An empirical test of the ergodic hypothesis: wealth distributions in the United States

Yonatan Berman[*], Ole Peters[†] and Alexander Adamou[‡§]

Abstract

Classical studies of wealth inequality make the "ergodic hypothesis" that rescaled wealth converges rapidly to a stationary distribution. This allows powerful analytical techniques but constrains models. We test the hypothesis and find it unsupported by data. In a simple model of a growing economy with wealth reallocation, we fit the reallocation parameter to historical United States wealth data. We find negative reallocation, *i.e.* from poorer to richer, for which no stationary distribution exists. When we find positive reallocation, convergence to the stationary distribution is slow. Both cases invalidate the ergodic hypothesis. Studies which make it should be treated with caution.

Keywords: Ergodic hypothesis, wealth inequality, stochastic processes.

JEL Codes: C62, D31, E01, P46.

[*]School of Physics and Astronomy, Tel-Aviv University, Tel-Aviv, Israel, yonatanb@post.tau.ac.il

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

[‡]London Mathematical Laboratory, London, UK, a.adamou@lml.org.uk

[§]We dedicate this study to the memory of Tony Atkinson, who provided invaluable comments to us in Oxford in October 2016. His encouragement was crucial to our work. YB also thanks him for his support through the Dan David Prize Scholarship for Young Researchers. We thank Jean-Philippe Bouchaud for sharing his expertise on the solutions of Equation (7) and Tobias Morville for his comments on the manuscript. YB acknowledges travel support and hospitality from the London Mathematical Laboratory.

1 Introduction

Paul Samuelson identified the "ergodic hypothesis" in the mindset of classical economic theorists, defining it as "a belief in unique long-run equilibrium independent of initial conditions" (Samuelson, 1968, pp. 11-12). Classically, the equilibrium is studied and the transient phenomena preceding it are ignored. This approach simplifies analysis greatly and is widespread. Here we present an empirical test of the validity of the ergodic hypothesis as applied to the distribution of wealth. We find it inconsistent with historical data from the United States (US), raising doubts about the conclusions of studies in which it is made.

Economics is often concerned with growth. A growing quantity cannot be ergodic in Samuelson's sense because it does not tend to an equilibrium value. Naïvely, then, the ergodic hypothesis seems an inappropriate modelling choice in this context, ruling out a swathe of analytical techniques. Although this is rarely stated explicitly, a common strategy to salvage these techniques is the following: find a transformation of the non-ergodic process that produces a meaningful ergodic observable. If such an ergodic observable can be derived, then classical techniques may still be useful. For instance, Peters and Gell-Mann (2016) show that utility theory can be viewed as an attempt to transform non-ergodic growing wealth into ergodic growth rates. Expectation values – which would otherwise be misleading – then quantify time-average growth of the decision-maker's wealth.

Studies of wealth distributions also employ this strategy. Individual wealth is modelled as a growing quantity. Dividing by the population average transforms this to a rescaled wealth, which is hypothesised to be ergodic. Formally, the distribution of the stochastic process for the rescaled wealth is assumed to converge to a unique and time-independent distribution. For example, Benhabib et al. (2011) "impose assumptions [...] that guarantee the existence and uniqueness of a limit stationary distribution" (p. 130). Similar formulations appear in Stiglitz (1969); Bewley (1977); Piketty and Saez (2013); De Nardi (2015); De Nardi et al. (2015); Jones (2015). These studies take advantage of the simplicity with which the stationary distribution can then be analyzed, *e.g.* to test the effects of policies encoded in model parameters.

Treating the stationary model distribution as representative of the empirical wealth distribution implies an assumption of fast convergence, *i.e.* that the actual distribution approaches its asymptotic form on a timescale shorter than that of relevant changes in economic conditions. For example, the top rate of income tax in the US was reduced from 91 percent to 70 percent between 1963 and 1965. The next change occurred about 20 years later, in 1982, when it was reduced to 50 percent. In this example, convergence would be fast if the model distribution were close to its asymptotic form after, say, 5 years. Convergence would be slow if this took, say, 100 years. Atkinson (1969, p. 137) argued that "the speed of convergence makes a great deal of difference to the way in which we think about the model". It determines the practical relevance of the stationary distribution (Atkinson, 1969; Cowell, 2014).

Recent research has marked a shift away from studying the stationary distribution in isolation. Gabaix et al. (2016); Berman et al. (2016); Kaymak and Poschke (2016); Berman and Shapira (2017) study the dynamics of wealth and income, rather than how they are distributed asymptotically. Some of these studies also consider convergence times (Gabaix et al., 2016; Berman and Shapira, 2017).

There is, however, an elephant in the room. To our knowledge, the validity of the ergodic hypothesis for rescaled wealth has never been tested empirically. Here we present such a test. For the test to be meaningful, a model is required that does not assume ergodicity from the outset. It must have regions of parameter space ("regimes") in which the ergodic hypothesis for rescaled wealth does and does not hold.

Our model satisfies this condition. Its basis is that individual wealth undergoes noisy exponential growth. Social structure is represented by a wealth reallocation mechanism in which a fraction of everyone's wealth is pooled and shared. The sign of the reallocation rate parameter determines whether a stationary distribution exists: if positive (corresponding to reallocation from richer to poorer) then it does; if negative (from poorer to richer) then it doesn't.

We estimate model parameters using three different datasets of historical wealth shares in the US population. These estimated parameters – and not an *a priori* assumption – tell us whether a stationary distribution exists. When it does, it has a Pareto tail, consistent with data and standard models (Pareto, 1897; Drăgulescu and Yakovenko, 2001). We evaluate convergence times, when convergence is possible, to assess whether the ergodic hypothesis is acceptable in practice.

Our findings invalidate the ergodic hypothesis. The fitted reallocation rate is not robustly positive for any dataset we analyze. Indeed, for one dataset we find it to be consistently negative for the last thirty years or so. We cannot overstate our surprise at this finding. Most theorists would consider a model in which individual wealths grow independently, *i.e.* with no reallocation, an extreme and unrealistic model of an advanced Western economy with socio-political institutions and infrastructure. We would expect to infer consistent positive reallocation from data for such an economy. We find the opposite: that, from the 1980s to the present, the US economy is best described in our model as one in which wealth is systematically reallocated from poorer to richer.

Other datasets yield reallocation rates for which a stationary distribution does exist, but with convergence times of decades or centuries. Whichever data are used, our analysis does not provide the unequivocal endorsement of the ergodic hypothesis that would justify its ubiquitous use in the field. Policy recommendations based on models which assume the existence of a stationary distribution and fast convergence may be ineffective in practice. Worse still, such inappropriately-constrained models may paint a misleading picture of reality, for example that with taxation and public spending our economies are positively redistributive. This could lead to policy prescriptions which run counter to policy goals.

We also extend our model to treat explicitly additive changes in wealth, akin to labor income and consumption. This allows us to answer the question whether rescaled wealth is inherently unstable, or whether it is inherently stable with increases in its inequality driven by increasingly unequal earnings. We find that earnings had only a small effect on the dynamics of the wealth distribution over the last century and that changes in their distribution do not explain adequately the observed instability in wealth.

Our contribution is threefold. Firstly, we develop a theoretical model which describes the dynamics of wealth and allows us to test empirically the validity of the ergodic hypothesis. Naturally, this cannot be done using models in which the ergodic hypothesis is implicit.

Secondly, we use the model to assess how fast the wealth distribution converges to its asymptotic form, when this exists. We find convergence times greater than the typical intervals between changes in policy and other determinants of the distribution. This casts doubt over the ability of models with fast convergence to interpret observed changes in inequality and yield effective policy.

Thirdly, our study extends a body of theory of economic systems without assuming ergodicity. Peters and Gell-Mann (2016) show that a sound decision theory can be developed by finding ergodic growth rates for non-ergodic wealth processes. In the present context, Adamou and Peters (2016) derive a fundamental measure of wealth inequality from such growth rates. Elsewhere, Peters (2011b) resolves the St. Petersburg paradox by maximizing individual performance over time and Peters (2011a) applies the same reasoning to derive an optimal leverage for stock-market investments. Peters and Adamou (2013) extend this work to derive constraints on price fluctuations in freely-traded assets, resolving the equity premium puzzle of Mehra and Prescott (1985).

Indiscriminate use of the ergodic hypothesis can mean that non-ergodic processes are analyzed with methods appropriate only for ergodic processes. This is a severe methodological problem, which we argue is behind many open problems in economics. Our work demonstrates the importance of establishing empirically and analytically which observables may be treated legitimately as ergodic.

2 Model

We call our model Reallocating Geometric Brownian Motion (RGBM). Individual wealth undergoes random multiplicative growth, modelled as Geometric Brownian Motion (GBM), and is reallocated among individuals by a simple pooling and sharing mechanism. Thus everyone's wealth is coupled to the total wealth in the economy. We view RGBM as a null model of an exponentially growing economy with social structure. It is intended to capture only the most general features of the dynamics of wealth.

RGBM has both ergodic and non-ergodic regimes, characterised by the sign of the reallocation rate parameter. Reallocation from richer to poorer produces an ergodic regime, in which wealths are positive, distributed with a Pareto tail, and confined around their mean value. Reallocation from poorer to richer produces a non-ergodic regime, in which the population splits into two classes, characterised by positive and negative wealths which diverge away from the mean. If the reallocation rate is zero, RGBM reduces to GBM, in which individual wealths grow independently and no social structure is represented.

2.1 Model positioning and definition

We use the literature review of De Nardi et al. (2015) to place our model in the context of existing work. They introduce "a simple accounting framework, originally due to Meade (1964)" (p. 9):

> "At birth, individuals are endowed with a, possibly individual-specific, fraction of the contemporaneous, aggregate capital stock. The only source of income in the economy is an idiosyncratic rate of return on individual wealth.
>
> In a given period, individual wealth, normalised by the average capital stock, accrues at the exponential rate
>
> $$r_t^i - g + s_t^i, \tag{1}$$

where r_t^i is the realised rate of return for individual i of current age t and s_t^i is the ratio between the individual's flow of (dis)saving and her wealth at the beginning of the period."

Our model is similar. We denote by $x_i(t)$ the wealth at time t of the i^{th} member of a population of N individuals. We define

$$y_i(t) \equiv \frac{x_i(t)}{\langle x(t) \rangle_N} \tag{2}$$

as the same individual's wealth rescaled by the population average,

$$\langle x(t) \rangle_N \equiv \frac{1}{N} \sum_{i=1}^{N} x_i(t). \tag{3}$$

Equation (1) implies the following differential equation for the rescaled wealth:

$$dy_i = y_i \left[r_i(t) - g + s_i(t) \right] dt. \tag{4}$$

Meade (1964) assumes that the population-average wealth grows exponentially at rate g. Under this assumption, individual wealth obeys

$$dx_i = x_i \left[r_i(t) + s_i(t) \right] dt. \tag{5}$$

The sum $r_i(t) + s_i(t)$ is the time-varying exponential growth rate of individual wealth. Meade (1964) interprets this in terms of idiosyncratic rates of return and saving. We offer a simpler interpretation by decomposing $\left[r_i(t) + s_i(t) \right] dt$ into:

- μdt – a common, constant part, representing economic growth; and
- $\sigma dW_i(t)$ – an idiosyncratic, random part, representing individual growth.

Taking $dW_i(t)$ to be the increment in a Wiener process (normally distributed with zero mean and variance dt) we arrive at GBM,

$$dx_i = x_i \left[\mu dt + \sigma dW_i(t) \right], \tag{6}$$

with drift μ and volatility σ. This is our basic model for the evolution of individual wealth in the absence of social structure. Unsurprisingly, Equation (6) is also the most influential model of asset price dynamics in financial economics.

Wealth evolves multiplicatively under GBM. There are no additive changes akin to labor income and consumption. This is unproblematic for large wealths, where additive changes are dwarfed by capital gains. For small wealths, however, wages and consumption are significant. Indeed, empirical distributions exhibit different regularities for low and high wealths (Drăgulescu and Yakovenko, 2001). There are other multiplicative effects at play. For instance, investing in one's health, housing, and education may be closer to a multiplicative investment than to wealth-depleting consumption. Nonetheless, we leave a red flag here to indicate that our model's realism is questionable for small wealths. We treat additive earnings explicitly in Section 6 and find that including them in a less parsimonious model of wealth accumulation does not alter fundamentally our conclusions.

We note that t in our framework denotes time, rather than the age of an agent. As Meade (1964) puts it, our agents "do not marry or have children or die or even grow old" (p. 41). Therefore, the individual in our setup is best imagined as a household or a family, *i.e.* some long-lasting unit into which personal events are subsumed.

Finally, we add to the system a term that is absent from many other models in the literature. This term represents redistributive social structure. We imagine, as the simplest social interaction, that each individual pays a fixed proportion of its wealth, $x_i \tau dt$, into a central pot ("contributes to society") and gets back an equal share of the pot, $\langle x \rangle_N \tau dt$, ("benefits from society"):

$$dx_i = x_i \left[(\mu - \tau) dt + \sigma dW_i(t) \right] + \langle x \rangle_N \tau dt. \tag{7}$$

A more complex model would treat the economy as a system of agents that interact with each other through a network of relationships. These relationships include trade in goods and services, employment, paying taxes, using centrally-organised infrastructure (roads, schools, a legal system, social security, scientific research, and so on), insurance schemes, wealth transfers through inheritance and gifts, and everything else that constitutes an economic network. It would be a hopeless task to produce an exhaustive list of all these interactions, let alone include them as model components. Instead we introduce a single parameter – the reallocation rate, τ – to represent their net effect. If τ is positive, the direction of net reallocation is from richer to poorer; if negative, it is from poorer to richer.

