Counterfactual Wealth Distributions*

MATTHIEU GOMEZ †

First version: March 2024
This version: December 2024
[latest version]

Abstract

This paper analyzes the response of the wealth distribution to changes in the economic environment. I show that the first-order response of top wealth shares to a shock at some horizon is determined by the average wealth response of individuals projected to be in the top percentile at this horizon. This simple result helps to isolate the key forces and empirical moments that determine the dynamics of top wealth shares in response to various counterfactual scenarios. I leverage this framework to revisit existing issues in the inequality literature, such as the effect of transitory and permanent shocks in asset returns on the wealth distribution.

*I wish to thank Dirk Krueger and Clara Martinez-Toledano for helpful suggestions, as well as seminar participants at Columbia, Penn, and Imperial College for comments. Patricio Goldstein and Jack Johnson provided excellent research assistance.

†Columbia University; mg3901@columbia.edu

Introduction

An important goal for policymaking is to understand what drives realized movements in the wealth distribution or how different economic policies would affect wealth inequality. One common way to answer these questions is to build models of the wealth distribution and use them to simulate the effect of policy changes on wealth inequality. One drawback of these models is that they can be fairly complex, making it challenging to discern the key forces within them that are responsible for specific outcomes.

To help strengthen the credibility of this structural approach, this paper develops a tractable and transparent framework to analyze the response of top wealth shares to various counterfactual scenarios. The key insight is that, in a large class of models, the first-order response of top wealth shares at some horizon h is determined by the average response of individuals projected to be in the top percentile at horizon h. This simple result allows me to derive closed-form formulas for the dynamics of top wealth shares in response to perturbations in the dynamics of individual wealth.

This framework sheds light on the mechanism behind the impulse response of top wealth shares to various economic shocks. This is useful for researchers who want to better understand the impulse responses implied by structural models or the ones estimated empirically using local projection methods. Conversely, it also allows researcher to compute more credible counterfactuals, by indicating the set of empirical statistics that a model needs to match to generate credible counterfactuals. Additionally, under common assumptions on consumption function, this framework can also serve as a simple accounting framework to trace out the dynamics of top wealth shares under various counterfactuals, improving existing methods used by reduced-form researchers that abstract from composition changes (e.g. Saez and Zucman, 2016, Martínez-Toledano, 2020).

The paper is organized in three parts. In the first part, I show that, in a wide range of models, the first-order response of the average wealth in a top percentile is simply given by the average wealth response of individuals in the top percentile. Put differently, while economic shocks typically generate composition changes in the top percentile, these changes are second-order for the response of the average wealth in top percentiles. To understand why, note that composition changes are equal to the product between the relative mass of individuals entering the top percentile and the average impact of these entrants on the average wealth in the top percentiles. Because both terms are first-order in the size of the perturbation, their product is second-order.

This result can be used to trace out the entire impulse response function of top wealth shares to shocks. Note, however, that the response of the average wealth in a top percentile at horizon h is equal to the average wealth response of individuals *who will be in the top percentile* at horizon h, rather than the individuals initially in the top percentile (i.e., at $h = 0$). Put differently, although composition changes in top percentiles due to small shocks are second-order, composition changes due to the normal passage of time are not.[1] Still, this does not substantially complicate the application of this formula as composition changes in the baseline economy are observable.

To examine the accuracy of this first-order approximation, I also derive a second-order approximation for the response of the average wealth in top percentiles. This second-order term, that captures the effect of composition changes, increases with the relative mass of individuals around the percentile threshold and the heterogeneity in the response of individual wealth. Hence, holding other things equal, one can expect the first-order approximation to be more accurate for distributions that are highly unequal and for policy changes that tend to affect all rich individuals in the same way.

Second, I use this approach to characterize the impulse response of top wealth shares to changes in the dynamics of individual wealth in a wide range of random-growth models. I first focus on the simple case in which individual wealth follows a random walk, subject to some population renewal (e.g. death or population growth). In this case, I obtain a simple formula expressing the response of top wealth shares to a permanent change in individual growth rates in terms of the distribution of ages. In particular, the long-run elasticity of top wealth shares to individual growth rates is given by the difference between the wealth-weighted average age of individuals in the top percentile and individuals in the overall economy.

This simple result survives a number of extensions. I first consider a type-specific change in growth rates: in this case, the concept of age must be redefined as the average number of years spent in that type. I then consider a model in which individual wealth follows a random walk with an additive term (e.g. unexpected bequests): in this case, the concept of age must be redefined as an index that discount previous periods based on the relative significance of the additive factor in individual wealth growth.

I then consider existing calibration of random-growth models in the literature and I demonstrate that there is a dramatic heterogeneity in the distribution of age across the wealth distribution, and, therefore, of the elasticity of top wealth shares to various shocks in individual growth rates (by an order of magnitude). This implies that these models predict very different responses for the exact same counterfactual.

Finally, in the last part of the paper, I leverage my analytical results to set up and calibrate a flexible model of wealth accumulation. I ensure that the model exactly matches the relative importance of labor income and bequest across the wealth distribution, as measured using data from the Survey of Consumer Finances and the Forbes 400 list. Indeed, the analytical results discussed above suggest that this plays a key role for the effect of counterfactual shocks on individual wealth on top wealth shares. I argue that the calibrated model provides credible and transparent answers to important questions in the inequality literature such as: What is the role of labor income inequality and bequests for wealth inequality? How would top wealth shares evolve if the average

[1] Such composition changes are non-infinitesimal even if the time horizon is infinitesimal, since individuals experience large idiosyncratic shocks as emphasized in Gomez (2023).

return on wealth were to increase? How does the wealth distribution respond to transitory or permanent shocks on the returns of specific assets?

Literature review. There is a growing literature trying to understand the effect of various changes in economic environment on wealth inequality. One strand of the literature has turned to large macro models that are able to capture the rich heterogeneity of the micro data (e.g. De Nardi, 2004, Kaymak and Poschke, 2016, Benhabib et al., 2019, Peter, 2019, Hubmer et al., 2021, Guvenen et al., 2023…). One drawback of these models is that they can be fairly complex, making it challenging to discern the forces within the models that are responsible for specific outcomes.

Another strand of the literature relies on more stylized and analytically tractable "random growth" (e.g. Wold and Whittle, 1957, Benhabib et al., 2011, Jones, 2015, Piketty and Zucman, 2015, Moll et al., 2022). While these models are able to obtain closed-form formula for the effect of changes in Pareto inequality, this is only informative about the asymptotic behavior of top wealth shares, both with respect to time (long-run) and with respect to wealth (right-tail). I will stress that these results can be very uninformative about the deviation of top wealth shares at a given time or at a given top percentile.

My analytical approach provides closed-form formulas for the transition dynamics of top wealth shares in response to (small) changes in the growth rate of individual wealth. These formulas improves on the existing literature, that typically resort to numerical method (Kolmogorov-Forward equation) or simulations to compute the response of top wealth shares to aggregate shocks. Moreover, my results on the transition path of top wealth shares complements Gabaix et al. (2016), Luttmer (2016), and Atkeson and Irie (2022), who characterize analytically the convergence of densities according to global norms (e.g. the integral of the absolute difference, or the difference in moments). For the case of a mechanical effect in the growth rate of wealth, I show that one key statistic is the gradient of age with respect to wealth. This relates this paper to Luttmer (2011) who computes this gradient in a model of firm growth and compares it with the data. This result also has an antecedent in Saez and Zucman (2019), who studies the effect of an increase in billionaire taxation on the average wealth of future billionaires. In particular, they express the long-run effect of a tax rate on the average number of years that fortunes spend at the top, which is similar to the type of results I obtain in Section 2.

Finally, the paper contributes a more reduced-form literature that computes counterfactuals on wealth inequality. One widespread approach, started with Saez and Zucman (2016), is to use an accounting decomposition of wealth in top percentiles assuming that the set of households in the top percentile remains the same over time (see Martínez-Toledano, 2020 or Mian et al., 2020 for applications). The accounting framework discussed in the last section of the paper can be seen as a version of this approach that takes seriously the effect of composition changes. While the difference between the two does not matter much in the short-run (say, a few years), it can lead to dramatically different results in the longer-run (say, a few decades). In spirit, my accounting framework is more closely related to the exercises done by Feiveson and Sabelhaus (2018) and Ozkan et al. (2023), who take into account the effect of changes in the economic environment via the lifetime trajectory of individuals in the top. The approach of my paper to compute counterfactuals is also related from an influential literature in labor economics, that focuses on quantifying the effect of changes in individual *characteristics* (e.g. gender, union membership, skills) on inequality using re-weighting methods (see, for instance DiNardo et al., 1995, Firpo et al., 2009 and Fortin et al., 2011). Instead, my paper focuses on quantifying the effect of small changes in the dynamics of individual wealth (e.g. income, returns, bequests), holding fixed the set of individual characteristics.[2]

At a general level, this paper contributes to the growing work employing first-order approximations to characterize economies with a large degree of micro heterogeneity — see Chetty (2009) and Kleven (2021) for a review of this approach in public finance, Costinot and Rodríguez-Clare (2014) in trade, and Baqaee and Rubbo (2023) or Auclert et al. (2018) in macroeconomics.

Roadmap. The paper is organized in three sections. Section 1 characterizes the response of the average wealth in a top percentile in terms of the average response of individual wealth in the top percentile. Section 2 uses this fact to analyze the response of top wealth shares to perturbations in the dynamics of individual wealth in random-growth models. Section 3 leverages these results to calibrate a model of wealth accumulation and use it to trace out the dynamics of top wealth shares under various counterfactual scenarios.

1 General framework

In this section, I study the response of top wealth (or income) shares to general changes in the economic environment. In Section 1.1, I emphasize that the response in the average wealth in a top percentile can always be decomposed into an intensive margin (the average wealth response of individuals in the top percentile) and an extensive margin (the response in the composition of individuals in the top percentile). I then use this decomposition to characterize the response of the average wealth in the top percentile for small perturbations in Section 1.2.

1.1 Large perturbation

Consider an economy indexed by time $t \in \mathbb{R}$ with a continuum of individuals indexed by i. Assume that the cumulative distribution function of wealth is absolutely continuous and increasing, with a finite average.[3] Denote $W_{i,t}$ the wealth of individual i at time t. Consider a given top percentile p (say, the top $p = 1\%$), and denote $Q_t(p)$ the quantile corresponding to the top percentile

[2] I discuss more precisely the relationship between the two approaches in Appendix A.3.

[3] This implies that the density of wealth is well defined and positive everywhere. I discuss how to extend the results to wealth distributions without density in Appendix A.

p at time t:

$$\mathbb{P}(W_{i,t} \geq Q_t(p)) = p \tag{1}$$

Denote $\overline{W}_t(p)$ the average wealth in the top percentile p:

$$\overline{W}_t(p) = \mathbb{E}[W_{i,t} | W_{i,t} \geq Q_t(p)]. \tag{2}$$

Consider a change in the economic environment compared to a baseline scenario (e.g. deviation in technology or asset prices) indexed by θ, starting from time $t = 0$. For any economic variable Y_t in the baseline economy (such as individual wealth, quantile, or average wealth in the top), denote $Y_t(\theta)$ its value in the perturbed economy indexed by θ and $\Delta Y_t = Y_t(\theta) - Y_t(0)$ its difference between the counterfactual and the baseline economies. It is important to keep in mind that Δ represents a vertical deviation (i.e. with respect to a perturbation) rather than a horizontal change (i.e. a change with respect to time). To visualize this, Figure 1 plots an example for the wealth path of some individual i over time, in the baseline economy and the perturbed one.

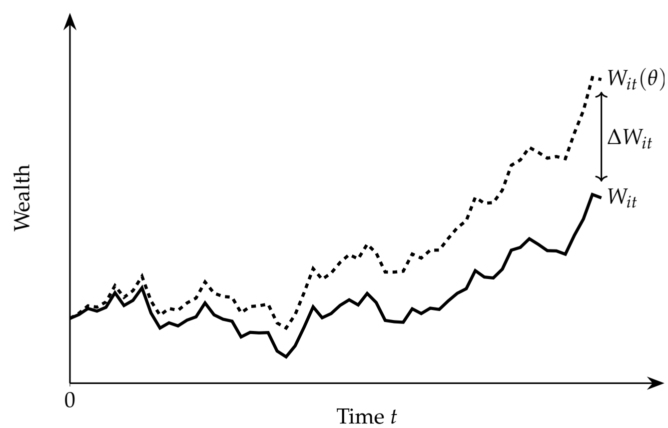

Figure 1: Baseline (solid) and perturbed (dashed) wealth path for some individual i

I am interested in the effect of this deviation on the average wealth in the top percentile. I will focus on the response of the average wealth in the top percentile for two reasons. First, it is typically the object measured and reported in existing models of inequality. Second, requirement for differentiability of the average wealth are typically lower relative to quantiles. Finally, note that the knowledge of the response of all top wealth shares is equivalent to the knowledge of the response of the whole distribution. By definition, the deviation in the average wealth in top percentile is:

$$\Delta \overline{W}_t(p) = \mathbb{E}\left[W_{i,t} + \Delta W_{i,t} | W_{i,t} + \Delta W_{i,t} \geq Q_t(p) + \Delta Q_t(p)\right] - \mathbb{E}[W_{i,t} | W_{i,t} \geq Q_t(p)].$$

One important fact is the perturbation typically changes not only the wealth of individuals initially the top percentile, but also the set of individuals in the top percentile (in the baseline versus perturbed economy). To disentangle between the two effects, it is useful to add and subtract $\mathbb{E}[W_{i,t} + \Delta W_{i,t} | W_{i,t} \geq Q_t(p)]$ in the right-hand side of this equation to rewrite the deviation in the average wealth in top percentile as:

$$\Delta \overline{W}_t(p) = \underbrace{\mathbb{E}\left[\Delta W_{i,t} | W_{i,t} \geq Q_t(p)\right]}_{\text{Intensive margin}} + \underbrace{\mathbb{E}\left[W_{i,t} + \Delta W_{i,t} | W_{i,t} + \Delta W_{i,t} \geq Q_t(p) + \Delta Q_t(p)\right] - \mathbb{E}\left[W_{i,t} + \Delta W_{i,t} | W_{i,t} \geq Q_t(p)\right]}_{\text{Extensive margin}}. \tag{3}$$

This equation expresses the deviation in the average wealth in a top percentile is the sum of two distinct terms. The first term, the "intensive" margin, corresponds to the average deviation in the wealth of individuals projected to be in the top percentile at time t in the baseline economy. The second term, the "extensive" margin, captures the effect of the perturbation on the composition of individuals in the top percentile. It is equal to the difference between the average perturbed wealth of individuals that are in the top percentile in the perturbed economy and the average perturbed wealth of those that were in the top percentile in the baseline economy. By definition, of top percentiles, this term is always greater or equal than zero.

