Growth and Redistribution: The Hedging Perspective*

Larry Samuelson
Yale University

Jakub Steiner
University of Zurich, CERGE-EI and CTS

December 6, 2024

Abstract

We investigate the impact of wealth redistribution on economic growth, building on Kelly's (1956) optimal investment portfolio theory. A growth-optimal policy redistributes wealth from 'lucky' overperforming individuals to underperforming ones, minimizing the systematic component of this redistribution in a myopic fashion. That is, the optimal policy minimizes the discrepancy between endowments and outcomes, counterfactually taking outcomes as independent of endowments. The myopia in this result follows from a decoupling argument that allows us to model the planner as independently choosing a growth-maximizing policy and a pattern of wealth circulation.

*This paper supersedes our preliminary working paper, Samuelson and Steiner (2023). We thank the editor and three referees for helpful comments and suggestions. We have also 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 and GAČR grant 24-10145S.

Growth and Redistribution: The Hedging Perspective

1 Introduction

We investigate a relatively underexplored aspect of the relationship among inequality, redistribution, and economic growth. Economic growth is inherently stochastic, with some fortunate individuals experiencing rapid wealth accumulation while others lag behind, even when controlling for productivity. Maximizing the long-run economic growth rate necessitates redistribution to hedge against these luck disparities.

To emphasize the hedging aspect of wealth redistribution, we set aside all the conventional modeling elements of economic growth. Incentive effects, production, technological progress, and other factors are all condensed into a black box that determines each individual's end-of-period wealth based on their initial wealth and the realization of a random state. Different individuals have distinct such *return functions*, reflecting persistent differences in their skills, education, opportunities, and other factors. How should a planner, capable of periodically redistributing wealth, optimally respond to individuals' ex-ante differences and varying degrees of luck?

Formally, we build upon Kelly's (1956) classical study of optimal private investment portfolios. Kelly's investor repeatedly allocates wealth to assets with uncertain returns to maximize the long-term growth rate of the portfolio's value. Kelly provides two main insights. First, a growth-optimal portfolio involves hedging—it comprises a variety of assets, including those with inferior expected returns, and is maintained by periodically redistributing wealth from unexpectedly overperforming to underperforming assets. Second, the optimal portfolio eliminates the systematic component of this redistribution, ensuring that the expected wealth share of each asset remains constant before and after rebalancing.

We view individuals as assets yielding uncertain returns. Again, the optimal policy periodically redistributes wealth to maintain a growth-maximizing allocation, akin to rebalancing a financial portfolio. To fit our economic context, though, we extend the standard portfolio choice setting in three ways. First, we allow the planner to control the state-contingent returns of individuals, subject to feasibility constraints, thereby capturing the effects of economic policies on individuals' varying productivities. Second, to address inequality concerns,

we impose constraints on the planner's choice of wealth allocation. Finally, we account for the interaction between returns and wealth allocations, acknowledging that redistribution can influence incentives and, consequently, individual returns.

To a first approximation, our results align with Kelly's insights on the growth-optimal portfolio. The growth-optimal economic policy involves hedging, allocating wealth among a variety of individuals, including those with inferior expected returns. Additionally, the growth-optimal policy routinely redistributes wealth from accidental economic winners to losers, thereby maintaining an optimal hedge. However, in our context, constraints may render the full elimination of systematic redistribution infeasible. Instead, the growth-optimal policy *myopically* minimizes this redistribution.

To formulate this result, we compare two types of wealth distributions. The *endowment* distribution refers to the wealth distribution periodically controlled by the planner through redistribution. The *outcome* distribution represents the wealth distribution that emerges after endowments have been augmented by returns, averaged over random states. Our main result states that the growth-optimal endowment distribution minimizes the Kullback-Leibler divergence from the induced outcome distribution, treating the latter as given.

Somewhat counterintuitively, this minimization is myopic; the optimality condition requires the planner to treat the outcome distribution as fixed and approximate it with an endowment distribution as closely as possible, while neglecting the dependence of outcomes on endowments. Acknowledging that the outcomes depend on the endowments could further reduce systematic redistribution, but would do so at the expense of the economy's growth.

Our analytical approach is novel. We express the economy's growth rate as a function of economic policy and the resulting pattern of wealth circulation, which we define to describe the flow of wealth among individuals via redistribution. Although wealth circulation depends on policy, our proof construction *decouples* them, enabling us to treat policy and the pattern of wealth circulation as separate controls determining the growth rate. The source of the myopia in our main result is then clear: maximizing the growth rate requires each control to be individually optimal when the other is fixed at its respective optimum. We finalize the result by linking the pattern of wealth circulation to the outcome distribution.

This decoupling admits an economic interpretation: we divide aggregate wealth into infinitesimal fractions, each of which stochastically circulates within

the economy, moving randomly among individuals through redistribution. These wealth fractions grow at different rates, depending on their distinct realized patterns of circulation. The fastest-growing fractions ultimately prevail and determine the economy's overall growth rate. Consequently, maximizing the economy's growth rate involves a joint optimization over two decoupled controls: the endowment distribution and the pattern of wealth circulation.

The following section reviews the literature. Section 3 subsequently sets up the model and presents the main result. Section 4 outlines our analytical approach, which traces the circulation of wealth driven by policies and achieves a decoupling between the policy control and wealth circulation. Example 5 in this section sketches an extension of the model, incorporating state inertia and practical redistribution policies. The online appendix provides additional results characterizing optimal policies and highlights connections between the planner's problem and the rational inattention problem, enabling the application of established solution methods from the rational inattention literature.

2 Literature

The theoretical literature identifies multiple channels through which inequality and redistribution can either promote or hinder growth, primarily explored within deterministic models that emphasize effects on productivity via incentives, resource allocation, and technological advancement. Aghion et al. (1999) and Mdingi and Ho (2021) provide surveys. The empirical literature also reports mixed results, as surveyed by Forbes (2000) and once more by Mdingi and Ho (2021). We abstract from most of the considerations appearing in the literature to focus on optimal hedging. We view our analysis as complementary, highlighting that the inherently stochastic nature of growth brings forth considerations related to inequality, redistribution, and growth that are often absent from the existing literature.

Mirrlees (1971) (see also the exposition in Mirrlees, 1986) is the foundational work for a substantial body of research on optimal redistribution. Mirrlees' framework maximizes utilitarian welfare within a static model, constrained by incentive compatibility due to agents' private information about their productivities.

Farhi and Werning (2013), Kapička (2013), and Golosov et al. (2016) extend the Mirrlees model to dynamic models, aligning it more closely with our tem-

poral framework. See Stantcheva (2020) for a comprehensive survey. In these models, the agent's private information evolves according to a Markov process, necessitating compensation for both the immediate and future implications of revealing this information.

