Aggregate Demand and the Top 1%

Adrien Auclert* Matthew Rognlie†

2 January 2016

Abstract

There has been a large rise in U.S. top income inequality since the 1980s. We merge a widely-studied model of the Pareto tail of labor incomes with a canonical model of consumption and savings to study the consequences of this increase for aggregate demand. Our model suggests that the rise of the top 1% may have led to a large increase in desired savings and can explain 0.5pp to 1pp of the fall in real interest rates. This effect cannot be attributed to differences in marginal propensities to consume, but arises from the combination of a wealth effect at the top and increased precautionary savings from declines lower in the income distribution.

*Stanford University and NBER. Landau Economics Building, 579 Serra Mall, Stanford, CA 94305. Phone: (650)723-2218. Email: aauclert@stanford.edu.

†Northwestern University and Princeton University. Julis Romo Rabinowitz Building, Princeton, NJ 08544. Phone: (919)358-1745. Email: mrognlie@princeton.edu.

Paper prepared for the 2017 AEA meetings, to be presented in the session: Heterogeneity and Power Laws in Macroeconomics. (Chair: Xavier Gabaix, Discussant: Erzo Luttmer.)

The rise of the top 1% of labor income earners since the 1980s has attracted a considerable amount of attention in the literature. The facts are by now well-established: the top of the labor income distribution is well described by a power law, whose tail appears to have fattened over time (see for example Piketty (2014)). However, the macroeconomic consequences of this rise in top income inequality remain unclear. In particular, since rich and poor have different consumption and savings patterns, higher inequality may have affected aggregate demand. But by how much?

A large literature on top income inequality relies on random growth processes to explain the tail of the income distribution. A recent literature further argues that changes in the fundamentals of this process can explain the observed rise in the top 1% since the 1980s (see for example Gabaix, Lasry, Lions and Moll (2016)). Building on the contributions of Huggett (1993) and Aiyagari (1994), in Auclert and Rognlie (2016) (henceforth AR) we develop a framework for mapping changes in income processes, first to changes aggregate demand (a partial equilibrium effect), and then to interest rates and output (a general equilibrium effect) depending on assumptions on monetary and fiscal policy. In this project, we combine these canonical frameworks to study the consequences of the rise in the top 1% of incomes. To be precise, we ask: if all that happened in the US between 1980 and today had been a fattening of the Pareto tail of the labor income distribution, holding its average constant (pure redistribution), what would the consequences for savings and interest rates have been?

1 Fitting the evolution of the top income distribution

The solid black line of figure 1 shows the evolution of the top 1% of US labor incomes (from data on wages and salaries) since the 1980s according to the World Top Incomes database. The share earned by the top 1% has roughly doubled, from 6.4% in 1980 to 11.1% today.[1] It is well-known that the upper tail of the income distribution follows a Power law, that is, that the fraction of income above any given, large enough level y is given by

$$\mathbb{P}\left(y_i \geq y\right) \propto y^{-\alpha} \tag{1}$$

If the entire distribution was Pareto, we could infer the tail coefficient α using the

[1]The most widely cited figures for the income share of the top 1% are higher than this, because they also include income on capital. (See for example Gabaix et al. (2016), from whom we take our data). Instead, we take the wage distribution as exogenous, and the model endogenously generates consequences for capital income.