Fitting τ to data will allow us to answer questions such as:

- what is the net reallocating effect of socio-economic structure on the wealth distribution?
- are observations consistent with the ergodic hypothesis that the rescaled wealth distribution converges to a stationary distribution?
- if so, how long does it take, after a change in conditions, for the rescaled wealth distribution to reach the stationary distribution?

To be clear, we do not want to separate different processes that affect the wealth distribution, as would be useful if we wanted to understand the effect of a specific process, say income tax, on economic inequality. We aim for the opposite – a model that summarises in as few parameters as possible everything that affects the wealth distribution, crucially without assuming ergodicity. For example, if effects of death and inheritance were separated out explicitly with additional parameters, rather than being included in τ, the fitted τ values would no longer tell us if the system as a whole is in an ergodic regime.

2.2 Model behaviour

Equation (7) is our model for the evolution of wealth with social structure and the basis for the empirical study that follows. It is instructive to write it as

$$dx_i = \underbrace{x_i \left[\mu dt + \sigma dW_i\left(t\right)\right]}_{\text{Growth}} \; \underbrace{- \;\; \tau(x_i - \langle x \rangle_N)dt}_{\text{Reallocation}} . \tag{8}$$

This can be thought of as GBM with a mean-reverting term like that of Uhlenbeck and Ornstein (1930) in physics and Vasicek (1977) in finance. This representation exposes the importance of the sign of τ. We discuss the two regimes in turn.

Positive τ For $\tau > 0$, wealth, x_i, reverts to the population average, $\langle x \rangle_N$. The large-sample approximation, $\langle x\left(t\right) \rangle_N \propto e^{\mu t}$, is valid[1] and yields a simple differential equation for the rescaled wealth,

$$dy_i = y_i \sigma dW_i\left(t\right) - \tau(y_i - 1)dt, \tag{9}$$

in which the common growth rate, μ, has been scaled out. The distribution of $y_i\left(t\right)$ can found by solving the corresponding Fokker-Planck equation (also known as the Kolmogorov forward equation). A stationary distribution exists with a Pareto tail, see Appendix B. It is known as the inverse gamma distribution and has probability density function,

$$\mathcal{P}\left(y\right) = \frac{\left(\zeta - 1\right)^{\zeta}}{\Gamma\left(\zeta\right)} e^{-\frac{\zeta - 1}{y}} y^{-(1+\zeta)}, \tag{10}$$

where $\zeta = 1 + (2\tau/\sigma^2)$ is the Pareto tail index, $\Gamma\left(\cdot\right)$ is the gamma function, and the index i has been dropped. Example forms of the stationary distribution are shown in Figure 1. The usual stylised facts are recovered: the larger σ (more randomness in the returns) and the smaller τ (less social cohesion), the smaller the tail index and the fatter the tail of the

[1] Strictly speaking, the large-sample approximation and resulting rescaled-wealth process, Equation (9), hold only for $\tau > \tau_c$. However, $\tau_c \approx 0$ for realistic model parameters and fits to data do not allow us to distinguish it from zero. Nonetheless, the derivation of τ_c is instructive, see Appendix A.

distribution. Moreover, the fitted τ values we obtain in Section 4 give typical ζ values between 1 and 2 for the different datasets analyzed, consistent with observed tail indices between 1.2 to 1.6 (Klass et al., 2006; Gabaix, 2009; Brzezinski, 2014; Vermeulen, 2017). Thus, not only does RGBM predict a realistic functional form for the distribution of rescaled wealth, but also it admits fitted parameters which match observed tail thicknesses. The inability to do the latter is a known shortcoming in models of earnings-based wealth accumulation, see Section 6.

Equation (9) and extensions of it have received much attention in statistical mechanics and econophysics (Bouchaud and Mézard, 2000; Bouchaud, 2015b). As a combination of GBM and an Ornstein-Uhlenbeck process, it is a simple and analytically tractable stochastic process. Liu and Serota (2017) provide an overview of the literature and known results.

Negative τ For $\tau < 0$ the model exhibits mean repulsion rather than reversion. The ergodic hypothesis is invalid and no stationary wealth distribution exists. The population splits into those above the mean and those below the mean. Whereas in RGBM with non-negative τ it is impossible for wealth to become negative, negative τ leads to negative wealth. No longer is total economic wealth a limit to the wealth of the richest individual because the poorest develop large negative wealth. The wealth of the rich in the population increases exponentially away from the mean, and the wealth of the poor becomes negative and exponentially large in magnitude, see Figure 2. Qualitatively, this echoes the findings that the rich are experiencing higher growth rates of their wealth than the poor (Piketty, 2014; Wolff, 2014) and that the cumulative wealth of the poorest 50 percent of the American population was negative during 2008–2013 (Rios-Rull and Kuhn, 2016; The World Wealth and Income Database, 2016).

Such splitting of the population is a common feature of non-ergodic processes. If rescaled wealth were an ergodic process, then individuals would, over long enough time, experience all parts of its distribution. People would spend 99 percent of their time as “the 99 percent” and

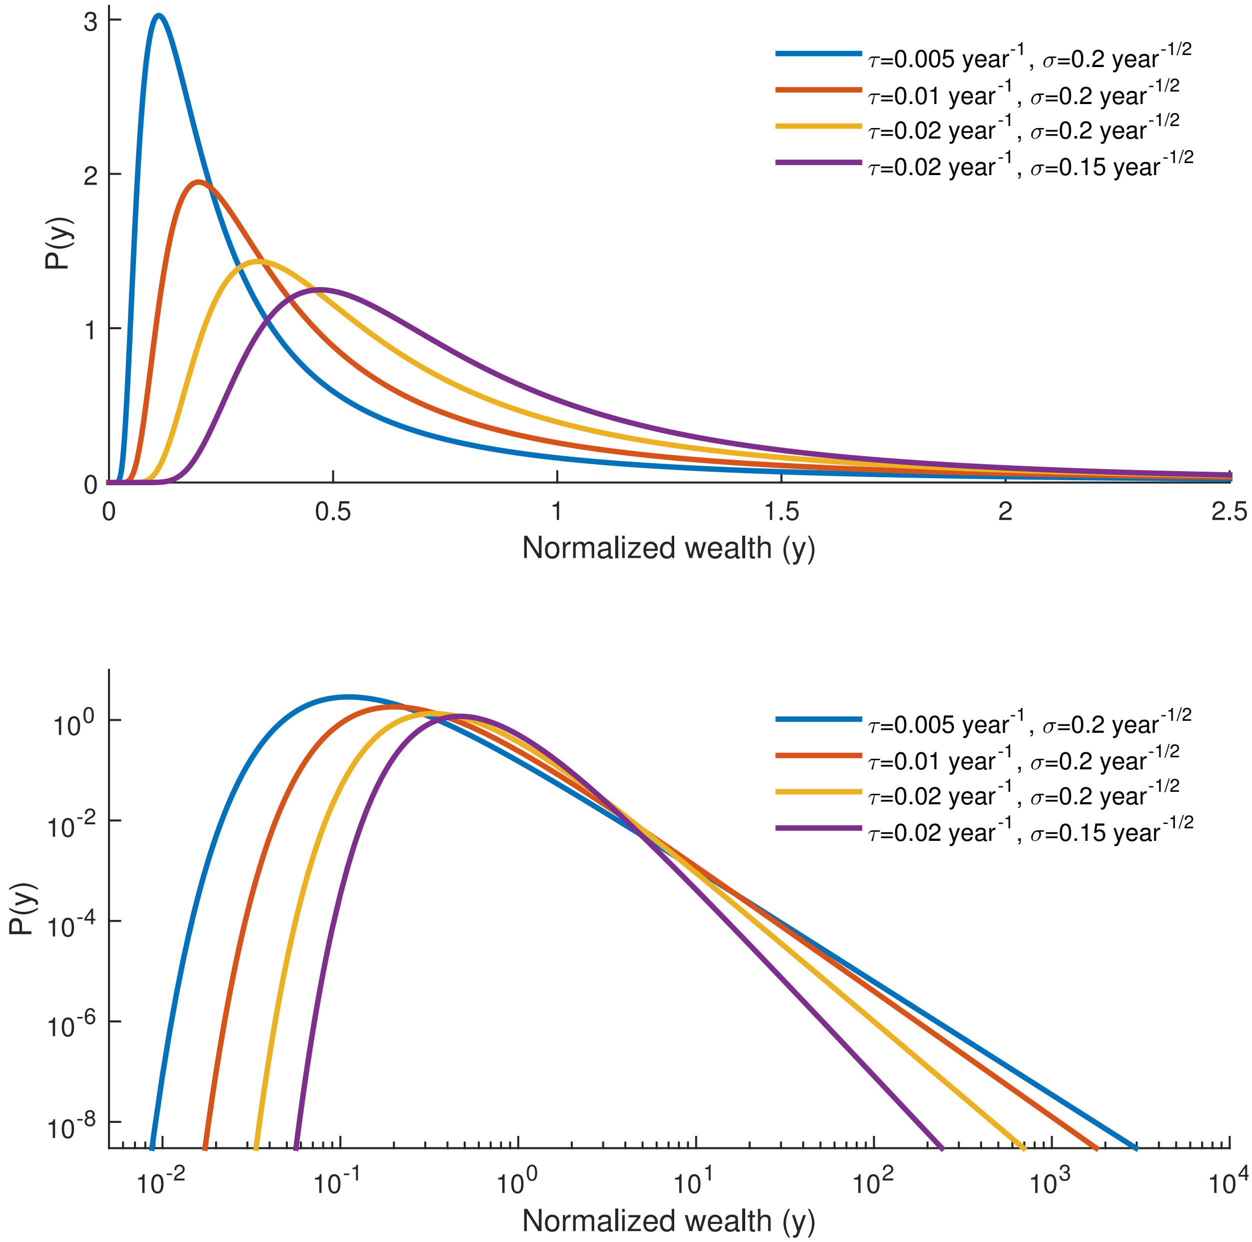

Figure 1: The stationary distribution for RGBM with positive τ. Top – linear scales; Bottom – logarithmic scales.

1 percent of their time as "the 1 percent". The social mobility that is, therefore, implicit in models that assume ergodicity might not exist in reality if that assumption is invalid. That inequality and immobility have been linked (Corak, 2013; Liu et al., 2013; Berman, 2017) may be unsurprising when both are viewed as consequences of non-ergodic wealth or income.

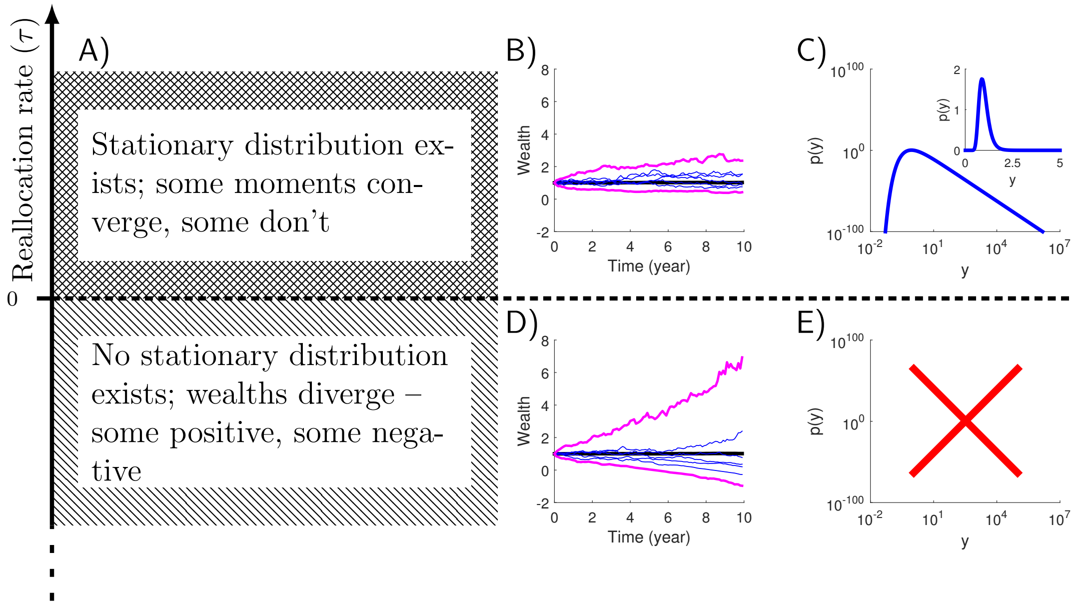

Figure 2: Regimes of RGBM. A) $\tau = 0$ separates the two regimes of RGBM. For $\tau > 0$, a stationary wealth distribution exists. For $\tau < 0$, no stationary wealth distribution exists and wealths diverge. B) Simulations of RGBM with $N = 1000$, $\mu = 0.021$ year^{-1} (presented after rescaling by $e^{\mu t}$), $\sigma = 0.14$ year$^{-1/2}$, $x_i(0) = 1$, $\tau = 0.15$ year^{-1}. Magenta lines: largest and smallest wealths, blue lines: five randomly chosen wealth trajectories, black line: sample mean. C) The stationary distribution to which the system in B) converges. Inset: same distribution on linear scales. D) Similar to B), with $\tau = -0.15$ year^{-1}. E) in the $\tau < 0$ regime, no stationary wealth distribution exists.

3 Data

3.1 Wealth share data

We analyze the wealth shares of the top quantiles of the US population, as estimated by three sources using different methods:

- The income tax method ("capitalization method") that uses information on capital income from individual income tax returns to estimate the underlying stock of wealth (Saez and Zucman, 2016; The World Wealth and Income Database, 2016). "If we can observe capital income $k = rW$, where W is the underlying value of an asset and r is the known rate of return, then we can estimate wealth based on capital income and capitalization factor $1/r$ defined using the appropriate choice of rate of return" (Kopczuk, 2015, p. 54). Data availability: the wealth shares of the top 5, 0.5, 0.1 and 0.01 percent for 1917–2012 and of the top 10 and 1 percent for 1913–2014 (annually).