1.2 Small perturbation

We now assume that individual wealth is differentiable with respect to the size of the deviation θ.

Assumption 1. *The function $\theta \to W_{i,t}(\theta)$ is in $\mathcal{C}^1$ for all individuals.*[4]

When this assumption is satisfied, we denote $\mathrm{d}W_{i,t} = W'_{i,t}(\theta)\,\mathrm{d}\theta$. This assumption restricts the set of economic deviations to the ones that generate smooth changes in individual wealth. For instance, this assumption allows for deviations in the economic environment that affect the size of individual wealth jumps; however, it rules out deviations that affect the probability of this jump happening. One way to think about this assumption is that it allows for deviations that change the realization of wealth across different state of natures, not for deviations that change the probability of these states happening.[5]

While the first assumption ensures that deviations in wealth are first-order in θ, the second assumption ensures that the mass of individuals that enter or exit is first-order in θ. Together, these two assumptions ensure that the extensive margin in (3) is second-order in θ.

Proposition 1 (First-Order Approximation). *At the first-order in θ, the deviation in the average wealth in a top percentile is:*

$$\mathrm{d}\overline{W}_t(p) = \mathbb{E}\left[\mathrm{d}W_{i,t} | W_{i,t} \geq Q_t(p)\right].$$

[4]Here, and in the rest of the paper, I use $\mathcal{C}^k$ to denote the set of functions that are differentiable $k-$th times with continuous derivatives.

[5]Still, this more general case could be handled using similar techniques as in Gomez (2023). However, in this case, composition changes would become first-order, as some individuals experience discontinuous changes in wealth in response to small changes in the economic environment

This proposition says that the first-order deviation in the average wealth in a top percentile is simply given by the average wealth deviation of individuals in the top percentile. Relative to the expression for the non-infinitesimal deviation (3), this proposition says that one can abstract from the extensive margin: the first-order effect of a shock on the average wealth in a top percentile can be computed *as if* the set of individuals in the top percentile remained the same in the perturbed economy. This is very convenient, both analytically (as one does not need to take into account re-ranking effects), but also empirically (as one only needs data on the effect of the perturbation on the wealth of individuals currently in the top)

Composition changes over time. This proposition characterizes the entire impulse response function of the wealth distribution, from its instantaneous response $t = 0$ to its long-run response $t \to \infty$. Despite the simplicity of the formula, note that there is an important subtlety: what matters for the response of the average wealth in the top percentile at some horizon t is the average wealth response of individuals who will be in the top percentile at time t, not the response of those who are initially in the top at $t = 0$. Mathematically:

$$\mathrm{d}\overline{W}_t(p) = \mathbb{E}\left[\mathrm{d}W_{i,t}|W_{i,t} \geq Q_t(p)\right] \neq \mathbb{E}\left[\mathrm{d}W_{i,t}|W_{i,0} \geq Q_0(p)\right]. \tag{4}$$

The impulse response of the average wealth in a top percentile is *not* the average impulse response of individuals initially in the top percentile. Formally, the difference between the two can be rewritten as:

$$\begin{aligned}\mathrm{d}\overline{W}_t(p) - \mathbb{E}\left[\mathrm{d}W_{i,t}|W_{i,0} \geq Q_0(p)\right] &= \mathbb{P}(W_{i,t} < Q_t(p)|W_{i,0} \geq Q_0(p)) \\ &\times \left(\mathbb{E}\left[\mathrm{d}W_{i,t}|W_{i,t} \geq Q_t(p), W_{i,0} \leq Q_0(p)\right] - \mathbb{E}\left[\mathrm{d}W_{i,t}|W_{i,t} \leq Q_t(p), W_{i,0} \geq Q_0(p)\right]\right).\end{aligned}$$

The difference between the two terms (the right-hand-side) is the product between the two terms: (i) the fraction of individuals in the top percentile at $t = 0$ who exit it between 0 and t and (ii) the difference between the wealth response of individuals that enter and those that exit between 0 and t. One can expect this wedge to increases with the horizon t, as more and more people exit the top percentile over time. To represent this fact, Figure 2 plots the fraction of individuals in the Forbes 400 that remains in the list after t years. As shown in the figure, this fraction quickly declines with the horizon: after ten years, only half of individuals originally in the top percentile remain in it. This emphasizes the importance of taking seriously the effect of composition changes over time when thinking about the effect of a perturbation at horizons longer than a few years. One way to summarize this discussion is that, while composition changes in the top percentile induced by the perturbation are first-order in θ, composition changes induced by the normal passage of time are not.[6]

[6]One could think that the reason for this fact is, while the size of the perturbation is infinitesimal, the horizon t is not. However, even in continuous-time, composition changes remain first-order with respect to the time horizon, as,

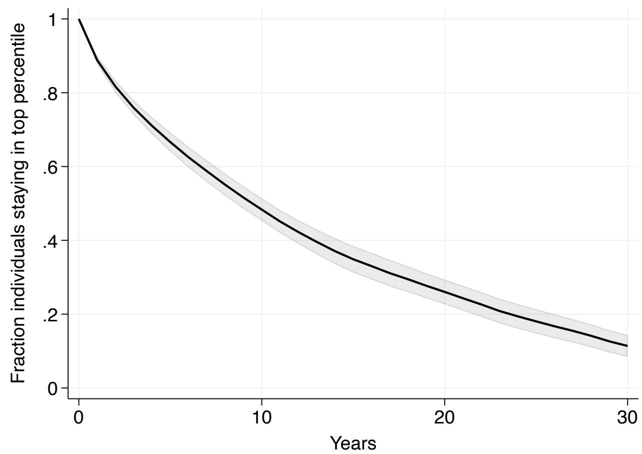

Figure 2: Fraction of individuals remaining in the Forbes 400 list over the horizon

Notes: The figure plots the fraction of individual in the Forbes 400 that remain in the top after t years. More precisely, for each t, I compute the average of a dummy variable equal to one if the individual remains in the top after t year, and zero if not. Confidence-intervals are obtained using two-way cluster-robust standard errors at the year and individual level.

Deviation in top wealth shares. I now extend Proposition 4 to derive a similar formula for the deviation in the *log* average wealth in top percentiles in terms of the deviations in *log* individual wealth. Dividing Proposition 1 by $\overline{W}_t(p)$ and rearranging gives:

$$\mathrm{d}\log \overline{W}_t(p) = \mathbb{E}^{W_{i,t}}\left[\mathrm{d}\log W_{i,t} | W_{i,t} \geq Q_t(p)\right]. \tag{5}$$

In words, the log deviation for the average wealth in the top percentile is the wealth-weighted average log deviation of individuals in the top percentile. Working with logarithms is particularly useful to express the deviation in top wealth shares, as $S_t(p) = p\overline{W}_t(p)/W_t(100\%)$ implies that the *log* deviation in the top wealth share is the difference between the *log* deviation for the average wealth in the top and in the economy, i.e.

$$\mathrm{d}\log S_t(p) = \mathbb{E}^{W_{i,t}}\left[\mathrm{d}\log W_{i,t} | W_{i,t} \geq Q_t(p)\right] - \mathbb{E}^{W_{i,t}}\left[\mathrm{d}\log W_{i,t}\right]$$

Deviation in quantiles. As shown in the proof of Proposition 1 in Appendix A, the idea that composition changes do not matter at the first order also works for quantiles; that is,

$$\mathrm{d}Q_t(p) = \mathbb{E}\left[\mathrm{d}W_{i,t} | W_{i,t} = Q_t(p)\right].$$

In the rest of the paper, however, I focus on the average wealth in top percentiles shares in the rest of the talk as they are typically easier to work with, both theoretically and empirically. But it is important to realize that the methods and formula that I obtain allow to characterize the deviation in the entire distribution of wealth over time.

typically, individual wealth paths are non differentiable with respect to time.

Higher-order terms. For now, we have focused on the first-order effect of a perturbation. In Appendix A.2, I study the deviation in the average wealth in the top percentile beyond the first-order. More precisely, I obtain the following second-order approximation:

$$\Delta\overline{W}_t(p) = \mathbb{E}[\Delta W_i | W_{i,t} \geq Q_t(p)] + \frac{1}{2}\frac{g_t(Q_t(p))}{p}\mathbb{V}[\Delta W_{i,t} | W_{i,t} = Q_t(p)] + o(\theta^2),$$

where $\mathbb{V}$ denotes the cross-sectional variance of the deviation in individual wealth. This expression provides a second-order approximation for the extensive margin, making it possible to assess the importance of second-order effects due to composition changes, both theoretically and empirically. In particular, note that the second-order term increases with the density of households at the percentile threshold and with the variance of wealth deviations at the percentile threshold, as both increase the importance of composition effects. Second-order terms are typically small for wealth distributions that are as unequal as in the data.

Relaxing the assumption of differentiability. For now, we have focused on the common case in which the deviation in wealth due to the perturbation is differentiable in the perturbation size θ (Assumption 1). In reality, some perturbation may generate non-differentiable or event discontinuous changes in individual wealth. For instance, a change in corporate taxes may change the decision of being a worker and an entrepreneur, generating discontinuous changes in individual wealth. I discuss the effect of such counterfactuals on top wealth shares in Appendix A.3.

2 Application to models of random-growth

The previous section expressed deviations of the average wealth in top percentiles in terms of deviations in the wealth of top individuals. I now use this result to characterize the impulse response of top wealth shares to shocks in individual wealth.

I focus on random growth models inequality, which are commonly used in the literature as they can generate realistic distributions for highly skewed economic variables such as wealth, income, firm size or city size. Focusing on statistical models of wealth allows me to maintain a fairly large degree of generality while abstracting from specific economic considerations (as discussed below, a lot of different economic models can give rise to random-growth models). Section 2.1 studies the case where individual wealth follows a random walk with some probability of wealth dissipation (e.g. death). Section 2.2 studies the case where individual wealth follows a random walk plus some additive shock (e.g. labor income). Finally, Section 2.3 discusses my results in the context of the existing literature.

2.1 Model with dissipation shocks

Setup. In this section, I assume that individual wealth follows a random walk until they die (or, more generally, suffer a dissipation shock), in which case their wealth resets to an initial value. More precisely, I assume that individual wealth evolves as a random walk

$$W_{i,t+1} = A_{i,t+1} W_{i,t}, \tag{6}$$

until they suffer a "dissipation shock", in which case wealth is "reset" to some positive random variable $B_{i,t+1}$. One natural interpretation of this dissipation shock is death or population growth, but the concept is more general.[7] I assume that both $A_{i,t}$ and $B_{i,t}$ are positive almost-surely, which implies that wealth remains positive. Importantly, and this is a key generalization relative to existing random-growth models, I allow the joint distribution of the initial wealth $B_{i,t}$, the growth rate of wealth, $A_{i,t}$ and of the dissipation shock to be arbitrarily correlated over time, and, potentially, to be time and individual specific.[8]

Such processes have also been used to model the law of motion of wealth (e.g. Piketty and Zucman, 2015, Jones, 2015, Moll et al., 2022, Gomez and Gouin-Bonenfant, 2024), income (e.g. Jones and Kim, 2018), and firm size (e.g. Luttmer, 2011). One micro-foundation for such process is that agents have homothetic utilities, earn stochastic returns on their wealth $R_{i,t+1}$, and no labor income, in which case consumption is proportional to wealth.[9]

I now consider a perturbation in the realized path of individual growth rates $(\mathrm{d}\log A_{i,s})_{0\leq s\leq t}$. The next proposition applies the general equation (5) obtained in the previous section to this particular setup to characterize the resulting deviation in the average wealth in the top percentile:

Proposition 2. *Consider an arbitrary perturbation in the realization of growth rates $(\mathrm{d}\log A_{i,s})_{1\leq s\leq t}$ in (6). The resulting deviation of the average wealth in the top percentile is:*

$$\mathrm{d}\log \overline{W}_t(p) = \mathbb{E}^{W_{i,t}}\left[\sum_{s=\max(1,t-a_{i,t}+1)}^{t} \mathrm{d}\log A_{i,s}\middle| W_{i,t} \geq Q_t(p)\right],$$

where $\mathbb{E}^{W_{i,t}}$ denotes the wealth-weighted expectation with respect to the wealth distribution and $a_{i,t}$ denotes the age of the individual (number of years since last dissipation shock).