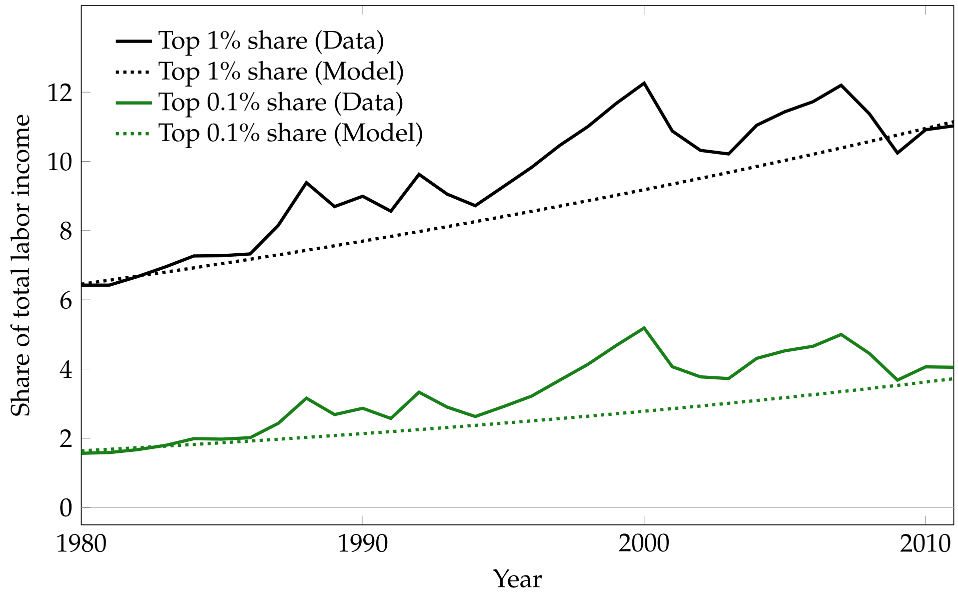

Figure 1: Top 1% and 0.1% labor income shares, US data and our model calibration

formula

$$\alpha = \frac{1}{1 - \frac{\log(\text{top 1\% share})}{\log(1\%)}} \tag{2}$$

This simple exercise delivers $\alpha_{1980} = 2.47$ and $\alpha_{\text{today}} = 1.91$. It turns out that infering α in this way also predicts the shape of the income distribution within the top 1% extremely well (for example the 0.1%, as shown by the green line of figure 1). We therefore maintain these as our baseline estimates of the Pareto tail of US wages.

It is widely known (see, for example, Harrison (1985)) that if the process for individual incomes y_{it} follows a geometric random walk with negative drift

$$d \log y_{it} = -\mu dt + \sigma dZ_{it} \tag{3}$$

where Z_{it} is a standard Brownian motion, and $\underline{y}$ is a lower reflecting barrier, then the stationary distribution y_{it} obeys (1) with

$$\alpha = \frac{2\mu}{\sigma^2} \tag{4}$$

To focus on the effects of a rise in tail inequality, we use this simple model of the income process to conduct quantitative experiments, studying the effect of a decline from $\alpha_{1980} = 2.47$ to $\alpha_{\text{today}} = 1.91$ and the accompanying surge in top incomes.

The large literature on earnings dynamics delivers useful orders of magnitude for

σ^2, the variance of the innovations to the permanent component of log US earnings. Estimates from Floden and Lindé (2001) to Heathcote, Perri and Violante (2010) suggests $\sigma^2 \in [0.01, 0.04]$ per year as reasonable values. The evidence is consistent with σ being either flat or rising modestly over time. We therefore set an initial value of $\sigma^2_{1980} = 0.02$, implying a modest negative drift in earnings of around 2.5% per year.

Through the lens of equation (4), we can interpret the fall in α from 1980 to today as some combination of changes in downward reversion μ and idiosyncratic volatility σ. Our experiments consider three combinations of μ and σ: we let

$$\sigma^2_{\text{today}} = \sigma^2_{1980} \left(\frac{\alpha_{1980}}{\alpha_{\text{today}}} \right)^k \quad \text{for } k = 0, 1, 2$$

and then let μ adjust to satisfy (4), adjusting the lower reflecting barrier $\underline{y}$ so as to hold average income constant in all three cases. When $k = 0$, the rise in income inequality comes from weaker downward drift: high incomes stay high for longer. When $k = 1$, the rise comes instead from higher volatility: there are more shocks to income, even though drift is constant. When $k = 2$, an even larger rise in volatility offsets a rise in drift: the right side of (3) is uniformly scaled up, implying larger movements in income across the board.

Starting from the initial stationary distribution in 1980, we progressively phase in the change in the earnings process in (3) so that the transition is complete today.[2] Figure 1 plots the resulting path for the share of top incomes in our model.

