Decision Theory and Stochastic Growth*

Arthur Robson
Simon Fraser University
arthur_robson@sfu.ca

Larry Samuelson
Yale University
larry.samuelson@yale.edu

Jakub Steiner
University of Zurich and CERGE-EI
jakubsteiner77@gmail.com

December 3, 2022

Abstract This paper examines connections between stochastic growth and decision problems. We use tools from the theory of large deviations to show that wishful thinking decision problems are equivalent to utility maximization problems, both of which are equivalent to growth maximization under idiosyncratic risk. Rational inattention problems are equivalent to growth-optimal portfolio problems, both of which are equivalent to growth maximization under aggregate risk. Stochastic growth generates extreme inequality, with nearly all wealth eventually held by those who happen to have faced empirical distributions that match the solution to the wishful thinking or rational inattention problem.

Keywords: stochastic growth, rational inattention, wishful thinking, Kelly criterion, large deviations.

Contents

*We have benefited from helpful discussions with Jaroslav Borovicka, Andrew Caplin, Mira Frick, Tai-Wei Hu, Ryota Iijima, Omer Tamuz. and Colin Stewart, and from comments at various presentations. We thank the editor (Navin Kartik) and three referees for helpful comments. This work was supported by ERC grant 770652.

Decision Theory and Stochastic Growth

1 Introduction

This paper derives and exploits connections between stochastic growth in large economies and individual decision problems.

We show that a decision model of wishful thinking is equivalent to standard expected utility maximization under a modified utility function, and that rational inattention problems are equivalent to growth-optimal portfolio problems. We then introduce the stochastic growth model, tracking wealth accumulation in a continuous population of agents that face idiosyncratic or aggregate risk. In the former case, each agent faces her own draw of the payoff state. In the latter case, all agents face the same draw. We show that expected utility maximization and wishful thinking are equivalent to growth maximization under idiosyncratic risk, while the growth-optimal portfolio problem and rational inattention optimization are equivalent to growth maximization under aggregate risk. Hence, in each of the following rows, the growth process and the two representations are equivalent:

Growth Process	Representations	
Idiosyncratic Risk	Expected Utility Maximization	Wishful Thinking
Aggregate Risk	Growth-Optimal Portfolio	Rational Inattention Problem

These results reflect a basic trade-off between the likelihood and the growth implications of a sequence of realizations of random shocks. Wealth accumulation involves a product of stochastic financial returns. The empirical distribution of these returns occasionally exhibits extremely unlikely deviations form the true distribution. Such large deviations have a negligible effect under an additive objective, such as maximizing the sum of flow utilities. However, they are of first-order importance under our multiplicative objective, because the extremely small probability of a large deviation can be compensated by an extremely large wealth advantage. We show that a standard tool from the theory of large deviations captures this trade-off, leading to straightforward proofs of the equivalences.

Tracking the population dynamics allows us to isolate the statistical properties of stochastic growth of wealth and to connect these to their decision-theoretic counterparts. First, stochastic growth generates extreme wealth concentration. Nearly all wealth is held by a fraction of the population that has enjoyed a specific "large deviation," and this fraction vanishes exponentially with time. For the growth process with idiosyncratic risk, the wealth concentration increases at the rate that equals the expenditure on belief distortion in the wishful thinking model. For the aggregate risk case, the wealth concentration increases at the rate that equals the expenditure on information acquisition in the rational inattention model. Second, the agents who end up owning nearly all wealth in the two processes enjoy the empirical distribution of payoff states given by large deviations. Under idiosyncratic risk, these winners of the growth process face an empirical distribution of states that equals the subjective distribution chosen by the wishful thinker. Similarly, the "winners" under aggregate risk face an empirical joint distribution of states and actions equal to that of the optimizing rationally inattentive decision maker.

One message of our results is that distinct processes may generate identical data. An analyst examining choice data may conclude that an individual is a wishful thinker rather than an expected utility maximizer (or vice versa). These representations of choice data reflect different views of a process of economic growth with idiosyncratic uncertainty. A perspective that focuses on the short-run growth of aggregate wealth corresponds to expected utility maximization, whereas an equivalent perspective that focuses on long-run growth of individuals' wealth corresponds to wishful thinking. Analogously, distinct but equivalent perspectives on the economic growth with aggregate uncertainty correspond to rational-inattention and growth-optimal portfolio problems. One might view these equivalences as a route to misspecification. We prefer a somewhat different interpretation—multiple equivalent representations are useful because they correspond to distinct, economically relevant aspects of the growth processes.

As an example of how these connections can be exploited, we show that rational inattention techniques can be used to derive comparative statics for the growth-optimal portfolio problem. We find that, under a regularity condition, the increase in the growth rate induced by a provision of a public signal equals the mutual information between the signal and the underlying state, and that the provision of the public signal reduces inequality. In a related result, we show that a decline of the growth rate caused by a misspecification of the prior distribution of the payoff shocks equals the Kullback-Leibler divergence of the true prior from the misspecified one. Strikingly, the value of information and cost of misspecification arising in the context of growth are universal in that they do not depend on the details of the underlying utility.

Growth processes with idiosyncratic and aggregate risks were introduced by Robson (1996). A subsequent literature, surveyed by Robson and Samuelson (2011), has explored this distinction. For example, Robson and Szentes (2014) examine the evolutionary foundations of time preferences, Steiner and Stewart (2016) examine the evolutionary foundations of distorted beliefs, and Sadowski and Sarver (2021) argue that evolution can select for ambiguity averse preferences. Like this literature, we use growth processes to provide insights into decision-theoretic models, while also using decision theory insights to derive statistical properties of growth processes.

Somewhat further afield, Dillenberger et al. (2017) show that the behavior of an expected utility maximizer is indistinguishable from that of an agent with a transformed utility function and motivated beliefs, without appealing to growth considerations. Benhabib and Bisin (2018) stress the propensity of stochastic growth to give rise to inequality, without making connections to decision theory. Heller and Robson (2021) (Section 6) also obtain a growing inequality result, but do not characterize this in terms of large deviations. Williams (2004) emphasizes the importance of large deviations for stochastic growth processes.

2 A Large Deviation Result

Our basic tool is an insight from large deviation theory. Consider a random variable $\mathbf{x}$ with a full-support distribution $p \in \Delta(X)$ on a finite set X.[1] Let f be a real-valued function on X. Let $D(q \parallel p) = \mathrm{E}_q \ln \frac{q(\mathbf{x})}{p(\mathbf{x})} \geq 0$ be the Kullback-Leibler divergence of the distributions q and p, commonly interpreted as the (pseudo) distance of the probability distribution q from the reference distribution p.[2] Dupuis and Ellis (1997) (Proposition 1.4.2) prove the following:

Lemma 1 (Donsker-Varadhan Lemma)**.**

$$\ln \mathrm{E}_p \, e^{f(\mathbf{x})} = \max_{q \in \Delta(X)} \left\{ \mathrm{E}_q \, f(\mathbf{x}) - D(q \parallel p) \right\}. \tag{1}$$

The maximum in (1) is attained, as the objective is a continuous function on a compact set.