In words, this proposition says that the log deviation of the average wealth in the top percentile

[7] We will discuss more precisely how to treat bequests and dynasties in Section 3. See Moll et al. (2022) for other interpretations.

[8] The only important condition for the results below is that (i) a perturbation to log wealth at time s transmit one-to-one to time t in the absence of dissipation shock and (ii) the resulting distribution of individual wealth at time t satisfies the smoothness assumptions described in the previous section.

[9] More precisely, the law of motion of wealth is $W_{i,t+1} = R_{i,t+1}W_{i,t} - C_{i,t+1}$. With homothetic utilities, consumption is linear in wealth $C_{i,t+1} = \rho W_{i,t}$ and so we obtain (6) with $A_{i,t+1} \equiv R_{i,t+1} - \rho$. Alternatively, one can allow for the existence of tradable labor income, in which case the same result follows after redefining $W_{i,t}$ to be the sum of financial wealth and human capital. I refer the reader to these papers for more details on such micro-foundations.

is the wealth-weighted average of the sum of the past perturbation of growth rates experienced by individuals in the top percentile. This expression highlights the cumulative nature of wealth, which depends on all prior innovations. One key point is that the deviation in $\log \overline{W}_t(p)$ depends on *all* the previous perturbations in the growth rate of these top individuals, not just the ones that happened when they already were in the top percentile.[10]

The expression obtained in Proposition 2 is very general. To understand the intuition, it is useful to consider a few specific cases. For the sake of simplicity, I assume that the baseline distribution is stationary.[11]

One-time uniform perturbation in growth rates. I first focus on the case of a one time uniform increase in individual growth rates at time $t = 1$ by $\mathrm{d}\mu$; that is, $\mathrm{d}\log A_{i,t} = \mathrm{d}\mu 1_{t=1}$. In this case, Proposition 2 implies:

$$\mathrm{d}\log \overline{W}_t(p) = \mathbb{P}^{W_{i,t}}\left(a_i \geq t \middle| W_i \geq Q(p)\right)\mathrm{d}\mu.$$

The key takeaway is that the effect a one-time uniform perturbation depends on the fraction of households in the top percentile with an age higher than t. In particular, note that at $t = 1$ we get $\mathrm{d}\mu$ while at $t \to \infty$, we get zero. Note that, in the absence of dissipation shocks (e.g. infinite-horizon model), the effect of the shock would be permanent. Dissipation shocks, hence, are key to generate transitory effect of transitory perturbations in wealth dynamics.

Permanent uniform perturbation in growth rates. I now focus on the case of a permanent uniform increase in individual growth rates by $\mathrm{d}\mu$; that is, $\mathrm{d}\log A_{i,t} = \mu\,\mathrm{d}\theta$. This increases the log wealth of individual i at time t by $\min(t, a_{i,t})\,\mathrm{d}\mu$, where $\min(t, a_{i,t})$ corresponds to the number of periods in which individual i experienced these higher growth rates. The resulting deviation for the average wealth in a top percentile is a weighted average of these individual deviations:

$$\mathrm{d}\log \overline{W}_t(p) = \mathbb{E}^{W_i}\left[\min(t, a_i) | W_i \geq Q(p)\right]\mathrm{d}\mu. \tag{7}$$

The key takeaway is that the distribution of ages in the top percentile acts as a "sufficient statistic" for the effect of a uniform increase in growth rates on the average wealth in a top percentile. Put differently, while the exact specification of the model (e.g.the distribution of growth rates $A_{i,t}$ and dissipation shocks, or the extent to which they vary over time, between ages, or across types of agents) matters for the counterfactual effect of a change in individual growth rate on the wealth distribution, it only matters through the effect that it has on the joint distribution of wealth and age in the baseline economy.

This equation implies the following expression for the instantaneous effect ($t = 0$) and long-

[10] This is related to our discussion of (4) in the previous section.

[11] All results would go through with a time-varying baseline distribution, but one would need to distinguish t and τ, the number of years since the start of the perturbations.

run effect ($t \to \infty$) of this permanent uniform increase in growth rates:

$$\mathrm{d}\log \overline{W}_1(p) = \mathrm{d}\mu$$
$$\lim_{t\to\infty} \mathrm{d}\log \overline{W}_t(p) = \mathbb{E}^{W_i}\left[a_i | W_i \geq Q(p)\right] \mathrm{d}\mu$$

In words, the short-run elasticity of the average wealth in the top percentile is one while its long-run one is given by the average age in the top percentile. The higher the average age in the top, the longer individuals in the top benefit from these higher growth rates, and, therefore, the larger the resulting deviation in their wealth. The flip-side of this observation is that, the higher the average in the top, the longer it takes for top wealth shares to converge to their long-term values. Note that, in the absence of dissipation shocks (e.g. infinite-horizon model), the relative change in the average wealth in top percentiles after t period would simply equal $t\,\mathrm{d}\mu$; that is, it would grow without bounds with the horizon. Dissipation shocks, hence, are key to generate finite long-run elasticities of average wealth in top percentiles to changes in individual growth rates.[12]

I now turn to the implication of these results for the log deviation of the top wealth *share* $S_t \equiv p\overline{W}_t(p)/W_t(100\%)$. As discussed in the previous section, taking logarithms and differentiating gives that the log deviation of the top wealth share is simply given by the difference in the log deviation of the average wealth in the top percentile and the log deviation of the average wealth in the economy. As a result (7) implies

$$\mathrm{d}\log S_t(p) = \Big(\mathbb{E}^{W_i}\left[\min(t, a_i) | W_i \geq Q(p)\right] - \mathbb{E}^{W_i}\left[\min(t, a_i)\right]\Big)\,\mathrm{d}\mu.$$

In particular, on impact (at $t = 0$), we get $\mathrm{d}\log S_1(p) = 0$; that is, there is no effect on a uniform change in growth rates on top wealth shares. In contrast, in the long-run, $\lim_{t\to\infty} \mathrm{d}\log S_t(p) = \left(\mathbb{E}^{W_i}\left[a_i|W_i \geq Q(p)\right] - \mathbb{E}^{W_i}\left[a_i\right]\right)\mathrm{d}\mu$: a uniform increase in growth rate affects top wealth shares through the gradient of age with wealth. In random-growth models, agents at the top of the wealth distribution tend to have higher ages (since dissipation shocks reverts wealth back to its initial value), and so this quantity is typically positive. However, we will see below how the gradient of age differs drastically across models.

Type-specific perturbation. I discuss another special case that is commonly encountered in the inequality literature: an increase in the growth rate of only one type of agents (e.g. Luttmer, 2012 or Gabaix et al., 2016). The resulting deviation in the average wealth in the top percentile at t depends on the (wealth-weighted) number of periods that individuals in the top percentile have spent in this specific type starting from time $t = 0$. In particular, the short-run elasticity of the average wealth in the top percentile to a type-specific increase in growth rates is the fraction of individuals in this top percentile that are currently of this type. Its long-run elasticity is the

[12] The idea that dissipation shocks (e.g. overlapping generation models) are important to generate finite elasticities of capital to interest rates is well-known in the literature. See, for instance, Piketty and Saez (2013) or Moll et al. (2022).

average number of periods that individuals in the top percentile have spent in this type over their lifetimes.

Note that the average wealth in a top percentile responds more to a uniform increase in individual growth rates. However, the opposite may hold for top wealth *shares*; in particular, an increase in the growth rate of the type of agents making it to the top of the distribution will tend to generate faster and larger changes in top wealth *shares* relative to a uniform increase in the growth rate of all agents. This distinction is typically very important quantitatively.

2.2 Generalized models

The model discussed in the previous section is very stylized: on the one hand, it assumes perfect persistence of wealth within each lifetime, while, on the other hand, there is zero persistence of wealth across lifetimes. I now show that similar formulas obtain in more complex models.

Models with additive shocks. I now assume that individual wealth evolves according to a Kesten (1973) process:

$$W_{i,t+1} = A_{i,t+1} W_{i,t} + B_{i,t+1}. \tag{8}$$

This type of process has been used to analyze the distribution of city size (e.g. Gabaix, 1999) as well as the distribution of wealth (e.g. Benhabib et al., 2011, Gabaix et al., 2016, Saez and Zucman, 2016). I assume that $A_{i,t} > 0$, $B_{i,t} \geq 0$, and I allow the distribution of $(A_{i,t+1}, B_{i,t+1})$ to be serially correlated, and to potentially be time and individual specific.[13]

To make this process a strict generalization of the model above, one can also allow the existence of a "dissipation shock", in which case wealth at time t is simply reset to $B_{i,t}$.

One micro-foundation for such process, in the context of wealth, is to consider a model in which individual wealth evolves as $W_{i,t+1} = R_{i,t+1} W_{i,t} + Y_{i,t+1} - C_{i,t+1}$, where $R_{i,t+1}$ denotes the return on existing asset, $Y_{i,t+1}$ denotes labor income, and $C_{i,t+1}$ denotes consumption. Assuming that consumption is linear in wealth; that is, $C_{i,t+1} = a_{i,t+1} W_{i,t} + b_{i,t+1}$, the law of motion of individual wealth takes the form (8) with $A_{i,t+1} \equiv R_{i,t+1} - a_{i,t+1}$ and $B_{i,t+1} \equiv Y_{i,t+1} - b_{i,t+1}$.[14]

I now consider a sequence of perturbation in the realization of growth rates $(\mathrm{d} \log A_{i,s})_{0 \leq s \leq t}$.[15] The next proposition applies the general equation (5) obtained in the previous section to this particular setup to characterize the resulting deviation in the average wealth in the top percentile.

[13] Again, the only condition is that the resulting distribution of individual wealth at time t satisfies the smoothness assumptions described in the previous section.

[14] Such a model is used in Benhabib et al. (2011), Gabaix et al. (2016), and Saez and Zucman (2016), among others. In Benhabib et al. (2011), agents have homothetic utility, live for one period, and are endowed with a mix of previous cohort wealth $W_{i,t}$ and human capital $Y_{i,t+1}$, and so optimal consumption is proportional to $W_{i,t} + Y_{i,t+1}$. In Gabaix et al. (2016), consumption is assumed to be proportional to wealth. In Saez and Zucman (2016), consumption is assumed to be proportional to the sum of labor and capital income.

[15] It would be trivial to also incorporate deviation in the realization of the additive process $(\mathrm{d} \log B_{i,s})_{0 \leq s \leq t}$.

Proposition 3. *Consider an arbitrary perturbation in the realization of growth rates* $(\mathrm{d}\log A_{i,s})_{0\leq s\leq t}$ *in* (8)*. The resulting deviation of the average wealth in the top percentile is*

$$\mathrm{d}\log \overline{W}_t(p) = \mathbb{E}^{W_{i,t}}\left[\sum_{s=0}^{t}\left(\prod_{u=s}^{t}\left(1-\frac{B_{i,u}}{W_{i,u}}\right)\right)\mathrm{d}\log A_{i,s}\middle| W_{i,t}\geq Q_t(p)\right]. \tag{9}$$

As in Proposition 2, this proposition says that the log deviation of the average wealth in the top percentile is the wealth-weighted average of the sum of the past perturbation of growth rates experienced by individuals currently in the top percentile. The important distinction, however, is that the extent to which past wealth shocks are discounted over time depends on the relative importance of the additive force $B_{i,t}$ in the level of wealth $W_{i,t}$. Intuitively, the higher the magnitude of $B_{i,t}$ relative to wealth, the lower the effect of past shocks in wealth on current wealth.

At the top of the wealth distribution, one can expect the additive force to converge to zero relative to wealth, which means that the model generates substantially the same result as the model with only dissipation shocks.

General statistical processes. We can generalize the results obtained in the main text to much more general statistical processes for wealth. More precisely, under regularity conditions, the deviation for log wealth admits a moving-average (MA) representation; that is,

$$\mathrm{d}\log W_{i,t} = \sum_{s=0}^{t}\frac{\partial \log W_{i,t}}{\partial \log W_{i,s}}\,\mathrm{d}\epsilon_{i,s}.$$

where $\mathrm{d}\epsilon_{i,s}$ denotes the direct effect of the perturbation on log wealth at time s. The process $\partial\log W_{i,t}/\partial\log W_{i,s}$ is called the "first-derivative process" for log wealth, and it reflects the extent to which shocks in wealth persist over time. The assumption for this moving average representation to hold is that this process decays quickly enough. Combining this equation with Proposition 1 gives the resulting change of the average wealth in a top percentile as the cumulative sum of past perturbations in relative wealth times the first-derivative process for log wealth:

$$\mathrm{d}\log \overline{W}_t(p) = \mathbb{E}^{W_{i,t}}\left[\sum_{s=0}^{t}\frac{\partial \log W_{i,t}}{\partial \log W_{i,s}}\,\mathrm{d}\epsilon_{i,s}|W_{i,t}\geq Q_t(p)\right].$$

Different statistical processes for wealth then lead to different expressions for this first-derivative process. Figure 3 plots the first-derivative process for the two models discussed above as a function of s. Both decay over time, reflecting the fact that the persistence of wealth shocks is less than one. In the first model, this is due to the dissipation shocks, while, in the second model, this is due to the presence of additive shocks.