2 Consequences for aggregate savings

In AR, we propose a model that allows us to map the consequences of changes in earnings processes such as the one discussed in the previous section onto changes in macroeconomic aggregates, building on the general equilibrium models of Huggett and Aiyagari. Our key observation is that the change in income process, holding macroeconomic aggregates constant, implies a change in *aggregate demand,* that is, in the pattern of desired consumption and savings over time. How macroeconomic imbalance between consumption and output is resolved in general equilibrium in turn

[2] Our discretization ensures that the process remains at its stationary distribution of states at each point in the transition, even as we progressively space out the grid points. We therefore do not have to face the slow underlying dynamics of the income process pointed out in Gabaix et al. (2016), and we are free to target any path for the top income share. See appendix A for details.

k	α	μ	σ^2	W^{PE}	r^{GE*}
Initial SS	2.47	0.024	0.02	2.98	4%
0	1.91	0.019	0.02	4.60	3.55%
1	1.91	0.024	0.026	5.32	3.35%
2	1.91	0.030	0.033	5.95	3.17%

* Assuming monetary policy that maintains full employment and fiscal policy that holds government spending and debt fixed.

Table 1: Main experiments

depends on fundamentals of the economy as well as monetary and fiscal policy rules, which multiply the initial effect on aggregate demand.

We calibrate our model to 1980. We assume that the initial steady state real interest rate is $r = 4\%$ and the income process is (a discretized version of) the process of section 1, held at its stationary 1980 distribution. Our agents are ex-ante identical consumers with common discount rate $\beta = 0.95$ and elasticity of intertemporal substitution $\sigma = \frac{1}{2}$. Their skills evolve stochastically according to the discretized process. They are paid a market wage per unit skill, implying that they are subject to idiosyncratic risk, and can only trade in non-contingent bonds and firm shares. The government uses an affine tax system to redistribute income at an earmarked rate of $\tau_r = 17.5\%$ and to pay for government spending ($\frac{G}{Y} = 20.6\%$) as well as interest on the debt ($\frac{B}{Y} = 27.1\%$). Firms produce using a Cobb-Douglas production function $Y = AK^{\alpha}L^{1-\alpha}$. We assume full employment ($L = 1$); the available capital stock is $\frac{K}{Y} = 271\%$. We calibrate this to include land, and hence calibrate the depreciation rate to $\delta = 2.9\%$. As a result the labor share is a counterfactually high $\alpha = 1 - (r + \delta)\frac{K}{Y} = 81.3\%$. (In AR we introduce an equity premium to resolve this discrepancy with the data.)[3]

We then assume that in 1980, agents learn that top incomes will rise, with the income process changing in the manner described in the previous section. We compute the new steady-states, as well as perfect-foresight transition paths, implied by each of our experiments. Table 1 shows that, if no other macroeconomic aggregate was affected, the increase in top income inequality would have resulted in a large rise in aggregate savings W^{PE}. Moreover, the magnitude of this rise depends on the exact underlying driver of the Pareto tail increase. The more idiosyncratic risk increases (the

[3]Our model, aided by our superstar income process and our high labor share assumption, does a good job at matching the wealth distribution until close to the top. (The 1% wealth share in 1980 is 23% in our calibration.). However, it quantitatively misses the rise in the top wealth share from 1980 to today: depending on simulations, the top 1% wealth share today is between 26% and 33% in our model, compared to the Saez-Zucman estimate of 40%.

