Economics from the Top Down

new ideas in economics and the social sciences

How Interest Rates Redistribute Income

Blair Fix

April 16, 2023

> [I]nterest rates are, first and foremost, a distributional variable that affects the income shares of various social groups.
>
> — Louis-Philippe Rochon & Mark Setterfield (2007)

When I read about monetary policy, I have a rule of thumb. Every time I see the phrase *interest rate*, I replace it with the term *wage rate*. Then I ask myself whether the discussion still makes sense. Often, it does not.

The reason I make this substitution is that in conceptual terms, the interest rate and the wage rate are similar: they are both *rates of return*. Wages are the return on employment. Interest rates are the return on credit.

Now, the important thing about rates of return is that when we change them, we are toying with the distribution of income. Hike wages and we send more income to workers. Hike the rate of interest and we send more income to creditors. Sure, the specifics of this redistribution are open for inquiry. But by definition, rates of return are 'distributional variables' — they determine how the income pie gets divvied up.

Back to my word substitution. When it comes to wages, the issue of distribution is typically front and center. That's why talk of a minimum-wage hike prompts businesses (and many economists) to complain about reduced profits. But when creditors hike the rate of interest, talk of income distribution is curiously absent. Instead, we get a barrage of macroeconomic jargon — terms like the 'natural rate of interest' and the 'non-accelerating inflation rate of unemployment'.

Why the discrepancy?

One possibility is that economists know something that we don't. Perhaps they've looked at the evidence and concluded that interest rates have a 'neutral' effect on the distribution of income.

Another possibility is that the macroeconomic jargon is mostly a distraction. In other words, like wages, the rate of interest is a 'distributional variable'. But it's one that mainstream economists prefer to ignore.

So which option is true? In this post, I let the evidence speak for itself. Looking at cross-country evidence, I find that interest rates are decidedly non-neutral. As interest rates rise, three things happen:

1. the interest share of income *increases*;
2. the labor share of income *decreases*;
3. income inequality *increases*.

In short, the evidence suggests that interest rates play a key role in the game of class warfare. And that makes sense. Interest, after all, is a rate of return. And when it comes to divvying up the income pie, rates of return are always zero sum.

How economists learned to ignore the distribution of income

Before we jump into the income-distribution data, it's worth reviewing some history. Among heterodox economists, the distribution of income is a hot-button topic. But among mainstream economists, it remains a secondary concern. Why?

To understand this apathy, we need to retrace the trajectory of economic thought. When the study of political economy got rolling in the 19th century, the distribution of income was front and center. For early political economists, what mattered most was *class-based* income. Why did land owners, capitalists, and workers receive their respective cuts? So crucial was this question that David Ricardo deemed it the 'principle problem in political economy'.

As the 19th century played out, thinkers like Karl Marx and Henry George recognized that distributing income involved obvious conflict. For example, if Alice the worker and Bob the capitalist both want a 60% cut of the pie, only one of them can get their way. In other words, the zero-sum nature of

class-based income necessitates class warfare. Or rather, it necessitates class warfare, supposing that the various classes want more than their current take.

Enter John Bates Clark. Seeing the instability involved in class struggle, Clark wanted to show that conflict was unnecessary. And he found an ingenious way to do it. With a few swift assumptions, Clark 'proved' that in a competitive market, each agent gets back the amount of wealth which that agent created. Thus was born the theory of 'marginal productivity', which tells everyone that they earn what they deserve. The message? In market economies, class struggle doesn't exist.

Of course, Clark's theory is based on assumptions that are obviously false. (For example, you must assume that society produces a *single* commodity.) So the real question is why economists settled on a theory that was clearly wrong.

Here's my take.

First, Clark's theory of marginal productivity told powerful people what they wanted to hear — namely, that the distribution of income is 'just'.

Second, Clark's approached gelled nicely with the emerging obsession with economic growth. In the mid 20th century, economists found that if they treated all of society like a single firm, they could use a production function to 'explain' economic growth.[1] Importantly, this function assumed not only that each class earned their 'marginal product', but that their income shares were *constant*. In other words, when modeling economic growth, economists could treat the distribution of income as a non-issue. And so they did.

