Growth and Likelihood*

Larry Samuelson
Yale University

Jakub Steiner
University of Zurich and CERGE-EI

July 31, 2023

Abstract

We examine a stochastic growth process that can alternatively be interpreted as a model of economic growth, financial portfolio management, statistical inference, or biological population growth. For the economic interpretation, we find that the growth-maximizing policy satisfies a meritocracy principle: it minimizes the discrepancy between the resource shares allocated to the agents and the agents' "merits." For the statistical interpretation, the setting is equivalent to a model of predictive coding, in which a misspecified system maximizes the fit of data. A consistency principle analogous to the meritocracy principle requires the optimal fit to minimize a degree of Bayes inconsistency.

1 Introduction

This paper examines optimal redistribution policies within an abstract model of stochastic economic growth. We introduce a concept of merit that outlines the various economic agents' contributions to the growth process. Our primary finding is that the growth-maximizing policy satisfies a meritocracy principle, minimizing the discrepancy between the allocation of aggregate wealth and the economic agents' merits.

Rare events, referred to as large deviations in probability theory, play an important role in stochastic growth. We visualize large deviations by introducing *paths of money*: stochastic descriptions of the circulation of wealth in the

*We have benefited from the comments of Mira Frick, Tomáš Gavenčiak, Ben Hébert, Ryota Iijima, Yuhta Ishii, Roman Kotecký, Jan Kulveit, Rava da Silveira, Ludvig Sinander, Colin Stewart and various seminar and workshop audiences. This work was supported by ERC grant 770652. We have used artificial intelligence GPT-4 for language editing.

economy. Certain fractions of wealth follow atypical paths that reach individuals mostly in those periods in which these individuals happen to enjoy high returns. The mass of initial dollars following such a lucky path diminishes exponentially with the considered time horizon, but such lucky dollars enjoy an extraordinary long-run return. Asymptotically, all wealth is generated by a specific path of money whose wealth accumulation swamps that of all other paths.

We represent growth maximization as a joint optimization over policies and money paths. We show that the growth-maximizing policy minimizes the Kullback-Leibler divergence between the policy and the path of money circulation it generates. Our meritocracy principle is then an applied corollary of this general consistency result.

Discussions of meritocracy typically focus on how an individual's initial condition, productivity, and luck should influence their wealth.[1] A meritocracy principle that induces maximal economic growth separates out all luck—nonpersistent stochastic elements of outcomes—since these do not predict future productivity. Perhaps surprisingly, however, the growth-maximizing meritocracy principle does not eliminate wealth transfers from assessments of merits. The principle rewards individuals based on a combination of transfers they enjoy and their productivity. The growth-maximizing redistribution policy is a fixed point that is generous to those who are both productive and receive generous treatment from this policy. We do not claim that our notion of merit justifies inequality or provides a prescription for redistribution, but only that it helps identify which policies enhance growth.

The fixed-point aspect of the principle implies that the growth-maximizing distribution of wealth is particularly sensitive to productivity shifts. An increase in an individual's productivity initiates a positive feedback loop: higher merit leads to a more favorable allocation of wealth, which in turn further boosts the individual's merit. In the context of financial portfolio management, we find that this feedback loop causes the growth-maximizing portfolio to exhibit a strong sensitivity to exogenous changes in asset returns.

Stochastic growth processes, involving the multiplication of random variables, arise in a number of seemingly disparate contexts. In economics, accumulated wealth is the product of random returns. Analogously, in statistics, the

[1] The current economic debate on meritocracy is primarily empirical and experimental. Almås et al. (2020), Cappelen et al. (2023) and Andre (2022) explore how third-party redistribution choices are influenced by productivity, luck, and initial conditions. Alesina and Angeletos (2005) propose a theoretical model where fairness attitudes are endogenous, resulting in multiple equilibria.

likelihood of a sample is the product of the likelihoods of a collection of random observations. In both cases, interest centers on the policy—a wealth allocation rule or a collection of estimated parameters—that maximizes the growth rate.

We demonstrate that the problem of finding a redistribution policy that maximizes economic growth is equivalent to a problem of predictive coding, studied at the intersection of statistics, cognitive science and machine learning. An economy that is optimized for maximal growth must ensure the highest possible degree of alignment between the distribution of wealth and the merit of its recipients. Similarly, a cognitive system designed for optimal data fitting must maximize the consistency between its prior belief and the expected posterior belief; the meritocracy principle translates in the cognitive domain to a constrained form of Bayes' consistency.

Section 3 presents the basic characterization of growth-maximizing policies in terms of merit in the context of economic growth. Section 4 introduces paths of money and provides the proof of our central merit-principle result. We then derive an additional merit-based characterization and examine an iterative procedure leading to optimality. Section 5 establishes connections to predictive coding. Section 6 discusses sufficient conditions and shows that rational inattention arises as a special case of our setting.

2 Related Literature

Rising inequality has directed renewed interest to the extensive literature on the relationship between inequality and growth. Theoretical models, reviewed in Aghion et al. (1995), identify avenues by which inequality can both enhance and inhibit growth. The empirical evidence is inconclusive, with studies like Barro (2000), Berg et al. (2018), and Saez and Zucman (2020) presenting conflicting results. Stiglitz (1969)'s neoclassical argument suggests that redistribution from the rich to the poor fosters growth by alleviating diminishing returns and credit constraints. Constraints on investment opportunities play a role in our model, while returns are stochastic and linear.

The information-theoretic literature, originating in Kelly (1956), considers the growth rate of financial portfolios. If one reinterprets the individuals in our growth problem as financial assets, wealth redistribution becomes analogous to financial hedging. The naive meritocracy principle generalizes Kelly's proportional betting strategy to accommodate general and endogenous returns and constraints on portfolio.

Our characterization of the growth-maximizing policy in terms of Kullback-Leibler divergence between two sets of distributions follows the approach of information geometry (Csiszar and Tusnady, 1984). Within the same approach, the so-called Blahut-Arimoto algoritm (Cover and Thomas, 2006, p. 334) characterizes the efficiency frontier of noisy communication in information theory and, analogously, the optimal stochastic choice function in rational inattention problems.

Our merit distribution concept is mathematically equivalent to a Bayesian update. Corollary 3 connects the merit distribution to variational Bayesian methods, as seen in Jordan et al. (1999). The predictive coding literature, such as Dayan et al. (1995) and Friston (2005), merges variational Bayesian updating with misspecified statistical learning, originating in Berk (1966) and White (1982). In Section 5, we show that maximizing economic growth is equivalent to the problem of predictive coding.

Recent work of Aridor et al. (2020, 2023) adapts the model of predictive coding—also known as variational autoencoder in the machine-learning literature—to capture single-person information processing. Instead, our work emphasizes the connection between the variational autoencoder and population-wide economic growth or financial hedging.

This paper builds upon our prior work in Robson et al. (2023) by introducing general constraints on economic policies, proving the myopic meritocracy principle, and characterizing growth-maximizing policies through a joint optimization involving money paths, as detailed in Proposition 3. Robson et al. (2023) showed that if there are no constraints on the allocation of wealth, then maximizing the growth rate of wealth is equivalent to selecting the optimal learning policy in a rational inattention model. Revisiting this result from our current perspective, we show that redistribution serves as an alternative to active learning about individuals' returns.

3 Stochastic Growth

3.1 Generalized Joint Distributions

We begin by introducing notation for a system of non-negative functions that we use throughout our analysis. Consider arbitrary finite sets I and Ω. The primitives of the system are a probability distribution on I, denoted by $\mathbf{q}(i)$, and a non-negative function $I \times \Omega \to \mathbb{R}_+$, denoted by $q(\omega \mid i)$. Despite the suggestive notation, $q(\omega \mid i)$ is not a conditional distribution over Ω, as it need not be normalized.

We adopt notation analogous to that of standard probability distributions and define a system of non-negative functions as follows:

$$\begin{aligned} q(i, \omega) &:= \mathbf{q}(i) q(\omega \mid i) \\ q(\omega) &:= \sum_i q(i, \omega) \\ \mathbf{q}(i \mid \omega) &:= \frac{q(i, \omega)}{q(\omega)}. \end{aligned}$$

We employ the same symbol q to denote any of the functions within this system, relying on the arguments of the function to differentiate between elements of the system. For example, $q(\omega \mid i)$ refers to the map $I \times \Omega \to \mathbb{R}_+$ as specified above, while $q(\omega)$ indicates the map $\Omega \to \mathbb{R}_+$, and so on. We refer to the system of functions generated by $\mathbf{q}(i)$ and $q(\omega \mid i)$ as a *generalized joint distribution*. When a function is a properly normalized probability distribution, we denote it in boldface. If these functions are all probability distributions, then we have the familiar case of a joint distribution ($\mathbf{q}(i, \omega)$) with accompanying marginal ($\mathbf{q}(i)$ and $\mathbf{q}(\omega)$) and conditional ($\mathbf{q}(\omega \mid i)$ and $\mathbf{q}(i \mid \omega)$) distributions.

3.2 The Growth Model

We consider a population comprised of a finite set I of individuals i. There is a finite set Ω of states ω. The population begins period 1 with a quantity of a perfectly divisible good that for concreteness we refer to as wealth. See Section 3.3 for alternative interpretations.

At the beginning of each period $t = 1, 2, \ldots$, nature independently draws a state according to an interior distribution $\mathbf{p}^0 \in \Delta(\Omega)$. A time invariant *allocation* $\mathbf{q} \in \Delta(I)$ assigns share $\mathbf{q}(i)$ of the current aggregate wealth to each individual $i \in I$ at the beginning of every period period t. Subsequently, each individual i generates a gross *return* per unit of her wealth which depends on the current state ω_t and the individual i. We denote the time invariant return function by $q(\omega_t \mid i)$; this notation will facilitate the statistical reinterpretation of the growth process.

If an individual i were in autarky, her long-run growth rate of wealth would be given by

$$\mathrm{E}_{\mathbf{p}^0} \ln q(\omega \mid i),$$

where the expectation is taken with respect to ω. The long-run growth rate of the aggregate wealth for a population can be expressed as

$$\mathrm{E}_{\mathbf{p}^0} \ln \Big(\sum_i \mathbf{q}(i) q(\omega \mid i) \Big) = \mathrm{E}_{\mathbf{p}^0} \ln \Big(\sum_i q(i, \omega) \Big) = \mathrm{E}_{\mathbf{p}^0} \ln q(\omega). \tag{1}$$

The allocation $\mathbf{q}(i)$, along with the return function $q(\omega \mid i)$, generates a generalized joint distribution q. The "marginal distribution" $q(\omega) = \sum_i q(i, \omega)$ represents the aggregate return in each state ω. By normalizing the aggregate wealth at the beginning of a period with state realization ω to one, the wealth of individual i at the end of the period is given by $q(i, \omega) = \mathbf{q}(i) q(\omega \mid i)$. We refer to the generalized distribution $q(i, \omega)$ as a *policy*.

We let $\mathcal{A} \subseteq \Delta(I)$ be the set of feasible allocations and let $\mathcal{R} \subseteq \mathbb{R}_+^{I \times \Omega}$ be the set of feasible returns. We then imagine a planner who chooses the allocation $\mathbf{q}(i)$ and the return $q(\omega \mid i)$ from a set $\hat{\mathcal{Q}} \subseteq \mathcal{A} \times \mathcal{R}$ of feasible pairs of allocations and returns to maximize the growth rate (1).

Since the growth rate depends on allocation and returns only through the policy $q(i, \omega)$, we also define the set of feasible policies

$$\mathcal{Q} = \Big\{ q(i, \omega) : q(i, \omega) = \mathbf{q}(i) q(\omega \mid i) \text{ for some } \big(\mathbf{q}(i), q(\omega \mid i)\big) \in \hat{\mathcal{Q}} \Big\}.$$

We assume that $\mathcal{Q}$ is compact and contains at least one policy resulting in a finite growth rate from equation (1); that is, at least one policy for which $q(\omega) > 0$ for all ω. The existence of an optimizer is then ensured.[2] We denote the *growth-maximizing* policy by $q^*(i, \omega) \in \mathcal{Q}$.

Our interest in the growth-maximizing policy is motivated by the next result, which states that this policy induces at least as high a growth rate as any other policy, even those that relax the stationarity restriction implicit in $q^*(i, \omega)$, along almost all sequences of states. Let a feasible *history-dependent* policy associate a policy $q(i, \omega) \in \mathcal{Q}$ with every history $(\omega_1, \ldots, \omega_{t-1})$, for every period t.