- The estate multiplier method that uses data from estate tax returns to estimate wealth for the upper tail of the wealth distribution (Kopczuk and Saez, 2004). "The basic idea is to think of decedents as a sample from the living population. The individual-specific mortality rate m_i becomes the sampling rate. If m_i is known, the distribution for the living population can be simply estimated by reweighting the data for decedents by inverse sampling weights $1/m_i$, which are called 'estate multipliers' " (Kopczuk, 2015, p. 53). Data availability: the wealth shares of the top 1, 0.5, 0.25, 0.1, 0.05 and 0.01 percent for 1916–2000 (annually, with several missing years).

- The survey-based method that uses data from the Survey of Consumer Finances (SCF) conducted by the Federal Reserve, plus defined-benefit pension wealth, plus the wealth of the members of the Forbes 400 (Bricker et al., 2016). Data availability: the wealth shares of the top 1 and 0.1 percent for 1989–2013 (for every three years).

These sources are based on different datasets and for different time periods. In the overlapping periods, they sometimes report markedly different wealth share estimates (see Figure 3).

Kopczuk (2015) reviewed the advantages and disadvantages of the different methods (see also the comment by Kopczuk on Bricker et al. (2016)). He observed that "the survey-based and estate tax methods suggest that the share of wealth held by the top 1 percent has not increased much in recent decades, while the capitalization method suggests that it has" (Kopczuk, 2015, p. 48).

Which method best reflects the recent trends in wealth inequality is a matter of ongoing debate. Each method suffers from bias. For example, the survey-based method suffers from

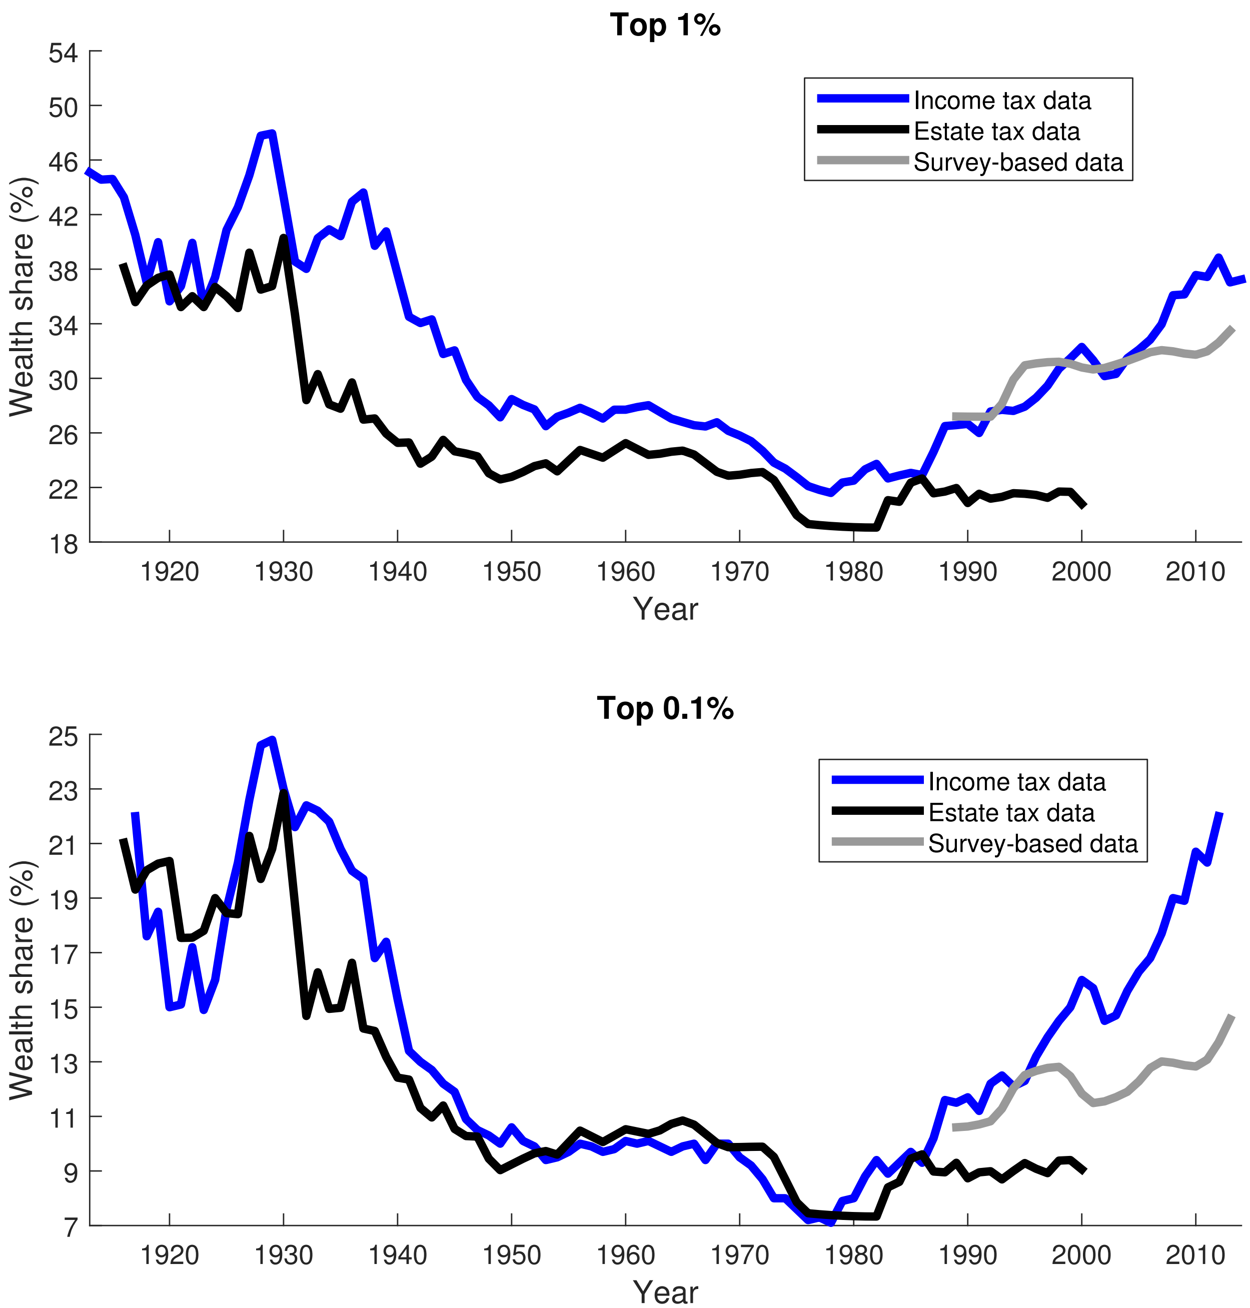

Figure 3: The top wealth shares in the US, 1913–2014. Sources – Saez and Zucman (2016); The World Wealth and Income Database (2016) (blue); Kopczuk and Saez (2004) (black); Bricker et al. (2016) (grey).

some underrepresentation of families who belong to the top end of the distribution. The income tax method suffers from some practical difficulties – "not all categories of assets generate capital income that appears on tax returns. [...] Owner-occupied housing does not generate annual taxable capital income" (Kopczuk, 2015, p. 54). The estate tax method suffers from the need to accurately estimate mortality rates for the wealthy, known to be lower than those for the rest of the population. We refer the reader to Kopczuk (2015); Bricker et al. (2016) for a thorough discussion. We analyze each data source separately.

3.2 Wealth Growth Rate

We find numerically that the results of our analysis do not depend on μ. This is because wealth shares depend only on the distribution of rescaled wealth and, for $\tau > 0$, it is possible to scale out μ completely from the wealth dynamic to obtain Equation (9) for rescaled wealth. The fitted $\tau < 0$ values we find are not large or persistent enough to make our simulations significantly μ-dependent. However, formally, since we allow negative τ, we must simulate Equation (7) and not Equation (9). This requires us to specify a value of μ, which we estimate as $\mu = 0.021 \pm 0.001$ year^{-1} by a least-squares fit of historical per-capita private wealth in the US (Piketty and Zucman, 2014) to an exponential growth curve.

3.3 Volatility

We must also specify the volatility parameter, σ, in Equation (7). In principle, this can vary with time. We have no access to real individual wealth trajectories, so we resort to estimating $\sigma(t)$ from other data. We find numerically that our results are not very sensitive to the details, so we need only a good "ballpark" estimate. We obtain that by assuming that the volatility in individual wealths tracks the volatility in the values of the companies that constitute the commercial and industrial base of the national economy. Therefore, for each year, we estimate $\sigma(t)$ as the standard deviation of daily logarithmic changes of the Dow Jones Industrial Average (Quandl, 2016), which we annualise by multiplying by $(250/\text{year})^{1/2}$. The values usually lie between 0.1 and 0.2 year$^{-1/2}$, with an average of 0.16 year$^{-1/2}$. Running our empirical analysis with constant σ in this range had little effect on our results (see Appendix C) so, for simplicity, we present the analysis using $\sigma(t) = 0.16$ year$^{-1/2}$ for all t.

Fitting σ to stock market data means that we have only one model parameter – the effective reallocation rate, $\tau(t)$ – to fit to the historical wealth shares.

4 Empirical Analysis

The goal of the empirical analysis is to estimate $\tau(t)$ from the historical wealth data, using RGBM as our model. This estimation allows us to address two main questions:

1. Is it valid to assume ergodicity for the dynamics of relative wealth in the US? For the ergodic hypothesis to be valid, fitted values of $\tau(t)$ would have to be robustly positive.

2. If $\tau(t)$ is indeed positive, how long does it take for the distribution to converge to its asymptotic form?

We fit a time series, $\tau(t)$, that reproduces the annually observed wealth shares in the three datasets (see Section 3): Income tax-based (Saez and Zucman, 2016; The World Wealth and Income Database, 2016), estate tax-based (Kopczuk and Saez, 2004) and survey-based (Bricker et al., 2016). The wealth share, S_q, is defined as the proportion of total wealth, $\sum_i^N x_i$, owned by the richest fraction q of the population, *e.g.* $S_{10\%} = 80$ percent means that the richest 10 percent of the population own 80 percent of the total wealth.

For an empirical wealth share time series, $S_q^{\mathrm{data}}(t)$, we proceed as follows.

– Step 1 Initialise N individual wealths, $\{x_i(t_0)\}$, as random variates of the inverse gamma distribution with parameters chosen to match $S_q^{\mathrm{data}}(t_0)$.

– Step 2 Propagate $\{x_i(t)\}$ according to Equation (7) over Δt, using the value of τ that minimises the difference between the wealth share in the modelled population, $S_q^{\mathrm{model}}(t+\Delta t, \tau)$, and $S_q^{\mathrm{data}}(t+\Delta t)$. We use the Nelder-Mead algorithm (Nelder and Mead, 1965).

– Step 3 Repeat Step 2 until the end of the time series.

We consider historical wealth shares of the richest $q = 10$, 5, 1, 0.5, 0.25, 0.1, 0.05 and 0.01 percent and obtain time series of fitted effective reallocation rates, $\tau_q(t)$, shown in Figure 4.

For each value of q we perform a run of the simulation for $N = 10^8$. Since in practice dW is randomly chosen, each run of the simulation will result in slightly different $\tau_q(t)$ values. However, we found that the differences between such calculations are negligible.

