ECONSTOR
Make Your Publications Visible.

A Service of
ZBW
Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics

Fix, Blair

Article — Published Version

Redistributing Income Through Hierarchy

Real-World Economics Review

Provided in Cooperation with:
The Bichler & Nitzan Archives

Suggested Citation: Fix, Blair (2021) : Redistributing Income Through Hierarchy, Real-World Economics Review, ISSN 1755-9472, World Economics Association, Bristol, Iss. 98, pp. 58-86, https://rwer.wordpress.com/2021/12/16/rwer-issue-no-98/

This Version is available at:
https://hdl.handle.net/10419/247778

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

http://creativecommons.org/licenses/by-nc-nd/4.0/

WWW.ECONSTOR.EU

Mitglied der Leibniz-Gemeinschaft

real-world economics review, issue no. 98
subscribe for free

Redistributing Income Through Hierarchy

Blair Fix [York University, Toronto, Canada]

Copyright: Blair Fix, 2021
You may post comments on this paper at
https://rwer.wordpress.com/comments-on-rwer-issue-no-98/

Abstract
Although the determinants of income are complex, the results are surprisingly uniform. To a first approximation, top incomes follow a power-law distribution, and the redistribution of income corresponds to a change in the power-law exponent. Given the messiness of the struggle for resources, why is the outcome so simple?
This paper explores the idea that the (re)distribution of top incomes is uniform because it is shaped by a ubiquitous feature of social life, namely *hierarchy*. Using a model first developed by Herbert Simon and Harold Lydall, I show that hierarchy can explain the power-law distribution of top incomes, including how income gets redistributed as the rich get richer.

To study income is to be perplexed

In a famous 1933 speech, John Maynard Keynes lamented his discontent with capitalism:

> "It is not intelligent, it is not beautiful, it is not just, it is not virtuous — and it doesn't deliver the goods. In short, we dislike it, and we are beginning to despise it. But when we wonder what to put in its place, we are extremely perplexed." (Keynes, 1933)

Today, we might attribute a similar sentiment to researchers who study the distribution of income. Heterodox economists agree that the current distribution of income is 'not virtuous', and that the dominant approach to understanding income (marginal productivity theory) 'doesn't deliver the goods'. But when we look for a better approach to understanding inequality, we are 'extremely perplexed'.

Like so many aspects of human society, the distribution of income is frustratingly complex — the joint result of ideology, politics, class struggle, and everything in between. Reviewing these complexities, Sandy Hager argues that it may be best to study inequality using a 'plurality of methodological approaches' (2020). I largely agree, but with one caveat. While the causes of inequality are surely complex, the outcome is not. Regardless of where we look, we find that top incomes follow a simple pattern: they are distributed according to a *power law*. That is, the probability of finding someone with income I is roughly proportional to $I^{-\alpha}$.

If the causes of income are complex, why can we model the result with a single parameter — the power-law exponent α? Moreover, why can we model income redistribution by shifting this parameter, and this parameter alone? Given the complexity of human society, the success of such a simple model seems unreasonable. How do the myriad of different forces driving inequality 'conspire' to create such a simple outcome?

One possibility is that the ultimate causes of inequality are indeed complex, but that they are mediated by a 'proximate' cause that is far simpler. If this mediator was ubiquitous, it could lead to the simple outcome that we observe (the power-law distribution of top incomes). So, what might this mediator be?

I propose that it is *hierarchy*. Although largely ignored by mainstream economics, hierarchy is a common feature of human life. It seems to be the default mode for organizing large groups. And its use appears to have spread with industrialization (Fix, 2021a).

The distinguishing feature of hierarchy is the chain of command, which concentrates power at the top. It is this feature, I propose, that mediates the distribution of top incomes. For a power-law to emerge, all we need is for income to increase (roughly) exponentially with hierarchical rank. Varying this rate of increase then causes a redistribution of top incomes. The result is a proximate explanation of inequality that locates the source of power-law distributions in the chain-of-command structure of hierarchies (Figure 1).

Figure 1: Hierarchy as a proximate cause of inequality.

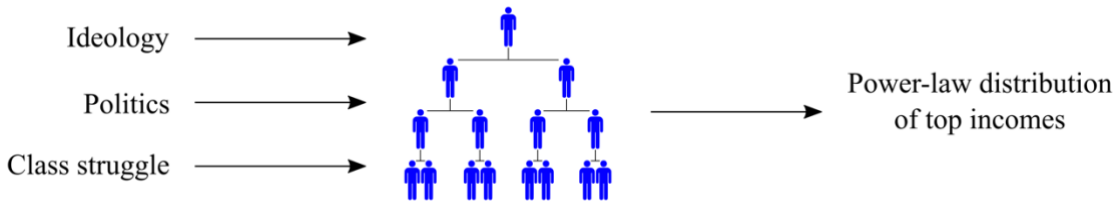

Although this focus on hierarchy does not explain the 'ultimate' cause of inequality, it dramatically changes the way we think about the problem. It is one thing to look at top incomes and wonder what is causing them to increase. It is quite another thing to understand that top incomes can be directly linked to the hierarchical pay structure of individual firms.

In the latter case, we realize that each firm is a microcosm of the distribution of income at large. Moreover, when we link top incomes to hierarchy, we are implicitly connecting the distribution of income to the power structure of society. The consequence is rather incendiary. When top incomes increase, it suggests that firm hierarchies are becoming more despotic.

The shape of top incomes

Before discussing how hierarchy relates to top incomes, we must cover some requisite knowledge about income and its (re)distribution. In the introduction to his 2014 treatise on inequality, Thomas Piketty observed:

> "Intellectual and political debate about the distribution of wealth has long been based on an abundance of prejudice and a paucity of fact." (Piketty, 2014)

Today, thanks in large part to Piketty's work, the 'paucity of facts' is no longer a problem (at least among people who are concerned with facts).[1] Many people know that income inequality has risen dramatically in recent decades. Matters came to a head during the Occupy movement when the term 'one-percenter' became a well-known put down (Di Muzio, 2015). The term alludes to the growing divide between the income of the majority (the bottom 99%) and the income of the elite (the top 1%).

Figure 2 shows this divide — the income share of the US top 1%. The U-shaped trend is now well known. After World War II, US inequality declined rapidly and then remained low for 30 years. But from the 1980s onward, inequality rose dramatically.

Figure 2: The fall and rise of US inequality.

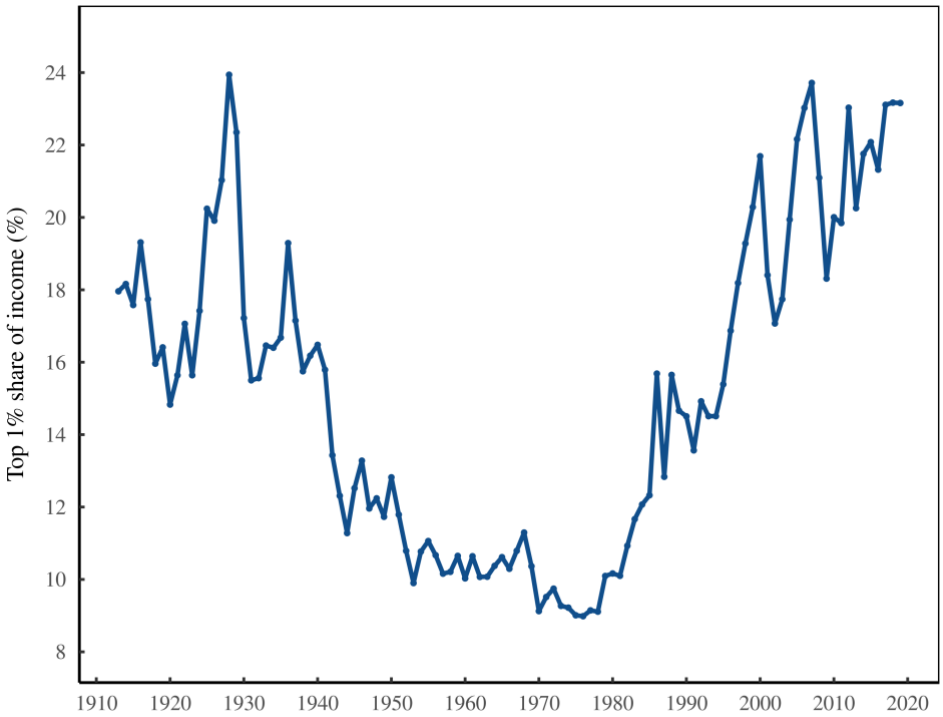

The timing of this rising inequality has eluded few observers. It corresponded with a seismic shift in US politics — a turn from the post-War expansion of the welfare state to the 'trickle down' policies of the

[1] Although Piketty popularized the study of top income shares, he built on the work of many researchers, including Atkinson & Harrison (1978), Atkinson & Bourguignon (2001), Atkinson & Piketty (2010), and Alvaredo, Atkinson, Piketty, & Saez (2013).

Reagan era. Given this conspicuous political shift, many researchers leap straight from the inequality evidence to a list of possible 'causes'.

I sympathize with this move but think that it is partially premature. Yes, we should look for correlates of inequality, of which there are many. (See, for instance, the work of Huber, Huo, & Stephens, 2017.) But we should also realize that looking only at the income share of a specific group (like the top 1%) gives a rather narrow window into the wider distribution of income.

Unfortunately, looking at the whole distribution of income takes some technical skills, which is likely why doing so is less popular than studying top income shares alone. Still, if we want to study growing inequality, we need to understand how *all* income is distributed.

Viewing the distribution of income in its entirety

In the interest of accessibility, I offer here a brief tutorial of how to visualize income distributions from top to bottom using log histograms. Readers familiar with this technique can skip to the next section.

The most basic way to visualize a distribution of income is to use a histogram. To construct a histogram, we put the data into size 'bins' and count how many observations occur within each bin. Then we plot the results.

Figure 3A shows a histogram of a hypothetical distribution of income. (For reference, this simulated society has about 10 million people, a median income of $30,000, and a top 1% income share of about 20%. It's intended as a scaled-down version of the modern United States.)