We develop a heuristic derivation of (1) that highlights its relevance for stochastic growth. Suppose that some quantity begins at value $\mathbf{s}_0 = 1$ and is then governed by the multiplicative process $\mathbf{s}_t = e^{f(\mathbf{x}_t)} \mathbf{s}_{t-1}$ for $t = 1, \ldots, T$. Assume that the values $\mathbf{x}_t$ are independently drawn from the distribution p.

We show that each side of (1) is an (at least approximate) expression for the growth rate of the expected value of $\mathbf{s}_T$. First, since the realizations of $\mathbf{x}_t$ are independently drawn, we have $\mathrm{E}\,\mathbf{s}_T = \left(\mathrm{E}_p \, e^{f(\mathbf{x})}\right)^T$. Taking a log and dividing by T to obtain a growth rate then gives the left side of (1):

$$\frac{1}{T} \ln \mathrm{E}\,\mathbf{s}_T = \ln \mathrm{E}_p \, e^{f(\mathbf{x})}. \tag{2}$$

Alternatively, let $\mathbf{x}^T = (\mathbf{x}_t)_{t=1}^T$ denote the random sequence of length T generated by drawing according to p in each period. We can compute $\mathrm{E}\,\mathbf{s}_T$ by averaging over such sequences. Let x^T stand for the sequence realization and let $q_{x^T}(x) = \frac{1}{T} \sum_{t=1}^T \mathbb{1}_{x_t = x}$, so that $q_{x^T} \in \Delta(X)$ is the empirical distribution for the sequence realization x^T. For each sequence realization x^T, $\mathbf{s}_T$ attains value

$$\exp\left[T \times \mathrm{E}_{q_{x^T}} \, f(\mathbf{x})\right]. \tag{3}$$

Ignoring divisibility concerns, the *large deviation principle* states that the log probability of observing an empirical probability $q \in \Delta(X)$ asymptotically approximates $-T \times I(q)$ as T grows, where $I : \Delta(X) \to \mathbb{R}_+$ is a *rate function*. In this case, taking the expectation of (3) over random sequences, pooling sequences with the same empirical distribution q, and applying the large

[1] We denote random variables in bold throughout.

[2] We use the usual convention that $0 \ln 0 = 0$.

deviation principle, we obtain

$$\begin{aligned}
\frac{1}{T}\ln \mathrm{E}\,\mathbf{s}_T &\approx \frac{1}{T}\ln \int_{q\in\Delta(X)} \exp\left[T\times\left(\mathrm{E}_q f(\mathbf{x}) - I(q)\right)\right] dq \\
&\approx \frac{1}{T}\ln\exp\left[T\times \max_{q\in\Delta(X)}\left\{\mathrm{E}_q f(\mathbf{x}) - I(q)\right\}\right] \\
&= \max_{q\in\Delta(X)}\left\{\mathrm{E}_q f(\mathbf{x}) - D(q \parallel p)\right\}. \qquad (4)
\end{aligned}$$

For the first approximate equality, observe that the contribution of each empirical distribution q to the expectation is the realized value of $\mathbf{s}_T$ (given by (3)) multiplied by the probability of q (approximately equal to $e^{-T\times I(q)}$, by the large deviation principle). For the second approximate equality, note that all these contributions are exponential functions of T. When T is large, then all the contributions are dominated by the exponential function with the largest exponent.[3] Finally, by Sanov's Theorem (Cover and Thomas (1999), Section 11.4), there is a rate function given by $I(q) = D(q \parallel p)$—the probability of observing an empirical sequence q when draws are taken from the distribution p decays exponentially with the length of the sequence at rate equal to the Kullback-Leibler divergence of the empirical distribution q from the actual distribution p.[4] Applying this result and simplifying gives the final equality.

Combining (2) and (4) gives an approximation of the desired result (1). The approximation becomes arbitrarily sharp at T grows, giving the equality in (1).

3 Two Equivalences

We establish two equivalences between pairs of decision problems. These problems are built on the following common foundation. An agent chooses action $a \in A$ and receives payoff $u(a, \boldsymbol{\theta})$, where the payoff state $\boldsymbol{\theta}$ is a random variable drawn from the distribution p that attains values in Θ. The sets A and Θ are finite for sake of simplicity. We call the pair (p, u) a *problem*.

[3] This argument is referred to as the Laplace principle in the large deviation theory.

[4] To build intuition, suppose n independent draws are taken from a coin whose probability of a head is p. Then letting $q = k/n$, the log of the probability of k heads, divided by n, is given by

$$\begin{aligned}
\frac{1}{n}\ln\frac{n!}{k!(n-k)!}p^k(1-p)^{n-k} &= \frac{1}{n}\left(\ln n! - \ln k! - \ln(n-k)! + k\ln p + (n-k)\ln(1-p)\right) \\
&= \frac{1}{n}\left(n\ln n - k\ln k - (n-k)\ln(n-k) + k\ln p + (n-k)\ln(1-p) + O(\ln n)\right) \\
&= -q\ln q - (1-q)\ln(1-q) + q\ln p + (1-q)\ln(1-p) + \frac{O(\ln n)}{n} \\
&\approx -D(q \parallel p),
\end{aligned}$$

where the second equality follows from an application of Stirling's formula ($\ln n! = n\ln n - n + O(\ln n)$).

3.1 Expected Utility and Wishful Thinking

We first show that expected utility maximization is behaviorally equivalent to wishful thinking. The *expected utility maximizer* in problem (p, u) chooses an action in

$$A^*_{eu}(p, u) = \arg\max_{a \in A} \mathrm{E}_p\, u(a, \boldsymbol{\theta}). \tag{5}$$

The *wishful thinker* in problem (p, u) solves

$$\max_{a \in A, q \in \Delta(\Theta)} \left\{ \mathrm{E}_q\, u(a, \boldsymbol{\theta}) - D(q \parallel p) \right\}. \tag{6}$$

That is, the wishful thinker chooses a subjective belief q at distortion cost $D(q \parallel p)$, chooses an action a, and enjoys subjective expectation $\mathrm{E}_q\, u(a, \boldsymbol{\theta})$ formed under the distorted belief. We let $A^*_{wt}(p, u)$ denote the set of actions that, together with some subjective belief, optimize the objective in (6). Caplin and Leahy (2019) have introduced this model as a psychologically plausible procedure.[5]

Proposition 1. *Wishful thinking with utility function $u(a, \theta)$ is behaviorally equivalent to expected utility maximization with the monotonically transformed utility function $U(a, \theta) = e^{u(a,\theta)}$:*

$$A^*_{wt}(p, u) = A^*_{eu}(p, U).$$

Thus, even though the two models appear to be conceptually and procedurally distinct, they have the same normative appeal. In fact, an analyst, who does not observe the decision maker's utility function cannot distinguish the two models based on choice data.