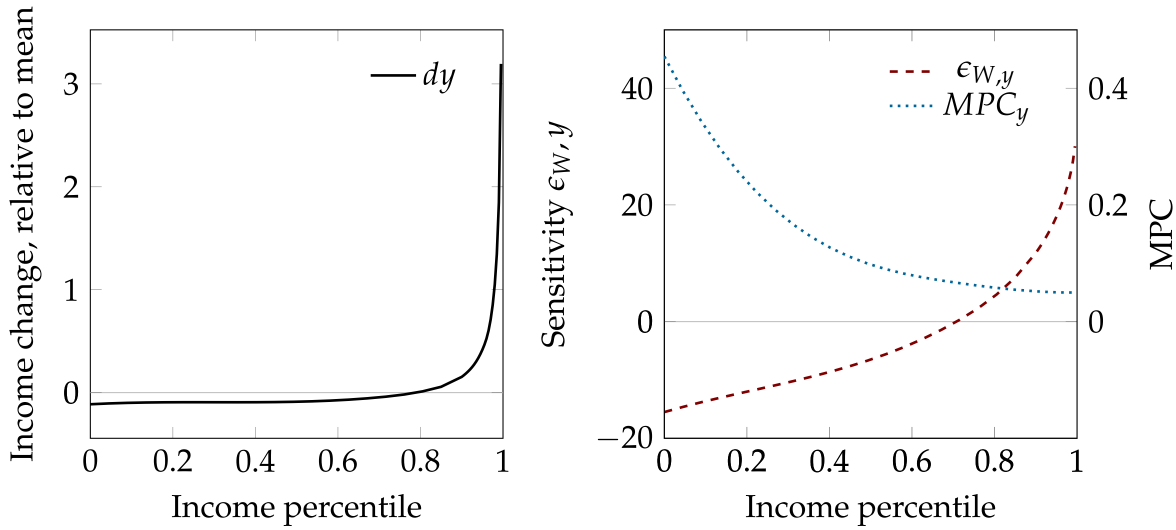

Figure 2: Redistribution from $\alpha : 2.47 \to 1.91$, and partial equilibrium sensitivities

higher k), the higher the increase in wealth; but note that there is a large increase in wealth increase even if σ stays constant. This suggests that precautionary saving, although important, is not the only effect at work; we now dig into this question further.

3 The role of redistribution

Conditional on aggregate income, a thicker Pareto tail means a rise in incomes at the top, at the expense of incomes at the bottom. The left panel of figure 2 shows the impact of our experiment, a decline in α from 2.47 to 1.91, for incomes at each percentile of the distribution, normalized by mean aggregate income. As expected, most of the rise in income occurs at the very top of the distribution.

The rise in savings: precautionary savings and wealth effect. The experiment $k = 2$ in table 1 has the special feature that the transition matrix between income quantiles is left unaltered, even as those quantiles become more dispersed. We can think of the experiment as a change in the income implied by each state in a Markov process, with the process itself being unchanged.

It is then possible to break down the increase in aggregate wealth $\frac{dW}{W}$ in the experiment into contributions from the change in each income level y, using the first-order approximation

$$\frac{dW}{W} = \text{Cov}\left(\epsilon_W, dy\right) \tag{5}$$

where $\epsilon_{W,y}$ is the increase in aggregate savings that would result if only the income of individuals earning y were raised, normalized by the fraction of such individuals.

The right panel of figure 2 shows that $\epsilon_{W,y}$ is monotonically increasing in y. For low y, $\epsilon_{W,y}$ is negative: therefore, when income at the bottom of the distribution goes down, individuals in the aggregate increase savings for precautionary reasons. For higher y, $\epsilon_{W,y}$ becomes positive: when income at the top of the distribution goes up, individuals also increase savings due to a wealth effect. Hence, in our experiment, both effects contribute positively to aggregate savings,[4] with the wealth effect being somewhat more important.

The role of MPC differences. It is often argued that the rise of the top 1% may depress aggregate demand because of they have lower marginal propensities to consume (MPCs). In AR, we show that this argument is correct if the change in income inequality is temporary.[5] We derive a sufficient statistic for the partial equilibrium effect ∂C on consumption in a given year, when an income inequality change takes place only in that year, which is:

$$\partial C = \text{Cov}\left(MPC, dy\right) \tag{6}$$

We then explain how (6) maps into the general equilibrium effect.

The right panel of figure 2 shows model-implied average MPCs by income percentile. MPCs decline with income, but the decline is much stronger near the bottom of the distribution (where MPCs are reasonably high) than at the top (where MPCs are consistently low). This brings down the covariance in (6), since the income changes in our experiment are most dispersed at the top of the distribution, where the variation in MPCs is limited.

Evaluating (6), the implied partial equilibrium consumption effect of a year-on-year increase in income inequality that thickens the Pareto tail from $\alpha = 2.47$ to 1.91, as in our experiment, is -1.8% of GDP. Furthermore, as demonstrated in AR, this partial equilibrium consumption effect is typically close in magnitude to the general equilibrium output effect.