Proposition 1 (Cover and Thomas (2006))**.** *Consider a feasible history-dependent*

[2] Suppose that for every policy $q(i, \omega)$, there exists some ω such that $q(\omega) = 0$. In this case, every policy results in a growth rate of $-\infty$. However, if there is a policy yielding a finite growth rate g, then the subset of $\mathcal{Q}$ that produces growth rates of at least g is compact. The objective function in equation (1) is continuous on this set, which ensures the existence of a maximizer.

policy and let random variable S_t denote the aggregate wealth accumulated under this policy in the first t periods. Let S_t^* denote the aggregate wealth accumulated by the growth-maximizing policy $q^*(i,\omega)$. Then,

$$\limsup_{t\to\infty}\left(\frac{\ln S_t}{t}-\frac{\ln S_t^*}{t}\right)\leq 0,$$

with probability 1.

The proof, in Appendix A.1, adapts the proof of Theorem 16.3.1 from Cover and Thomas (2006) to our setting.

3.3 Applications

3.3.1 Economic Growth

Our leading application interprets the model in terms of economic growth. Since a society that grows faster than other societies will eventually dominate, we envision a growth-maximizing planner who implements unmodelled tax, redistribution, industrial and other policies to achieve a stationary wealth distribution $\mathbf{q}(i)$ and gross returns $q(\omega \mid i)$. Naturally, the planner is constrained in both aspects. Equity constraints may restrict feasible allocations $\mathbf{q}(i)$ and technological constraints on physical and human capital formation will restrict the return function $q(\omega \mid i)$. The feasibility of allocations and returns is typically interconnected. The choice of the stationary wealth distribution may affect incentives and thus the feasibility of return functions, yielding a constraint set $\hat{\mathcal{Q}}$ that is not a product set. Unlike much of the economic research that derives the set $\hat{\mathcal{Q}}$ from microfoundations, we consider $\hat{\mathcal{Q}}$ as a primitive and develop results applicable to all such sets.[3]

3.3.2 Financial Portfolios

A special interpretation of our model arises when the set $\mathcal{R}$ is a singleton, so that the returns $q(\omega \mid i)$ are fixed. We can then interpret I as a set of financial assets, with asset i securing return $q(\omega \mid i)$ in state ω. The planner is an investor who allocates shares $\mathbf{q}(i)$ of her wealth to assets i at the beginning of each period, with the goal of maximizing the growth rate of the value of

[3]A straightforward reinterpretation allows us to incorporate consumption in our model of economic growth. Assume that individual i consumes a share $c(i,\omega)\in[0,1]$ of her allocation in state ω and generates gross return $r(i,\omega)$ on her residual wealth. We can then apply our analysis with return function $q(\omega \mid i)=(1-c(i,\omega))r(i,\omega)$.

her portfolio. Equivalently, the investor is an expected-utility maximizer with logarithmic utility over the end-of-period wealth. Kelly (1956) introduced this problem, assuming no restrictions on the feasible portfolios, i.e., $\mathcal{A} = \Delta(I)$. When $\mathcal{A} \subsetneq \Delta(I)$, we incorporate constraints on portfolios that may arise for, say, regulatory reasons.

The following example illustrates our framework in a well-known setting that admits a simple solution known as proportional betting.

Example 1 (Proportional Betting)**.** In each period, an investor allocates her wealth across assets indexed by $i \in I$. Each asset is a bet on an exclusive event. The bet on the realized event has gross return r, while all other bets pay zero. To formalize this, let $\Omega = I$, let the state ω represent the realized event, let $\omega \sim \mathbf{p}^0 \in \Delta(I)$, and define returns as $q(\omega \mid i) = r\mathbb{1}_{\omega=i}$.

It is easy to demonstrate that in this special case the growth-maximizing allocation $\mathbf{q}^*(i) = \mathbf{p}^0(i)$ distributes wealth proportionally to each bet's probability of success. This hedging strategy avoids concentrating all aggregate wealth on an asset with the highest expected payoff, as such a policy would eventually result in the loss of all wealth and a growth rate of $-\infty$. △

3.3.3 Statistical Learning

Our model also applies to a statistician who aspires to infer a data-generating distribution. In a first step, let us rephrase the well-known asymptotic learning characterization of White (1982) and Berk (1966) as growth optimization.

The likelihood of a sample $(\omega_\tau)_{\tau=1}^t$ drawn from distribution $\mathbf{p}^0(\omega)$, when evaluated under the hypothesis that the data are drawn from distribution $\mathbf{q}(\omega)$, is approximately

$$\prod_\omega \mathbf{q}(\omega)^{\mathbf{p^0}(\omega)t} \quad = \quad \exp\Big[t\sum_\omega \mathbf{p^0}(\omega)\ln \mathbf{q}(\omega)\Big],$$

with the approximation almost surely becoming arbitrarily sharp as t grows. The growth rate $\sum_\omega \mathbf{p^0}(\omega)\ln \mathbf{q}(\omega)$ on the right side equals, up to a sign and constant term, the Kullback-Leibler divergence between the true and hypothesised distribution. The estimator thus converges to the hypothesis minimizing this divergence, as shown for Bayesian and Maximum-likelihood estimators by Berk (1966) and White (1982), respectively.

Building on the connection between growth and statistical learning, we now draw a formal equivalence between our model and predictive coding, an influential framework from the fields of machine learning and cognitive science.[4]

In the predictive coding problem, an artificial or biological system samples a *signal* ω from a distribution $\mathbf{p}^0(\omega) \in \Delta(\Omega)$. The system aims to form a belief about a latent random variable i correlated with the signal, referred to as the *cause*, with support in I. The system considers a compact set $\mathcal{Q} \subseteq \Delta(I \times \Omega)$ of joint distributions $\mathbf{q}(i, \omega)$ where each $\mathbf{q}(i, \omega)$ corresponds to a hypothesis $\mathbf{q}(\omega) = \sum_i \mathbf{q}(i, \omega)$ over the signal distribution. The system selects one such joint distribution, $\mathbf{q}^*(i, \omega)$, that best fits the signal distribution. Specifically, the system is assumed to select a joint distribution that maximizes the growth rate of the likelihood of the signal sample as the sample expands:

$$\mathbf{q}^*(i, \omega) \in \underset{\mathbf{q}(i,\omega) \in \mathcal{Q}}{\arg\max} \sum_{\omega} \mathbf{p}^0(\omega) \ln \sum_i \mathbf{q}(i, \omega); \tag{2}$$

note that this objective coincides with (1).[5] Upon observing a realization of the signal ω, the system seeks to forms a conditional belief $\mathbf{q}^*(i \mid \omega)$ from $\mathbf{q}^*(i, \omega)$.

The predictive coding problem is often motivated as either a model of how the human brain processes information or as a template for the design of an artificial information-processing system. The set $\mathcal{Q}$ represents the system's initial knowledge and its constraints in information processing. Section 5 uses our growth-theoretic techniques to characterize the *best fit* $\mathbf{q}^*(i, \omega)$ and the Bayesian beliefs $\mathbf{q}^*(i \mid \omega)$.

3.3.4 Biological Growth

Stochastic growth also occurs in an evolutionary context. Let each $i \in I$ represent a phenotype and each $\omega \in \Omega$ an environmental shock. Each phenotype is a physical or behavioral characteristic, and the return $q(\omega \mid i)$ identifies the net number of descendants of phenotype i in environment ω. As in the literature on so-called bet-hedging strategies in biology (Kussell and Leibler, 2005), a genotype can be modeled as a mixed strategy $\mathbf{q}(i)$ over phenotypes i, capturing the fact that a given genotype can manifest itself as different characteristics in different individuals. Evolution selects a genotype $\mathbf{q}(i)$ with the highest population growth rate for a given distribution $\mathbf{p}_0(\omega)$ of the environmental shocks.

[4] Another term used is predictive processing and, in the context of machine learning, variational autoencoder. Early contributions to predictive coding can be found in Dayan et al. (1995) and Friston (2005), while a comprehensive literature review is available in Clark (2013). See Aridor et al. (2020, 2023) for recent economic applications.

[5] Note that problem (2) is equivalent to the minimization of the Kullback-Leibler divergence between the true and hypothesised signal distribution.

3.4 The Naive Meritocracy Principle

We define the *merit distribution* induced by policy $q(i, \omega)$ to be

$$\mathbf{m}_q(i) = \mathrm{E}_{\mathbf{p}^0}\, \mathbf{q}(i \mid \omega), \tag{3}$$

where the expectation is with respect to ω. To interpret the merit distribution, note that

$$\mathbf{q}(i \mid \omega) = \frac{q(i, \omega)}{q(\omega)} = \frac{\mathbf{q}(i) q(\omega \mid i)}{\sum_j \mathbf{q}(j) q(\omega \mid j)}. \tag{4}$$

Taking the expectation over ω gives individual i's merit $\mathbf{m}_q(i)$, which is thus the probability that a dollar randomly drawn from the total wealth at the end of a period has been generated by individual i.

As implied by its name, the merit distribution aims to quantify the relative significance of each individual's contributions to the long-term growth rate. The following corollary of Lemma 2 below illustrates how an individual's merit is indicative of their importance for growth. Suppose that returns of individual j increase by a multiplicative factor $\exp(\varepsilon_j)$ in each state ω while returns of other individuals are unmodified. Let $r^*(\varepsilon_j)$ denote the optimized growth rate of aggregate wealth.

Corollary 1. *The derivative of the optimized growth rate is proportional to individual j's merit:*

$$\frac{dr^*}{d\varepsilon_j}(0) = \mathbf{m}_{q^*}(j).$$

See Appendix A.2 for the proof.

Our concept of merit combines the allocation the individuals enjoy with the individuals' productivities—an individual can attain high merit by being highly productive, but also simply by enjoying a relatively high allocation. In an extreme example, an allocation $\mathbf{q}(1) = 1$ ensures individual 1 (and only 1) has maximal merit, even though other individuals may be more productive.

Let the allocation $\mathbf{q}^*(i)$ and returns $q^*(\omega \mid i)$ jointly maximize the aggregate growth rate (1). Let $\mathcal{A}^* = \{\mathbf{q}(i) : (\mathbf{q}(i), q^*(\omega \mid i)) \in \hat{\mathcal{Q}}\}$ be the set of allocations to which the planner can deviate without altering the optimized returns.

The following result establishes a necessary condition for the growth-maximizing policy. If a policy maximizes the aggregate growth rate, it must also minimize the disparity measured by the Kullback-Leibler (KL) divergence between the merit and allocation distributions.[6] Section 4.4 proves:

[6]Recall that KL-divergence is often interpreted as a measure of dissimilarity between two

Proposition 2 (Naive Meritocracy Principle)**.** *If the allocation* $\mathbf{q}^*(i)$ *and return function* $q^*(\omega \mid i)$ *jointly maximize the aggregate growth rate* (1)*, then* $\mathbf{q}^*(i)$ *minimizes the KL-divergence from the induced merit distribution:*

$$\mathbf{q}^*(i) \in \underset{\mathbf{q}(i)\in\mathcal{A}^*}{\arg\min}\, \mathrm{KL}\left(\mathbf{m}_{q^*}(i) \parallel \mathbf{q}(i)\right). \tag{5}$$

The growth-maximizing policy thus coincides with an outcome of a naive meritocracy. It is a meritocracy because the allocation $\mathbf{q}^*(i)$ comes as close as feasibility allows to the merit distribution $\mathbf{m}_{q^*}(i)$, meaning it adheres as closely as possible to the credo "to each according to his merit." It is a naive meritocracy because the planner fixes the growth-maximizing policy, calculates the induced merit distribution, and then asks whether she can get closer to this merit distribution by altering the allocation *while holding fixed the returns and merit distribution.* The planner does not in general minimize the KL-divergence between the allocation and the merit distribution across all feasible policies. One could reduce this distance by recognizing that altering the policy will alter the induced merit distribution, but such further "fairness improvement" would reduce growth.

Andre (2022) introduces the notion of a *shallow meritocracy*—one in which people judge the merit of others' behavior, without making allowances for the fact that different circumstances may induce different behavior. Andre reports experimental evidence consistent with such "shallow" judgments. Similarly, our merit distribution ranks individuals accordingly to their expected contribution to aggregate wealth, without distinguishing whether a high contribution reflects an individual who is highly productive or an individual endowed with a copious amount of wealth.

