

A Schumpeterian Model of Top Income Inequality

Charles I. Jones

Stanford GSB and NBER

Jihee Kim *

KAIST

June 17, 2014 – Version 0.7
Preliminary and Incomplete

Abstract

Top income inequality rose sharply in the United States over the last 30 years but increased only slightly in economies like France and Japan. Why? This paper explores a model in which the interplay between entrepreneurial efforts to grow existing market share and the creative destruction associated with outside innovation lead top incomes to obey a Pareto distribution. The extent to which entrepreneurs, interpreted very broadly, can grow the markets for their existing ideas is a key determinant of top income inequality, while the creative destruction by which one entrepreneur replaces another restrains inequality. Any changes in the economy that lead entrepreneurs to exert more effort at building market share will increase top inequality. Examples might include the creation of the world wide web, a reduction in top tax rates, a decline in misallocation, and possibly globalization. On the other hand, policies that restrain the growth of market share by entrepreneurs can limit inequality. Differences in these considerations across countries may help explain the varied patterns of top income inequality that we see in the data.

*We are grateful to Xavier Gabaix, Mike Harrison, Pete Klenow, Ben Moll, Chris Tonetti, Alwyn Young, and seminar participants at the AEA annual meetings, CREI, the Federal Reserve Bank of San Francisco, Princeton, Stanford, USC, and Zurich for helpful comments.

1. Introduction

As documented extensively by Piketty and Saez (2003) and Atkinson, Piketty and Saez (2011), top income inequality — such as the share of income going to the top 1% or top 0.1% of earners — has risen sharply in the United States since around 1980. The pattern in other countries is different and heterogeneous, however. For example, top inequality rose only slightly in France and Japan. Why? What economic forces explain the rise in top income inequality in the United States? And why has top inequality risen by much less in countries like France and Japan?

It is well-known that the upper tail of the income distribution follows a power law. One way of thinking about this is to note that income inequality is fractal in nature, as we document more carefully below. In particular, the following questions all have essentially the same answer: What fraction of the income going to the top 10% of earners accrues to the top 1%? What fraction of the income going to the top 1% of earners accrues to the top 0.1%? What fraction of the income going to the top 0.1% of earners accrues to the top 0.01%? The answer to each of these questions — which turns out to be around 40% in the United States today — is a simple function of the parameter that characterizes the power law. Therefore *changes* in top income inequality naturally involve changes in the power law parameter. This paper considers a range of economic explanations for changes in this parameter. We remain far from being able to point to a single explanation, say, for explaining the rise in top income inequality in the United States. Instead, we present a model that highlights different economic forces that could potentially contribute to such an explanation.

The model we develop uses the Pareto-generating mechanisms that researchers like Gabaix (1999) and Luttmer (2007) have used in other contexts. Gabaix studied why the distribution of city populations is Pareto with its key parameter equal to unity; Luttmer studies why the distribution of employment by firms has the same structure. It is worth noting that both cities and firm sizes exhibit substantially more inequality than top incomes (power law inequality for incomes is around 0.6, as we show below, versus around 1 for city populations and firm employment). Our approach therefore is slightly different: why incomes are Pareto and why Pareto

inequality is changing over time, rather than why a power law inequality measure is so close to unity.¹ The basic insight in this literature is that exponential growth, tweaked slightly, can deliver a Pareto distribution for outcomes. The tweak is needed for the following reason. Suppose that city populations (or incomes or employment by firms) grow exponentially at 2% per year plus some random normally-distributed shock. In this case, the log of population would follow a normal distribution with a variance that grows over time. To keep the distribution from spreading out forever, we need some kind of “tweak.” For example, a constant probability of death will suffice to render the distribution stationary, as would some minimum size for the city.

In the model we develop below, researchers create new ideas — new songs, best-selling books, smartphone apps, financial products, surgical techniques, or even new ways of organizing a law firm. The random growth process corresponds to the way entrepreneurs increase their productivity and build market share for their new products. The growth rate of this process is tied to entrepreneurial effort, and anything that raises this effort, resulting in faster growth of idiosyncratic productivity, will raise top income inequality. The “death rate” in our setup is naturally tied to creative destruction: researchers invent new ideas that make the previous state-of-the-art surgical technique or best-selling iPad application obsolete. A higher rate of creative destruction restrains the random growth process and results in lower top income inequality. In this way, the interplay between existing entrepreneurs growing their markets and the creative destruction associated with new ideas determines top income inequality.

This paper proceeds as follows. Section 2 presents some basic facts of top income inequality, emphasizing that the rise in the United States is accurately characterized by a change in the power law parameter. Section 3 considers a simple model to illustrate the basic mechanism in the paper. The next two sections then develop the model, first with an exogenous allocation of labor to research and then

¹These papers in turn build on a large literature on such mechanisms outside economics. For example, see Reed (2001), Mitzenmacher (2004), and Malevergne, Saichev and Sornette (2010). Gabaix (2009) and Luttmer (2010) have excellent surveys of these mechanisms, written for economists.

more fully with an endogenous allocation of labor. Section 6 uses the IRS public use panel of tax returns to estimate several of the key parameters of the model, illustrating that the mechanism can lead to rising top inequality. A final main section studies numerical examples to illustrate several additional quantitative possibilities of the framework.

1.1. Skill-Biased Technical Change?

Before moving on, it is worth pausing to consider a simple, familiar explanation in order to understand why it is incomplete: skill-biased technical change. For example, if the distribution of skill is Pareto and there is a rise in the return to skill, does this raise top inequality? The answer is no, and it is instructive to see why.

Suppose that income for highly-skilled person i is $y_i = \bar{w}x_i^\alpha$, where x_i is person i 's skill and \bar{w} is the wage per unit of skill. If the distribution of skill across people is Pareto with inequality parameter η_x , then the income distribution will be Pareto with inequality parameter $\eta_y = \alpha\eta_x$. That is, if $\Pr[x_i > x] = x^{-1/\eta_x}$, then $\Pr[y_i > y] = (\frac{y}{\bar{w}})^{-1/\eta_y}$. An increase in \bar{w} — a skill-biased technical change that increases the return to skill — shifts the entire distribution right but does not change Pareto inequality η_y . Take for granted for the moment that a changing η_y is what the data require — this will be explained in more detail in the next section. The point is that a simple story of skill-biased technical change is not enough.

On the other hand, if the exponent α were to rise over time, this would lead to a rise in Pareto inequality. But this requires something more than just a simple skill-biased technical change story. Moreover, even a rising α would leave unexplained the question of why the underlying skill distribution is Pareto. The remainder of this paper can be seen as explaining why x is Pareto and what economic forces might cause α to change over time or differ across countries.

1.2. The existing literature


A number of other recent papers contribute to our understanding of the dynamics of top income inequality. Piketty, Saez and Stantcheva (2014) and Rothschild and

Scheuer (2011) explore the possibility that the decline in top tax rates has led to a rise in rent seeking, leading top inequality to increase. Philippon and Reshef (2009) focus explicitly on finance and the extent to which rising rents in that sector can explain rising inequality; see also Bell and Van Reenen (2010). Bakija, Cole and Heim (2010) and Kaplan and Rauh (2010) note that the rise in top inequality occurs across a range of occupations; it is not just focused in finance or among CEOs, for example, but includes doctors and lawyers and star athletes as well. Benabou and Tirole (2013) discuss how competition for the most talented workers can result in a “bonus culture” with excessive incentives for the highly skilled. Haskel, Lawrence, Leamer and Slaughter (2012) suggest that globalization may have raised the returns to superstars via a Rosen (1981) mechanism. There is of course a much larger literature on changes in income inequality throughout the distribution. Katz and Autor (1999) provide a general overview, while Autor, Katz and Kearney (2006), Gordon and Dew-Becker (2008), and Acemoglu and Autor (2011) provide more recent updates.

Lucas and Moll (2014) explore a model of human capital and the sharing of ideas that gives rise to endogenous growth. Perla and Tonetti (2014) study a similar mechanism in the context of technology adoption by firms. These papers show that if the initial distribution of human capital or firm productivity has a Pareto upper tail, then the ergodic distribution also inherits this property and the model can lead to endogenous growth, a result reminiscent of Kortum (1997). The Pareto distribution, then, is more of an “input” in these models rather than an outcome.

The most closely-related papers to this one are Benhabib, Bisin and Zhu (2011), Nirei (2009), Aoki and Nirei (2013), Moll (2012), Piketty and Saez (2012), and Piketty and Zucman (2014). These papers study economic mechanisms that generate endogenously a Pareto distribution for wealth, and therefore for capital income. The mechanism responsible for random growth in these papers is either the asset accumulation equation (which naturally follows a random walk when viewed in partial equilibrium) or the capital accumulation equation in a neoclassical growth model. The present paper differs most directly by focusing on labor income rather than wealth. Since much of the rise in top income inequality in the United States is due to labor income — e.g. see Piketty and Saez (2003) — this focus is worthwhile.

Figure 1: Top Income Inequality in the United States and France



Source: World Top Incomes Database.

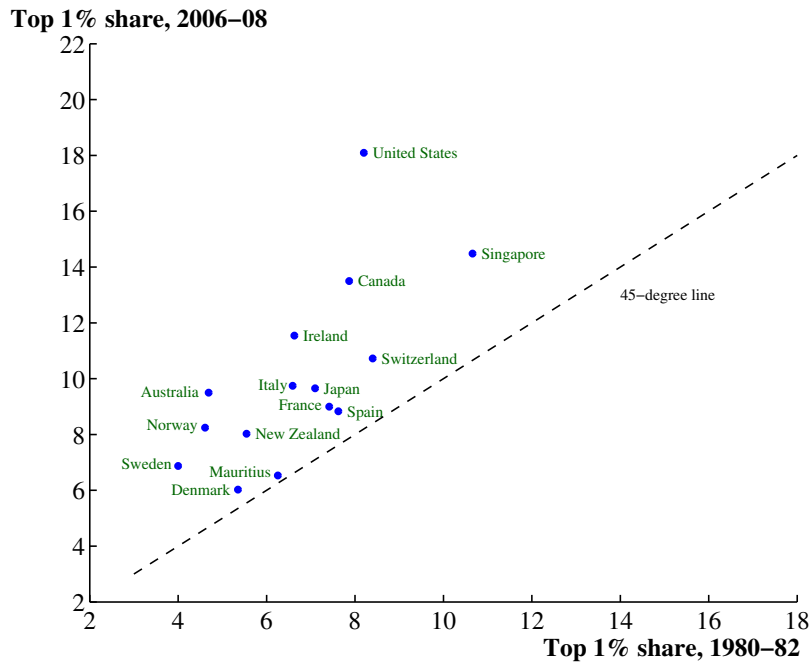
As mentioned earlier, this paper is not able to conclusively distinguish between various explanations that might contribute to the different patterns of top income inequality that we see around the world. Instead, the goal is less ambitious. Drawing on Pareto-generating mechanisms that have been used to study Zipf's Law for cities and the size distribution of firms, we hope to provide a baseline model that explains why top income inequality obeys a power law and suggest some of the forces that may contribute to the heterogeneity we see across countries and over time.

2. Some Basic Facts

Figures 1 and 2 show some of the key facts about top income inequality that have been documented by Piketty and Saez (2003) and Atkinson, Piketty and Saez (2011). For example, the first figure shows the large increase in top inequality for the United States since 1980, compared to the relative stability of inequality in France.