We maximize the growth rate rather than utilitarian welfare, and simplify the interaction between redistribution and output by imposing an exogenous feasibility constraint and assuming all information is public. We employ a dynamic model but abstract from serial correlations, except in the extension outlined in Example 5.

Our focus on hedging implies a close link between our work and the portfolio choice literature. Kelly (1956) provides the foundational analysis of portfolio choice, introducing the 'Kelly criterion' to maximize the long-run growth rate of private wealth. Analogously, we maximize the growth rate, replacing assets with individuals. Our model extends the standard portfolio choice framework in three aspects: it incorporates constraints on wealth distributions, enables the planner to influence returns as well as wealth shares, and allows for interactions between wealth shares and returns.

Most of the portfolio choice models maximize lifetime expected utility (e.g., Merton, 1973; Merton and Samuelson, 1974). Related to our constrained planner, Campbell and Sigalov (2022) restrict the investor in Merton's framework by requiring continual consumption of the expected return of her portfolio. Given our emphasis on economic growth rather than individual investment strategies, we maximize the long-run aggregate growth rate instead of utility and abstract from explicit consumption modeling. Similarly to Campbell and Sigalo, we can capture a requirement for the economy to consume a fraction of aggregate wealth each period—perhaps for entitlement spending—by formally allocating a proportion of wealth each period to an additional 'government' individual who generates zero return.

Our extensions of Kelly's framework necessitate the derivation of new optimality conditions, which we accomplish by leveraging a novel connection between stochastic growth and variational inference. This literature, originating with Jordan et al. (1999) and Kingma and Welling (2013), characterizes the maximum likelihood estimation of statistical models involving both observable and latent variables and highlights a link between estimation and the updating of beliefs about latent variables. We establish a formal equivalence between the variational inference problem and growth maximization with redistribution. Under this newly established analogy, policy is analogous to an estimate, whereas

the induced pattern of wealth circulation parallels the updates.

This paper is one of our three projects that highlight the role of large deviations—atypical sequences of draws—in stochastic growth processes. In Robson et al. (2023), we noted that the wealth concentration resulting from stochastic growth without redistribution can be analyzed through an equivalence to a rational-inattention problem. This paper examines the impact of redistribution on growth and exploits the equivalence between growth with redistribution and the variational inference problem. In Samuelson and Steiner (2024), we focus on inference, utilizing the growth connection to provide a microfounded interpretation of the variational inference problem.

3 Economic Growth

3.1 Growth Maximization

The population consists of a finite set I of individuals i, and there is a finite set Ω of states ω. The economy begins in period 1 with an initial aggregate endowment $W_1 = 1$ of perfectly divisible wealth. We consider wealth to encompass physical, financial, human, and social capital, which are aggregated into a single, malleable quantity at this level of abstraction.

A time-invariant *endowment distribution* $p \in \Delta(I)$ allocates a share $p(i)$ of the current aggregate wealth to each individual $i \in I$ at the beginning of every period $t = 1, 2, \ldots$. Afterwards, in each period, nature independently draws a state ω_t according to an interior distribution $q^0 \in \Delta(\Omega)$. Subsequently, each individual i earns a nonnegative gross return per unit of her endowment. The individual return depends on the current state ω_t and the individual i; we denote the time-invariant *return function* as $r(i, \omega_t)$. The resulting aggregate wealth is then redistributed at the beginning of the next period according to p, and so on, generating aggregate wealth

$$W_t = W_{t-1} \left(\sum_i p(i) r(i, \omega_t) \right) = \prod_{s=1}^{t} \left(\sum_i p(i) r(i, \omega_s) \right)$$

at the end of each period t. Taking the logarithm and applying the law of large numbers to the realized sequence of states yields the expectation in the following expression. Consequently, the growth rate of the economy, $\ln W_t / t$,

almost surely converges, with

$$\lim_{t\to\infty} \frac{\ln W_t}{t} = \mathrm{E}_{q^0(\omega)} \ln \Big(\sum_i p(i) r(i,\omega) \Big). \tag{1}$$

A planner selects a policy pair $(p(i), r(i,\omega))$ from a set $\mathcal{P} \subseteq \Delta(I) \times \mathbb{R}_+^{I\times\Omega}$ of feasible policy pairs in order to maximize the growth rate (1).[1]

The set $\mathcal{P}$ is compact and includes at least one policy pair that yields a finite growth rate. The existence of an optimizer is then ensured.[2] The *growth-maximizing* policy pair is denoted by $p^*(i)$ and $r^*(i,\omega)$.

Example 1 (Expected Return vs. Growth Rate)**.** Maximizing the long-run growth rate of aggregate wealth, given in (1), differs from maximizing the expected aggregate return $\mathrm{E}_{q^0(\omega)} \sum_i p(i) r(i,\omega)$. For a stark illustration of the difference, consider two individuals and two states, $I = \Omega = \{1, 2\}$. The exogenous returns are specified as $r(i,\omega) = 2 \times \mathbb{1}_{i=\omega}$, implying that each individual i becomes bankrupt whenever her unfavourable state $\omega_t \neq i$ occurs. Maximization of the expected return would dictate allocating all wealth to the individual i^* with higher probability $q^0(i^*)$ of being productive, but such policy ensures the loss of all aggregate capital (and zero wealth thereafter) as soon as $\omega_t \neq i^*$. By contrast, the growth-maximizing policy allocates wealth in proportion to each individual's probability of avoiding bankruptcy, $p^*(i) = q^0(i)$, thereby ensuring a non-negative long-run growth rate. ▲

Naturally, societies face constraints in both elements of the policy pair. Equity constraints may restrict feasible endowments $p(i)$, while technological constraints on capital formation may restrict the returns $r(i,\omega)$. The feasibility of endowments and returns is typically interconnected. The choice of the redistribution scheme may affect incentives and, consequently, the feasibility of return

[1] Under our leading interpretation of the model, $i \in I$ indexes an individual's identity. However, the setting and results are also compatible with an alternative interpretation where i represents an ascending person's rank in the beginning-of-period wealth distribution. In this case, the feasible set $\mathcal{P}$ ensures the monotonicity of $p(i)$, and the return function depends on a person's wealth rank rather than identity. Under this alternative interpretation, the policy maintains a stationary wealth distribution but does not necessarily ensure a stationary wealth share for each individual. We thank a referee who suggested this interpretation. Yet another reinterpretation of the model incorporates consumption. Let individual i consume a share $c(i,\omega) \in [0,1]$ of her endowment in state ω and generate a gross return $r(i,\omega)$ on her residual wealth. Then our analysis applies with return function $r'(i,\omega) = (1 - c(i,\omega))r(i,\omega)$.

[2] If there exists a policy pair yielding a finite growth rate g, then the subset of $\mathcal{P}$ of policy pairs that produce growth rates of at least g is nonempty and compact. The objective function in equation (1) is continuous on this set, ensuring the existence.

functions, resulting in a constraint set $\mathcal{P}$ that is not a product set. Unlike much of the economic research that derives the set $\mathcal{P}$ from microfoundations, we treat $\mathcal{P}$ as a primitive and develop results applicable across all such sets.