Note that, in a model in which portfolio returns increase with wealth (as in the model with non-homothetic preferences in Gaillard et al., 2023), the first-derivative process for log wealth may be higher than one for some time, before decaying due to death or some other negative force.

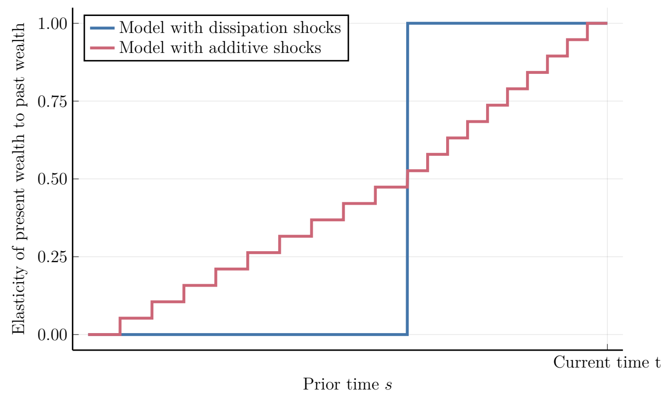

Figure 3: Elasticity of present wealth (time t) to prior wealth (time s) in random-growth models.

Notes: The figure plots the elasticity of wealth at time t with respect to wealth at time $s < t$, $\partial \log W_{i,t}/\partial \log W_{i,s}$, as a function of s. In the model where wealth follows a random walk with dissipation shocks (6), this elasticity is given by $\partial \log W_{i,t}/\partial \log W_{i,s} = 1_{s \geq t-a_{i,t}}$. In the model where wealth follows a random walk with additive shocks (8), this elasticity is given by $\partial \log W_{i,t}/\partial \log W_{i,s} = \prod_{u=s+1}^{t}(1 - B_{i,s}/W_{i,s})$. The area under these curves determines the elasticity of current wealth to a uniform change in past growth rates.

2.3 Comparison with results on tail indices

The core result of this section was the derivation of closed-form formulas for the impulse response of top wealth shares to perturbations in the dynamics in individual wealth. I now discuss how these formulas relate to the existing results in the literature on random-growth models. In particular I focus on two types of results: (i) analytical characterization of the tail index of a distribution, and (ii) analytical characterization of the transition dynamics of a distribution toward its steady-state.

Results on tail indices. I now briefly emphasize the advantage of my approach relative to the existing literature on inequality, which tends to focus on analytical characterization for tail indices. When wealth follows a random growth model of the type discussed in Section 2.1 or 2.2, one can often show that, under additional condition of stationarity, the distribution of wealth has a right Pareto tail; that is,[16]

$$\log \overline{W}(p) \sim -(1/\zeta) \times \log p \quad \text{as } p \to 0. \tag{10}$$

where ζ denotes the tail index of the wealth distribution. Consider a permanent perturbation in the distribution of these growth rates starting from time $t = 0$. In this case, one can typically show that the perturbed wealth distribution remains Pareto and obtain an analytical formula for the resulting deviation in the tail index (see Appendix B.1 for examples). This analytical characteriza-

[16] Here, and in the rest of the paper, I use $f(p) \sim g(p)$ as $p \to 0$ to denote $\lim_{p \to 0} f(p)/g(p) = 1$.

tion implies the following characterization for top wealth shares:[17]

$$\lim_{t\to\infty} \mathrm{d}\log \overline{W}_t(p) \sim -\mathrm{d}\log\zeta \times \log \overline{W}(p) \quad \text{as } p \to 0, \tag{11}$$

The key takeaway is that, while a deviation in the tail index is informative on the asymptotic behavior of top wealth shares, it only characterizes the asymptotic deviation in top wealth shares in the long-run ($t \to \infty$) and in the right tail ($p \to 0$).[18] This is much less precise than the expressions obtained in Proposition 2 or Proposition 3, that characterize the deviation of the average wealth in the top at *any* given time horizon t and at *any* given top percentile p. In fact, I discuss in Appendix B.1 how focusing on the response of the tail index can be very uninformative about the actual response of top wealth shares, even in the long-run. In particular, my analytical results characterize the response of top wealth shares to the wide range of perturbations that affect top wealth shares without affecting tail indices (e.g., perturbations that only impact individuals for a finite amount of time periods) and they do not require that the baseline wealth distribution is stationary nor that it obeys a Pareto shape.

Results on transition dynamics. Luttmer (2012) and Gabaix et al. (2016) study the transition dynamics of the wealth distribution following (non-infinitesimal) changes in the law of motion of individual wealth. However, these papers are only able to characterize the rate of decay of moments of the wealth distribution or of its L^1 norm. In contrast, my results directly characterize the dynamics of top wealth shares.[19]

The results in my paper sheds light on the main findings of these papers, which is that,, while random-growth models with uniform types take a long time to converge after a uniform perturbation in growth rates, models where some agents grow very fast do a better job. My results shed light on why: as seen in Equation 7, the lower the age at the top, the faster the average wealth converges to its long-run value. However, the flipside of this is that, the lower the age at the top, the less there is an increase in the average wealth in the top percentile relative to the economy. Put differently, the average age at the top controls *both* the transition speed of a top percentile following a permanent deviation in growth rates and its long-run deviation.

Hence, this result links two important observations in the inequality literature: Gabaix et al. (2016)'s emphasis that standard random growth models generate slow transition dynamics that are too slow relative to the data and Luttmer (2011)'s emphasis that they imply an average age in top percentile that is too high relative the data.

[17]Formally, (10) implies $\lim_{t\to\infty} \mathrm{d}\log \overline{W}_t(p) \sim -(\zeta^{-2}\,\mathrm{d}\zeta)\log p \sim -(\zeta^{-1}\,\mathrm{d}\zeta)\log \overline{W}(p)$ as $p \to 0$.

[18]If the wealth distribution is *exactly* Pareto, then there is a one-to-one mapping between the tail index and the *level* of top wealth shares. However, most existing models of wealth inequality only imply that the wealth distribution has a Pareto *tail*, consistently with the data. One exception is models where individual wealth follows a random walk with a reflecting barrier at some lower level of wealth.

[19]One downside of this approach, however, it that it only characterizes small deviations in the wealth distribution.

Determinants of age. The formulas obtained in this section represent the dynamics of top wealth shares in terms of endogenous quantities (the distribution of age across the wealth distribution, rather than the distribution of growth). I explore the determinants of these endogenous quantities in Appendix B.1. I show that, in a large class of random-growth models, the average age increases linearly with log wealth, with a slope equal to the inverse of the derivative of the cumulant generating function (CGF) of the growth rate of wealth taken at the tail index. Intuitively, the more convex the CGF is, the higher the dispersion in growth rates across individuals, and, therefore, the less time it takes for the most successful individuals to reach top percentiles. In particular, a model in which the growth rate of individual wealth has a high skewness or kurtosis (potentially because some individuals experience higher growth rates over long periods of time, a la Luttmer, 2011) is a model in which the luckiest households reach the right tail of the distribution pretty quickly, and, therefore, in which top wealth shares respond less to a given increase in individual growth rates.

2.4 Comparison with calibrated models.

I now report the distribution of age in existing calibrated models in the literature. I focus on examining three models of inequality, which are all particular instances of the model with dissipation shocks discussed in Section 2.1: Gabaix et al. (2016), Moll et al. (2022) and Gomez and Gouin-Bonenfant (2024).

Table 1: Age distribution in existing models

	Top 100%	Top 0.01%
Moll et al. (2022)	26	134
Gabaix et al. (2016)	37	85
Gomez and Gouin-Bonenfant (2024)	10	22

Notes: The table reports the wealth-weighted average effect of a 1pp. uniform increase for individual wealth, which corresponds to the wealth weighted average age in each top percentile (i.e. number of periods since last dissipation shock), see Proposition 2.

The table reports the average age of households across the wealth distribution. The average "age" of households in the top 0.01% (number of periods since dissipation shock) ranges from 22 in Gomez and Gouin-Bonenfant (2024) to 134 in Moll et al. (2022). Given the formulas obtained above, this implies that the long-run elasticity of top wealth shares to a uniform change in the growth rate growth rate of individuals differs by an order of magnitude across models.

The key takeaway is that, despite the fact that these models match similar cross-sectional moments about the wealth distribution (e.g. level of top wealth shares), they have vastly different implications on the relationship between wealth and age, and therefore, on the effect of a given increase in individual growths rate on inequality.

3 Application to calibrated models of wealth inequality

Up to now, my framework described the impulse response of top shares for any distribution. Because I focused on statistical models of inequality, these results can apply to distribution of labor income, city size, or wealth distribution. I now focus on the wealth distribution. Here, I use my framework to dissect the effect of asset price shocks on the wealth distribution in a state-of-the-art model with heterogeneous agents.

3.1 Model description

I consider the following law of motion of wealth for household i between t and $t+1$ is:

$$W_{i,t+1} = R_{i,t+1}W_{i,t} + B_{i,t+1} + Y_{i,t+1} - C_{i,t+1}, \tag{12}$$

where $R_{i,t+1} > 0$ denotes the return on wealth, $Y_{i,t+1} > 0$ denotes labor income, $B_{i,t+1} > 0$ denotes bequests received, and $C_{i,t+1}$ denotes consumption. With some probability, individuals die, at which point their wealth is redistributed as inheritance to their children.

Linearizing this equation and solving backward gives the deviation in log wealth at some point t as the cumulative sum of past perturbations in returns, labor income, and bequests.

Proposition 4. *The effect of a one-time asset price shock on the average wealth in the top percentile is given by:*

$$\mathrm{d}\log \overline{W}_t(p) = \mathbb{E}^{W_i}\left[\left(\prod_{s=\max(0,t-a_i-1)}^{t}\left(1 - \frac{B_{i,s}+Y_{i,s}-C_{i,s}}{W_{i,s}} + \mathrm{MPC}_{i,c}\right)\right)\alpha_{i,0}\middle| W_i \geq Q(p)\right]\mathrm{d}\mu,$$

where $\mathrm{MPC}_{i,s} \equiv \partial C_{i,s}/\partial W_{i,s}$ *denote the average and marginal propensities to consume out of wealth, respectively.*

This proposition expresses the deviation in consumption at some present time as a sum of the direction deviations in previous wealth. That the persistence of these perturbations on wealth depend on the relative importance of labor and bequest on wealth, as well as the marginal propensity to consume out of wealth. Assuming that the remaining perturbation are not causally related, this expression allows one to trace out the effect of returns, labor income or bequests using the historical path of the labor income and bequest to wealth ratios of households.

There is a normal decay over time due to the importance of labor income relative to financial wealth. There is also a discrete discounting that happens when receiving bequest: receiving a large inheritance decreases the effect of the household own returns received prior to this inheritance while increasing the effect of the parents own returns. Finally, note that the model implies that the persistence of returns is very small for individuals at the bottom of the wealth distribution, but very high for individuals at the top of the wealth distribution (for which labor income represents an infinitesimal fraction of current wealth.)

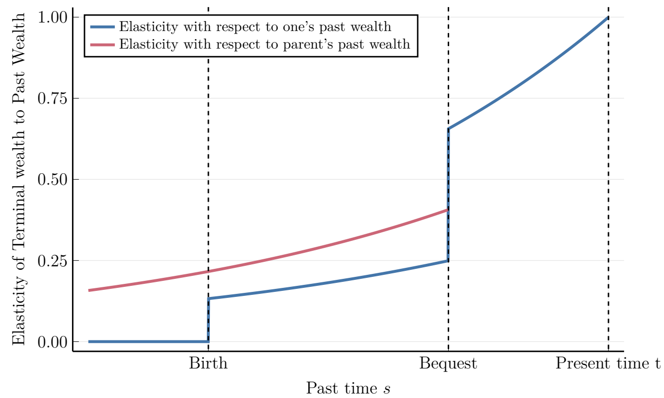

Figure 4: Elasticity of present wealth (time t) to a shock in prior wealth (time s) in the model.

Notes: The figure plots the elasticity of wealth of household i at time t with respect to wealth of household j at time $s < t$, $\partial \log W_{i,t} / \partial \log W_{js}$, as a function of s. In blue is the case $j = i$ (household's own wealth), and in red is the case where j leaves a bequest to i (parents' wealth). The sum of the area under these curves can be interpreted as the elasticity of current wealth to a uniform change in past growth rates.

3.2 Data construction

Proposition 4 allows us to compute the effect of changes in returns on changes in household wealth as long as we can observe the historical path of their labor income, inheritance and wealth. I now discuss the data I use to construct a measure of the lifetime path of labor income, wealth, bequest received by households in top percentiles for the U.S. from a combination of the Survey of Consumer Finances (SCF) from 1989 to 2016 and Forbes 400 in years.

Since the SCF does not report the previous income and wealth for individuals currently in top percentiles, I use a "synthetic cohort" approach where I construct the average income and wealth across top percentiles within each cohort, and I then construct the lifetime path of each agent by assuming that agents remain in their relative percentiles over their lifetimes (as in the "pseudo-panel" approach of Feiveson and Sabelhaus, 2019). I set the labor income of households in Forbes 400 to zero. I define the age of a household as the age of the household head minus 25.