I have put individual incomes into bins that are $2000 wide. On the vertical axis, I have plotted the number of people within each bin. Each point represents the person count, plotted at the midpoint of the income bin. This representation of a histogram, which connects bin counts with a line, is sometimes called a 'frequency polygon'. But for ease of reference, I will simply call it a 'histogram'.

Our Figure 3A histogram does not look like the familiar 'bell curve'. Rather, it has a 'fat' right tail that continues far past the chart's income cutoff of $100,000. This fat tail is a ubiquitous feature of distributions of income and is the face of inequality in histogram form. It tells us that some individuals earn far more than the average person.

Figure 3: Three ways to visualize a distribution of income.

Using a simulated distribution of income, this figure shows three ways of visualizing the distribution with a histogram. Panel A shows the standard form with income bins of constant size. The problem here is that the rich are 'off the chart'. Panel B uses log-spaced bins, with both the bins and counts plotted on log scales. We see the power-law tail of top incomes on the right side. Panel C normalizes the histogram so that it is comparable to different samples of income.

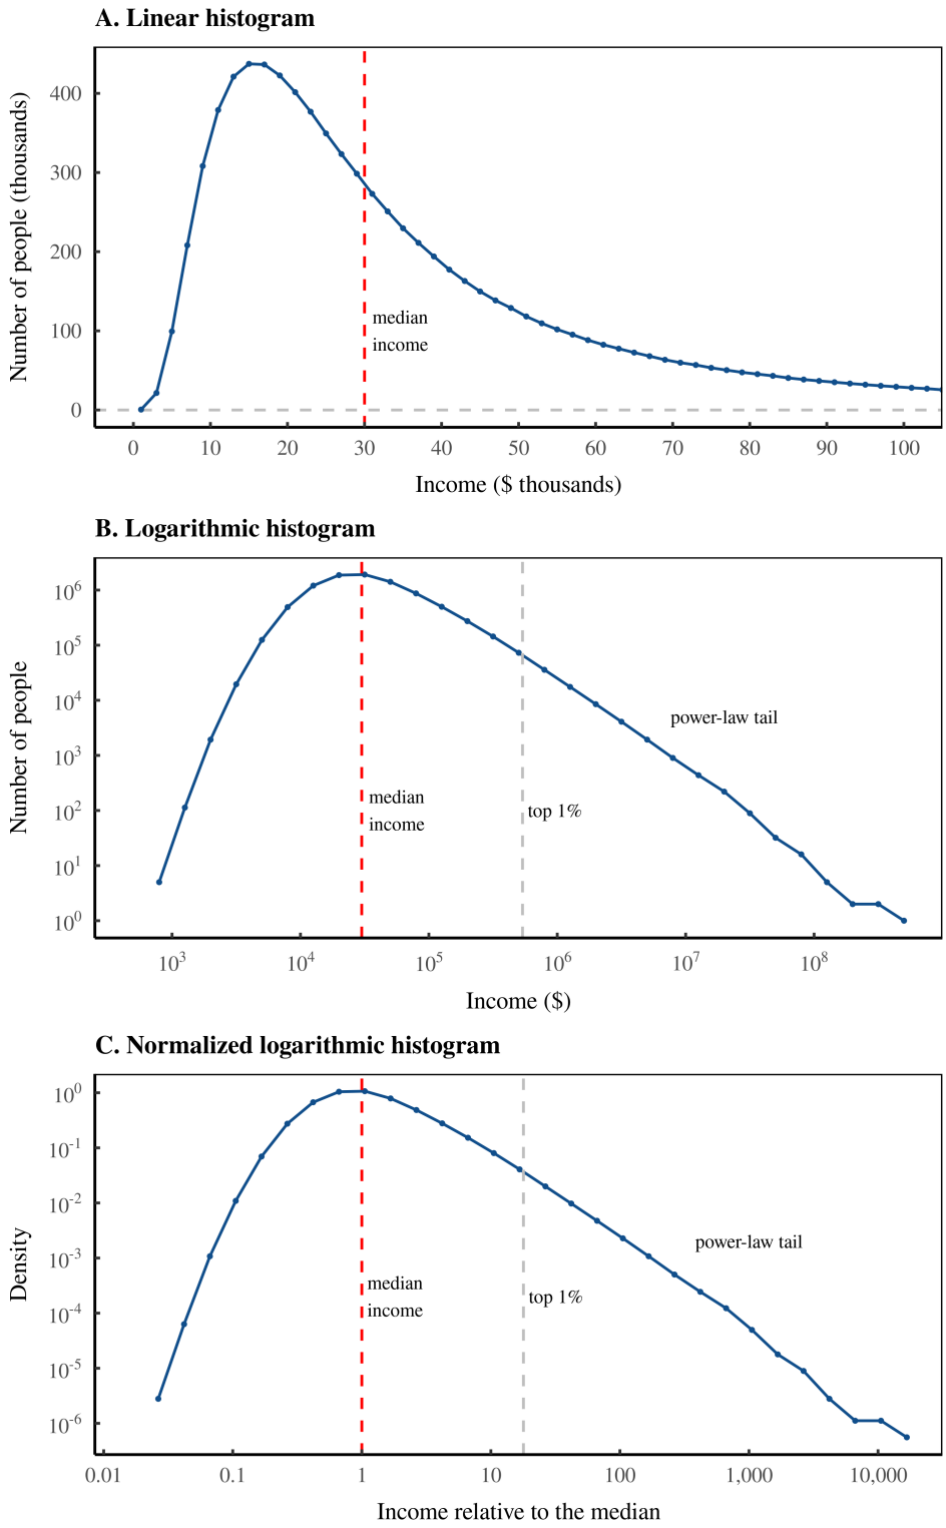

The problem with our standard histogram is that we cannot see the rich — they are literally off the chart. To visualize the distribution of top incomes, we need a different approach. The best option is to move to a logarithmic histogram.

A log histogram uses income bins that are logarithmically spaced. For instance, the first bin might go from \$1 to \$10, the second from \$10 to \$100, the third from \$100 to \$1000, and so on.[2] By using log spacing, we can reach enormous incomes with relatively few bins. The key is that we then plot both the bins and the corresponding counts on logarithmic scales. In the resulting logarithmic histogram, shown in Figure 3B, we can see the rich and the poor alike. The poor are on the left, with incomes that are far smaller than the median. And the rich are on the right, with incomes that are far larger than the median.

In our log histogram, we can also see a key feature of top incomes: they tend to be distributed according to a power law.[3] A power law is a type of distribution in which the probability of finding a person with income I is proportional to that income, raised to some exponent α:

$$P(I) = c \cdot I^{-\alpha} \qquad (1)$$

Power law distributions have the interesting feature that if we plot their logarithmic histogram (as we have in Figure 3B), we get a straight line. The reason is beautifully simple. When we take the logarithm of both sides of Equation 1, we get a linear relation whose slope is $-\alpha$:

$$\log P(I) = \log c - \alpha \cdot \log I \qquad (2)$$

So, the fact that the right tail of our log histogram looks like a straight line means that top incomes roughly follow a power law.

If we wish to compare the distribution of income at different points in time (or between different countries) there is one last step: we must 'normalize' the histogram. To do that we convert incomes from dollar values to relative values. In Figure 3C, I compare all incomes to the median. Next, we normalize the histogram counts so that they are unaffected by sample size. I do that in Figure 3C by converting bin counts to a 'probability density'. This transformation defines the vertical scale so that the area under the histogram sums to 1.

Although our normalized histogram looks identical to the un-normalized version, it now has standardized axes. That means we can compare different distributions of income.

[2] Instead of using log-spaced bins, another option is to use linear bins but count the frequency of log(income). The results will be the same.

[3] The power-law distribution of top incomes (and wealth) was discovered at the turn of the 20th century by Vilfredo Pareto (1897). For a sample of subsequent confirmations of Pareto's discovery, see Di Guilmi, Gaffeo, & Gallegati (2003), Clementi & Gallegati (2005), Coelho, Richmond, Barry, & Hutzler (2008), Toda (2012), and Atkinson (2017).

Income redistribution in the United States

Now that the reader has the requisite knowledge, we are ready to look at the distribution of US income in its entirety. Figure 4 shows the US distribution of income in 1970 and 2007. I have chosen these years because they are the dates of minimum (1970) and maximum (2007) inequality in recent US history. The change in the distribution of income is easy to spot.

Figure 4: The US distribution of income, 1970 and 2007.

Using the log-histogram technique outlined in Figure 3, I plot here the distribution of US income when inequality was at a minimum (1970) and a maximum (2007). For sources and methods, see the Appendix.

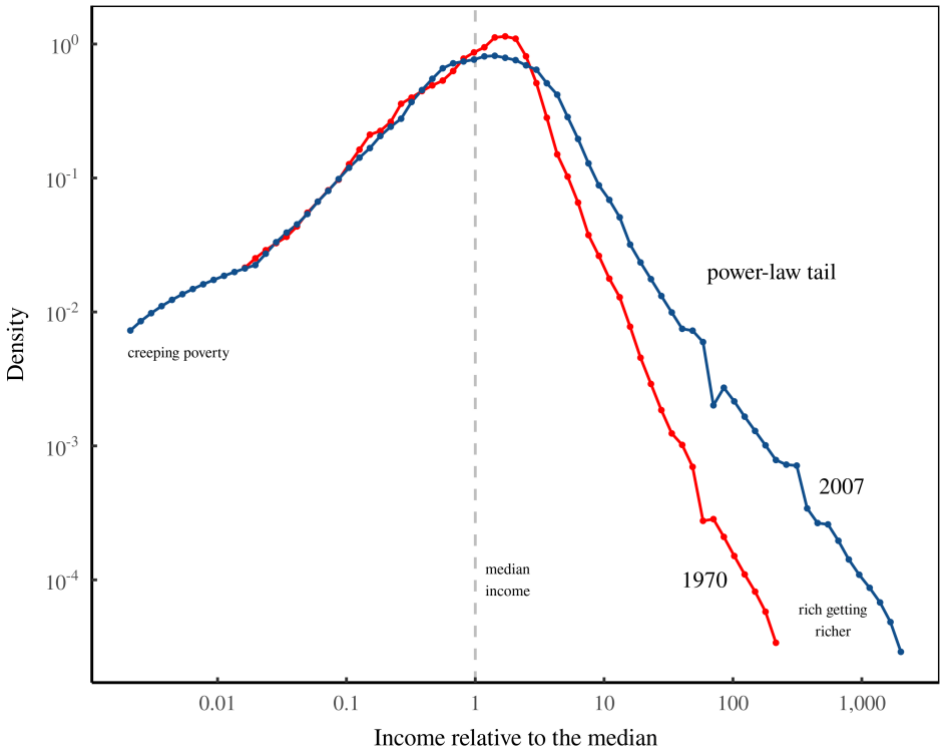

Let us start, however, with what did *not* change between 1970 and 2007. To spot a lack of change, look for locations where the two histograms overlap. In Figure 4, we can see that this overlap occurs below the median income, where the two histograms are nearly identical. This similarity tells us that for the bottom half of Americans, little has changed (in terms of relative income) over the last 4 decades.