Proof of Proposition 1. Since the logarithm is an increasing function, the expected utility maximizer in problem $(p, U) = (p, e^u)$ solves

$$\max_{a \in A} \ln \mathrm{E}_p\, e^{u(a, \boldsymbol{\theta})}.$$

By the Donsker-Varadhan Lemma, this is equivalent to

$$\max_{a \in A, q \in \Delta(\Theta)} \left\{ \mathrm{E}_q\, u(a, \boldsymbol{\theta}) - D(q \parallel p) \right\},$$

where the last optimization is the problem of the wishful thinker, as needed. $\square$

Strzalecki (2011) (pages 56–57), in the course of providing an axiomatic foundation for the multiplier preferences of Hansen and Sargent (2001) and Hansen et al. (2006), uses a similar argument to establish an equivalence between multiplier preferences and expected utility maximization.[6]

[5] Little (2021) studies a special case of the wishful thinking model in which the agent makes no action choice.

[6] Strzalecki (2011) shows that in an Anscombe and Aumann (1963) setting, multiplier preferences are the intersection of variational preferences (Maccheroni et al. (2006)) and second-order expected utility preferences (Ergin and

3.2 Growth-Optimal Portfolios and Rational Inattention

Now interpret A as a set of assets. The asset $a \in A$ yields the growth rate $u(a,\theta)$ and hence the gross return $e^{u(a,\theta)}$ in state θ. The portfolio that distributes an investor's wealth in proportions $\alpha(a)$ across assets $a \in A$ then enjoys the growth rate $\ln\left(\mathrm{E}_\alpha\, e^{u(\mathbf{a},\theta)}\right)$ in each state θ, where we formalize the portfolio as a random variable $\mathbf{a}$ drawn from distribution $\alpha \in \Delta(A)$. Following a large literature originating in Kelly (1956) and surveyed in Christensen (2005), we say that a mixed strategy $\alpha^* \in \Delta(A)$ is a *growth-optimal portfolio* for (p,u) if

$$\alpha^* \in \underset{\alpha \in \Delta(A)}{\arg\max}\, \mathrm{E}_p \ln\left(\mathrm{E}_\alpha\, e^{u(\mathbf{a},\boldsymbol{\theta})}\right). \tag{7}$$

The solution to (7) maximizes the expected growth rate of the value of the portfolio. Appealingly, an investor who rebalances her investment proportions according to α^* every period achieves at least as high an asymptotic growth rate of wealth as any other strategy; see e.g. Cover and Thomas (1999) (Theorem 16.3.1). See Section 4 for the evolutionary reinterpretation.

We show that the growth-optimal portfolio problem (7) is equivalent to an associated *rational inattention problem*. In the latter problem, as introduced by Matějka and McKay (2015), a decision maker chooses a state-contingent stochastic choice rule $q(a \mid \theta) \in \Delta(A)$ for each state θ and solves

$$\max_{q \in \Delta(A)^\Theta} \left\{\mathrm{E}_{p,q}\, u(\mathbf{a},\boldsymbol{\theta}) - I_{p,q}(\mathbf{a};\boldsymbol{\theta})\right\}, \tag{8}$$

where the optimization is over rules $q = (q(a \mid \theta))_{a,\theta}$. Matějka and McKay motivate the objective in (8) as an information acquisition problem: a more informed choice rule q achieves a higher expected payoff but incurs a higher information cost. The authors set this cost to the mutual information

$$I_{p,q}(\mathbf{a};\boldsymbol{\theta}) = H(\boldsymbol{\theta}) - H(\boldsymbol{\theta} \mid \mathbf{a}),$$

where H stands for (conditional) entropy.[7]

A marginal action distribution $q^*(a)$ is *induced by the solution of the rational inattention problem* if $q^*(a) = \mathrm{E}_p\, q^*(a \mid \boldsymbol{\theta})$ and the rule $q^*(a \mid \theta)$ solves (8).

Proposition 2. *An action distribution is induced by the solution of the rational inattention problem* (8) *if and only if it is the growth-optimal portfolio* (7)*.*

Again, the equivalence between the two problems holds despite their procedural differences: an investor in (7) cannot acquire information while the decision maker in (8) can. This proposition is mathematically equivalent to Lemma 2 in Matějka and McKay (2015), who offer a proof based on a

---

Gul (2009)). In a setting with only subjective uncertainty, as examined here, multiplier preferences are equivalent to expected utility maximization. More precisely, the multiplier problem $\max_{a \in A} \min_{q \in \Delta(\Theta)} \mathrm{E}_q\, u(a,\theta) + D(q \parallel p)$ (interpreted as calling for the maximizing agent to choose an action a, knowing that nature will then choose a belief q at cost $D(q \parallel p)$ that minimizes the agent's expected utility) has the same solution as the utility maximization problem $\max_{a \in A} \mathrm{E}_{p \in \Delta(\Theta)} -e^{-u(a,\theta)}$.

[7] Hence, $H(\boldsymbol{\theta}) = -\sum_\theta p(\theta) \ln p(\theta)$ and $H(\boldsymbol{\theta} \mid \boldsymbol{a}) = \sum_a q(a) H(\boldsymbol{\theta} \mid \boldsymbol{a} = a)$.

characterization (their Theorem 1) of the solution of the rational inattention problem. Matějka and McKay (2015) do not draw the connections to the investment problem or to the growth process. Section 4.2 and 4.3 illustrate how these connections can be useful.

Proof of Proposition 2. For each state $\theta \in \Theta$ and any strategy α, applying the Donsker-Varadhan Lemma to the random action $\mathbf{a}|\theta$ gives

$$\ln\left(\mathrm{E}_\alpha\, e^{u(\mathbf{a},\theta)}\right) = \max_{q(a|\theta)\in\Delta(A)} \left\{\mathrm{E}_{q(a|\theta)}\, u(\mathbf{a},\theta) - D\big(q(a \mid \theta) \parallel \alpha(a)\big)\right\}.$$

Taking an expectation over $\boldsymbol{\theta}$ with respect to p and then maximizing over $\alpha \in \Delta(A)$, the left side of this equation becomes the growth-optimal portfolio problem (7), which is accordingly equivalent to

$$\max_{\alpha\in\Delta(A),q\in\Delta(A)^\Theta} \left\{\mathrm{E}_{p,q}\, u(\mathbf{a},\boldsymbol{\theta}) - \mathrm{E}_p\, D\big(q(a \mid \boldsymbol{\theta}) \parallel \alpha(a)\big)\right\}. \tag{9}$$