Finally, I construct a measure of bequest received across top percentiles by using the total bequest distributed over their lifetimes as reported in SCF (which reports the amount and the year of the three biggest inheritance or in-vivo transfer received over the household's lifetime). Finally, to compute the average labor income to wealth ratio of parents, I assume that, with probability half, they are in the same top percentile (within their cohorts) as their children, while, with probability half, they are randomly drawn from the population. For households in Forbes, I classify individuals manually into three categories: self-made, heirs, and in-between. I assume that the inheritance to wealth ratio at the time of inheritance is zero for self-made fortunes (60% of observations), 1 for heirs (24% of observations), and 0.5 for people in between (16% of observations).

3.3 One-time increase in asset returns

I now use the methodology to compute the impulse response of top wealth shares to a one-time increase in asset returns (as opposed to a permanent increase in returns). One key result is that these shocks tend to have very persistent effect on top wealth shares.

Figure 5a reports the effect of a one-time 10% increase in equity returns. The short-run impact at $t = 0$ is large, which reflects the fact that households in top percentiles are more exposed to equity returns than the average household in the economy. My estimates are consistent with the reduced-form exposures from Kuhn et al., 2020 and Gomez, 2017. The key contribution of this methodology is to evaluate the longer-run effect of these asset returns on top wealth shares. In particular, one striking fact is how persistent the response is for Forbes 400. This reflects the fact that a substantial portion of individuals in Forbes 400 inherit their wealth; as a result, the impact of increased equity returns extend beyond the lifetimes of the original equity holders.

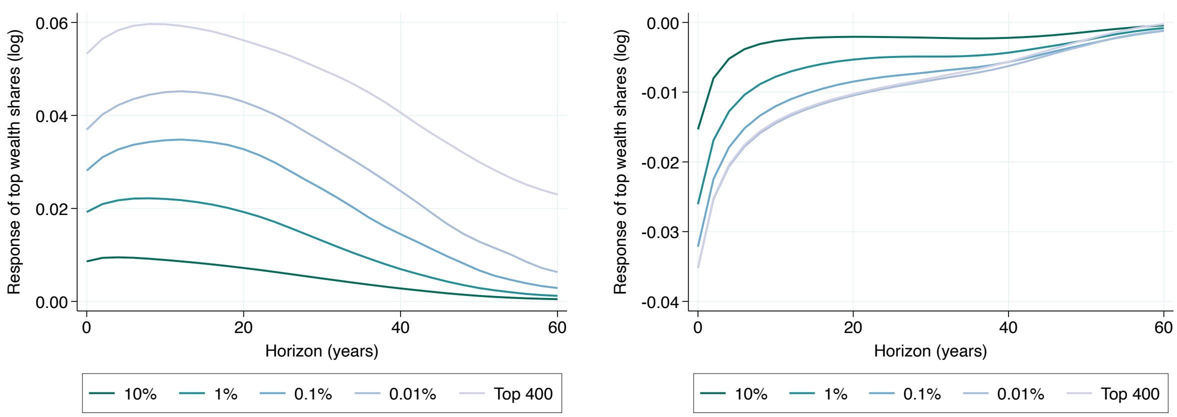

(a) 10pp. increase in equity returns

(b) 10pp. increase in housing returns

Figure 5: Effect of a one-time increase in asset returns on top wealth shares

Figure 5b reports the effect of a one-time 10% increase in housing returns. Top wealth shares drop, which reflects the fact that top households hold less housing than the rest of the distribution (see, for instance, Kuhn et al., 2020, Greenwald et al., 2022, Martínez-Toledano, 2020, ...). However, the effect of an increase in housing returns is much less persistent than the effect of equity returns. This reflects the fact that wealth shocks tend to be much less persistent at the bottom of the wealth distribution, as a larger fraction of the wealth for the average household in the economy comes from capitalized labor income, which dampens the effect of past returns on capital on current wealth (see Proposition 4).

3.4 Permanent increase in asset returns

To visualize the data construction, Table 2 reports the implied long-run elasticity of wealth to changes in individual growth rate, computed recursively from (4). If individuals receive no inheritance and if labor income is an infinitesimal fraction of their wealth, this corresponds literally to the age of individuals, as in the model with dissipation shocks discussed in Section 2.1. Relative to this baseline, the presence of labor income tends to decrease the persistence of wealth shocks while the presence of bequests tends to increase them.

Table 2: Long-run elasticity of wealth to returns implied by the data

	Household age (Baseline)	Taking into account bequests	Taking into account bequests & labor income
Top 100%	33	37	26
Top 10%	34	38	28
Top 1%	35	38	31
Top 0.1%	36	40	35
Top 0.01%	37	44	39
Top 400	39	59	49

Notes: The table reports summary statistics on the wealth-weighted distribution of age in top percentiles of the U.S. economy. Data from the Survey of Consumer Finances and Forbes 400 from 1989 to 2016.

The first column of the table reports the elasticity from (4) that one would obtain without taking into account inheritance of labor income; i.e., defining the age of the household as the household head's age minus 25, as discussed above. Note that the gradient of average age across the distribution is almost zero. This almost zero gradient implies that without bequests and labor income, changes in individual returns would not affect top wealth shares. The second column of the table adjusts for bequests, i.e. reports the elasticity obtained from (4) while still setting $Y_{i,u} = 0$. This increases the elasticity of wealth at the top since richer households tend to receive larger bequests. The third column of the table then also accounts for the relative importance of labor income by reporting the full formula (4). This tends to decrease the average elasticity of wealth at the bottom since labor income constitutes a large fraction of their wealth every period. Overall, I find that the gradient of "age" becomes much steeper once one accounts for inheritances and labor income.

I now briefly compare my estimates for the long-run elasticity of individual wealth to changes in individual growth rates to the range of estimates implied by the existing literature discussed in Section 3. In average, my estimates are lower then existing calibrations. One way to interpret this evidence is that Moll et al. (2022) and Gabaix et al. (2016) effectively overestimate the role of dynastic wealth in top percentiles while Gomez and Gouin-Bonenfant (2024) underestimate it relative to the U.S. data.

As summarized by Piketty and Zucman (2015), most models of wealth inequality imply that an increase in the average return of wealth increases wealth inequality (or, at least, the thickness of its right tail). However, there is still substantial uncertainty on how large the effect of an increase in the average return on wealth on top wealth inequality actually is. I now use (4) to compute the response of top wealth shares to a permanent increase in asset returns.

Figure 6a plots the result for a 1pp uniform increase in the growth rate of wealth. In the first few years, the effect of an increase in the return on wealth is almost zero. As discussed in Section 2.1, this occurs because all individuals in the economy experience this higher growth rate; hence, top wealth *shares* do not react. Still, top wealth shares start increasing after a few years. There is "fanning-out" of top wealth shares which reflects the fact that agents at the top of the wealth distribution benefit relatively more from higher growth rates, as reported in Table 2. The underlying reason is that the effect of a given increase in asset returns is much more persistent for agents at the top of the wealth distribution than for the average agent in the economy. Overall, I find that, after 60 years, a 1pp increase in the return on individual wealth increases the share of wealth owned by the top 0.01% by 0.1 log points. Note, however, that this increase is relatively small compared to the overall rise in top wealth inequality since 1980, which amounts to approximately 1 log point (Saez and Zucman, 2022)

Figure 6b plots the increase in the average return of only one type of assets, public and private equity. Relative to the previous experience, this increases the return of each agent by 1pp times the share of wealth invested in equity. Relative to the previous experiment, we observe an immediate increase in top wealth shares, reflecting the fact that individuals at the top of the wealth distributions disproportionately hold equity. We can see that, relative to the previous experiment, top wealth share increase faster and with a bigger magnitude. Overall, a 1pp increase in the return on equity increases the top 0.01% by 0.2 log points, which is twice as much as the effect of a 1pp increase in the overall return on wealth.

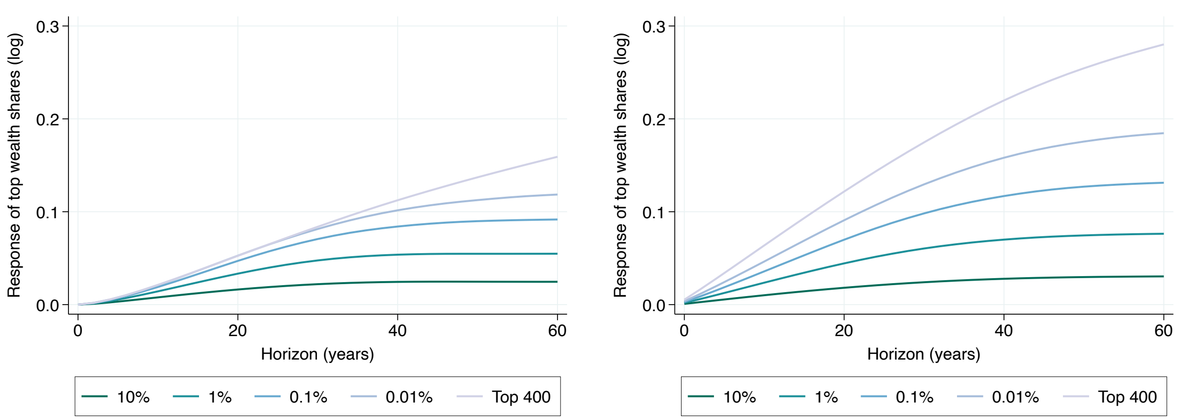

(a) 1pp. increase in returns on all assets

(b) 1pp. increase in returns on equity

Figure 6: Effect of a permanent increase in returns on top wealth shares

Conclusion

The key insight of this paper is that, for a wide range of counterfactuals, the first-order response of top wealth shares at some horizon t is determined by the wealth response at time t of individuals projected to be in the top percentile at time t. I leverage this insight to obtain clean formulas for the impulse response of top wealth shares in a wide range of random-growth models. I then use these results as a diagnostic tool to analyze counterfactuals implied by existing models of wealth inequality. I also use them as as an accounting framework to compute the dynamics of top wealth shares under various counterfactual scenarios.

Overall, this paper provides a flexible and transparent methodology to analyze the effects of economic shocks on wealth inequality. This methodology could help both reduced-form and structural approaches to strengthen the credibility of computing responses of the wealth distribution to counterfactual changes in the economic environment.

References

Atkeson, Andrew G and Magnus Irie, "Rapid Dynamics of Top Wealth Shares and Self-Made Fortunes: What Is the Role of Family Firms?," *American Economic Review: Insights*, 2022, *4* (4), 409–24.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub, "Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models," *Econometrica*, 2021, *89* (5), 2375–2408.

__ , Matthew Rognlie, and Ludwig Straub, "The intertemporal keynesian cross," Technical Report, National Bureau of Economic Research 2018.

Baqaee, David and Elisa Rubbo, "Micro propagation and macro aggregation," *Annual Review of Economics*, 2023, *15*, 91–123.

Beare, Brendan K and Alexis Akira Toda, "Determination of Pareto exponents in economic models driven by Markov multiplicative processes," *Econometrica*, 2022, *90* (4), 1811–1833.

Benhabib, Jess, Alberto Bisin, and Mi Luo, "Wealth distribution and social mobility in the US: A quantitative approach," *American Economic Review*, 2019, *109* (5), 1623–47.

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

Chetty, Raj, "Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods," *Annu. Rev. Econ.*, 2009, *1* (1), 451–488.

Costinot, Arnaud and Andrés Rodríguez-Clare, "Trade theory with numbers: Quantifying the consequences of globalization," in "Handbook of international economics," Vol. 4, Elsevier, 2014, pp. 197–261.

de Saporta, Benoite, "Tail of the Stationary Solution of the Stochastic Equation $Y_{n+1} = aY_n + b_n$ with Markovian coefficients," *Stochastic Processes and their Applications*, 2005, *115*, 1954–1978.

DiNardo, John, Nicole Fortin, and Thomas Lemieux, "Labor market institutions and the distribution of wages, 1973-1992: A semiparametric approach," 1995.

Feiveson, Laura and John Sabelhaus, "How does intergenerational wealth transmission affect wealth concentration?," 2018.

— **and** — , "Lifecycle patterns of saving and wealth accumulation," 2019.

Firpo, Sergio, Nicole M Fortin, and Thomas Lemieux, "Unconditional quantile regressions," *Econometrica*, 2009, *77* (3), 953–973.

Fortin, Nicole, Thomas Lemieux, and Sergio Firpo, "Decomposition methods in economics," in "Handbook of labor economics," Vol. 4, Elsevier, 2011, pp. 1–102.

Gabaix, Xavier, "Zipf's Law for Cities: an Explanation," *Quarterly Journal of Economics*, 1999, pp. 739–767.

— **, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll**, "The dynamics of inequality," *Econometrica*, 2016, *84* (6), 2071–2111.

Gaillard, Alexandre, Christian Hellwig, Philipp Wangner, and Nicolas Werquin, "Consumption, Wealth, and Income Inequality: A Tale of Tails," 2023.

Gomez, Matthieu, "Asset prices and wealth inequality," *Working Paper*, 2017.

— , "Decomposing the growth of top wealth shares," *Econometrica*, 2023, *91* (3), 979–1024.