Figure 2 shows the dynamics of top income inequality for a range of countries.


Figure 2: Top Income Inequality around the World, 1980-82 and 2006–2008



Note: Top income inequality has increased since 1980 in most countries for which we have data. The size of the increase varies substantially, however. Source: World Top Incomes Database.

The horizontal axis shows the share of aggregate income going to the top 1% of earners, averaged between 1980 and 1982, while the vertical axis shows the same share for 2006–2008. Nearly all the economies for which we have data lie above the 45-degree line: that is, top income inequality has risen almost everywhere. The rise is largest in the United States, South Africa, and Norway, but substantial increases are also seen elsewhere, such as in Ireland, Portugal, Singapore, Italy, and Sweden. Japan and France exhibit smaller but still noticeable increases. For example, the top 1% share in France rises from 7.4% to 9.0%.

Figure 3: The Composition of the Top 0.1 Percent Income Share



Note: The figure shows the composition of the top 0.1 percent income share. Source: These data are taken from the “data-Fig4B” tab of the September 2013 update of the spreadsheet appendix to Piketty and Saez (2003).

2.1. The Role of Labor Income

As discussed by Atkinson, Piketty and Saez (2011), a substantial part of the rise in U.S. top income inequality represents a rise in labor income inequality, particularly if one includes “business income” (i.e. profits from sole proprietorships, partnerships and S-corporations) in the labor income category. Figure 3 shows an updated version of their graph for the period since 1950, supporting this observation.

Because the model in this paper is based on labor income as opposed to capital income, documenting the Pareto nature of labor income inequality in particular is also important. It is well known, dating back to Pareto (1896), that the top portion of the income distribution can be characterized by a power law. That is, at high levels, the income distribution is approximately Pareto. In particular, if Y is a random

variable denoting incomes, then, at least above some high level (i.e. for $Y \geq y_0$)

$$\Pr [Y > y] = \left(\frac{y}{y_0} \right)^{-\xi}, \quad (1)$$

where ξ is called the “power law exponent.”

Saez (2001) documents that wage and salary income from U.S. income tax records in the early 1990s is well-described by a Pareto distribution. Figure 4 replicates his analysis for 1980 and 2005. In particular, the figures show mean wage income above some threshold as a ratio to the threshold itself. If wage income obeys a Pareto distribution like that in (1), then this ratio should equal the constant $\frac{\xi}{\xi-1}$, *regardless of the threshold*. That is, as we move to higher and higher income thresholds, the ratio of average income above the threshold to the threshold itself should remain constant.² Figure 4 shows that this property holds reasonably well in 1980 and 2005, and also illustrates that the ratio has risen substantially over this period, reflecting the rise in top wage income inequality.

2.2. Fractal Inequality and the Pareto Distribution

An important property of Pareto distributions is that they exhibit a fractal pattern of top inequality. To see this, let $\tilde{S}(p)$ denote the share of income going to the top p percentiles. For the Pareto distribution defined in equation (1) above, this share is given by $(p/100)^{1-1/\xi}$. A larger power-law exponent, ξ , is associated with lower top income inequality. It is therefore convenient to define the “power-law inequality” exponent as

$$\eta \equiv \frac{1}{\xi} \quad (2)$$

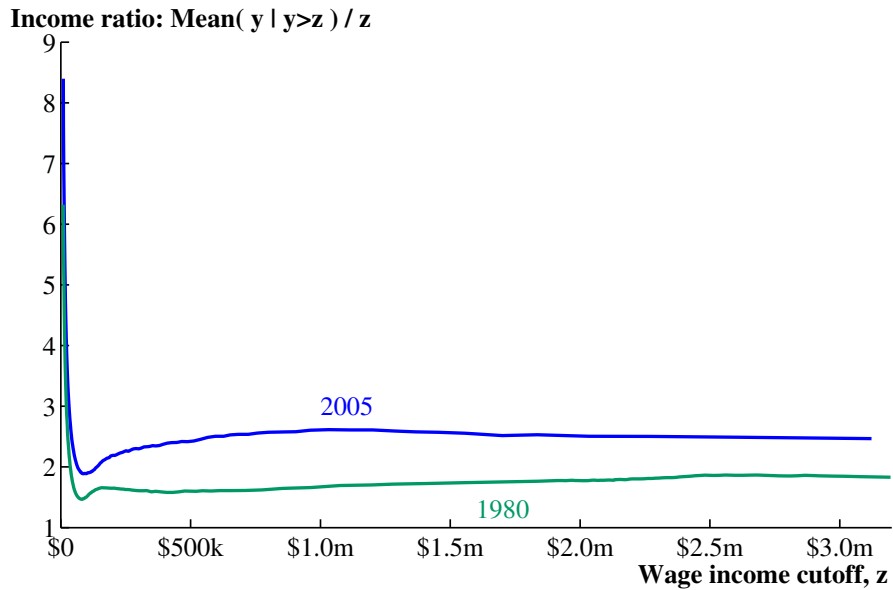
so that

$$\tilde{S}(p) = \left(\frac{100}{p} \right)^{\eta-1}. \quad (3)$$

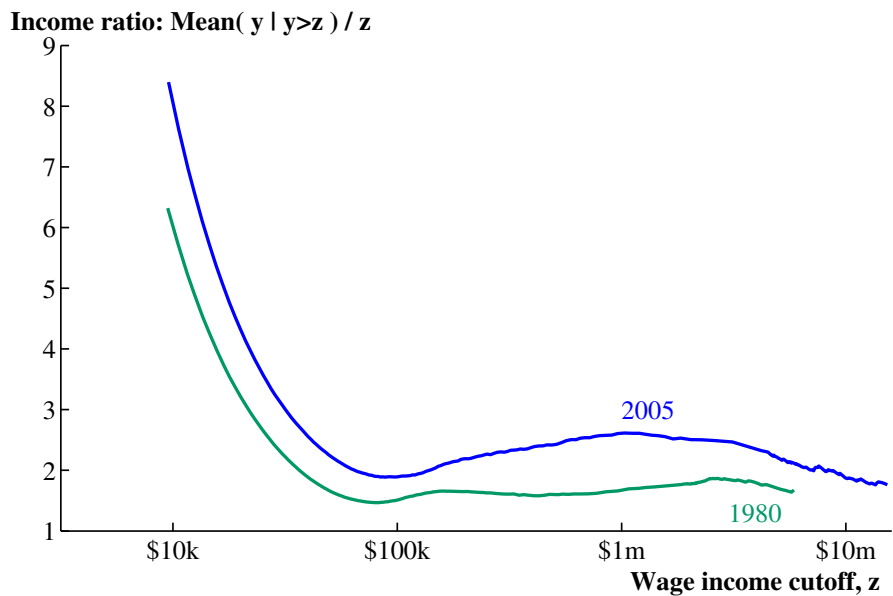
For example, if $\eta = 1/2$, then the share of income going to the top 1% is $100^{-1/2} = .10$. However, if $\eta = 3/4$, the share going to the top 1% rises sharply to $100^{-1/4} \approx 0.32$.

²This follows easily from the fact that the mean of a Pareto distribution is $\frac{\xi y_0}{\xi-1}$ and that the conditional mean just scales up with the threshold.

Figure 4: The Pareto Nature of Wage Income



(a) Linear scale, up to \$3 million



(b) Log scale

Note: The figures plot the ratio of average wage income above some threshold z to the threshold itself. For a Pareto distribution with Pareto inequality parameter η , this ratio equals $1/1 - \eta$. Saez (2001) produced similar graphs for 1992 and 1993 using the IRS public use tax files available from the NBER at www.nber.org/taxsim-notes.html. The figures here replicate these results using the same data source for 1980 and 2005.

To see the fractal structure of top inequality under a Pareto distribution, let $\tilde{S}(a, b)$ denote the fraction of income earned by the top b percent of people that actually goes to the top a percent. For example, $\tilde{S}(1, 10)$ answers the question, “Of the income going to the top 10 percent of earners, what fraction actually goes to the top 1 percent?” If income follows a Pareto distribution, then

$$\begin{aligned}\tilde{S}(a, b) &= \frac{\tilde{S}(a)}{\tilde{S}(b)} \\ &= \left(\frac{b}{a}\right)^{\eta-1}\end{aligned}$$

It is often useful — certainly given the data in the World Top Incomes Database — to look at shares where b is 10 times larger than a . In this case, let’s define $S(a) \equiv \tilde{S}(a, 10 \cdot a)$. For example, $S(1)$ is the fraction of income going to the top 10% that actually accrues to the top 1%, and $S(.1)$ is the fraction of income going to the top 1% that actually goes to the top 1 in 1000 earners. Under a Pareto distribution,

$$S(a) = 10^{\eta-1} \tag{4}$$

and


$$\log_{10} S(a) = \eta - 1. \tag{5}$$

Notice that this last result holds for all values of a , or at least for all values for which income follows a Pareto distribution. This means that top income inequality obeys a *fractal* pattern: the fraction of the Top 10 percent’s income going to the Top 1 percent is the same as the fraction of the Top 1 percent’s income going to the Top 0.1 percent, which is the same as the fraction of the Top 0.1 percent’s income going to the Top 0.01 percent.

Not surprisingly, top income inequality is well-characterized by this fractal pattern.³ Figure 5 shows the $S(a)$ shares directly. At the very top, the fractal prediction holds remarkably well, and $S(.01) \approx S(.1) \approx S(1)$. This is another indication that

³Others have noticed this before. For example, see Aluation.wordpress.com (2011).

Figure 5: Fractal Inequality of U.S. Wage Income



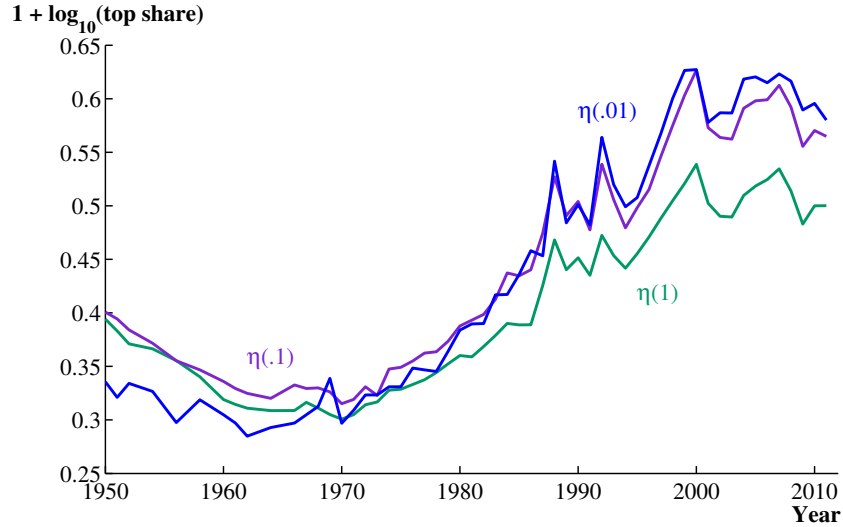
Note: $S(a)$ denotes the fraction of income going to the top 10a percent of earners that actually goes to the top a percent. For example, $S(1)$ is the share of the Top 10%'s income that accrues to the Top 1%. Source: Underlying wage income shares are from the September 2013 update of the Piketty and Saez (2003) spreadsheet appendix.

wage income is well-described as a Pareto distribution. Prior to 1980, the fractal shares are around 25 percent: one quarter of the Top X percent's income goes to the Top $X/10$ percent. By the end of the sample in 2010, this fractal share is closer to 40 percent.