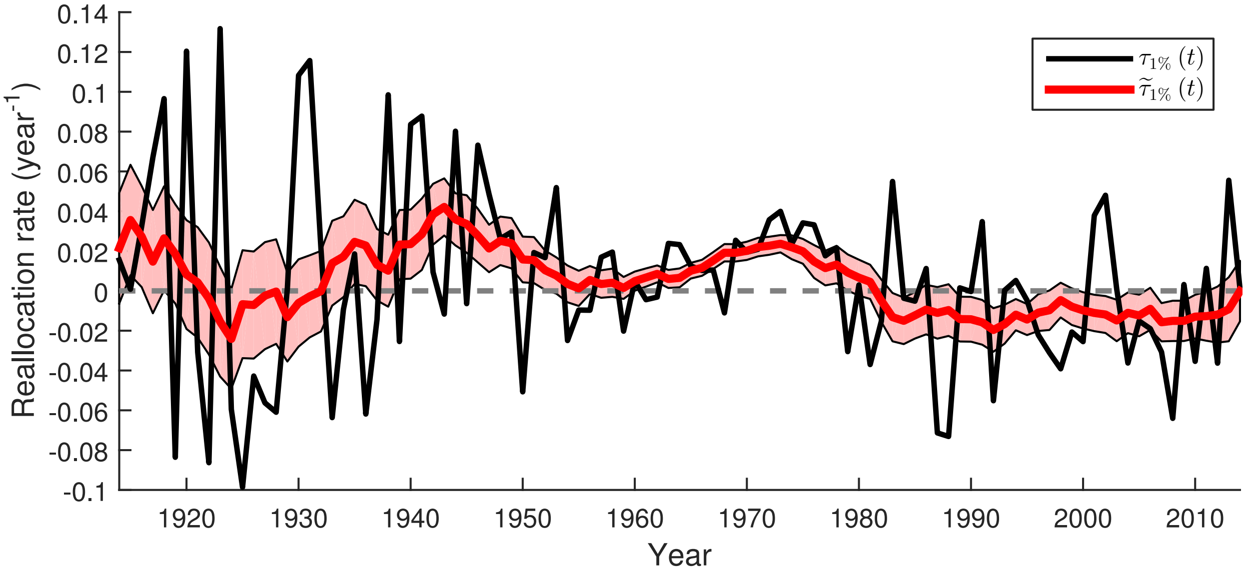

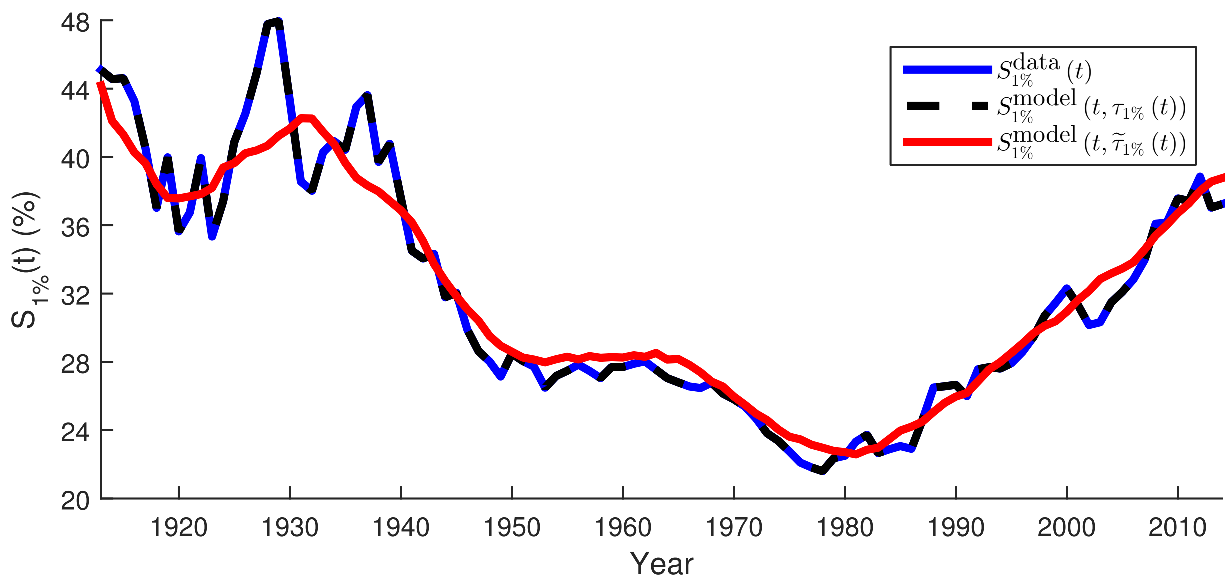

Figure 4: Fitted effective reallocation rates. Calculations done using $\mu = 0.021$ year^{-1} and $\sigma = 0.16$ year$^{-1/2}$. Top: $\tau_{1\%}(t)$ (black) and $\widetilde{\tau}_{1\%}(t)$ (red). Translucent envelopes indicate one standard error in the moving averages. Bottom: $S_{1\%}^{\mathrm{DATA}}$ (blue), $S_{1\%}^{\mathrm{MODEL}}$ based on the annual $\tau_{1\%}(t)$ (dashed black), based on the 10-year moving average $\widetilde{\tau}_{1\%}(t)$ (red).

Figure 4 (top) shows large annual fluctuations in $\tau_q(t)$. We are interested in longer-term changes in reallocation driven by structural economic and political changes. To elucidate these we smooth the data by taking a central 10-year moving average, $\widetilde{\tau}_q(t)$, where the window is truncated at the ends of the time series. To ensure the smoothing does not introduce artificial biases, we reverse the procedure and use $\widetilde{\tau}_q(t)$ to propagate the initially inverse gamma-distributed $\{x_i(t_0)\}$ and determine the wealth shares $S_q^{\text{model}}(t)$. The good agreement with $S_q^{\text{data}}(t)$ suggests that the smoothed $\widetilde{\tau}_q(t)$ is meaningful, see Figure 4 (bottom).

For the income tax method wealth shares (Saez and Zucman, 2016), the effective reallocation rate, $\widetilde{\tau}(t)$, has been negative – *i.e.* from poorer to richer – since the mid-1980s. This holds for all of the inequality measures we derived from this dataset.

For the survey-based wealth shares (Bricker et al., 2016), we observe briefer periods in which $\widetilde{\tau}(t) < 0$. The same is true for the estate tax data (Kopczuk and Saez, 2004), see Figure 5. When $\tau(t)$ is positive, relevant convergence times are very long compared to the time scales of policy changes, namely at least several decades.

All three datasets indicate that making the ergodic hypothesis is an unwarranted restriction on models and analyses. The hypothesis makes it impossible to observe and reason about the most dramatic qualitative features of wealth dynamics, such as rising inequality, negative reallocation, negative wealth, and social immobility.

5 Convergence times

In the ergodic regime it is possible to calculate how fast the wealth shares of different quantiles converge to their asymptotic value. We do this numerically. Starting with a population of equal wealths and assuming $\mu = 0.021$ year^{-1}, $\sigma = 0.16$ year$^{-1/2}$, and $\tau = 0.04$ year^{-1}, we let the system equilibrate for 3000 years, long enough for the distribution to reach its asymptotic form to numerical precision. We then create a "shock", by changing τ to a different "shock value", and allow the system to equilibrate again for 3000 years, see

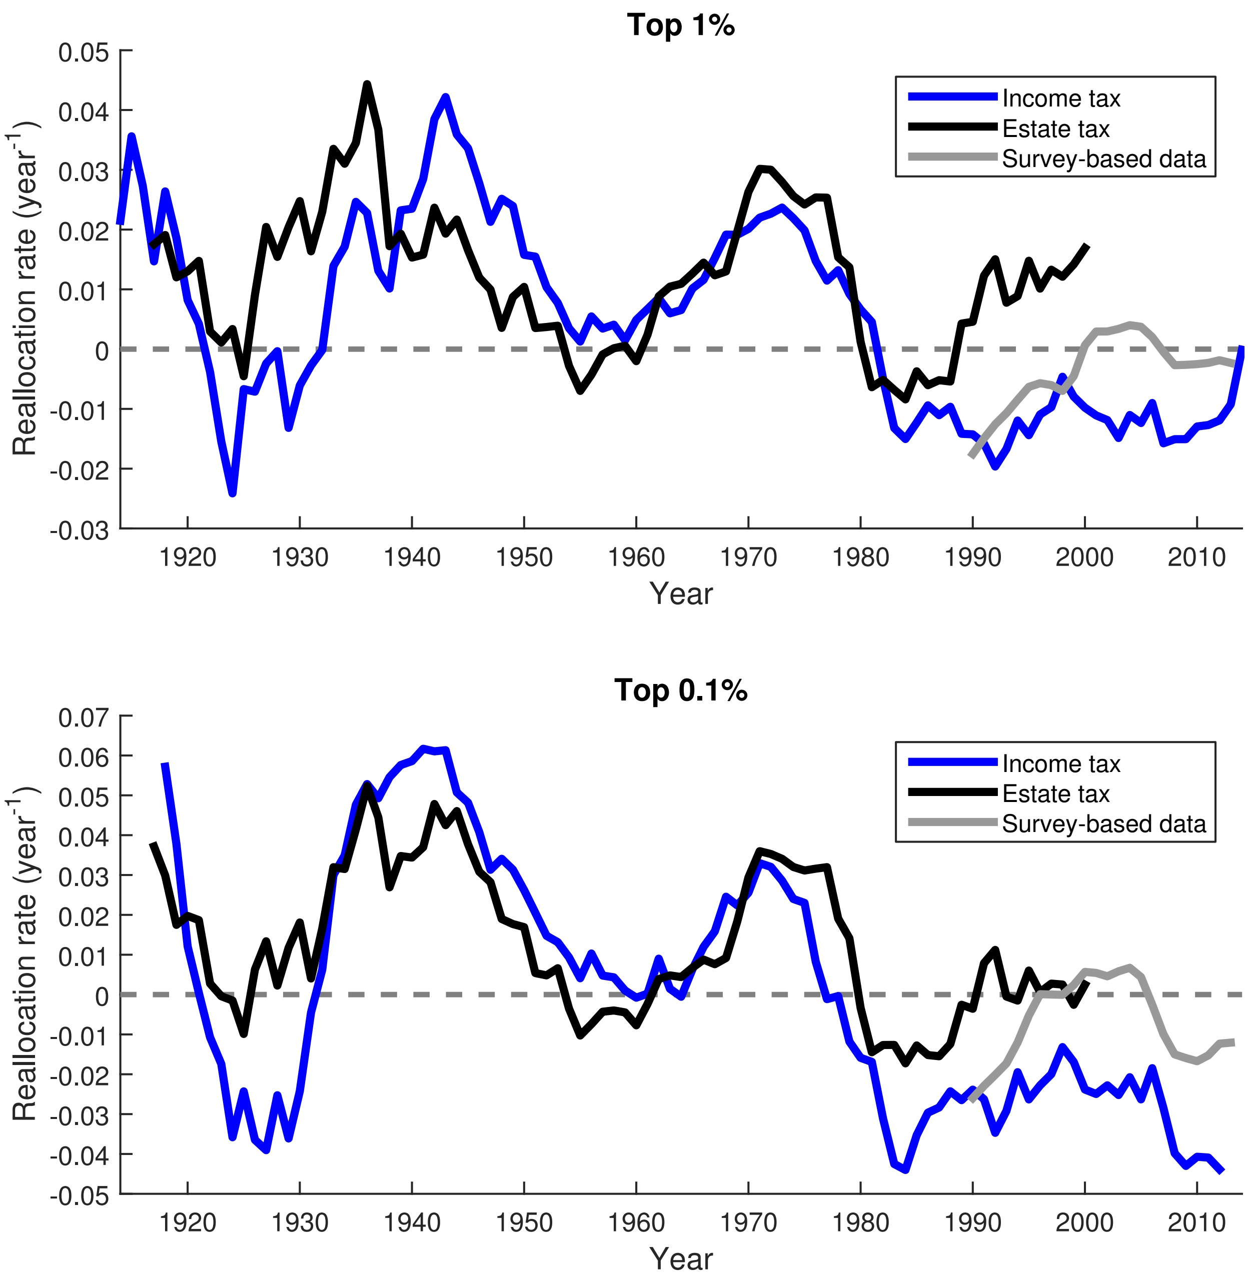

Figure 5: Effective reallocation rates for different datasets.

top panel of Figure 6. Following the shock, the wealth shares converge to their asymptotic values. We fit this convergence numerically with an exponential function and interpret the inverse of the exponential convergence rate as the convergence time. The bottom panel of Figure 6 shows the convergence times versus the shock value of τ.

In addition, it is possible to calculate the convergence time of the variance of the sta-

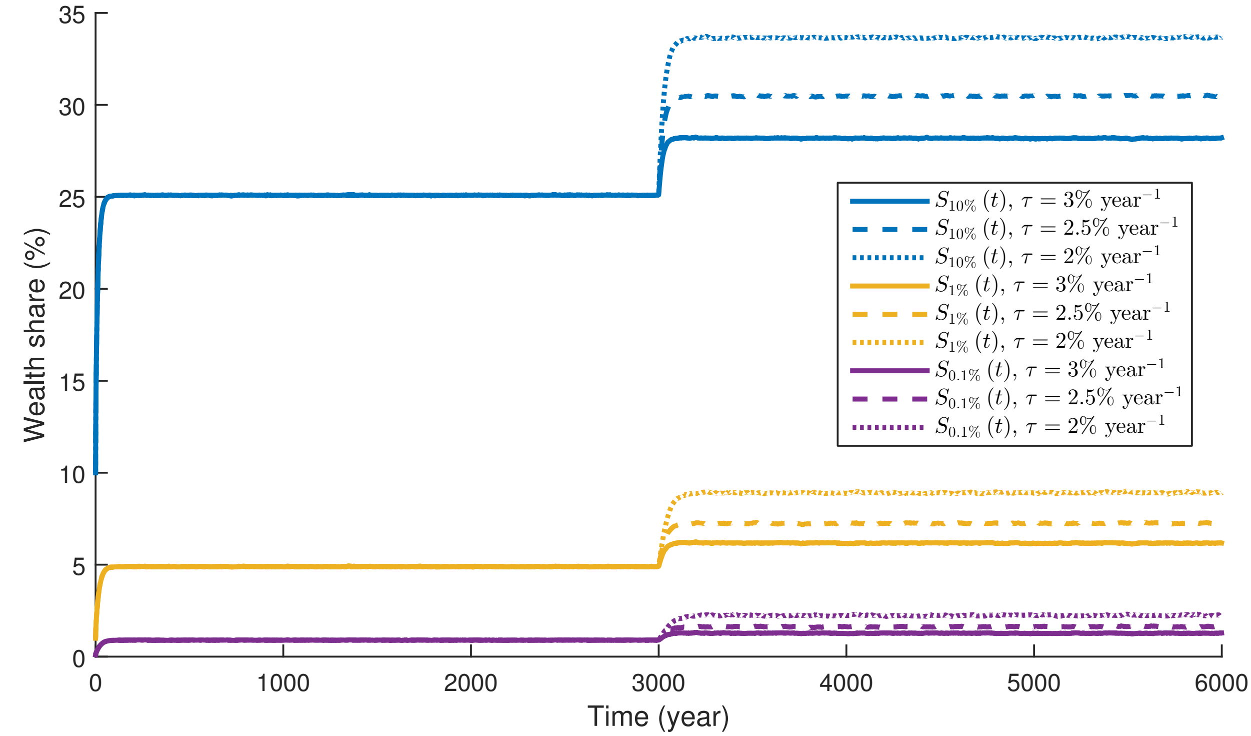

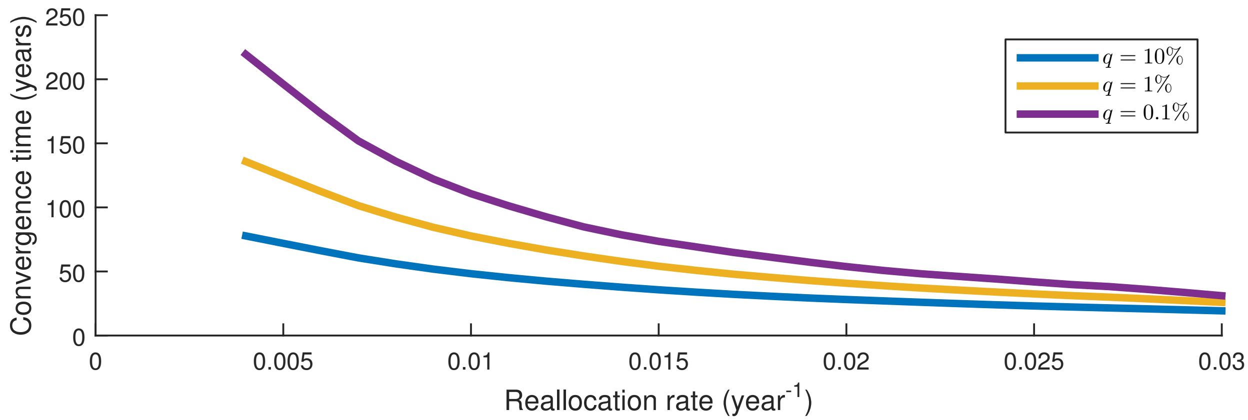

Figure 6: Wealth share convergence time. Top: The convergence of the wealth share for $q = 10$ percent (blue), $q = 1$ percent (yellow) and $q = 0.1$ percent (purple) following a change in the value of τ from 0.04 year^{-1} to 0.03 year^{-1} (solid), 0.025 year^{-1} (dashed) and 0.02 year^{-1} (dotted). Bottom: The wealth share exponential convergence time for $q = 10$ percent (blue), $q = 1$ percent (yellow) and $q = 0.1$ percent (purple) as a function of τ.