— **and Émilien Gouin-Bonenfant**, "Wealth inequality in a low rate environment," *Econometrica*, 2024, *92* (1), 201–246.

Greenwald, Daniel L, Matteo Leombroni, Hanno Lustig, and Stijn Van Nieuwerburgh, "Financial and total wealth inequality with declining interest rates," 2022. Working Paper.

Guvenen, Fatih, Gueorgui Kambourov, Burhan Kuruscu, Sergio Ocampo, and Daphne Chen, "Use it or lose it: Efficiency and redistributional effects of wealth taxation," *The Quarterly Journal of Economics*, 2023, *138* (2), 835–894.

Hubmer, Joachim, Per Krusell, and Anthony A Smith Jr, "Sources of US wealth inequality: Past, present, and future," *NBER Macroeconomics Annual*, 2021, *35* (1), 391–455.

Jones, Charles I, "Pareto and Piketty: The macroeconomics of top income and wealth inequality," *The Journal of Economic Perspectives*, 2015, *29* (1), 29–46.

— **and Jihee Kim**, "A Schumpeterian model of top income inequality," *Journal of political Economy*, 2018, *126* (5), 1785–1826.

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.

Kesten, Harry, "Random difference equations and renewal theory for products of random matrices," 1973.

Kleven, Henrik J, "Sufficient statistics revisited," *Annual Review of Economics*, 2021, *13*, 515–538.

Kuhn, Moritz, Moritz Schularick, and Ulrike I Steins, "Income and wealth inequality in America, 1949–2016," *Journal of Political Economy*, 2020, *128* (9), 3469–3519.

Luttmer, Erzo GJ, "On the mechanics of firm growth," *The Review of Economic Studies*, 2011, *78* (3), 1042–1068.

— , "Slow convergence in economies with firm heterogeneity," *Federal Reserve Bank of Minneapolis working paper*, 2012, (696).

— , "Further notes on micro heterogeneity and macro slow convergence," *Working Paper*, 2016.

Martínez-Toledano, Clara, "House price cycles, wealth inequality and portfolio reshuffling," *WID. World Working Paper*, 2020, (2020/02).

Mian, Atif R, Ludwig Straub, and Amir Sufi, "The saving glut of the rich," 2020. Working paper.

Moll, Benjamin, Lukasz Rachel, and Pascual Restrepo, "Uneven growth: automation's impact on income and wealth inequality," *Econometrica*, 2022, *90* (6), 2645–2683.

Nardi, Mariacristina De, "Wealth inequality and intergenerational links," *The Review of Economic Studies*, 2004, *71* (3), 743–768.

Ozkan, Serdar, Joachim Hubmer, Sergio Salgado, and Elin Halvorsen, "Why Are the Wealthiest So Wealthy? A Longitudinal Empirical Investigation," *A Longitudinal Empirical Investigation (March 9, 2023)*, 2023.

Peter, Alessandra, "Owning Up: Closely Held Firms and Wealth Inequality," *JMP Stan*, 2019.

Piketty, Thomas, *Capital in the twenty-first century*, Harvard University Press, 2014.

__ **and Emmanuel Saez**, "A theory of optimal inheritance taxation," *Econometrica*, 2013, *81* (5), 1851–1886.

__ **and Gabriel Zucman**, "Wealth and Inheritance in the Long Run," *Handbook of Income Distribution*, 2015, *2*, 1303–1368.

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. Updated series (February 2022 vintage) can be downloaded from https://gabriel-zucman.eu/usdina/.

__ **and** __ , "Progressive wealth taxation," *Brookings Papers on Economic Activity*, 2019, *2019* (2), 437–533.

__ **and** __ , "Top wealth in america: A reexamination," Technical Report, National Bureau of Economic Research 2022.

Touchette, Hugo, "The large deviation approach to statistical mechanics," *Physics Reports*, 2009, *478* (1-3), 1–69.

Wold, Herman OA and Peter Whittle, "A model explaining the Pareto distribution of wealth," *Econometrica, Journal of the Econometric Society*, 1957, pp. 591–595.

Appendix

A Appendix for Section 1

A.1 Proof of Proposition 1

Proof of Proposition 1. The proof is in two steps. In the first, I relate the deviation of density of wealth to the deviations in individual wealth. For any function $f \in \mathcal{C}^1$ defined on $\mathbb{R}$ that decays quickly enough with wealth, the cross-sectional average of $f(W_{i,t})$ in the counterfactual economy can be expressed in two ways:

$$\int_{\mathbb{R}} f(W)\left(g_t(W) + \Delta g_t(W)\right) dW = \int_{\mathbb{R}} \mathbb{E}[f(W_{i,t} + \Delta W_{i,t})|W_{i,t} = W] g_t(W) dW,$$

where g_t denotes the density of wealth at time t in the economy and d denotes the usual differential operator with respect to wealth (*not* the differential operator with respect to the size of the perturbation). Subtracting both sides by the cross-sectional average of $f(W_{i,t})$ in the baseline economy gives:

$$\int_{\mathbb{R}} f(W)\Delta g_t(W) dW = \int_{\mathbb{R}} \mathbb{E}[f(W_{i,t} + \Delta W_{i,t}) - f(W_{i,t})|W_{i,t} = W] d\mathbb{P}(W_{i,t} \leq W).$$

Dividing by θ and taking the limit $\theta \to 0$ gives

$$\int_{\mathbb{R}} f(W)\, \mathrm{d}g_t(W) dW = \int_{\mathbb{R}} f'(W)\mathbb{E}[\mathrm{d}W_{i,t}|W_{i,t} = W] g_t(W) dW. \tag{13}$$

Integrating by parts gives

$$\int_{\mathbb{R}} f'(W)(-\,\mathrm{d}G_t(W)) dW = \int_{\mathbb{R}} f'(W)\mathbb{E}[\mathrm{d}W_{i,t}|W_{i,t} = W] g_t(W) dW.$$

Since this holds for any continuous function f', this must hold at every point, which implies

$$\mathrm{d}G_t(W) = -\mathbb{E}[\mathrm{d}W_{i,t}|W_{i,t} = W] g_t(W) \tag{14}$$

where G_t denotes the cumulative distribution function for wealth at time t. Note that this equation is similar in spirit to Kolmogorov-Forward equation, except that instead of characterizing the derivative of the CDF with respect to time t, it characterizes the derivative of the CDF with respect to the perturbation size θ.

In the second step of the proof, I use this result to obtain the deviation in the average wealth in the top percentile. Applying the implicit function theorem on the definition of the top quantile (1) gives:

$$\mathrm{d}Q_t(p) = -\frac{1}{g_t(Q_t(p))} \int_{Q_t(p)}^{\infty} \mathrm{d}g_t(W) dW. \tag{15}$$

Differentiating the definition of the average wealth in the top percentile (2) gives:

$$\begin{aligned}
\mathrm{d}\overline{W}_t(p) &= \frac{1}{p}\int_{Q_t(p)}^{\infty} W\, \mathrm{d}g_t(W)dW - \frac{1}{p}Q_t(p)g_t(Q_t(p))\,\mathrm{d}Q_t(p) \\
&= \frac{1}{p}\int_{Q_t(p)}^{\infty} (W - Q_t(p))\,\mathrm{d}g_t(W)dW \\
&= -\frac{1}{p}\int_{Q_t(p)}^{\infty} \mathrm{d}G_t(W)dW.
\end{aligned}$$

where the last line uses integration by part. Substituting out $\mathrm{d}G_t(W)$ using (14) gives the result

$$\mathrm{d}\overline{W}_t(p) = \frac{1}{p}\int_{Q_t(p)}^{\infty} \mathbb{E}[\mathrm{d}W_{i,t}|W_{i,t} = W]g_t(W)dW.$$

Additionally, combining (15) with (14) gives a similar formula for the infinitesimal deviation in quantiles:

$$\mathrm{d}Q_t(p) = \mathbb{E}[\mathrm{d}W_{i,t}|W_{i,t} = Q_t(p)].$$

□

A.2 Higher-order terms

Non-infinitesimal deviation. I now derive a second-order approximation for the deviation in the average wealth in the top percentile. One can integrate the expression obtained in Proposition 1 to obtain an expression for a non-infinitesimal change in policy θ

$$\Delta\overline{W}_t(p) = \int_{u=0}^{\theta} \mathbb{E}\left[\mathrm{d}W_{i,t}(u)|W_{i,t}(u) \geq Q_t(p,u)\right]\mathrm{d}u. \tag{16}$$

This expression differs from the first-order approximation obtained in Proposition 1 as the set of individuals in the top percentile (i.e., for which $W_{i,t}(u) \geq Q_t(u)$) changes along the policy path.

Second-order approximation. One can approximate the integral (16) using the trapezoid rule to obtain a *second-order* approximation for the counterfactual change in top shares:

$$\Delta\overline{W}_t(p) = \frac{1}{2}\Big(\mathbb{E}\left[\Delta W_{i,t}|W_{i,t} \geq Q_t(p)\right] + \mathbb{E}\left[\Delta W_{i,t}|W_{i,t} + \Delta W_{i,t} \geq Q_t(p) + \Delta Q_t(p)\right]\Big) + o(\theta^2). \tag{17}$$

This second-order approximation expresses the change in the average wealth in the top percentile as the average of two terms: the average wealth deviation of households who are in the percentile in the baseline economy, $\mathbb{E}\left[\Delta W_{i,t}|W_{i,t} \geq Q_t(p)\right]$ and the average wealth deviation of households who are in top percentile in the perturbed economy, $\mathbb{E}\left[\Delta W_{i,t}|W_{i,t} + \Delta W_{i,t} \geq Q_t(p) + \Delta Q_t(p)\right]$. It is straightforward to rewrite (17) to obtain:[20]

$$\Delta\overline{W}_t(p) = \mathbb{E}\left[\Delta W_{i,t} | W_{i,t} \geq Q_t(p)\right] + \frac{1}{2}\Big(\mathbb{E}\left[\Delta W_{i,t} | W_{i,t} + \Delta W_{i,t} \geq Q_t(p) + \Delta Q_t(p)\right] - \mathbb{E}\left[\Delta W_{i,t} | W_{i,t} \geq Q_t(p)\right]\Big) + o(\theta^2).$$

The first term corresponds to the intensive margin defined in (3); hence, the second term corresponds to a leading-order approximation of the extensive margin defined in (3). It is equal to half of the difference between the average deviation of people that enter the top percentile in the perturbed economy, relative to the baseline economy, and people that exit it. The next proposition gives an analytical expression for this term in terms of the variance of the deviation in wealth across individuals.

Proposition 5 (Second-Order Approximation)**.** *Assume that the function $\theta \to W_{i,t}(\theta)$ is in $\mathcal{C}^2$ and that $W \to (\mathbb{E}[\partial W_{i,t}], \mathbb{E}[\partial W_{i,t}^2])$ is continuous. Then, the deviation in the average wealth in a top percentile is:*

$$\Delta\overline{W}_t(p) = \underbrace{\mathbb{E}[\mathrm{d}W_i + \frac{1}{2}\mathrm{d}^2 W_i | W_{i,t} \geq Q_t(p)]}_{\textit{Intensive margin}} + \underbrace{\frac{1}{2}\frac{g_t(Q_t(p))}{p}\mathbb{V}[\mathrm{d}W_i | W_{i,t} = Q_t(p)] + o(\theta^2)}_{\textit{Extensive margin}},$$

where $\mathbb{V}$ denotes the cross-sectional variance of the deviation in individual wealth.

Proof of Proposition 5. Differentiating (13) at the second order in θ gives

$$\int_{\mathbb{R}} f(W)\,\mathrm{d}^2 g_t(W) dW = \int_{\mathbb{R}} \Big(f'(W)\mathbb{E}[\mathrm{d}^2 W_{i,t} | W_{i,t} = W] + \frac{1}{2} f''(W)\mathbb{E}[(\mathrm{d}W_{i,t})^2 | W_{i,t} = W] \Big) g_t(W) dW.$$

Integrating by parts the right hand side gives

$$\int_{\mathbb{R}} f(W)\,\mathrm{d}^2 g_t(W) dW = \int_{\mathbb{R}} f(W) \Big\{ -\partial_W \left(\mathbb{E}[\mathrm{d}^2 W_{i,t} | W_{i,t} = W] g_t(W)\right) + \frac{1}{2}\partial_{WW}\Big(\mathbb{E}[(\mathrm{d}W_{i,t})^2 | W_{i,t} = W] g_t(W)\Big) \Big\} dW.$$

[20] Comparing this term with the exact extensive margin defined in (3) implies that the reminder — the term denoted by $o(\theta^2)$ — is:

$$\mathbb{E}\left[W_{i,t} + \frac{1}{2}\Delta W_{i,t} \Big| W_{i,t} + \Delta W_{i,t} \geq Q_t(p) + \Delta Q_t(p)\right] - \mathbb{E}\left[W_{i,t} + \frac{1}{2}\Delta W_{i,t} \Big| W_{i,t} \geq Q_t(p)\right].$$

This term, which is third-order in θ, corresponds to the average wealth of people that enter the top percentile (averaged before and after the perturbation) minus the average wealth of people that exit the top percentile. Relative to the second-order term, this term captures the asymmetry of the deviation in average wealth for entering and exiting agents: we can expect this term to be close to zero when the density of agents is close to constant around the threshold and the distribution of wealth deviation is symmetric.