It is now a straightforward calculation that (9) is equivalent to the rational inattention problem. In particular, the chain rule for Kullback-Leibler divergence implies that

$$\begin{aligned} D\big(p(\theta)q(a \mid \theta) \parallel p(\theta)\alpha(a)\big) &= D\big(p(\theta) \parallel p(\theta)\big) + \mathrm{E}_p\, D\big(q(a \mid \boldsymbol{\theta}) \parallel \alpha(a)\big) \\ &= D\big(q(a) \parallel \alpha(a)\big) + \mathrm{E}_{q(a)}\, D\big(q(\theta \mid \mathbf{a}) \parallel p(\theta)\big), \end{aligned}$$

where the left side is the divergence of two joint distributions, $q(a) = \mathrm{E}_p\, q(a \mid \boldsymbol{\theta})$ is the marginal action distribution for prior p and rule $q(a \mid \theta)$, and $q(\theta \mid a)$ is the conditional state distribution for a given action a. Since $D\big(p(\theta) \parallel p(\theta)\big) = 0$, we have

$$\mathrm{E}_p\, D\big(q(a \mid \boldsymbol{\theta}) \parallel \alpha(a)\big) = D\big(q(a) \parallel \alpha(a)\big) + \mathrm{E}_{q(a)}\, D\big(q(\theta \mid \mathbf{a}) \parallel p(\theta)\big).$$

Thus, the expected divergence in (9) is for a given rule $q(a \mid \theta)$ minimized by $\alpha(a) = q(a)$. Hence, expression (9) and in turn the growth-optimal portfolio problem (7) are equivalent to

$$\max_{q\in\Delta(A)^\Theta} \left\{\mathrm{E}_{p,q}\, u(\mathbf{a},\boldsymbol{\theta}) - \mathrm{E}_{q(a)}\, D\big(q(\theta \mid \mathbf{a}) \parallel p(\theta)\big)\right\}.$$

This last problem is equivalent to the rational inattention problem (8) as needed because

$$I_{p,q}(\boldsymbol{\theta};\mathbf{a}) = \mathrm{E}_{q(a)}\, D\big(q(\theta \mid \mathbf{a}) \parallel p(\theta)\big).$$

□

An analyst with access to the type of data typically associated with a revealed-preference exercise—a collection of problems and selected alternatives—would be unable to distinguish wishful thinking or rational inattention from their behaviorally equivalent counterparts. Richer data would give the analyst more power. The state dependent stochastic choice data of Caplin and Dean (2015) would make the information acquisition of rational inattention apparent. In a similar vein, data that included objective and elicited subjective beliefs could detect wishful thinking.

4 A Population Perspective

We have phrased the problems of Section 3 as problems of a single decision maker. We now consider a continuum of agents who repeatedly experience draws of the state. We consider two cases. In one case, each agent faces an independent draw, which leads to expected utility maximization. In another case, all agents face the same draw, which leads to the growth-optimal portfolio problem with hedging over the finite action set occurring on the aggregate level.

Time is discrete, with the length δ of the periods $t = 1, 2, \ldots$ in this section normalized to $\delta = 1$. In our leading interpretation, we normalize the size of the population and the initial wealth of each agent to one. Each agent chooses a private action $a \in A$ each period and enjoys wealth $w_t = e^{u(a,\boldsymbol{\theta}_t)} w_{t-1}$ at the end of each period t, where $\boldsymbol{\theta}_t$ is drawn from $p \in \Delta(\Theta)$, independently across periods.[8] Alternatively, we can interpret the process as demographic growth. We normalize the *initial* population size to one and interpret a as a characteristic a parent transmits to her offspring, with $e^{u(a,\theta)}$ then being the attendant expected number of surviving offspring and w_t the size of the population at time t.

In the *idiosyncratic* uncertainty variant, the payoff shocks $\boldsymbol{\theta}_t$ are independent across agents. We impose a law of large numbers on the continuous population. Hence, if all agents choose an action a in all periods, then the population wealth in period t is $\left(\mathrm{E}_p\, e^{u(a,\boldsymbol{\theta})}\right)^t$ with certainty. We seek an action that maximizes the growth rate of the population's wealth, and hence solves[9]

$$\max_{a \in A} \ln \mathrm{E}_p\, e^{u(a,\boldsymbol{\theta})}. \tag{10}$$

Since the logarithmic transformation does not affect the maximizer, this problem is equivalent to the expected utility maximization problem (5) with utility $U(a,\theta) = e^{u(a,\theta)}$.

In the *aggregate* uncertainty variant of the growth process, we assume that all agents in each period t are subject to the same period-dependent $\boldsymbol{\theta}_t$. We let each agent randomize their action according to a mixed strategy $\alpha \in \Delta(A)$ independently across agents and periods. We again assume that a law of large numbers applies to the continuous population. Hence, if the share of the agents choosing an action a is $\alpha(a)$, then the population's aggregate wealth is multiplied by factor $\mathrm{E}_\alpha\, e^{u(\mathbf{a},\theta)}$ in each state θ. We seek a strategy α^* that maximizes the long run growth rate of the population's wealth, that is, the solution of the growth-optimal portfolio problem (7):

$$\max_{\alpha \in \Delta(A)} \mathrm{E}_p \ln \left(\mathrm{E}_\alpha\, e^{u(\mathbf{a},\boldsymbol{\theta})} \right).$$

[8] Our investor relentlessly reinvests all of her wealth. Allowing her to consume a fixed proportion of her income in each period would add a constant to the growth objective, while preserving the results. Total reinvestment is appropriate in a biological interpretation of the growth process.

[9] Allowing for mixed actions, period-dependent, or history-dependent choice is inconsequential under idiosyncratic risk since an optimal history-independent and time-independent pure action always exists.

Our implicit assumption in adopting these objectives is that strategies maximizing the long run aggregate wealth of the population will dominate the population. This will be the case if these strategies are deliberately chosen by maximizing agents. Alternatively, this prevalence will emerge, no matter how strategies are chosen, if market forces cause wealth maximizing strategies to displace other strategies (cf. Blume and Easley (1992)), or if more successful strategies are imitated. Wealth maximizing strategies will automatically dominate if one samples behavior according to wealth, or if the growth process is interpreted as population growth (cf. Tanny (1981)).

Under idiosyncratic risk, every finite sequence of states is realized for some segment of the population and nearly all aggregate growth occurs along atypical sequences of states. We characterize these atypical sequences by applying the Donsker-Varadhan Lemma to the stochasticity of states, leading to the wishful thinking representation. Under aggregate risk, the growth rate is maximized along the typical sequences of shocks, but some fractions of the population enjoy large deviations in the sequences of their actions. Proposition 3.2 characterizes these atypical sequences by applying the Donsker-Varadhan Lemma to the stochasticity of actions, leading to the rational inattention representation.

4.1 Idiosyncratic Uncertainty

This subsection connects the equivalence between expected utility maximization and wishful thinking established in Proposition 1 to growth under idiosyncratic risk. We show that a wealth-weighted sampling of the population's realized *states* reproduces the optimally distorted distribution q^* from the wishful thinking problem. Population wealth becomes concentrated on a vanishingly small subpopulation whose empirical distribution of realized states matches q^* and whose size shrinks at a rate equal to the belief distortion expense in the wishful thinking representation.