tionary distribution (and other cumulants and moments of interest). In the ergodic regime the stationary distribution has a finite variance only if $\tau > \sigma^2/2$ (Liu and Serota, 2017). Convergence of the actual variance to the stationary variance occurs exponentially over a timescale $1/(2\tau - \sigma^2)$. Figure 7 shows the convergence times for different values of σ. See

Appendix D for more details.

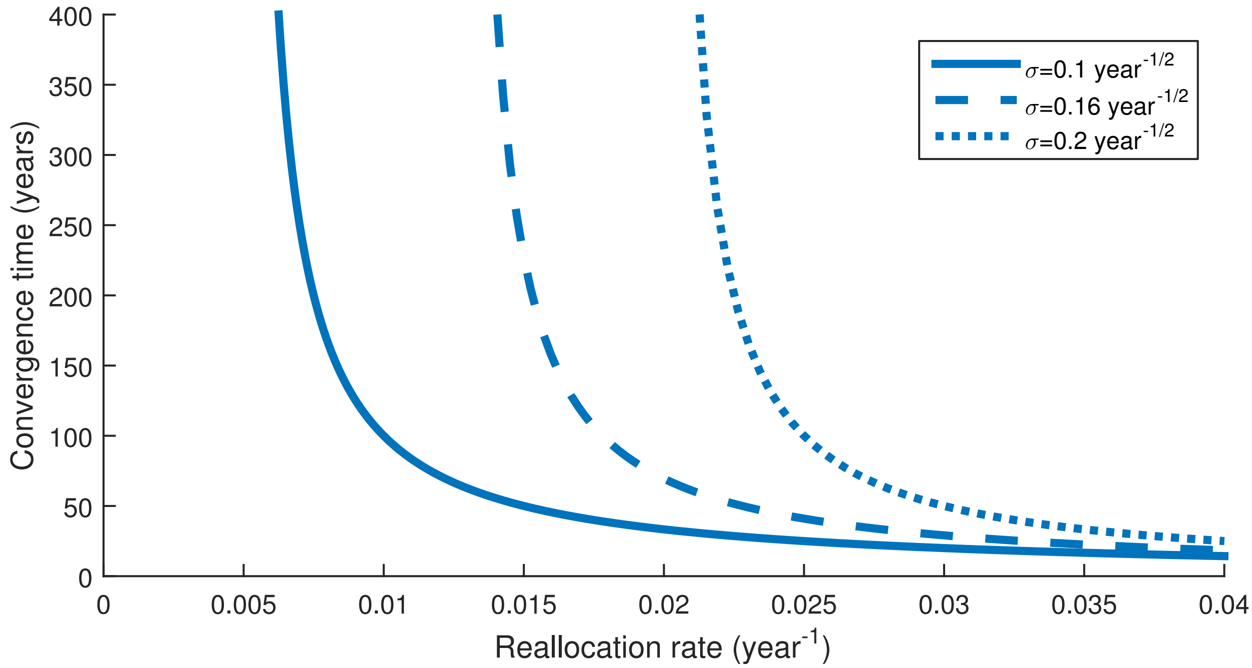

Figure 7: VARIANCE CONVERGENCE TIME

Convergence times for wealth shares and variance are long, ranging from a few decades to several centuries. This implies that empirical studies which assume ergodicity and fast convergence will be inconsistent with the data. To test this, we simulate such a study by performing a different RGBM parameter fit. We find the reallocation rates, $\tau_q^{\text{eqm}}(t)$, that generate stationary distributions consistent with observed wealth shares. In other words, we assume instantaneous convergence.

Figure 8 contrasts $\tau_{1\%}^{\text{eqm}}(t)$ assuming ergodicity with $\widetilde{\tau}_{1\%}(t)$ without assuming ergodicity (using the income tax method dataset). If convergence were always possible and fast, then the two values would be identical within statistical uncertainties. They are not. In addition, the generally large discrepancies between the wealth inequality implied by $\tau_{1\%}^{\text{eqm}}(t)$ (bottom panel, Figure 8, green line) and as observed (bottom panel, Figure 8, blue line) indicate that the wealth distribution does not stay close to its asymptotic form. This means that the long convergence times we calculate are a practical methodological problem for conventional studies.

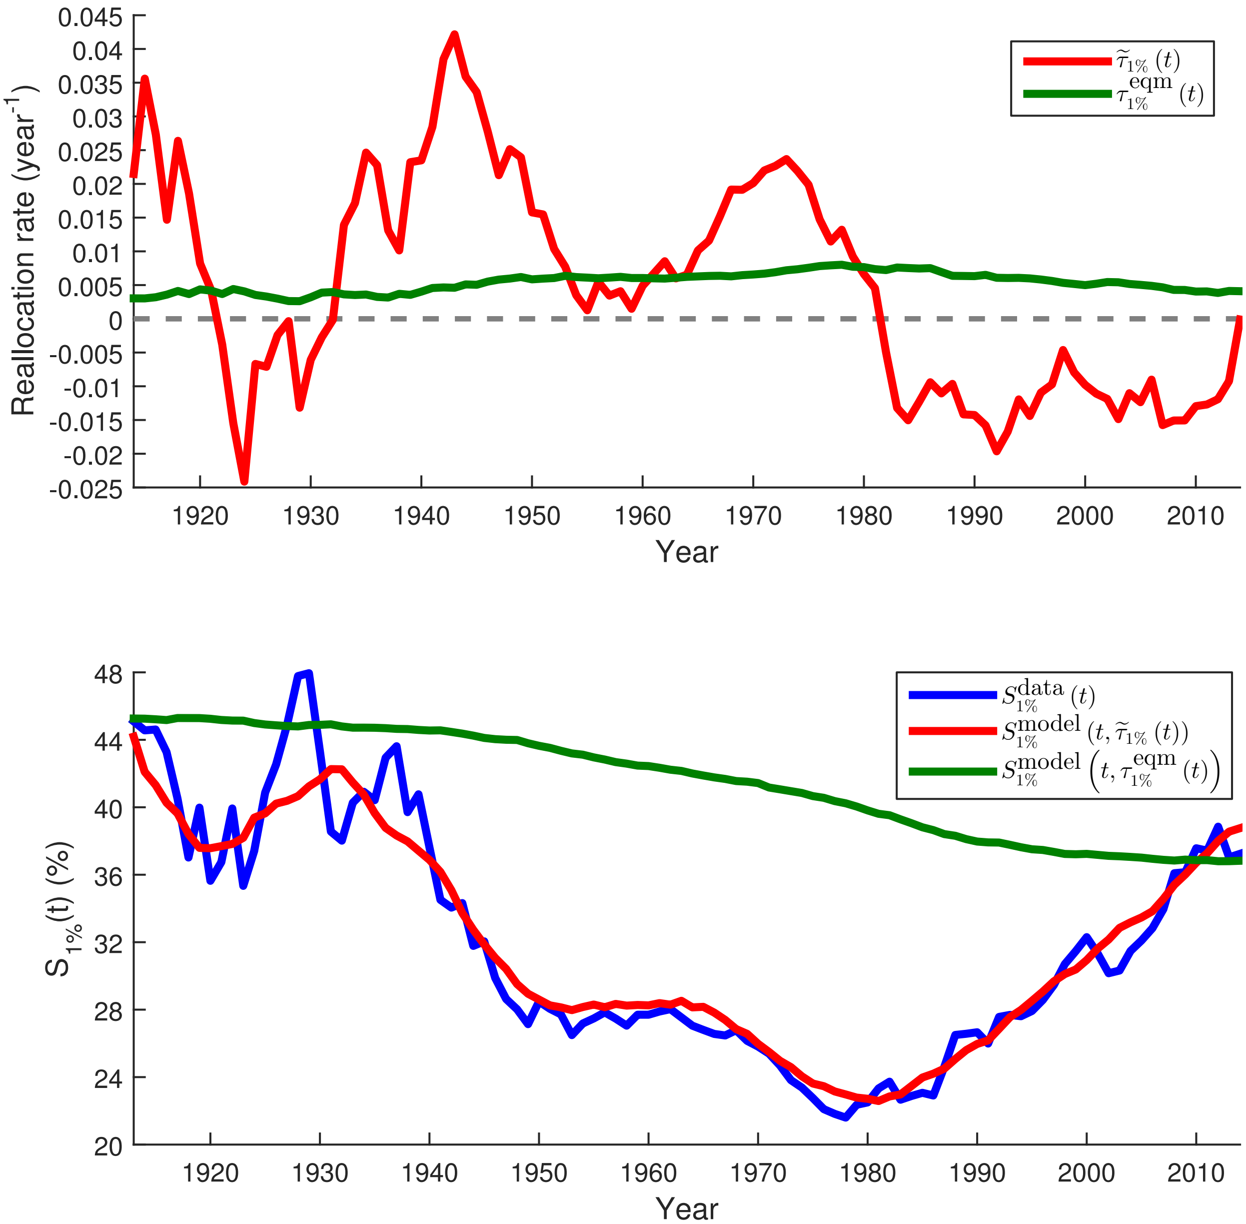

Figure 8: Comparison of dynamic and equilibrium reallocation rates. Top: $\widetilde{\tau}_{1\%}(t)$ (red, same as in the top of Figure 4). $\tau_{1\%}^{\mathrm{EQM}}(t)$ (green), defined such that $\lim_{t' \to \infty} S_{1\%}^{\mathrm{MODEL}}\left(t', \tau_{1\%}^{\mathrm{EQM}}(t)\right) = S_{1\%}^{\mathrm{DATA}}(t)$. It is impossible by design for this value to be negative. The significant difference between the red and green lines demonstrates that the fast convergence assumption is invalid for the problem under consideration. Bottom: $S_{1\%}^{\mathrm{DATA}}$ (blue), $S_{1\%}^{\mathrm{MODEL}}$ based on the 10-year moving average $\widetilde{\tau}_{1\%}(t)$ (red), based on $\tau_{1\%}^{\mathrm{EQM}}(t)$ (green). The reallocation rates found under the fast convergence assumption generate model wealth shares which bear little relation to reality.

6 The effect of earnings

In a recent review, Benhabib et al. (2017, p. 593) remark that "the literature has largely emphasised the role of earnings inequality in explaining wealth inequality" and point to empirical failures of this approach. The most popular models, in which agents accumulate wealth through stochastic earnings and precautionary savings, predict a positive correlation between earnings and wealth inequality, which is not observed in cross-country data, and struggle "to reproduce the thick right tail of the wealth distribution observed in the data" (Benhabib et al., 2017, p. 593). The latter failure is also noted by Hubmer et al. (2016). Benhabib et al. (2017, p. 595) conclude that "other factors, like stochastic idiosyncratic returns on wealth" must be at play.

Consistent with these observations, RGBM models wealth accumulation as a multiplicative process. Changes in individual wealth are composed of terms proportional to either individual wealth or average wealth, in effect generating stochastic idiosyncratic returns. Additive changes akin to labor income and consumption are not treated explicitly. Instead their effects are wrapped into the reallocation rate, τ. While this parsimony makes the model tractable – in that we can write down the rescaled wealth distribution, Equation (10), for positive τ – it makes it impossible to disentangle the various real-world drivers of the observed increase in wealth inequality. It is, therefore, legitimate to ask whether wealth is inherently unstable (because it is reallocated negatively) or whether the increase in wealth inequality is really due to changes in the earnings distribution. If the latter, then the negative reallocation found in RGBM would be an artefact of an underspecified model and an unreliable indicator that rescaled wealth is non-ergodic.

To ensure we are not fooling ourselves, we check this by adding to Equation (7) a term representing earnings, for which data are available. We find that earnings have had a small effect on the dynamics of the wealth distribution over the last century. Since about 2000 this may have contributed to the increase in wealth inequality, but in general the effect has been stabilizing. This implies that the purely multiplicative dynamics of wealth, *i.e.* excluding additive earnings, have exhibited greater negative reallocation than fits to RGBM might suggest.

