

The Economic Determinants of Top Income Inequality

Charles I. Jones

What I've been spending my time on for the last couple years is something that I think was very much in the Gary Becker tradition of economics, which is if you want to understand something, then an important part of that understanding is being able to write down models of rational economic agents, where the phenomenon that you're studying emerges in the model. You ought to be able to build an economy in which the phenomenon happens. So I've been thinking about the kind of assumptions we need to make in order to build economies where top income inequality rises dramatically in the United States after being flat for thirty years before, but stays pretty flat in France.

Figure 3.1 shows a picture that we're all familiar with when talking about income inequality: the income share of the top 0.1 percent starting from a common value in the US and France from 1950 to 1980 of around 2.5 percent or 3 percent, rising quite dramatically in the US up to 8 or 9 percent. And in contrast, the top income share stays quite low in France.

Figure 3.2 shows you the top 1 percent share in the early eighties on the horizontal axis, and the top 1 percent share in the mid-2000s on the vertical axis, together with a 45-degree line. And this chart uses the data that Emmanuel Saez and Thomas Piketty put together. What you can see in this graph—which I think is quite remarkable—is that the rise in income inequality is not just a US phenomenon. It's happening everywhere in the world for which we

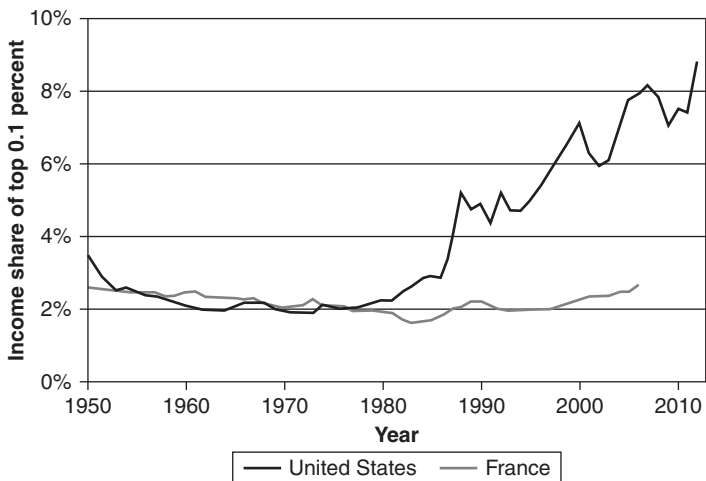


FIGURE 3.1. Top income inequality in the United States and France

Source: Charles I. Jones and Jihee Kim, "A Schumpeterian Model of Top Income Inequality," *NBER Paper No. 20637*, October 2014, figure 1

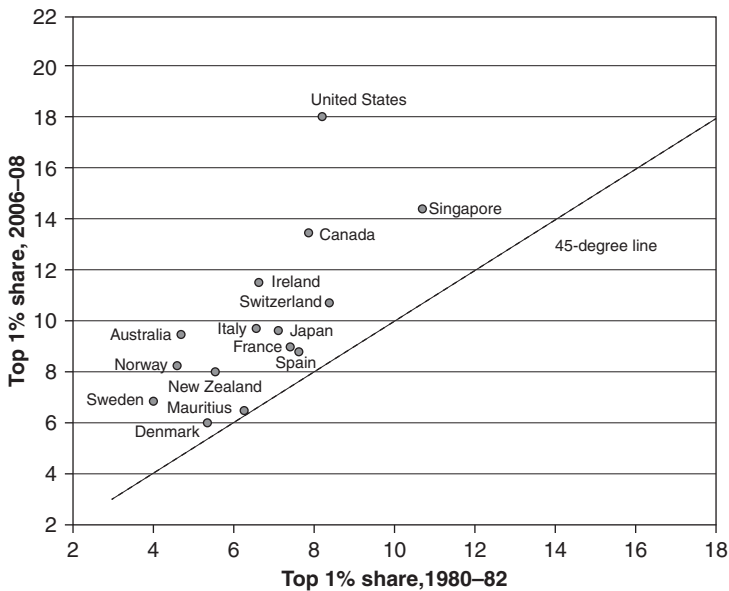


FIGURE 3.2. Top income inequality around the world

Source: Charles I. Jones and Jihee Kim, "A Schumpeterian Model of Top Income Inequality," *NBER Paper No. 20637*, October 2014, figure 3

have data over these two periods. Yet there's heterogeneity in the extent of the rise. You see an increase in top income and inequality even in Sweden, even in France, and it's just the magnitude of the inequality differences that stand out. But inequality is rising everywhere, which suggests to me there is something broad-based going on, and that appealing to narrow features of the US economy, for example, is likely to give you an incomplete explanation of the phenomenon.

