Aggregate Demand and the Top 1%

Adrien Auclert Matthew Rognlie*

13 January 2017

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 a 0.45pp to 0.85pp decline in long-run real interest rates. This effect arises from the combination of a wealth effect at the top and increased precautionary savings from declines lower in the income distribution.

*Auclert: Rognlie: Stanford University: Landau Economics Building, 579 Serra Mall, Stanford, CA 94305. Phone: (650)723-2218. Email: aauclert@stanford.edu. Rognlie: Northwestern University and Princeton University. Julis Romo Rabinowitz Building, Princeton, NJ 08544. Phone: (609)258-7447. Email: mrognlie@princeton.edu. We thank Ben Moll and our discussant Erzo Luttmer for helpful comments.

Paper prepared for the 2017 AEA meetings and 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 public debate, as well as in academia. 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, Thomas 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 through which channels, and 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 Xavier Gabaix, Jean-Michel Lasry, Pierre-Louis Lions and Benjamin Moll 2016). Building on the contributions of Mark Huggett (1993) and S. Rao Aiyagari (1994), in Adrien Auclert and Matthew Rognlie (2016) (henceforth AR) we develop a framework for mapping changes in income processes, first to changes in 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: suppose that all that happened between 1980 and today had been the observed fattening of the Pareto tail of the labor income distribution, with average income held constant (i.e., income had been redistributed in a way consistent with Pareto's law). What would the consequences for savings and interest rates have been? Our results are consistent with the view that this change led to a large increase in desired savings, which in turn may have reduced long-run equilibrium interest rates by 45 to 85 basis points.

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

[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.

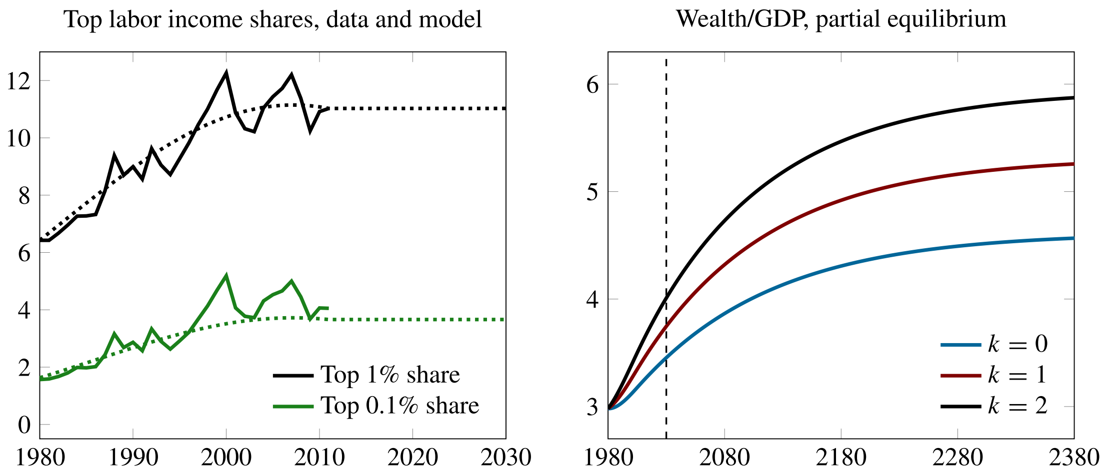

Figure 1: Top labor income shares and model-implied partial equilibrium wealth/GDP paths

power law; that is, that the fraction of individuals making more than any given, large enough level of income 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 formula

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

This simple exercise delivers $\alpha_{1980} = 2.48$ and $\alpha_{\text{today}} = 1.92$. It turns out that inferring α in this way also predicts the shape of the income distribution *within* the top 1% extremely well—for example, the top 0.1% share, 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 also widely known (see, for example, J. Michael 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 $\alpha_{1980} = 2.48$ to $\alpha_{\text{today}} = 1.92$ 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 Martin Floden and Jesper Lindé (2001) to Jonathan Heathcote, Fabrizio Perri and Giovanni L. Violante (2010) suggest $\sigma^2 \in [0.01, 0.04]$ 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. Our preferred experiment, however, is $k = 2$, which scales up the right-hand side of (3) uniformly. This has the natural interpretation of an increase in the return to skill: if individual skills e_{it} follow the process in (3), and income is $y_{it} = e_{it}^{\xi_t}$, then a rise in the 'relative skill price' ξ_t by the factor $\alpha_{1980}/\alpha_{\text{today}}$ delivers our $k = 2$ outcome.

