,({}^1\!/_2\boldsymbol{L} + {}^1\!/_2\boldsymbol{N}) = 4 - \frac{4\beta}{7}$
- $\mathrm{CE}_\beta\,({}^1\!/_2\boldsymbol{M} + {}^1\!/_2\boldsymbol{N}) = \frac{1}{2\cdot(8-\beta)} \cdot \left(64 - 16\beta + \beta^2 - \beta^3\right)$

and hence Agent β is indifferent between $\boldsymbol{L}$ and $\boldsymbol{M}$ but is 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 using expected utility.

option $\boldsymbol{y}$ yielding either a low realization ℓ or a high realization h. We run the following scenario for the distribution of risky lottery 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:[4]
  - 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 and the safe alternative: $q = {}^{\ell}/_{\mu}$.
  - r is the ratio between safe alternative and the high realization: $r = {}^{\mu}/_{h}$.

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 & safe alternatives change from one generation to another.

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 rates achieved in each generation). We evaluate the performance of each distribution of preferences according to the relative growth rate loss, i.e., according to the percentage of the optimal growth rate that is lost with this distribution of preferences. The optimal homogeneous population in which all agents have logarithmic utility looses 2.1% of the optimal growth rate. Heterogeneous CRRA populations reduce this loss substantially to less than 1% (which is better than what can be achieved by mixed-average populations). Moreover, heterogeneous CRRA populations in which β is distributed uniformly (or a according to a unimodal distribution) achieve an additional reduction of this loss into only 0.15%.

Robustness check In order to check the robustness of our results, we tested the following other parameter distributions (30 additional distributions in total):

- Taking the probability, $\Pr[\boldsymbol{y} = h]$, and the two ratios, ${}^{GM(\boldsymbol{y})}/_{\mu}$ and ${}^{\mu}/_{\mathbb{E}[\boldsymbol{y}]}$, to be three i.i.d. uniformly distributed random numbers in $[0,1]$.

[4]Equivalently, p, ℓ, h are independent given μ and 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 $\langle 0, 1\rangle$ ($F(x) = 1 - {}^{\mu}/_{x}$; $f(x) = {}^{\mu}/_{x^2}$).

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
		$\mathbb{E}[\alpha]$	$GM(gr_\alpha(\boldsymbol{y}, \mu))$	$1 - \frac{GM(gr_\alpha(\boldsymbol{y},\mu))}{GM(gr_{\alpha^\star}(\boldsymbol{y},\mu))}$
Optimal	(Corollary 1)	0.500	1.425	0.00%
Homogeneous populations	Always risky	1.0	0.995	30.2%
	Always safe	0.0	1.000	29.8%
	Risk neutral	0.644	1.315	7.7%
	Logarithmic utility (CRRA_1)	0.499	1.395	2.1%
	Harmonic utility (CRRA_2)	0.357	1.317	7.6%
Heterogeneous CRRA populations	Uniform	0.500	1.423	0.14%
	Unimodal	0.500	1.423	0.15%
	Bimodal	0.0500	1.421	0.29%
	Positively skewed	0.551	1.412	0.93%
	Negatively skewed	0.448	1.412	0.94%
Heterogeneous mixed-average populations	Uniform	0.530	1.405	1.4%
	Unimodal	0.535	1.398	1.9%
	Bimodal	0.524	1.410	1.0%
	Positively skewed	0.577	1.376	3.5%
	Negatively skewed	0.491	1.405	1.4%

- Conditioning the two distributions on the event $[GM(\boldsymbol{y}) \leqslant \mu \leqslant \mathbb{E}[\boldsymbol{y}]]$ and the events $\left[\frac{i-1}{k} \leqslant \frac{\mu - GM(\boldsymbol{y})}{\mathbb{E}[\boldsymbol{y}] - GM(\boldsymbol{y})} \leqslant \frac{i}{k}\right]$ 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 with a simple distribution of the relative risk parameter (uniform and unimodal).

5 Discussion

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

Empirical predictions Our model suggests that natural selection had endowed the population with (1) risk-averse preferences and (2) heterogeneity in the level of risk aversion s.t. 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 typically tested in the empirical literature.[5] Chiappori & Paiella (2011) rely on large panel data to show that the elasticity of the risky asset share to wealth is small and statistically insignificant, which supports our first prediction of people having constant relative risk aversion utilities. Halek & Eisenhauer (2001) 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 (Halek & Eisenhauer, 2001, Figure 1).