Example 2 (Portfolio Rebalancing)**.** In the context of financial portfolio management, portfolio $\mathbf{q}(i \mid \omega)$ represents the share of an investor's wealth allocated to each asset i at the end of a period, after the returns in state ω have been realized. The distribution $\mathbf{m}_q(i) = \mathrm{E}_{\mathbf{p}_0}\, \mathbf{q}(i \mid \omega)$ then corresponds to the end-of-period portfolio, averaged across state realizations. The naive meritocracy principle asserts that the growth-maximizing portfolio $\mathbf{q}^*(i)$, to which

distributions. Given two probability distributions $\mathbf{p}$ and $\mathbf{q}$, the KL-divergence is defined as

$$\mathrm{KL}\left(\mathbf{p}(x) \parallel \mathbf{q}(x)\right) := \sum_{x\in X} \mathbf{p}(x) \ln \frac{\mathbf{p}(x)}{\mathbf{q}(x)}.$$

We adopt the standard convention that $0 \ln 0 = 0$.

the investor rebalances at the beginning of each period, minimizes the wedge $\mathrm{KL}\left(\mathbf{m}_{q^*}(i) \parallel \mathbf{q}(i)\right)$ among all feasible portfolios. In a special case, when the investor is free to choose any portfolio $\mathbf{q}(i)$ from $\mathcal{A} = \Delta(I)$, this principle simplifies to the observation made in Cover and Thomas (2006), Section 16.2, which states that no predictable rebalancing occurs at the optimum: $\mathbf{q}^*(i) = \mathbf{m}_{q^*}(i)$. △

The following numerical example illustrates the interaction between redistribution, access to investment opportunities, and economic growth.

Example 3 (Inequality and Growth)**.** Suppose there are four states, labeled ω_1, ω_2, ω_3, and ω_4, with prior probabilities 0.25, 0.275, 0.225 and 0.25. There are five individuals, labeled 1, 2, 3, 4 and 5. The set $\mathcal{R}$ of feasible return functions is a singleton, with the returns $q(\omega \mid i)$ fixed throughout, given by

		State			
		ω_1	ω_2	ω_3	ω_4
	1	4	1	1	1
	2	0	3	1	1
Individual	3	0	1	3	1
	4	0	1	1	3
	5	1.2	1.2	1.2	1.2

Individual 1 has an autarkic growth rate of $.25 \ln 4 = 0.35$ and individual 5 has an autarkic growth rate of $\ln 1.2 = 0.18$. For each of the other three individuals, there is a state with a zero return, ensuring their autarkic growth rates of $-\infty$. In the absence of redistribution, individual 1's share of the economy's wealth will converge to one.

Now let the planner choose an allocation $\mathbf{q}(i)$ without any restrictions, $\mathcal{A} = \Delta(\{1, 2, 3, 4, 5\})$. The growth-maximizing allocation is (approximately)

$$\begin{aligned}
\mathbf{q}^*(1) &= 0.63 \\
\mathbf{q}^*(2) &= 0.19 \\
\mathbf{q}^*(3) &= 0.063 \\
\mathbf{q}^*(4) &= 0.13 \\
\mathbf{q}^*(5) &= 0.
\end{aligned}$$

Computation of the merit distribution according to its definition (3) reveals that $\mathbf{m}_{q^*}(i) = \mathbf{q}^*(i)$ for all individuals i in this case. This stems from the naive meritocracy principle, which, given the planner's unconstrained choice of allocation, necessitates an allocation that perfectly matches the merit distribution.

The growth rate of aggregate wealth is 0.4, which exceeds the autarkic growth rates of all individuals. The growth-maximizing allocation is akin to an investment portfolio that is optimally hedged against risk. Individual 1 enjoys the highest share, reflecting the fact that she does very well in one state, and faces bankruptcy in none. Individuals 2, 3 and 4 receive positive shares of wealth, even though they are surely headed for ruin under autarky, because each fares relatively well in one of the states and the allocation successfully hedges against the individuals' risks. Their shares are ranked according to the relative likelihoods of the states in which they exhibit their high return. The exclusion of an individual, in this case of the individual 5, is not a coincidence. Relying on an analogy with the rational inattention problem, we show in Section 6.2 that when the allocation is unrestricted and the returns are fixed, then the set of individuals who receive positive wealth shares is no larger than the state space.

We now introduce constraints on the allocation. Let a "poverty-level" constraint entitle each individual to at least 0.1 of aggregate wealth, and also an "equity" constraint impose an upper bound on the Gini coefficient.[7] The optimal allocation and the corresponding merit distribution are then

$$\begin{aligned}
\mathbf{q}^*(1) &= .53 & \qquad \mathbf{m}_{q^*}(1) &= .55\\
\mathbf{q}^*(2) &= .14 & \qquad \mathbf{m}_{q^*}(2) &= .15\\
\mathbf{q}^*(3) &= .102 & \qquad \mathbf{m}_{q^*}(3) &= .10\\
\mathbf{q}^*(4) &= .12 & \qquad \mathbf{m}_{q^*}(4) &= .11\\
\mathbf{q}^*(5) &= .10 & \qquad \mathbf{m}_{q^*}(5) &= .08,
\end{aligned}$$

with growth rate 0.38. Individual 1 receives a yet smaller share, in an effort to reduce the Gini coefficient (though individual 1's wealth still grows faster than in autarky). The allocation to individual 5 binds against the poverty constraint. The allocation of individual 3 rises slightly above the poverty constraint, with

[7] The Gini coefficient for the wealth shares of a population of n agents is defined as the relative mean difference in shares,

$$\frac{1}{2n^2\overline{q}}\sum_{i=1}^{n}\sum_{j=1}^{n}|\mathbf{q}(i)-\mathbf{q}(j)|,$$

where $\overline{q}$ is the mean wealth share in the population (equal to $1/5$ in our population of size $n = 5$).

the slack helping to reduce the Gini coefficient.

The merit distribution no longer equals the allocation, but the naive meritocracy principle (5) continues to hold. The allocation $\mathbf{q}^*(i)$ minimizes the wedge $\mathrm{KL}(\mathbf{m}_{q^*} \parallel \mathbf{q})$ among all feasible allocations $\mathbf{q}$.

Redistribution in this setting serves to direct resources to valuable investment opportunities. One might expect redistribution to favor individuals with high autarkic growth rates, but we have seen that this link is far from precise. Growth-maximizing redistribution directs resources to individuals who have relatively high returns in states that are (i) relatively likely and (ii) in which aggregate output is relatively low. The merit distribution induced by the growth-optimal policy summarizes both these aspects. △

4 Paths of Money

4.1 The Growth-Maximizing Policy

We prove the naive meritocracy principle, as well as derive further results, based on the auxiliary characterization of the growth-maximizing policy in the next proposition. The proposition makes use of an extension of the KL-divergence that applies the standard formula (see footnote 6) to pairs $(\mathbf{p}, q)$, where $\mathbf{p}$ is a distribution and q a generalized distribution; that is, $\mathrm{KL} : \Delta(X) \times \mathbb{R}_+^X \to \mathbb{R} \cup \{\infty\}$. We continue to interpret $\mathrm{KL}(\mathbf{p} \parallel q)$ as a degree of discrepancy between $\mathbf{p}$ and q (although it may now take on negative values). Indeed, letting $\mathbf{q}(x) = q(x)/\sum_{y\in X} q(y)$ stand for the normalization of $q(x)$, we have

$$\mathbf{q} \in \operatorname*{arg\,min}_{\mathbf{p}\in\Delta(X)} \mathrm{KL}(\mathbf{p} \parallel q).$$

The distribution closest to a generalized distribution is thus its normalization, extending the fact that the KL-divergence of two distributions is minimized when these are equal.

Proposition 3. *A policy $q^*(i, \omega)$ maximizes the growth rate* (1) *if and only if* *it solves*

$$\begin{aligned} \min_{\mathbf{p}(i,\omega),q(i,\omega)} \quad & \mathrm{KL}\left(\mathbf{p}(i,\omega) \parallel q(i,\omega)\right) \\ \text{s.t.} \quad & \mathbf{p}(\omega) = \mathbf{p}^0(\omega) \\ & q(i,\omega) \in \mathcal{Q} \\ & \mathbf{p}(i,\omega) \in \Delta(I \times \Omega), \end{aligned}$$

together with some minimizer $\mathbf{p}^*(i,\omega)$*. Additionally,* $\mathbf{p}^*(i,\omega)$ *equals the merit distribution on the margin:*

$$\mathbf{m}_{q^*}(i) = \mathbf{p}^*(i) = \sum_{\omega} \mathbf{p}^*(i,\omega).$$

The proposition describes the growth-maximizing policy and its resulting merit distribution as the "closest pair" of two (generalized) distributions belonging to separate sets. Sections 4.2–4.4 offer a proof, intuition, and interpretation of the control $\mathbf{p}(i,\omega)$.

4.2 Dominant Frequency

We start with a technical result. The next lemma is closely related to the Donsker-Varadhan Lemma (Dupuis and Ellis, 1997, Proposition 1.4.2).[8] Our formulation of the lemma and the proof are designed to clarify how the consistency optimization, such as that in Proposition 3, relates to growth.

Lemma 1 (Dominant Frequency)**.** *Let* X *be a finite set. The following holds for any* $q : X \to \mathbb{R}_{++}$*:*

$$\ln \sum_{x \in X} q(x) = - \min_{\mathbf{p} \in \Delta(X)} \mathrm{KL}(\mathbf{p} \parallel q). \tag{6}$$

We sketch the argument here and relegate the formal proof to Appendix A.3.

Proof sketch. The first step is to bring growth into the picture by noting that $\ln \sum_x q(x)$ is the growth rate of the multiplicative process that takes value $\left(\sum_x q(x)\right)^t =: y^t$ for $t = 1, 2, \ldots$. Next, one can write an equivalent expression

[8] The Donsker-Varadhan Lemma states that $\mathrm{E}_{\mathbf{f}} \exp u(x) = \max_{\mathbf{p}} \{\mathrm{E}_{\mathbf{p}}\, u(x) - \mathrm{KL}(\mathbf{p} \parallel \mathbf{f})\}$. Substitution $q(x) = \mathbf{f}(x) \exp u(x)$ and application of this equality implies equation (6).

for y^t by considering all possible t-length sequences of elements from X. For each t-sequence in the set X^t, we calculate the product $\prod_{\tau=1}^{t} q(x_\tau)$ of the consecutive terms, and then sum across sequences giving.

$$\ln \sum_{x \in X} q(x) = \frac{1}{t} \ln \sum_{(x_1,\ldots,x_t)} \prod_{\tau=1}^{t} q(x_\tau).$$

In effect, the growth process y^t is the sum of constituent multiplicative processes, each of which corresponds to a particular sequence of values drawn from X.

We group the t-sequences into equivalence classes, where sequences in an equivalence class display the same frequencies $\mathbf{p}(x) \in \Delta(X)$ for all $x \in X$. For each equivalence class $\mathbf{p}$, we can associate a growth rate, that we can attribute to two sources. Firstly, every sequence within an equivalence class has its own growth rate, which depends solely on the frequency $\mathbf{p}$ and hence is constant across members of the equivalence class. Secondly, as t increases, so does the number of sequences in the equivalence class $\mathbf{p}$.

Over time, even minor differences in growth rates compound into overwhelming differences in accumulated values. Consequently, in the long run, the equivalence class with the highest growth rate determines the growth rate of y^t, while the contributions of all other equivalence classes fall into insignificance.

Which equivalence class will emerge as dominant? A distribution $\mathbf{p}$, that is concentrated on values of x for which $q(x)$ is large, generates sequences with a high growth rate. However, distributions $\mathbf{p}$ that are "random," meaning they are dispersed throughout X rather than concentrated, benefit from having numerous sequences whose frequencies match $\mathbf{p}$. The dominating equivalence class $\mathbf{p}$ balances these two desiderata. We demonstrate through straightforward algebra that each equivalence class achieves a growth rate equal to $-\mathrm{KL}(\mathbf{p} \parallel q)$. The growth rate, $\ln y$, is then determined by the equivalence class whose frequencies minimize this divergence, as stated in (6). That is, the dominating frequency $\mathbf{p}$ maximizes a measure of consistency with q.

4.3 Growth Rate Characterization

We now return to economic growth and recall from (1) that a policy $q(i, \omega)$ induces growth rate $\sum_\omega \mathbf{p}^0(\omega) \ln \sum_i q(i, \omega)$. By applying Lemma 1 to each state ω individually, we obtain the following characterization of the growth rate.

Lemma 2. *A policy $q(i,\omega)$ induces growth rate*

$$-\min_{\mathbf{p}(i,\omega)} \quad \mathrm{KL}\left(\mathbf{p}(i,\omega) \parallel q(i,\omega)\right) - \mathrm{H}\left(\mathbf{p}^0(\omega)\right) \tag{7}$$

$$\text{s.t.} \quad \mathbf{p}(\omega) = \mathbf{p}^0(\omega).$$

See Appendix A.4 for the proof. The term $-\mathrm{H}\left(\mathbf{p}^0(\omega)\right)$ is independent of the control variable and thus does not affect the optimizer. It represents the negative impact of prior state stochasticity on growth.