This rise in top income inequality shown in Figure 5 can be related directly to the power-law income inequality exponent using equation (5). Or, put another way, the *change* in the fractal shares is precisely equal to the change in the PL inequality exponent:

$$\Delta \log_{10} S(a) = \Delta \eta. \quad (6)$$

The corresponding Pareto inequality measures are shown in Figure 6.

Figure 6: The Power-Law Inequality Exponent η , United States

Note: $\eta(a)$ is the inequality power law exponent obtained from the fractal inequality wage income shares in Figure 5 assuming a Pareto distribution. See equation (5) in the text.

2.3. Summary

Here then are the basic facts related to top income inequality that we'd like to be able to explain. Between 1960 and 1980, top income inequality was relatively low and stable in both the United States and France. Since around 1980, however, top inequality has increased sharply in countries like the United States, Norway, and Portugal, while it has increased only slightly in others, including France and Japan. Finally, labor income is well-described by a Pareto distribution, and rising top income inequality is to a great extent associated with rising labor income inequality. Changing top income inequality corresponds to a change in the power-law inequality exponent. The remainder of this paper develops and analyzes a model to help us understand these facts.

3. A Simple Model of Top Income Inequality

It is well-known that exponential growth and Pareto distributions are tightly linked, and this link is at the heart of the basic mechanism in this paper. To illustrate this point in the clearest way, we begin with a simple, stylized model.⁴

When a person first becomes a top earner (“entrepreneur”), she earns income y_0 . As long as she remains a top earner, her income grows over time at rate μ , so the income of a person who’s been a top earner for x years — think of x as “entrepreneurial experience” — is $y(x) = y_0 e^{\mu x}$. The growth rate μ can be thought of as the return to entrepreneurial experience.

People do not remain top earners forever. Instead, there is a constant probability δ per unit of time (more formally, a Poisson process) that an existing entrepreneur is displaced. If this occurs, the existing entrepreneur drops out of the top, becoming a “normal” worker, and is replaced by a new entrepreneur who starts over at the bottom of the ladder and earns y_0 .

What fraction of people in this economy have income greater than some level y ? The answer is simply the fraction of people with at least $x(y)$ years of experience, where

$$x(y) = \frac{1}{\mu} \log \left(\frac{y}{y_0} \right). \quad (7)$$

With a Poisson replacement process, it is well-known that the distribution of experience for a given individual follows an exponential distribution, i.e. $\Pr[\text{Experience} > x] = e^{-\delta x}$. Let’s take for granted for the moment that the stationary distribution of experience across a population of entrepreneurs will be this same exponential distribution; we’ll return to consider this step in more detail shortly. Then, the remain-

⁴Gabaix (2009) presents a version of the model developed below in the context of Zipf’s Law for cities and attributes the original development of this model for that application to Steindl (1965).

der of the argument is straightforward:

$$\begin{aligned}
 \Pr [\text{Income} > y] &= \Pr [\text{Experience} > x(y)] \\
 &= e^{-\delta x(y)} \\
 &= \left(\frac{y}{y_0} \right)^{-\frac{\delta}{\mu}}
 \end{aligned} \tag{8}$$

which is a Pareto distribution!

The power law exponent for income in this model is then $\xi_y = \delta/\mu$, and power law inequality is given by

$$\eta_y = \frac{\mu}{\delta}. \tag{9}$$

In this model, top income inequality can change for two reasons. First, an increase in the returns to experience, μ , among top earners will widen the distribution of top incomes: the higher is the growth rate of incomes, the higher is the ratio of top incomes to the income of a new entrepreneur. Second, an increase in the “death rate” δ will reduce top income inequality, as top entrepreneurs have less time during which to build their advantage.

What are the economic determinants of μ and δ , and why might they change over time? Answering these questions is one of the goals of the full model that we develop below.

3.1. Intuition

The logic of the simple model provides useful intuition about *why* the Pareto result emerges. First, in equation (7), the log of income is proportional to experience. This is a common and natural assumption. For example, in models where income grows exponentially over time, income and time are related in this way. Or in labor economics, log income and experience are linked in Mincer-style equations. Next, the distribution of experience is exponential. This is a property of a Poisson process with a constant arrival rate. Putting these two pieces together, log income has an exponential distribution. But this is just another way of saying that income has a Pareto distribution.

3.2. The Stationary Distribution of Experience

It is helpful to show the argument that the stationary distribution of experience is exponential. The reason is that this illustrates a simple version of the Kolmogorov forward equation, which is used later in solving for inequality in the full stochastic model.

Let $F(x, t)$ denote the distribution of experience at time t , and consider how this distribution evolves over some discrete time interval Δt :⁵

$$F(x, t + \Delta t) - F(x, t) = \underbrace{\delta \Delta t (1 - F(x, t))}_{\text{inflow from above } x} - \underbrace{[F(x, t) - F(x - \Delta x, t)]}_{\text{outflow as top folks age}}$$

Dividing both sides by $\Delta t = \Delta x$ and taking the limit as the time interval goes to zero yields:

$$\frac{\partial F(x, t)}{\partial t} = \delta(1 - F(x, t)) - \frac{\partial F(x, t)}{\partial x} \quad (10)$$

One can continue with this equation directly. But for what comes later, it is more useful to take the derivative of both sides of this equation with respect to x . Letting $f(x, t) := \frac{\partial F(x, t)}{\partial x}$ denote the pdf,

$$\frac{\partial f(x, t)}{\partial t} = -\delta f(x, t) - \frac{\partial f(x, t)}{\partial x}. \quad (11)$$

This equation is the non-stochastic version of the Kolmogorov forward equation that we will see later. The intuition underlying this equation is easiest to see in the version involving the cdf, equation (10), which just involves the inflows and outflows mentioned earlier.

Finally, to solve for the stationary distribution, we seek $f(x)$ such that $\frac{\partial f(x, t)}{\partial t} = 0$. Therefore

$$0 = -\delta f(x) - \frac{df(x)}{dx}.$$

Integrating this equation twice yields the result that the stationary distribution is

⁵This equation drops a term involving both Δt and Δx , as it goes to zero later when we take limits.

exponential. In particular,

$$F(x) = 1 - e^{-\delta x}.$$

4. A Schumpeterian Model of Top Income Inequality

The simple model illustrates in a reduced-form fashion the main mechanism at work in this paper. In our full model, we develop a theory in which the economic determinants of μ and δ are apparent, and we consider what changes in the economic environment could be responsible for the range of patterns we see in top income inequality across countries. In the full model, effort by entrepreneurs influences the growth of their incomes. In addition, this process is assumed to be stochastic, which allows us to better match up the model with micro data on top incomes. Finally, the death rate is made to be endogenous by tying it to the process of creative destruction in a Schumpeterian endogenous growth model. The setup seems to capture some of the key features of top incomes: the importance of entrepreneurial effort, the role of creative destruction, and the centrality of “luck” as some people succeed beyond their wildest dreams while others fail.

4.1. Entrepreneurs and Market Share

An entrepreneur is a monopolist with the exclusive right to sell a particular variety, in competition with other varieties. We interpret this statement quite broadly. For example, think of a Silicon Valley startup, an author of a new book, a new rock band, an athlete just making it to the pro’s, or a doctor who has invented a new surgical technique. Moreover, we do not associate a single variety with a single firm — the entrepreneur could be a middle manager in a large company who has made some breakthrough and earned a promotion.

When a new variety is first introduced, it has a low quality/productivity, denoted by x , which in turn will correspond to a low market share. The entrepreneur then expends effort to improve x . We explain later how x affects firm productivity and profitability. For the moment, however, it is sufficient to just assume that the en-

trepreneur's income is proportional to x , as it will be in general equilibrium. Note that we are recycling notation: this x does not measure experience as it did in the simple model of Section 2.

Given an x , the entrepreneur maximizes the expected present discounted value of utility, $u(c, \ell) = \log c_t + \beta \log \ell_t$, subject to the following constraints:

$$c_t = \psi_t x_t \tag{12}$$

$$e_t + \ell_t + \tau = 1 \tag{13}$$

$$dx_t = \mu(e_t)x_t dt + \sigma x_t dB_t \tag{14}$$

$$\mu(e) = \phi e \tag{15}$$

For simplicity, we do not allow entrepreneurs to smooth their consumption and instead assume that consumption equals income, which in turn is proportional to productivity x . The factor of proportionality, ψ_t , is exogenous to the individual's actions and is the same for all entrepreneurs; it is endogenized in general equilibrium shortly. The entrepreneur has one unit of time each period, which can be used for effort e or leisure ℓ or it can be wasted, in amount τ . This could correspond to time spent addressing government regulations and bureaucratic red tape, for example.

Equation (14) describes how effort improves idiosyncratic productivity x through a geometric Brownian motion. The average growth rate of productivity is $\mu(e) = \phi e$, where ϕ is a technological parameter converting effort into growth. dB_t denotes the standard normal increment to the Brownian motion.

Finally, there is a Poisson creative destruction process by which the entrepreneur loses her monopoly position and is replaced by a new entrepreneur. This occurs at the (endogenized in general equilibrium) rate δ . In addition, there is an exogenous piece to destruction as well, which occurs at a constant rate $\bar{\delta}$.

The Bellman equation for the entrepreneur is

$$\rho V(x_t, t) = \max_e \log \psi_t + \log x_t + \beta \log(\Omega - e_t) + \frac{\mathbb{E}[dV(x_t, t)]}{dt} + (\delta + \bar{\delta})(V^w(t) - V(x_t, t)) \tag{16}$$

subject to (14), where $\Omega \equiv 1 - \tau$ and $\frac{\mathbb{E}[dV(x_t, t)]}{dt}$ is short-hand for the Ito calculus terms, i.e. $\frac{\mathbb{E}[dV(x_t, t)]}{dt} = \mu(e_t)x_t V_x(x_t, t) + \frac{1}{2}\sigma^2 x_t^2 V_{xx}(x_t, t) + V_t(x_t, t)$. $V(x, t)$ is the expected utility of an entrepreneur with quality x . The flow of the value function depends on the “dividend” of utility from consumption and leisure, the “capital gain” associated with the expected change in the value function, and the possible loss associated with creative destruction, in which case the entrepreneur becomes a worker with expected utility V^w .

The first key result describes the entrepreneur’s choice of effort. (Proofs of all propositions are given in the [appendix](#)).

Proposition 1 (Entrepreneurial Effort): *Entrepreneurial effort solves the Bellman problem in equation (16) and along the balanced growth path is given by*

$$e^* = 1 - \tau - \frac{1}{\phi} \cdot \beta(\rho + \delta + \bar{\delta}). \quad (17)$$

This proposition implies that entrepreneurial effort is an increasing function of the technology parameter ϕ but decreases whenever τ , β , ρ , δ , or $\bar{\delta}$ are higher.