Note, however, that such a enormous shift in the income distribution is unlikely to

[4]Observe that both the dy and the $\epsilon_{W,y}$ lines in figure 2 cross zero around the same income percentile.

[5]However, we show that MPCs play no role in explaining long-run aggregate demand when the change in income inequality is long-lasting.

be transitory, or to happen within a single year. More realistic transitory shocks to the right tail of the income distribution, like the year-to-year fluctuations in figure 1, will have commensurately smaller impacts.

Clear empirical evidence on the pattern of MPCs by income is lacking, especially at the very top. Available empirical evidence lends support to the view that MPCs decline with income, although probably not as strongly as what is implied by figure 2. Hence, our -1.8% number is very likely to be an upper bound of the consumption effect of a rise in the top 1%.

4 Macroeconomic effects

The framework in AR allows us to map the effect on savings discussed in section 3 on to an effect on output and interest rates depending on assumptions about monetary and fiscal policy. If fiscal policy holds the level of government debt and spending constant, and monetary policy lowers interest rate to ensure full employment, then the effect on equilibrium interest rates is given in table 1. The interest rate effect is substantial, amounting to somewhere between 45 and 85 basis points depending on the source of the change in the Pareto tail. As shown in appendix B, the model suggests this decline is very progressive and only fully realized after around 100 years. Hence, the full macroeconomic effects of the rapid rise in income inequality we have observed may take decades to filter through, providing one force for depressed equilibrium interest rates going forward.

While the rise of the top 1% is unlikely to have affected aggregate demand because of MPC differences, it may well have affected it because of the resultant increase in desired savings, via a combination of a wealth effect and higher precautionary savings. In turn, this may have contributed to pushing the economy to the zero lower bound. In AR, we show that once the zero lower bound starts binding, further increases in inequality can be extremely deleterious, and can potentially lead to secular stagnation.

References

Aiyagari, S. Rao, "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, August 1994, *109* (3), 659–684.

Auclert, Adrien and Matthew Rognlie, "Inequality and Aggregate Demand," December 2016.

Champernowne, D. G., "A Model of Income Distribution," *The Economic Journal*, June 1953, *63* (250), 318–351.

Floden, Martin and Jesper Lindé, "Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?," *Review of Economic Dynamics*, April 2001, *4* (2), 406–437.

Gabaix, Xavier, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll, "The Dynamics of Inequality," *Econometrica*, November 2016, *84* (6), 2071–2111.

Harrison, J. Michael, *Brownian motion and stochastic flow systems* Wiley series in probability and mathematical statistics, New York: Wiley, 1985.

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante, "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967–2006," *Review of Economic Dynamics*, January 2010, *13* (1), 15–51.

Huggett, Mark, "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, September 1993, *17* (5–6), 953–969.

Piketty, Thomas, *Capital in the Twenty-First Century*, first edition edition ed., Cambridge Massachusetts: Belknap Press, March 2014.

A Discretization of the income process

We assume that log income follows a random walk with negative drift

$$d \log y_{it} = -\mu dt + \sigma dZ_{it} \tag{7}$$

with a reflecting barrier at the lower bound $\underline{y}$. This is known to produce a Pareto stationary distribution with shape parameter $2\mu/\sigma^2$. For tractability, the general equilibrium Huggett-Aiyagari model that we use, adapted from Auclert and Rognlie (2016), has a discrete state space for exogenous incomes. Hence, it is necessary to choose some discretization for (7).

To do so, we adapt a simple process from Champernowne (1953), which produces a discretized Pareto distribution.[6] Assume that log income $x_{it} = \log y_{it}$ can take the values $\{\underline{x} + aj\}$ for $j \geq 0$. Also assume that $x_{it} > \underline{x}$ follows a continuous-time Markov process with a transition rate of u to $x_{it} + a$ and d to $x_{it} - a$, for some $0 < u < d$, where all other transition rates being zero. For $x_{it} = \underline{x}$, assume that the only permissible transition is to $x_{it} + a$, with rate u.