The proof involves decomposing the growth of aggregate wealth into a sum of exponential processes. Each constituent process corresponds to a distinct value of the control variable $\mathbf{p}(i,\omega)$. Among these constituent processes, the fastest-growing one outpaces all others, resulting in the optimization expressed in (7). To gain intuition, imagine each dollar drawn from the initial stock of perfectly divisible aggregate wealth as founding a separate dynasty of subsequent wealth, identifying the wealth in each subsequent period descended from the original dollar. In each period, each such dynasty finds itself in the hands of an individual randomly drawn from distribution $\mathbf{q}(i)$ and in a random state of the economy drawn from $\mathbf{p}^0(\omega)$, and its wealth multiplies by the return $q(\omega \mid i)$. Over the course of time, a distribution $\mathbf{p}(i,\omega)$ describes the frequency of individuals and states occupied by this dynasty. We refer to $\mathbf{p}(i,\omega)$ as the *path* of the dynasty. Every such path satisfies the consistency condition $\mathbf{p}(\omega) = \mathbf{p}^0(\omega)$ for all $\omega \in \Omega$ in the long run. For every such consistent path and a sufficiently long finite time horizon, a positive measure of dynasties exists whose frequencies of visiting (i,ω) approximate that path $\mathbf{p}(i,\omega)$.

For each path $\mathbf{p}(i,\omega)$ and any given time horizon, we define the wealth of the path by summing up the wealth of all dollar dynasties whose realized frequencies of visiting (i,ω) match $\mathbf{p}(i,\omega)$. The growth rate of wealth varies across different paths. We use Lemma 1 to show that for a path $\mathbf{p}(i,\omega)$ under policy $q(i,\omega)$, the growth rate is given by

$$-\mathrm{KL}\left(\mathbf{p}(i,\omega) \parallel q(i,\omega)\right),$$

which indicates that the wealth of a path consistent with policy $q(i,\omega)$ grows relatively quickly. The growth rate of a path can be broken down into two channels. First, paths that deviate from the allocation $\mathbf{q}(i)$ by visiting individuals i with distinct frequencies $\mathbf{p}(i \mid \omega) \neq \mathbf{q}(i)$ in states ω involve atypical draws, causing the measure of dollar dynasties following such paths to decrease over time; the more significant the deviation from the allocation, the faster the decrease. Secondly, the wealth of a dollar dynasty grows rapidly if it frequently visits pairs (i, ω) with high returns $q(\omega \mid i)$.

As the wealth of each path grows exponentially, the wealth of the fastest-growing path ultimately surpasses that of all other paths, determining the overall growth rate of the economy. Since growth defined in (1) is evaluated alongside a single sequence of states, the wealth of the economy equals the wealth of the winning path divided by the number of state sequences with frequencies $\mathbf{p}^0$. This normalization is captured by the term $-\mathrm{H}(\mathbf{p}^0)$ in (7).

4.4 Growth-Maximizing Path and Policy

The growth-maximizing path solving (7) has a close connection to the merit distribution defined in (3). We expand the definition of the latter as follows: Given a policy $q(i, \omega)$, we define a *merit joint distribution* as

$$\mathbf{m}_q(i, \omega) = \mathbf{p}^0(\omega)\mathbf{q}(i \mid \omega), \tag{8}$$

which represents the probability that a dollar, randomly sampled from the aggregate wealth at the end of a period, is held by individual i in state ω. The merit distribution from (3) is then the marginal distribution $\mathbf{m}_q(i) = \sum_\omega \mathbf{m}_q(i, \omega)$.

Given a policy $q(i, \omega)$, let the *best response* $\mathbf{p}_q^*(i, \omega)$ be the path that solves the optimization problem (7).

Lemma 3. *For any policy $q(i, \omega)$, the best response $\mathbf{p}_q^*(i, \omega)$ equals the merit distribution $\mathbf{m}_q(i, \omega)$.*

Specifically, reminiscent of Kelly's proportional betting, the growth-maximizing path visits individuals in proportion to their merits.

The proof, provided in Appendix A.5, is based on the observation that the conditional distribution most consistent with $\mathbf{q}(i \mid \omega)$ is given by $\mathbf{p}_q^*(i \mid \omega) = \mathbf{q}(i \mid \omega)$.

Proof of Proposition 3. Adding optimization over policies, the result immediately follows from Lemma 2 and Lemma 3. □

We are now prepared to prove the naive meritocracy principle from Proposition 2. In short, it follows from the fact that the growth-maximizing policy must maximize consistency with the growth-maximizing path of money.

Proof of Proposition 2. By Proposition 3, the growth-maximizing policy satisfies

$$q^*(i,\omega) \in \underset{q(i,\omega)\in\mathcal{Q}}{\arg\min}\, \mathrm{KL}\left(\mathbf{p}_{q^*}(i,\omega) \parallel q(i,\omega)\right)$$

$$= \underset{q(i,\omega)\in\mathcal{Q}}{\arg\min}\, \mathrm{KL}\left(\mathbf{m}_{q^*}(i,\omega) \parallel q(i,\omega)\right).$$

Applying the chain rule to decompose the KL-divergence of (generalized) joint distributions, we can rewrite this last objective as follows:[9]

$$\mathrm{KL}\left(\mathbf{m}_{q^*}(i) \parallel \mathbf{q}(i)\right) + \sum_i \mathbf{m}_{q^*}(i)\, \mathrm{KL}\left(\mathbf{m}_{q^*}(\omega \mid i) \parallel q(\omega \mid i)\right). \quad (9)$$

By setting the returns $q(\omega \mid i)$ to $q^*(\omega \mid i)$ and restricting optimization of the last inline expression to control of allocations $\mathbf{q}(i) \in \mathcal{A}^*$, we obtain the statement in (5). □

It is tempting to think that to maximize growth, a meritocracy concept should allocate wealth according to individual productivities, while disregarding initial allocations. The proof of Proposition 2 provides an intuition for the "naivete" of the meritocracy principle that, instead, rewards individuals proportionally to a combination of their productivities and allocated wealth. This is a consequence of the growth-maximizing policy maximizing the growth rate of the "winning" path of money. This winning money path visits individuals based on a combination of their (i) productivities and (ii) allocated wealth shares. The latter is essential for the growth rate of the path, as the number of dollar dynasties deviating significantly from the allocation diminishes rapidly. The growth rate of the winning path is best enhanced by a fixed-point policy that rewards individuals based on a blend of their productivities and allocated wealth.

The fixed-point aspect of the naive meritocracy principle has notable implications for comparative statics, making merit and the growth-maximizing allocation seemingly oversensitive to variations in individuals' productivities. A boost in productivity of a specific individual directly enhances that individual's merit. Additionally, though, a positive feedback-loop arises via the adaptation of the allocation.

To illustrate, consider a scenario where individual j experiences an increase

[9] In Appendix A.4, Lemma 4 verifies that this chain rule remains valid for the extended definition of KL-divergence.

in their exogenous return in state $\tilde{\omega}$. Initially, let us examine the elasticity of the state-contingent merit $\mathbf{m}_q(j \mid \tilde{\omega})$ given by equation (4), keeping the allocation $\mathbf{q}(i)$ constant. Straightforward calculation gives elasticity

$$\frac{\partial \mathbf{m}_q(j \mid \tilde{\omega})/\mathbf{m}_q(j \mid \tilde{\omega})}{\partial q(\tilde{\omega} \mid j)/q(\tilde{\omega} \mid j)} = 1 - \mathbf{m}_q(j \mid \tilde{\omega}).$$

Now suppose, instead, that the planner adjusts the allocation $\mathbf{q}^*(i)$ to maximize the growth rate. The elasticity then becomes

$$\frac{d\mathbf{m}_q(j \mid \tilde{\omega})/\mathbf{m}_q(j \mid \tilde{\omega})}{dq(\tilde{\omega} \mid j)/q(\tilde{\omega} \mid j)} = 1 - \mathbf{m}_q(j \mid \tilde{\omega}) + \xi_{jj} - \sum_i \mathbf{m}_q(i \mid \tilde{\omega})\xi_{ij},$$

where $\xi_{ij} = \frac{d\mathbf{q}^*(i)/\mathbf{q}^*(i)}{dq(\tilde{\omega}|j)/q(\tilde{\omega}|j)}$ is the elasticity of growth-maximizing allocation for individual i with respect to j's return in state $\tilde{\omega}$. Consequently, the planner favors j and treats the remaining individuals less favorably, which further boosts j's merit and allocation beyond the direct impact. In the realm of financial portfolio management, the indirect effects contribute to sensitivity of the growth-maximizing portfolio to exogenous variations in the distribution of returns.

4.5 A Second Meritocracy Principle

The merit principle presented in Proposition 2 held returns fixed and established a necessary condition satisfied by the growth-maximizing *allocation*. Here we derive a similar consistency-based necessary condition satisfied by the growth-maximizing *returns*. We find that the growth-maximizing policy distributes returns for each individual across states in a manner that aligns as closely as possible with the individual's contributions to economic growth in those states. The result again follows quickly from Proposition 3.

Fix an individual j. Given a policy $q(i, \omega)$, we define j's contribution to growth in each state ω as $\mathbf{m}_q(\omega \mid j) = \mathbf{m}_q(j, \omega)/\mathbf{m}_q(j)$, where $\mathbf{m}_q(i, \omega)$ is the merit joint distribution from (8). In simpler terms, $\mathbf{m}_q(\omega \mid j)$ represents the probability that the current state is ω given that a dollar sampled from the aggregate wealth at the end of a period was produced by individual j (under the policy q). Normalizing an individual's total contribution to growth to one, we interpret $\mathbf{m}_q(\omega \mid j)$ as the distribution of individual j's contribution to economic growth across states.

Let

$$\mathcal{R}_{q^*,j} = \Big\{ q(\omega \mid j) : \Big(\mathbf{q}^*(i), \big(q^*(\omega \mid k)\big)_{k \neq j}, q(\omega \mid j) \Big) \in \hat{\mathcal{Q}} \Big\}$$

represent the set of return functions for individual j to which the planner can deviate from the optimal policy without altering the allocation and returns of others.

Proposition 4. *If policy $q^*(i, \omega)$ maximizes the growth rate and an individual j has positive merit, $\mathbf{m}_{q^*}(j) > 0$, then j's return function $q^*(\omega \mid j)$ minimizes the KL-divergence from the induced conditional merit distribution $\mathbf{m}_{q^*}(\omega \mid j)$:*

$$q^*(\omega \mid j) \in \underset{q(\omega \mid j) \in \mathcal{R}_{q^*,j}}{\arg\min} \ \mathrm{KL}\left(\mathbf{m}_{q^*}(\omega \mid j) \parallel q(\omega \mid j) \right). \tag{10}$$

Similar to the first naive meritocracy principle, this result treats $\mathbf{m}_{q^*}(\omega \mid j)$ as given; the optimization in (10) does not account for the influence of variations in j's returns on j's merit.

Proof of Proposition 4. Recall from the proof of Proposition 2 that the growth-maximizing policy minimizes (9). When fixing the allocation to $\mathbf{q}^*(i)$ and returns to $q^*(\omega \mid k)$ for individuals $k \neq j$, the remaining optimization over $q(\omega \mid j)$ leads to the result. □

Example 4 (Optimal Returns)**.** Consider a single individual $i = 1$ and two states $\omega \in \{1, 2\}$. Interest then centers on how individual 1 distributes her returns $q(\omega \mid 1) \geq 0$ across the two states to maximize the growth rate. Assume a feasibility constraint $q(1 \mid 1) + q(2 \mid 1) = 1$.

The growth maximization problem for this autarkic individual is given by:

$$\max_{q(\omega \mid 1)} \sum_{\omega} \mathbf{p}^0(\omega) \ln q(\omega \mid 1),$$

subject to the feasibility constraint. This is equivalent to:

$$\min_{q(\omega \mid 1)} \mathrm{KL}\left(\mathbf{p}^0(\omega) \parallel q(\omega \mid 1) \right),$$

subject to the same constraint, since the two objectives differ only in sign and a term independent of the control. This minimization is equivalent to the second meritocracy principle in (10) because, when there is only one individual, $\mathbf{m}_{q^*}(\omega \mid$

$1) = \mathbf{p}^0(\omega)$.[10] If the states are equally likely, then the individual equalizes her state-contingent returns.