4.2. The Stationary Distribution of Entrepreneurial Income

We assume there are a continuum of entrepreneurs of unit measure at any point in time, corresponding to the usual unit measure of product varieties. The initial distribution of idiosyncratic productivity x is given by $f_0(x)$, and the distribution evolves according to the geometric Brownian motion process given above. Entrepreneurs can be displaced in one of two ways. Endogenous creative destruction (the Poisson process at rate δ) leads to replacement by a new entrepreneur who inherits the existing quality x ; hence the distribution is not mechanically altered by this form of destruction. In large part, this is a simplifying assumption; otherwise one has to worry about the extent to which the higher baseline quality of the new entrepreneur trades off with the higher x that the previous entrepreneur has accumulated. We treat the exogenous destruction at rate $\bar{\delta}$ differently. In this case, existing entrepreneurs are replaced by new “young” entrepreneurs with some base

quality level x_0 . Exogenous destruction could correspond to the actual death or retirement of existing entrepreneurs, or it could stand in for policy actions by the government: one form of misallocation may be that the government appropriates the patent from an existing entrepreneur and gives it to a new favored individual. Finally, it simplifies the analysis to assume that x_0 is the minimum possible productivity: there is a “reflecting barrier” at x_0 ; this assumption could be relaxed.

We’ve set up the stochastic process for x so that we can apply a well-known result in the literature for generating Pareto distributions.⁶ If a variable follows a Brownian motion, like x above, the density of the distribution $f(x, t)$ satisfies a Kolmogorov forward equation. This is the stochastic generalization of an equation like (11) that we saw in the simple model at the start of the paper. In particular, the density satisfies

$$\frac{\partial f(x, t)}{\partial t} = -\bar{\delta}f(x, t) - \frac{\partial}{\partial x} [\mu(e^*)xf(x, t)] + \frac{1}{2} \cdot \frac{\partial^2}{\partial x^2} [\sigma^2 x^2 f(x, t)] \quad (18)$$

If a stationary distribution, $\lim_{t \rightarrow \infty} f(x, t) = f(x)$ exists, it therefore satisfies

$$0 = -\bar{\delta}f(x) - \frac{d}{dx} [\mu(e^*)xf(x)] + \frac{1}{2} \cdot \frac{d^2}{dx^2} [\sigma^2 x^2 f(x)] \quad (19)$$

Guessing that the Pareto form $f(x) = Cx^{-\xi-1}$ solves this differential equation, one obtains the following result:

Proposition 2 (The Pareto Income Distribution): *The stationary distribution of (normalized) entrepreneurial income is given by*

$$F(x) = \left(\frac{x}{x_0} \right)^{-\xi^*} \quad (20)$$

where

$$\xi^* = -\frac{\tilde{\mu}^*}{\sigma^2} + \sqrt{\left(\frac{\tilde{\mu}^*}{\sigma^2} \right)^2 + \frac{2\bar{\delta}}{\sigma^2}} \quad (21)$$

and $\tilde{\mu}^* \equiv \mu(e^*) - \frac{1}{2}\sigma^2 = \phi(1 - \tau) - \beta(\rho + \delta^* + \bar{\delta}) - \frac{1}{2}\sigma^2$. *Power-law inequality is therefore given by $\eta^* \equiv 1/\xi^*$.*

⁶For more detailed discussion, see Reed (2001), Mitzenmacher (2004), Gabaix (2009), and Luttmer (2010). Malevergne, Saichev and Sornette (2010) is closest to the present setup.

The word “normalized” in the proposition refers to the fact that the income of an entrepreneur with productivity x is $\psi_t x$. Aggregate growth occurs via the ψ_t term, as discussed when we turn to general equilibrium, while the distribution of x is what is stationary. Finally, we put a “star” on δ as a reminder that this value is determined in general equilibrium as well.

Comparative statics: Taking δ^* as exogenous for the moment, the comparative static results are as follows: power-law inequality, η^* , increases if effort is more effective at growing market share (a higher ϕ), decreases if the time endowment is reduced by government policy (a higher τ), decreases if entrepreneurs place more weight on leisure (a higher β), and decreases if either the endogenous or exogenous rates of creative destruction rise (a higher δ^* or $\bar{\delta}$).⁷

The analysis so far shows how one can endogenously obtain a Pareto-shaped income distribution. We’ve purposefully gotten to this result as quickly as possible while deferring our discussion of the general equilibrium in order to draw attention to the key economic forces that determine top income inequality.

Next, however, we flesh out the rest of the general equilibrium: how productivity x enters the model, how x affects entrepreneurial income (the proportionality factor ψ_t), and how creative destruction δ^* is determined.

4.3. Production and General Equilibrium

The remainder of the setup is a relatively conventional model of endogenous growth with quality ladders and creative destruction, in the tradition of Aghion and Howitt (1992). A fixed population of people choose to be basic laborers, researchers (searching for a new idea), or entrepreneurs (who have found an idea and are in the process of improving it).

A unit measure of varieties exist in the economy, and varieties combine to pro-

⁷The effect of σ^2 on power-law inequality is more subtle. If $\eta^* > \mu^*/\delta^*$, then a rise in σ^2 increases η^* , while if $\eta^* < \mu^*/\delta^*$, then the effect goes in the other direction.

duce a single final output good:

$$Y = \left(\int_0^1 Y_i^\theta di \right)^{1/\theta}. \quad (22)$$

Each variety is produced by an entrepreneur using a production function that exhibits constant returns to the only nonrivalrous input, labor L_i :

$$Y_i = \gamma^{n_t} x_i^\alpha L_i. \quad (23)$$

The productivity in variety i 's production function depends on two terms. The first captures aggregate productivity growth. The variable n_t measures how far up the quality ladder the variety is, and $\gamma > 1$ is the step size. For simplicity, we assume that a researcher who moves a particular variety up the quality ladder generates spillovers that move all varieties up the quality ladder: in equilibrium, every variety is on the same rung of the ladder. (This just avoids us having to aggregate over varieties at different positions on the ladder.)

The second term is the key place where an entrepreneur's effort decisions pays off: labor productivity depends on x_i^α . In this sense, x_i most accurately captures the entrepreneur's productivity in making variety i . It could equivalently be specified as quality, and, as usual, variety i 's market share is increasing in x_i .

The key resource constraint in this environment involves labor:

$$L_t + R_t + 1 = \bar{N}, \quad L_t \equiv \int_0^1 L_{it} di \quad (24)$$

A fixed measure of people, \bar{N} , are available to the economy. People can work as the raw labor making varieties, or as researchers, R_t , or as entrepreneurs — of which there is always just a unit measure, though their identities can change. It is convenient to define $\bar{L} \equiv \bar{N} - 1$.

Researchers discover new ideas through a Poisson process with arrival rate λ per researcher. Research is undirected and a successful discovery increases the productivity of a randomly chosen variety by a proportion $\gamma > 1$. Once the research is suc-

cessful, the researcher becomes the entrepreneur of that variety, replacing the old entrepreneur by endogenous creative destruction. In addition, as explained above, the new idea generates spillovers that raise productivity in all other varieties as well.

The flow rate of innovation is therefore

$$\dot{n}_t = \lambda R_t \quad (25)$$

and this also gives the rate of creative destruction:

$$\delta_t = \dot{n}_t. \quad (26)$$

4.4. The Allocation of Resources

There are 12 key quantities in this economic environment: $Y, Y_i, x_i, L_i, L, R, n, \delta, e_i, c_i, \ell_i, \psi$. The entrepreneur's choice problem laid out earlier pins down c, ℓ , and e for each entrepreneur. Production functions determine Y, Y_i, n , and δ . This leaves us needing to determine R, L_i , and ψ .

It is easiest to do this in two stages. Conditional on a choice for R , standard equilibrium analysis can easily pin down the other variables, and the comparative statics can be calculated analytically. So to begin, we focus on a situation in which the fraction of people working as researchers is given exogenously: $R/\bar{L} = \bar{s}$. Later, we let markets determine this allocation as well and provide numerical results.

We follow a standard approach to decentralizing the allocation of resources. The final goods sector is perfectly competitive, while each entrepreneur engages in monopolistic competition in selling their varieties. Each entrepreneur is allowed by the patent system to act as a monopolist and charges a markup over marginal cost given by $1/\theta$. In equilibrium, then, the following conditions hold:

Proposition 3 (Output, Wages, and Profits): *Let w denote the wage per unit of raw labor, and let π_i denote the profit earned by the entrepreneur selling variety i . Assume now and for the rest of the paper that $\alpha = (1 - \theta)/\theta$.⁸ The equilibrium with*

⁸This is merely a simplifying assumption that makes profits a linear function of x_i . It can be relaxed with a bit more algebra.

monopolistic competition leads to

$$Y_t = \gamma^{n_t} X_t^\alpha L_t \quad (27)$$

$$w_t = \theta \gamma^{n_t} X_t^\alpha \quad (28)$$

$$\pi_{it} = (1 - \theta) \gamma^{n_t} X_t^\alpha \left(\frac{x_{it}}{X_t} \right) \quad (29)$$

where $X_t \equiv \int_0^1 x_{it} di$ is the mean of the x distribution across entrepreneurs.

According to the proposition, aggregate output is an increasing function of the mean of the idiosyncratic productivity distribution, X . More inequality (a higher η) therefore has a level effect on output in this economy. Notice that this benefits workers as well by increasing the wage for raw labor. The entrepreneur's profits are linear in idiosyncratic productivity, x_i .

We can now determine the value of ψ_t , the parameter that relates entrepreneurial income to x . Entrepreneurs earn the profits from their variety, π_{it} . In the entrepreneur's problem, we previously stated that the entrepreneur's income is $\psi_t x_{it}$, so these two equations define ψ_t as

$$\psi_t = (1 - \theta) \gamma^{n_t} X_t^{\alpha-1}. \quad (30)$$

Finally, we can now determine the overall growth rate of the economy along a balanced growth path. Once the stationary distribution of x has been reached, $X = \frac{x_0}{1-\eta}$ is constant. Since L is also constant over time, the aggregate production function in equation (27) implies that growth in output per person is $\dot{n}_t \log \gamma = \lambda \bar{s} \bar{L} \log \gamma$ if the allocation of research is given by $R/\bar{L} = \bar{s}$. This insight pins down the key endogenous variables of the model, as shown in the next result.

Proposition 4 (Growth and inequality in the \bar{s} case): *If the allocation of research is given exogenously by $R/\bar{L} = \bar{s}$ with $0 < \bar{s} < 1$, then along a balanced growth path, the growth of final output per person, g_y , and the rate of creative destruction are given by*

$$g_y^* = \lambda \bar{s} \bar{L} \log \gamma \quad (31)$$

$$\delta^* = \lambda \bar{s} \bar{L}. \quad (32)$$

Power-law inequality is then given by *Proposition 2* with this value of δ^* .

As a reminder, from *Proposition 2*, power-law inequality is given by

$$\eta^* \equiv 1/\xi^*, \text{ where } \xi^* = -\frac{\tilde{\mu}^*}{\sigma^2} + \sqrt{\left(\frac{\tilde{\mu}^*}{\sigma^2}\right)^2 + \frac{2\bar{\delta}}{\sigma^2}} \quad (33)$$

and $\tilde{\mu}^* \equiv \mu(e^*) - \frac{1}{2}\sigma^2 = \phi(1 - \tau) - \beta(\rho + \delta^* + \bar{\delta}) - \frac{1}{2}\sigma^2$.

4.5. Growth and Inequality: Comparative Statics

In the setup with an exogenously-given allocation of research, the comparative static results are easy to see, and these comparative statics can be divided into those that affect top income inequality only, and those that also affect economic growth.