The stationary distribution is Pareto. The stationary distribution of this process is a discretized Pareto distribution. Indeed, if π_j denotes the mass of individuals in state $j \in [0, J]$ in the stationary distribution, stationarity requires that the entering and exiting flows are equalized in each state, i.e.

$$\begin{aligned} u\pi_{j-1} + d\pi_{j+1} &= (u+d)\,\pi_j \qquad j \geq 1 \\ d\pi_1 &= u\pi_0 \end{aligned}$$

whose solution, enforcing $\sum_{j\geq 0} \pi_j = 1$, is

$$\pi_j = \left(1 - \left(\frac{u}{d}\right)\right)\left(\frac{u}{d}\right)^j$$

and therefore in particular, for any $x_j = \underline{x} + aj$

$$\Pr\left(x_{it} \geq x_j\right) = \left(\frac{u}{d}\right)^j = e^{\left(x_j - \underline{x}\right)\frac{\log\left(\frac{u}{d}\right)}{a}}$$

[6] For better tractability and a closer approximation to the continuous-time random walk, we start by formulating the process in continuous time, unlike discrete time as in Champernowne (1953), and then derive the implied discrete time transition matrix.

and hence for any $y_j = e^{x_j}$

$$\Pr\left(y_{it} \geq y_j\right) = \left(\frac{y}{y_j}\right)^{\frac{-\log\left(\frac{u}{d}\right)}{a}}$$

We recognize this as the CDF of a discretized Pareto distribution with scale parameter (minimum value) $\underline{y} = e^{\underline{x}}$ and shape parameter

$$\alpha = \frac{-\log\left(\frac{u}{d}\right)}{a} \tag{8}$$

Drift and volatility. Given u and d, both the drift and squared volatility of this process are constant, with

$$\begin{aligned} \mu &= a(d-u) \\ \sigma^2 &= a^2(d+u) \end{aligned}$$

Inverting this relationship, for given μ and σ^2, we have

$$u = \frac{1}{2}\left(\frac{\sigma^2}{a^2} - \frac{\mu}{a}\right) \tag{9}$$

$$d = \frac{1}{2}\left(\frac{\sigma^2}{a^2} + \frac{\mu}{a}\right) \tag{10}$$

Plugging into (8) and simplifying gives

$$\alpha = \frac{1}{a}\log\left(\frac{1-\frac{a\mu}{\sigma^2}}{1+\frac{a\mu}{\sigma^2}}\right) \tag{11}$$

Note that in the limit $a \to 0$, $\log\left(1 \pm \frac{a\mu}{\sigma^2}\right) \sim \frac{a\mu}{\sigma^2}$, so that (11) reduces to $\alpha = \frac{2\mu}{\sigma^2}$. This is exactly the formula for α in (4). Hence, as the discretization becomes finer, the relationship between μ, σ, and α approaches that of our idealized income process, the geometric random walk with negative drift and a lower reflecting barrier.

Calibrating the process. Given that $\underline{x}$ has to be chosen to achieve our normalization $\mathbb{E}\left[y_{it}\right] = 1$, our process has three free parameters (a, u, d).

The Lorenz curve $L(u)$ of a Pareto with shape α is $1 - L(u) = (1-u)^{1-\frac{1}{\alpha}}$. As mentioned in the text, given a value for the top 1% share, we can then back out the implied Pareto α using

$$\alpha = \frac{1}{1 - \frac{\log(\text{top 1\% share})}{\log(1\%)}}$$

Calibrating α on the basis of the top 1% share in 1980, (8) then provides one restriction on (a, u, d). Our calibration of $\sigma^2_{1980} = 0.02$ provides another. Given a, these jointly pin down u and d, which in turn imply $\mu = a(d - u)$, by

$$\begin{aligned} \frac{u}{d} &= e^{-a\alpha} \\ u + d &= \frac{\sigma^2}{a^2} \end{aligned}$$

The remaining choice is a. This is made primarily on computational grounds. Since we require finitely many states for computation, it is necessary to truncate the set $\{\underline{x} + aj\}$ at some maximum $\underline{x} + aJ$. To avoid truncation bias, we pick aJ high enough such that only 0.001% of aggregate income is earned at or above this state in the ideal Pareto distribution for our initial calibration, writing $e^{-(\alpha-1)aJ} = 10^{-5}$. It follows that $aJ = 7.83$, so that the maximum income state is approximately 2500 times higher than the minimum income state. Since the algorithm in AR is $O(J^2)$, we set $J = 40$ for reasonable computation time, implying $a \approx 0.2$.

