No. 24

The Two Growth Rates of the Economy

Alexander Adamou Yonatan Berman Ole Peters

September 2020

The Two Growth Rates of the Economy

Alexander Adamou*

Yonatan Berman[†]

Ole Peters^{‡ §}

September 22, 2020

Abstract

Economic growth is measured as the rate of relative change in gross domestic product (GDP) per capita. Yet, when incomes follow random multiplicative growth, the ensemble-average (GDP per capita) growth rate is higher than the time-average growth rate achieved by each individual in the long run. This mathematical fact is the starting point of ergodicity economics. Using the atypically high ensemble-average growth rate as the principal growth measure creates an incomplete picture. Policymaking would be better informed by reporting both ensemble-average and time-average growth rates. We analyse rigorously these growth rates and describe their evolution in the United States and France over the last fifty years. The difference between the two growth rates gives rise to a natural measure of income inequality, equal to the mean logarithmic deviation. Despite being estimated as the average of individual income growth rates, the time-average growth rate is independent of income mobility.

Keywords: economic growth, inequality, ergodicity economics

*London Mathematical Laboratory, a.adamou@lml.org.uk

[†]London Mathematical Laboratory and Stone Center on Socio-Economic Inequality, y.berman@lml.org.uk

[‡]London Mathematical Laboratory and Santa Fe Institute, o.peters@lml.org.uk

[§]The code to generate all figures is available at <https://bit.ly/DDPcodes>.

1 Introduction

Ergodicity economics is a perspective that emerges when we ask explicitly whether stochastic processes used as economic models are ergodic (Peters, 2019). If a stochastic process is ergodic, its ensemble average¹ is identical to its time average.² An important counterexample is multiplicative growth with random shocks, a widely used model of growth in economics and elsewhere. Here, the ensemble average grows exponentially at a deterministic rate, which we call the *ensemble-average growth rate*. In the long-time limit, any individual trajectory also grows exponentially at a deterministic rate, which we label the *time-average growth rate*. Crucially, these two rates differ: the ensemble-average growth rate is greater than the time-average growth rate.

In this paper we explore what ergodicity economics tells us about measures of national economic performance. We find that the two growth rates described above correspond to two types of economic growth that are discussed in various guises in the literature.

The ensemble-average growth rate is closely related to the most quoted measure of national economic performance, namely to the growth rate of gross domestic product (GDP) per capita. GDP is the total income generated by the economic activity of a national population, so the growth rate of GDP per capita is simply the growth rate of average income. For large populations, this coincides with the ensemble-average growth rate. The GDP per capita growth rate is a so-called *plutocratic* growth measure in the following sense: it is approximately a weighted average of individual income growth rates, where the weight assigned to each individual is proportional to their income.³

The time-average growth rate is representative of individual growth, being the growth rate realised with certainty by each income in the long-time limit. Unlike the ensemble-average growth rate, it is rarely measured or stated. The time-average growth rate can be estimated as the average of the individual income growth rates. This makes it a *democratic* growth measure, because each person is assigned an equal weight in the average. To keep the symmetry between ensemble averages and time averages, and to provide a handy nomenclature, we introduce the *democratic domestic product* (DDP) per capita as the quantity that grows at the time-average growth rate.⁴

The ensemble-average (GDP per capita) growth rate may measure changes in a nation’s total productive capacity, but it is an unsuitable measure of the typical economic experience of the nation’s citizens. Conversely, the time-average (DDP per capita) growth rate may reflect how citizens expe-

¹The ensemble average of a stochastic process $Y(t)$ is the large- N limit of the average value over N realisations. This limit is also called the expected value, or expectation value, or first moment, or mean, denoted $\langle Y(t) \rangle \equiv \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N y_i(t)$.

²The time average of a stochastic process $Y(t)$ is the long-time limit of the average value along a trajectory $y(t)$, denoted $\bar{Y} \equiv \lim_{T \rightarrow \infty} \frac{1}{T} \int_t^T y(s) ds$.

³This nomenclature is standard. For example, Milanovic (2005) uses the adjective “plutocratic” to describe statistics with this property. We make the term mathematically precise in Section 2.

⁴Existing names for the same concept are the “people’s growth rate” (Milanovic, 2005; Saez and Zucman, 2019) and the “population-weighted growth rate” (Milanovic, 2005). As far as we know, the distinction between the two growth rates was first drawn by Chenery et al. (1974), when discussing different ways to weight quintile average income growth rates to obtain a single indicator of economic performance.

rience the economy, but it is not a good measure of what’s happening to total productive capacity. The tension between these two growth rates gives rise to a natural measure of inequality (Adamou and Peters, 2016). If the ensemble-average growth rate exceeds the time-average growth rate, *i.e.* if GDP per capita grows faster than DDP per capita, then inequality is increasing. If the converse is true, then inequality is decreasing. This suggests defining inequality simply as the quantity which grows at the difference between the two rates, ensemble-average minus time-average. We call this inequality measure the *ergodicity gap*. Under multiplicative growth, it corresponds to an established inequality metric: the *mean logarithmic deviation* (MLD), also known as the second index of Theil (1967). At perfect equality, the ergodicity gap vanishes, and GDP per capita and DDP per capita are equal.

Computing the time-average (DDP per capita) growth rate as the average of individual growth rates requires, *prima facie*, hard-to-obtain panel data. However, we show that this growth rate and, by extension, the ergodicity gap (MLD) depend only on the initial and terminal income distributions. They do not depend on the reordering of individuals that occurs in-between, *i.e.* they are independent of rank mobility. This has the practical benefit of allowing their estimation using cross-sectional income data.