Among the American poor, though, there is one conspicuous difference between 1970 and 2007: in the latter year, the social safety net had been removed. This removal appears in Figure 4 as a leftward extension of the blue histogram into ever-more diminutive incomes. This is creeping poverty in histogram form. Today, many Americans earn less than 1% of the median income — something that was not true in 1970.

While creeping US poverty is worth studying, it is not the subject of this paper. Instead, I am concerned with the right-side of the histogram. Here we can see the egregious redistribution of top incomes. Between 1970 and 2007, the American rich got richer … much richer. Whereas in 1970, no one earned more than a few hundred times the median income, by 2007, a handful of Americans earned more than *1000* times the median.

It is easy to marvel at the absurd size of top US incomes. But here I am more concerned with the *uniformity* of income redistribution. As expected, top US incomes (roughly) follow a power-law distribution, evident as the straight right tail in both distributions. What is fascinating is that despite the complex reasons for growing US inequality, to a first approximation, all that changed between 1970 and 2007 is the *slope* of the distribution tail.

This simple result deserves an explanation. Why can we model the messy business of the rich getting richer by turning a single dial — the power-law exponent of top incomes?

Income redistribution among all countries

Before we conclude that the rich getting richer is a simple process, we ought to look at more data. It could be, for instance, that the United States is a uniquely simple case, and that elsewhere, the redistribution of income is more complicated.

To test this possibility, let's look at income redistribution in every country for which there is suitable data. Using data from the World Inequality Database, Figure 5 plots the income-redistribution trends for 176 different countries covering the years 1900 to 2019.

Figure 5: As top income shares grow, the income-distribution tail gets fatter.

This figure visualizes income redistribution among countries. Each line indicates the path through time of a particular country. The vertical axis shows the country's top 1% share of income. The horizontal axis shows estimates for the power-law exponent of top incomes (fitted to the top 1% of incomes). As top income shares increase, the power-law exponent tends to decline, indicating that the distribution tail gets fatter. For sources and methods, see the Appendix.

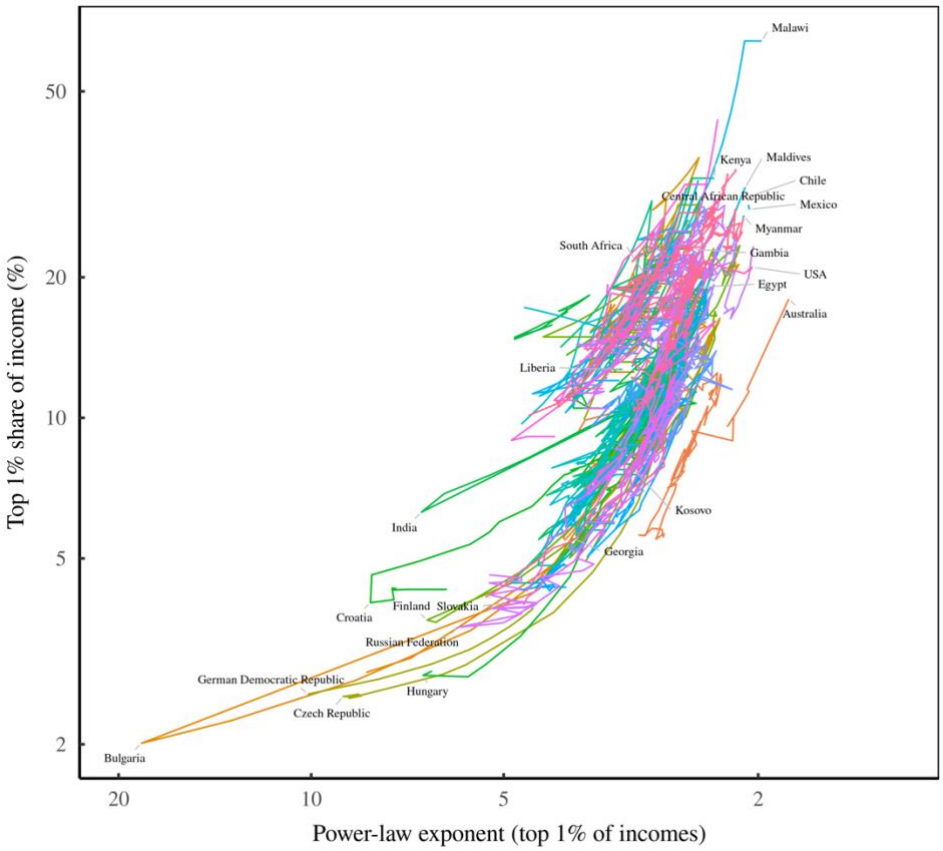

Rather than show the complete distribution for each country (in each year), I have plotted the top 1% income share against the power-law exponent of top incomes. To reiterate, this exponent measures the slope of the income distribution tail. A smaller exponent indicates a fatter tail. (For power-law fitting methods, see the Appendix.)

If income redistribution was a messy, heterogeneous process, we would expect no clear relation between top income shares and the power-law exponent of top incomes. But that is not what we find. Instead, we see in Figure 5 a very clear relation. Growing top income shares are associated with a decline in the power-law exponent of top incomes. In other words, there is startling uniformity in the way that societies redistribute income.

Generating power laws

To understand the distribution of top incomes, we need to understand more about power laws. Where do they come from? How are they generated?

Although the causal mechanisms may appear complex, the mathematical mechanisms for generating power laws are surprisingly simple. I will discuss two main routes. (For a review of mechanisms for generating power laws, see Mitzenmacher, 2004.)

The first route to a power law is through income *dynamics*. Suppose an individual starts out with annual income I. Over time, their income grows and shrinks for reasons that we do not understand. But what we do know is that this income change can be modelled as a random number. After t years, the person's new income is the product of successive random growth rates, g:

$$I_t = I_1 \cdot g_1 \cdot g_2 \cdot \ldots \cdot g_t \qquad (3)$$

Now suppose that everyone's income behaves the same way: it is the product of a series of random growth rates. After many growth iterations, the resulting distribution of income will follow a lognormal distribution — a fact discovered by Robert Gibrat (1931).

To get a power-law distribution, we introduce one more requirement: a lower 'wall' that limits the smallness of incomes. If anyone's income gets below this lower threshold, it gets 'reflected' in the opposite direction. After many growth iterations, income will be distributed according to a power law.

This 'stochastic' model of income was first articulated by David Champernowne (1953). While the model's mathematics are beyond dispute, many political economists find its appeal to 'randomness' troubling. After all, incomes have definite causes (or so we believe). But to be fair to the Champernowne model, it does not claim that income dynamics are *actually* random, only that we can model them as such.

The Champernowne model tells us that we can understand the power-law distribution of top incomes without knowing anything about the complexities of human behavior. All that we need are general assumptions about the dynamics of income. I find this result fascinating because it is counter-intuitive. Yet it is also underwhelming because it does not tell us why people earn what they do. For that reason, I will focus on a second route to power laws — a route that can be tied to social structure.

The second route to a power law comes from merging two different exponential functions. Suppose two variables, x and y, are both exponential functions of a third variable, t:

$$x = e^{a \cdot t} \qquad (4)$$

$$y = e^{b \cdot t} \qquad (5)$$

If we combine these two functions and eliminate t, we find that x and y are related by a power law:[4]

$$y = x^{b/a} \qquad (6)$$

So, we can create a power law by merging two exponential functions. The question is, why would such functions apply to income? The answer, I propose, is simple. These are the equations that describe income in a hierarchy.

Power-laws via hierarchy

Hierarchies are perhaps the dominant feature of our working lives. Yet paradoxically, they rarely enter into mainstream theories of income distribution. Fortunately, a handful of researchers have explored the distributional consequences of hierarchy. I build on their work here.

To my knowledge, the first person to explicitly model income within a hierarchy was the polymath Herbert Simon (1957). Simon noted that hierarchies are government by a chain of command in which each superior controls multiple subordinates. The consequence is that the number subordinates one controls increases exponentially with rank. At the same time, income within a hierarchy tends to increase exponentially with rank. Combining these two exponential functions gives a power law.

Simon, though, was not interested in the power-law distribution of top incomes. Instead, he was interested in another power law — the fact that CEO pay scales with the power of firm size:

$$\text{CEO pay} \propto (\text{Firm size})^{D} \qquad (7)$$

Simon argued that this scaling (which was discovered by David Roberts in 1956), stemmed from hierarchy. It was caused by merging the exponential growth of subordinates (with hierarchical rank) and the exponential growth of pay (with hierarchical rank).

[4] Here are the algebraic steps. First, take the logarithm of both functions and solve for t:

$$t = \frac{1}{a}\log x$$

$$t = \frac{1}{b}\log y$$

Next, combine the two equations to eliminate t:

$$\log(y) = \frac{b}{a}\log(x)$$

Note that $\frac{b}{a}\log(x)$ is equivalent to $\log x^{b/a}$. Therefore,

$$y = x^{b/a}$$

Although largely ignored by mainstream economists, Simon's reasoning remains sound. In fact, we can extend it to every member of the hierarchy (not just CEOs). As Figure 6 indicates, relative income within hierarchies scales with the number of subordinates one controls. For ease of reference, I give 'the total number of subordinates' a shorthand name. I call it 'hierarchical power', defined as:

$$\text{hierarchical power} = 1 + \text{number of subordinates} \qquad (8)$$

Across a wide variety of institutions, relative income appears to scale with hierarchical power.