Consider a large time horizon t. Under idiosyncratic risk, any empirical distribution q of states is enjoyed by a fraction of population of approximate size $e^{-t\times D(q\|p)}$ up to period t (again ignoring divisibility concerns). If the population chooses an action a, then agents from the subpopulation experiencing empirical distribution q have enjoyed growth rate $\mathrm{E}_q\, u(a,\boldsymbol{\theta})$ and hold aggregate wealth $\exp\left[t\times\left(\mathrm{E}_q\, u(a,\boldsymbol{\theta}) - D(q\parallel p)\right)\right]$ at the end of the period t. The first term in the exponent describes the growth rate of per capita wealth of this subpopulation, while the second describes the size of the subpopulation.

The law of large numbers can be misleading when examining the distribution of wealth. Some sequences of realized states will be exponentially rare, but these sequences may beget exponentially large wealth. The aggregate wealth of the entire population will be approximated by the wealth of the subpopulation whose sequence of realized states appropriately balances these two forces, i.e., that has enjoyed the empirical distribution q^* that maximizes $\mathrm{E}_q\, u(a,\boldsymbol{\theta}) - D(q\parallel p)$. An ever-growing fraction of the population's wealth, approaching one in the limit, will be held by this ever-shrinking subpopulation, whose size becomes negligible in the limit, while the wealth of the residual population also becomes negligible. To secure a nonnegligible wealth share, an agent needs to be lucky, in order to experience the optimal large deviation of shocks, and needs to be prepared, in the sense of choosing an action that maximizes the growth rate under this deviation.

To make these remarks precise, define the *sampled* distribution as

$$\tilde{q}(\theta) = \frac{p(\theta)e^{u(a^*,\theta)}}{\mathrm{E}_p\, e^{u(a^*,\boldsymbol{\theta})}}, \tag{11}$$

where a^* is the optimal action from the idiosyncratic growth problem (10). To interpret the sampled distribution, think of sampling a dollar uniformly from the population wealth at the end of any given period. The numerator is proportional to aggregate wealth generated by state θ and the denominator is proportional to the total wealth. Hence, $\tilde{q}$ is the distribution of the payoff state θ of the dollar's owner from this period. A naive frequentist who learns the distribution of states by sampling agents proportionally to their wealth learns the sampled distribution. Alternatively, a population biologist who took an unweighted sample of survivors would learn the sampling distribution.

Corollary 1. *The sampled distribution $\tilde{q}$ equals the distorted distribution q^* chosen by the wishful thinker.*

Proof. The first-order conditions in the wishful thinking optimization problem,

$$\begin{aligned} \max_{q} \quad & \left\{ \sum_{\theta\in\Theta} q(\theta)u(a^*,\theta) - D(q \parallel p) \right\} \\ \text{s.t.:} \quad & \sum_{\theta\in\Theta} q(\theta) = 1, \end{aligned}$$

have a unique solution, given by

$$q^*(\theta) = \frac{p(\theta)e^{u(a^*,\theta)}}{\mathrm{E}_p\, e^{u(a^*,\boldsymbol{\theta})}}, \tag{12}$$

which coincides with the sampled distribution $\tilde{q}$ from (11). □

For intuition for Corollary 1, think of sampling a dollar in period t, tracing the predecessor dollar whose investment generated the terminal dollar, and tracing the predecessor's predecessor, and so on. The payoff shocks experienced by the owner of these dollars are independently drawn from the sampled distribution $\tilde{q}$. By the law of large numbers, the empirical state distribution experienced by the owner of the random dollar sampled at t is likely to be near $\tilde{q}$ when t is large. But for large t, nearly all wealth is owned by agents who have enjoyed the optimal large deviation q^*, which is also the distorted distribution of the wishful thinker, ensuring $\tilde{q} = q^*$.

The next result describes a debiasing procedure that arises endogenously in our setting and is akin to the robust-control approach of Hansen and Sargent (2011).

Corollary 2. *Let a^* be the optimal action and q^* be the sampled distribution in the growth process with prior p and growth rates $u(a,\theta)$. Then,*

$$p \in \arg\min_{p'} \left\{ \mathrm{E}_{p'}\, u(a^*,\boldsymbol{\theta}) + D(p' \parallel q^*) \right\}. \tag{13}$$

In Hansen and Sargent (2011), an agent faces a generic decision problem under uncertainty and observes the distribution q^*. The agent assumes that the true distribution of payoff states is within a Kullback-Leibler-divergence ball of distributions around q^*, and selects the worst case distribution from this ball. Our problem (13) coincides with this worst case optimization relaxed by the Lagrange method with the Lagrange multiplier normalized to one.

To prove Corollary 2, we fix the optimal action a^* and trace the wealth of agents from period t backwards. The states are distributed according to the sampled distribution q^* and the wealth of each agent shrinks by $e^{u(a^*,\boldsymbol{\theta}_t)}$ each period; that is, it grows with rate $-u(a^*,\boldsymbol{\theta}_t)$. The proof consists of characterizing the sampled distribution of this backward process.

Proof of Corollary 2. Consider a growth process with prior distribution $q^*(\theta)$ and growth rates $u'(a,\theta) = -u(a^*,\theta)$. By (12), its sampled distribution is

$$\frac{q^*(\theta)e^{u'(a^*,\theta)}}{\mathrm{E}_{q^*}\, e^{u'(a^*,\boldsymbol{\theta})}} = \frac{q^*(\theta)e^{-u(a^*,\theta)}}{\mathrm{E}_{q^*}\, e^{-u(a^*,\boldsymbol{\theta})}} = p(\theta)e^{u(a^*,\theta)}e^{-u(a^*,\theta)} = p(\theta).$$

By Corollary 1, the sampled distribution of this growth process is in

$$\arg\max_{p'\in\Delta(\Theta)}\{-\mathrm{E}_{p'}\, u(a^*,\boldsymbol{\theta}) - D(p' \parallel q^*)\} = \arg\min_{p'\in\Delta(\Theta)}\{\mathrm{E}_{p'}\, u(a^*,\boldsymbol{\theta}) + D(p' \parallel q^*)\}.$$

Hence, p is the minimizer of the problem on the right side, as needed. $\square$

The wishful thinking representation is informative of an additional property of the stochastic growth process. A dollar sampled at any given period has originated from an investment with expected return $\mathrm{E}_{q^*}\, u(a^*;\boldsymbol{\theta})$ given by the sampled distribution q^*. However, the growth rate of the aggregate wealth is only $\mathrm{E}_{q^*}\, u(a^*;\boldsymbol{\theta}) - D(q^* \parallel p)$. This difference in growth rates implies that wealth is becoming increasingly concentrated. Indeed, nearly all wealth is concentrated among the subpopulation that has enjoyed the optimal deviation, and its size vanishes exponentially at rate $D(q^* \parallel p)$. Thus:

Corollary 3. *The distortion expense $D(q^* \parallel p)$ from the wishful thinking optimization identifies the degree of wealth concentration in the idiosyncratic growth process, as all wealth becomes concentrated in a subpopulation whose size shrinks at rate $D(q^* \parallel p)$.*

4.2 Aggregate Uncertainty

We now assume that in each period a single payoff state θ_t is drawn that applies to all individuals. We show that a wealth-weighted sampling of the population's realized *actions* reproduces the solution of the rational inattention problem. The population wealth becomes concentrated on a subpopulation whose size exponentially vanishes at rate equal to the information expense in the associated rational inattention problem.