Our contribution is twofold. First, we relate the ensemble-average and time-average growth rates of ergodicity economics to the two types of growth measure discussed in the literature. We use them to define a fundamental measure of inequality. To better inform policymaking, we propose that national statistical authorities compile and report all three measures. Second, we illustrate potential uses for these statistics by studying their strikingly different evolutions in the United States and France over the past fifty years.

The existence of the two growth measures raises the possibility of political disconnect. To a central planner, say the government in Washington, DC, the most visible figure is GDP. It is an accurate measure of the nation’s productive capacity and central planners are right to be interested in it. But the situation may arise, and data suggest it has arisen, where ensemble-average growth is consistently faster than time-average growth. Central planners may see ‘the economy’ doing well while, on the ground, individuals are doing less well. If allowed to persist, this situation can lead to a disconnect between the central government and the typical citizen. Citizens may feel left behind, excluded from economic prosperity, or even lose faith in their institutions and press, who promulgate statistics that paint a picture of prosperity unrecognisable to most.

2 The two growth rates of the economy

Ergodicity economics places great emphasis on the functional form of the growth rate (Peters and Adamou, 2018; Peters, 2019). In simple terms, a growth rate extracts the scale parameter of time in a growth function. Different types of growth correspond to different growth functions and, therefore, to different growth rates. For example, a quantity that grows linearly in time according

to the growth function $y = rt$, has its growth rate extracted by the operation $dy/dt = r$; while under an exponential growth function, $y = \exp(rt)$, the growth rate is extracted by $d \ln y/dt = r$. In both cases, a constant is obtained from a time-varying function. At a deeper level, and generalising to growth with fluctuations, the role of the growth rate is to extract a stationary mathematical object – specifically an ergodic observable – from a model of a growing observable. This is achieved by identifying an ergodicity transformation, which maps the observable of interest (here, income) to an observable whose rate of change is ergodic. Finding this transformation is often a powerful step in ergodicity economics, as we shall see in Section 2.3. In the model of income we use here, we show that the transformation is simple (the logarithm) and we do not dwell on its deeper meaning.

2.1 Random multiplicative growth

In random multiplicative growth, income obeys Gibrat’s law of proportionate effect (Gibrat, 1931), *i.e.* incomes are multiplied by random variables at successive time steps. Variants of this model have been studied since the 1930s. We recommend Aitchison and Brown (1957, Ch. 11) for a review of early research and Gabaix et al. (2016) for up-to-date coverage. We consider random multiplicative growth a null model of income, which captures two dominant effects: that incomes fluctuate; and that they do so in proportion to themselves.⁵

We consider a model population of N people, labelled $i = 1 \dots N$, whose incomes, $x_i(t)$, are functions of time. Income is measured in units of currency per unit time, *e.g.* dollars per year, and is always positive. Income at one observation time is the income at the previous observation time multiplied by a random variable,

$$x(t + \delta t) = x(t)(1 + \epsilon), \quad (2.1)$$

where δt is the time between observations and ϵ is the random variable. (2.1) defines a stochastic process which generates trajectories of income,

$$x(t + T\delta t) = x(t) \prod_{\tau=1}^T (1 + \epsilon_{\tau}), \quad (2.2)$$

where ϵ_{τ} is the realisation of ϵ corresponding to the τ^{th} observation. We assume ϵ is stationary, *i.e.* that its distribution does not change in time, and that realisations at different times τ are independent. Note that this makes the stochastic process ϵ_{τ} ergodic: the time average and ensemble average of ϵ_{τ} are trivially identical.

The rate of change of income over the single time step in (2.1) is $\delta x/\delta t = x(t)\epsilon/\delta t$. This is not

⁵Changes in income, such pay rises and cuts, are typically expressed as percentages per year, as are changes in regular outgoings, like housing and transport costs. We do not suggest this captures all aspects of reality that determine people’s incomes, just the most important ones. We note that predictions about the distribution of income under random multiplicative growth can be tested empirically (Guvenen et al., 2017) and that extensions to general growth processes are available (Peters and Adamou, 2018).

ergodic because it depends on $x(t)$.⁶ Therefore, we introduce an ergodicity transformation, *i.e.* a function of $x_i(t)$ whose rate of change is ergodic. We do this by requiring the change in the function to be independent of $x(t)$. Taking the logarithm of (2.1) turns the product of income and income multiplier into a sum, which we rearrange as

$$\ln x(t + \delta t) - \ln x(t) = \ln(1 + \epsilon). \quad (2.3)$$

This is the desired result: the left-hand side is the additive change of transformed income; the right-hand side is a stationary random variable with no x -dependence. The rate of change, $g = \delta \ln x(t) / \delta t$, is now ergodic, $\langle g \rangle = \bar{g}$, and we call it the ergodic growth rate.⁷

Since this analysis does not depend on the length of time, δt , between observations, we write the growth rate of individual income over a general time period from time t to time $t + \Delta t$ as

$$g_i = \frac{\Delta \ln x_i(t)}{\Delta t}, \quad (2.6)$$

where $\Delta \ln x_i(t) = \ln x_i(t + \Delta t) - \ln x_i(t)$ is the change in the logarithm of the i^{th} income.

2.2 Ensemble-average growth rate

We denote by $\langle x(t) \rangle_N$ the population-average income,

$$\langle x(t) \rangle_N \equiv \frac{1}{N} \sum_{i=1}^N x_i(t), \quad (2.7)$$

which we can compute from observations of real incomes (or estimate from partial data). In a model of income dynamics, in which we can generate as many model incomes as we like, the ensemble-average income, $\langle x(t) \rangle$, is the $N \rightarrow \infty$ limit of the population-average income,

$$\langle x(t) \rangle \equiv \lim_{N \rightarrow \infty} \langle x(t) \rangle_N. \quad (2.8)$$

In practice, the national populations we study are large enough for differences between modelled

⁶The ensemble-average is $\langle \delta x / \delta t \rangle = \langle x(t) \rangle \langle \epsilon \rangle / \delta t$. The time-average is the average value of $\delta x / \delta t$ along an income trajectory, which will diverge or be zero depending on whether $x(t)$ grows or decays.