We now let individual 1 be joined by an individual 2, which alters the marginal contributions of individual 1 to growth across states. This alteration is reflected in the change of $\mathbf{m}_{q^*}(\omega \mid 1)$ and $q^*(\omega \mid 1)$ between the two scenarios.

We continue with the case of two equiprobable states. We fix the allocation to be the egalitarian one, with $\mathbf{q}(1) = \frac{1}{2} = \mathbf{q}(2)$, and fix individual 2's returns to be $q(1 \mid 2) = \frac{1}{4}$ and $q(2 \mid 2) = \frac{3}{4}$. Due to the symmetry of the setting, the growth-maximizing return for individual 1 complements that of individual 2: $q^*(\omega \mid 1) = 1 - q(\omega \mid 2)$.

The meritocracy principle (10) continues to be satisfied: To compute $\mathbf{m}_{q^*}(\omega \mid 1)$, we note that the growth-maximizing policy for this setting is a properly normalized joint probability distribution with marginal distribution $\mathbf{q}^*(\omega) = \mathbf{p}^0(\omega)$. The merit distribution $\mathbf{m}_{q^*}(i, \omega)$ minimizes KL-divergence to $\mathbf{q}^*(i, \omega)$, subject to being equal to $\mathbf{p}^0(\omega)$ on the margin. Hence, in this case, $\mathbf{m}_{q^*}(i, \omega) = \mathbf{q}^*(i, \omega)$ and hence $\mathbf{m}_{q^*}(\omega \mid 1) = \mathbf{q}^*(\omega \mid 1)$. Thus, once more, the growth-maximizing return function of individual 1 minimizes the KL-divergence from $\mathbf{m}_{q^*}(\omega \mid 1)$. $\triangle$

4.6 Learning the Growth-Maximizing Policy

We now examine a series of straightforward policy adjustments, which myopically optimize a fairness measure. The resulting sequence of policies converges to the growth-maximizing policy. This finding may be interpreted as offering a numerical optimization algorithm, and also as an indication that the myopic application of basic fairness principles can guide societies towards growth maximization.

Let $\mathbf{m}_0(i, \omega)$ be an arbitrary joint distribution from the interior of $\Delta(I \times \Omega)$

[10] Recall that $\mathbf{m}_{q^*}(\omega \mid i)$ is the state distribution conditioned on the event that a random dollar has been produced by individual i. This event is uninformative in autarky, resulting in $\mathbf{m}_{q^*}(\omega \mid 1) = \mathbf{p}^0(\omega)$.

such that $\mathbf{m}_0(\omega) = \mathbf{p}^0(\omega)$. Then, define recursively:

$$\begin{aligned} q_k(i,\omega) \in \underset{q(i,\omega)}{\arg\min} \quad & \mathrm{KL}\left(\mathbf{m}_k(i,\omega) \parallel q(i,\omega)\right) \\ \text{s.t.} \quad & q(i,\omega) \in \mathcal{Q}, \end{aligned} \tag{11}$$

$$\begin{aligned} \mathbf{m}_k(i,\omega) \in \underset{\mathbf{p}(i,\omega)}{\arg\min} \quad & \mathrm{KL}\left(\mathbf{p}(i,\omega) \parallel q_{k-1}(i,\omega)\right) \\ \text{s.t.} \quad & \mathbf{p}(\omega) = \mathbf{p}^0(\omega). \end{aligned} \tag{12}$$

Iteration (12) determines the merit distribution induced by the policy $q_{k-1}(i,\omega)$. Iteration (11) identifies the "fairest" feasible policy given the merit distribution $\mathbf{m}_k(i,\omega)$. By employing the chain rule for the KL-divergence as in (9), we can express each policy adjustment (11) as a constrained optimization aiming to satisfy a combination of both meritocracy principles.

Proposition 5. *Suppose that $\mathcal{Q}$ is convex and includes a policy q such that $q(i,\omega) > 0$ for all pairs (i,ω). Then, $q_k(i,\omega)$ converges to the growth-maximizing policy $q^*(i,\omega)$, and $\mathbf{m}_k(i,\omega)$ converges to the merit distribution $\mathbf{m}_{q^*}(i,\omega)$ induced by this growth-maximizing policy.*

The proof in Appendix A.6 relies on a convergence result by Csiszar and Tusnady (1984) that uses methods of information geometry.

5 Predictive Coding

We now return to the statistical application as outlined in Section 3.3.3. The predictive coding literature regards the problem of finding the best fit $\mathbf{q}^*(i,\omega)$ and the resulting Bayesian belief $\mathbf{q}^*(i \mid \omega)$ as computationally prohibitive, either for the human brain or for artificial systems. The literature proposes that, instead of such brute force, the system fits the data indirectly by adopting two distinct models and maximizing their mutual consistency.

The *generative model* is a joint distribution $\mathbf{q}(i,\omega) \in \mathcal{Q} \subset \Delta(I \times \Omega)$ and the *recognition model* is a typically distinct joint distribution $\mathbf{p}(i,\omega)$, such that its marginal distribution satisfies $\mathbf{p}(\omega) = \mathbf{p}^0(\omega)$. The generative model represents the system's internal model of the distribution over pairs of causes and signals. On the other hand, the recognition model specifies the system's interpretation of the sampled signals. The marginal distribution $\mathbf{p}(\omega) = \mathbf{p}^0(\omega)$ represents the sampled signal distribution. When a signal ω is observed, the system forms a belief $\mathbf{p}(i \mid \omega)$ about the cause. This belief is allowed to be arbitrary, allowing for any joint distribution $\mathbf{p}(i, \omega)$ subject to $\mathbf{p}(\omega) = \mathbf{p}^0(\omega)$.

Proposition 3 applied to the statistical context provides a formal foundation for the central result of the predictive coding literature; see e.g. Dayan et al. (1995).

Corollary 2. *The joint distribution* $\mathbf{q}^*(i, \omega)$ *is the best fit if and only if it solves the following optimization jointly with some recognition model* $\mathbf{p}^*(i, \omega)$*:*

$$\begin{aligned} \min_{\mathbf{p}(i,\omega),\mathbf{q}(i,\omega)} \quad & \mathrm{KL}\left(\mathbf{p}(i, \omega) \parallel \mathbf{q}(i, \omega)\right) \\ s.t. \quad & \mathbf{p}(\omega) = \mathbf{p}^0(\omega) \\ & \mathbf{q}(i, \omega) \in \mathcal{Q}. \end{aligned} \tag{13}$$

Furthermore, the recognition model $\mathbf{p}^*(i, \omega)$ *that jointly solves this optimization specifies the Bayesian belief derived from the best fit:* $\mathbf{p}^*(i \mid \omega) = \mathbf{q}^*(i \mid \omega)$*.*

For the connection between the two applications, observe that in both contexts of economic growth and statistical inference one aims to maximize the long-run growth rate of the sequence of random variables:

$$\prod_{\tau=1}^{t} \sum_i \mathbf{q}(i) q(\omega_\tau \mid i).$$

In the economic context, $\mathbf{q}(i)$ represents the allocation of wealth while in the inference context, it is the generative prior distribution of causes i. Similarly, $q(\omega \mid i)$ represents either a return in economics or a conditional signal likelihood in inference. Thus, the inline expression represents the evolution of aggregate wealth under the policy $q(i, \omega) = \mathbf{q}(i)q(\omega \mid i)$ in economics, while it is the sample likelihood under the generative model $\mathbf{q}(i, \omega)$ in inference. Finally, $\mathbf{p}(i, \omega)$ serves as either a path of money in the context of economic growth or a recognition model in the other context.

In predictive coding, the counterpart to the naive meritocracy principle asserts that the system maximizes a degree of Bayesian consistency as follows.

Corollary 3. *The optimal generative prior distribution* $\mathbf{q}^*(i)$ *maximizes con-*

sistency with the recognition distribution $\mathbf{p}^*(i)$*:*

$$\mathbf{q}^*(i) \in \underset{\mathbf{q}(i)\in\mathcal{A}^*}{\arg\min} \operatorname{KL}\left(\mathbf{p}^*(i) \parallel \mathbf{q}(i)\right),$$

where $\mathcal{A}^* = \left\{\mathbf{q}(i) : \left(\mathbf{q}(i), \mathbf{q}^*(\omega \mid i)\right) \in \hat{\mathcal{Q}}\right\}$ *is the set of generative priors* $\mathbf{q}(i)$ *that are feasible given the optimized likelihood function* $\mathbf{q}^*(\omega \mid i)$*.*

Interpreting $\mathbf{p}^*(i \mid \omega)$ as the system's posterior belief about the cause i upon observing signal ω, $\mathbf{p}^*(i)$ stands for the average of these posteriors when the signals are sampled from $\mathbf{p}^0(\omega)$. The corollary states that the system aligns its generative prior belief $\mathbf{q}(i)$ over the causes as close as possible to $\mathbf{p}^*(i)$.

An analogy can also be drawn between the optimization procedures for economic growth and predictive coding. As we demonstrated for economic growth in Section 4.6, an iterative process that sequentially adapts the policy to match the induced merit converges to the growth-maximizing policy. Similarly, in machine learning, iterative adaptations of the generative and recognition models that myopically maximize their mutual consistency are examined, which converge to the best fit.

The predictive coding literature establishes the characterization in (13) using a variational argument.[11] Our approach, on the other hand, emphasizes the growth aspect. To review the growth-based argument in the context of predictive coding, consider a system that has adopted a generative model $\mathbf{q}(i, \omega)$ and observes a sample of signals $(\omega_\tau)_{\tau=1}^t$ drawn from $\mathbf{p}^0(\omega)$. The likelihood of this sample is the sum of the likelihoods of all samples of pairs $(i_\tau, \omega_\tau)_{\tau=1}^t$ that coincide with the observed sample on the margin. The sum of likelihoods of all such samples of pairs with an empirical distribution $\mathbf{p}(i, \omega)$ grows at the rate

$$-\operatorname{KL}\left(\mathbf{p}(i,\omega) \parallel \mathbf{q}(i,\omega)\right) - \operatorname{H}\left(\mathbf{p}^0(\omega)\right).$$

Maximization of the growth rate over the empirical distribution and the generative model then leads to (13).

[11] Using the chain rule, this variational argument points out that the objective in (13) weakly exceeds the objective in (2) and that the two objectives coincide when $\mathbf{p}(i \mid \omega) = \mathbf{q}(i \mid \omega)$.

6 Further Results

6.1 Sufficient as Well as Necessary Conditions

Proposition 3 implies that a growth-maximizing policy $q^*(i,\omega)$ and its corresponding merit distribution $\mathbf{m}_{q^*}(i,\omega)$ constitute a fixed point:

$$q^*(i,\omega) \in \arg\min_{q(i,\omega)} \quad \mathrm{KL}\left(\mathbf{m}_{q^*}(i,\omega) \parallel q(i,\omega)\right) \qquad \text{s.t.} \quad q(i,\omega) \in \mathcal{Q}, \tag{14}$$

$$\mathbf{m}_{q^*}(i,\omega) \in \arg\min_{\mathbf{p}(i,\omega)} \quad \mathrm{KL}\left(\mathbf{p}(i,\omega) \parallel q^*(i,\omega)\right) \qquad \text{s.t.} \quad \mathbf{p}(\omega) = \mathbf{p}^0(\omega). \tag{15}$$

This fixed-point condition is necessary but not sufficient for growth maximization. The following example demonstrates that multiple fixed points may exist, some of which do not correspond to the growth-maximizing policy.

Example 5 (Multiplicity)**.** Consider an economic-growth process in which each individual i has exogenous return function $q(\omega \mid i)$ and the planner chooses an unconstrained allocation $\mathbf{q}(i) \in \mathcal{A} = \Delta(I)$. Then, for each subset $J \subsetneq I$, there exists a fixed point satisfying (14)–(15) that excludes all individuals $j \in J$ by allocating them zero wealth shares. Such a fixed point can be constructed by solving for the growth-maximizing allocation $\mathbf{q}^*_{I\setminus J}(i)$ of the restricted problem that considers individuals $I \setminus J$ only and then extending it to I by allocating zero shares to all individuals $j \in J$. The resulting allocation indeed generates a fixed point in the original setting with individuals in I: the induced merit distribution satisfies $\mathbf{m}_{q^*}(j,\omega) = \mathbf{p}^0(\omega)\mathbf{q}^*(j \mid \omega) = 0$ for all $j \in J$ and all $\omega \in \Omega$. Thus, the individuals from J have zero merit. The policy that solves (14) against this merit distribution $\mathbf{m}_{q^*}$ then allocates shares $\mathbf{q}^*_{I\setminus J}(i)$ to $i \in I \setminus J$ and zero shares to all individuals from J, as needed for the fixed point. $\triangle$

The following results, proven in Appendix A.7 for our general setting, can be used to determine whether a candidate fixed point represents a growth-maximizing policy.