Since Pareto created his distribution more than a century ago, it's been appreciated that the top of the income distribution looks like a Pareto distribution. If you want to think about a model of top income inequality, you're inevitably drawn to consider Pareto distributions. What figure 3.3 shows is that top incomes in the US are consistent with the Pareto distribution. Suppose we pick a cutoff, let's say \$500,000 a year. Now consider all the taxpayers who make more than \$500,000 a year, and figure out what the average of their incomes is. Look at that as a ratio to the \$500,000

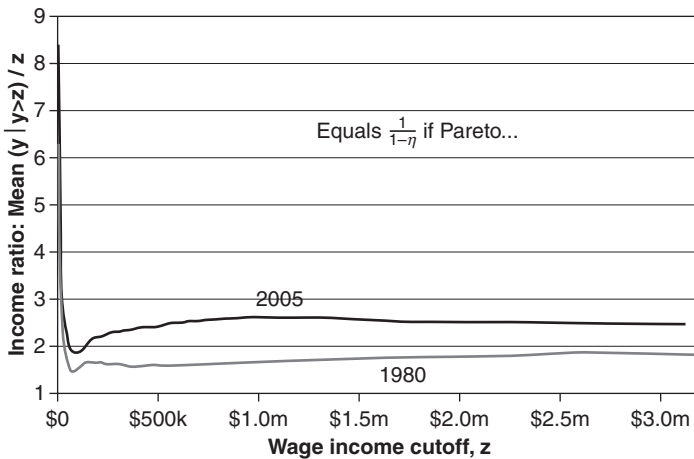


FIGURE 3.3. The Pareto nature of labor income

Source: Charles I. Jones and Jihee Kim, "A Schumpeterian Model of Top Income Inequality," *NBER Paper No. 20637*, October 2014, figure 5

cutoff. It turns out that number is about two and a half in 2005. Now move the cutoff. Look at everyone above a million, everyone above a million and a half, look at the average of the incomes above the cutoff, and divide it by the cutoff. For a Pareto distribution, it's a constant, and that's what you see in this graph. Moreover, in 1980, that same statistic is flat at a value of sort of one and a half or one and three quarters, once you get above, say, \$250,000 or \$300,000. In 2005, instead, the ratio is flat at a number more like two and a half.

The notion that top incomes are Pareto jumps out at you from the data. That's true for labor income. Let me spend just a minute on connecting Pareto distributions with the kind of numbers that Piketty and Saez and others have made famous now, which are the share of income going to the top 1 percent.

It turns out, there's a one-to-one mapping between a key parameter of the Pareto distribution and those top 1 percent shares. A Pareto distribution says, what fraction of people have incomes above Y ? Well, that's just Y raised to some power. And if you let $S(P)$ be the share of income going to the top P percent, so $S(1)$ is the top 1 percent share, it turns out that's just $100/P$ raised to some power. And the key power there is this parameter called η , it's a key parameter of the Pareto distribution, and basically if η is a half, then the share of income going to the top 1 percent is 10 percent. If η is three quarters, the share of income going to the top 1 percent is a third. And you can move that thing around and can get different shares. The point is that, if you want to write down a model where we can make sense of the data, what you're going to need is to write down a model where the data generate a Pareto distribution for incomes, and where the exponent in that Pareto distribution is changing. That's what you need to get out of a model. That's the kind of thing I'm thinking about.

One of the neat things about the Pareto distribution, and one of the neat things about the income data, is that it has a fractal

pattern. So if you ask, “What share of the top 10 percent’s income goes to the top 1 percent?” that’s some number. Which share of the top 1 percent’s income goes to the top tenth? That’s some number. What share of the top tenth goes to the top hundredth? That’s some number. It turns out, with a Pareto distribution, they are all the same number. And as top income and inequality rise, that same number is rising over time.