The restriction to stationary policy pairs is without loss of optimality in settings with iid payoff states because the optimal stationary policy achieves at least as high a growth rate as any history-dependent policy. Let a feasible *history-dependent* policy map every history $(\omega_1, \ldots, \omega_{t-1})$ to a policy pair from $\mathcal{P}$. Applying Theorem 16.3.1 of Cover and Thomas (2006) to the current setting establishes:

Proposition 1 (Cover and Thomas (2006))**.** *Let the random variable W_t denote the aggregate wealth accumulated under a feasible history-dependent policy over the first t periods. Let W_t^* denote the aggregate wealth accumulated under the growth-maximizing stationary policy pair. Then, almost surely,*

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

3.2 Wealth Distributions

Our main result contrasts two wealth distributions. First, the endowment distribution $p(i)$, as introduced above and controlled by the planner, specifies the wealth shares at the beginning of each period. Second, we denote the wealth distribution at the end of each period as

$$o_{p,r}(i \mid \omega) = \frac{p(i)r(i,\omega)}{\sum_{j\in I} p(j)r(j,\omega)}, \tag{2}$$

which specifies the wealth shares of individuals i at the end of each period in which the state $\omega_t = \omega$. Furthermore,

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

denotes the long-run end-of-period wealth distribution, representing the probability that a dollar, randomly selected from the aggregate wealth at the end of a random period, belongs to individual i. We refer to $o_{p,r}(i \mid \omega)$ and $o_{p,r}(i)$ as the (state-contingent) *outcome distributions*.

Our main result provides a necessary condition for the growth-maximizing policy pair. It must minimize a measure of systematic redistribution, while taking the outcome distribution as given. To formulate this, we recall that

the Kullback-Leibler divergence between two probability distributions, $o(i)$ and $p(i)$, is defined as:

$$\mathrm{KL}\left(o(i) \parallel p(i)\right) := \sum_{i \in I} o(i) \ln \frac{o(i)}{p(i)},$$

which quantifies the discrepancy between the two distributions.[3]

3.3 Main Result

The optimal policy pair must satisfy a fixed-point condition. To state this condition, let $\mathcal{E}^* = \{p(i) : (p(i), r^*(i,\omega)) \in \mathcal{P}\}$ be the set of endowment distributions to which the planner can deviate, starting from the optimal policy pair, without altering the optimized return function $r^*(i,\omega)$. Section 4 proves:

Proposition 2 (Myopically Minimal Systematic Redistribution)**.** *If the endowment distribution $p^*(i)$ and the return function $r^*(i,\omega)$ jointly maximize the aggregate growth rate, then $p^*(i)$ minimizes the* KL*-divergence from the induced outcome distribution:*

$$p^*(i) \in \underset{p(i) \in \mathcal{E}^*}{\arg\min}\, \mathrm{KL}\left(o_{p^*,r^*}(i) \parallel p(i)\right). \tag{4}$$

In other words, the planner myopically minimizes the systematic redistribution of wealth. Among endowment distributions that are feasible in combination with the optimized return function, the planner selects the one most closely aligned with the outcome distribution. This optimization is myopic, meaning the planner proceeds as if the outcome distribution $o_{p,r^*}(i)$ is fixed at $o_{p^*,r^*}(i)$ and independent of the control $p(i)$. As a result, the planner generally does not choose the endowments $p(i)$ that minimize $\mathrm{KL}\left(o_{p,r^*}(i) \parallel p(i)\right)$. Such minimization would further decrease this measure of systematic redistribution by exploiting the link between the endowment distribution and the induced outcome distribution. However, this would generally decrease the growth rate.

A simple case arises when the planner is unconstrained in her choice of an endowment distribution. Then, the growth-optimal policy pair fully eliminates systematic redistribution.[4]

[3] We use the standard convention $0 \ln 0 = 0$.

[4] Recalling that $o_{p^*,r^*}(i) = \mathrm{E}_{q^0(\omega)} \frac{p^*(i) r^*(i,\omega)}{\sum_j p^*(j) r^*(j,\omega)}$, the corollary implies that when $p(i)$ is unconstrained, the expected relative return $\mathrm{E}_{q^0(\omega)} \frac{r^*(i,\omega)}{\sum_j p^*(j) r^*(j,\omega)}$ is equalized across all

Corollary 1 (Unconstrained Endowments)**.** *Assume that any endowment distribution $p(i) \in \Delta(I)$ is feasible along with each of the feasible return functions $r(i, \omega)$. Then, the optimal policy pair eliminates systematic redistribution:*

$$p^*(i) = o_{p^*, r^*}(i). \tag{5}$$

Cover and Thomas (2006, Section 16.2) derive the optimality condition (5) from the first-order conditions in a special case involving exogenous return functions. Proposition 2 extends this classical result to cases where constraints prevent the full elimination of systematic redistribution and where the planner controls the return function.

The following example illustrates statements of Corollary 1 and Proposition 2.

Example 2 (Illustrations of the Main Result)**.** Consider three individuals $i \in \{1, 2, 3\}$ and four states $\omega \in \Omega = I \cup \{g\}$. In each period, each individual either has a log-productivity of one, $\ln r(i, \omega_t) = 1$, or faces bankruptcy, in which case $r(i, \omega_t) = 0$. No individual experiences bankruptcy in the good state $\omega = g$; however, in state $\omega = i \in I$, all individuals except i face bankruptcy. The a priori probabilities of the individuals' favorable states are $q^0(1) = .01$, $q^0(2) = .02$, and $q^0(3) = .07$, leading to $q^0(g) = .9$.

Let us start with an unconstrained planner who can select any endowment distribution $p(i) \in \Delta(I)$. Corollary 1 states that such planner eliminates systematic redistribution. As we derive in Example 3, this implies a simple analytical solution. The planner selects the endowment shares

$$p^*(i) = q^0(i)/(1 - q^0(g)), \tag{6}$$

proportionally to the probabilities of states favoring the individuals. Calculating the outcome distribution generated by this optimal policy according to (2) confirms that $p^*(i) = o_{p^*}(i)$, as required.

The above uneven endowments may be deemed unacceptable. Accordingly, impose a constraint on the endowment distribution by limiting its Theil inequality index to no more than $\ln 3 - 1 \approx .099$.[5] As a result, the constrained-

individuals with positive endowments. See the online appendix for a link to the optimality conditions outlined by Caplin and Dean (2013) for the rational inattention problem with entropic cost.