First, a technological change that increases ϕ will increase top income inequality in the long run. This corresponds to anything that increases the effectiveness of entrepreneurs in building the market for their product. A canonical example of such a change might be the rise in the World Wide Web. For a given amount of effort, the internet allows successful entrepreneurs to grow their markets much more quickly than before, and we now see many examples now of firms that go from being very small to very large quite quickly. Such a change is arguably not specific to any particular economy but rather common to the world. This change can be thought of as contributing to the overall rise in top income inequality throughout most economies, as was documented back in *Figure 2*.

Interestingly, this technological change has no effect on the long-run growth rate of the economy, at least as long as \bar{s} is held fixed. The reason is instructive about how the model works. In the long run, there is a stationary distribution of x . Some varieties are extraordinarily successful, while most are not. Even though an increase in ϕ increases the rate of growth of x for all entrepreneurs, this only serves to widen the stationary distribution. There is a level effect on overall GDP (working through X), but no growth effects. Long-run growth only comes about through the arrival

of new ideas, not through the productivity growth associated with enhancing an existing idea. (And this is true despite the presence of a linear differential equation for x !)

The parameters τ and β also affect top income inequality without affecting growth when \bar{s} is held constant. An increase in τ corresponds to a reduction in the time endowment available to entrepreneurs — an example of such a policy might be the red tape and regulations associated with starting and maintaining a business. With less time available to devote to the productive aspects of running a business, the distribution of x and therefore the distribution of entrepreneurial income is narrowed and top income inequality declines. A similar result obtains if two economies differ with respect to β . An economy where preferences are such that entrepreneurs put more weight on leisure will spend less time building businesses and feature lower top income inequality in the long run.

The two key parameters in the model that affect both growth and top income inequality are λ and \bar{s} , and they work the same way. If researchers are more productive at finding new ideas (a higher λ) or if simply a larger fraction of the labor works in research, the long-run growth rate will be higher — a traditional result in Schumpeterian growth models. Here, however, there will also be an effect on top income inequality. In particular, faster growth means more creative destruction — a higher δ . This means that entrepreneurs have less time to build successful businesses, and this reduces top income inequality in the stationary distribution.

These are the basic comparative statics of top income inequality. Notice that a rise in top income inequality can be the result of either favorable changes in the economy — a new technology like the World Wide Web — or unfavorable changes — like policies that protect existing entrepreneurs from creative destruction.

5. Endogenizing R&D

We now endogenize the allocation of labor to research (s). This allocation is pinned down by the following condition: ex ante, people are indifferent between being a worker and being a researcher.

A worker earns a wage that grows at a constant rate and simply consumes this labor income. The worker's value function is therefore

$$\rho V^w(t) = \log w_t + \frac{dV^w(t)}{dt} \quad (34)$$

A researcher searches for a new idea. If successful, the researcher becomes an entrepreneur. If unsuccessful, we assume the researcher still earns a wage $\bar{m}w$, where \bar{m} is a parameter governing the amount of social insurance for unsuccessful research (this wage payment is financed by a lump-sum tax on entrepreneurs).

The value function for a researcher at time t is

$$\rho V^R(t) = \log(\bar{m}w_t) + \frac{dV^R(t)}{dt} + \lambda (\mathbb{E}[V(x, t)] - V_t^R) + \bar{\delta}_R (V(x_0, t) - V^R(t)). \quad (35)$$

The first two terms on the right-hand side capture the basic consumption of an unsuccessful entrepreneur and the capital gain associated with wage growth. The last two terms capture the successful transition a researcher makes to being an entrepreneur when a new idea is discovered. This can happen in two ways. First, with Poisson flow rate λ the researcher innovates, pushing the research frontier forward by the factor γ , and replaces some randomly-selected existing entrepreneur. Alternatively, the researcher may benefit from the exogenous misallocation process: at rate $\bar{\delta}_R \equiv \bar{\delta}/R$, the researcher replaces a randomly-chosen variety and becomes a new entrepreneur with productivity x_0 .

Finally the indifference condition $V^w(t) = V^R(t)$ determines the allocation of labor as summarized in the following proposition.

Proposition 5 (Allocation of Labor): *In the stationary general equilibrium, the allocation of labor to research, s , is determined by the following two equations:*

$$s^* = 1 - \frac{L^*}{L} \quad (36)$$

$$\log L^* = \log \frac{\theta}{1-\theta} - \beta \log(\Omega - e^*) - \frac{\mu(e^*) - \frac{1}{2}\sigma^2}{\rho + \delta^* + \bar{\delta}} - \frac{(\rho + \delta^* + \bar{\delta}) \log \bar{m} + \lambda \eta^*}{\lambda + \bar{\delta}_R} - \log(1 - \eta^*) \quad (37)$$

Table 1: Equations Characterizing the Stationary General Equilibrium

Drift of log x	$\tilde{\mu}^* = \phi(1 - \tau) - \beta(\rho + \delta^* + \bar{\delta}) - \frac{1}{2}\sigma^2$
Pareto inequality	$\eta^* = 1/\xi^*, \quad \xi^* = -\frac{\tilde{\mu}^*}{\sigma^2} + \sqrt{\left(\frac{\tilde{\mu}^*}{\sigma^2}\right)^2 + \frac{2\bar{\delta}}{\sigma^2}}$
Creative destruction	$\delta^* = \lambda s^* \bar{L}$
Growth	$g^* = \delta^* \log \gamma$
Research share	$s^* = 1 - \frac{L^*}{L}$
Allocation of labor	$\log L^* = \log \frac{\theta}{1-\theta} - \beta \log(1 - \tau - e^*) - \frac{\tilde{\mu}^*}{\rho + \delta^* + \bar{\delta}}$ $- \frac{(\rho + \delta^* + \bar{\delta}) \log \bar{m} + \lambda \eta^*}{\lambda + \delta_R} - \log(1 - \eta^*)$

where e^* and η are given in *Proposition 1* and *Proposition 2*, respectively.

The key equations that describe the stationary general equilibrium are then shown in Table 1. As we've discussed earlier the endogenous creative destruction has a positive impact on growth and reduces inequality. And the number of researchers determine the intensity of creative destruction. However, it is not easy to discuss the comparative statics as there's no closed-form algebraic solution for s^* . Instead, in the next section we will show numerically how each parameter affects growth and inequality in the stationary general equilibrium.

The model features transition dynamics away from steady state. The key reason for this is that the initial distribution of x can be anything, and the evolution of this distribution affects aggregate variables through X_t and δ_t . All numerical results that follow will therefore be of the form of comparing steady states. It would be valuable

to solve the transition dynamics of this model, but this brings in many well-known complications (e.g. Krusell and Smith (1998)) and is beyond the scope of the present paper.

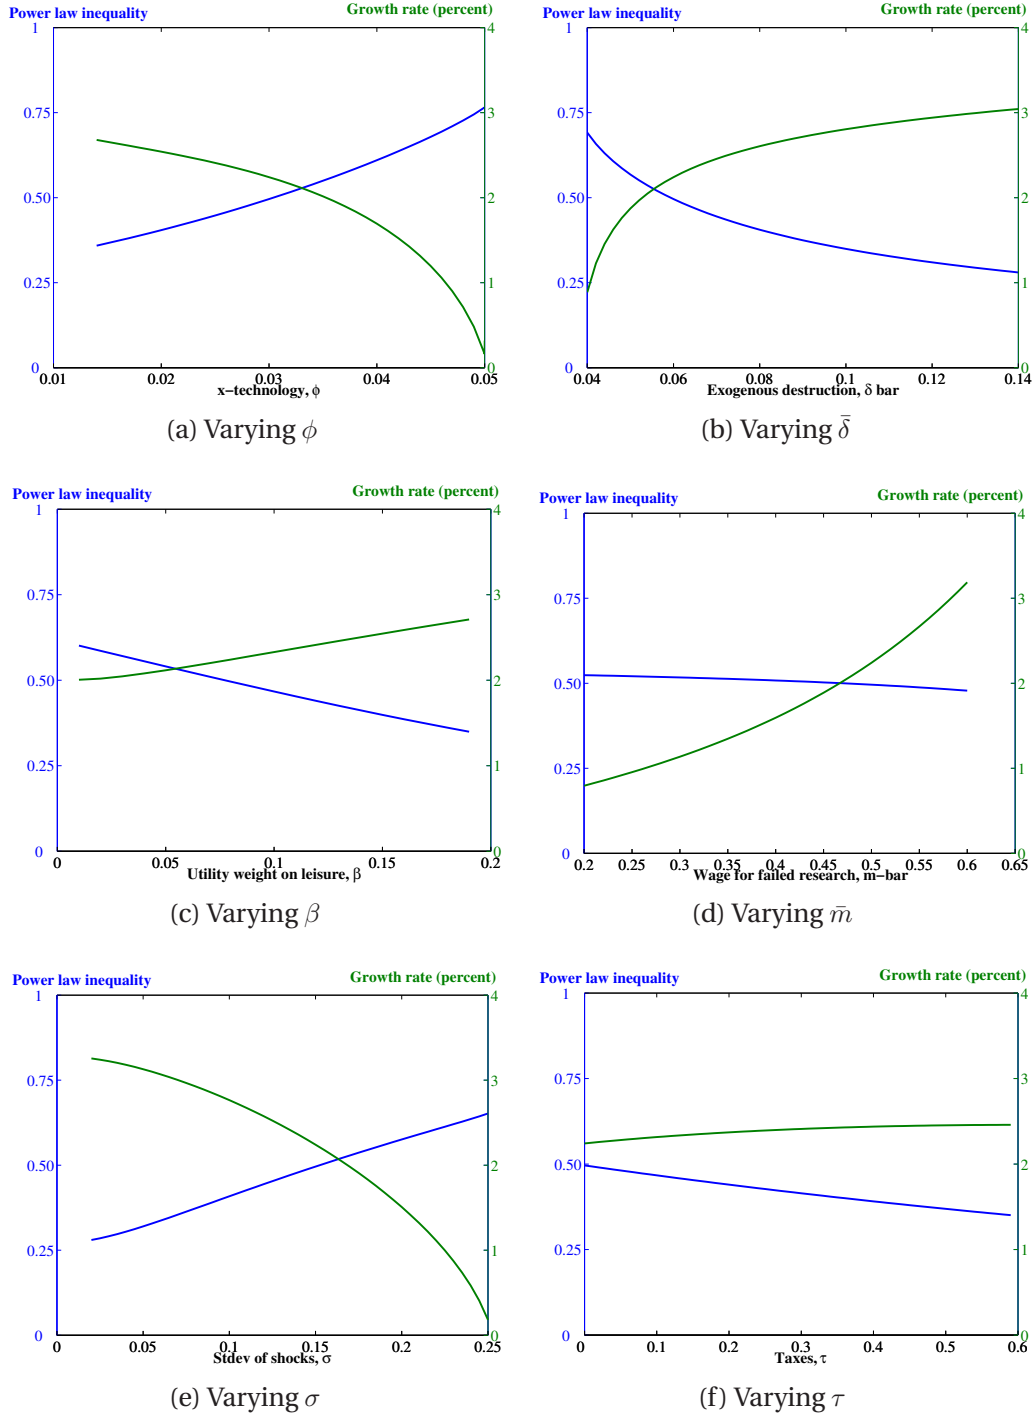
5.1. Comparative Statics

Figure 7 shows numerical results when $s \equiv R/\bar{L}$ is endogenously determined. The effects on Pareto inequality are similar to those from the exogenous \bar{s} case. Now, however, we can also study the effects on economic growth, and those are quite interesting. For example, consider the effect of an increase in the technology parameter ϕ , shown in Figure 7a: an increase in ϕ raises Pareto inequality, as discussed earlier, but — perhaps surprisingly — causes a decline in the growth rate of GDP per person. Similar results occur throughout Figure 7: parameter changes that increase Pareto inequality tend to reduce economic growth.