Figure 6: Within hierarchies, income grows with hierarchical power.

This figure shows evidence from a variety of institutions indicating that relative income within hierarchies scales with 'hierarchical power'. In the case-study firms and the US military, income is measured relative to the average in the bottom hierarchical rank. Each point indicates the average hierarchical power within a rank. For CEOs, income is measured relative to the average pay within the firm. I assume the CEO commands the firm, meaning their hierarchical power is equivalent to the firm's total employment. For sources and methods, see the Appendix.

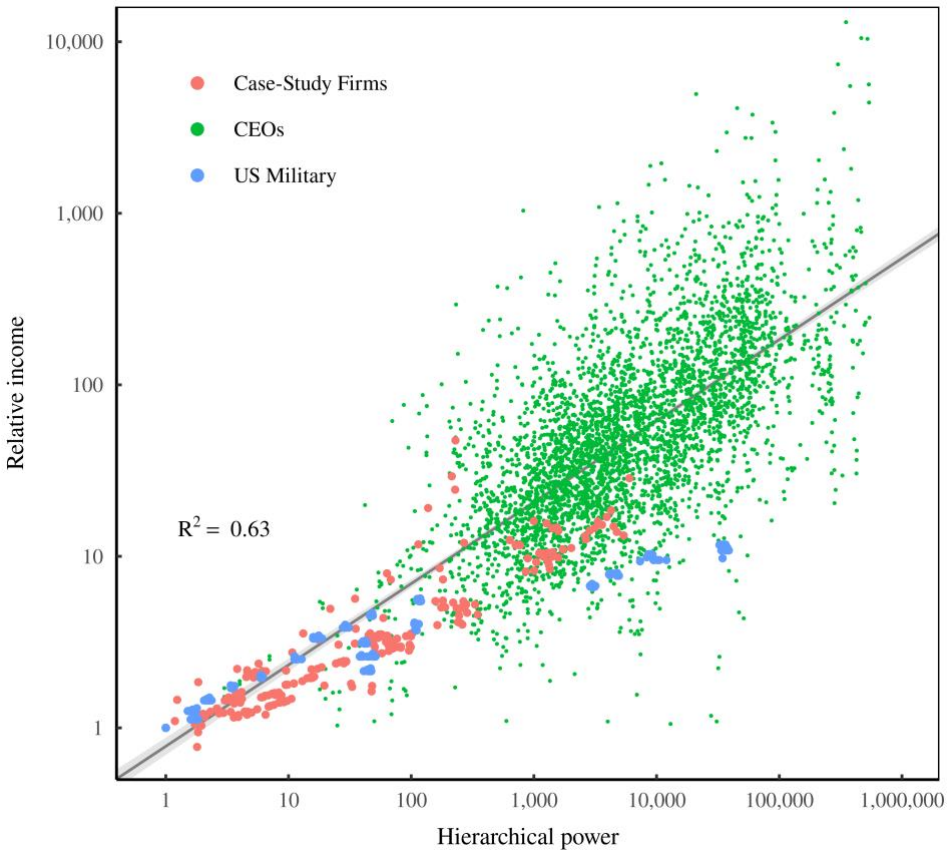

Two years after Herbert Simon published his results, Harold Lydall (1959) realized that the same model of hierarchy could explain the power-law distribution of top incomes. The mechanism was exactly the same — the merger of two exponential functions. (Interestingly, Lydall appears to have been unaware of Simon's work.)

Like Simon, Lydall assumed that income grows exponentially with hierarchical rank. That gives exponential function number one. The second function comes from the number of people within each rank. As we move up the hierarchy, the number of people within each rank declines exponentially — a consequence of the nested chain of command. By merging these two exponential functions, Lydall showed that hierarchy could create a power-law distribution of income.

Because Simon and Lydall's pioneering research was completed a half century ago, one would think that today there would be a burgeoning literature on the distributional consequences of hierarchy. Sadly, this is not the case. Instead, shortly after Simon and Lydall published their work, the study of income distribution became dominated by human capital theory, which focused on personal traits and neglected 'structural' explanations of income (Fix, 2021b). And so today, we know little about how hierarchy affects the distribution of income.

Despite the historical neglect, I think focusing on hierarchy is a promising way to understand income (Fix, 2018, 2019b, 2020). And as I discuss below, I think it is also a promising way understand income *redistribution*.

A sign from CEOs

To understand how income redistribution relates to hierarchy, I propose that we return to where Herbert Simon started: with CEOs. Over the last 40 years, the relative pay of US CEOs has increased dramatically. The timing of this pay explosion aligns tightly with rising US inequality. Figure 7 shows the trend.

Figure 7: Increasing US inequality corresponds with a growing CEO pay ratio.

The CEO pay ratio is calculated by dividing the pay of CEOs in the 350 top US firms (ranked by sales) by the average income of workers in the corresponding industry. For sources and methods, see the Appendix.

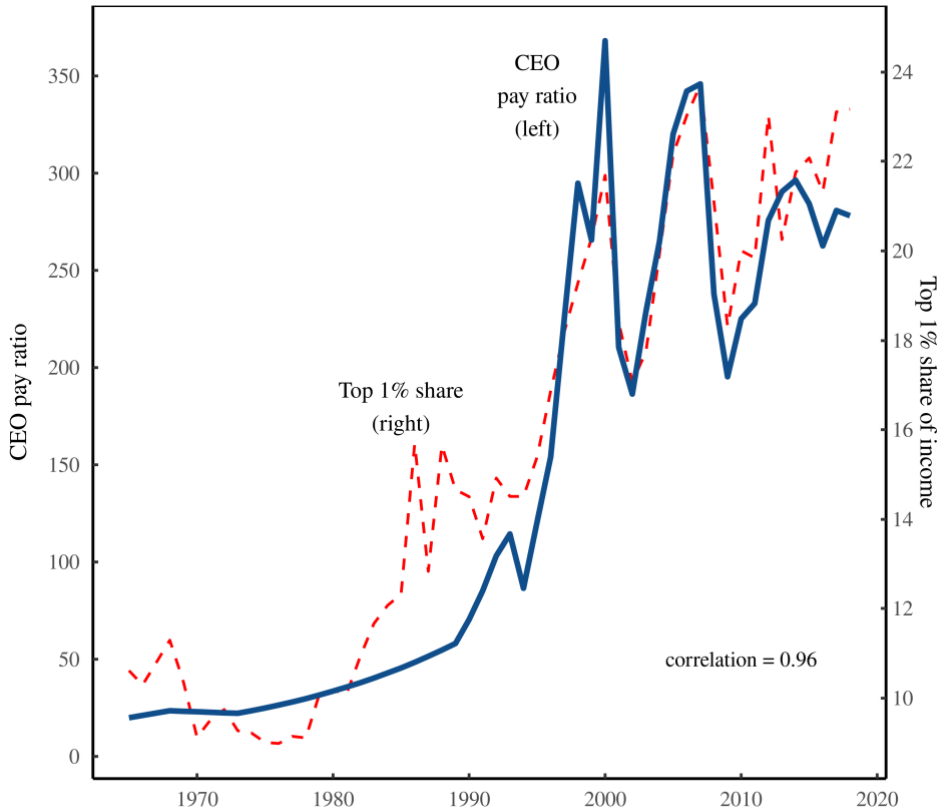

The obvious conclusion, reached by many observers, is that runaway CEO pay is related to runaway inequality. Interestingly, however, there have been few attempts to generalize this finding into a model of income distribution.

The way to do this, I believe, is by treating CEOs as canaries in the coal mine. I propose that the exploding pay of CEOs is part of a wider redistribution of income within hierarchies. It is evidence that US firms are becoming more *despotic*.

I use the word 'despotic' in both a general sense (as in the abuse of power) and in a more technical sense, as follows. A key feature of hierarchies is that they concentrate power at the top — a feature that inevitably creates problems. Yes, rulers *can* use their power to benefit the group. But they can also use their power to enrich themselves. The more they do so, the more 'despotic' the hierarchy.

Importantly, despotism is not just a game for rulers. It is a game played by everyone in the hierarchy. The result, I propose, is that the more despotic the hierarchy becomes, the more rapidly income will increase with hierarchical power. It makes sense, then, to use the scaling of income with hierarchical power, D, as a measure of the 'degree of hierarchical despotism'. The greater the value of D, the more despotic the hierarchy.

$$\text{relative income} \propto (\text{hierarchical power})^{D} \qquad (9)$$

To frame this idea, let's return to the empirical evidence. In Figure 8, I have replotted (as grey points) the empirical trend between relative income and hierarchical power (the trend originally shown in Fig. 6). Over top of this data, I show scaling relations for different values of D.

Figure 8: How the degree of hierarchical despotism, D, affects income.

Grey points replot empirical data from Fig. 6. Colored lines indicate the (hypothetical) scaling of income with hierarchical power for different values for D — the degree of hierarchical despotism. For sources and methods, see the Appendix.