Proposition 6 (Sufficient and Necessary Condition)**.** *Suppose that the set $\mathcal{Q}$ of feasible policies is convex. A policy $q^*(i,\omega)$ maximizes the growth rate if and* *only if*

$$\mathrm{E}_{\mathbf{p}^0(\omega)} \frac{q(\omega)}{q^*(\omega)} \leq 1 \textit{ for all } q(i,\omega) \in \mathcal{Q}, \tag{16}$$

where $q^*(\omega) = \sum_i q^*(i,\omega)$ *and* $q(\omega) = \sum_i q(i,\omega)$ *are the aggregate returns in state* ω *under the policies* q^* *and* q*, respectively.*

Corollary 4. *If an interior policy* $q^*(i,\omega) > 0$ *from convex* $\mathcal{Q}$ *together with the induced merit distribution* $\mathbf{m}_{q^*}(i,\omega)$ *satisfy the fixed point* (14)–(15)*, then the policy maximizes the growth rate.*

The intuition for the proposition is that the left side of (16) is the linearization of the objective in the growth-rate maximization problem (1) around a candidate solution to (1). Since the growth-rate maximization objective is concave, the linearization around optimum does not affect the set of maximizers. The corollary shows that a interior solution to the fixed-point problem solves the requisite first-order conditions.

To illustrate the proposition and the corollary, we revisit Example 5.

Example (Example 5 continued)**.** Since the allocation is unconstrained, with $\mathcal{A} = \Delta(I)$, the naive meritocracy principle implies that the growth-maximizing allocation equals the marginal merit distribution: $\mathbf{q}^*(i) = \mathbf{m}_{q^*}(i)$. If such fixed point is interior, then it corresponds to the growth-maximizing policy.

Let us now allow for non-interior allocations. Proposition 6 indicates that the allocation $\mathbf{q}^*(i)$ optimizes the growth rate if and only if

$$\sum_{\omega} \frac{\mathbf{p}^0(\omega)}{q^*(\omega)} q(\omega \mid i) \quad = \quad 1 \textit{ if } \mathbf{q}^*(i) > 0, \tag{17}$$

$$\sum_{\omega} \frac{\mathbf{p}^0(\omega)}{q^*(\omega)} q(\omega \mid i) \quad \leq \quad 1 \textit{ if } \mathbf{q}^*(i) = 0. \tag{18}$$

The two inline conditions assert that marginal contributions to growth must be equal across individuals who possess positive shares of wealth, and that excluded individuals must have lower marginal contributions to growth. △

6.2 Rational Inattention

The naive meritocracy principle simplifies in situations where the planner faces no constraints in choosing an allocation. In such cases, the principle mandates that the optimal allocation equals the merit distribution. In this subsection, we leverage this simple implication of the principle in the absence of allocation constraints and demonstrate how our characterization of growth-maximizing policies encompasses some existing results from the rational-inattention literature.

We study here the setting from Example 5 with unconstrained allocation and exogenously fixed returns. In Robson et al. (2023) we have investigated this particular case and established its equivalence to the rational inattention problem originally presented by Matějka and McKay (2015). We review here this equivalence using the framework of the current paper.

Matějka and McKay (2015) present a problem of information acquisition for a single person as follows:

$$\begin{aligned} \max_{\mathbf{r}(i,\omega)} \quad & \left\{ \mathrm{E}_{\mathbf{r}(i,\omega)}\, u(i,\omega) - \mathrm{I}_{\mathbf{r}(i,\omega)} \right\} \\ \text{s.t.} \quad & \mathbf{r}(\omega) = \mathbf{p}^0(\omega). \end{aligned} \tag{19}$$

Their decision-maker has a utility function $u(i,\omega)$ and aims to select a joint distribution $\mathbf{r}$ over pairs (i,ω) that maximizes the expected utility minus the cost of acquiring information, which is measured by the mutual information $\mathrm{I}_{\mathbf{r}(i,\omega)}$. In this literature, the variable $i \in I$ represents an action and $\omega \in \Omega$ is an uncertain payoff state.

Robson et al. select an allocation $\mathbf{q}(i) \in \Delta(I)$ that maximizes the growth rate in an economy with exogenous returns set to $q(\omega \mid i) = \exp u(i,\omega)$; that is, $u(i,\omega)$ denotes the log-return. Proposition 3 of the current paper implies that the growth-maximizing policy $q^*(i,\omega)$ and the induced merit distribution $\mathbf{m}_{q^*}(i,\omega)$ in this growth problem jointly solve (where the merit distribution is the maximizer over the paths $\mathbf{p}(i,\omega)$):

$$\begin{aligned} \min_{\mathbf{p}(i,\omega),q(i,\omega)} \quad & \mathrm{KL}\left(\mathbf{p}(i,\omega) \parallel q(i,\omega)\right) \\ \text{s.t.} \quad & \mathbf{p}(\omega) = \mathbf{p}^0(\omega) \\ & \mathbf{q}(i) = \Delta(I) \\ & q(\omega \mid i) = \exp u(i,\omega). \end{aligned} \tag{20}$$

Proposition 7 (Robson et al (2023))**.** *The joint distribution* $\mathbf{r}^*(i,\omega)$ *solves the rational inattention problem* (19) *if and only if it equals the merit distribution*

$\mathbf{m}_{q^*}(i, \omega)$ *from problem* (20).

To understand the statement in Proposition 7, consider an economy under the growth-maximizing allocation. Sample a dollar from the total wealth at the end of a random period, and take note of the dollar's owner i and the current economic state ω. A selection effect causes a correlation between the owner i and the state ω, as individuals with high returns in the state ω are oversampled. The proposition asserts that the resulting joint distribution from this wealth-weighted sampling, $\mathbf{m}_{q^*}(i, \omega)$, perfectly matches the joint distribution that a decision-maker with a utility function $u(i, \omega)$ and the entropy-based cost of acquiring information would opt for.

The redistribution of wealth thus acts as a replacement of active learning about individuals' returns. To understand why, consider again the "dynasties of dollars" introduced in Section 4.3. Our planner allocates the wealth shares $\mathbf{q}(i)$ without the benefit of any information about the realized returns of the individuals. However, for any given finite horizon, some fraction of the dynasties benefits from a favorable empirical correlation between its owner i and the payoff state ω, resulting in an extraordinary growth rate for such lucky dynasties. The downside of this advantageous correlation is that the fraction of the dynasties that enjoy such luck decreases exponentially with the considered time horizon. The growth-maximizing path involves a compromise level of correlation that balances the growth advantage of the correlation against the cost stemming from the dwindling number of dynasties that enjoy such a correlation. This cost-benefit trade-off is equivalent to the trade-off between the benefit and cost of information in the rational inattention problem with entropy-based cost.

Our sufficient and necessary optimality conditions (17)–(18), derived for the setting with unrestricted allocation, coincide with the sufficient and necessary optimality conditions for the rational inattention problem from Caplin and Dean (2015). In addition, the equivalence to the rational inattention problem implies a simple bound on the number of individuals who enjoy positive shares of aggregate wealth under the unconstrained growth-maximizing policy: For the setting with unrestricted allocation, there exists a growth-maximizing policy such that the number of individuals i who receive positive shares $\mathbf{q}^*(i) > 0$ is at most equal to the size of the state space $|\Omega|$. This observation follows for the analogous bound on the number of actions chosen with positive probabilities in the rational inattention problem derived by Caplin and Dean (2013).[12]

[12] Caplin and Dean derive the bound by a concavification argument: The value of the rational

7 Discussion

This paper arbitrages across fields. We point out that a system adapted to maximize growth must optimize consistency with outcomes it generates. The result provides guidance on redistribution in the context of economic growth and a foundation for predictive coding. In the economic setting, the consistency relation manifests as a meritocracy principle, while it delivers constrained Bayes consistency in the statistical context.

Our proof techniques emphasize the relevance of rare events, as described in large deviations theory. Essentially, the growth-maximizing redistribution policy assists the most productive path of money circulation in the economy, rather than focusing directly on productive individuals. This emphasis on the path of money circulation, rather than individuals, clarifies how the naive meritocracy principle arises. Additionally, we find the growth-based approach helpful for understanding and interpreting the predictive coding model. We hope that demonstrating its equivalence with the growth process will offer pedagogical guidance for economists interested in the predictive coding framework.

Our model of redistribution and growth is highly stylized. One natural extension of our model for future research would incorporate serial correlations of shocks and Markovian redistribution policies. The current model's restriction of the redistribution policy to history-independent allocation is without loss under serially independent shocks. However, if shocks are serially correlated, Markovian redistribution policies, such as proportional taxes, become advantageous because currently productive individuals are likely to remain productive. We conjecture that the current characterization of the optimal policy extends to a Markovian setting, with the Markovian redistribution policy aiming to maximize consistency with a Markovian path of money circulation.

A Appendix

A.1 Proof of Proposition 1

Proof. Starting with unit wealth, let $S = q(\omega)$ be the random variable representing the aggregate wealth accumulated in one period by a generic policy

inattention problem is given by the hyperplane tangent to the decision-maker's value function adjusted for information cost. By Carathéodory's theorem, this hyperplane can be defined by at most $|\Omega|$ tangency points that correspond each to an action chosen with a positive probability.

$q(i,\omega) \in \mathcal{Q}$, and let $S^* = q^*(\omega)$ be the random variable representing the aggregate wealth accumulated in one period by the growth-maximizing policy. Then,

$$\mathrm{E}_{\mathbf{p}^0} \frac{S}{S^*} = \sum_{\omega} \mathbf{p}^0(\omega) \frac{q(\omega)}{q^*(\omega)} \leq 1, \tag{21}$$

where the expectation is with respect to ω and the inequality follows from the necessary and sufficient optimality condition on $q^*(i,\omega)$ from Proposition 6. The inequality (21) extends Theorem 16.2.2 from Cover and Thomas (2006) to our setting. (Cover and Thomas, unlike us, consider unrestricted allocations and exogenous returns).

The remainder of the proof is taken from the proof of Theorem 16.3.1 in Cover and Thomas (2006) with no significant changes. Using (21) repeatedly, we obtain $\mathrm{E}[S_t/S_t^*] \leq 1$, where S_t and S_t^* represent wealth accumulated in the first t periods under the generic and the optimal policy. Then, by Markov's inequality (Billingsley, 2012, p. 85), we have

$$\Pr\left(\frac{S_t}{S_t^*} \geq c_t\right) \leq \frac{1}{c_t}.$$

Hence,

$$\Pr\left(\frac{1}{t} \ln \frac{S_t}{S_t^*} \geq \frac{1}{t} \ln c_t\right) \leq \frac{1}{c_t}.$$

Setting $c_t = t^2$, we obtain

$$\sum_{t=1}^{\infty} \Pr\left(\frac{1}{t} \ln \frac{S_t}{S_t^*} \geq \frac{2 \ln t}{t}\right) \leq \sum_{t=1}^{\infty} \frac{1}{t^2} < \infty.$$

Finally, by the Borel-Cantelli Lemma, for almost every sequence of states ω_t, there exists T such that $\frac{1}{t} \ln \frac{S_t}{S_t^*} < \frac{2 \ln t}{t}$ for $t > T$, which implies the statement in the proposition. □

A.2 Proof of Corollary 1

Proof. Denote ε_j simply by ε. The set of the feasible policies is, for each ε,

$$\mathcal{Q}_\varepsilon = \left\{q_\varepsilon(i,\omega) : q_\varepsilon(i,\omega) = q(i,\omega) e^{\mathbb{1}_{i=j}\varepsilon} \text{ and } q(i,\omega) \in \mathcal{Q}\right\}.$$

By Lemma 2, the optimized growth rate $r^*(\varepsilon)$ equals

$$-\min_{\mathbf{p}(i,\omega),q(i,\omega)} \quad \mathrm{KL}\left(\mathbf{p}(i,\omega) \parallel q(i,\omega)e^{\mathbb{1}_{i=j}\varepsilon}\right) - \mathrm{H}\left(\mathbf{p}_0(\omega)\right)$$

$$\begin{aligned} s.t. \quad & \mathbf{p}(\omega) = \mathbf{p}^0(\omega) \\ & q(i,\omega) \in \mathcal{Q}. \end{aligned}$$