⁷The ensemble-average is

$$\langle g \rangle = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \frac{\delta \ln x_i(t)}{\delta t} = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \frac{\ln(1 + \epsilon_i)}{\delta t}, \quad (2.4)$$

where i labels different trajectories of the income process, (2.2), at a single point in time. The time-average is

$$\bar{g} = \lim_{T \delta t \rightarrow \infty} \frac{1}{T \delta t} \sum_{\tau=0}^{T-1} \frac{\delta \ln x(t + \tau \delta t)}{\delta t} \delta t = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{\tau=0}^{T-1} \frac{\ln(1 + \epsilon_\tau)}{\delta t}, \quad (2.5)$$

where τ labels different times of a single trajectory of the income process, (2.2). Since ϵ is stationary and its realisations are independent, the two averages have the same limiting value.

values of $\langle x(t) \rangle$ and $\langle x(t) \rangle_N$ to be negligible (see Appendix A for details). Therefore, we shall refer to both quantities as the *ensemble-average income*, using the subscript N to distinguish the finite-population statistic from that of the infinite ensemble.

The ensemble-average income, being the total personal income of a population divided by the number of people, resembles the gross domestic product (GDP) per capita. We shall make this association explicit throughout the text.

The ensemble-average income (GDP per capita) grows at the rate

$$g_{\langle} \equiv \frac{\Delta \ln \langle x(t) \rangle}{\Delta t}. \quad (2.9)$$

or, for a finite population, at the rate

$$g_{\langle}_N \equiv \frac{\Delta \ln \langle x(t) \rangle_N}{\Delta t}. \quad (2.10)$$

We call this the *ensemble-average growth rate* of income (GDP per capita growth rate).⁸

If changes in income are small, then we can approximate the ensemble-average (GDP per capita) growth rate as a weighted average of individual growth rates

$$g_{\langle}_N \approx \sum_{i=1}^N w_i g_i, \quad (2.15)$$

where each person's growth rate is weighted by their share of the total income, $w_i = x_i / \sum_{i=1}^N x_i$. The ensemble-average (GDP per capita) growth rate is, therefore, a *plutocratic* measure of growth (Milanovic, 2005), in the sense that each person's weight in the average is proportional to their income.

⁸What we call the ensemble-average growth rate (GDP per capita growth rate) is the growth rate of the ensemble average (GDP per capita) and, we stress, not the ensemble average of the growth rate.

⁹For small income changes, $\Delta \ln x_i \approx \Delta x_i / x_i$ and

$$\frac{\Delta \ln \langle x(t) \rangle_N}{\Delta t} \approx \frac{\Delta \langle x \rangle_N / \langle x \rangle_N}{\Delta t} \quad (2.12)$$

$$= \frac{\sum_{i=1}^N \Delta x_i}{\sum_{i=1}^N x_i \Delta t} \quad (2.13)$$

$$= \sum_{i=1}^N \frac{x_i}{\sum_{i=1}^N x_i} \frac{\Delta x_i / x_i}{\Delta t} \quad (2.14)$$

$$\approx \sum_{i=1}^N w_i \frac{\Delta \ln x_i}{\Delta t}, \quad (2.15)$$

where $w_i = x_i / \sum_{i=1}^N x_i$ as required.

2.3 Time-average growth rate

Whereas the ensemble-average growth rate requires a model where we can average over indefinitely many individuals, the *time-average growth rate* of income requires just one modelled individual but an indefinitely long income trajectory. It is defined as the $\Delta t \rightarrow \infty$ limit of the individual growth rate in (2.6),

$$\bar{g} \equiv \lim_{\Delta t \rightarrow \infty} \frac{\Delta \ln x_i(t)}{\Delta t}. \quad (2.16)$$

Conceptually, therefore, the time-average growth rate measures the typical growth of an individual, experienced in a single income trajectory unfolding over time.

Alternatively, $\bar{g}$ can be expressed as a time average if we partition the period Δt into T shorter periods of fixed duration δt :

$$\bar{g} = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{\tau=0}^{T-1} \frac{\delta \ln x_i(t + \tau \delta t)}{\delta t}. \quad (2.17)$$

Thus we see that the two different procedures – computing a long-time growth rate and averaging a short-time growth rate over many time steps – are equivalent. This convenient property follows from our definition of the growth rate as the (additive) rate of change of a function of the process.

Of course, we cannot compute any infinite-time limit from data of real incomes. However, we can exploit the ergodicity of the growth rate, $\bar{g} = \langle g \rangle$, derived in Section 2.1, to express the time-average growth rate as an ensemble average,

$$\bar{g} = \left\langle \frac{\Delta \ln x(t)}{\Delta t} \right\rangle, \quad (2.18)$$

which for a large population we can estimate as a finite- N average,

$$\bar{g}_N \equiv \left\langle \frac{\Delta \ln x(t)}{\Delta t} \right\rangle_N. \quad (2.19)$$

(2.19) is our practical definition of the time-average growth rate of income.