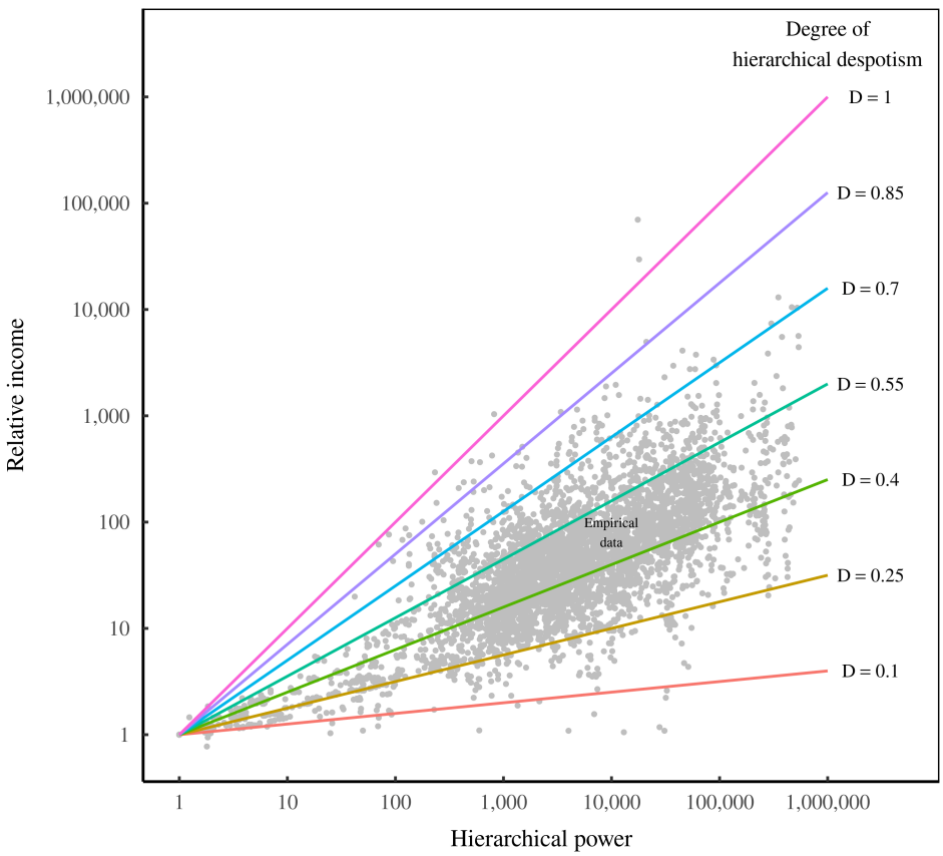

In large hierarchies, the value of D affects top incomes dramatically. For instance, when $D = 0.1$, a CEO with one million subordinates will earn only about 4 times more than a bottom-ranked worker. But when $D = 1$, the same CEO will earn a *million* times more than an entry-level employee.

US CEOs as canaries of hierarchical despotism

Based on the scatter in the empirical data (in Fig. 8), it seems clear that the 'degree of despotism' can vary between hierarchies. The question is, can the average degree of despotism also vary over time?

To answer this question definitively, we would need time-series data for the hierarchical pay structure of many different firms. Since such data does not exist, I propose a rougher approach: we use CEOs as despotism 'canaries'. Among US CEOs, we know that income scales with hierarchical power (where the CEO's hierarchical power is measured by firm size). What we do not know, though, is how this relation has changed with time.

To investigate this question, Figure 9 plots data for US CEO pay in two years: 1992 and 2007. In both years, the CEO pay ratio tends to increase with hierarchical power. Yet the rate of this increase differs. In 2007, CEO pay scaled more steeply with hierarchical power than it did in 1992. If CEOs are 'canaries' for a larger trend within firms, this result hints that US firms have become more despotic.

Figure 9: Changing hierarchical despotism among US CEOs.

This figure plots the relation between the CEO pay ratio and hierarchical power for US CEOs. I assume that CEOs command their respective firms, meaning their hierarchical power is equivalent to the firm's employment. Data for 1992 is shown as red triangles. Data for 2007 is shown as blue circles. Lines indicate the trend line, which indicates the 'degree of hierarchical despotism', D. The evidence suggests that US firms have grown more despotic over the period shown. For sources and methods, see the Appendix. Note 1: By 1992, the pay ratio of US CEOs had already increased significantly from its low point in the 1970s. Unfortunately, the data used here (from Execucomp) begins in 1992, so we cannot observe 'hierarchical despotism' in earlier years. Note 2: I estimate hierarchical despotism, D, using a regression that is fixed through the point (1, 1). Although it is usually inadvisable to force a regression through a fixed point, this is a special circumstance. By definition, when a firm has 1 member, that person has a hierarchical power of 1. And since there is only one member, the 'CEO pay ratio' is by definition 1. It follows that the relation between the CEO pay ratio and hierarchical power must go through the point (1, 1).

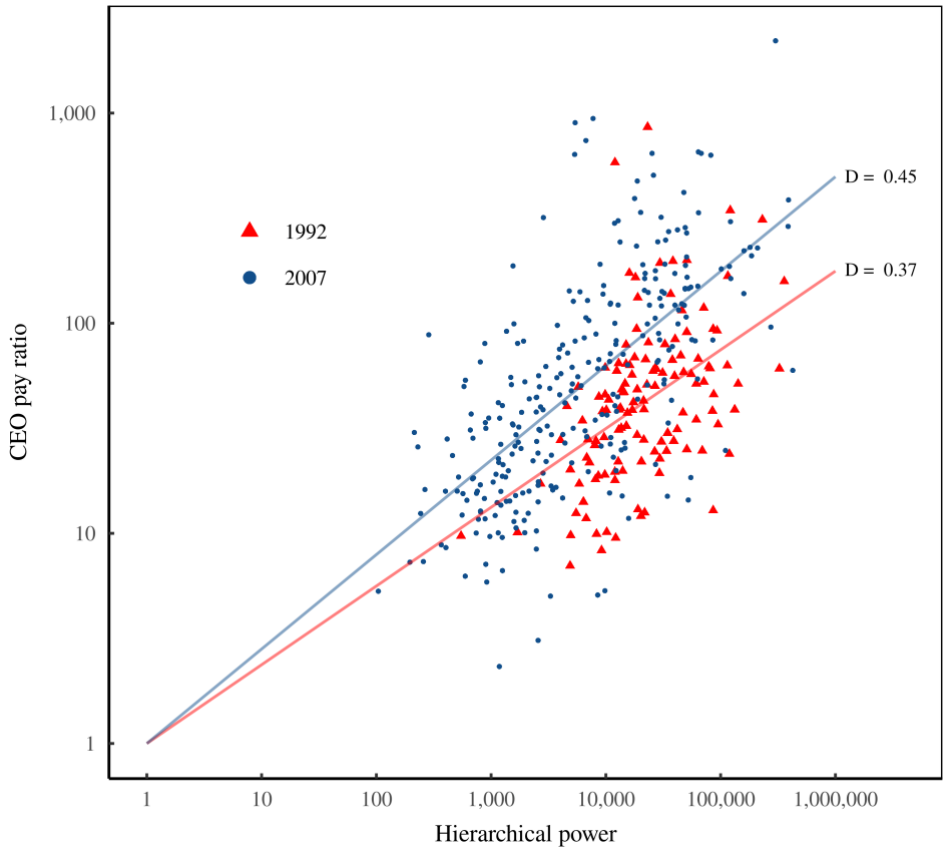

The next question is — does changing hierarchical despotism correspond with growing inequality? To test this possibility, we can generalize the method shown in Figure 9. In each year between 1992 and 2019, we regress the relative pay of US CEOs onto their hierarchical power. The result is a time-series estimate of the average degree of hierarchical despotism among US firms.

We want to know whether this changing despotism relates to rising inequality. The evidence, shown in Figure 10, suggests that it does. As my estimates for hierarchical despotism rise, so does the income share of the US top 1%.

Figure 10: Increasing despotism among US CEOs correlates with growing US inequality.

This figure generalizes the regression shown in Fig. 9. In each year between 1992 and 2019, I regress the pay ratio of US CEOs onto their hierarchical power. The slope of this regression is D, the estimated 'degree of hierarchical despotism' within these firms. Here, I show that this degree of despotism correlates with growing US inequality, as measured by the income share of the top 1%. For sources and methods, see the Appendix.

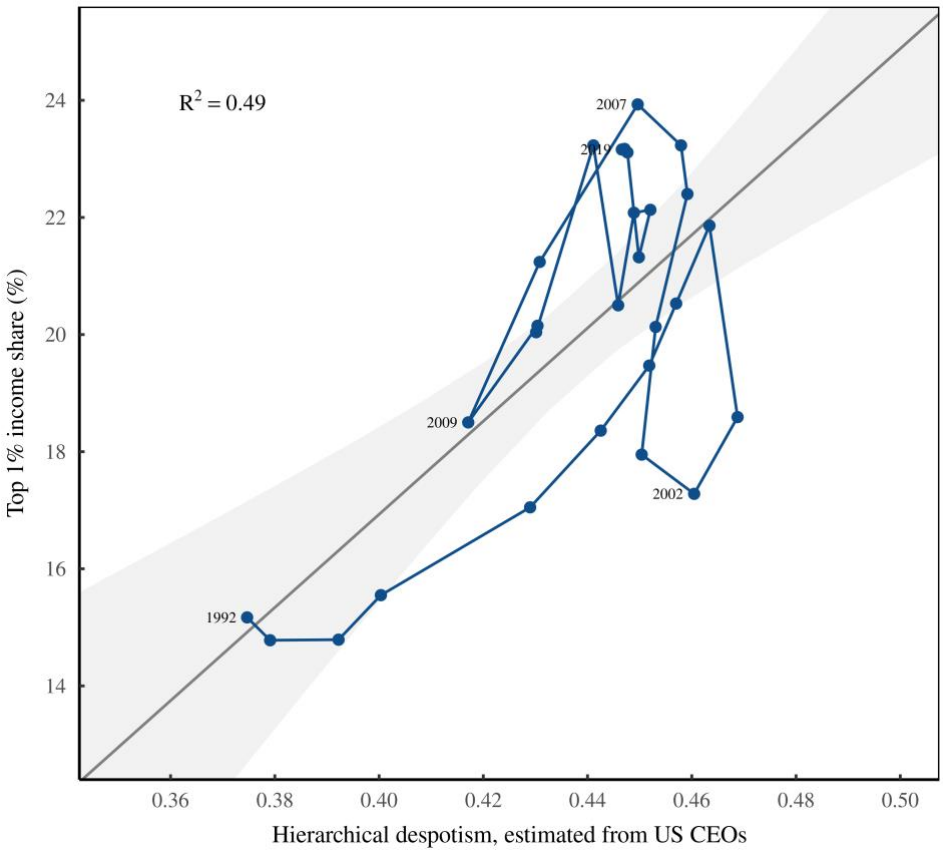

If US CEOs are indeed 'canaries' in the hierarchy, this evidence suggests that rising US inequality has been driven by growing despotism within firms. Ultimately, I would like to test this incendiary idea directly by peering into corporate hierarchies. But since big corporations are unlikely to open up their payroll structure anytime soon, we are forced to further test this idea using a more indirect route. On that note, let us return to the modelling work of Herbert Simon and Harold Lydall.