We separate earnings out as follows, in a model we call Earnings Geometric Brownian Motion (EGBM):

$$dx_i = x_i \left[\left(\mu(t) - \tau^{\mathrm{EGBM}}(t) \right) dt + \sigma(t)\, dW_i(t) \right] + \langle x \rangle_N \tau^{\mathrm{EGBM}}(t)\, dt + \epsilon_i(t)\, dt. \tag{11}$$

The notation is similar to RGBM, with the addition of the individual earnings term $\epsilon_i(t)$. We consider earnings after spending and after tax, as they directly change wealth. This changes the role of the reallocation rate, as reflected in the notation: τ^{EGBM} reflects changes in wealth for reasons other than additive income and consumption.

Using the same techniques as for the RGBM fits, we fit a time series, $\tau^{\mathrm{EGBM}}(t)$, that reproduces annually observed wealth shares in the US (for the income tax method dataset). This requires us to find individual earnings in Equation (11), which we do as follows. We assume that the earnings are log-normally distributed (see for example Pinkovskiy and Sala-i-Martin (2009)) and then fit the two parameters of the distribution every year. This is done using the mean earnings (calculated as total disposable national income multiplied by the personal savings rate and divided by the population size, as reported by Piketty and Zucman (2014)) and the top 1% income share data (as reported in The World Wealth and Income Database (2016)).

Since the joint wealth-earnings distribution is not known fully, we analyze two bounding scenarios. First, a "correlated" scenario in which earnings are perfectly correlated with wealth. Every year, the richest individual has the highest earnings, the second richest individual has the second highest earnings, and so on. This is the least stabilizing possibility. Second, an "anti-correlated" scenario in which the poorest individual has the highest earnings, the second poorest individual has the second highest earnings, and so on. This is the most stabilizing possibility. We fit $\tau^{\mathrm{EGBM}}(t)$ for each scenario, and compare the results to the fitted $\tau(t)$ values for which the effects of earnings have not been disentangled from other effects. See Figure 9.

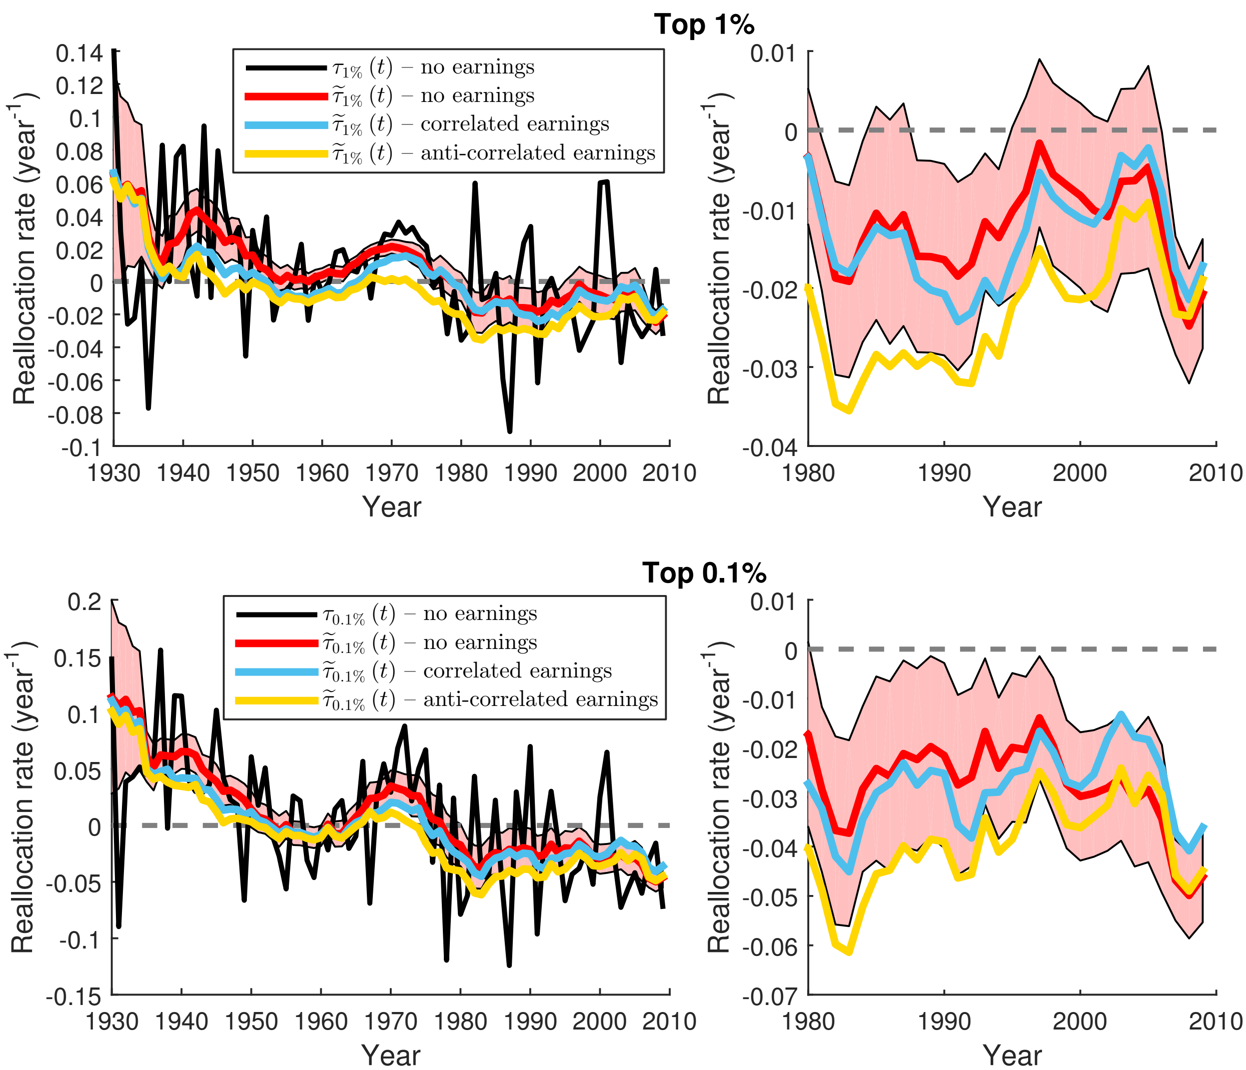

Figure 9: THE EFFECTIVE REALLOCATION RATES FOR THE RGBM AND EGBM MODELS. THE RED TRANSLUCENT ENVELOPES INDICATE ONE STANDARD ERROR IN THE MOVING AVERAGES OF THE RGBM FITTED REALLOCATION RATES.

In general, disentangling earnings has a small effect on the fitted reallocation rates. For the correlated scenario – which is the more realistic of the two because the correlation between wealth and earnings is positive (Rios-Rull and Kuhn, 2016) – the fitted reallocation rates $\tau^{\mathrm{EGBM}}(t)$ are within the statistical error of $\tau(t)$, into which the effects of earnings have been subsumed.

When earnings have a stabilizing effect, the fitted reallocation rate is lower under EGBM than under RGBM. Even in the correlated scenario, which is the least stabilizing, earnings have a stabilizing effect for years 1930–2000. During 2000–2010, we find that earnings destabilise the wealth distribution in the correlated scenario. However, since the difference between the fitted values is small, we can conclude that the contribution of earnings to the reported increase in wealth inequality from 1985 to 2010 is negligible. Moreover, the earnings contribution cannot explain the invalidity of the ergodic hypothesis that constitutes our main finding.

We suspect the main reason for the small effect of earnings is the low ratio of earnings to wealth. Figure 10 shows the ratio of average earnings (after spending) to average wealth in the US during the period under study. Typically this ratio is around 1%, at which level earnings play only a secondary role in the dynamics of wealth. This echoes the observations of Piketty (2014); Piketty and Zucman (2014); Berman and Shapira (2017). Under such conditions, it is unsurprising that models reliant on earnings as the primary mechanism of wealth accumulation fail to resemble reality (Hubmer et al., 2016; Benhabib et al., 2017). By including a multiplicative growth mechanism for wealth, RGBM avoids these pitfalls. Not only does it predict the Pareto tail of the stationary wealth distribution, when it exists, but also it allows for fitted tail exponents that reproduce the observed tail thickness (see Section 2.2). Retreating to the simplicity of RGBM is, therefore, justified.

7 Conclusions

Studies of economic inequality often assume ergodicity of relative wealth. This assumption also goes under the headings of equilibrium, stationarity, or stability (Adamou and Peters, 2016). Specifically, it is assumed that:

1. the system can equilibrate, *i.e.* a stationary distribution exists to which the observed distribution converges in the long-time limit; and

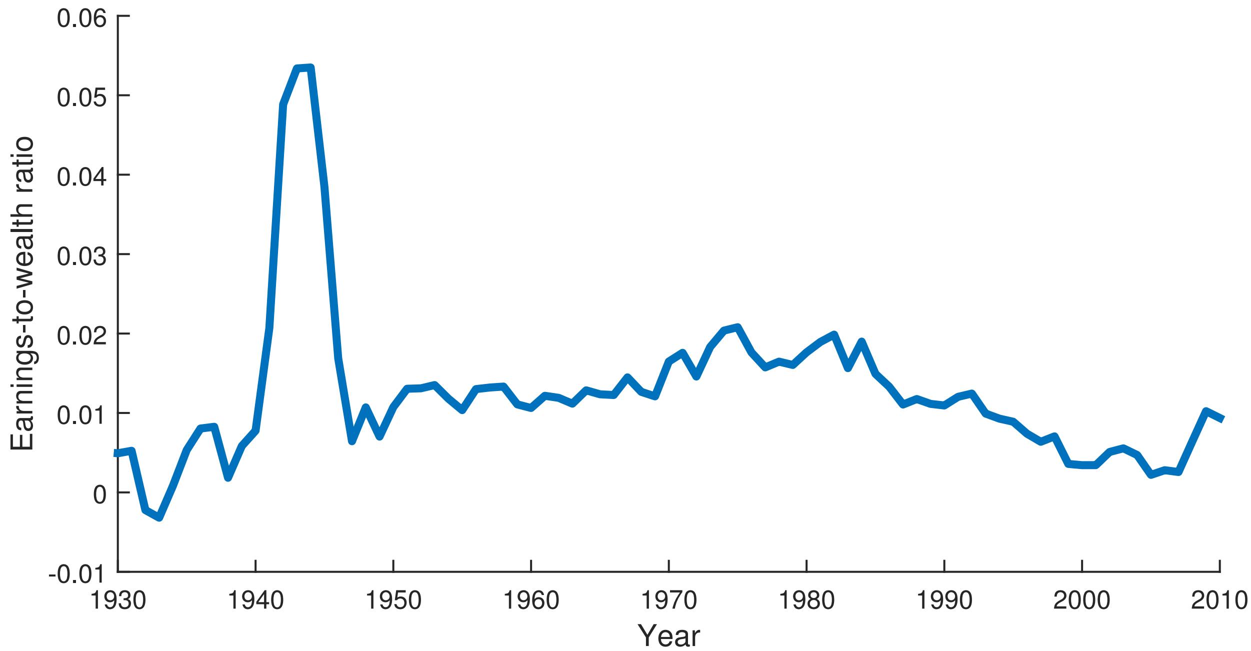

Figure 10: Time series of the ratio of average earnings after spending to average private wealth in the US. Data from Piketty and Zucman (2014).

2. the system equilibrates quickly, *i.e.* the observed distribution gets close to the stationary distribution after a time shorter than other relevant timescales, such as the time between policy changes.

Assumption 2 is often left unstated, but it is necessary for the stationary (model) distribution to resemble the observed (real) distribution. This matters because the stationary distribution is often a key object of study – model parameters are found by fitting the stationary distribution to observed inequality, and effects of various model parameters on the stationary distribution are explored.

We do not assume ergodicity. Fitting τ in RGBM allows the data to speak without constraint as to whether the ergodic hypothesis is valid. We find it to be invalid because:

A. We observe negative τ values in all datasets analyzed, most notably using the income tax method, especially since about 1980. The wealth distribution is non-stationary and inequality increases for as long as these conditions prevail.

B. When we observe positive τ, the associated convergence times are mostly of the order of decades or centuries, see Figure 5 and Figure 6 (bottom). They are much longer than the periods over which economic conditions and policies change – they are the timescales of history rather than of politics.

The ergodic hypothesis precludes what we find. Item A above corresponds to reallocation that moves wealth from poorer to richer individuals, which is inconsistent with the ergodic hypothesis. In this sense the ergodic hypothesis is a set of blindfolds, hiding from view the most dramatic economic conditions. For the most recent data, the system is in a state best described by non-ergodic RGBM, $\tau < 0$ (Saez and Zucman, 2016; The World Wealth and Income Database, 2016) or $\tau \approx 0$ (Bricker et al., 2016). Therefore, each time we observe the wealth distribution, we see a snapshot of it either in the process of diverging or very far from its asymptotic form. It is much like a photograph of an explosion in space: it will show a fireball whose finite extent tells us nothing of the eventual distance between pieces of debris.

We also find that changes in the earnings distribution do not provide an adequate alternative explanation of the described dynamics of the wealth distribution. Although earnings have become more unequal over the recent decades in which wealth inequality has increased, their effect on the wealth distribution has been small and generally stabilizing rather than destabilizing. Treating earnings explicitly in our model does not change fundamentally our conclusions.

The economic phenomena that trouble theorists most – such as diverging inequality, social immobility, and the emergence of negative wealth – are difficult to reproduce in a model that assumes ergodicity. In our simple model, this is easy to see: in the ergodic regime, $\tau > 0$, our model cannot reproduce these phenomena at all. One may be tempted to conclude that their existence is a sign of special conditions prevailing in the real world – collusion and conspiracies. But if we admit the possibility of non-ergodicity, $\tau \leq 0$, it becomes clear that these phenomena can easily emerge in an economy that does not actively guard against them.