Our ability to estimate the time-average growth rate as an ensemble average is worth pondering. At any moment in time there exists a long income trajectory that would result from experiencing the ensemble of short-time ergodic growth rates in sequence. We speculate that this momentary estimate of the long-time growth rate corresponds closely to how an economy feels to an individual. It tells the individual what he or she can expect in the long run if the balance of risks and opportunities stays the same.¹⁰

¹⁰Another interpretation of $\bar{g}_N$ is the growth rate a hypothetical person would experience if they spent $1/N$ of each year growing like every person in the population. Since $\bar{g}_N = \langle g \rangle_N$, we have

$$x(t) \exp(\bar{g}_N \Delta t) = x(t) \exp \left(\sum_{i=1}^N \frac{g_i \Delta t}{N} \right). \quad (2.20)$$

This everyman growing at $\bar{g}_N$, unlike his lucky compatriot growing at $g_{\langle \rangle_N}$, shares in all of the experiences of the

The time-average growth rate is an average of individual growth rates, in which each person's growth rate has equal weight. This makes it a *democratic* measure of growth, as explained in Section 1, and we also call it the DDP per capita growth rate.

DDP per capita is the quantity, $y(t)$, that grows at the time-average growth rate, *i.e.* $y(t + \Delta t) = y(t) \exp(\bar{g}_N \Delta t)$. Rearranging gives the DDP per capita as

$$y(t) = \exp(\langle \ln x(t) \rangle_N), \quad (2.21)$$

which is the geometric mean income and has the same units as income.

(2.19) demonstrates that knowledge of individual growth rates is unnecessary for computing $\bar{g}_N$. We need only the values of $\langle \ln x \rangle_N$ at the start and end of the observation period. This has practical importance, since reliable panel data, required for estimating individual growth rates, are harder to obtain than data on the initial and terminal distributions. It also means that $\bar{g}_N$ is independent of income mobility, *i.e.* of the reordering that takes place between observations.

2.4 Inequality

The ensemble-average (GDP per capita) growth rate and time-average (DDP per capita) growth rate are conceptually different. The former is the growth rate of the average (or, multiplying by N , the total) income in a large population. The latter is the growth rate achieved by a single member of the population over long time. Under multiplicative growth, and in general, these growth rates do not have the same value.

We can show this by a simple thought experiment: suppose all N incomes in a population are equal at time t , *i.e.* $x_i(t) = x(t)$ for all i , and that they evolve to different incomes, $x_i(t + \Delta t)$, at time $t + \Delta t$. The ensemble-average (GDP per capita) growth rate is

$$g_{\langle \rangle_N} = \frac{\ln \langle x(t + \Delta t) \rangle_N - \ln \langle x(t) \rangle_N}{\Delta t} \quad (2.22)$$

and the time-average (DDP per capita) growth rate is

$$\bar{g}_N = \frac{\langle \ln x(t + \Delta t) \rangle_N - \langle \ln x(t) \rangle_N}{\Delta t}. \quad (2.23)$$

Since the logarithm is a concave function, the famous result of Jensen (1906) implies that the difference between the growth rates is non-negative, *i.e.*

$$g_{\langle \rangle_N} - \bar{g}_N = \frac{1}{\Delta t} [\ln \langle x(t + \Delta t) \rangle_N - \langle \ln x(t + \Delta t) \rangle_N] \geq 0, \quad (2.24)$$

vanishing only if the $x_i(t + \Delta t)$ are equal.

national economy in a single year. A truly representative agent.

This thought experiment is illuminating: we started with perfect income equality, moved to a state of income inequality, and found that $g_{\langle \rangle_N} > \bar{g}_N$. In general, if the ensemble-average growth rate is higher than the time-average growth rate (GDP per capita grows faster than DDP per capita), then the current high earners grew faster than the low earners over the observation window, and inequality increased. If the converse is true, then inequality decreased.

In this picture, a natural measure of inequality (of which myriad measures exist in the literature) is the quantity which grows at the difference between the two growth rates (Adamou and Peters, 2016). We call this measure *the ergodicity gap* and denote it by $J(t)$ to highlight its connection to Jensen (1906). We have

$$\frac{\Delta J}{\Delta t} = g_{\langle \rangle_N} - \bar{g}_N = \frac{\Delta \ln \langle x \rangle_N}{\Delta t} - \frac{\Delta \langle \ln x \rangle_N}{\Delta t}, \quad (2.25)$$

which, integrated over time, gives the ergodicity gap as

$$J(t) = \ln \langle x(t) \rangle_N - \langle \ln x(t) \rangle_N, \quad (2.26)$$

up to an additive constant. $J(t)$ is an established income inequality measure, known as the *mean logarithmic deviation* (MLD) and also as the second Theil index (Theil, 1967).¹¹

2.5 Example: geometric Brownian motion

The results above are valid for any large population in which incomes are assumed to follow random multiplicative growth. Here we study a simple example of such dynamics: geometric Brownian motion (GBM). We do this to show how the quantities we have defined – the ensemble-average (GDP per capita) growth rate, the time-average (DDP per capita) growth rate, and the ergodicity gap (MLD) – behave in an analytically solved case.

GBM is a continuous-time model of income in which relative changes are normally distributed. Each income $x_i(t)$ obeys the Itô stochastic differential equation,

$$dx_i = x_i(\mu dt + \sigma dW_i), \quad (2.27)$$

where μ and σ are the drift and volatility parameters, dx_i is the infinitesimal change in income, dt is the infinitesimal change in time, and $dW_i \sim \mathcal{N}(0, dt)$ is the infinitesimal increment in a Wiener process. Because the model is well known, we only list the key model properties here, leaving derivations to Appendix B.