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] The left panel of figure 1 plots the resulting paths for the share of top incomes in our model as dashed lines.

[2] As pointed out by Gabaix et al. (2016), with the skill price interpretation it is possible to target any transition path for the top Pareto coefficient. See appendix A for details on our income transition dynamics.

k	α	μ	σ^2	W^{PE}	r^{GE*}
Initial SS	2.48	0.024	0.02	2.98	4%
0	1.92	0.019	0.02	4.61	3.55%
1	1.92	0.024	0.026	5.31	3.35%
2	1.92	0.031	0.033	5.94	3.16%

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

Table 1: Main experiments

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 just discussed, 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 employment and interest rates 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, depends on fundamentals of the economy as well as monetary and fiscal policy rules, which multiply the initial effect on aggregate demand. We now briefly describe the key elements of the model.

The economy is populated by infinitely-lived, ex-ante identical agents with common discount rate β and elasticity of intertemporal substitution $\nu = \frac{1}{2}$. Their skills evolve stochastically according to the process in (3), initially held at its stationary 1980 distribution. Pre-tax earnings are proportional to skills. We assume that a rate $\tau_r = 17.5\%$ is applied to these earnings, with proceeds redistributed in a lump-sum fashion. Agents face no aggregate uncertainty and can only trade in an asset delivering a constant return r, subject to maintaining positive net worth at all times. We calibrate the model to 1980, setting $r = 4\%$ and finding β such that the ratio of aggregate wealth to overall post-tax labor income is consistent with its 1980 value. This completes the description of the partial equilibrium model.

To compute general equilibrium outcomes, we further assume that the overall wealth-income ratio is the sum of government debt $\frac{B}{Y} = 27.1\%$ and capital $\frac{K}{Y} = 271\%$. The government taxes labor income to finance spending $\frac{G}{Y} = 20.6\%$ and interest on the debt. Firms produce using a Cobb-Douglas production function $Y = AK^{\alpha}L^{1-\alpha}$, with capital depreciating at an annual rate of $\delta = 2.9\%$. We assume that monetary policy maintains full employment at all times, and that fiscal policy responds to developments in inequality by maintaining the levels of spending G and debt B constant. [3]

We 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 employment and interest rates had remained constant, the increase in top income inequality would have resulted in a large rise in aggregate savings W^{PE}. The left panel of Figure 1 further shows that this increase in aggregate wealth would have taken centuries to materialize, reflecting the slow process through which income inequality accumulates to determine wealth inequality. Moreover, the magnitude of this rise depends on the exact underlying driver of the Pareto tail increase. The more idiosyncratic risk increases (the 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 2.48 to 1.92, 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. Following our preferred interpretation, we can think of the experiment as a change in the income implied by each state in a Markov process for skills, 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

[3] Our model, aided by our fat-tailed income process, does a good job at matching the wealth distribution until close to the top, but it misses the shape of the tail as well as the dynamics of wealth inequality documented in the data. See Appendix C for details.