Since this holds for any continuous function f, this must hold pointwise; hence,

$$\mathrm{d}^2 g_t(W) = -\partial_W \left(\mathbb{E}[\mathrm{d}^2 W_{i,t} | W_{i,t} = W] g_t(W)\right) + \partial_{WW} \left(\mathbb{E}[(\mathrm{d}W_{i,t})^2 | W_{i,t} = W] g_t(W)\right) \tag{18}$$

Now, differentiating the deviation of top wealth shares at the second-order in θ gives

$$\begin{aligned}
\mathrm{d}^2 \overline{W}_t(p) &= \frac{1}{p} \int_{Q_t(p)}^{\infty} (W - Q_t(p))\, \mathrm{d}^2 g_t(W) dW - \frac{1}{p} \mathrm{d}Q_t(p) \int_{Q_t(p)}^{\infty} \mathrm{d}g_t(W) dW \\
&= \frac{1}{p} \int_{Q_t(p)}^{\infty} (W - Q_t(p))\, \mathrm{d}^2 g_t(W) dW - \frac{g_t(Q_t(p))}{p} \left(\frac{1}{g_t(Q_t(p))} \int_{Q_t(p)}^{\infty} \mathrm{d}g_t(W) dW\right)^2.
\end{aligned}$$

Combining with (14) and (18) gives

$$\begin{aligned}
\mathrm{d}^2 \overline{W}_t(p) =& \frac{1}{p} \int_{Q_t(p)}^{\infty} (W - Q_t(p)) \Big\{ -\partial_W \left(\mathbb{E}[\mathrm{d}^2 W_{i,t} | W_{i,t} = W] g_t(W)\right) \\
&\qquad + \partial_{WW} \left(\mathbb{E}[(\mathrm{d}W_{i,t})^2 | W_{i,t} = W] g_t(W)\right) \Big\} dW \\
&- \frac{g_t(Q_t(p))}{p} \left(\frac{1}{g_t(Q_t(p))} \int_{Q_t(p)}^{\infty} \left(-\partial_W \left(\mathbb{E}[\mathrm{d}W_{i,t} | W_{i,t} = W] g_t(W)\right)\right) dW\right)^2
\end{aligned}$$

Integrating by parts gives

$$\begin{aligned}
\mathrm{d}^2 \overline{W}_t(p) =& \frac{1}{p} \int_{Q_t(p)}^{\infty} \mathbb{E}[\mathrm{d}^2 W_{i,t} | W_{i,t} = W] g_t(W) dW + \frac{g_t(Q_t(p))}{p} \mathbb{E}\left[(\mathrm{d}W_{i,t})^2 | W_{i,t} = W\right] \\
&- \frac{g_t(Q_t(p))}{p} \mathbb{E}\left[\mathrm{d}W_{i,t} | W_{i,t} = W\right]^2 \\
=& \mathbb{E}[\mathrm{d}^2 W_{i,t} | W_{i,t} \geq Q_t(p)] + \frac{g_t(Q_t(p))}{p} \mathbb{V}[\mathrm{d}W_i | W_{i,t} = Q_t(p)].
\end{aligned} \tag{19}$$

Now, a second-order Taylor approximation for $\Delta \overline{W}_t$ gives:

$$\Delta \overline{W}_t(p) = \mathrm{d}\overline{W}_t(p) + \frac{1}{2} \mathrm{d}^2 \overline{W}_t(p) + o(\theta^2)$$

Substituting out the expressions for $\mathrm{d}\overline{W}_t$ obtained in Proposition 1 and for $\mathrm{d}^2 \overline{W}_t$ obtained in (19) gives:

$$\Delta \overline{W}_t(p) = \mathbb{E}\left[\mathrm{d}W_{i,t} + \frac{1}{2} \mathrm{d}^2 W_{i,t} | W_{i,t} \geq Q_t(p)\right] + \frac{1}{2} \frac{g_t(Q_t(p))}{p} \mathbb{V}[\mathrm{d}W_i | W_{i,t} = Q_t(p)] + o(\theta^2),$$

which concludes the proof. □

Log deviation. We can divide Proposition 5 by $\overline{W}_t(p)$ to obtain the following second-order approximation for the *relative* deviation in the average wealth in a top percentile:

$$\frac{\Delta\overline{W}_t(p)}{\overline{W}_t(p)} = \underbrace{\mathbb{E}^{W_{i,t}}\left[\frac{\Delta W_{i,t}}{W_{i,t}}\middle|W_{i,t} \geq Q_t(p)\right]}_{\text{Intensive margin } O(\theta)} + \underbrace{\frac{1}{2}\frac{g_t(Q_t(p))Q_t(p)^2}{p\overline{W}_t(p)}\mathbb{V}\left[\frac{\Delta W_i}{W_{i,t}}\middle|W_{i,t} = Q_t(p)\right]}_{\text{Extensive margin } O(\theta^2)} + o(\theta^2). \quad (20)$$

Note that this expression is similar in spirit to the core result of Gomez (2023), who characterizes the instantaneous growth rate of the average wealth in the top percentile in terms of the dynamics of individual wealth (the key difference is that, here, I focus on changes w.r.t. θ rather than with respect to time).

In particular, note that, when the wealth distribution in the baseline economy has a Pareto tail, we have $g_t(Q_t(p))Q_t(p)/(p\overline{W}_t(p)) \to \zeta - 1$ where ζ is the tail exponent of the wealth distribution. Hence, holding other things equal, second-order effects are lower when the baseline wealth distribution is higher. Intuitively, a higher level of inequality reduces both churning around the percentile threshold as well as the effect of this churning on the average wealth beyond the top percentile. Quantitatively, $\zeta \approx 1.5$ for the U.S. wealth distribution, which implies that while the first-term (the intensive margin) is equal to the average relative deviation in wealth, the second-term (the extensive margin) is equal to a fourth of the variance of the relative deviation in wealth. Hence, we can expect this second-order effects to be relatively small.

One can rewrite (20) to obtain a second-order approximation for the deviation in the *log* average wealth in a top percentile:

$$\begin{aligned}
\Delta \log \overline{W}_t(p) &= \log\left(1 + \frac{\Delta\overline{W}_t(p)}{\overline{W}_t(p)}\right) \\
&= \log\left(1 + \mathbb{E}^{W_{i,t}}\left[\frac{\Delta W_{i,t}}{W_{i,t}}\middle|W_{i,t} \geq Q_t(p)\right] + \frac{1}{2}\frac{g_t(Q_t(p))Q_t(p)^2}{p\overline{W}_t(p)}\mathbb{V}\left[\frac{\Delta W_i}{W_{i,t}}\right] + o(\theta^2)\right) \\
&= \log\left(1 + \mathbb{E}^{W_{i,t}}\left[\frac{\Delta W_{i,t}}{W_{i,t}}\middle|W_{i,t} \geq Q_t(p)\right]\right) + \frac{1}{2}\frac{g_t(Q_t(p))Q_t(p)^2}{p\overline{W}_t(p)}\mathbb{V}\left[\frac{\Delta W_{i,t}}{W_{i,t}}\right] + o(\theta^2) \\
&= \underbrace{\log \mathbb{E}^{W_{i,t}}\left[e^{\Delta \log W_{i,t}}\middle|W_{i,t} \geq Q_t(p)\right]}_{\text{Intensive margin } O(\theta)} + \underbrace{\frac{1}{2}\frac{g_t(Q_t(p))Q_t(p)^2}{p\overline{W}_t(p)}\mathbb{V}\left[\Delta \log W_{i,t}\right]}_{\text{Extensive margin } O(\theta^2)} + o(\theta^2).
\end{aligned}$$

Note that the "intensive" margin $\log \mathbb{E}^{W_{i,t}}[e^{\Delta \log W_{i,t}}|W_{i,t} \geq Q_t(p)]$ is typically larger than the change in the average log wealth $\mathbb{E}^{W_{i,t}}[\Delta \log W_{i,t}|W_{i,t} \geq Q_t(p)]$ (the first-order term in Equation 5) due to Jensen inequality: since log is a concave function, the logarithm of the average deviation is always higher than the average of the log deviation.

A.3 Case where individual wealth is non-differentiable in the size of the perturbation

In the main text, we have focused on the common case in which the deviation in wealth due to the perturbation is differentiable in the perturbation size θ (Assumption 1). In reality, some perturbation may generate non-differentiable or event discontinuous changes in individual wealth. For instance, a change in corporate taxes may change the decision of being a worker and an entrepreneur, generating discontinuous changes in individual wealth.

Still, one can still characterize the first-order response of the average wealth in the top percentile as long as the set of individuals for which $\theta \to W_{i,t}(\theta)$ is not differentiable is, itself, infinitesimal. Under this assumption, the average wealth in the top percentile is still differentiable in θ and we have:

$$\mathrm{d}\overline{W}_t(p) = \mathrm{d}\mathbb{E}\left[(W_{i,t}(\theta) - Q_t(p))^+\right].$$

This expression is related to Firpo et al. (2009), who focuses on estimating the effect of changes in individual characteristics on inequality.

We can get something closer to Proposition 1 by distinguishing between the set of individuals with differentiable and non differentiable wealth. Formally, denote $\mathcal{ND}$ the set of individuals for which $\Delta W_{i,t}$ is not differentiable at 0. The first-order response in the average wealth in the top percentile is:

$$\mathrm{d}\overline{W}_t(p) = \mathbb{E}\left[\mathrm{d}W_{i,t} | i \notin \mathcal{ND}\right] + \lim_{\theta \to 0} \frac{\Delta\mathbb{E}\left[(W_{i,t}(\theta) - Q_t(p))^+\right]}{\Delta\theta} \mathbb{P}(i \in \mathcal{ND})$$

A.4 Extension to more general cumulative-distributive function

A.5 Comparison with differentiated Kolmogorov-Forward equation

I now briefly discuss my results relative to an alternative approach, which would study the effect of perturbations in the economy on the wealth distribution by differentiating the Kolmogorov-Forward equation (for instance, Auclert et al., 2021). Suppose that, in the baseline economy, the distribution of wealth evolves according to

$$G_{t+1} = \mathbb{T}_{t+1} G_t$$

where G_t denotes the CDF of the wealth distribution and $\mathbb{T}_{t+1}$ denotes a linear operator encoding the law of motion of individual wealth between t and $t+1$. Differentiating the previous equation with respect to a perturbation of wealth dynamics starting from $t = 0$.

$$\mathrm{d}G_{t+1} = \mathbb{T}_{t+1}\,\mathrm{d}G_t + (\mathrm{d}\mathbb{T}_{t+1})\,G_t$$

Solving this recurrence relationship backward gives:

$$\mathrm{d}G_t = \sum_{s=0}^{t} \left(\prod_{u=s+1}^{t} \mathbb{T}_u \right) (\mathrm{d}\mathbb{T}_s) G_{s-1}$$

The term $(\mathrm{d}\mathbb{T}_s)G_{s-1}$ corresponds to the change in the distribution of wealth at time s due to the perturbation in the law of motion of individual wealth $\mathrm{d}\mathbb{T}_s$. The effect of this deviation at time s on the distribution at time t is then mediated by the operator $\prod_{u=s+1}^{t} \mathbb{T}_u$, which captures how an initial distribution dissipates over time. While this equation can be useful to compute first-order deviations in the wealth distribution in terms of first-order deviations in the linear operators $\mathbb{T}$, is it less useful to characterize concretely top wealth shares (which is the object we observe in the data).

B Appendix for Section 2

Proof of Proposition 2. In this model, individual wealth is positive almost sure. Hence, taking logs and differentiating (6) with respect to an arbitrary sequence of perturbations for the growth rate of wealth $(\mathrm{d}\log A_{i,t})_{s \leq t}$ gives the following recurrence relation for $t \geq 0$:

$$\mathrm{d}\log W_{i,t+1} = \begin{cases} \mathrm{d}\log A_{i,t+1} + \mathrm{d}\log W_{i,t+1} & \text{if } a_{i,t+1} \geq 1 \\ 0 & \text{if } a_{i,t+1} = 0 \end{cases}$$

Iterating backward gives the perturbation of individual wealth at time t as the cumulated sum of all previous shocks in the growth rate of wealth:

$$\mathrm{d}\log W_{i,t} = \sum_{s=\max(0, t-a_{i,t}+1)}^{t} \mathrm{d}\log A_{i,s}.$$

Plugging this expression into (5) gives the result. □

Proof of Proposition 3. In this model, individual wealth is positive almost sure. Hence, taking logs and differentiating (8) with respect to an arbitrary sequence of perturbations for the growth rate of wealth $(\mathrm{d}\log A_{i,t})_{s \leq t}$ gives the following recurrence relation for $t \geq 0$:

$$\mathrm{d}\log W_{i,t+1} = \left(1 - \frac{B_{i,t+1}}{W_{i,t+1}}\right) (\mathrm{d}\log W_{i,t} + \mathrm{d}\log A_{i,t+1})$$

Iterating backward gives the perturbation of individual wealth at time t as the cumulated sum of all previous shocks in the growth rate of wealth:

$$\mathrm{d}\log W_{i,t} = \sum_{s=0}^{t} \left(\prod_{u=s}^{t} \left(1 - \frac{B_{i,u}}{W_{i,u}}\right)\right) \mathrm{d}\log A_{i,s}.$$

Plugging this expression into (5) gives the result.