Under GBM, the ensemble-average (GDP per capita) growth rate of income is

$$g_{\langle \rangle} = \frac{\Delta \ln \langle x \rangle}{\Delta t} = \mu, \quad (2.28)$$

¹¹In general, the form of the ergodicity gap, like the form of the ergodic growth rate, depends on the transformation required to make income changes ergodic. Here, under multiplicative dynamics, this transformation is logarithmic.

and the time average growth rate (DDP per capita growth rate) of income is

$$\bar{g} = \frac{\Delta \langle \ln x \rangle}{\Delta t} = \mu - \frac{\sigma^2}{2}. \quad (2.29)$$

Thus, the ensemble-average income (GDP per capita) grows faster than a single income trajectory in the long-time limit. This serves as a concrete example of the more abstract result in (2.24) and it has an important corollary: using the ensemble-average (GDP per capita) growth rate as the principal growth measure can create a misleading impression. In GBM, it is a systematic overestimate of how individual citizens experience the economy, as captured by the time-average (DDP per capita) growth rate.

Figure 1 shows the distribution of income and growth at the ensemble-average (GDP per capita) and time-average (DDP per capita) growth rates under GBM. The difference, both conceptual and quantitative, between these growth rates is the central observation in this study.

Figure 1: The distribution of income in GBM. The blue line is the ensemble-average income (GDP per capita), while the red line is the DDP per capita, which grows at the time-average growth rate. In the left panel, the thin dark lines are 200 simulated individual income trajectories. In the right panel the shaded areas represent one (darker shade) and two (lighter shade) standard deviations from the DDP per capita. The parameters used are $\mu = 0.05 \text{ year}^{-1}$ and $\sigma = 0.2 \text{ year}^{-\frac{1}{2}}$. We assume the initial income is $\$10^4/\text{year}$.

The ergodicity gap (MLD), $J(t)$, grows at the difference between the ensemble-average (GDP per capita) and time-average (DDP per capita) growth rates. In the GBM model, this difference is

$$\frac{dJ}{dt} = g_{\langle} - \bar{g} = \frac{\sigma^2}{2}, \quad (2.30)$$

such that

$$J(t) = \frac{\sigma^2 t}{2}, \quad (2.31)$$

up to an additive constant. As the distribution of income broadens with time, so income inequality – as measured by the ergodicity gap (MLD) – increases without bound.

GBM is a model in which members of an ensemble grow their incomes exponentially and randomly, without ever interacting. In particular, there is no reallocation of income between members, as might happen in a real economy with social institutions (Berman, Peters and Adamou, 2020). It paints a stark picture in which the aggregate income grows systematically faster than the typical income, an unintuitive result made possible only by ever-increasing inequality. Such phenomena are inherently interesting to economists, policymakers, and indeed citizens. They become easier to detect when the ensemble-average (GDP per capita) growth rate, time-average (DDP per capita) growth rate, and ergodicity gap (MLD) are reported.

3 Growth rates in the United States and France

To demonstrate the differences between $\bar{g}_N$ and $g_{\langle \rangle_N}$, we use income data from the United States and France over the last 50 years.¹² Figure 2 presents the evolution of $\bar{g}_N$ and $g_{\langle \rangle_N}$ in these two countries. It also shows the evolution of GDP per capita and DDP per capita.

In the United States, ensemble-average income growth (GDP per capita growth) and time-average income growth (DDP per capita growth) were similar between the mid-1960s and the early 1980s. Between 1966 and 1980, the ensemble-average income increased by 13% and the time-average income increased by 16%. Since 1980, there has been a clear divergence. Ensemble-average growth totalled +40% between 1980 and 2014, whereas time-average growth totalled +20%. This makes ensemble-average growth (meaning GDP per capita growth) a misleading quantity to follow in the United States, if the wellbeing of individuals is of interest.

In France, by contrast, the two growth rates were much closer. The ensemble-average (GDP per capita) growth rate does not create a misleading picture of the typical economic experience of France's citizens and it closely resembles the time-average (DDP per capita) growth rate. This means that a multiplicative growth model with random shocks may be inappropriate in this case. Effects that lead to divergence of incomes may be weaker here, or effects that lead to convergence may be stronger. We study a model that can capture such effects in Berman, Peters and Adamou (2020).

¹² $g_{\langle \rangle_N}$ is easily estimated by using the average national income for $\langle x \rangle_N$ in (2.22). Calculating $\bar{g}_N$ using (2.23) requires estimating $\langle \ln x \rangle_N$, which is done using income quantile data (see Appendix C).

Figure 2: The evolution of the two growth rates of the economy. Top: the ensemble-average (blue) and the time-average (red) growth rates in the United States (left) and in France (right). Bottom: the evolution of ensemble-average income (GDP per capita) and DDP per capita in the United States (left) and in France (right). All data were taken from [The World Inequality Database \(2020\)](#). We note that the time-average growth rate is estimated by considering the changes in the geometric mean income each year and not by averaging over all individual growth rates (see Appendix A for more details). We also truncated the bottom decile of the distribution (see Appendix D for more details).

Figure 2 shows that in the United States the ergodicity gap remained stable between 1966 and the early 1980s. Since then it has grown. In France, the ergodicity gap has been roughly stable since 1970. Figure 3 depicts the evolution of the top 10% income share and the ergodicity gap in the United States and in France over the last 50 years. It demonstrates that both measures evolve similarly in both countries.

4 Discussion

Economic growth is characterised by two growth rates: the ensemble-average (GDP per capita) growth rate, which is the growth rate of the average income; and the time-average (DDP per

Figure 3: The evolution of the top 10% income share and the ergodicity gap in the United States (left) and in France (right). The correlation between the top 10% share and the ergodicity gap was 0.94 in the United States and 0.79 in France. All data were taken from [The World Inequality Database \(2020\)](#).

capita) growth rate, achieved by each individual in the long run. When incomes follow random multiplicative growth, the ensemble-average (GDP per capita) growth rate is greater than the time-average (DDP per capita) growth rate, and inequality – measured as the ergodicity gap (MLD) – increases indefinitely. From the perspective of ergodicity economics, these results are unsurprising: they are direct consequences of the non-ergodicity of the stochastic process for income.