To understand this result, recall that the growth rate of the economy is determined by the fraction of people who decide to enter the research process, prospecting for the possibility of become successful entrepreneurs. On the one hand, an increase in ϕ makes it easier for entrepreneurs to build up their market share, which tends to make research more attractive. However, from the standpoint of a researcher who has not yet discovered a new idea, another effect dominates. The positive technological improvement from a rising ϕ raises average wages in the economy, both for workers and for unsuccessful researchers. However, it also increases the inequality among successful researchers, making the research process itself more risky. Moreover, because this risk is Pareto in nature, with thick tails, it is a *large* increase in risk. Our researchers are risk-averse individuals with log utility, and the result of this risk aversion is that a rise in ϕ results in a smaller fraction of people becoming researchers, which lowers the long-run growth rate in this endogenous growth model.

One can, of course, imagine writing down the model in a different way. For example, if research is undertaken by risk-neutral firms, then this effect would not be present. Ultimately, this question must be decided by empirical work. Our model,

Figure 7: Numerical Examples: Endogenous s



Note: The baseline parameter values in these examples are $\rho = .03$, $\bar{L} = 7.25$, $\theta = 2/3$, $\gamma = e^{2/5}$, $\lambda = .07$, $\phi = .03$, $\beta = .08$, $\sigma = 0.15$, $\bar{s} = .0985$, and $\tau = 0$. These values will be discussed in more detail in Section 7.

however, makes it clear that this additional force is present, so that increases in Pareto inequality that result from positive technological changes need not increase the rate of growth.

6. Micro Evidence

To what extent is our model consistent with empirical evidence? There are several ways to answer this question. Some relate to previous empirical work on income dynamics. Some relate to evidence we provide using public use micro data from the U.S. Internal Revenue Service. And some relate to evidence that we hope others with better access to administrative data can conceivably provide in the future.

The first point to make is that the basic stochastic process for incomes assumed in our model — a geometric random walk with positive drift — is the canonical data generating process estimated in an extensive empirical literature on income dynamics. Meghir and Pistaferri (2011) surveys this literature, highlighting prominent examples such as MaCurdy (1982), Abowd and Card (1989), Topel and Ward (1992), Baker and Solon (2003), and Meghir and Pistaferri (2004). There are of course exceptions and some papers prefer alternative specifications, with the main one being the “heterogeneous income profiles” which allow for individual-specific means and returns to experience and often find a persistence parameter less than one; for example, see Lillard and Willis (1978), Baker (1997), and Guvenen (2007, 2009). While debate continues within this literature, it is fair to say that a fundamental benchmark is that the log of income features a random walk component. In that sense, the basic data generating process we assume in this paper has solid micro-econometric foundations.

With unlimited access to micro data, our model makes some clear predictions that could be tested. In particular, one could estimate the stochastic process for incomes around the top of the income distribution. In addition to the geometric random walk with drift, one could estimate the creative destruction parameters — to what extent do high income earners see their incomes drop by a large amount in a short time? Guvenen, Ozkan and Song (2013) provide evidence for precisely this

effect, stating “[I]ndividuals in higher earnings percentiles face persistent shocks that are more negatively skewed than those faced by individuals that are ranked lower, consistent with the idea that the higher an individual’s earnings are, the more room he has to fall” (p. 20).

Beyond estimating this stochastic process, one could also see how the process differs before and after 1980 in the United States and how it differs between the United States and other countries. For example, one would expect the positive drift of the random walk to be higher for top incomes after 1980 than before. And one would expect this drift to be higher in the United States in the 2000s than in France; there could also be differences in the creative destruction parameters between countries and over time that could be estimated.

Access to panel micro data on top income earners in the U.S. over time and in other countries is typically very hard to gain. However, we can make some progress with the U.S. Internal Revenue Service public use tax model panel files created by the Statistics of Income Division from 1979 to 1990, hosted by the NBER.⁹

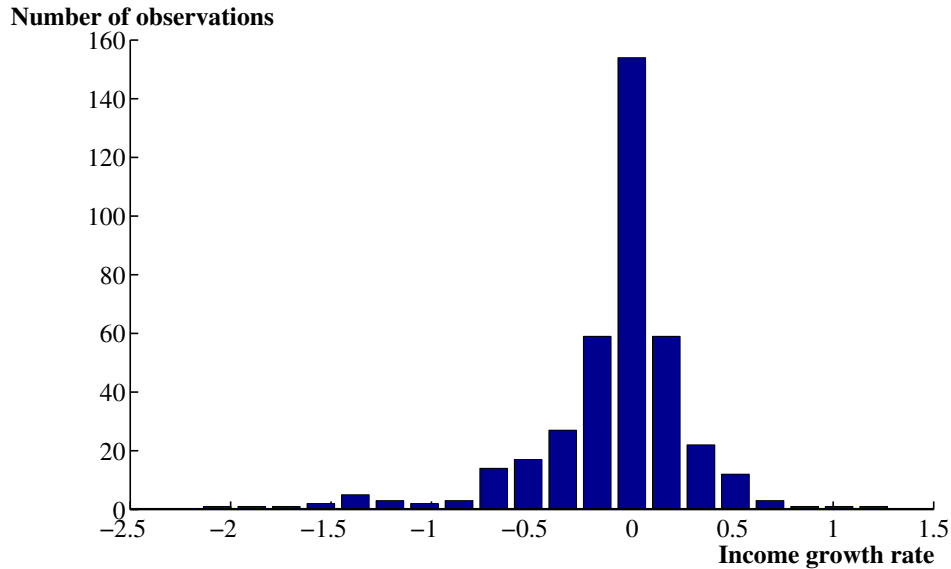
6.1. Estimating the Top Income Process for the United States

Motivated by our theory, consider the following empirical model for top incomes:

$$\log y_t = \begin{cases} \log y_{t-1} + \tilde{\mu}_t + \sigma_t \epsilon_t & \text{with probability } 1 - \delta_t^e \\ \nu_t - \log y_{t-1} & \text{with probability } \delta_t^e \end{cases} \quad (38)$$

where $\epsilon_t \sim N(0, 1)$ and $\nu_t \sim N(\mu_{\delta,t}, \sigma_{\delta,t}^2)$ are independent random variables. That is, with probability $1 - \delta_t^e$, the top income follows a geometric random walk with drift $\tilde{\mu}_t$. Alternatively, with probability δ_t^e , the top income earner reverts to being a “normal” worker and draws a new income from a lognormal distribution. (The “e” superscript in δ^e denotes the “empirical” model; we discuss how δ^e relates to the creative destruction parameter below.) According to this model, if the individual remains a top earner, then the distribution of income growth rates is normal. How-

⁹See <http://www.nber.org/taxsim-notes.html>.

Figure 8: The Distribution of Top Income Growth Rates between 1988 and 1989


Note: This figure shows a histogram of the change in log wage and salary income for tax units consisting of married taxpayers filing jointly that are in the Top 5 percent of tax units in 1988 and for which we also observe income in 1989.

ever, the destruction shock results in a potentially large downward shift in incomes, causing the growth rate distribution to be left-skewed. As noted above, Guvenen, Ozkan and Song (2013) provide evidence of this skewness.

We can see the same thing in our data from the IRS public use panel. Restricting our sample to tax units that involve married taxpayers filing jointly in two consecutive years, an example of this growth rate distribution for tax units starting in the 95th percentile and above in 1988 is shown in Figure 8. Quantitatively, the left-skewness of the distribution of growth rates helps us to identify δ_t^e , while the center part of the distribution helps us to identify $\tilde{\mu}$ and σ .

Our data and estimation are discussed in more detail in Appendix A. In brief, the empirical model in equation (38) can be viewed as a 2-state mixture model of normal random variables for each cross-section at a point in time. Therefore, for each year between 1980 and 1990, we use the cross section of growth rates from a two-year panel of tax returns to estimate the 5 parameters of this mixture model using

Figure 9: Estimates of μ from U.S. IRS Public Use Panel



Note: This figure shows the estimates of μ when using the Top 5 (blue/grey) and Top 10 (green/tan) percentiles of wage/salary income from the IRS public use panel data. 95 percent confidence intervals obtained with 1000 bootstrap iterations are shown.

the “EM” algorithm. We compute 95 percent confidence intervals using a 1000-draw bootstrap. Finally, we do this for two cut-offs of “top” incomes, corresponding to incomes in the Top 5 percent and incomes in the Top 10 percent, respectively.


6.2. Results

Figures 9 through 11 shows the point estimates and 95 percent confidence intervals for various key parameters of the model. In each figure, the blue line and the grey bands correspond to estimates using the Top 5 percent of incomes, while the green line and the tan bands correspond to estimates from the Top 10 percent.

Overall, the key takeaways from this estimation are


- The drift parameter $\mu \equiv \tilde{\mu} + \frac{1}{2}\sigma^2$ shows clear evidence of increasing during the 1980s, especially in the first half of the decade. For example, using the Top 5 percent data, the value of μ rises from -0.05 in 1980 to 0.036 in 1985 and then 0.00 in 1990.

Figure 10: Estimates of σ from U.S. IRS Public Use Panel



Note: Note: This figure shows the estimates of the standard deviation parameter for the innovation to the Brownian motion, σ , when using the Top 5 (blue/grey) and Top 10 (green/tan) percentiles of wage/salary income from the IRS public use panel data. 95 percent confidence intervals obtained with 1000 bootstrap iterations are shown.

Figure 11: Estimates of δ^e from U.S. IRS Public Use Panel



Note: This figure shows the estimates of the overall destruction rate $\delta^e = \delta + \bar{\delta}$ when using the Top 5 (blue/grey) and Top 10 (green/tan) percentiles of wage/salary income from the IRS public use panel data. 95 percent confidence intervals obtained with 1000 bootstrap iterations are shown.


- The negative point estimates for μ in the early 1980s may reflect the large recession during those years.
- The confidence intervals can be wide. Partly this reflects the changing sample sizes. For example, for the Top 5 percent estimation, the sample size in 1980 is 943 observations, falling to a low of 183 observations in 1983, and then ending with 359 observations in 1990.
- As shown in Figure 10, the standard deviation of the innovation, σ , shows some evidence of increasing, but the confidence intervals are wide. For example, using the Top 5 percent data, the value is 0.14 in 1980 and rises to 0.20 in 1990.
- The estimates of δ_t^e are shown in Figure 11. There is no large trend in the values, and the mean value is 0.108, albeit with a wide confidence interval.