Returning to the Simon-Lydall model

In the 1950s, Simon and Lydall both used a simple model of hierarchy to explain the power-law behavior of top incomes. Simon showed how hierarchy could explain why CEO pay scales with firm size. And Lydall demonstrated that hierarchy could create a power-law distribution of income.

The key feature of the Simon-Lydall model is the 'span of control', which is assumed to be constant. The 'span' determines how many direct subordinates each superior controls. If the span is constant throughout the group, we get hierarchies that look like the ones shown in Figure 11. A large span of control creates a 'flat' hierarchy. A small span of control creates a 'steep' hierarchy.

Figure 11: The Simon-Lydall model of hierarchy.

In the Simon-Lydall model, hierarchies are assumed to have a constant span of control. A large span creates a 'flat' hierarchy (left). A small span creates a 'steep' hierarchy. For visualization purposes, I show here the actual chain of command within each hierarchy. However, the Simon-Lydall model only simulates aggregate membership within each rank. For model equations, see the Appendix.

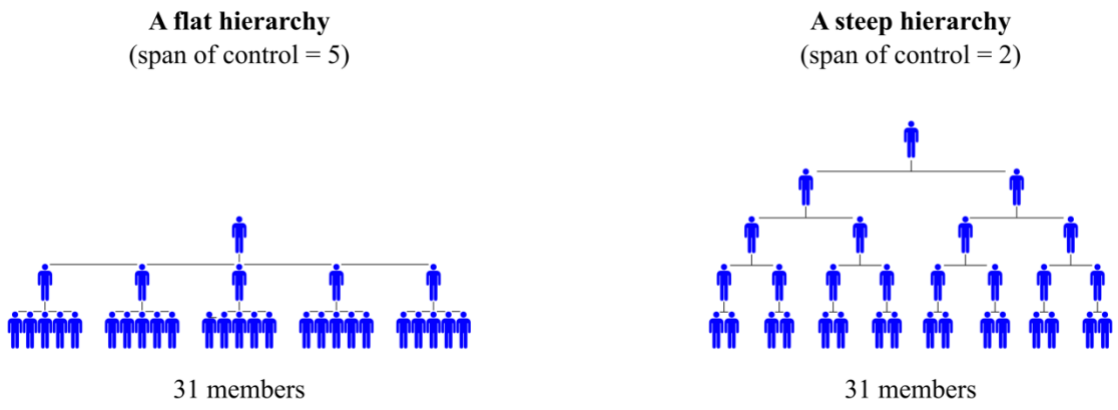

The second key element of the Simon-Lydall model is that income increases exponentially with hierarchical rank. Merge this exponential function with the exponential behavior of the chain of command, and out pop power laws. In what follows, I generalize the Simon-Lydall model to understand how hierarchy affects the distribution of top incomes.

Unlike Simon and Lydall (who used analytic methods), I will build a numerical model. The model starts not with hierarchies, but with the size distribution of firms. Empirical evidence suggests that firm sizes are distributed according to a power law (Axtell, 2001). Based on this observation, I simulate a size distribution of firms by drawing random numbers from a discrete power-law distribution. The simulation is designed to roughly match the size distribution of firms in the United States.

The next step is to use the Simon-Lydall model to give each firm a hierarchical structure. Each individual in the firm is assigned a hierarchical rank, and from this rank we calculate their hierarchical power. (For the model equations, see the Appendix.)

I then model individual income as a function of hierarchical power. To make the model realistic, I introduce stochastic 'noise' into the power-income relation:

$$\text{income} = \text{noise} \cdot (\text{hierarchical power})^{D}$$

The output of the model is a simulated distribution of income. What we want to understand, from the model, is how the degree of hierarchical despotism, D, affects the distribution of top incomes.

Figure 12 shows my results. I have plotted here the distribution of income (using a log histogram) for three iterations of the hierarchy model. Each iteration uses a different value for D. As expected, the model produces a power-law distribution of top incomes, evident as the straight line in the right tail. (Note that when D is small, the income 'noise' dominates the distribution of income, so we do not get a power law.)

Figure 12: In a model of hierarchy, increasing hierarchical despotism fattens the income-distribution tail.

This figure shows results from my implementation of the Simon-Lydall model of hierarchy. In the model, income is assumed to scale with hierarchical power, where the scaling rate is D (a rate which I call the 'degree of hierarchical despotism'). Varying D changes the distribution of top incomes. A larger value of D causes the distribution tail to get 'fatter'. For sources and methods, see the Appendix.

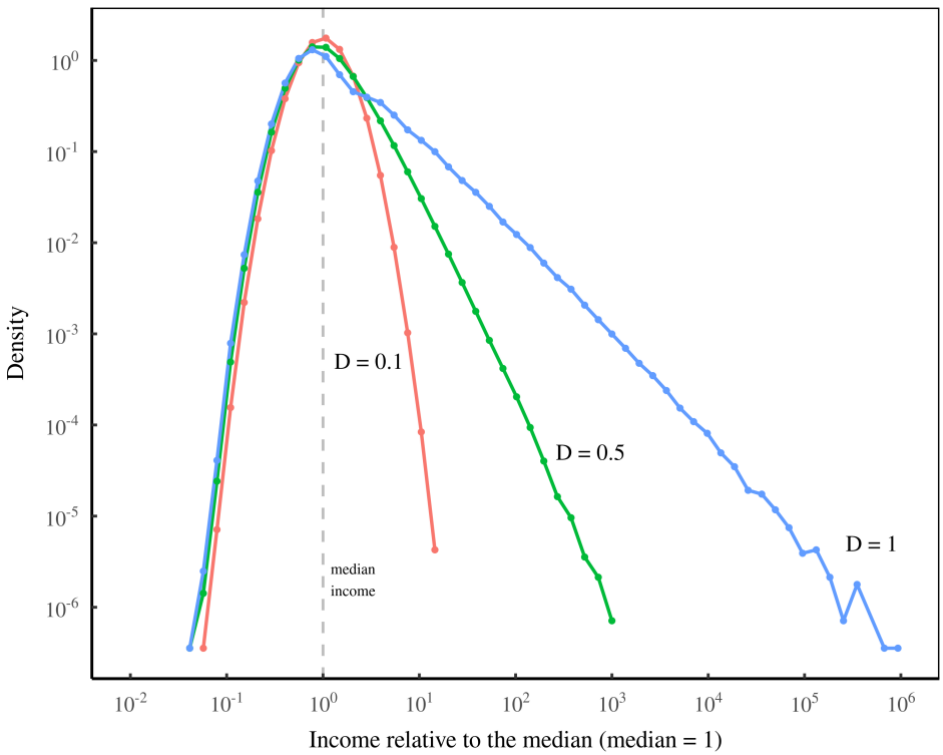

What we are interested in is how the distribution of top incomes is affected by hierarchical despotism. On that front, the results are clear. Increasing hierarchical despotism 'fattens' the distribution tail. In short, it makes the rich get richer in a highly uniform way.

To summarize the evidence thus far, we know the following:

1. The United States has grown more unequal over the last 4 decades (Fig. 2);
2. This growing inequality occurred via a 'fattening' of the income distribution tail (Fig. 4);
3. Growing inequality is associated with a dramatic increase in US CEO pay (Fig. 7);
4. Like the redistribution of top incomes, the pay increases of US CEOs have an underlying uniformity: the rate at which income scales with hierarchical power seems to have increased (Fig. 9);
5. This increasing 'hierarchical despotism' among US CEOs correlates with rising US inequality (Fig. 10), suggesting that US hierarchies have become more despotic.
6. When we put changing hierarchical despotism into a model of hierarchy, we find that it produces a 'fattening' of the income distribution tail (Fig. 12).

All in all, this evidence strongly hints that hierarchy lies at the root of US income redistribution. But perhaps the US is a unique case. To test this possibility, the last step of the puzzle is to see if the hierarchy model can explain the redistribution of income observed across countries.

Recall from Figure 5 that across a wide swath of countries, greater inequality is associated with a smaller power-law exponent among top incomes. Figure 13 replots this data in grey. On top of the empirical data, I plot the trend produced by the hierarchy model. Each colored point represents a model iteration, with color indicating the degree of hierarchical despotism. As we ramp up despotism, the hierarchy model cuts through the middle of the path tracked by real-world countries.

Figure 13: Changing the degree of despotism within modelled hierarchies reproduces international trends in income redistribution.

Grey lines show the empirical trend within countries — the top 1% share of income plotted against the power-law exponent of top incomes. (The empirical data is replotted from Figure 5.) Colored points show iterations of the hierarchy model. By varying the degree of hierarchical despotism within hierarchies, the model reproduces the trend observed across countries. This result suggests that the redistribution of income consists largely of a change in hierarchical despotism. For sources and methods, see the Appendix.