Analogously to the sampled distribution from Subsection 4.1, we define the *sampled choice rule*

$$\tilde{q}(a \mid \theta) = \frac{\alpha^*(a) e^{u(a,\theta)}}{\mathrm{E}_{\alpha^*}\, e^{u(\mathbf{a},\theta)}} \text{ for all } \theta, \tag{14}$$

where the mixed strategy α^* is the maximizer of the growth-optimal portfolio problem (7). We write $\tilde{q}$ for the state-contingent system $(\tilde{q}(a \mid \theta))_{\theta \in \Theta}$.

For interpretation again note that the owner of a dollar uniformly selected from the wealth of population at the end of any period t in which $\boldsymbol{\theta}_t = \theta$ has chosen action a with probability $\tilde{q}(a \mid \theta)$. While the actions chosen at the beginning of any period are uncorrelated with the payoff state, the sampled choice rule oversamples successful actions and thus features correlation induced by the growth process itself. In the biological interpretation, sampling of surviving offspring yields the sampled choice rule

Corollary 4. *A rule solves the rational inattention problem if and only if it is the sampled rule* (14).

Proof. We have shown in Proposition 2 that α^* equals the marginal action distribution induced by the choice rule that solves the rational inattention problem. Matějka and McKay (2015) show in their Theorem 1 that given $q^*(a) = \alpha^*(a)$ the optimal choice rule $q^*(a \mid \theta)$ satisfies (14). □

The equivalence in Corollary 4 arises even though the agents in the growth process cannot engage in information acquisition, in contrast to the decision maker in the rational inattention problem. Again, the equivalence can be understood in terms of large deviation theory. For any finite time horizon, various fractions of the population in the growth process with aggregate risk experience empirical correlations between their actions and states by luck. Some of these random correlations boost growth. Asymptotically, nearly all population wealth is concentrated among the subpopulation that has enjoyed sufficiently favorable but not too unlikely deviations. The tradeoff between the growth advantage of the deviation and its rarity is equivalent to the tradeoff between the benefit and cost of information in the rational inattention problem.

An empirical researcher who samples agents proportionally to their wealth learns the joint distribution of actions and states that coincides with that of the rationally inattentive decision maker. The researcher might be tempted to conclude that the agents acquire partial information about the states in a cost-benefit calculation, even though the observed correlation pattern stems from uninformed hedging and selection bias.

Again, the information-acquisition expense from the rational inattention problem corresponds to wealth concentration in the growth process:

Corollary 5. *Asymptotically, almost all wealth is held by a fraction of population whose size shrinks exponentially at rate* $I_{p;q^*}(\boldsymbol{\theta}; \boldsymbol{a})$.

Aggregate wealth becomes concentrated on the subpopulation that has experienced the optimal large deviation in their sequence of actions. The members of this subpopulation have been blessed by incredible luck, and appear to have based their choices on (partial) information about the realized sequence of states.

4.3 Optimal Hedging

The growth-optimal portfolio involves hedging. For example, suppose an action matching the state gives a positive return, while an action that does not match the state gives a gross return of zero. Though matching the a priori most likely state maximizes expected return, it also ensures bankruptcy. Instead, growth is maximized by matching each state with probability equal to its prior probability. See Bergstrom (2014), Robson (1996) and Robson and Samuelson (2011) for the evolutionary context.

Solving for optimal hedging is in general difficult. The equivalence between the growth-optimal portfolio and the rational inattention problems allows us to transfer the methods derived for the latter problem to the first one. To illustrate this, we use rational inattention insights to derive sharp characterizations of the impact of public information and belief misspecification on growth and inequality. The characterizations are based on an analysis of posterior beliefs. These naturally arise in the rational inattention setting but their relevance for the growth process is less immediate.

Our baseline problem is the growth-optimal portfolio problem (7) in which agents have no information beyond their prior. In our first application, we extend this baseline by letting agents condition their actions on a public signal. The payoff states $\boldsymbol{\theta}_t$ and public signals $\mathbf{y}_t$ are drawn from a known joint distribution $p(\theta, y)$, independently across periods. The agents optimize over stochastic choice rules $\alpha(a \mid y)$ that specify conditional action distributions for each signal realization y. The *growth-optimal portfolio problem with public information* is then

$$\max_{\alpha} \mathrm{E}_{p(\theta,y)} \ln \left(\mathrm{E}_{\alpha(a|\mathbf{y})} \, e^{u(\mathbf{a},\boldsymbol{\theta})} \right).$$

We impose a regularity condition on the public signal. Let $q^*(a \mid \theta)$ be the sampled choice rule in the baseline problem (7) without public information and define posterior distributions $q^*(\theta \mid a) = q^*(a \mid \theta)\frac{p(\theta)}{q^*(a)}$ for each action in the support of α.[10] We require that the public signal is not too informative:

Regularity condition 1: The conditional state distribution $p(\theta \mid y)$ is in the convex hull of $q^*(\theta \mid a)$, $a : q^*(a) > 0$, for each realization y.

We let V^* and V^{**} denote the growth rates achieved in the settings without and with information, respectively. We continue to write $\alpha^*(a)$ for the optimal strategy in the baseline setting and write $\alpha^{**}(a \mid y)$ for the optimal rule in the setting with public information; $\alpha^{**}(a) = \mathrm{E}\, \alpha^{**}(a \mid \mathbf{y})$ is the marginal action distribution.

Proposition 3. *Under the regularity condition 1:*

[10] These posteriors can be interpreted in the growth context as the state distributions conditional on a dollar randomly sampled at the end of any period originating from action a.

3.1 The provision of public signal increases the growth rate of wealth by the mutual information between the payoff state and the signal:

$$V^{**} - V^{*} = I(\boldsymbol{\theta}; \mathbf{y}).$$

3.2 The marginal action distribution is unaffected by the information provision:

$$\alpha^{*}(a) = \alpha^{**}(a) \text{ for all } a \in A.$$

3.3 The provision of public information diminishes the exponential rate of wealth concentration.

In our second application, we keep the same baseline and compare it to a setting in which the agents hold a biased prior belief p' and receive no information beyond the prior. Letting $\alpha_p^* \in \Delta(A)$ be the optimal strategy for the prior p,

$$L\left(p, p'\right) = \mathrm{E}_p \ln \mathrm{E}_{\alpha_p^*} e^{u(\mathbf{a},\boldsymbol{\theta})} - \mathrm{E}_p \ln \mathrm{E}_{\alpha_{p'}^*} e^{u(\mathbf{a},\boldsymbol{\theta})}$$

is the reduction of the growth rate caused by the optimization for the misspecified prior p' when the objective prior is in fact p.