A useful way to summarize the movements in these parameters is to use our theoretical model to compute their implications for long-run top income inequality. More specifically, for each year, we use the parameter estimates from that year to compute the value of η implied by equation (33). There are two caveats to this calculation. First, the value of η is actually the value that would apply in the long run, not immediately, as this is a steady-state result. However, bearing this caveat in mind, the calculation as a summary still seems useful. Second, given values for μ and σ , the calculation requires a value of $\bar{\delta}$. We obtain this by noting that $\bar{\delta} = \delta^e - \delta$. That is, we simply need to subtract the endogenous rate of creative destruction from the point estimate. For illustration, we assume a constant value of $\delta = .05$.

Figure 12 then shows the estimates of η implied by this calculation, together with bootstrapped confidence intervals. The striking result shown there is that the empirical estimates imply a substantial rise in top income inequality. For example, using the Top 5 percent of incomes, η rises from 0.14 to 0.56.¹⁰ One way to get a simple feel for these numbers is to return to the simple model at the start of the paper. There, recall, $\eta = \mu/\delta$. So a value of μ that rises from 0.01 to 0.05 with a value of δ around 0.10 is roughly consistent with these facts. Alternatively, the rise in η

¹⁰The confidence intervals are 0.10 to 0.20 in 1980 and 0.33 to 1.20 in 1990.

Figure 12: Implied Long-Run Pareto Inequality from the IRS Public Use Panel



Note: This figure shows the estimates of the implied steady-state Pareto inequality, η , based on the parameter values estimated for the stochastic income process. The calculation assumes the endogenous creative destruction rate δ equals 0.05. Estimates correspond to using the Top 5 (blue/grey) and Top 10 (green/tan) percentiles of wage/salary income from the IRS public use panel data. 95 percent confidence intervals obtained with 1000 bootstrap iterations are shown.

for the U.S. data is from around 0.4 to 0.6, which would correspond to μ rising from 0.04 to 0.06, for example.

The evidence from the IRS public use panel of tax returns is obviously limited in many ways. Nevertheless, the evidence from estimation using these data is revealing. It suggests that the basic model used in this paper is capable of explaining a substantial rise in U.S. top income inequality. Access to the restricted administrative tax data from the IRS and from tax authorities in other countries could be used to make additional progress and is a clear priority for future research.

7. Numerical Examples

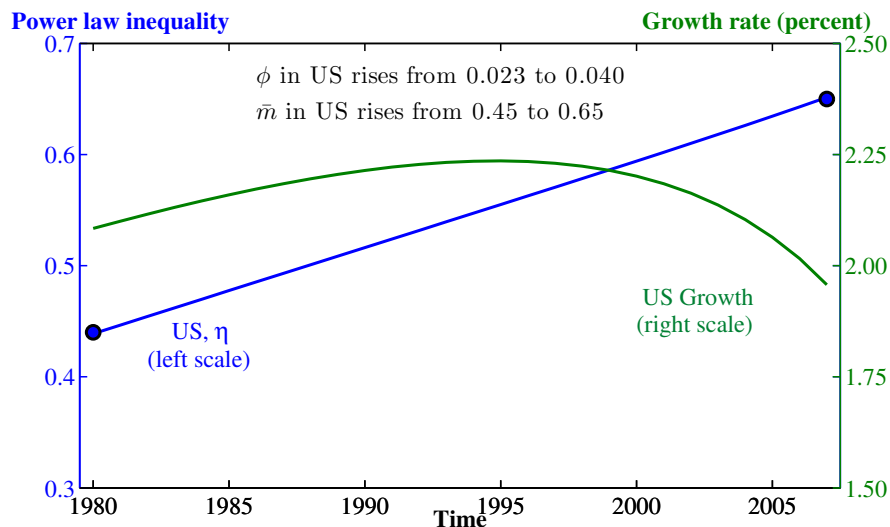
We now provide three numerical examples to illustrate the ability of the Schumpeterian model to match the levels and changes of top income inequality in the United States and France. Where possible, our baseline parameter values are chosen to be consistent with the empirical estimates from the IRS panel data from the previous section. For example, we assume that σ for the U.S. examples rises from 0.15 to 0.20, broadly consistent with the evidence in Figure 10. We assume $\bar{\delta} = 0.06$ and $\gamma = e^{2/5}$, so that $\delta + \bar{\delta} = 0.11$ when the economy's growth rate is 2 percent, broadly consistent with the evidence in Figure 11. Finally, the values of ϕ we consider deliver a values of μ ranging from 0.0078 to 0.0195, broadly consistent with the evidence in Figure 9.

These remain examples, however, for three main reasons. First, the U.S. panel evidence we have on μ , σ , and δ^e applies only to the years 1980 to 1990, whereas we attempt in this section to shed light on inequality movements all the way until 2007. Second, the “reduced-form” empirical evidence is insufficient to identify the underlying structural parameters of the model. As one simple example, movements in both ϕ and τ can deliver changes in μ over time, as we show below. Finally, for reasons discussed above, our numerical exercises do not consider transition dynamics and instead report a sequence of steady states. Nevertheless, the aim is to show that the Schumpeterian model can deliver changes in top income inequality consistent with both the overall inequality data and with the underlying micro data on top incomes.

7.1. Matching U.S. Inequality

Figure 13 shows a numerical example that illustrates one way in which the model can match the time series behavior of top income inequality in the United States. We start with a set of baseline parameters, $\rho = 0.03$, $\bar{L} = 7.25$, $\tau = 0$, $\theta = 2/3$, $\beta = 0.08$, $\gamma = e^{2/5}$, $\lambda = 0.07$, $\sigma = 0.15$, $\bar{\delta} = 0.06$, $\phi = .0233$, $\bar{m} = 0.45$, which match U.S. Pareto inequality in 1980. Then we assume that ϕ , the technology parameter converting entrepreneurial effort into growth in a variety's productivity, rises over time. We calibrate the change in ϕ to match the rise in the U.S. top inequality between

Figure 13: Numerical Example: Matching U.S. Inequality



Note: Baseline parameter values are given in the text. Over time in this simulation, ϕ rises linearly from 0.0233 to 0.0395, σ rises linearly from 0.15 to 0.20, and \bar{m} rises linearly from 0.45 to 0.65. The rise in \bar{m} is required to keep growth rates from falling sharply with the rise in ϕ .

1980 and 2007.


Although the rise in ϕ matches the changes in inequality, it has a negative impact on growth. An increase in ϕ leads to a rise in entrepreneurial effort and to a rise in inequality. Because prospective researchers are risk averse, more inequality among entrepreneurs reduces the number of researchers, thereby limiting overall growth. To offset this, we vary another parameter in the U.S. in Figure 13. In particular, we assume that the subsidy to research, \bar{m} , rises from 0.45 to 0.65, which stabilizes U.S. growth on average.¹¹

7.2. Matching Inequality in France

Next, we illustrate the ability of the model to simultaneously match the patterns of top income inequality in both the U.S. and France. For the United States, we assume

¹¹There are other ways to achieve this same end. For example, one could let \bar{L} increase instead, possibly reflecting globalization.

Figure 14: Numerical Example: U.S. and France




Note: The baseline parameter values described in the text are used; for example, the U.S. parameter values are the same as in Figure 13. In France, we assume ϕ rises in exactly the same way. To offset the large increase in inequality that would be implied, we assume $\bar{\delta}$ rises from 0.07 to 0.084.

the same parameter values just described. For France, we assume the same rise in ϕ applies, for example capturing the fact that the World Wide Web is a worldwide phenomenon. We know from our U.S. example that this would produce a large rise in top inequality, other things equal, so some other change is needed to offset the rise in France. For this example, we consider a rise in $\bar{\delta}$. In particular, a linear increase in $\bar{\delta}$ from 0.07 to 0.084 (versus a constant value of 0.06 for the United States) allows us to match the evidence on French inequality, as shown in Figure 14.

7.3. Taxes

As already discussed, these exercises are simply numerical examples, and there are other ways to match the change in top inequality in the model in ways that are broadly consistent with the data. In this example, we consider changes in the tax parameter τ . A literal interpretation of this parameter is that it is a tax on the time endowment of entrepreneurs — e.g. lost time due to “red tape.” It is tempting to

Figure 15: Numerical Example: Taxes and Inequality



Note: The figure shows the results of linear declines in the tax parameter τ on top income inequality and on economy-wide average income growth.

wonder about the effects of marginal labor income tax rates. However, because consumption enters in log form in the utility function (to obtain analytic tractability), the income and substitution effects from a labor income tax cancel exactly and a labor income tax leaves equilibrium effort — and therefore top income inequality — unchanged in this setup. The time tax τ is suggestive of what labor income taxes might imply in a richer framework.¹²

To what extent can changes in τ account for the differential patterns of top income inequality that we see in the United States and France? This question is considered in Figure 15. The baseline parameter values used in this example are identical to those used in Figure 14, with two exceptions. First, we raise the (constant) level of ϕ to 0.03 to match the level of U.S. inequality in 1980, when the initial tax rate is positive. Second, we consider linear declines in τ for both the U.S. and France. Having τ in the United States fall from 39 percent to 0 percent nearly generates the

¹²Kim (2013) considers labor income taxes in a model in which goods are used to accumulate human capital, so that labor income taxes can affect top income inequality, and finds that tax changes can account for some a modest portion of the changes in inequality over time and across countries.

U.S. inequality numbers. And a more modest decline from 55 percent to 30 percent can generate the French inequality facts. The U.S. tax change by itself, however, leads to a large *decline* in overall income growth, suggesting that something else must also be going on. As before, we offset this growth effect with a rise in \bar{m} , in this case from 0.47 to 0.64.

8. Conclusion

A model in which entrepreneurs expend effort to enhance the productivity of their existing ideas while researchers seek new ideas to replace incumbents in a process of creative destruction generates a Pareto distribution for labor incomes at the top of the income distribution. Moreover, it suggests economic forces that can give rise to changes in top income inequality. Forces that increase the effort of existing entrepreneurs in improving their products — or that increase the efficiency of their effort — can increase top inequality. Forces that enhance creative destruction can decrease top inequality. Changes in these forces over time or differences in their strength across countries can potentially explain the patterns of top income inequality that we see in the data.

A Appendix: Estimating the U.S. Stochastic Income Process

A1. Data

The data we use are the U.S. Internal Revenue Service public use tax model panel files created by the Statistics of Income Division from 1979 to 1990, hosted by the NBER. See <http://www.nber.org/taxsim-notes.html>. We restrict our sample to tax units that involved married taxpayers filing jointly in two consecutive years and convert nominal values to 2012 constant dollars using the consumer price index. Our income measure is field 11 of the data, corresponding to wages and salaries. For each pair of consecutive years, we record initial income and the change in log income for each tax unit. This constitutes our main data used in the estimation.

Two sets of estimates are made, based on different cutoffs for the meaning of “top incomes.” In the main case, we use only tax units in the Top 5 percent of our income measure. As a robustness check, we also consider a Top 10 percent cutoff.

A2. Estimation

For each set of consecutive years, we use our cross-section of income levels and growth rates to estimate the stochastic process in equation (38). Hence, our five parameters $\tilde{\mu}_t$, σ_t , δ_t^e , $\mu_{\delta,t}$, and $\sigma_{\delta,t}$ are indexed by time. The empirical model is a 2-state mixture model of normal random variables for each cross-section. Therefore, for each year between 1980 and 1990, we use the cross-section of initial incomes and growth rates to estimate the parameters using the “EM” algorithm, as described by Hamilton (1994) and especially Elkan (2010). We compute 95 percent confidence intervals using 1000-draw bootstraps in each year. Complete point estimates and sample sizes are shown in Tables 2 and 3.