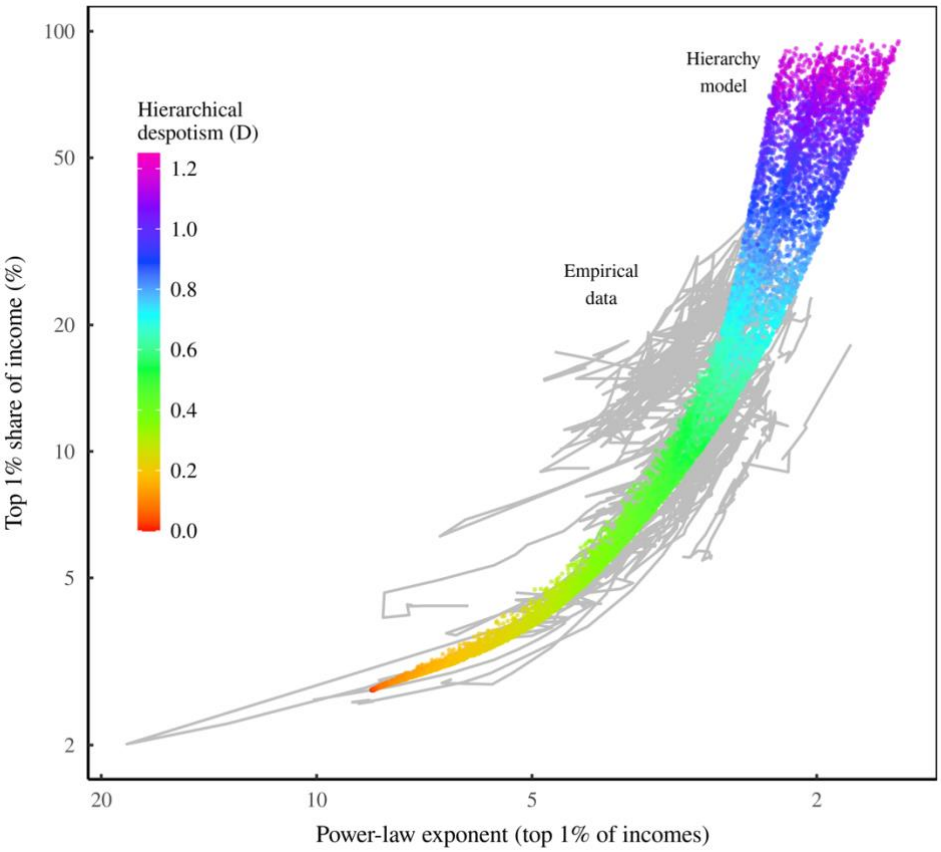

Having noted the model's success, there are a few caveats. First, the model cannot reproduce the low levels of inequality observed in countries like Soviet-era Bulgaria (bottom left of Figure 13). That is because even when we remove all returns to hierarchical rank, there is still income 'noise', which generates inequality. We could change this noise if we desired. But to keep the model as simple as possible, I leave the noise function constant.

Second, the hierarchy model assumes a constant size distribution of firms, similar to the distribution found in the United States. In the real world, the firm size distribution varies both across countries and across time within countries. (See Fix, 2017 for details.) A more complex model could incorporate this firm-size variation.

Finally, in the Simon-Lydall model, the span of control is a free parameter. In the model used here, I let the span vary randomly between 1.2 and 13 — a range consistent with what we know from case studies of hierarchy. (See the appendix in Fix, 2019b for a review.) In the real world, we expect the span of

control to vary between firms and possibly between societies. Such patterns could be incorporated into a more complex model. That said, the span of control has a weak effect on inequality — far weaker than the effect of hierarchical despotism. (See Figure 14.)

To summarize, my model of hierarchy is highly stylized, neglecting many elements of the real world. But its purpose is not to be ultra-realistic, but instead, to isolate the effects of hierarchical despotism. And these effects are clear — increasing hierarchical despotism makes the rich get richer in much the same way as they do in the real world.

Conclusions

Despite the complexities of human life, the distribution of top incomes follows a remarkably uniform pattern. To a first approximation, top incomes are distributed according to a power law. And when income gets redistributed, this power law changes. In short, it seems that we can model the rich getting richer with a single parameter — the power-law exponent α. Such simplicity deserves an explanation.

The reason top incomes follow a uniform pattern, I have argued, is not because income has an ultimately simple cause. Instead, it is because the complex forces that shape income pass through a ubiquitous feature of human organization: *hierarchy*. Thus, I propose that hierarchy is a proximate cause of both the distribution of top incomes, and the uniformity with these incomes get redistributed when the rich get richer.

We have known since Lydall's work in the 1950s that hierarchy can produce a power-law distribution of top incomes. The more complex model used here confirms Lydall's result. I also find that by varying the rate that income increases with hierarchical rank, we vary the distribution of top incomes in much the same way as we observe in the real world. This result suggests that growing inequality is caused by a redistribution of income within hierarchies. Importantly, evidence from CEOs points at the same trend — namely, that growing inequality is associated with hierarchies becoming more 'despotic'.

Appealing to hierarchy, I have admitted, does not explain the root cause of inequality. To do that, we would need to explain why income within hierarchies scales the way it does (something that I do not attempt here). So, in a sense, the hierarchy model of income merely kicks the causal can: it explains one parameter (the power-law exponent of top incomes) in terms of another parameter (the degree of despotism within hierarchies).

Still, I consider that progress. It suggests that we can better understand the causes of inequality by studying the command structure of firms.

Acknowledgements

This work was supported in part by the following individuals: Pierre, Norbert Hornstein, Rob Rieben, Tom Ross, James Young, Tim Ward, Mike Tench, Hilliard MacBeth, Grace and Garry Fix, John Medcalf, Fernando, Joe Clarkson, Michael Defibaugh, Steve Keen, Robin Shannon, and Brent Gulanowski.

Appendix

Source data and code for this paper are available at the Open Science Framework: https://osf.io/h98gn/.

Top income shares

Data for top income shares comes from the World Inequality Database (WID). For the long-term trend in US inequality (Fig. 2), I use the average of series sfiinc992t and sfiinc999t. These series are the closest to the measurements presented in Piketty (2014). International data (Fig. 5) is from WID series sptinc992j.

US income density

To estimate the density function for the US distribution of income (Fig 4), I use income threshold data from series WID tfiinc999t. This series reports the income thresholds for various income percentiles. From these thresholds, I first construct the cumulative distribution of US income. Then I take the derivative of this function to estimate the density curve.

Estimating power-law exponents

To estimate the power-law exponent of the top 1% of incomes, I use the method outlined in Virkar & Clauset (2014). They describe a maximum-likelihood function for fitting power-laws to binned data. The required data is:

1. bin thresholds;
2. counts within each bin.

The WID series tptinc992j provides the needed data. It reports income thresholds for various income percentiles. I use the various percentiles as the 'bins'. The percentile income thresholds are therefore the bin thresholds. And the bin count is simply the income percentile itself (i.e., the portion of the population it represents).

The caveat is that any data can be 'fitted' with a power-law exponent. But this does not mean that the data itself is distributed according to a power law.

US CEO pay ratio

Data for the US CEO pay ratio (Fig. 7) is from the Economic Policy Institute (Mishel & Wolf, 2019). I have plotted data in which stock options are measured using 'realized gains'. For why this is the most appropriate way to measure stock-option income see Hopkins & Lazonick (2016).

Relative income vs. hierarchical power

Data for the relative income within hierarchies (Fig. 6) is from a variety of sources:

- Case-Study Firms: Data is from Audas, Barmby, & Treble (2004); Baker, Gibbs, & Holmstrom (1993); Dohmen, Kriechel, & Pfann (2004); Lima (2000); Morais & Kakabadse (2014); Treble, Van Gameren, Bridges, & Barmby (2001). For details about these studies, see the appendix in Fix (2019b).
- CEOs: The data covers the years 2006–2019 and includes CEOs across many countries (but mostly within the US). CEO pay data is from Execucomp, series TOTAL_ALT2. I estimate the CEO's hierarchical power from firm size — Compustat series EMP. I plot, in Fig. 6, the CEO's income relative to the average employee. I estimate average income in the firm by dividing employment expenses (Compustat series XLR) by firm employment. (Compustat series EMP). For more details, see Fix (2020).
- Note that the CEO data is not strictly comparable to the other series in Fig. 6 because it measures pay relative to the firm *average*. All other series, however, measure pay relative to the average in the *bottom rank* of the hierarchy.
- US military: Data is from annual demographics reports (*Demographics: Profile of the Military Community*) between 2010 and 2019. I exclude warrant officers from the data. I calculate the pay within each rank as the average of the minimum and maximum pay by years of experience. For details, see Fix (2019a).

Hierarchical despotism of US CEOs

The CEO data used in Figures 9 and 10 is slightly different than the CEO data used in Fig. 6. For one thing the Fig. 9-10 includes only US CEOs. But more importantly, the Fig. 9-10 data measures CEO pay using Execucomp series TDC1, rather than series TOTAL_ALT2. The latter series offers a better accounting of stock-option income (using realized gains). But it begins in 2006. In contrast, series TDC1 uses the (more dubious) Black-Scholes method to estimate stock option income. However, data for TDC1 extends back to 1992.

Hierarchy model

The hierarchy model used in this paper is based on equations derived independently by Herbert Simon (1957) and Harold Lydall (1959). In this model, hierarchies have a constant span of control. We assume that there is one person in the top rank. The total membership in the hierarchy is then given by the following geometric series:

$$N_T = 1 + s + s^2 + \cdots + s^{n-1} \qquad (10)$$

Here n is the number of ranks, s is the span of control, and N_T is the total membership. Summing this geometric series gives:

$$N_T = \frac{1 - s^n}{1 - s} \qquad (11)$$

In my model of hierarchy, the input is the hierarchy size N_T and the span of control s. To model the hierarchy, we must first estimate the number of hierarchical ranks n. To do this, we solve the equation above for n, giving:

$$n = \left\lfloor \frac{\log[1 + N_T(s-1)]}{\log(s)} \right\rfloor \qquad (12)$$

Here $\lfloor\rfloor$ denotes rounding down to the nearest integer. Next we calculate N_1 — the employment in the bottom hierarchical rank. To do this, we first note that the firm's total membership N_T is given by the following geometric series:

$$N_T = N_1\left(1 + \frac{1}{s} + \frac{1}{s^2} + \cdots + \frac{1}{s^{n-1}}\right) \qquad (13)$$

Summing this series gives:

$$N_T = N_1\left(\frac{1 - 1/s^n}{1 - 1/s}\right) \qquad (14)$$

Solving for N_1 gives:

$$N_1 = N_T\left(\frac{1 - 1/s}{1 - 1/s^n}\right) \qquad (15)$$

Given N_1, membership in each hierarchical rank h is:

$$N_h = \left\lfloor \frac{N_1}{s^{h-1}} \right\rfloor \qquad (16)$$

Sometimes rounding errors cause the total employment of the modeled hierarchy to depart slightly from the size of the original input value. When this happens, I add/subtract members from the bottom rank to correct the error.