We require that the misspecification is not too large.

Regularity condition 2: The misspecified belief p' is in the convex hull of $q^*(\theta \mid a)$, $a : q^*(a) > 0$.

Proposition 4. *Under the regularity condition 2, misspecification of prior belief decreases the growth rate of wealth by $L(p, p') = D(p \parallel p')$.*

The value of information and the loss from misspecification in a static single-person decision problem depend on the specific payoff function. In contrast, the implications of information and misspecification are strikingly simple in the context of growth; they are independent of the specific payoff function (under the regularity conditions). Moreover, the provision of (weak enough) public information introduces correlation between actions and states in the growth-optimal portfolio problem, but does not change the marginal action distribution. Thus, (weakly informative) public persuasion campaigns are ineffective in achieving state-independent goals in this setting.

We use the next lemma to prove both propositions. Let the value function defined by $V^*(p) = \max_\alpha \mathrm{E}_p \ln \left(\mathrm{E}_\alpha e^{u(\mathbf{a},\boldsymbol{\theta})}\right)$ map each prior p to the optimal growth rate in the baseline setting.

Lemma 2. *The value function restricted to the convex hull of optimal posteriors from the baseline setting satisfies $V^*(p) = g(p) - H(p)$, where g is a linear function.*

The lemma follows from the "locally invariant posteriors" property of Caplin et al. (2022).

Proof. Caplin et al. (2022) observe that the rational inattention problem (8) is equivalent to

$$\begin{aligned} \max_{\mathbf{r}} \quad & \mathrm{E}\left[v(\mathbf{r}) + H(\mathbf{r}) - H(p)\right] \\ \text{s.t.:} \quad & \mathrm{E}\,\mathbf{r} = p, \end{aligned}$$

where $\mathbf{r} \in \Delta(\Delta(\Theta))$ is a random posterior and $v(r) = \max_{a \in A} \mathrm{E}_r\, u(a, \boldsymbol{\theta})$ for a realized posterior r. The term $\mathrm{E}\left[H(p) - H(\mathbf{r})\right]$ is the mutual information for the random posterior $\mathbf{r}$. The constraint is the Bayes plausibility condition.

Caplin, Dean and Leahy point out that the support of the optimal random posterior equals the set of beliefs that support the tangent hyperplane of the concavification of $v(\mathbf{r}) + H(\mathbf{r})$ above the prior belief. An implication, dubbed the locally invariant posteriors property by these authors, is that if $\mathbf{r}^*$ is the optimal random posterior for prior p, then the optimal random posterior for any prior p' in the convex hull of support of $\mathbf{r}^*$ has the same support as $\mathbf{r}^*$. Letting $\mathbf{r}^*(p)$ be the optimal random posterior for a prior p, this implies that $\mathrm{E}\left[v\left(\mathbf{r}^*(p)\right) + H\left(\mathbf{r}^*(p)\right)\right]$ is linear in p. Thus, the lemma holds with $g(p) = \mathrm{E}\left[v\left(\mathbf{r}^*(p)\right) + H\left(\mathbf{r}^*(p)\right)\right]$. □

The locally invariant posteriors property allows us to reformulate our regularity conditions. They are equivalent to requirements that information provision, respectively misspecification, are sufficiently small so that they do not affect the set of actions that are chosen with positive probability.

Proof of Proposition of 3. The growth rate achieved under provision of the public signal y is a convex combination of the growth rates achieved at priors $p(\theta \mid y)$:

$$V^{**} = \mathrm{E}\, V^*\left(p(\theta \mid \mathbf{y})\right).$$

Claim [3.1] follows:

$$\begin{aligned} V^{**} - V^*(p) &= \mathrm{E}\, V^*\left(p(\theta \mid \mathbf{y})\right) - V^*(p) \\ &= H(p) - \mathrm{E}\, H\left(p(\theta \mid \mathbf{y})\right) \\ &= I(\boldsymbol{\theta}; \mathbf{y}), \end{aligned}$$

where the expectations are with respect to the signal $\mathbf{y}$. The second equality follows from Lemma 2: Since the $p(\theta \mid \mathbf{y})$ are in the convex hull of the optimal posteriors of the baseline problem by the regularity condition, $V^*(p) = -H(p)$ up to an irrelevant linear term. Indeed, the linear term in V^* does not affect the difference because $\mathrm{E}\, p(\theta \mid \mathbf{y}) = p(\theta)$.

Claim [3.2] follows from the fact that the prior probability of reaching any of the posteriors, and hence the probability of any of the actions, is unaffected by the signal provision. For Claim [3.3], note that the signal provision decreases the rate at which wealth concentrates by $I(\boldsymbol{\theta}; \mathbf{y}) \geq 0$. □

Proof of Proposition of 4. The growth rate $\mathrm{E}_p \ln \mathrm{E}_{\alpha^*(p')}\, e^{u(\mathbf{a}, \boldsymbol{\theta})}$ is a linear function of p, the graph of which is a hyperplane tangent to $V^*(p)$ at p'. The linearity follows from the fact that it is an expectation with respect to p and the tangency follows from the optimality of $\alpha^*(p')$ at prior p'. Thus, $L\left(p, p'\right)$ is the error of the linear approximation of V^* around p' evaluated at p. By Lemma 2, and since the linear term of V^* does not affect this error, $L\left(p, p'\right)$ is the error of the linear approximation of $-H(p)$ around p' evaluated at p. Thus, $L(p, p') = D(p \parallel p')$, because Kullback-Leibler divergence is a Bregman divergence associated with function $-H$ (Bregman (1967)). $\square$

Cabrales et al. (2013) show that signal structures are ranked by mutual information in a class of static investment problems. Barron and Cover (1988) establish that in the absence of our regularity conditions, the characterizations derived here continue to be upper bounds on the value of information and on loss from misspecification, respectively.[11] Frick et al. (2021) provide a distinct problem-independent characterization of the loss from misspecification.

5 Discussion

Beginning with Robson (1996), a large literature (discussed in Robson and Samuelson (2011)) has examined models of biological growth subject to idiosyncratic or aggregate environmental risks. The basic finding of this literature, reflecting the concavity of the log function that appears in (7), is that evolution will select for greater aversion to aggregate than to idiosyncratic risk.