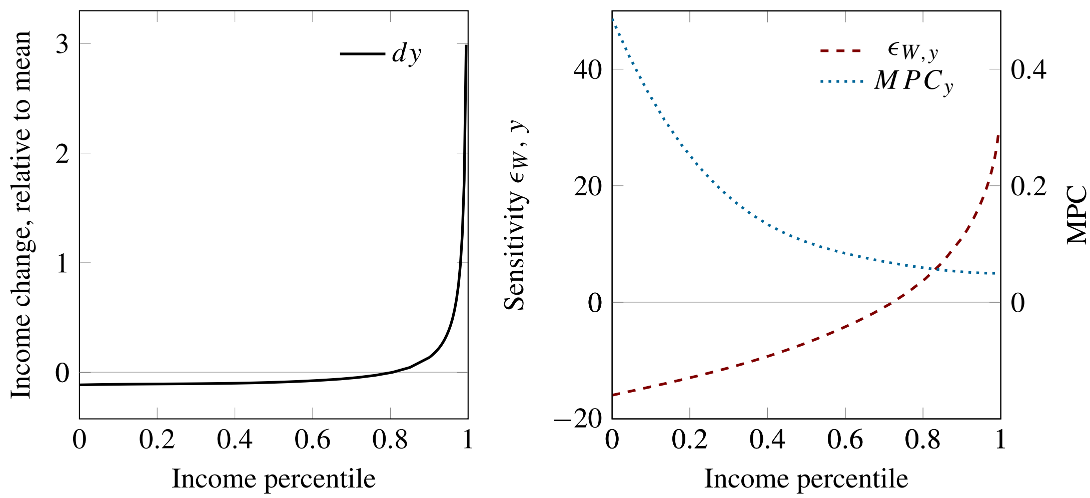

Figure 2: Redistribution from $\alpha : 2.48 \rightarrow 1.92$, and partial equilibrium sensitivities

approximation (see appendix D)

$$\frac{dW}{W} = \mathrm{Cov}\left(\epsilon_{W,y}, dy\right) \tag{5}$$

where $\epsilon_{W,y}$ is the proportional 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 the rich 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

[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.

that year, which is:

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

We then show that (6) is a key determinant of 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 $\alpha = 2.48$ to 1.92, 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 large shift in the income distribution is unlikely to 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 onto 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 rates to ensure full employment, then the effect on equilibrium interest rates is given in table 1. The long-term 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, this decline is protracted 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 damaging, and can potentially lead to secular stagnation.

References

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

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

Benhabib, Jess, Alberto Bisin, and Shenghao Zhu. 2015. "The Wealth Distribution in Bewley Economies with Capital Income Risk." *Journal of Economic Theory*, 159, Part A: 489–515.

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

De Nardi, Mariacristina, and Giulio Fella. 2017. "Saving and Wealth Inequality." *CEPR Discussion Paper No 11746*.

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

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

Harrison, J. Michael. 1985. *Brownian Motion and Stochastic Flow Systems.* New York:Wiley.

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

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

Piketty, Thomas. 2014. *Capital in the Twenty-First Century.* . First ed., Cambridge, MA:Belknap Press.

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

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 D. G. 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 $L\left(u\right)$ of a Pareto with shape α is $1 - L\left(u\right) = \left(1-u\right)^{1-\frac{1}{\alpha}}$. 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 2.48 to 1.92 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).

When computing transition dynamics, it is also necessary to specify how the pa-

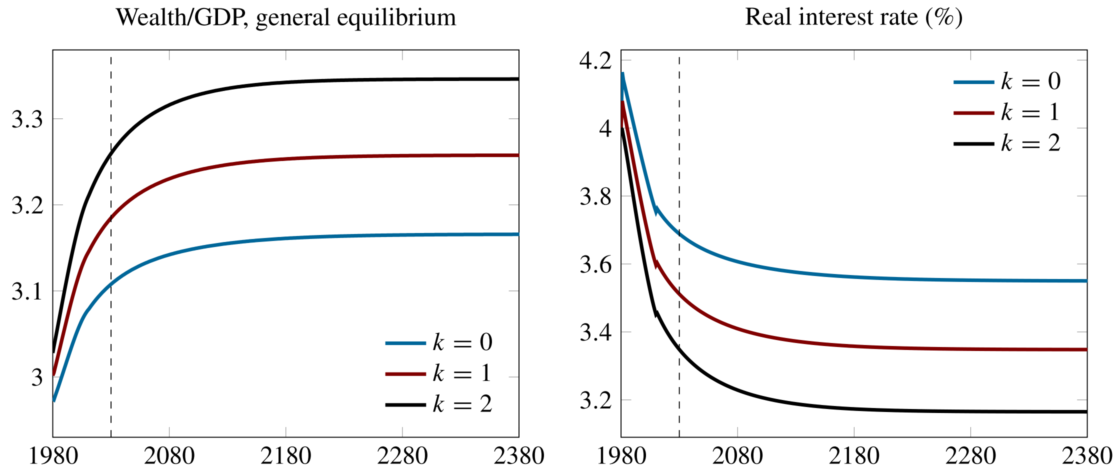

Figure 3: Simulated paths for W_t^{PE} and r_t^{GE}

rameters (a, μ, σ) adjust over time to the new steady state. We continue to require that $a\mu/\sigma^2$ is unchanged at all times, so that $a \propto \alpha^{-1}$ in every year. We choose a quadratic trend of a from 1980 to 2011, such that a is consistent with the actual top 1% income share in 1980 and 2011, and minimizes the average square difference between the model and data for the top 1% share across intermediate years. We then assume that σ^2 follows the same convergence path from 1980 and 2011, and infer the implied μ from $\mu \propto \sigma^2/a$. From 2011 onward, we assume all parameters (a, μ, σ) are constant at their new steady-state values.