To get these Pareto distributions, to get a theory about why that Pareto exponent is changing, is relatively straightforward. And people have been thinking about this in many contexts for a number of years. The key insight is that exponential growth and Pareto distributions tend to go hand-in-hand. When you have one, you often get the other.

Imagine we have some entrepreneurs. Suppose you’re out there, sitting in your mom’s basement eating ramen noodles, trying to write an iPhone app. You finally write the new iPhone app, it gets posted on the iTunes store, and you get some income, some people start buying it. But initially, there aren’t that many people buying it. Your initial income is low. And then, as you’ve worked hard, as you’ve continued marketing your app, your income rises, and it rises exponentially at some growth rate. That’s the exponential growth.

The second piece of the intuition is that you need exponential growth to occur for an exponentially distributed amount of time. That’s a fancy way of saying something that’s really simple, which is, say there’s some probability that your business dies every period. By death in this entrepreneurial sense, we just mean somebody else comes along, and they were writing an iPhone app in their basement, and now they kick you off the top hundred list of the iTunes app store, and they’re the next Angry Birds, and you’re the old Angry Birds. There’s some constant probability that this happens every period. It turns out, if you put these two things together, you get a Pareto distribution. And the key Pareto exponent turns out to be the ratio of the growth rate to the death rate.

The way this works in a graph is fairly straightforward. Let x_0 denote your initial sales of the iPhone app. You're working hard. You're causing your sales to grow and it bounces around because it's kind of a random process, you have good days and bad days, good weeks and bad weeks. The more effort you put in, the faster is the growth. The faster the growth, the wider the distribution. There is going to be a bigger gap between the top person and the newest entrepreneur. So the faster the rate of exponential growth, the wider this inequality could be. And then the death rate pushes you down and kicks you out, and lets someone else start over. The higher is the death rate, well, that's going to restrain the extent to which the graph line can get away from the bottom. Exponential growth and creative destruction, which is what death corresponds to are the two forces that are operating: exponential growth pushes inequality up, while creative destruction pushes inequality down.

What we need in a richer model is a theory of what determines that growth rate, what economic forces affect the rate at which an entrepreneur's income grows, and what economic forces affect the death rate, or creative destruction. It's other people trying to come

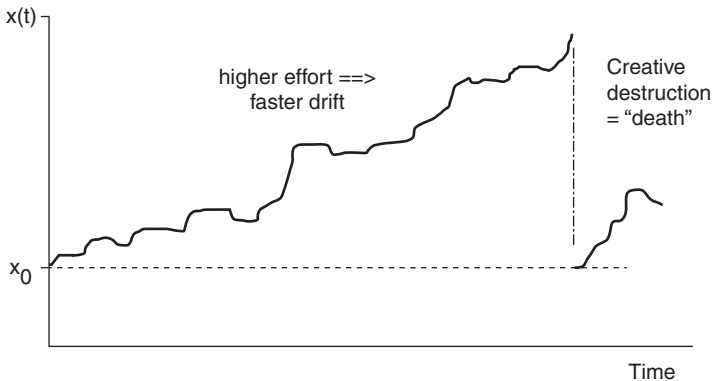


FIGURE 3.4. Basic mechanism: random growth with death \rightarrow Pareto

Source: Charles I. Jones and Jihee Kim, "A Schumpeterian Model of Top Income Inequality," *NBER Paper No. 20637*, October 2014, figure 9

up with new ideas and kick you out of the market and take over your position.

These same kinds of forces apply to wealth inequality. Piketty's book has a lot of discussion about how wealth inequality depends on $r - g$, and we saw some graphs like that earlier today. It turns out—and Piketty's book is based on a model like this at some level— $r - g$ is the growth part of the wealth. If you ask, how fast is wealth growing? Well, the interest rate is the key part of that wealth, and if you're looking at wealth normalized by income to get a stationary distribution, then you're going to get wealth growing at $r - g$, and then you might incorporate taxes, and you might take into account that wealth doesn't just grow. People eat some of it, too. And so, in some sense, the growth rate is $r - g$ minus taxes, minus some consumption. And then the effective death rate has to do with births and deaths. Those deaths can be creative-destruction kind of deaths, or they could be literal deaths. Piketty makes a point that the population growth rate is an important part of the wealth distribution. This same logic gives you the kind of model that he's talking about.