To map the process to discrete time, we take the matrix exponential to convert the transition *rate* matrix Σ to a Markov transition matrix Π: $\Pi = e^{\Sigma}$.

Changing the distribution. We consider a decline in α from 2.47 to 1.91 to match the rise in the income share of the top 1% since 1980. Using (11), there are various combinations of changes in (a, μ, σ) that can replicate this decline. To ensure that the truncation stays accurate, we vary $a \propto \alpha^{-1}$, such that the maximum state $\underline{x} + aJ$ remains at the same percentile of the Pareto distribution. (It follows that all other states remain at the same percentiles as well.)

Given $a \propto \alpha^{-1}$, it is clear from (11) that $a\mu/\sigma^2$ must be unchanged. The three experiments in table 1, also discussed in the main text, represent different choices of μ and σ^2 that accomplish this: either (1) $\mu \propto \alpha$ and σ constant ($k = 0$), (2) μ constant and $\sigma^2 \propto \alpha^{-1}$ ($k = 1$), or (3) $\mu \propto \alpha^{-1}$ and $\sigma^2 \propto \alpha^{-2}$ ($k = 2$). The third choice yields unchanged u and d in (9) and (10), and this produces an unchanged transition matrix that is particularly useful for the decomposition (5).

B Simulated paths for W^{PE} and r^{GE}

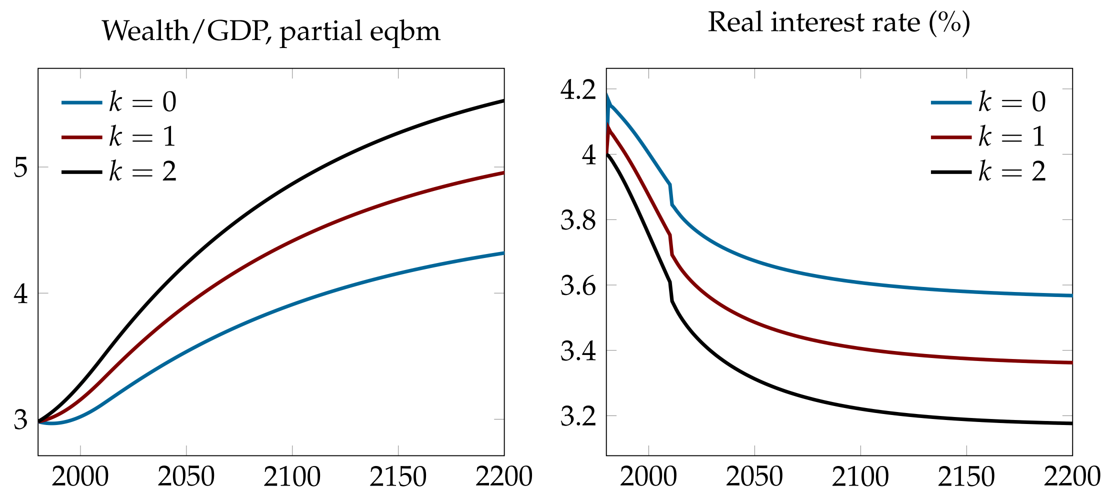

Figure 3: Simulated paths for W_t^{PE} and r_t^{GE}

Figure 3 plots the perfect-foresight transition dynamics of our model, from its 1980 to its long-run new steady state after the rise in the Pareto tail of the income distribution. As discussed in the main text, we phase in the rise in inequality over the first 30 years. The rapid decline in interest rates around year 2010 on the right panel reflects the end of this phase in, with income inequality staying constant thereafter, but interest rates continuing to decline.