[5] The Theil index for the endowments $p(i)$ of three individuals is given by $\frac{1}{3}\sum_i \frac{p(i)}{\overline{p}} \ln \frac{p(i)}{\overline{p}} = -\operatorname{H}\big(p(i)\big) + \ln 3$, where the average endowment $\overline{p} = 1/3$. Therefore, this inequality constraint

optimal endowment distribution becomes more egalitarian, yielding $p^*(1) \approx .192$, $p^*(2) \approx .264$, and $p^*(3) \approx .544$. This constrained planner systematically redistributes wealth, creating the wedge between endowment and outcome distributions, with $o_{p^*,r}(1) \approx .183$, $o_{p^*,r}(2) \approx 0.258$, and $o_{p^*,r}(3) \approx .559$. Treating this outcome distribution as fixed, the planner cannot match it with the endowment distribution due to the inequality constraint. Nonetheless, by Proposition 2, the planner myopically minimizes systematic redistribution by selecting an endowment distribution p that minimizes $\mathrm{KL}\left(o_{p^*,r}(i) \parallel p(i)\right)$, subject to the inequality constraint. Indeed, numerically minimizing this objective yields the constrained growth-optimal $p^*(i)$, as required.

The myopia is optimal. For contrast, consider a planner who minimizes systematic redistribution, measured by the divergence $\mathrm{KL}\left(o_{p,r}(i) \parallel p(i)\right)$ between the outcome and endowment distributions, subject to the same inequality constraint, but accounting for the effect of endowments on outcomes. This planner selects $\tilde{p}(1) \approx .197$, $\tilde{p}(2) \approx .257$, and $\tilde{p}(3) \approx .546$, achieving a slightly lower divergence but also a lower growth rate than p^*. ▲

4 Analysis

This section proves the main result and offers intuition.

4.1 Generalized Distributions

To facilitate the use of tools from information theory, we relabel the notation for the policy pairs. We retain $p(i)$ to denote the endowment distribution. However, instead of $r(i,\omega)$, we now write $p(\omega \mid i)$, a function $I \times \Omega \to \mathbb{R}_+$, for the return function. We treat $p(\omega \mid i)$, $i \in I$, as generalized conditional distributions (which need not be normalized to have unit mass) and introduce

$$p(i,\omega) := p(i)p(\omega \mid i),$$

which we again treat as a generalized joint distribution. Note that $p(i,\omega)$ is a sufficient summary statistic of the policy pair with respect to the aggregate growth rate (1), because this rate depends solely on this product. Therefore, we refer to $p(i,\omega)$ as the *policy*. The set $\tilde{\mathcal{P}}$ of feasible policies comprises those

is equivalent to requiring that the entropy $\mathrm{H}\left(p(i)\right)$ be at least 1.

$p(i,\omega) = p(i)p(\omega \mid i)$ that can be expressed as the product of a feasible policy pair $\big(p(i), p(\omega \mid i)\big) \in \mathcal{P}$.

We extend the standard definition of KL-divergence to take any well-normalized joint distribution $q(i,\omega)$ and any policy $p(i,\omega)$ as arguments (noting that it can now attain negative values). It is straightforward to verify that the usual chain rule continues to hold:

$$\mathrm{KL}\left(q(i,\omega) \parallel p(i,\omega)\right) = \mathrm{KL}\left(q(i) \parallel p(i)\right) + \sum_i q(i)\,\mathrm{KL}\left(q(\omega \mid i) \parallel p(\omega \mid i)\right),$$

with the summation over $i \in \mathrm{supp}(q(i)) \cap \mathrm{supp}(p(i))$. An analogous equality holds when the roles of i and ω are reversed.

4.2 Wealth Circulation

We analyze the economy's growth rate by tracking the circulation of wealth among individuals. To introduce wealth circulation informally, we imagine each dollar drawn from the initial stock of perfectly divisible aggregate wealth as founding a separate dynasty of subsequent wealth. In each period, every such dynasty is held by an individual i, randomly drawn from the endowment distribution $p(i)$, and in a random aggregate state ω of the economy, independently drawn from $q^0(\omega)$. Consequently, the dynasty's wealth is multiplied by the return $p(\omega \mid i)$. Without loss of generality, we assume that these dynasties remain separate and do not mix. That is, all wealth generated from the compounding returns of each dollar initially founding a dynasty circulates within the economy as separate, infinitesimal fractions of the aggregate wealth, each representing a distinct dynasty.

Consider a sufficiently large time horizon, allowing us to abstract from any issues related to divisibility. Up to this horizon, let the dynasty-specific distribution $q(i,\omega)$ represent the frequency of individuals and states occupied by the dynasty; we refer to $q(i,\omega)$ as the dynasty's pattern of circulation, or simply, its pattern. Since the states ω_t are aggregate, all patterns must satisfy the consistency condition $q(\omega) = q^0(\omega)$ in the long run. Accordingly, we formally define a *pattern* as any joint distribution $q(i,\omega) \in \Delta(I \times \Omega)$ that satisfies $q(\omega) = q^0(\omega)$. For each pattern, there are numerous distinct dynasties characterized by this pattern, each corresponding to a distinct permutation of $(i_t, \omega_t)_t$. Note that a dynasty's pattern serves as a sufficient statistic with respect to its wealth.

Over any finite horizon, dynasties differ in the frequencies with which they

have been in the hands of each individual i, as well as in how their allocations to various individuals correlate with those individuals' random productivities.[6] By chance, a small subset of lucky dynasties spend a disproportionate fraction of periods in the hands of individuals when these individuals are particularly productive. As it turns out, these fortunate dynasties generate effectively all of the economy's wealth.

For each pattern $q(i,\omega)$, we define its wealth as the aggregate wealth of all dynasties whose realized patterns correspond to $q(i,\omega)$. We show that the wealth of the pattern $q(i,\omega)$ under the policy $p(i,\omega)$ grows exponentially at a long-run rate given by

$$\begin{aligned} \gamma(p,q) &= \mathrm{E}_{q(i,\omega)} \ln p(\omega \mid i) - \mathrm{E}_{q^0(\omega)} \mathrm{KL}\left(q(i \mid \omega) \parallel p(i)\right) \\ &= -\mathrm{KL}\left(q(i,\omega) \parallel p(i,\omega)\right) + \text{const.} \end{aligned} \tag{7}$$

The first equality breaks down the growth rate of a pattern into two distinct components. First, the growth rate of each separate dynasty adhering to the pattern $q(i,\omega)$ is given by its long-run average log-return $\mathrm{E}_{q(i,\omega)} \ln p(\omega \mid i)$. Second, patterns that deviate from the endowment distribution $p(i)$ by assigning distinct state-contingent frequencies $q(i \mid \omega) \neq p(i)$ to individuals are atypical, leading to a gradual decline in the measure of dynasties following such patterns over time. The measure of these dynasties, and hence their initial wealth, diminishes at the rate given by $\mathrm{E}_{q^0(\omega)} \mathrm{KL}\left(q(i \mid \omega) \parallel p(i)\right)$. The greater the deviation from the endowment distribution, the faster the decline.[7] The wealth of all dynasties following a given pattern then grows at a rate given by the difference between the two rates.