I think I've given you the basic spirit of the model. We've got researchers out there, spending time in their mom's basement, eating ramen noodles, trying to write iPhone apps. When they are successful, they become entrepreneurs, and then they're building their market share. The rate at which they build their market share is going to be a top determinant of income inequality. And then the rate at which other people are out there, trying to kick them off their pedestal, that's going to be another determinant.

I don't want you to think of this model as being just about iPhone apps. When I say entrepreneurs and researchers, I mean it to apply more broadly, and you can think of different applications of how this might work. Surgeons are trying to create new surgical techniques. You're known as being the best surgeon for this kind of heart replacement surgery, and then you can sell your idea to the

rest of the world. You build your market share by being successful and by advertising, and then someone else comes along trying to come up with a better surgical technique. Lawyers are similar, as are rock stars and musicians. Alan Krueger wrote a nice paper about the data for rock stars and top income inequality. But even middle managers are trying to come up with new business innovations that allow them to expand their segment of the business and get the attention of the higher-ups.

This story applies more broadly. And in some sense, connecting with what Josh Rauh was saying, one of the many useful insights of his research is that it's not just finance, it's not just CEOs, it's not just managers, it's not just athletes. It's all these people together. For top income inequality, you're not going to get one story in which it's only the hedge-fund managers. Part of it is the hedge-fund managers. But there are a lot of other people in the top who are seeing their incomes grow, and who are part of the top tenth of a percent or the top 1 percent. So you need a story that's broader. And that's what I'm trying to get at with this line of argument.

What are the economic determinants of top income inequality according to a model like this? Well, it turns out there are several of them. In a sense, one of the problems is there are too many. It's hard to get data on them, and it's hard to quantify and say, "This is really a thing that's driving change." Instead, the model highlights the kind of forces that could be driving top income inequality.

The first one to highlight is technology. I mentioned earlier that one of the key forces that increases top income inequality is anything that raises the growth rate of the entrepreneurs. Think about the World Wide Web. In the old days, when you had a great idea, you sold it to the people around you. Maybe that was the end or maybe you went around and sold it to a broader audience, but it was remarkably hard to start selling to a larger and larger market. With the advent of the World Wide Web, it's

become increasingly easy to grow your market very rapidly. And so your ability to grow rapidly, the returns to your effort if you're trying to grow the market rapidly, have increased enormously, and that's a worldwide force. When I look at that early graph, and I see top income inequality rising all over the world, it makes me think of technological change, and in my model, this is the kind of technological change that could lead to a rise in top income inequality.

Some of the other things are less obvious and maybe more interesting because they're less obvious. The second one is subsidies to research. Think about what a subsidy to research does in a model like this. Well, if you pay people to sit in their mom's basement writing iPhone apps, you subsidize more research. You get more people looking for new ideas, that's going to lead to more of this creative destruction. They're going to kick more people out. The death rate is going to go up. That is going to prevent the gap between the top people and the new entrepreneurs from expanding as far. That's actually going to reduce top income inequality in this setting. Research subsidies, anything that increases creative destruction, is actually going to limit top income inequality.

Conversely, you can think about blocking innovation. One of the things in the political economy of this that firms might try to do, is they may say, "We want to protect our market. We don't want to let people kick us out." And so you lobby, you set up mechanisms to protect your market. Anything that does that is going to limit creative destruction. That's going to increase top income inequality. And that would be a bad thing. So you can see there are good things and bad things. I see technology as a good thing, subsidies to research probably at the margin a good thing; blocking innovation will be a bad thing. So when you see top income inequality going up, it could be a good thing, it could be a bad thing. It depends on the forces that are behind it. That's one of the lessons that I've learned from this model.