The difference between these growth rates is consequential for policymakers. At present, the most widely used summary statistic of national economic growth is the ensemble-average (GDP per capita) growth rate. This captures growth in the nation’s total productive capacity – referred to colloquially as ‘the size of the economy’ – and is useful. However, in the model, it is systematically higher than the time-average (DDP per capita) growth rate, which reflects the typical individual’s experience of the economy. If these circumstances were mirrored in reality, preoccupation of policymakers with ensemble-average (GDP per capita) growth would create a worrying disconnect between them and the public. There are reasons to believe this situation is present-day reality in the United States, see Figure 2 and [Boushey and Clemens \(2018\)](#); [Saez and Zucman \(2019\)](#).

We conclude with some interpretations of these fundamental growth rates. The ensemble-average (GDP per capita) growth rate is the growth rate of average income. By contrast, the time-average (DDP per capita) growth rate is the average growth rate of income. The former is approximately the income-weighted average of individual growth rates and we can think of it, therefore, as a ‘one dollar, one vote’ measure of economic growth (strictly speaking, one dollar per year). The latter is an equally weighted average of individual growth rates and we can think of it as a ‘one person, one vote’ measure. Both have their place in the analysis of economic growth.

References

- Adamou, Alexander, and Ole Peters.** 2016. “Dynamics of Inequality.” *Significance*, 13(3): 32–35.
- Aitchison, John, and James A. C. Brown.** 1957. *The Lognormal Distribution*. Cambridge University Press.
- Berman, Yonatan, Ole Peters, and Alexander Adamou.** 2020. “Wealth Inequality and the Ergodic Hypothesis: Evidence from the United States.” *Journal of Income Distribution*. *Forthcoming*.
- Boushey, Heather, and Austin Clemens.** 2018. “Disaggregating Growth: Who Prospers when the Economy Grows.” Washington Center for Equitable Growth Equitable Growth March.
- Chenery, Hollis, Montek S. Ahluwalia, C. L. G. Bell, John H. Duloy, and Richard Jolly.** 1974. *Redistribution with Growth: Policies to Improve Income Distribution in Developing Countries in the Context of Economic Growth*. London: Oxford University Press.
- Gabaix, Xavier, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll.** 2016. “The Dynamics of Inequality.” *Econometrica*, 84(6): 2071–2111.
- Gibrat, Robert.** 1931. *Les Inégalités Économiques*. Librairie du Recueil Sirey.
- Guvenen, Fatih, Sam Schulhofer-Wohl, Jae Song, and Motohiro Yogo.** 2017. “Worker Betas: Five Facts about Systematic Earnings Risk.” *American Economic Review: Papers & Proceedings*, 107(5): 398–403.
- Jensen, Johan L. W. V.** 1906. “Sur les fonctions convexes et les inégalités entre les valeurs moyennes.” *Acta Mathematica*, 30: 175–193.
- Milanovic, Branko.** 2005. *Worlds Apart: Measuring International and Global Inequality*. Princeton University Press.
- Peters, Ole.** 2019. “The Ergodicity Problem in Economics.” *Nature Physics*, 15: 1216–1221.
- Peters, Ole, and Alexander Adamou.** 2018. “The Time Interpretation of Expected Utility Theory.” [arXiv:1801.03680](https://arxiv.org/abs/1801.03680).
- Saez, Emmanuel, and Gabriel Zucman.** 2019. *The Triumph of Injustice: How the Rich Dodge taxes and How to Make Them Pay*. WW Norton & Company.
- Theil, Henri.** 1967. *Economics and Information Theory*. North-Holland Publishing Company.
- The World Inequality Database.** 2020. “USA and France Income Data.” <http://wid.world/data/>, Accessed: 02/23/2020.

A Finite populations

In Section 2 we defined the ensemble-average and time-average growth rates of income. By “ensemble-average growth rate” we mean the growth rate of the ensemble average of income (not the ensemble average of the growth rate). In the random multiplicative growth model, this is:

$$g_{\langle} \equiv \frac{\Delta \ln \langle x(t) \rangle}{\Delta t}. \quad (\text{A.1})$$

The time-average growth rate is the $\Delta t \rightarrow \infty$ limit of growth rate of individual income, which we express as the ensemble average of the ergodic growth rate:

$$\bar{g} \equiv \left\langle \frac{\Delta \ln x(t)}{\Delta t} \right\rangle. \quad (\text{A.2})$$

These are deterministic quantities, which are properties of the model. They can be expressed in terms of model parameters, as we do for GBM in Section 2.5.

For a finite population of N individuals, we use the finite- N versions of these growth rates:

$$g_{\langle N} = \frac{\Delta \ln \langle x(t) \rangle_N}{\Delta t}; \quad \bar{g}_N = \left\langle \frac{\Delta \ln x(t)}{\Delta t} \right\rangle_N. \quad (\text{A.3})$$

These growth rates are statistics that can be computed from observations of a real population. In the model, they are random variables.

Here we verify that, in a model population of the size of a typical nation, the finite- N growth rates in (A.3) are good approximations of their $N \rightarrow \infty$ limits. Our aim is to establish whether these limiting growth rates would be relevant in practice, if the incomes of a real finite population were well described by random multiplicative growth. Of course, unmodelled effects – such as reallocation of income and correlation of income changes – may be important in reality, and conclusions drawn from the limiting growth rates of a model may be inapt. In our empirical study in Section 3, we do not use asymptotic results: rather we estimate finite- N growth rates and watch them evolve over time.