The second equality in (7) follows from straightforward algebra. This rearrangement indicates that the closer the pattern $q(i,\omega)$ aligns with the policy $p(i,\omega)$, the greater the resulting growth rate.

As the wealth of each pattern grows exponentially, the wealth of the fastest-growing pattern eventually dwarfs that of all others. Consequently, the overall growth rate of the economy under a given policy is determined by the pattern that minimizes its divergence from the policy. The maximum feasible growth rate of the economy is thus achieved by the pattern-policy pair that minimizes

[6] A probability theorist would refer to atypical patterns of circulation as 'large deviations' from the generating distribution $p(i)q^0(\omega)$.

[7] This follows from Sanov's theorem, which implies that the chance of drawing a sample with an empirical frequency $q(i)$ from a distribution $p(i)$ decreases at the rate of $\mathrm{KL}\left(q(i) \parallel p(i)\right)$ as the sample expands.

their divergence. The following result, which is proven in the appendix, formalizes this intuition:

Proposition 3. *A policy $p^*(i,\omega)$ maximizes the economy's growth rate if and only if it solves*

$$\begin{aligned} \min_{q(i,\omega),p(i,\omega)} \quad & \mathrm{KL}\left(q(i,\omega) \parallel p(i,\omega)\right) \\ \text{s.t.} \quad & p(i,\omega) \in \tilde{\mathcal{P}} \\ & q(i,\omega) \in \Delta(I \times \Omega) \\ & q(\omega) = q^0(\omega), \end{aligned} \tag{8}$$

alongside a corresponding minimizer $q^*(i,\omega)$.

This optimization problem appears in a different context in the seminal machine learning paper by Kingma and Welling (2013), where it is introduced as a variational approach to maximum likelihood estimation. In the generative task studied by Kingma and Welling, a machine aims to learn a joint distribution $p(i,\omega) \in \tilde{\mathcal{P}}$ over the latent variable i and the observable variable ω that best matches a large sample of observed data, $(\omega_1, \ldots, \omega_n)$. In this context, the set $\tilde{\mathcal{P}} \subset \Delta(I \times \Omega)$ represents the collection of joint distributions that the machine considers and is capable of expressing. The asymptotic maximum likelihood estimator is given by the optimizer $p^*(i,\omega)$ from Problem (8), where $q^0(\omega)$ corresponds to the true data generating distribution. The accompanying optimizer $q^*(i,\omega)$ characterizes Bayesian updates derived from the estimated p^*: $q^*(i \mid \omega) = p^*(i \mid \omega)$.[8]

The equivalence between growth maximization and maximum likelihood estimation may seem surprising. The connection arises because maximum likelihood estimation can be interpreted as a competition among growth processes. Specifically, the likelihood of each considered hypothesis $p(i,\omega)$ evolves as a stochastic growth process, with a (negative) stochastic growth rate $\ln p(\omega_t) = \ln \sum_i p(i,\omega_t)$ in each period t. Consequently, as the sample size increases, the likelihood of the hypothesis grows at the long-run rate $\mathrm{E}_{q^0(\omega)} \ln \sum_i p(i,\omega)$. Asymptotically, maximum likelihood estimation selects the hypothesis with the highest growth rate, analogous to a planner selecting the growth-maximizing policy.

[8] We explore this cognitive interpretation of Problem (8) in Samuelson and Steiner (2024). In a similar vein, Gabaix and Graeber (2024) represent cognition in a macroeconomically inspired model.

Our proof of Proposition 3 is distinct from the standard arguments in the variational inference literature that are based on convex analysis. Instead, we focus on the dynamic aspect of stochastic growth and draw on intuitions from large deviations theory.

4.3 Growth-Maximizing Pattern and Policy

To prove the main result in Proposition 2, we establish an identity between two conceptually distinct objects: the outcome distribution and the growth-maximizing pattern of circulation.

We extend the definition of the outcome distribution by introducing a *joint* outcome distribution, induced by the policy $p(i,\omega)$, and defined as

$$o_p(i,\omega) := q^0(\omega) o_p(i \mid \omega), \tag{9}$$

where $o_p(i \mid \omega)$ is the state-contingent outcome distribution defined in (2). That is, $o_p(i,\omega)$ is the probability that a dollar, sampled from the aggregate wealth at the end of a random period, is held by individual i in state ω. The outcome distribution, as defined in (3), is then given by the marginalization $o_p(i) = \sum_\omega o_p(i,\omega)$.

Given a policy $p(i,\omega)$, let

$$q_p^*(i,\omega) \in \underset{q(i,\omega)\in\Delta(I\times\Omega)}{\arg\min} \quad \mathrm{KL}\left(q(i,\omega) \parallel p(i,\omega)\right) \tag{10}$$
$$\text{s.t.} \quad q(\omega) = q^0(\omega)$$

denote the fastest-growing pattern.

Lemma 1. *For any given policy $p(i,\omega)$, the fastest-growing pattern $q_p^*(i,\omega)$ is identical to the joint outcome distribution $o_p(i,\omega)$.*

Proof. Recall that the policy $p(i,\omega)$ is a generalized (non-normalized) joint distribution, on which we perform standard probability-theory operations. We formally define the conditional policy as $p(i \mid \omega) = \frac{p(i,\omega)}{\sum_i p(i,\omega)}$. Observe that $p(i \mid \omega) = o_p(i \mid \omega)$, where $o_p(i \mid \omega)$ is the state-contingent outcome distribution, as previously defined in (2).

Applying the chain rule and utilizing the constraint $q(\omega) = q^0(\omega)$, we express

the objective from (10) as

$$\mathrm{KL}\left(q(i,\omega) \parallel p(i,\omega)\right) = \mathrm{KL}\left(q^0(\omega) \parallel p(\omega)\right) + \sum_{\omega} q^0(\omega)\,\mathrm{KL}\left(q(i \mid \omega) \parallel p(i \mid \omega)\right).$$

Minimizing with respect to $q(i \mid \omega)$ yields $q_p^*(i \mid \omega) = p(i \mid \omega)$ for each ω. The lemma then follows from the definition of $o_p(i,\omega)$ in (9):

$$o_p(i,\omega) = q^0(\omega) o_p(i \mid \omega) = q^0(\omega) p(i \mid \omega) = q^0(\omega) q_p^*(i \mid \omega) = q_p^*(i,\omega).$$

□

Proposition 3 and Lemma 1 jointly establish that the optimal policy and the associated outcome distribution are two optimized, decoupled controls. The proof of the main result is presented below.