Another example is misallocation. Imagine that in some country when entrepreneurs develop new ideas, the officials are more likely not to protect the intellectual property rights, or to take that property right and give it to a crony friend. In that case, there's more destruction. The returns to you from coming up with a new idea are lower. The returns to you from growing your market share once you have a new idea are lower, because it's going to be stolen, and then that's going to increase destruction. That is going to reduce top income inequality. Finally, consider taxes on entrepreneurial effort. Once you become one of the entrepreneurs here, a key determinant of top income inequality is the rate at which you're growing your profits. That depends on your effort. You have to work hard to do that, and anything that distorts your effort is going to affect the growth rate, and therefore affect top income inequality. And here it goes the way you naturally think. Higher taxes mean entrepreneurs put forth less effort and that lowers income inequality.

The other interesting thing that wasn't obvious in this model is that growth and inequality tend to move in opposite directions. At a casual level, it's tempting to look at the data and say, if inequality were driven by technological changes, for example, you'd expect that to increase growth and increase inequality together. When you look at the data, you don't see a correlation between growth and inequality that way.

It turns out that in my model, there are some forces that tend to make growth and inequality move in opposite directions. It's not clear they have to go in the same direction. There are two reasons. The first is one I've highlighted already, that faster growth means more creative destruction. But that means there's less time for inequality to grow—due to the higher death rate—and so you lower inequality. Entrepreneurs may work less hard to grow their market. The second force is actually less intuitive, which is that with greater inequality, research is riskier. Think about tech-

nological change like the World Wide Web. Sure, that raises the returns to being an entrepreneur, because you can now build your market share more easily. You would think that tends to make research more attractive, because if you happen to be successful as a researcher, that's going to allow you to be this great entrepreneur. That force is there.

It's also the case that if you decide to be a worker, working for one of those great entrepreneurs, you benefit from that technological change as well. Technological change raises the wages of entrepreneurs, but it also raises the wages of the people who are working for the entrepreneurs. So that doesn't shift people one direction or the other. That doesn't say you do more research or less research. What turns out to shift research here is the risk effect. If there's more inequality among entrepreneurs, the research process is riskier. You don't know if you're going to be the wildly successful one or just the normal one. So research is riskier. And risk-averse researchers tend to do less research, and that could actually decrease growth and decrease creative destruction. Greater inequality from, say, the World Wide Web could decrease growth. The linkages between growth and inequality are much more subtle than I appreciated before.

Let me conclude by giving you plausible explanations in these models for inequality in France and in the United States, the graph that I began with. In rising US inequality and, to some extent, rising worldwide inequality, there's lots of evidence that technologies are a part of the story. The ability to sell to larger markets, the ability to grow your market share because of information technology, has got to be part of the story. That comes through very clearly in a model like this. That's a worldwide phenomenon, not just something about the United States, and so we'd expect it to raise inequality everywhere.

Second thing: lower taxes on top income. We've seen a decline in the top marginal tax rate in the US. It's possible that that increases

effort by entrepreneurs and increases inequality, having ambiguous effects on growth in this framework. One has to be a little careful with this statement, however. The way taxes affect effort in our models is not nearly as clear as I think it is in the data. The data, I think, are pretty clear that there's an effect there. However, because substitution and income effects tend to cancel, it's a little harder to get taxes to affect top inequality in a model like this.

Then, what about France? I played with a lot of different models as I've been working this project for the last three years. And a lot of times I would get situations where rising inequality in the US is good because of technological change, and the fact that France hasn't let their inequality rise is bad, because they're resisting that technological change. It turns out in this framework that I have now, it's more subtle than that. That could be what's going on. It could be that France is delaying the adoption of good technologies, and that would clearly be a bad thing. Or it could be there's increased misallocation. Maybe they're killing off entrepreneurs more quickly, and that's causing inequality not to rise with technology.

On the other hand, there could be efficiency-enhancing explanations. Maybe it's the case that France is subsidizing research. We tend to do too little research in lots of settings. If France subsidizes research, that's going to result in more creative destruction. A higher death rate is going to kill people off, and that's going to lower inequality. So that could be going on in France. Or, it could be that France has reduced the blocking of innovations. Maybe in France these older firms were really protected. They didn't let competitors come in. Maybe France has relaxed this protection so that there's less blocking of innovations. That means there's more creative destruction and, again, that would tend to restrain inequality. Looking at the data through the lens of this model, the bottom line is it's not clear what's going on. We have stories where France could be doing the right thing or France could be doing a bad thing, and similar for the US as well.