We can contribute to the above insight by making the period length δ explicit, effectively changing $u(a, \theta)$ to $\delta u(a, \theta)$. The optimization under the idiosyncratic risk then becomes equivalent to

$$\max_{a \in A, q \in \Delta(\Theta)} \left\{ \mathrm{E}_q\, u(a, \boldsymbol{\theta}) - \frac{1}{\delta} D(q \parallel p) \right\},$$

and the optimization under the aggregate risk becomes

$$\max_{q \in \Delta(A)^\Theta} \left\{ \mathrm{E}_{p,q}\, u(\mathbf{a}, \boldsymbol{\theta}) - \frac{1}{\delta} I_{p,q}(\mathbf{a}; \boldsymbol{\theta}) \right\}.$$

Since the weight, $1/\delta$, of the two information-theoretic terms explodes as the period length vanishes, the choices arising in both problems converge to the choices made in the expected utility maximization $\max_a \mathrm{E}_p\, u(a, \boldsymbol{\theta})$ for generic priors p. Hence, the growth-optimal portfolio involves no hedging when periods are sufficiently short. Robatto and Szentes (2017) and Robson and Samuelson (2019) discuss alternative conditions under which optimal choices under idiosyncratic and aggregate risk do or do not coincide.[12]

A natural extension of our setting would allow for serial correlations of the payoff states. We conjecture that the optimal choice arising under serial correlations can be represented by the dynamic extension of the rational inattention problem studied in Steiner et al. (2017). When the payoff states are serially correlated, then optimal hedging may involve endogenous serial action correlations because wealth accumulation is enhanced by serial correlations of the growth rates.

[11] Our regularity condition is automatically satisfied in a special setting in which the investor's wealth vanishes unless she correctly guesses the state (the so-called horse race setting) because the optimal posteriors are fully informed in this case. The value of information has been characterized for this special case by Kelly (1956).

[12] In the limit here, the distribution of unit-time growth rates converges to a point mass. The finance literature (Merton, 1969, 1975) considers an alternative limit with unit-time growth rates remaining non-degenerate and incremental growth rates approaching Brownian motion. An alternative approach uses Poisson processes; the model of Robatto and Szentes (2017) fits into this latter framework.

Such serially correlated hedging may lead to endogenous inertia of behavior, relevant for instance in the macroeconomic context of sticky prices.

References

Anscombe, F. J. and R. J. Aumann (1963). A definition of subjective probability. *Annals of Mathematical Statistics 34*(1), 199–205.

Barron, A. R. and T. M. Cover (1988). A bound on the financial value of information. *IEEE Transactions on Information Theory 34*(5), 1097–1100.

Benhabib, J. and A. Bisin (2018). Skewed wealth distributions: Theory and empirics. *Journal of Economic Literature 56*(4), 1261–91.

Bergstrom, T. C. (2014). On the evolution of hoarding, risk-taking, and wealth distribution in nonhuman and human populations. *Proceedings of the National Academy of Sciences 111*(Supplement 3), 10860–10867.

Blume, L. and D. Easley (1992). Evolution and market behavior. *Journal of Economic theory 58*(1), 9–40.

Bregman, L. M. (1967). The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR computational mathematics and mathematical physics 7*(3), 200–217.

Cabrales, A., O. Gossner, and R. Serrano (2013). Entropy and the value of information for investors. *American Economic Review 103*(1), 360–77.

Caplin, A. and M. Dean (2015). Revealed preference, rational inattention, and costly information acquisition. *American Economic Review 105*(7), 2183–2203.

Caplin, A., M. Dean, and J. Leahy (2022). Rationally inattentive behavior: Characterizing and generalizing Shannon entropy. *Journal of Political Economy 130*(6), 1676–1715.

Caplin, A. and J. Leahy (2019). Wishful thinking. Technical report, National Bureau of Economic Research.

Christensen, M. M. (2005). On the history of the growth optimal portfolio. Technical report, University of Southern Denmark.

Cover, T. M. and J. A. Thomas (1999). *Elements of Information Theory*. John Wiley & Sons.

Dillenberger, D., A. Postlewaite, and K. Rozen (2017). Optimism and pessimism with expected utility. *Journal of European Econonic Association 15*(5), 1158–1175.

Dupuis, P. and R. S. Ellis (1997). *A Weak Convergence Approach to the Theory of Large Deviations*. John Wiley & Sons.

Ergin, H. and F. Gul (2009). A theory of subjective compound lotteries. *Journal of Economic Theory 144*(3), 899–929.

Frick, M., R. Iijima, and Y. Ishii (2021). Welfare comparisons for biased learning. Technical report, Yale University, Yale University and Penn State University.

Hansen, L. P. and T. J. Sargent (2001). Robust control and model uncertainty. *American Economic Review 91*(2), 60–66.

Hansen, L. P. and T. J. Sargent (2011). *Robustness*. Princeton university press.

Hansen, L. P., T. J. Sargent, G. A. Turmuhambetova, and N. Williams (2006). Robust control and model uncertainty. *Journal of Economic Theory 128*(1), 45–90.

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

Kelly, J. (1956). A new interpretation of information rate. *The Bell System Technical Journal*.

Little, A. T. (2021). Detecting motivated reasoning. Technical report, OSF Preprints, July 15.

Maccheroni, F., M. Marinacci, and A. Rustichini (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica 74*(6), 1447–1498.

Matějka, F. and A. McKay (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review 105*(1), 272–298.

Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *The review of Economics and Statistics*, 247–257.

Merton, R. C. (1975). Optimum consumption and portfolio rules in a continuous-time model. In *Stochastic optimization models in finance*, pp. 621–661. Elsevier.

Robatto, R. and B. Szentes (2017). On the biological foundation of risk preferences. *Journal of Economic Theory 172*, 410–422.

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

Robson, A. J. and L. Samuelson (2011). The evolutionary foundations of preferences. *Handbook of Social Economics 1*, 221–310.

Robson, A. J. and L. Samuelson (2019). Evolved attitudes to idiosyncratic and aggregate risk in age-structured populations. *Journal of Economic Theory 6*(181), 44–81.

Robson, A. J. and B. Szentes (2014). A biological theory of social discounting. *American Economic Review 104*(11), 3481–3497.

Sadowski, P. and T. Sarver (2021). An evolutionary perspective on updating risk and ambiguity preferences. Technical report, Duke University.

Steiner, J. and C. Stewart (2016). Perceiving prospects properly. *American Economic Review 107*(7), 1601–1631–553.

Steiner, J., C. Stewart, and F. Matějka (2017). Rational inattention dynamics: Inertia and delay in decision-making. *Econometrica 85*(2), 521–553.

Strzalecki, T. (2011). Axiomatic foundations of multiplier preferences. *Econometrica 79*(1), 47–73.

Tanny, D. (1981). On multitype branching processes in a random environment. *Advances in Applied Probability 13*(3), 464–497.

Williams, N. (2004). Small noise asymptotics for a stochastic growth model. *Journal of Economic Theory 119*(2), 271–298.