Proof of Proposition 2. Let $p^*(i,\omega)$ be the growth-maximizing policy and $q_{p^*}^*(i,\omega)$ the associated fastest-growing pattern. By Proposition 3, the policy p^* maximizes the growth rate of this pattern:

$$\begin{aligned} p^*(i,\omega) &\in \operatorname*{arg\,min}_{p(i,\omega)\in\tilde{\mathcal{P}}} \mathrm{KL}\left(q_{p^*}^*(i,\omega) \parallel p(i,\omega)\right) \\ &= \operatorname*{arg\,min}_{p(i,\omega)\in\tilde{\mathcal{P}}} \mathrm{KL}\left(o_{p^*}(i,\omega) \parallel p(i,\omega)\right), \end{aligned}$$

where the equality follows from Lemma 1.

Applying the chain rule, we can express the objective as:

$$\mathrm{KL}\left(o_{p^*}(i) \parallel p(i)\right) + \sum_{i\in\mathrm{supp}(o_{p^*}(i))\cap\mathrm{supp}(p(i))} o_{p^*}(i)\,\mathrm{KL}\left(o_{p^*}(\omega \mid i) \parallel p(\omega \mid i)\right).$$

By setting the returns $p(\omega \mid i)$ to the optimized returns $p^*(\omega \mid i)$ and restricting the optimization of the last inline expression to the control of the endowment distribution $p(i)$ within $\mathcal{E}^*$, we obtain the statement in (4). □

4.4 Additional Examples

We now present three examples that focus on characterizing the optimal policy. The first example illustrates how Proposition 3 characterizes the optimal policy and the resulting pattern of wealth circulation. The analysis leverages the specific structure of the example, resulting in a simple pattern. The online

appendix extends the analytical approach from this example to more general settings. Specifically, it shows that the planner's problem—when the return function is fixed and the endowment distribution is unconstrained—is mathematically equivalent to the rational inattention problem with entropic information costs, as introduced by Matějka and McKay (2015).

Example 3 (Analysis)**.** We now revisit the unconstrained setting from Example 2 to illustrate the application of Proposition 3. First, we fix a possibly suboptimal endowment distribution $p(i)$ and compute the corresponding growth-maximizing pattern $q_p^*(i,\omega)$. Since $q(\omega) = q^0(\omega)$ is fixed, we only need to find the conditional distributions $q_p^*(i \mid \omega)$ for each $\omega \in \Omega$. Recall that in states $\omega \in I$, all individuals except $i = \omega$ become bankrupt, with $r(j,\omega) = 0$ for $j \neq \omega$. Hence, $q_p^*(i \mid \omega) = \mathbb{1}_{i=\omega}$ for all $\omega \in I$. In other words, the growth-maximizing pattern avoids bankruptcy in any states $\omega \in I$ where bankruptcy could arise, as failing to do so would result in a growth rate of negative infinity. In the state $\omega = g$, the growth-maximizing pattern $q^*(i \mid \omega) = p(i)$ aligns with the endowment distribution, as all individuals receive the same return in this state, making any deviation from $p(i)$ unbeneficial. Therefore, given the endowment distribution $p(i)$, the outcome distribution $o_p(i)$ is given by

$$o_p(i) = \mathrm{E}_{q^0(\omega)}\, q_p^*(i \mid \omega) = q^0(i) + q^0(g)p(i).$$

Proposition 3 requires the optimal policy to minimize divergence from the pattern $q_{p^*}^*$. Specifically, when the endowment distribution is unconstrained, Corollary 1 applies, yielding $p^*(i) = o_{p^*}(i) \equiv q_{p^*}^*(i)$, which implies that

$$p^*(i) = q^0(i) + q^0(g)p^*(i),$$

thereby delivering the characterization given in (6). ▲

We now present an extension where the justifiability of redistribution depends on past payoff states, thereby introducing history dependence to the policy constraint.

Example 4 (History-Dependent Constraints)**.** We let the set $\mathcal{P}_{\omega_{t-1}}$, representing the policies $p(i,\omega)$ feasible in period t, depend on the preceding state ω_{t-1}. For instance, the planner may need to select an endowment distribution $p(i)$ at the beginning of period t that maintains the ranking of individuals' returns from the end of period $t-1$, consistent with the meritocracy notion. For illustration,

consider an environment with sets $I = \{1, 2, 3\}$ and $\Omega = \{1, 2, 3, g\}$, as well as returns $p(\omega \mid i) \equiv r(i, \omega)$ from Example 2. Given that returns are fixed here, we abuse notation by associating the set $\mathcal{P}_{\omega_{t-1}}$ of feasible policies with the set of feasible endowment distributions. Recall that in states $\omega_{t-1} \in \{1, 2, 3\}$, all individuals except $i = \omega_{t-1}$ went bankrupt at $t-1$. Thus, following these states, the feasible set is

$$\mathcal{P}_{\omega_{t-1}} = \{p(i) \in \Delta(I) : p(j) \geq p(k) \text{ for } j = \omega_{t-1} \text{ and all } k \neq \omega_{t-1}\},$$

and the set $\mathcal{P}_g = \Delta(I)$ permits any endowment distribution since all individuals have received the same return when $\omega_{t-1} = g$.

The following separability property extends our results to environments with state-dependent policy constraints. Define $\tilde{\gamma}(p) = \mathrm{E}_{q^0(\omega)} \ln \sum_i p(i, \omega)$ as the growth rate achieved by applying a policy $p(i, \omega)$ in all rounds. Since the states are serially independent, a conditional policy $p_{\omega_{t-1}}$ achieves a growth rate of $\mathrm{E}_{q^0(\omega)} \tilde{\gamma}(p_\omega)$. Consequently, a planner facing state-dependent feasibility sets $\mathcal{P}_{\omega_{t-1}}$ selects, in each period t, the policy that is optimal in the main setting with $\tilde{\mathcal{P}} = \tilde{\mathcal{P}}_{\omega_{t-1}}$.

For the specification from Example 2, the planner is unconstrained following $\omega_{t-1} = g$ and thus chooses the unconstrained optimal endowment distribution $(.1, .2, .7)$. After state $\omega_{t-1} = 1$, the planner must allocate the largest share to individual 1, resulting in the constrained optimal endowment distribution $(.4, .2, .4)$. After state $\omega_{t-1} = 2$, individual 2 must receive the largest share, yielding the constrained optimal endowment distribution $(.1, .45, .45)$. Finally, following state $\omega_{t-1} = 3$, the constraint is not binding and thus the planner chooses $(.1, .2, .7)$. ▲

Although our basic model assumes that productivities are serially independent and the planner can periodically redistribute the entire aggregate wealth, applications may require incorporating serially correlated productivities and more realistic redistribution schemes. The following example sketches an extension of our framework, characterizing how the optimal flat tax rate depends on productivity inertia.[9]

[9] While a class of settings from this paper is homeomorphic to the static rational inattention problem from Matějka and McKay (2015) (see the online appendix), we conjecture that a natural extension of this class, which generalizes Example 5, is homeomorphic to the dynamic rational inattention problem from Steiner et al. (2017).