References

Adamou, Alexander and Ole Peters, "Dynamics of Inequality," *Significance*, 2016, *13* (3), 32–35.

Atkinson, Anthony B., "The Timescale of Economic Models: How Long is the Long Run?," *The Review of Economic Studies*, 1969, *36* (2), 137–152.

Benhabib, Jess, Alberto Bisin, and Mi Luo, "Earnings Inequality and Other Determinants of Wealth Inequality," *American Economic Review*, 2017, *107* (5), 593–597.

_ , _ , and Shenghao Zhu, "The Distribution of Wealth and Fiscal Policy in Economies with Finitely Lived Agents," *Econometrica: Journal of the Econometric Society*, 2011, *79* (1), 123–157.

Berman, Yonatan, "Understanding the Mechanical Relationship between Inequality and Intergenerational Mobility," 2017. Available at: `http://papers.ssrn.com/abstract=2796563`.

_ and Yoash Shapira, "Revisiting r > g—The Asymptotic Dynamics of Wealth Inequality," *Physica A: Statistical Mechanics and its Applications*, 2017, *467*, 562–572.

_ , Eshel Ben-Jacob, and Yoash Shapira, "The Dynamics of Wealth Inequality and the Effect of Income Distribution," *PLOS ONE*, 2016, *11* (4), e0154196.

Bewley, Truman, "The Permanent Income Hypothesis: A Theoretical Formulation," *Journal of Economic Theory*, 1977, *16* (2), 252–292.

Bouchaud, Jean-Philippe, "Note on Mean-field Wealth Models and the Random Energy Model," 2015. Private communication.

_ , "On Growth-Optimal Tax Rates and the Issue of Wealth Inequalities," *Journal of Statistical Mechanics: Theory and Experiment*, 2015, *2015* (11), P11011.

_ and Marc Mézard, "Wealth Condensation in a Simple Model of Economy," *Physica A: Statistical Mechanics and its Applications*, 2000, *282* (4), 536–545.

Bricker, Jesse, Alice Henriques, Jacob Krimmel, and John Sabelhaus, "Measuring Income and Wealth at the Top Using Administrative and Survey Data," *Brookings Papers on Economic Activity*, 2016, pp. 261–331.

Brzezinski, Michal, "Do Wealth Distributions Follow Power Laws? Evidence from Rich Lists," *Physica A: Statistical Mechanics and its Applications*, 2014, *406*, 155–162.

Corak, Miles, "Income Inequality, Equality of Opportunity, and Intergenerational Mobility," *The Journal of Economic Perspectives*, 2013, *27* (3), 79–102.

Cowell, Frank A., "Piketty in the Long Run," *The British Journal of Sociology*, 2014, *65* (4), 708–720.

Drăgulescu, Adrian and Victor M. Yakovenko, "Exponential and Power-Law Probability Distributions of Wealth and Income in the United Kingdom and the United States," *Physica A: Statistical Mechanics and its Applications*, 2001, *299* (1), 213–221.

Gabaix, Xavier, "Power Laws in Economics and Finance," *Annual Review of Economics*, 2009, *1* (1), 255–294.

_ **, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll**, "The Dynamics of Inequality," *Econometrica: Journal of the Econometric Society*, 2016, *84* (6), 2071–2111.

Hubmer, Joachim, Per Krusell, and Anthony A. Smith Jr., "The Historical Evolution of the Wealth Distribution: A Quantitative-Theoretic Investigation," 2016. Working Paper 23011, NBER.

Jones, Charles I., "Pareto and Piketty: The Macroeconomics of Top Income and Wealth Inequality," *The Journal of Economic Perspectives*, 2015, *29* (1), 29–46.

Kaymak, Barış and Markus Poschke, "The Evolution of Wealth Inequality Over Half a Century: The Role of Taxes, Transfers and Technology," *Journal of Monetary Economics*, 2016, *77*, 1–25.

Klass, Oren S., Ofer Biham, Moshe Levy, Ofer Malcai, and Sorin Solomon, "The Forbes 400 and the Pareto Wealth Distribution," *Economics Letters*, 2006, *90* (2), 290–295.

Kopczuk, Wojciech, "What Do We Know about the Evolution of Top Wealth Shares in the United States?," *The Journal of Economic Perspectives*, 2015, *29* (1), 47–66.

_ **and Emmanuel Saez**, "Top Wealth Shares in the United States, 1916–2000: Evidence from Estate Tax Returns," *National Tax Journal*, 2004, *57* (2), 445–487.

Liu, Kang, Nicholas Lubbers, William Klein, Jan Tobochnik, Bruce Boghosian, and Harvey Gould, "The Effect of Growth on Equality in Models of the Economy," *arXiv:1305.0794*, 2013.

Liu, Z. and Rostislav A. Serota, "Correlation and Relaxation Times for a Stochastic Process with a Fat-tailed Steady-state Distribution," *Physica A: Statistical Mechanics and its Applications*, 2017, *474*, 301–311.

Meade, James E., *Efficiency, Equality and The Ownership of Property*, London, UK: Allen & Unwin, 1964.

Mehra, Rajnish and Edward C. Prescott, "The Equity Premium: A puzzle," *Journal of Monetary Economics*, 1985, *15* (2), 145–161.

Nardi, Mariacristina De, "Quantitative Models of Wealth Inequality: A Survey," 2015. Working Paper 21106, NBER.

_ , Giulio Fella, and Fang Yang, "Piketty's Book and Macro Models of Wealth Inequality," 2015. Working Paper 21730, NBER.

Nelder, John A. and Roger Mead, "A Simplex Method for Function Minimization," *The Computer Journal*, 1965, *7* (4), 308–313.

Pareto, Vilfredo, *Cours d'économie Politique*, Lausanne, Switzerland: F. Rouge, 1897.

Peters, Ole, "Optimal Leverage from Non-ergodicity," *Quantitative Finance*, 2011, *11* (11), 1593–1602.

_ , "The Time Resolution of the St Petersburg Paradox," *Phil. Trans. R. Soc. A*, December 2011, *369* (1956), 4913–4931.

_ **and Alexander Adamou**, "Stochastic Market Efficiency," 2013. Working Paper 13-06-022, SFI.

_ **and Murray Gell-Mann**, "Evaluating Gambles Using Dynamics," *Chaos*, February 2016, *26* (2), 23103.

_ **and William Klein**, "Ergodicity Breaking in Geometric Brownian Motion," *Physical Review Letters*, March 2013, *110* (10), 100603.

Piketty, Thomas, *Capital in the Twenty-first Century*, Cambridge, MA: Harvard University Press, 2014.

_ **and Emmanuel Saez**, "A Theory of Optimal Inheritance Taxation," *Econometrica: Journal of the Econometric Society*, 2013, *81* (5), 1851–1886.

— **and Gabriel Zucman**, "Capital is Back: Wealth-Income Ratios in Rich Countries, 1700–2010," *The Quarterly Journal of Economics*, 2014, *129* (3), 1255–1310.

Pinkovskiy, Maxim and Xavier Sala-i-Martin, "Parametric Estimations of the World Distribution of Income," 2009. Working Paper 15433, NBER.

Quandl, "Dow Jones Industrial Average," `http://www.quandl.com/data/BCB/UDJIAD1-Dow-Jones-Industrial-Average` 2016. Accessed: 04/19/2016.

Rios-Rull, Jose-Victor and Moritz Kuhn, "2013 Update on the U.S. Earnings, Income, and Wealth Distributional Facts: A View from Macroeconomics," *Quarterly Review*, 2016, (April), 1–75.

Saez, Emmanuel and Gabriel Zucman, "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data," *The Quarterly Journal of Economics*, 2016, *131* (2), 519–578.

Samuelson, Paul A., "What Classical and Neoclassical Monetary Theory Really Was," *The Canadian Journal of Economics*, 1968, *1* (1), 1–15.

Stiglitz, Joseph E., "Distribution of Income and Wealth Among Individuals," *Econometrica: Journal of the Econometric Society*, 1969, *37* (3), 382–397.

The World Wealth and Income Database, "USA Top 10% and 1% and Bottom 50% Wealth Shares, 1913–2014," `http://wid.world/data/` 2016. Accessed: 12/26/2016.

Uhlenbeck, George E. and Leonard S. Ornstein, "On the Theory of the Brownian Motion," *Physical Review*, Sep 1930, *36* (5), 823–841.

Vasicek, Oldrich, "An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics*, 1977, *5* (2), 177–188.

Vermeulen, Philip, "How Fat is the Top Tail of the Wealth Distribution?," *Review of Income and Wealth*, 2017.

Wolff, Edward N., "Household Wealth Trends in the United States, 1983–2010," *Oxford Review of Economic Policy*, 2014, *30* (1), 21–43.

A The derivation of self-averaging time t_c

In transforming the differential equation for wealth, x, into a differential equation for rescaled wealth, y, we use the approximation $\langle x(t)\rangle_N = \langle x(0)\rangle_N e^{\mu t}$. In other words we assume that the population-average wealth grows like the expectation value of wealth.

It is known that this approximation is invalid for long times (Peters and Klein, 2013). Specifically, over long times $\langle x(t)\rangle_N$ grows at the exponential rate $\mu - \sigma^2/2$, whereas the expectation value, $\langle x(t)\rangle$, grows at the exponential rate μ.

This raises the question for how long our approximation is valid. The answer depends on the sample size N, as it must because the expectation value is just the $N \to \infty$ limit of the population average. To assess whether the fluctuations in $\langle x(t)\rangle_N$ are important, we compare the variance of $\langle x(t)\rangle_N$ to the expectation value squared (equivalent to comparing the standard deviation to the expectation value). If the variance is smaller than the expectation value squared, then the approximation is acceptable. If this is not the case, then we cannot use this approximation.

The calculations that use the approximation relate to properties of the stationary distribution. This exists for τ above some positive threshold, *i.e.* with sufficiently strong reallocation. The coupling of wealth trajectories through reallocation lengthens the timescale over which the population average resembles the expected wealth. Therefore, we are on safe ground if we can show that the timescale on which the approximation is valid when $\tau = 0$ is longer than practically relevant timescales. This is a sufficient condition for the approximation to be valid when $\tau > 0$.

This means we work with a population of N independent GBMs, which makes the computation of the variance easy. GBM is log-normally distributed,

$$\ln [x(t)] \sim \ln \mathcal{N} \ln \left(x(0) + \left[\mu - \frac{\sigma^2}{2} \right] t, \sigma^2 t \right). \tag{12}$$

From this, it follows that the expectation value of a single trajectory grows as

$$\langle x(t)\rangle = x(0)\, e^{\mu t}, \tag{13}$$

and the variance grows as

$$V[x(t)] = x(0)^2\, e^{2\mu t}\left[e^{\sigma^2 t} - 1\right]. \tag{14}$$

Because the wealth trajectories are independent, the variance of an average over N trajectories is one-N^{th} of the variance for the individual trajectory. We are now in a position to compare standard deviation and average as follows

$$\frac{V\left[\langle x(t)\rangle_N\right]}{\langle\langle x(t)\rangle_N\rangle^2} = \frac{e^{\sigma^2 t} - 1}{N}. \tag{15}$$

So long as this is less than one, $\langle x(t)\rangle$ is a good approximation for $\langle x(t)\rangle_N$. Rearranging and taking $N \gg 1$ gives an upper bound on the time for which this approximation holds:

$$t < t_c \equiv \frac{\ln N}{\sigma^2}. \tag{16}$$

For typical parameter values in our model, $N = 10^8$ and $\sigma = 0.16\,\text{year}^{-1/2}$, we find $t_c \approx 700$ years. It turns out that, strictly speaking, the minimum value of τ required for the stationary distribution to exist is not zero but proportional to the inverse of this timescale (Bouchaud, 2015a), *i.e.*

$$\tau > \tau_c \equiv \frac{\sigma^2}{2\ln N}. \tag{17}$$

In essence, the inequality-increasing effects of multiplicative growth drive wealths apart on the timescale t_c, whereas the inequality-reducing effects of reallocation drive wealths back together on the timescale $1/\tau$. Thus, τ_c marks the point at which reallocation overcomes the forces that drive wealths apart, leading to a stationary distribution of rescaled wealth. In our case $\tau_c \approx 0.0007$ year^{-1}, which is in practice indistinguishable from zero.

The timescale t_c was derived assuming that everyone starts out equally, which is not generally the case. If the initial distribution of wealth is very unequal, then the variance of $\langle x \rangle_N$ will be dominated by the fluctuations experienced by the wealthiest individuals, and the approximation $\langle x \rangle_N \approx \langle x(0) \rangle_N e^{\mu t}$ becomes invalid more quickly. We confirmed numerically that the effect is negligible for our study: our results are indistiguishable whether we simulate dx (which requires an estimate for μ), or dy (where μ does not appear but the above-mentioned approximation is made).

B The derivation of the stationary distribution

We start again with the SDE for the rescaled wealth,

$$dy = \sigma\, y\, dW - \tau\,(y-1)\, dt. \tag{18}$$

This is an Itô equation with drift term $A = \tau(y-1)$ and diffusion term $B = y\sigma$.