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

Figure 3 plots the perfect foresight transition dynamics of the general equilibrium version of our model (with the benchmark experiment $k = 2$), from its initial steady state to its long-run new steady state after the rise in the Pareto tail of the income distribution. As discussed in the main text and elsewhere in the appendix, we phase in the rise in inequality from 1980 to 2011. The slight discontinuity in interest rates around 2011 on the right panel reflects the end of this phase-in, with income inequality staying constant thereafter, but interest rates continuing to decline.

Note, however, that convergence is much faster than the convergence of partial equilibrium wealth in figure 1: whereas wealth/GDP is only a fraction of the way toward its new steady state in partial equilibrium by 2030 (dashed line) in figure 1, it is a majority of the way toward its new steady state in general equilibrium by 2030 in figure 3.

The long-run increase in wealth/GDP is also much less dramatic in general equilibrium: even when $k = 2$, it increases by slightly over 10% in figure 3, whereas partial equilibrium wealth/GDP in figure 1 nearly doubles. The increase is much smaller in general equilibrium because the endogenous fall in interest rates discourages households from accumulating so much wealth. In our calibration, this margin is more elastic with respect to interest rates than the capital demand margin on the production side of the economy, and therefore bears most of the burden of equilibriation.[7]

C Results for the wealth distribution

Our model also endogenously generates a household wealth distribution. Although it is not the primary focus of this paper, it is also edifying to study this distribution, especially at the top, and how it varies as we make the income distribution more concentrated.

Figure 4 shows the shares of wealth and income held by the top 1% and 0.1% in the model. Several features are apparent:

a) Wealth is endogenously much more concentrated than income. In our 1980 calibration, the top 1% wealth share is 23.6% and the top 0.1% wealth share is 6.5%. (For comparison, these are just slighly below the Emmanuel Saez and Gabriel Zucman (2016) estimates, which are respectively 24.3% and 8.0% for that year.)

b) A positive, permanent shock to income concentration induces an upward movement in wealth concentration. The induced wealth concentration effect is, in fact, somewhat larger in percentage point terms, although not in relative terms. Following a rise in the top 1% pre-tax income share from 6.4% to 11.0% (1.72x), the top 1% wealth share goes from 23.6% to 33.6% in general equilibrium (1.42x), and from 23.6% to 27.0% (1.14x) in partial equilibrium.

c) On its own, the increase in income inequality is not able to match the trends in wealth concentration documented by Saez and Zucman (2016). Our model first

[7] The 1980 level of wealth/GDP is differs in the left panel figure 3 for each choice of k. This is because we are displaying perfect foresight paths *after* the shock becomes known. Due to capital adjustment costs, the anticipated general equilibrium adjustment of r leads to changes in the valuation q of capital from its pre-shock steady state. Although r declines in the long run for every k, they initially increase by enough in the $k = 0$ case to offset that, leading to a decline on impact in wealth/GDP from its previous steady-state level of 3. In the $k = 2$ case, by contrast, the long-run decline in r dominates, leading to an increase in wealth/GDP on impact.