The chain rule (Lemma 4) implies:

$$\mathrm{KL}\left(\mathbf{p}(i,\omega) \parallel q(i,\omega)e^{\mathbb{1}_{i=j}\varepsilon}\right) = \mathrm{KL}\left(\mathbf{p}(i) \parallel \mathbf{q}(i)\right) + \sum_i \mathbf{p}(i)\,\mathrm{KL}\left(\mathbf{p}(\omega \mid i) \parallel q(\omega \mid i)e^{\mathbb{1}_{i=j}\varepsilon}\right).$$

The envelope theorem then implies that

$$\frac{dr^*}{d\varepsilon}(0) = \mathbf{p}^*(j),$$

where $\mathbf{p}^*(j) = \sum_\omega \mathbf{p}^*(j,\omega)$ and $\mathbf{p}^*(i,\omega)$ is the optimizer jointly with $q^*(i,\omega)$. We have used that $\partial_\varepsilon \mathrm{KL}\left(\mathbf{p}(\omega \mid i) \parallel q(\omega \mid i)e^{\mathbb{1}_{i=j}\varepsilon}\right) = -\mathbb{1}_{i=j}$. The corollary then follows from the fact that $\mathbf{m}_{q^*}(j) = \mathbf{p}^*(j)$ by Lemma 3. □

A.3 Proof of Lemma 1

Proof. We interpret the expression of interest, $\ln \sum_x q(x)$, as the growth rate of the multiplicative process $\left(\sum_x q(x)\right)^t$, writing

$$\ln \sum_x q(x) = \frac{1}{t} \ln \Big(\sum_x q(x) \Big)^t,$$

with $t \in \mathbb{N}$. Denoting a generic sequence of length t as $\mathbb{x} = (x_1, \ldots, x_t) \in X^t$, we rewrite the sum on the right side as a sum over all such sequences, giving

$$\ln \sum_x q(x) = \frac{1}{t} \ln \sum_{\mathbb{x} \in X^t} \prod_{x \in X} q(x)^{\mathbf{p}_{\mathbb{x}}(x)t}, \tag{22}$$

where $\mathbf{p}_{\mathbb{x}} \in \Delta(X)$, defined by

$$\mathbf{p}_{\mathbb{x}}(x) = \frac{1}{t} \sum_{\tau=1}^{t} \mathbb{1}_{x_\tau = x},$$

is the empirical distribution of the sequence $\mathbb{x}$. We rearrange the summand from (22) as follows,[13]

$$\begin{aligned}
\prod_{x \in X} q(x)^{\mathbf{p}_{\mathbb{x}}(x)t} &= \exp\left[t \times \sum_x \mathbf{p}_{\mathbb{x}}(x) \ln q(x)\right] \\
&= \exp\left[t \times \left(\sum_x \mathbf{p}_{\mathbb{x}}(x) \ln \mathbf{p}_{\mathbb{x}}(x) - \sum_x \mathbf{p}_{\mathbb{x}}(x) \ln \frac{\mathbf{p}_{\mathbb{x}}(x)}{q(x)}\right)\right] \\
&= \exp\left[-t \times \big(\,\mathrm{H}\left(\mathbf{p}_{\mathbb{x}}\right) + \mathrm{KL}\left(\mathbf{p}_{\mathbb{x}} \parallel q\right)\big)\right]. \qquad (23)
\end{aligned}$$

Since all sequences with the same empirical distribution generate the same value of the last expression in (23), we can substitute (23) into (22) to obtain

$$\ln \sum_x q(x) = \frac{1}{t} \ln \sum_{\mathbf{p} \in \Delta_t} \exp\left[-t \times \big(\mathrm{H}(\mathbf{p}) + \mathrm{KL}(\mathbf{p} \parallel q)\big)\right] n(\mathbf{p}), \qquad (24)$$

where $\Delta_t \subset \Delta(X)$ is the set of the empirical distributions that can be generated by sequences $\mathbb{x}$ of length t, and $n : \Delta_t \to \mathbb{N}$ returns for each empirical distribution $\mathbf{p}$ the number $n(\mathbf{p})$ of sequences of length t that generate such an empirical distribution. (We suppress dependence of $n(\cdot)$ on t in our notation.)

The number of the sequences that generate an empirical distribution $\mathbf{p}$ can be approximated as follows.

$$\frac{1}{(t+1)^{|X|}} \exp[t \times \mathrm{H}(\mathbf{p})] \leq n(\mathbf{p}) \leq \exp[t \times \mathrm{H}(\mathbf{p})];$$

for all $\mathbf{p} \in \Delta_t$. See Theorem 11.1.3 in Cover and Thomas (2006) for these bounds. Substituting this into (24), we obtain the bounds

$$\begin{aligned}
\frac{1}{t} \ln \sum_{\mathbf{p} \in \Delta_t} \exp\left[-t \times \mathrm{KL}(\mathbf{p} \parallel q)\right] - \frac{|X| \ln(t+1)}{t} \quad &\leq \quad \ln \sum_x q(x) \leq \\
&\quad\ \frac{1}{t} \ln \sum_{\mathbf{p} \in \Delta_t} \exp\left[-t \times \mathrm{KL}(\mathbf{p} \parallel q)\right].
\end{aligned}$$

Since the term $\frac{|X| \ln(t+1)}{t}$ vanishes as t diverges, we have proven that

$$\ln \sum_x q(x) = \lim_{t \to \infty} \frac{1}{t} \ln \sum_{\mathbf{p} \in \Delta_t} \exp\left[-t \times \mathrm{KL}(\mathbf{p} \parallel q)\right]. \qquad (25)$$

[13] If $\mathbf{q}$ is a distribution, then each summand is the sample likelihood.

To conclude the proof we use (25) to establish lower and upper bounds for $\ln \sum_x q(x)$ that both equal $-\min_{\mathbf{p}\in\Delta(X)} \mathrm{KL}(\mathbf{p} \parallel q)$, as needed. For the lower bound, we replace the sum on the right of (25) by its largest summand. To this end, letting $\mathbf{p}_t^* \in \arg\min_{\mathbf{p}\in\Delta_t} \mathrm{KL}(\mathbf{p} \parallel q)$, we have

$$\begin{aligned}
\ln \sum_x q(x) &\geq \lim_{t\to\infty} \frac{1}{t} \ln \exp\left[-t \times \mathrm{KL}(\mathbf{p}_t^* \parallel q)\right] \\
&= -\lim_{t\to\infty} \mathrm{KL}(\mathbf{p}_t^* \parallel q) \\
&= -\min_{\mathbf{p}\in\Delta(X)} \mathrm{KL}(\mathbf{p} \parallel q).
\end{aligned}$$

For the upper bound observe,

$$\begin{aligned}
\ln \sum_x q(x) &= \lim_{t\to\infty} \frac{1}{t} \ln \sum_{\mathbf{p}\in\Delta_t} \exp\left[-t \times \mathrm{KL}(\mathbf{p} \parallel q)\right] \\
&= -\min_{\mathbf{p}\in\Delta(X)} \mathrm{KL}(\mathbf{p} \parallel q) + \lim_{t\to\infty} \frac{1}{t} \ln \sum_{\mathbf{p}\in\Delta_t} \exp\left[-t \times \big(\mathrm{KL}(\mathbf{p} \parallel q) - \min_{\mathbf{p}\in\Delta(X)} \mathrm{KL}(\mathbf{p} \parallel q)\big)\right] \\
&\leq -\min_{\mathbf{p}\in\Delta(X)} \mathrm{KL}(\mathbf{p} \parallel q) + \lim_{t\to\infty} \frac{1}{t} \ln(t+1)^{|X|} \\
&= -\min_{\mathbf{p}\in\Delta(X)} \mathrm{KL}(\mathbf{p} \parallel q),
\end{aligned}$$

where we have used the fact that the summands in the sum from the second line are bounded between 0 and 1. Thus, the sum from the second line is bounded from above by the number of the empirical distributions, $|\Delta_t|$, attainable by sequences of the length t. The number of such distributions is bounded by $(t+1)^{|X|}$ because for each $x \in X$, $\mathbf{p}(x)$ attains values $\frac{0}{t}, \frac{1}{t}, \ldots, \frac{t}{t}$. □

A.4 Proof of Lemma 2

We first verify that the standard result, the so-called chain rule, continues to hold for the extended definition of the KL-divergence with $\mathrm{KL} : \Delta(X) \times \mathbb{R}_+^X \to \mathbb{R} \cup \{\infty\}$.

Lemma 4 (Chain Rule)**.** *Consider any distribution $\mathbf{p}(x, y)$, generalized distribution $q(x, y)$ on $X \times Y$ and a pair $q(x)$ and $q(y \mid x)$ such that $q(x, y) = q(x)q(y \mid x)$. Assume $\mathbf{p}(x) > 0$ for x in a set $X^* \subseteq X$. Then,*

$$\mathrm{KL}\left(\mathbf{p}(x, y) \parallel q(x, y)\right) = \mathrm{KL}\left(\mathbf{p}(x) \parallel q(x)\right) + \sum_{x\in X^*} \mathbf{p}(x)\, \mathrm{KL}\left(\mathbf{p}(y \mid x) \parallel q(y \mid x)\right).$$

Proof.

$$\begin{aligned}
\mathrm{KL}\left(\mathbf{p}(x,y) \parallel q(x,y)\right) &= \sum_{(x,y)\in X\times Y} \mathbf{p}(x,y) \ln \frac{\mathbf{p}(x,y)}{q(x,y)} \\
&= \sum_{(x,y)\in X^*\times Y} \mathbf{p}(x)\mathbf{p}(y \mid x) \ln \frac{\mathbf{p}(x)\mathbf{p}(y \mid x)}{q(x)q(y \mid x)} \\
&= \sum_{x\in X^*} \mathbf{p}(x) \ln \frac{\mathbf{p}(x)}{q(x)} + \sum_{x\in X^*}\sum_{y\in Y} \mathbf{p}(x)\mathbf{p}(y \mid x) \ln \frac{\mathbf{p}(y \mid x)}{q(y \mid x)} \\
&= \mathrm{KL}\left(\mathbf{p}(x) \parallel q(x)\right) + \sum_{x\in X^*} \mathbf{p}(x)\, \mathrm{KL}\left(\mathbf{p}(y \mid x) \parallel q(y \mid x)\right),
\end{aligned}$$

where we used for the third equality that $\mathbf{p}(x,y)$ is a distribution and thus $\sum_y \mathbf{p}(y \mid x) = 1$. □

Proof of Lemma 2. Fix ω and observe that

$$\begin{aligned}
\ln \sum_i q(i,\omega) &= -\min_{\mathbf{p}(i|\omega)\in\Delta(I)} \mathrm{KL}\left(\mathbf{p}(i \mid \omega) \parallel q(i,\omega)\right) \\
&= -\min_{\mathbf{p}(i|\omega)\in\Delta(I)} \mathrm{KL}\left(\mathbf{p}(i \mid \omega) \parallel \mathbf{q}(i \mid \omega)q(\omega)\right) \\
&= -\min_{\mathbf{p}(i|\omega)\in\Delta(I)} \left\{\mathrm{KL}\left(\mathbf{p}(i \mid \omega) \parallel \mathbf{q}(i \mid \omega)\right) - \ln q(\omega)\right\},
\end{aligned}$$

where the minimization is over state-contingent distributions $\mathbf{p}(i \mid \omega) \in \Delta(I)$. We treat here $q(i,\omega)$ as function of i only, with a fixed parameter ω. The first inline equality follows from Lemma 1. The last equality follows from the fact that

$$\mathrm{KL}(\mathbf{p} \parallel \lambda q) = \mathrm{KL}(\mathbf{p} \parallel q) - \ln \lambda$$

for any distribution $\mathbf{p}$, generalized distribution q and $\lambda > 0$.