Example 5 (Inertia)**.** We now extend the model to incorporate serially correlated productivities and policies akin to flat tax rates. The payoff state now evolves in a stationary Markov chain $q^0(\omega_t \mid \omega_{t-1})$. Rather than selecting an endowment distribution $p(i)$, the planner now chooses a redistribution rule, represented by a Markov chain $p(i \mid j)$, which specifies the share of wealth from each individual j that is redistributed to each individual i at the end of each period. For example, a flat tax rate τ with uniform redistribution of the tax revenue results in $p(i \mid j) = \tau/|I| + (1-\tau)\mathbb{1}_{i=j}$. The redistribution rule and the return function jointly induce a stochastic evolution of wealth levels:

$$w_{i,t+1} = \sum_{j \in I} p(i \mid j) w_{j,t} r(j, \omega_t),$$

where the planner's objective is to maximize the long-run growth rate.

We recall the concept of wealth dynasties from Section 4.2 and let them circulate across individuals according to the chain $p(i_t \mid i_{t-1})$, which replaces the iid draws from the endowment distribution. Since the dynasties circulate across individuals independently of the state evolution, the pair (i_t, ω_t), which tracks the individual and the state, evolves according to the Markov chain

$$m(i_{t+1}, \omega_{t+1} \mid i_t, \omega_t) = p(i_{t+1} \mid i_t) q^0(\omega_{t+1} \mid \omega_t).$$

Again, for a finite time horizon, some dynasties, by chance, follow a pattern described by a chain $q(i_{t+1}, \omega_{t+1} \mid i_t, \omega_t)$ that differs from m. The aggregate wealth of dynasties following the pattern q grows at the rate

$$\gamma(q, p) = \mathrm{E}_{\overline{q}} \ln r(i_t, \omega_t) - \mathrm{E}_{\overline{q}} \mathrm{KL}\left(q(i_{t+1}, \omega_{t+1} \mid i_t, \omega_t) \| m(i_{t+1}, \omega_{t+1} \mid i_t, \omega_t)\right),$$

where $\overline{q}(i_t, \omega_t)$ is the stationary distribution of the chain q. The first term on the right represents the growth rate of each dynasty along the pattern q. The second term represents the divergence rate between chains q and m, capturing the rate at which the measure of dynasties adhering to pattern q declines over time. For a given policy p, the economy's growth rate is given by $\gamma(q, p)$, maximized over all q that match $q^0(\omega_t \mid \omega_{t-1})$ on the margin. Thus, the planner's problem involves jointly maximizing $\gamma(q, p)$ over feasible patterns of circulation q and policies p.

For brevity, consider a setting with two individuals and two states, $I = \Omega = \{1, 2\}$, with returns given by $r(i, \omega) = 2 \times \mathbb{1}_{i=\omega}$. In this setting, each individual

i doubles her wealth in her productive state $\omega = i$ and becomes bankrupt when $\omega \neq i$, as in Example 1. In this case, the state evolves according to the Markov chain $q^0(\omega_t \mid \omega_{t-1}) = \sigma$ when $\omega_t = \omega_{t-1}$, where $\sigma \in [1/2, 1)$ represents the degree of state inertia. The planner selects a tax rate $\tau \in [0, 1]$ and redistributes the tax revenue evenly across both individuals, inducing the redistribution rule $p(i \mid i) = \nu$ and $p(i \mid j) = 1 - \nu$ for $i \neq j$, where $\nu = 1 - \tau/2$.

In this setting, the optimal pattern q^* is straightforward: a dynasty must be owned by the productive individual each period to survive, which implies $i_t = \omega_t$ almost surely under the optimal q^*. Then,

$$\mathrm{KL}\left(q^*\big(i_{t+1}, \omega_{t+1} \mid i_t, \omega_t\big) \| m\big(i_{t+1}, \omega_{t+1} \mid i_t, \omega_t\big)\right)$$

takes the value $-\ln \nu$ (in case $i_{t+1} = \omega_{t+1} = \omega_t = i_t$) or $-\ln(1 - \nu)$ (in case $i_{t+1} = \omega_{t+1} \neq \omega_t = i_t$), because q^* enforces $i_{t+1} = \omega_{t+1}$ with certainty, while under m, $i_{t+1} = \omega_{t+1}$ with probabilities ν and $1 - \nu$, respectively, in the two scenarios. Given the probabilities σ and $1-\sigma$ for these two cases, a redistribution rule $p_\nu(i \mid j)$ yields a growth rate of

$$\gamma(q^*, p_\nu) = \ln 2 + \sigma \ln \nu + (1 - \sigma) \ln(1 - \nu).$$

Maximizing over ν results in the optimal value $\nu^* = \sigma$, which corresponds to a tax rate $\tau^* = 2 - 2\sigma$. When $\sigma = 1/2$, indicating that states are independent, the planner optimally confiscates all wealth each period and redistributes it uniformly, aligning with the baseline analysis for this symmetric setting. However, for $\sigma > 1/2$, the planner takes advantage of productivity inertia by redistributing only a smaller fraction of wealth, anticipating continued high productivity from currently wealthy individuals. ▲

5 Summary

We have deliberately set aside the usual factors that mediate the impact of redistribution on economic growth, to emphasize a relatively unexplored factor stemming from the stochastic nature of growth. Optimizing growth necessitates regular redistribution. This optimization is forward-looking, with the optimal policy regularly redistributing incidental wealth gains and losses. Nevertheless, our main result underscores the adverse impact of *systematic* redistribution on growth. The growth-optimal policy myopically minimizes systematic redistri-

bution, treating economic outcomes as independent of the policy. Specifically, the growth-optimal policy does not globally minimize systematic redistribution. This counterintuitive myopia arises from the fixed-point property of the optimal policy, which we derive through a decoupling approach that has an intuitive economic interpretation. Example 5 outlines an extension that incorporates realistic redistribution policies and accounts for inertial environments.