We use GBM, presented in Section 2.5 and in greater detail in Appendix B, as our model of income, where $g_{\langle} = \mu$ and $\bar{g} = \mu - \sigma^2/2$. For the ensemble-average growth rate, we use numerical simulations with $\mu = 0.02$ per year and σ^2 ranging from 0 to 0.09 per year.¹³ We let N change from 2 to 10^7 . For each value of σ and N , we simulate one time step (equivalent to one year) in a GBM starting from a lognormal distribution, and we compute $g_{\langle N}$. We wish to show that $g_{\langle N}$ is distributed tightly enough that it is usually closer to μ than to $\mu - \sigma^2/2$. We do this by repeating the simulation 10^4 times for each value of σ and N , creating a sample of $g_{\langle N}$ values. We then check what fraction of the sample satisfies $g_{\langle N} > \mu - \sigma^2/4$. The left panel in Figure 4 shows this fraction as a function of σ and N . It is above 90% for typical national populations and realistic values of σ .

¹³Realistically, for wealth, σ^2 is less than 0.04 per year in the US (Berman, Peters and Adamou, 2020).

Figure 4: Left: The fraction of $g_{\langle \rangle_N}$ values closer to μ than to $\mu - \sigma^2/2$ as a function of σ and N in 10^4 realisations of GBM; Right: The probability of $\bar{g}_N$ being closer to $\mu - \sigma^2/2$ than to μ as a function of σ and N , see (A.5).

We conduct a similar analysis for the time-average growth rate, now requiring that $\bar{g}_N$ is usually closer to $\mu - \sigma^2/2$ than to μ . Here the analysis is simpler, since in the model $\bar{g}_N$ is just the sample mean of a normally distributed growth rate. Its sampling distribution is

$$\bar{g}_N \sim \mathcal{N}\left(\mu - \frac{\sigma^2}{2}, \frac{\sigma^2}{N\Delta t}\right), \quad (\text{A.4})$$

and we are interested in the probability $P(\bar{g}_N < \mu - \sigma^2/4)$. For the normal distribution this is

$$\frac{1}{2} + \frac{1}{2} \operatorname{erf}\left(\sqrt{\frac{\sigma^2 N \Delta t}{32}}\right). \quad (\text{A.5})$$

The right panel in Figure 4 shows that this probability is above 90% for typical national populations and realistic values of σ .

B Geometric Brownian motion

In Section 2.5 we use GBM as an example model of income dynamics. Here we present derivations of the results described therein.

We recall the basic multiplicative model of income evolution in (2.1),

$$x(t + \delta t) = x(t)(1 + \epsilon), \quad (\text{B.1})$$

where δt is the time between income observations and ϵ is a random variable.

In GBM, periodic income multipliers are independent and identically distributed normal random

variables. We introduce this using finite-size increments. Income evolves as

$$x(t + \delta t) = x(t)(1 + \mu\delta t + \sigma\zeta\sqrt{\delta t}), \quad (\text{B.2})$$

where ζ is a standard normal variate, $\zeta \sim \mathcal{N}(0, 1)$. The model has two parameters, the drift μ and the volatility σ . Here $\mu\delta t$ is the common deterministic component of the income multiplier, and $\sigma\zeta\sqrt{\delta t}$ is the idiosyncratic stochastic part.

In effect, we have chosen $\epsilon \sim \mathcal{N}(\mu\delta t, \sigma^2\delta t)$ in (B.1). This scaling of mean and variance with time ensures that the random variable for one long time step has a distribution consistent with that obtained by summing successive random variables over smaller time steps of equal total duration. This matters because results should not depend on how we partition time.

Defining the change in income at time t as

$$\delta x = x(t + \delta t) - x(t), \quad (\text{B.3})$$

we can write (B.2) as

$$\delta x = x(t)(\mu\delta t + \sigma\zeta\sqrt{\delta t}). \quad (\text{B.4})$$

We now take the limit $\delta t \rightarrow 0$ to obtain an Itô stochastic differential equation for income $x(t)$,

$$dx = x(\mu dt + \sigma dW). \quad (\text{B.5})$$

Thus we arrive at (2.27) in the main text.

(B.5) is solved by using Itô's formula to find the stochastic differential equation obeyed by $\ln x(t)$. This is a Brownian motion with drift,

$$d \ln x = \left(\mu - \frac{\sigma^2}{2} \right) dt + \sigma dW, \quad (\text{B.6})$$

which is integrated to give the solution for a single trajectory of the stochastic process,

$$\ln x(t + \Delta t) = \ln x(t) + \left(\mu - \frac{\sigma^2}{2} \right) \Delta t + \sigma W(\Delta t), \quad (\text{B.7})$$

or

$$x(t + \Delta t) = x(t) \exp \left[\left(\mu - \frac{\sigma^2}{2} \right) \Delta t + \sigma W(\Delta t) \right]. \quad (\text{B.8})$$

Δt is the period of time over which the process runs and $W(\Delta t)$ is a Wiener process, with distribution $\mathcal{N}(0, \Delta t)$.

Distribution Unlike the general multiplicative process in (B.1), we see from (B.7) that the change in logarithmic income is normally distributed at all times rather than just for long time.

The distribution is

$$\ln x(t + \Delta t) - \ln x(t) \sim \mathcal{N} \left(\left(\mu - \frac{\sigma^2}{2} \right) \Delta t, \sigma^2 \Delta t \right) \quad (\text{B.9})$$

and it is time-dependent. In particular, its variance grows linearly in Δt .