Taking expectation with respect to ω gives that a policy $q(i,\omega)$ induces the growth rate

$$
\begin{aligned}
& \sum_{\omega} \mathbf{p}^0(\omega) \ln \sum_i q(i,\omega) \\
= \quad & -\sum_{\omega} \mathbf{p}^0(\omega) \min_{\mathbf{p}(i|\omega)\in\Delta(I)} \mathrm{KL}\left(\mathbf{p}(i \mid \omega) \parallel \mathbf{q}(i \mid \omega)\right) + \sum_{\omega} \mathbf{p}^0(\omega) \ln q(\omega) \\
= \quad & -\sum_{\omega} \mathbf{p}^0(\omega) \min_{\mathbf{p}(i|\omega)} \mathrm{KL}\left(\mathbf{p}(i \mid \omega) \parallel \mathbf{q}(i \mid \omega)\right) - \mathrm{KL}\left(\mathbf{p}^0(\omega) \parallel q(\omega)\right) - \mathrm{H}\left(\mathbf{p}^0(\omega)\right) \\
= \quad & -\min_{\mathbf{p}(i,\omega)\in\Delta(I\times\Omega)} \left\{ \mathrm{KL}\left(\mathbf{p}(\omega) \parallel q(\omega)\right) + \sum_{\omega} \mathbf{p}(\omega)\, \mathrm{KL}\left(\mathbf{p}(i \mid \omega) \parallel \mathbf{q}(i \mid \omega)\right) \right\} - \mathrm{H}\left(\mathbf{p}^0(\omega)\right) \\
& \qquad\qquad \text{s.t.} \quad \mathbf{p}(\omega) = \mathbf{p}^0(\omega).
\end{aligned}
$$

Applying the chain rule, Lemma 4, to the last objective function gives the result. □

A.5 Proof of Lemma 3

Proof. Using the chain rule, we can write the objective from optimization in (7) as

$$
\begin{aligned}
& \mathrm{KL}\left(\mathbf{p}(i,\omega) \parallel q(i,\omega)\right) \\
= \quad & \mathrm{KL}\left(\mathbf{p}(\omega) \parallel q(\omega)\right) + \sum_{\omega} \mathbf{p}(\omega)\, \mathrm{KL}\left(\mathbf{p}(i \mid \omega) \parallel \mathbf{q}(i \mid \omega)\right) \\
= \quad & \mathrm{KL}\left(\mathbf{p}^0(\omega) \parallel q(\omega)\right) + \sum_{\omega} \mathbf{p}^0(\omega)\, \mathrm{KL}\left(\mathbf{p}(i \mid \omega) \parallel \mathbf{q}(i \mid \omega)\right),
\end{aligned}
$$

where we have used for the second equality that $\mathbf{p}(\omega)$ is constrained to equal $\mathbf{p}^0(\omega)$. Minimization of this objective with respect to $\mathbf{p}(i \mid \omega)$ for each ω implies that $\mathbf{p}_q^*(i \mid \omega) = \mathbf{q}(i \mid \omega)$. The lemma then follows from the definition of $\mathbf{m}_q(i,\omega)$ in (8): $\mathbf{m}_q(i,\omega) = \mathbf{p}^0(\omega)\mathbf{q}(i \mid \omega) = \mathbf{p}_q^*(i,\omega)$. □

A.6 Proof of Proposition 5

Proof. The result follows from Theorem 3 and the subsequent Remark in Csiszar and Tusnady (1984). Their result requires the condition that (in their notation) $p_0 \in P$ is positive for exactly those $x \in X$ for which there exist $p \in P$ and $q \in Q$ such that $p(x)q(x) > 0$. The condition is satisfied in our setting because $\mathbf{m}_0(i,\omega)$ (the counterpart of their p_0) is assumed to be interior and $\mathcal{Q}$ is assumed to contain an interior $q(i,\omega)$ (giving the counterpart of their p and q with $p(x)q(x)$ interior). Their result ensures that our $q_k(i,\omega)$ converges to the growth-maximizing policy for those (i,ω) such that $\mathbf{m}^*(i,\omega) > 0$. In our specific setting, $q_k(i,\omega)$ converges also for those (i,ω) such that $\mathbf{m}^*(i,\omega) = 0$, because, $\mathbf{m}_k(i \mid \omega) = \mathbf{q}_{k-1}(i \mid \omega)$ (by Lemma 3) and hence $\mathbf{q}_k(i \mid \omega)$ converges to 0 and also $q_k(i,\omega) \to 0$. □

A.7 Proof of Proposition 6 and Corollary 4

Proof of Proposition 6. Fix individual j and state $\tilde{\omega}$. The derivative of the growth rate from (1) with respect to the policy is

$$\begin{aligned} \partial_{q(j,\tilde{\omega})} \sum_{\omega} \mathbf{p}^0(\omega) \ln\Big(\sum_i q(i,\omega)\Big)\Big|_{q=q^*} &= \frac{\mathbf{p}^0(\tilde{\omega})}{\sum_i q^*(i,\tilde{\omega})} \\ &= \frac{\mathbf{p}^0(\tilde{\omega})}{q^*(\tilde{\omega})}. \end{aligned}$$

Since the optimization objective $\sum_\omega \mathbf{p}^0(\omega) \ln(\sum_i q(i,\omega))$ is concave in q and the feasible set $\mathcal{Q}$ is convex, the set of maximizers is unaffected by the linearization of the objective around the optimum. Thus, $q^*(i,\theta)$ is a growth-maximizing policy if and only if it maximizes $\sum_{i,\omega} \frac{\mathbf{p}^0(\omega)}{q^*(\omega)} q(i,\omega)$ on $\mathcal{Q}$. Summing up the latest objective across i leads to the inequality from (16). □

Proof of Corollary 4. If the policy $q^*(i,\omega)$ is positive, then

$$\begin{aligned} \partial_{q(j,\tilde{\omega})} \mathrm{KL}\left(\mathbf{m}_{q^*}(i,\omega) \parallel q(i,\omega)\right)\big|_{q=q^*} &= -\frac{\mathbf{m}_{q^*}(j,\tilde{\omega})}{q^*(j,\tilde{\omega})} \\ &= -\frac{\mathbf{p}^0(\tilde{\omega})\mathbf{q}^*(j \mid \tilde{\omega})}{q^*(\tilde{\omega})\mathbf{q}^*(j \mid \tilde{\omega})} \\ &= -\frac{\mathbf{p}^0(\tilde{\omega})}{q^*(\tilde{\omega})}, \end{aligned}$$

where we have used that the merit distribution induced by policy q^* is $\mathbf{m}_{q^*}(i,\omega) = \mathbf{p}^0(\omega)\mathbf{q}^*(i \mid \omega)$ for the second equality and positivity of $\mathbf{q}^*(j \mid \tilde{\omega})$ for the third equality. Since the objective in the minimization problem (14) is convex, $q^*(i,\omega)$ that solves (14) also maximizes

$$\sum_{i,\omega} \frac{\mathbf{p}^0(\omega)}{q^*(\omega)} q(i,\omega) = \mathrm{E}_{\mathbf{p}^0(\omega)} \frac{q(\omega)}{q^*(\omega)}$$

on $\mathcal{Q}$. Hence, $q^*(i,\omega)$ satisfies the necessary and sufficient condition from Proposition 6. $\square$

A.8 Proof of Proposition 7

Proof. Since the allocation $\mathbf{q}(i) \in \mathcal{A} = \Delta(I)$ is unconstrained, the naive meritocracy principle from Proposition 2 implies that $\mathbf{q}^*(i) = \mathbf{m}_{q^*}(i)$ (which equals $\mathbf{p}^*(i)$). This simplifies the objective in (20) as follows.

$$\begin{aligned}
\mathrm{KL}\left(\mathbf{p}(i,\omega) \parallel q(i,\omega)\right) &= \mathrm{KL}\left(\mathbf{p}(i) \parallel \mathbf{q}(i)\right) + \sum_{i \in I^*} \mathbf{p}(i)\,\mathrm{KL}\left(\mathbf{p}(\omega \mid i) \parallel q(\omega \mid i)\right) \\
&= \sum_{i \in I^*} \mathbf{p}(i)\,\mathrm{KL}\left(\mathbf{p}(\omega \mid i) \parallel q(\omega \mid i)\right) \\
&= -\sum_{i \in I^*} \mathbf{p}(i)\Big(\mathrm{E}_{\mathbf{p}(\omega|i)}\, u(i,\omega) + \mathrm{H}\left(\mathbf{p}(\omega \mid i)\right)\Big) \\
&= -\left(\mathrm{E}_{\mathbf{p}(i,\omega)}\, u(i,\omega) - \mathrm{I}_{\mathbf{p}(i,\omega)}\right) - \mathrm{H}\left(\mathbf{p}^0(\omega)\right),
\end{aligned}$$

where $I^* \subseteq I$ is the set of i with $\mathbf{p}(i) > 0$. We have used the chain rule for the first equality and the fact that $\mathrm{KL}\left(\mathbf{p}(i) \parallel \mathbf{q}(i)\right) = 0$ at optimum for the second equality. The third equality follows from the definitions of the KL-divergence and the entropy and from the assumed relationship $q(\omega \mid i) = \exp u(i,\omega)$. The last equality follows from the fact that mutual information $\mathrm{I}_{\mathbf{p}(i,\omega)} = \mathrm{H}\left(\mathbf{p}(\omega)\right) - \sum_i \mathbf{p}(i)\,\mathrm{H}\left(\mathbf{p}(\omega \mid i)\right)$ and $\mathbf{p}(\omega) = \mathbf{p}^0(\omega)$.

We have demonstrated that, after optimization of $\mathbf{q}(i)$, the objectives in the minimization problem (20) and the rational-inattention problem (19) differ only in the sign and the term $\mathrm{H}\left(\mathbf{p}^0(\omega)\right)$, where this term is independent of the controls, as needed. $\square$

References

Aghion, P., E. Caroli, and C. Garca-Pealosa (1995). Inequality and economic growth: The perspective of the new growth theories. *Journal of Economic Literature 37*(4), 1615–1660.

Alesina, A. and G.-M. Angeletos (2005). Fairness and redistribution. *American Economic Review 95*(4), 960–980.

Almås, I., A. W. Cappelen, and B. Tungodden (2020). Cutthroat capitalism versus cuddly socialism: Are Americans more meritocratic and efficiency-seeking than scandinavians? *Journal of Political Economy 128*(5), 1753–1788.

Andre, P. (2022). Shallow meritocracy. Technical report, SSRN Discussion Paper Series, Discussion Paper 318.

Aridor, G., F. Brechi, and M. Woodford (2020). Adaptive efficient coding: A variational auto-encoder approach. Technical report, Columbia University. bioRxiv preprint doi: https://doi.org/10.1101/2020.05.29.124453.

Aridor, G., R. A. da Silveira, and M. Woodford (2023). Information-constrained coordination of economic behavior. Technical report, Northwestern Univeristy, ENS Paris and University of Basel, and Columbia University.

Barro, R. J. (2000). Inequality and growth in a panel of countries. *Journal of Economic Growth 5*(1), 5–32.

Berg, A., J. D. Ostry, C. G. Tsangarides, and Y. Yakhshilikov (2018). Redistribution, inequality, and growth: New evidence. *Journal of Economic Growth 23*, 259–305.

Berk, R. H. (1966). Limiting behavior of posterior distributions when the model is incorrect. *The Annals of Mathematical Statistics 37*(1), 51–58.

Billingsley, P. (2012). *Probability and Measure* (Anniversary ed.). New York: John Wiley and Sons.

Caplin, A. and M. Dean (2013). Behavioral implications of rational inattention with Shannon entropy. Technical report, National Bureau of Economic Research.

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

Cappelen, A. W., K. O. Moene, S.-E. Skjelbred, and B. Tungodden (2023). The merit primacy effect. *The Economic Journal 133*(651), 951–970.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences 36*(3), 181–204.

Cover, T. M. and J. A. Thomas (2006). *Elements of Information Theory* (Second ed.). New York: John Wiley and Sons.

Csiszar, I. and G. Tusnady (1984). Information geometry and alternating minimization procedures. *Statistics and Decisions, Dedewicz 1*, 205–237.

Dayan, P., G. E. Hinton, R. M. Neal, and R. S. Zemel (1995). The Helmholtz machine. *Neural Computation 7*(5), 889–904.

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

Friston, K. (2005). A theory of cortical responses. *Philosophical transactions of the Royal Society B: Biological sciences 360*(1456), 815–836.

Jordan, M. I., Z. Ghahramani, T. S. Jaakkola, and L. K. Saul (1999). An introduction to variational methods for graphical models. *Machine Learning 37*, 183–233.

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

Kussell, E. and S. Leibler (2005). Phenotypic diversity, population growth, and information in fluctuating environments. *Science 309*(5743), 2075–2078.

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.

Robson, A., L. Samuelson, and J. Steiner (2023). Decision theory and stochastic growth. *American Economic Review: Insights*. Forthcoming.

Saez, E. and G. Zucman (2020). The rise of income and wealth inequality in America: Evidence from distributional macroeconomic accounts. *Journal of Economic Perspectives 34*(4), 3–26.

Stiglitz, J. E. (1969). Distribution of income and wealth among individuals. *Econometrica 37*(3), 382–397.

White, H. (1982). Maximum likelihood estimation of misspecified models. *Econometrica 50*(1), 1–25.