Once the hierarchy has been constructed, income (I) is a function of hierarchical power:

$$I = N(\overline{P}_h)^D \qquad (17)$$

Here D is the 'degree of hierarchical despotism' — a free parameter that determines how rapidly income grows with hierarchical power. N is statistical noise generated by drawing random numbers from a lognormal distribution. (The noise function generates inequality equivalent to a Gini index of about 0.2.) $\overline{P}_h$ is the average hierarchical power (per person) associated with rank h. It is defined as:

$$\overline{P}_h = 1 + \overline{S}_h \qquad (18)$$

where $\overline{S}_h$ is the average number of subordinates per member of rank h:

$$\overline{S}_h = \sum_{i=1}^{h-1} \frac{N_i}{N_h} \qquad (19)$$

The model is implemented numerically in C++, using the Armadillo linear algebra library (Sanderson & Curtin, 2016). For R users, I have created R functions implementing the model, available at Github:

- github.com/blairfix/hierarchy
- github.com/blairfix/hierarchical_power

Size distribution of firms

The input into the hierarchy algorithm is a size distribution of firms generated from a discrete power law distribution with $\alpha = 2$. The resulting distribution is similar to that found in the modern United States. See Fix (2020) for details.

The span of control

In the hierarchy model, the span of control is a free parameter. I let it vary between a low of 1.2 and a high of 13. As Figure 14 shows, this variation has a small effect on the power-law distribution of top incomes. Instead, the effect is dominated by the degree of hierarchical despotism.

Figure 14: In the hierarchy model, the span of control weakly affects the power-law distribution of top incomes.

Points represent different iterations of the hierarchy model, with the degree of hierarchical despotism shown on the horizontal axis. The vertical axis shows the resulting power-law exponent of top incomes. Color indicates the span of control, which has a weak effect on top incomes.

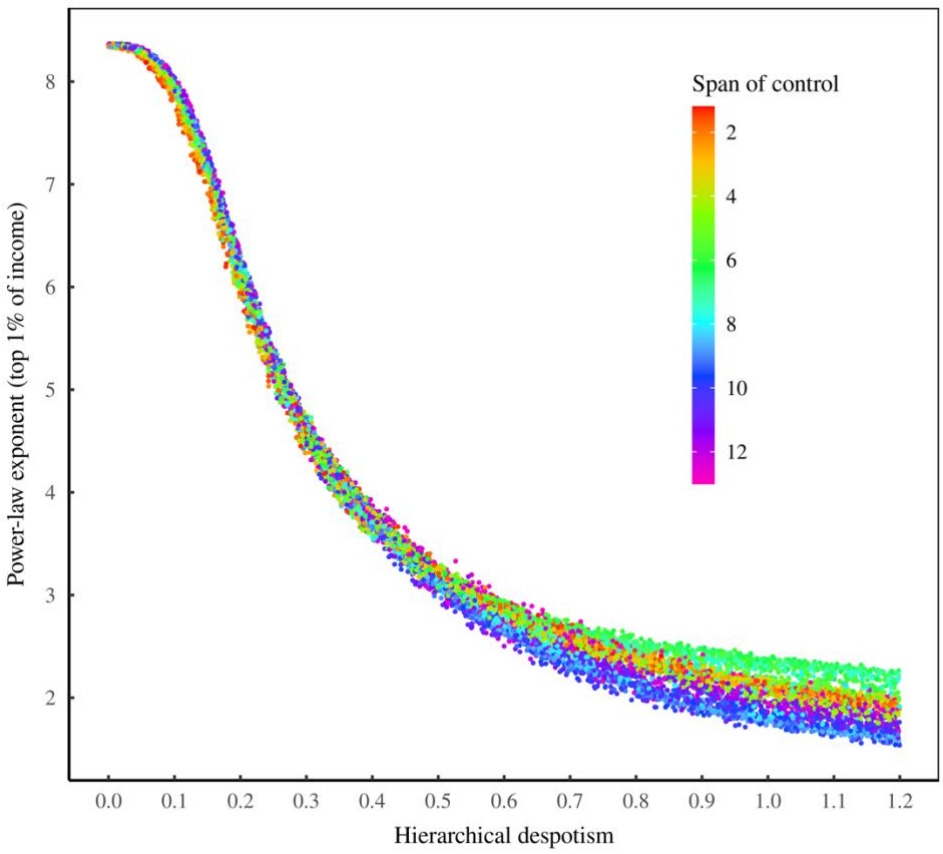

References

Alvaredo, F., Atkinson, A. B., Piketty, T., & Saez, E. (2013). The top 1 percent in international and historical perspective. *The Journal of Economic Perspectives*, *27*(3), 3–20.

Atkinson, A. B. (2017). Pareto and the upper tail of the income distribution in the UK: 1799 to the present. *Economica*, *84*(334), 129–156.

Atkinson, A. B., & Harrison, A. J. (1978). *Distribution of personal wealth in Britain*. Cambridge Univ Pr.

Atkinson, A., & Bourguignon, F. (2001). Income distribution. In *International encyclopedia of the social and behavioral sciences (economics/public and welfare economics)* (pp. 7265–7271). Amsterdam: Elsevier.

Atkinson, A. B., & Piketty, T. (2010). *Top incomes: A global perspective*. New York: Oxford University Press.

Audas, R., Barmby, T., & Treble, J. (2004). Luck, effort, and reward in an organizational hierarchy. *Journal of Labor Economics*, *22*(2), 379–395.

Axtell, R. L. (2001). Zipf distribution of US firm sizes. *Science*, *293*, 1818–1820.

Baker, G., Gibbs, M., & Holmstrom, B. (1993). Hierarchies and compensation: A case study. *European Economic Review*, *37*(2-3), 366–378.

Champernowne, D. G. (1953). A model of income distribution. *The Economic Journal*, *63*(250), 318–351.

Clementi, F., & Gallegati, M. (2005). Power law tails in the Italian personal income distribution. *Physica A: Statistical Mechanics and Its Applications*, *350*(2-4), 427–438.

Coelho, R., Richmond, P., Barry, J., & Hutzler, S. (2008). Double power laws in income and wealth distributions. *Physica A: Statistical Mechanics and Its Applications*, *387*(15), 3847–3851.

Di Guilmi, C., Gaffeo, E., & Gallegati, M. (2003). Power law scaling in world income distribution. *Economics Bulletin*.

Di Muzio, T. (2015). *The 1% and the rest of us: A political economy of dominant ownership*. Zed Books Ltd.

Dohmen, T. J., Kriechel, B., & Pfann, G. A. (2004). Monkey bars and ladders: The importance of lateral and vertical job mobility in internal labor market careers. *Journal of Population Economics*, *17*(2), 193–228.

Fix, B. (2017). Energy and institution size. *PLOS ONE*, *12*(2), e0171823.

Fix, B. (2018). Hierarchy and the power-law income distribution tail. *Journal of Computational Social Science*, *1*(2), 471–491.

Fix, B. (2019a). How hierarchy can mediate the returns to education. *Economics from the Top Down*. https://economicsfromthetopdown.com/2019/12/20/how-hierarchy-can-mediate-the-returns-to-education/

Fix, B. (2019b). Personal income and hierarchical power. *Journal of Economic Issues*, *53*(4), 928–945.

Fix, B. (2020). How the rich are different: Hierarchical power as the basis of income size and class. *Journal of Computational Social Science*, 1–52.

Fix, B. (2021a). Economic development and the death of the free market. *Evolutionary and Institutional Economics Review*, 1–46.

Fix, B. (2021b). The rise of human capital theory. *Real-World Economics Review*, (95), 29–41.

Gibrat, R. (1931). *Les inegalites economiques*. Recueil Sirey.

Hager, S. B. (2020). Varieties of top incomes? *Socio-Economic Review*, *18*(4), 1175–1198.

Hopkins, M., & Lazonick, W. (2016). The mismeasure of mammon: Uses and abuses of executive pay data. *Institute for New Economic Thinking*, *Working Paper No. 49*, 1–60.

Huber, E., Huo, J., & Stephens, J. D. (2017). Power, policy, and top income shares. *Socio-Economic Review*, *0*(0), 1–23. https://doi.org/10.1093/ser/mwx027

Keynes, J. M. (1933). National self-sufficiency. *Studies: An Irish Quarterly Review*, *22*(86), 177–193.

Lima, F. (2000). Internal labor markets: A case study. *FEUNL Working Paper*, *378*.

Lydall, H. F. (1959). The distribution of employment incomes. *Econometrica: Journal of the Econometric Society*, *27*(1), 110–115.

Mishel, L., & Wolf, J. (2019). CEO compensation has grown 940% since 1978: Typical worker compensation has risen only 12% during that time. *Economic Policy Institute*, *171191*. https://www.epi.org/publication/ceo-compensation-2018/

Mitzenmacher, M. (2004). A brief history of generative models for power law and lognormal distributions. *Internet Mathematics*, *1*(2), 226–251.

Morais, F., & Kakabadse, N. K. (2014). The corporate Gini index (cgi) determinants and advantages: Lessons from a multinational retail company case study. *International Journal of Disclosure and Governance*, *11*(4), 380–397.

Pareto, V. (1897). *Cours d'economie politique* (Vol. 1). Librairie Droz.

Piketty, T. (2014). *Capital in the twenty-first century*. Cambridge: Harvard University Press.

Roberts, D. R. (1956). A general theory of executive compensation based on statistically tested propositions. *The Quarterly Journal of Economics*, *70*(2), 270–294.

Sanderson, C., & Curtin, R. (2016). Armadillo: A template-based C++ library for linear algebra. *Journal of Open Source Software*, *1*(2), 26.

Simon, H. A. (1957). The compensation of executives. *Sociometry*, *20*(1), 32–35.

Toda, A. A. (2012). The double power law in income distribution: Explanations and evidence. *Journal of Economic Behavior & Organization*, *84*(1), 364–381.

Treble, J., Van Gameren, E., Bridges, S., & Barmby, T. (2001). The internal economics of the firm: Further evidence from personnel data. *Labour Economics*, *8*(5), 531–552.

Virkar, Y., & Clauset, A. (2014). Power-law distributions in binned empirical data. *The Annals of Applied Statistics*, *8*(1), 89–119.

Author contact: blairfix@gmail.com

SUGGESTED CITATION:
Fix, Blair (2021) "Redistributing Income Through Hierarchy" *real-world economics review*, issue no. 98, 14 December, pp. 58-86, http://www.paecon.net/PAEReview/issue98/Fix98.pdf

You may post and read comments on this paper at https://rwer.wordpress.com/comments-on-rwer-issue-no-98/