A Proof of Proposition 3

A policy $p(i,\omega)$ generates a growth rate of aggregate wealth in each state ω given by:

$$\begin{aligned}\ln \sum_i p(i,\omega) &= \frac{1}{t} \ln \left(\sum_i p(i,\omega) \right)^t \\ &= \frac{1}{t} \ln \sum_{\mathbb{i} \in I^t} \prod_i p(i,\omega)^{q_{\mathbb{i}}(i) t},\end{aligned}$$

where $t \in \mathbb{N}$ is an arbitrary time horizon, $\mathbb{i} \in I^t$ is a sequence $(i_1, \ldots, i_t)$ and $q_{\mathbb{i}}(i) = \frac{1}{t} \sum_{\tau=1}^{t} \mathbb{1}_{i_\tau = i}$ is the empirical distribution of the sequence $\mathbb{i}$. Since all sequences $\mathbb{i}$ with the same empirical distribution generate the same value of the inline summand, we obtain the growth rate in state ω:

$$\ln \sum_i p(i,\omega) = \frac{1}{t} \ln \left(\sum_{q(i|\omega) \in \Delta_t} n_t\big(q(i \mid \omega)\big) \prod_i p(i,\omega)^{q(i|\omega) t} \right), \tag{11}$$

where $\Delta_t \subset \Delta(I)$ is the set of the empirical distributions that can be generated by sequences $\mathbb{i}$ of length t, and $n_t : \Delta_t \to \mathbb{N}$ maps each q to the number $n_t(q)$ of sequences of length t that generate such an empirical distribution. Using the entropy $\mathrm{H}(q)$, Theorem 11.1.3 in Cover and Thomas (2006) approximates this function as follows:

$$\frac{1}{(t+1)^{|I|}} \exp[t \times \mathrm{H}(q)] \leq n_t(q) \leq \exp[t \times \mathrm{H}(q)], \tag{12}$$

for all $q \in \Delta_t$.

Substituting these bounds into (11) yields:

$$\begin{aligned}&\frac{1}{t} \ln \sum_{q(i|\omega) \in \Delta_t} \exp\left[t \times \left(\mathrm{E}_{q(i|\omega)} \ln p(i,\omega) + \mathrm{H}\left(q(i \mid \omega)\right) \right) \right] - \frac{|I| \ln(t+1)}{t} \\ \leq\ & \ln \sum_i p(i,\omega) \qquad (13) \\ \leq\ & \frac{1}{t} \ln \sum_{q(i|\omega) \in \Delta_t} \exp\left[t \times \left(\mathrm{E}_{q(i|\omega)} \ln p(i,\omega) + \mathrm{H}\left(q(i \mid \omega)\right) \right) \right],\end{aligned}$$

which further simplifies to:

$$\begin{aligned} & \max_{q(i|\omega)\in\Delta_t} \left\{ \mathrm{E}_{q(i|\omega)} \ln p(i,\omega) + \mathrm{H}\left(q(i \mid \omega)\right) \right\} - \frac{|I| \ln(t+1)}{t} \\ \leq \quad & \ln \sum_i p(i,\omega) \\ \leq \quad & \max_{q(i|\omega)\in\Delta_t} \left\{ \mathrm{E}_{q(i|\omega)} \ln p(i,\omega) + \mathrm{H}\left(q(i \mid \omega)\right) \right\} + \frac{|I| \ln(t+1)}{t}. \end{aligned}$$

For the lower bound, we replaced the sum in lower bound from (13) with its maximal summand. For the upper bound, we replaced each summand in the upper bound from (13) by the maximal summand and observed that there are at most $(t+1)^{|I|}$ summands. This holds because the distributions $q(i \mid \omega) \in \Delta_t$ takes values in $\{\frac{0}{t}, \frac{1}{t}, \ldots, \frac{t}{t}\}$ for each $i \in I$; hence $|\Delta_t| \leq (t+1)^{|I|}$ (Cover and Thomas, 2006, Theorem 1.1.1).

Since $\frac{\ln(t+1)}{t}$ vanishes as t diverges, taking the limit $t \to \infty$ yields:[10]

$$\ln \sum_i p(i,\omega) = \max_{q(i|\omega)\in\Delta(I)} \left\{ \mathrm{E}_{q(i|\omega)} \ln p(i,\omega) + \mathrm{H}\left(q(i \mid \omega)\right) \right\}. \tag{14}$$

Taking the expectation over $\omega \sim q^0(\omega)$ and straightforward algebraic steps yield the characterization of the aggregate growth rate induced by a policy $p(i,\omega)$:

$$\begin{aligned} \mathrm{E}_{q^0(\omega)} \ln \sum_i p(i,\omega) = \quad & \max\nolimits_{q(i,\omega)\in\Delta(I\times\Omega)} \quad \gamma\big(q(i,\omega) \parallel p(i,\omega)\big) \\ & \text{s.t. } q(\omega) = q^0(\omega), \end{aligned}$$

where the growth rate $\gamma\big(q(i,\omega) \parallel p(i,\omega)\big)$ is defined in (7). Finally, optimization over the policy $p(i,\omega)$ yields the proposition.

References

Aghion, P., E. Caroli, and C. García-Peñalosa (1999). Inequality and economic growth: The perspective of the new growth theories. *Journal of Economic Literature 37*(4), 1615–1660.

Campbell, J. Y. and R. Sigalov (2022). Portfolio choice with sustainable spend-

[10] Equation (14) follows from the Donsker-Varadhan Lemma. We derive it from first principles to provide deeper intuition regarding wealth circulation.

ing: A model of reaching for yield. *Journal of Financial Economics 143*(1), 188–206.

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

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

Farhi, E. and I. Werning (2013). Insurance and taxation over the life cycle. *Review of Economic Studies 80*(2), 596–635.

Forbes, K. J. (2000). A reassessment of the relationship between inequality and growth. *American Economic Review 90*(4), 869–887.

Gabaix, X. and T. Graeber (2024). The complexity of economic decisions. Working paper 24-049, Harvard Business School.

Golosov, M., M. Troshkin, and A. Tsyvinski (2016). Redistribution and social insurance. *American Economic Review 106*(2), 359–386.

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.

Kapička, M. (2013). Efficient allocations in dynamic private information economies with persistent shocks: A first-order approach. *Review of Economic Studies 80*(3), 1027–1054.

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

Kingma, D. P. and M. Welling (2013). Auto-encoding variational Bayes. arxiv preprint arxiv:1312.6114, Cornell University.

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.

Mdingi, K. and S.-Y. Ho (2021). Literature review on income inequality and economic growth. *MethodsX 8*, 101402.

Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica: Journal of the Econometric Society 41*(5), 867–887.

Merton, R. C. and P. A. Samuelson (1974). Fallacy of the log-normal approximation to optimal portfolio decision-making over many periods. *Journal of financial economics 1*(1), 67–94.

Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *The Review of Rconomic Studies 38*(2), 175–208.

Mirrlees, J. A. (1986). The theory of optimal taxation. In K. J. Arrow and M. D. Intriligator (Eds.), *Handbook of Mathematical Economics*, Volume 3, pp. 1197–1249. Elsevier.

Robson, A., L. Samuelson, and J. Steiner (2023). Decision theory and stochastic growth. *American Economic Review: Insights 5*(3), 357–376.

Samuelson, L. and J. Steiner (2023). Growth and likelihood. CEPR Discussion Paper 18339, Centre for Economic Policy Research.

Samuelson, L. and J. Steiner (2024). Constrained data-fitters. Working paper, Yale University and University of Zurich, CERGE-EI, and CTS.

Stantcheva, S. (2020). Dynamic taxation. *Annual Review of Economics 12*, 801–831.

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