B Appendix: Proofs of the Propositions

To be provided...

Table 2: Parameter Estimates (Top 5 percentile cutoff)

Year	Basic parameters					Implied values		N
	$\tilde{\mu}$	σ	δ^e	μ_δ	σ_δ	μ	η	
1980	-0.0603	0.1415	0.1263	10.425	1.228	-0.0503	0.1408	943
1981	-0.0489	0.1776	0.0685	10.353	1.213	-0.0332	0.2905	932
1982	-0.0215	0.1842	0.1153	10.102	1.095	-0.0045	0.3711	197
1983	0.0093	0.1377	0.1499	9.732	1.388	0.0188	0.3583	183
1984	-0.0142	0.1235	0.1178	10.181	1.431	-0.0066	0.2465	194
1985	0.0291	0.1206	0.1787	10.580	0.707	0.0363	0.3763	198
1986	-0.0030	0.1608	0.1016	10.020	1.979	0.0100	0.4724	205
1987	-0.0209	0.2261	0.0927	10.305	1.186	0.0047	0.5669	198
1988	-0.0118	0.2414	0.0879	10.113	1.238	0.0174	0.7350	381
1989	-0.0276	0.2447	0.0763	9.343	1.174	0.0023	0.6643	388
1990	-0.0236	0.2050	0.0750	10.159	2.323	-0.0026	0.5592	359

Note: See text for details. For the implied parameter values, $\mu \equiv \tilde{\mu} + \frac{1}{2}\sigma^2$ and η is computed from the other parameter values using equation (33) with $\bar{\delta} = \delta^e - \delta$ assuming $\delta = .05$.

Table 3: Parameter Estimates (Top 10 percentile cutoff)

Year	Basic parameters					Implied values		N
	$\tilde{\mu}$	σ	δ^e	μ_δ	σ_δ	μ	η	
1980	-0.0540	0.1248	0.1263	10.433	1.036	-0.0462	0.1230	1874
1981	-0.0357	0.1471	0.0796	10.450	1.209	-0.0249	0.2509	1844
1982	-0.0253	0.1647	0.0888	10.204	1.112	-0.0117	0.3493	388
1983	0.0132	0.1347	0.1261	10.135	1.244	0.0223	0.4431	371
1984	-0.0015	0.1275	0.0940	10.213	1.641	0.0066	0.4128	385
1985	0.0162	0.1137	0.1524	10.556	0.813	0.0226	0.3424	403
1986	0.0045	0.1316	0.1022	10.196	1.700	0.0132	0.4525	411
1987	-0.0135	0.1940	0.0776	10.408	1.052	0.0053	0.6160	396
1988	-0.0030	0.1629	0.1185	10.379	1.439	0.0103	0.4188	776
1989	-0.0105	0.1633	0.1136	9.874	1.221	0.0029	0.3827	784
1990	-0.0180	0.1660	0.0916	10.325	1.867	-0.0042	0.3982	737

Note: See notes to Table 2.

References

- Abowd, John M. and David Card, "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, March 1989, 57 (2), 411–45.
- Acemoglu, Daron and David Autor, "Skills, Tasks and Technologies: Implications for Employment and Earnings," in O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics*, Vol. 4, Elsevier, June 2011, chapter 12, pp. 1043–1171.
- Aghion, Philippe and Peter Howitt, "A Model of Growth through Creative Destruction," *Econometrica*, March 1992, 60 (2), 323–351.
- Aluation.wordpress.com, "Fractals, fractiles and inequality," Technical Report August 2011.
- Aoki, Shuhei and Makoto Nirei, "Pareto Distributions and the Evolution of Top Incomes in the U.S," MPRA Paper 47967, University Library of Munich, Germany July 2013.
- Atkinson, Anthony B., Thomas Piketty, and Emmanuel Saez, "Top Incomes in the Long Run of History," *Journal of Economic Literature*, 2011, 49 (1), 3–71.
- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney, "The Polarization of the U.S. Labor Market," *American Economic Review*, May 2006, 96 (2), 189–194.
- Baker, Michael, "Growth-Rate Heterogeneity and the Covariance Structure of Life-Cycle Earnings," *Journal of Labor Economics*, April 1997, 15 (2), 338–75.
- and Gary Solon, "Earnings Dynamics and Inequality among Canadian Men, 1976-1992: Evidence from Longitudinal Income Tax Records," *Journal of Labor Economics*, April 2003, 21 (2), 267–288.
- Bakija, Jon, Adam Cole, and Bradley Heim, "Jobs and Income Growth of Top Earners and the Causes of Changing Income Inequality: Evidence from U.S. Tax Return Data," working paper, Indiana University November 2010.
- Bell, Brian and John Van Reenen, "Bankers' pay and extreme wage inequality in the UK," Open Access publications from London School of Economics and Political Science <http://eprints.lse.ac.uk/>, London School of Economics and Political Science 2010.
- Benabou, Roland and Jean Tirole, "Bonus Culture: Competitive Pay, Screening, and Multi-tasking," Working Paper 18936, National Bureau of Economic Research April 2013.

- Benhabib, Jess, Alberto Bisin, and Shenghao Zhu, “The Distribution of Wealth and Fiscal Policy in Economies With Finitely Lived Agents,” *Econometrica*, 01 2011, 79 (1), 123–157.
- Elkan, Charles, “Mixture Models,” March 2010. UCSD manuscript.
- Gabaix, Xavier, “Zipf’s Law for Cities: An Explanation,” *Quarterly Journal of Economics*, August 1999, 114 (3), 739–767.
- , “Power Laws in Economics and Finance,” *Annual Review of Economics*, 2009, 1 (1), 255–294.
- Gordon, Robert J. and Ian Dew-Becker, “Controversies about the Rise of American Inequality: A Survey,” NBER Working Papers 13982, National Bureau of Economic Research, Inc May 2008.
- Guvenen, Fatih, “Learning Your Earning: Are Labor Income Shocks Really Very Persistent?,” *American Economic Review*, June 2007, 97 (3), 687–712.
- , “An Empirical Investigation of Labor Income Processes,” *Review of Economic Dynamics*, January 2009, 12 (1), 58–79.
- , Serdar Ozkan, and Jae Song, “The Nature of Countercyclical Income Risk,” Technical Report, Board of Governors of the Federal Reserve System (U.S.) September 2013. working paper.
- Hamilton, James D., *Time Series Analysis*, Princeton, NJ: Princeton University Press, 1994.
- Haskel, Jonathan, Robert Z. Lawrence, Edward E. Leamer, and Matthew J. Slaughter, “Globalization and U.S. Wages: Modifying Classic Theory to Explain Recent Facts,” *Journal of Economic Perspectives*, Spring 2012, 26 (2), 119–40.
- Kaplan, Steven N. and Joshua Rauh, “Wall Street and Main Street: What Contributes to the Rise in the Highest Incomes?,” *Review of Financial Studies*, 2010, 23 (3), 1004–1050.
- Katz, Lawrence F. and David H. Autor, “Changes in the wage structure and earnings inequality,” in O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics*, Vol. 3 of *Handbook of Labor Economics*, Elsevier, June 1999, chapter 26, pp. 1463–1555.
- Kim, Jihee, “The Effect of the Top Marginal Tax Rate on Top Income Inequality,” 2013. KAIST, unpublished paper.

- Kortum, Samuel S., "Research, Patenting, and Technological Change," *Econometrica*, 1997, 65 (6), 1389–1419.
- Krusell, Per and Anthony A. Smith, "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, October 1998, 106 (5), 867–896.
- Lillard, Lee A. and Robert J. Willis, "Dynamic Aspects of Earning Mobility," *Econometrica*, September 1978, 46 (5), 985–1012.
- Lucas, Robert E. and Benjamin Moll, "Knowledge Growth and the Allocation of Time," *Journal of Political Economy*, February 2014, 122 (1), 1–51.
- Luttmer, Erzo G.J., "Selection, Growth, and the Size Distribution of Firms," *Quarterly Journal of Economics*, 08 2007, 122 (3), 1103–1144.
- , "Models of Growth and Firm Heterogeneity," *Annual Review Economics*, 2010, 2 (1), 547–576.
- MaCurdy, Thomas E., "The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis," *Journal of Econometrics*, 1982, 18 (1), 83–114.
- Malevergne, Y., A. Saichev, and D. Sornette, "Zipf's law and maximum sustainable growth," Quantitative Finance Papers 1012.0199, arXiv.org December 2010.
- Meghir, Costas and Luigi Pistaferri, "Income Variance Dynamics and Heterogeneity," *Econometrica*, 01 2004, 72 (1), 1–32.
- and —, "Earnings, Consumption and Life Cycle Choices," in O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics, Volume 4*, Elsevier, 2011, chapter 9, pp. 773–854.
- Mitzenmacher, Michael, "A Brief History of Generative Models for Power Law and Lognormal Distributions," *Internet Mathematics*, 2004, 1 (2).
- Moll, Benjamin, "Inequality and Financial Development: A Power-Law Kuznets Curve," 2012. Princeton University working paper.
- Nirei, Makoto, "Pareto Distributions in Economic Growth Models," IIR Working Paper 09-05, Institute of Innovation Research, Hitotsubashi University July 2009.
- Pareto, Vilfredo, *Cours d'Economie Politique*, Geneva: Droz, 1896.

- Perla, Jesse and Christopher Tonetti, "Equilibrium Imitation and Growth," *Journal of Political Economy*, February 2014, 122 (1), 52–76.
- Philippon, Thomas and Ariell Reshef, "Wages and Human Capital in the U.S. Financial Industry: 1909-2006," NBER Working Papers 14644, National Bureau of Economic Research, Inc January 2009.
- Piketty, Thomas and Emmanuel Saez, "Income Inequality In The United States, 1913–1998," *Quarterly Journal of Economics*, February 2003, 118 (1), 1–39.
- and — , "A Theory of Optimal Capital Taxation," NBER Working Papers 17989, National Bureau of Economic Research, Inc April 2012.
- and Gabriel Zucman, "Wealth and Inheritance in the Long Run," April 2014. forthcoming in the Handbook of Income Distribution.
- , Emmanuel Saez, and Stefanie Stantcheva, "Optimal Taxation of Top Labor Incomes: A Tale of Three Elasticities," *American Economic Journal: Economic Policy*, February 2014, 6 (1), 230–71.
- Reed, William J., "The Pareto, Zipf and other power laws," *Economics Letters*, 2001, 74 (1), 15–19.
- Rosen, Sherwin, "The Economics of Superstars," *American Economic Review*, December 1981, 71 (5), 845–58.
- Rothschild, Casey and Florian Scheuer, "Optimal Taxation with Rent-Seeking," NBER Working Papers 17035, National Bureau of Economic Research, Inc May 2011.
- Saez, Emmanuel, "Using Elasticities to Derive Optimal Tax Rates," *Review of Economic Studies*, 2001, 68, 205–229.
- Steindl, Joseph, *Random Processes and the Growth of Firms*, New York: Hafner, 1965.
- Topel, Robert H and Michael P Ward, "Job Mobility and the Careers of Young Men," *The Quarterly Journal of Economics*, May 1992, 107 (2), 439–79.