Exponential growth Income grows exponentially. The growth rate,

$$g = \frac{\Delta \ln x}{\Delta t}, \quad (\text{B.10})$$

of a single income over the period Δt is a random variable, which we can extract from (B.7) as

$$g = \mu - \frac{\sigma^2}{2} + \frac{\sigma W(\Delta t)}{\Delta t}. \quad (\text{B.11})$$

This has distribution

$$g \sim \mathcal{N} \left(\mu - \frac{\sigma^2}{2}, \frac{\sigma^2}{\Delta t} \right). \quad (\text{B.12})$$

From these results we derive the ensemble-average and time-average growth rates in a GBM, by taking the limits $N \rightarrow \infty$ and $\Delta t \rightarrow \infty$ respectively.

Ensemble-average growth rate We recall the definition of $g_{\langle \rangle}$ in (2.9),

$$g_{\langle \rangle} \equiv \frac{\Delta \ln \langle x(t) \rangle}{\Delta t}. \quad (\text{B.13})$$

We evaluate this by taking the ensemble average of both sides of (B.8) to get

$$\langle x(t + \Delta t) \rangle = \langle x(t) \rangle \exp(\mu \Delta t), \quad (\text{B.14})$$

the growth rate of which we obtain by inspection as

$$g_{\langle \rangle} = \mu. \quad (\text{B.15})$$

Time-average growth rate The time-average growth rate is the long-time limit of the individual growth rate in (B.10), *i.e.*

$$\bar{g} \equiv \lim_{\Delta t \rightarrow \infty} \{g\}. \quad (\text{B.16})$$

The variance of g in (B.12) tends to zero as $\Delta t \rightarrow \infty$. Therefore, the growth rate converges to the mean of the distribution in this limit, *i.e.*

$$\bar{g} = \mu - \frac{\sigma^2}{2}. \quad (\text{B.17})$$

C Estimating $\bar{g}_N$ with quantile data

The calculation of the time-average (DDP per capita) growth rate in (2.23) requires us to estimate $\langle \ln x \rangle_N$ at different points in time. In practice we do this using quantile average data. Instead of directly computing

$$\bar{g}_N = \frac{\Delta \langle \ln x \rangle_N}{\Delta t}, \quad (\text{C.1})$$

which would be possible if all N incomes were available, we compute

$$\bar{g}_N \approx \frac{\frac{1}{Q} \sum_{q=1}^Q \Delta \ln x^{(q)}}{\Delta t}, \quad (\text{C.2})$$

where Q is the number of quantiles (*e.g.* $Q = 100$ for percentiles) and $x^{(q)}$ is the average value of x within the q^{th} quantile.

Clearly, the finer the division into quantiles is, the more accurate the approximation becomes. To validate the values of $\bar{g}_N$ shown in Figure 2, we compare them to values for $\bar{g}_N$ obtained using the lower and upper thresholds of each quantile as the quantile average. These give crude upper and lower bounds for $\bar{g}_N$. Figure 5 shows very little difference between $\bar{g}_N$ as calculated in (C.2) and these bounds. We conclude that the sensitivity of our results to division into quantiles is low, and that (C.2) with percentile averages is a good approximation of the time-average growth rate in practice.

Figure 5: The DDP per capita using percentile data. Top: the time-average growth rate (solid red) in the United States (left) and in France (right) estimated using (C.2) based on percentile averages, lower percentile thresholds (dashed black), and upper percentile thresholds (dotted black). Bottom: the evolution of the DDP per capita in the United States (left) and in France (right) based on percentile averages (solid red), lower percentile thresholds (dashed black) and upper percentile thresholds (dotted black). All data were taken from [The World Inequality Database \(2020\)](#).

D Truncation of data

In Section 2 we show that DDP per capita can be computed as the geometric mean of income. As such, it requires incomes to be positive and is sensitive to uncertainty in the measurement of low incomes. Therefore, we present results in Figure 2 using truncated data, where the bottom 10% of incomes are excluded from the calculation of growth rates and average incomes.

Figure 6 explains why we chose to truncate at this percentile. It shows how the growth of GDP per capita and DDP per capita in the United States and France depend on the fraction of incomes excluded from their computation. Truncating at percentile 1 means that all percentiles are included, while truncating at percentile 100 means that only the top percentile is included (in which case GDP per capita and DDP per capita coincide by construction). In between, DDP per capita growth changes rapidly when the point of truncation varies from percentile 1 to 10, and more slowly after

that. This reflects the well-known fact that geometric means are sensitive to changes in small-value data. The arithmetic mean, used to compute GDP per capita, is less sensitive to truncation. Specifically, it shows no major change in truncation-dependence around percentile 10.

Figure 6: The cumulative growth of income depending on point of truncation in the United States and France. The figure depicts the cumulative growth of the GDP per capita (blue) and DDP per capita (red) in the United States between 1966 to 2014 (left), and in France between 1970 to 2015 (right), as a function of the point of truncation. When the truncation is done at percentile 1, all percentiles are included. When it is done at percentile 100, only the top percentile is included and the GDP per capita and DDP per capita coincide. All data were taken from [The World Inequality Database \(2020\)](#).

While truncation has a quantitative effect on the evolution of the GDP and DDP per capita presented in Figure 2, it does not change qualitatively the message. Figure 7 presents the GDP and DDP per capita and their growth rates in the United States and in France without truncating the distribution. We draw similar conclusions about the differences between the United States and France as we do from the truncated data in Figure 2.

Figure 7: The evolution of the two growth rates of the economy. Top: the ensemble-average (blue) and the time-average (red) growth rates in the United States (left) and in France (right). Bottom: the evolution of ensemble-average income (GDP per capita) and DDP per capita in the United States (left) and in France (right). All data were taken from [The World Inequality Database \(2020\)](#). Compare with Figure 2 in which the bottom decile of incomes was excluded from the data in all years.