To conclude: what are the policy implications? From my standpoint, the policy implications are relatively unclear in that I don't know which of these things is driving the rise in inequality. On the other hand, the one thing I'll say—that again was something that emerged fairly clearly from the kind of models I was playing with—is that policies that encourage research in this framework tend to lead to more growth and more creative destruction. Creative destruction tends to restrain inequality. And in general in these kinds of models, policies that encourage research are a good idea. And so even if inequality is not something we care about directly, I think it's the case that if you adopt policies that encourage research, one byproduct will be restraining the rise in inequality.

Question and Answer Session

QUESTION: I was thinking a little bit about the interpretation of death and I was wondering whether one of the ways that there can be death in the model is people moving outside of the country. And one example is that anecdotally wealthy people have been moving out of France, perhaps to London, perhaps to the United States, because of maybe either perceptions about the tax code or perceptions about the business climate. So, for example, the population of the United Kingdom has increased over the last ten years by about 9 percent. The percentage increase of French people living in the UK has been about 50 percent. I was wondering if one can potentially use your model to explain different countries' changes in income inequality by also considering what drives people to stay and go. It may be about perceptions about tax policy. It may be about perceptions of the business climate.

JONES: My coauthor on this research project did some work on taxes and top income inequality. I was struck by how low the top

marginal tax rates were in France and how much they'd fallen. When she looked at feeding in tax rates into the model, you get a lot less action out than I would have expected. My casual impression is tax rates were high in the 1970s everywhere, and they came down a lot in the US, and maybe not so much in France, and that's right in line with the movements in inequality. You can get something along those lines. But it was not nearly as much as my casual impression led me to believe. But certainly part of the explanation for why top inequality hasn't risen so much in some of those economies could easily be related to the fact that people are leaving to avoid higher taxes and particularly the threat of higher taxes or the perception they may be higher in the future.

QUESTION: Chad, your model does focus on those creative destruction roles so importantly. And maybe one of the things that could be going on here in thinking about it as you presented it, is technology preventing creative destruction against your own product. It could be changing the product or improving. And you can see that in some social media. I wonder if that might be something you can comment on or test for it at some point?

JONES: Yes, I think that's a great point. Anything that blocks creative destruction, and to me I think businesses have lots of incentives to try to block that, would work along those lines. One of the facts that is related is shown in a graph that I didn't have time to put up. It was something that Steve Davis and John Haltiwanger have worked on: business dynamism in the US. And the fact that I find completely striking but which is consistent with this kind of connection is job creation and job destruction. You see a downward trend in that since 1980 in the US. If you ask, what fraction of employment is due to young firms, there's a downward trend, and it's true across industries.

It's remarkably robust. And that would be another thing that could lead to a rise in inequality in this kind of setting.

QUESTION: Given the horizon a model like yours would likely have, why would you make such a stark distinction between level and growth effects? Because a lot of level effects are going to look like growth effects when you actually go to the data with this. The other thing I wanted to think about in that regard is linking this up to what we had in the previous paper, which is a lot more first-generation people, all these other things, which in some sense would say destruction is on the rise. And I would also go back to what I said before, which is all the market changes and other things, and how those would feed into your model. The guy who's in the basement can get funded a lot faster than he could compared with 1950. When Ray Kroc took on McDonald's, it took a long time to develop. In today's world, Five Guys can expand a lot faster because of all the improvements in capital markets. How would that fit into your model?

CHAD JONES: I think that's right. On the level effect and the growth effects, that's exactly right. It turns out in this model, solving for the transition dynamics is not easy. It's a model of heterogeneous agents, so it's got the usual kind of problems, so that's why I haven't done it.

On the McDonald's and the Five Guys and the dynamism point, I'm with you. Casual introspection, certainly living in Silicon Valley, you have to think: what's a more dynamic place? Things are more dynamic than they have ever been.