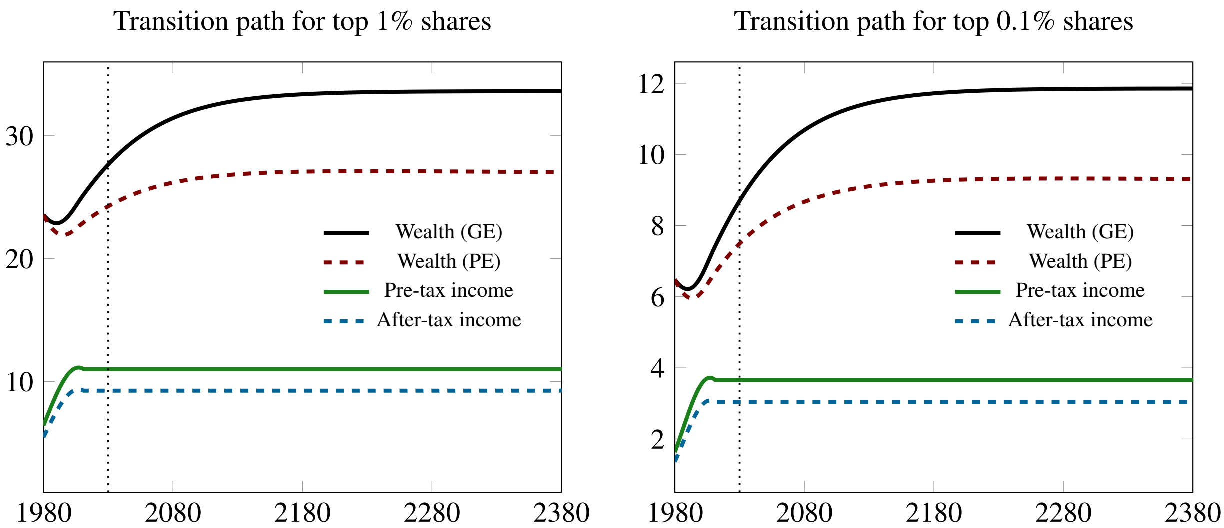

Figure 4: Simulated concentration of wealth and income in the model.

misses the level: their latest (2012) numbers indicate a top 1% share of 41.8%, and a top 0.1% share of 22.0%. By contrast our *steady-state* general equilibrium numbers are, respectively, 33.6% and 11.8%. It also misses the trend: by 2011, when the rise in the top 1% income share is entirely phased in, wealth concentration in our simulations has barely moved at all from its original level. Slightly less than half of the transition occurs in the first 50 years. This delayed convergence has echoes of the slow convergence obtained for the Pareto tail of income in Gabaix et al. (2016), although here the dynamics arise from the consumption-savings decisions of households rather than from a stochastic income process.

d) Steady-state wealth is more concentrated in our general equilibrium experiments than in our partial equilibrium experiments. This is because interest rates fall in general equilibrium, and, perhaps suprisingly given that lower real interest rates make wealth accumulation more difficult, declining interest rates lead to higher wealth concentration in our model.

Figure 5 shows a different feature of the wealth and income distributions: the tail Pareto parameter. This is inferred from our simulation results by comparing shares held by the top 0.01% and the top 0.001%, and then using the Pareto identity

$$1 - \frac{1}{\alpha} = \frac{\log\left(\frac{0.01\%\text{ share}}{0.001\%\text{ share}}\right)}{\log\left(\frac{0.01\%}{0.001\%}\right)} \tag{12}$$

Transition path for implied tail Pareto parameters

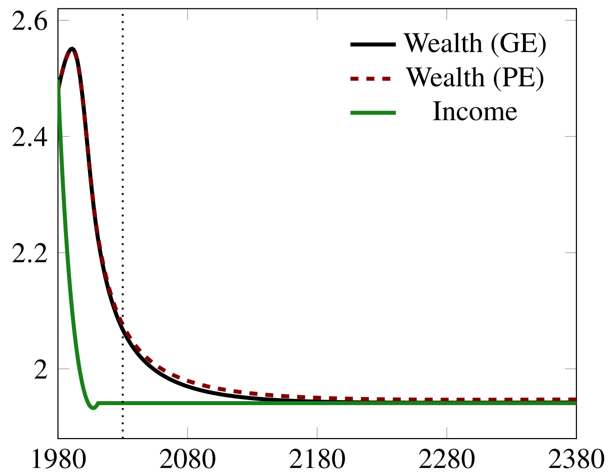

Figure 5: Simulated dynamics of tail thickness in the model.

The comparison of the top 0.01% and 0.001% shares is chosen because the wealth distribution is roughly Pareto this far out in the tail, and because approximation error starts becoming more significant at higher quantiles.

Figure 5 provides numerical confirmation of an analytical result in Jess Benhabib, Alberto Bisin and Shenghao Zhu (2015), which finds that in the stationary steady state of standard Bewley model without idiosyncratic return risk, the wealth distribution has a Pareto tail with the same parameter as the income distribution. The figure also confirms that convergence to a thicker tail for wealth occurs more slowly than for income.