□

B.1 Relation to existing results on tail indices

I now briefly compare my approach to the existing literature that focuses on characterizing tail indices.

Tail index. I now add a certain number of assumption to the model to ensure that the distribution in the baseline economy is stationary with a Pareto tail. Assume that individual wealth follows the law of motion (6), that the probability of death faced by each individual is constant over time and equal to δ, and that the distribution of $\log A_{i,t}$ is i.i.d. Then, it is well known that the distribution of wealth has a Pareto tail with tail index ζ given by the unique positive number satisfying:

$$\log \mathbb{E}\left[A_{i,t}^{\zeta}\right] = \delta.$$

Similarly, if individual wealth follows the law of motion (8), that the distribution of $A_{i,t}$ is i.i.d. with $\mathbb{E}[\log A_{i,t}] < 0$, and that the distribution of $Y_{i,t}$ is i.i.d with thin tails, then it is well known that the distribution of wealth has a Pareto tail with tail index ζ given by the unique positive number satisfying:

$$\log \mathbb{E}\left[A_{i,t}^{\zeta}\right] = 0.$$

Deviation in tail indices. In both cases, consider a *permanent* perturbation in the law of motion of individual wealth. Differentiating the expressions above tells us that, if the perturbed wealth distribution is also Pareto, then the resulting deviation in the tail index is:

$$\mathrm{d}\log\zeta = -\frac{\mathbb{E}^{A_{i,t}^{\zeta}}\left[\mathrm{d}\log A_{i,t}\right]}{\mathbb{E}^{A_{i,t}^{\zeta}}\left[\log A_{i,t}\right]}.$$

where $\mathbb{E}^{A_{i,t}^{\zeta}}$ denotes the distribution of growth rates "tilted" by $A_{i,t}^{\zeta}$.

Physical interpretation of tilted distribution. As discussed in Touchette (2009), the distribution of growth rates that is exponentially tilted by $A_{i,t}^{\zeta}$ can be interpreted as the distribution of prior growth rates conditional on being the top. Indeed, combining the previous equation with Proposition 2 implies

$$\mathbb{E}^{W_{i,t}}\left[\sum_{s=t-a_{i,t}+1}^{t} \mathrm{d}\log A_{i,s}|W_{i,t}\geq Q_t(p)\right] \sim \frac{\mathbb{E}^{A_{i,t}^{\zeta}}\left[\mathrm{d}\log A_{i,t}\right]}{\mathbb{E}^{A_{i,t}^{\zeta}}\left[\log A_{i,t}\right]}\log \overline{W}_t(p) \quad \text{as } p\to 0.$$

In the particular case of a uniform increase in growth rate, this implies

$$\mathbb{E}^{W_{i,t}}\left[a_{i,t}|W_{i,t}\geq Q_t(p)\right] \sim \frac{\log \overline{W}}{\mathbb{E}^{A_{i,t}^{\zeta}}\left[\log A_{i,t}\right]} \text{as } p\to 0.$$

This expression shows that the lifetime average "speed" of individuals in the right tail is asymptotically equal to the average growth rate of wealth in the distribution that is tilted by $A_{i,t}^{\zeta}$; that is, $\mathbb{E}^{A_{i,t}^{\zeta}}\left[\log A_{i,t}\right]$.

One can easily see that $\mathbb{E}^{A_{i,t}^{\zeta}}\left[\log A_{i,t}\right]$ can also be seen as the derivative of the CGF function of log growth rates at ζ; that is,

$$\mathbb{E}^{A_{i,t}^{\zeta}}\left[\log A_{i,t}\right] = \frac{\partial \log\left[A_{i,t}^{\zeta}\right]}{\partial \zeta}.$$

Now, it is well known that the CGF can be written as a Taylor expansion of cumulants; that is

$$\log \mathbb{E}[A_{i,t}^{\zeta}] = \zeta\mu + \frac{1}{2}\zeta^2\sigma^2 + \frac{1}{6}\zeta^3\cdot \text{skewness}\cdot\sigma^3 + \frac{1}{24}\zeta^4\cdot\text{kurtosis}\cdot\sigma^4 + \dots$$

where μ, σ, skewness, kurtosis to the mean, standard deviation, skewness, and kurtosis of $\log A_{i,t}$. Hence, the 'speed' of individuals reaching the right tail of the wealth distribution can be rewritten as the derivative of that expression with respect to ζ:

$$\mathbb{E}^{W_{i,t}}\left[a_{i,t}|W_{i,t}\geq Q_t(p)\right] \sim \frac{\log \overline{W}}{\mu + \zeta\sigma^2 + \frac{1}{2}\zeta^2\cdot\text{skewness}\cdot\sigma^3 + \frac{1}{6}\zeta^3\cdot\text{kurtosis}\cdot\sigma^4 + \dots} \quad \text{as } p\to 0.$$

Hence, for a given tail index, a higher mean, variance, skewness and kurtosis lead to a faster wealth path to top percentiles. Note that, in the case in which log growth rates are serially correlated over time, these cumulants should be understood as the derivatives of the scaled cumulating generating function.[21]

Serial correlation. All these results can be easily extended to models with serial correlations between individual growth rates (or models with type-specific distributions of growth rates), which is an important pattern in the data. In this type of model, Saporta (2005) and Beare and Toda (2022)

[21] That is, $\zeta \to \lim_{T\to\infty}\frac{1}{T}\log\mathbb{E}\left[\prod_{t=1}^{T} A_{i,t}^{\zeta}\right]$.

show that the tail index is characterized by

$$\lim_{T\to\infty}\frac{1}{T}\log\mathbb{E}\left[\prod_{t=1}^{T}A_{i,t}^{\zeta}\right]=\delta.$$

Differentiating this expression gives:

$$\mathrm{d}\log\zeta=-\lim_{T\to\infty}\frac{\mathbb{E}^{\prod_{t=1}^{T}A_{i,t}^{\zeta}}\left[\sum_{t=1}^{T}\mathrm{d}\log A_{i,t}\right]}{\mathbb{E}^{\prod_{t=1}^{T}A_{i,t}^{\zeta}}\left[\sum_{t=1}^{T}\log A_{i,t}\right]}.$$

This type of expression is obtained by Gomez and Gouin-Bonenfant (2024). Note that this equation implies that perturbating uniformly the growth rate of all agents in the economy generates the same change in the tail index as perturbing the growth rate of only one type of agents (the type reaching the right tail of the distribution). Yet, we discussed in Section 2 how an increase in the growth rate of agents that end up at the top of the wealth distribution typically generates much larger changes in top wealth shares relative to a uniform increase in the growth rate of all agents in the economy. This type of distinction cannot be picked up by tail indices, as they only characterize the asymptotic deviation in top wealth shares as $p\to 0$. This point illustrated that focusing on changes in tail indices can sometimes be very unhelpful to characterize the actual movements in top wealth shares.

C Appendix for Section 3

C.1 Proof of Proposition 4

Proof of Proposition 4. Consider a one-time perturbation in return, income and bequest at time $s<t$. Given (12), this creates the following perturbation in wealth:

$$\begin{aligned}\mathrm{d}\log W_{i,t+1}=&\left(1-\frac{Y_{i,t+1}+B_{i,t+1}-C_{i,t+1}}{W_{i,t+1}}\right)(\mathrm{d}\log R_{i,t+1}+\mathrm{d}\log W_{i,t})\\&+\frac{Y_{i,t+1}}{W_{i,t+1}}\mathrm{d}\log Y_{i,t+1}+\frac{B_{i,t+1}}{W_{i,t+1}}\mathrm{d}\log B_{i,t+1}-\frac{C_{i,t+1}}{W_{i,t+1}}\mathrm{d}\log C_{i,t+1}.\end{aligned}$$

Now, assuming that consumption is realized after return is realized but before income and bequest are and defining MPC as derivative of consumption w.r.t. post return wealth, we have:

$$-\frac{C_{i,t+1}}{W_{i,t+1}}\mathrm{d}\log C_{i,t+1}=-\mathrm{MPC}_{i,t+1}\times\frac{W_{i,t}}{W_{i,t+1}}\mathrm{d}\log W_{i,t}$$

Combining with the previous equation gives

$$\begin{aligned} \mathrm{d}\log W_{i,t+1} &= \left(1 - \frac{Y_{i,t+1} + B_{i,t+1} - C_{i,t+1}}{W_{i,t+1}}\right) \mathrm{d}\log R_{i,t+1} + \frac{Y_{i,t+1}}{W_{i,t+1}} \mathrm{d}\log Y_{i,t+1} + \frac{B_{i,t+1}}{W_{i,t+1}} \mathrm{d}\log B_{i,t+1} \\ &= \left(1 - \frac{Y_{i,t+1} + B_{i,t+1} - C_{i,t+1} - \mathrm{MPC}_{i,t+1} \times W_{i,t}}{W_{i,t+1}}\right) \mathrm{d}\log W_{i,t} \end{aligned}$$

In the continuous-time limit we can approximate $\mathrm{MPC}_{i,t+1} \frac{W_{i,t}}{W_{i,t+1}} \approx \mathrm{MPC}_{i,t+1}$, which gives the result.

□

C.2 Equalizing average asset returns across asset classes

As pointed out in Piketty (2014), one contributor to wealth inequality, beyond the average return on assets, is also the fact that richer households tend to own asset classes with higher average returns. One open question, however, is to assess magnitude of this effect. Put differently, how much lower would top wealth shares be if the average return was equalized across equity, housing, and debt? My framework provides a simple way to answer this question. Figure C1 plots the result of computing a counterfactual where I assume that the average return of every asset class (debt, equity, and housing) gets equalized starting from some time $t = 0$. I find that, over the long-run, the wealth share of the top 0.01% would decrease by 0.4 log points after sixty years (i.e. a decrease of approximately 30%). The drop is even more pronounced for households in Forbes 400, who almost hold 100% of their wealth in equity (see Gomez, 2017 for more evidence).

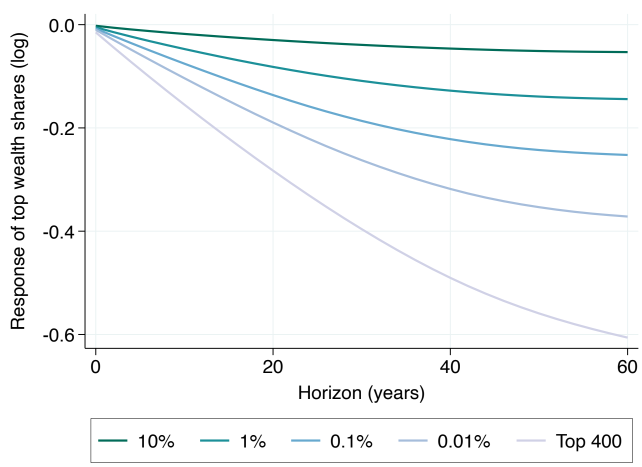

Figure C1: Effect of equalizing returns across asset classes on top wealth shares

D Comparison with synthetic accounting framework.

A commonly used approach in the literature is that, after specifying a model of wealth dynamics such as (4), is to assume that this applies for a "synthetic" individual in the top percentile; that is, that there is no composition changes in the top percentile over time.

Seen through the framework, this approach can be interpreted as assuming that there is perfect bequest (i.e. each individual in the top percentile inherits from parents in the same top percentile).

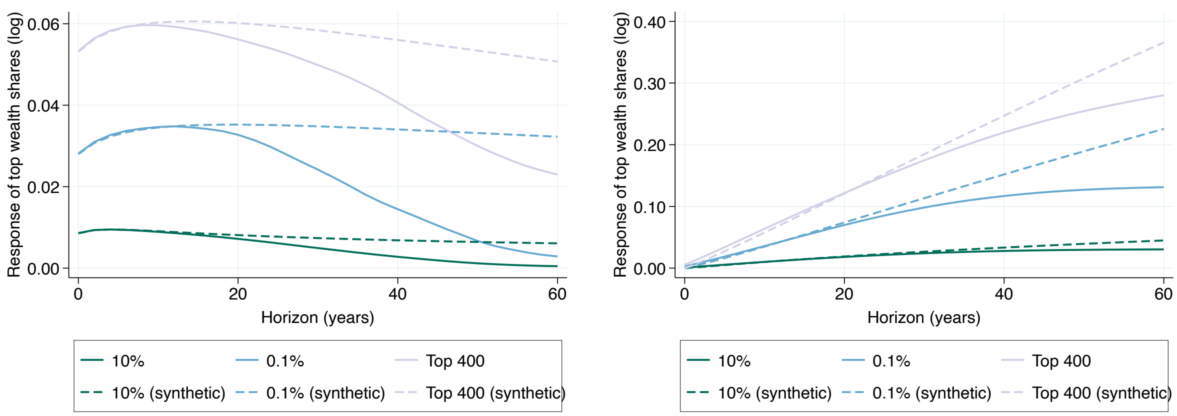

(a) 10pp. one-time increase in equity returns
(b) 1pp. permanent increase in equity returns

Figure D2: Comparison to naive accounting framework

Since this is not the case in the data, this type of framework this will tend to increase the persistence at the top, and generate large long-run elasticity of top wealth shares (and in particular, infinite ones if $Y/W \approx 0$. Int he short-run, however, there is no big difference between the two (as most of the people remain in the top over a small time horizon).