Risk neutrality with respect to idiosyncratic risk: We implicitly assume that agents are risk neutral with respect to idiosyncratic risk over the number of offspring. This assumption is plausible, as it is well-known that risk neutrality with respect to idiosyncratic risk maximizes the growth rate (see, e.g., Robson, 1996). In particular, if there are multiple safe alternatives (with idiosyncratic risk) it is optimal for all agents to choose the alternative with the highest expectation.

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 identically distributed risky alternatives $\boldsymbol{y}^1, .., \boldsymbol{y}^n$, then it is not hard to show that it is optimal for the agents to divide equally among these alternatives, and this will 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.

[5] A non-linear 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 a shift to the right of the $[0, 2]$ interval.

Difficulties with applying the existing representations in our setup: In what follows we briefly discuss why it is difficult to apply the three existing representations in our setup, which highlights the advantage of our new representation. The assumption that agents cannot hedge their own personal risk (but, rather each agent must make a binary choice between $\boldsymbol{y}$ and μ) does not allow the population to achieve the optimal growth rate by either the first representation (all agents having logarithmic utility, as discussed after Fact 1) or the third representation (each agent applies bet hedging). The assumption that an agent cannot condition her play on the aggregate behavior, prevents the population from achieving the optimal growth rate by relying on the second representation (agents maximizing expected relative fitness, where the calculation of relative fitness crucially depends on the aggregate behavior). We think that both assumptions have been plausible in our evolutionary past: in many cases risky alternatives (such as, foraging technique) require long training and specialization that makes it difficult for an agent to hedge risk by dividing her time between several different alternatives. Moreover, it seems plausible that agents in prehistoric times would not know the aggregate behavior in the population.

Random expected utility Out 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., Gul & Pesendorfer, 2006), namely that 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 key result that aggregate risk yields a lower growth rate than idiosyncratic risk has been applied in the existing literature to derive three representations for the optimal risk preferences of agents: logarithmic utility, relative fitness, and bet hedging. We argue that all of these representations are difficult to implement in a plausible setup in which each agent faces a binary choice between a risky alternative and a safe alternative, and agents do not know the aggregate choices in the population. We present a new representation, according to which nature induces the agents with a distribution of preferences to choose between a risky alternative and a safe alternative.

We show that in any any distribution of preferences which induces the maximal long-run growth rate, 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, that is, indifferent between obtaining a risky lottery and obtaining its arithmetic mean for sure. Moreover, 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).

A Optimal Monotone Distribution Does Not Have an Expected Utility Representation

Consider the following five lotteries

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

Then, the median agent prefers $\boldsymbol{L}$ over $\boldsymbol{M}$ and prefers $\boldsymbol{Y} = 1/2\boldsymbol{M} + 1/2\boldsymbol{N}$ over $\boldsymbol{X} = 1/2\boldsymbol{L} + 1/2\boldsymbol{N}$. But this is a violation of the *Independence Axiom* of vNM and in particular, the preference of the median agent cannot be represented using expected utility.[6]

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

References

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

Brennan, Thomas J., & Lo, Andrew W. 2012. An Evolutionary Model of Bounded Rationality and Intelligence. *PLOS ONE*, **7**(11), 1–8.

[6]By Proposition 2, the certainty equivalent value of the median agent for a lottery $\boldsymbol{y} \in \mathscr{Y}$ is the unique solution to $\mathbb{E}\left[\left(\frac{\boldsymbol{y}}{\boldsymbol{y}-x} - 1/2\right)^{-1} \,\middle|\, \boldsymbol{y} \neq x\right] = 0$.

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

Cohen, Dan. 1966. Optimizing reproduction in a randomly varying environment. *Journal of theoretical biology*, **12**(1), 119–129.

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

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

Grafen, Alan. 1999. 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.

Gul, Faruk, & Pesendorfer, Wolfgang. 2006. Random expected utility. *Econometrica*, **74**(1), 121–146.

Halek, Martin, & Eisenhauer, Joseph G. 2001. Demography of risk aversion. *Journal of Risk and Insurance*, 1–24.

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

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

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

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

Netzer, Nick. 2009. Evolution of time preferences and attitudes toward risk. *American Economic Review*, **99**(3), 937–55.

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

Robatto, Roberto, & Szentes, Balázs. 2017. On the biological foundation of risk preferences. *Journal of Economic Theory*, **172**, 410–422.

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

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

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

Robson, Arthur J, & Samuelson, Larry. 2011. The evolutionary foundations of preferences. *Pages 221–310 of:* Benhabib, Jess, Bisin, Alberto, & Jackson, Matthew (eds), *The Handbook of Social Economics*, vol. 1. Elsevier.

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