Another important lesson emerges from the contrast between figures 4 and 5. In figure 4, wealth appears substantially more concentrated than income, even though in figure 5 both distributions have the same steady-state Pareto tail parameters. In short, a larger fraction of wealth than income is held by the top 1%, but in the model the shape of the distribution *within* the top 1% (and, even more so, higher percentiles) is roughly the same for wealth and income.

Taken as a whole, these results suggest that models in the Aiyagari (1994) tradition may be able to match a substantial component of wealth inequality when augmented with income processes that generate Pareto tails. They are unable, however, to match phenomena that are specific to the right tail of the wealth distribution. These include the Pareto shape parameter of the tail, which Saez and Zucman (2016) indicates to be 1.43 in 2011.[8] The model here also has difficulty matching recent increases in wealth inequality that have been largely confined to the tail, also documented in Saez and Zucman (2016). To obtain these features of the wealth distribution, it is likely necessary to add additional elements to the model—for instance, idiosyncratic return risk, entrepreneurship, or bequests. Benhabib, Bisin and Zhu (2015) and Mariacristina De Nardi and Giulio Fella (2017) outline some of the possibilities.

D Obtaining our decompositions (5) and (6)

Recall that our discretization of the income process implies a $J \times J$ Markov transition matrix Π, together with levels of incomes $y_1, \ldots y_J$ where $y_j = e^{\underline{y}+a_j}$, and that in our $k = 2$ experiment the transition matrix Π is fixed as we change the income levels y_j are changed.

Write $W(\Theta, \Pi, y_1, \ldots y_J)$ for the steady-state level of wealth generated by our partial-equilibrium household model, given parameters $\Theta = (\beta, \nu, r, \tau_r)$, together with Π and y_j. Given fixed Θ, Π, a first-order Taylor expansion of W yields a total change in W equal to

$$dW = \sum_{j=1}^{J} \frac{\partial W_{y_j}}{\partial y_j} dy_j + o(\|dy\|) \tag{13}$$

Write π_j for the weight of y_j in the stationary distribution induced by Π, then (13) implies that

$$\frac{dW}{W} = \sum_{j=1}^{J} \pi_j \frac{\partial W_{y_j}/\partial y_j}{\pi_j W} dy_j + o(\|dy\|)$$

Write $\epsilon_{Wj} \equiv \frac{\partial W_{y_j}/\partial y_j}{\pi_j W}$, and drop higher-order terms for notational simplicity. This yields

$$\frac{dW}{W} = \mathbb{E}\left[\epsilon_{Wj} dy_j\right] = \mathrm{Cov}\left(\epsilon_{Wj}, dy_j\right) + \mathbb{E}\left[\epsilon_{Wj}\right] \mathbb{E}\left[dy_j\right]$$

But notice that by construction our income process $\mathbb{E}\left[dy_j\right] = 0$, so we obtain (5).

Similarly, consider the change in consumption at date 0 induced by a change in income at date 0 alone. We can write date-0 aggregate consumption as $C_0(\Theta, \Pi, \mathbf{y}; y_{01}, \ldots y_{0J})$, where now $\mathbf{y} = (y_1, \ldots y_J)$ represents income in each state at all future dates and is held

[8] This is inferred analogously to (12) by comparing the 0.1% and 0.01% shares for that year, which were 20.3% and 10.1%, respectively.

fixed. Then a first-order Taylor expansion of C_0 with respect to this date-0 change is

$$dC_0 = \sum_{j=1}^{J} \pi_j \frac{\partial C_{y_{0j}}}{\pi_j \partial y_{0j}} dy_{0j} + o\left(\|dy_0\|\right)$$

But note that

$$C_0 = \sum_{j=1}^{J} \pi_j \int c_0\left(b, y_j\right) d\Psi_j\left(b\right)$$

where $c_0\left(b, y_j\right)$ is the date-0 policy function of agents with wealth level b and income level j, and $\Psi_j\left(b\right)$ is the density function for wealth b conditional on income being y_j. Hence $\frac{\partial C_{y_{0j}}}{\pi_j \partial y_{0j}} = \int MPC_j\left(b\right) d\Psi_j\left(b\right)$, the average marginal propensity to consume of agents with income level j, which we simply write MPC_j. This delivers

$$dC_0 = \mathbb{E}\left[MPC_j dy_{0j}\right]$$

and noting once again that $\mathbb{E}\left[dy_{0j}\right] = 0$, we obtain (6).