Such equations imply ordinary second-order differential equations that describe the evolution of the pdf, called Fokker-Planck equations. The Fokker-Planck equation describes the change in probability density, at any point in (relative-wealth) space, due to the action of the drift term (like advection in a fluid) and due to the diffusion term (like heat spreading). In this case, we have

$$\frac{dp\,(y,t)}{dt} = \frac{\partial}{\partial y}\left[Ap\,(y,t)\right] + \frac{1}{2}\frac{\partial^2}{\partial y^2}\left[B^2 p\,(y,t)\right]. \tag{19}$$

The steady-state Fokker-Planck equation for the pdf $p\,(y)$ is obtained by setting the time derivative to zero,

$$\frac{\sigma^2}{2}\left(y^2 p\right)_{yy} + \tau\left[(y-1)\,p\right]_y = 0. \tag{20}$$

Positive wealth subjected to continuous-time multiplicative dynamics with non-negative reallocation can never reach zero. Therefore, we solve Equation (20) with boundary condition $p\,(0) = 0$ to give

$$p\,(y) = C\,(\zeta)\, e^{-\frac{\zeta-1}{y}}\, y^{-(1+\zeta)}, \tag{21}$$

where

$$\zeta = 1 + \frac{2\tau}{\sigma^2} \tag{22}$$

and

$$C\,(\zeta) = \frac{(\zeta-1)^{\zeta}}{\Gamma\,(\zeta)}, \tag{23}$$

with the gamma function $\Gamma\,(\zeta) = \int_0^\infty x^{\zeta-1} e^{-x}\, \mathrm{d}x$. The distribution has a power-law tail as

$y \to \infty$, resembling Pareto's often confirmed observation that the frequency of large wealths tends to decay as a power law. The exponent of the power law, ζ, is called the Pareto parameter and is one measure of economic inequality.

C The effect of fixed versus time-varying σ

Our analysis requires us to set the volatility parameter, σ, in the RGBM model. We estimate this as a time-varying quantity, $\sigma(t)$, using data from the American stock market. For each year in the analysis, we define $\sigma(t)$ as the standard deviation of the daily logarithmic changes in the Dow Jones Industrial Average (DJIA) for that year (Quandl, 2016), annualised by multiplying by $(250/\text{year})^{1/2}$. We present this in Figure 11.

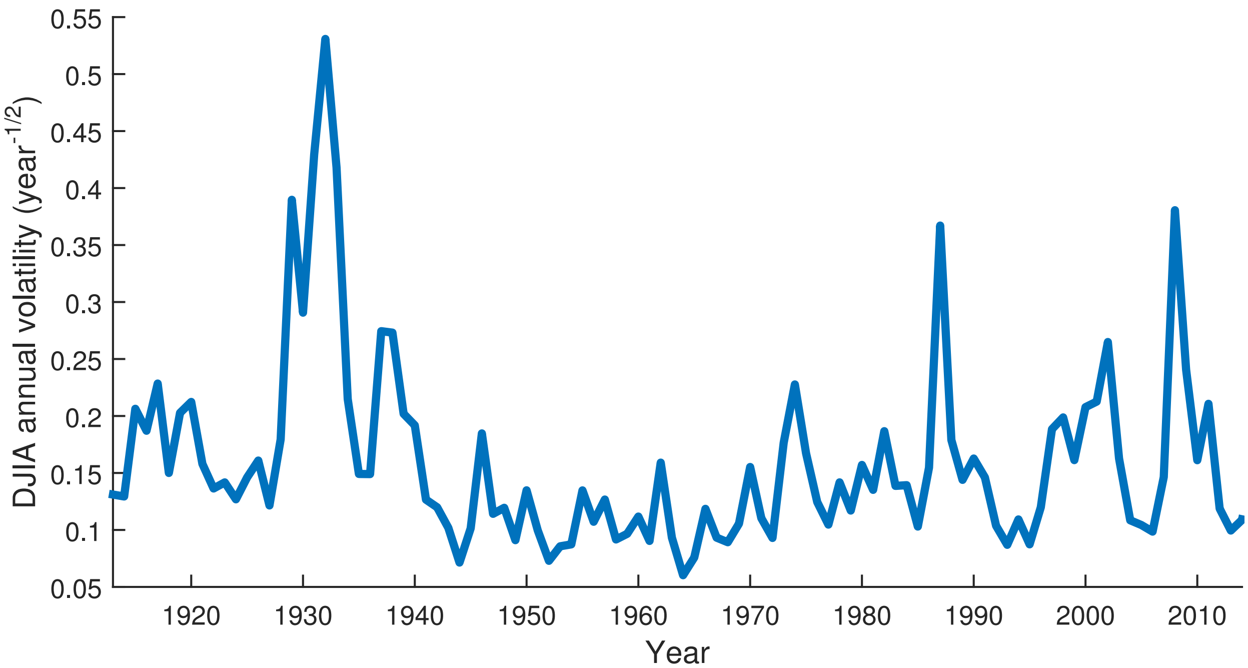

Figure 11: Annualised $\sigma(t)$ estimated from daily changes in the DJIA. Data taken from Quandl (2016).

In the results presented in Section 4, we used a fixed value – $\sigma = 0.16$, the average of the time-varying $\sigma(t)$ – for simplicity. In order to show this has little effect on the estimated reallocation rates, a comparison is presented in Figure 12.

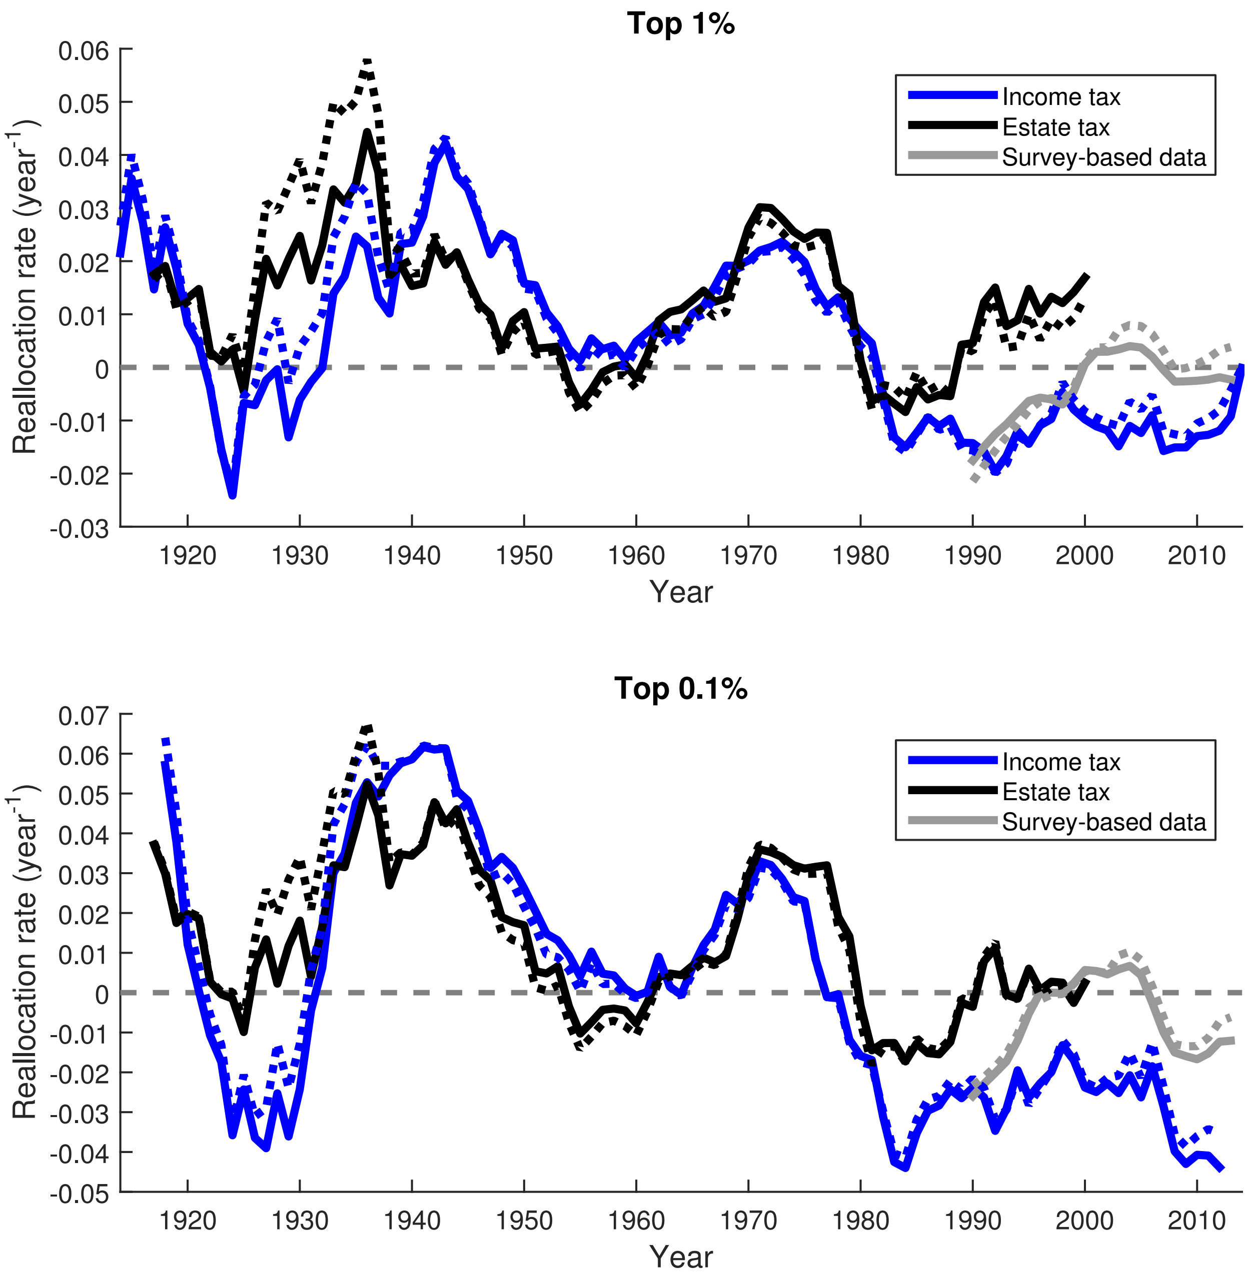

Figure 12: Smoothed effective reallocation rates with fixed and time-varying σ. Solid lines – fixed σ; Dotted lines – time-varying σ.

D The derivation of the variance convergence time

Our key finding is that under currently prevailing economic conditions it is not safe to assume the existence of stationary wealth distributions in models of wealth dynamics. Nevertheless, we present some results for the regime of our model where a stationary distribution exists. The full form of the distribution is derived in Appendix B. Because it has a power-law tail for large wealths, only the lower moments of the distribution exist, while higher moments diverge. Below, we derive a condition for the convergence of the variance and calculate its convergence time.

The variance of y is a combination of the first moment, $\langle y \rangle$ (the average), and the second moment, $\langle y^2 \rangle$:

$$V(y) = \left\langle y^2 \right\rangle - \langle y \rangle^2 \tag{24}$$

We thus need to find $\langle y \rangle$ and $\langle y^2 \rangle$ in order to determine the variance. The first moment of the rescaled wealth is, by definition, $\langle y \rangle = 1$.

To find the second moment, we start with the SDE for the rescaled wealth:

$$dy = \sigma y \, dW - \tau (y - 1) \, dt. \tag{25}$$

This is an Itô process, which implies that an increment, df, in some (twice-differentiable) function $f(y, t)$ will also be an Itô process, and such increments can be found by Taylor-expanding to second order in dy as follows:

$$df = \frac{\partial f}{\partial t} dt + \frac{\partial f}{\partial y} dy + \frac{1}{2} \frac{\partial^2 f}{\partial y^2} dy^2 \,. \tag{26}$$

We insert $f(y, t) = y^2$ and obtain

$$d\left(y^2\right) = 2y dy + (dy)^2 \,. \tag{27}$$

We substitute dy in Equation (27), which yields terms of order dW, dt, dW^2, dt^2, and $(dW\,dt)$. The scaling of Brownian motion allows us to replace dW^2 by dt, and we ignore $o\,(dt)$ terms. This yields

$$d\left(y^2\right) = 2\sigma y^2\,dW - \left(2\tau - \sigma^2\right) y^2\,dt + 2\tau y\,dt$$

Taking expectations on both sides, and using $\langle y \rangle = 1$, produces an ordinary differential equation for the second moment:

$$\frac{d\langle y^2\rangle}{dt} = -\left(2\tau - \sigma^2\right)\langle y^2\rangle + 2\tau \tag{28}$$

with solution

$$\langle y\,(t)^2\rangle = \frac{2\tau}{2\tau - \sigma^2} + \left(\langle y\,(0)^2\rangle - \frac{2\tau}{2\tau - \sigma^2}\right) e^{-\left(2\tau - \sigma^2\right)t}\,. \tag{29}$$

The variance $V\,(t) = \langle y\,(t)^2\rangle - 1$ therefore follows

$$V\,(t) = V_\infty + (V_0 - V_\infty)\, e^{-\left(2\tau - \sigma^2\right)t}\,, \tag{30}$$

where V_0 is the initial variance and

$$V_\infty = \frac{2\tau}{2\tau - \sigma^2}\,. \tag{31}$$

V converges in time to the asymptote, V_∞, provided the exponential in Equation (30) is decaying. This can be expressed as a condition on τ

$$\tau > \frac{\sigma^2}{2}\,. \tag{32}$$

Clearly, for negative values of τ the condition cannot be satisfied, and the variance (and inequality) of the wealth distribution will diverge. In the regime where the variance exists,

$\tau > \sigma^2/2$, it also follows from Equation (30) that the convergence time of the variance is $1/\left(2\tau - \sigma^2\right)$.

As τ increases, increasingly high moments of the distribution become convergent to some finite value. The above procedure for finding the second moment (and thereby the variance) can be applied to the k^{th} moment, just by changing the second power y^2 to y^k, and any other cumulant can therefore be found as a combination of the relevant moments. For instance Liu and Serota (2017) also compute the third cumulant.