Of course, the larger backdrop is that following World War II, humanity massively increased its consumption of resources — a pattern that had nothing to do with economists' models and everything to do with the exploitation of fossil fuels. But like everyone else, economists came along for the ride. And so they got obsessed with economic growth, and learned to ignore the distribution of income. Figure 1 illustrates this ideological shift as it is written in English word frequency.

[1] Actually, the history of economic growth theory is just as embarrassing as the history of Clark's theory of marginal productivity. Like Clark, neoclassical growth theorists realized that to make their theory work, they had to assume that society produced a single commodity. Worse, critics later showed that economists' favorite production function — the Cobb-Douglas function — was a thinly veiled version of an accounting identity used to define the national accounts. In other words, the function 'worked' because it was a tautology. For more details, see 'Economic Growth Theory … Bah Humbug!'.

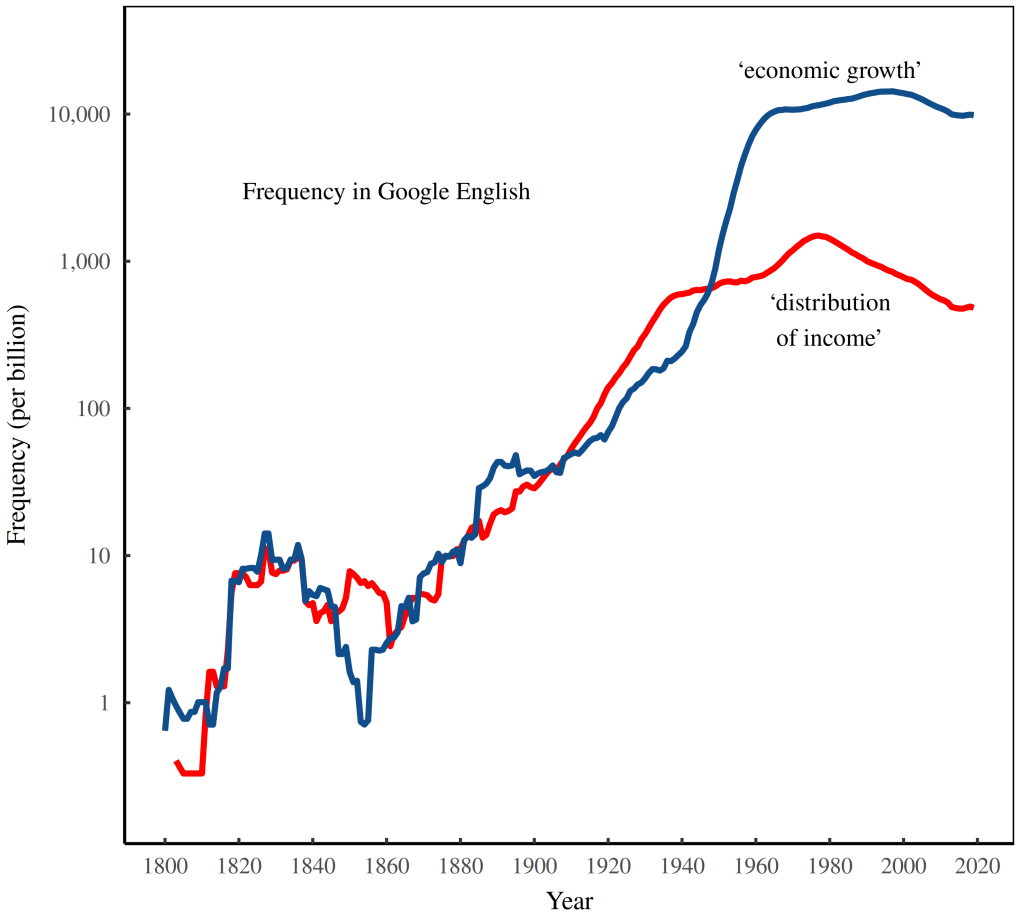

Figure 1: Learning to ignore the distribution of income

Using the Google English corpus, this figure charts the frequency of the phrases 'economic growth' and 'distribution of income'. Until the 1940s, the two phrases were equally popular. But during the post-WWII boom, economists learned to love growth and ignore distribution. Sources and methods

---

Looking to the coming century, my guess is that economists will eventually rekindle their interest in the distribution of income … but only because they will be forced to. As the fossil-fuel era wanes, economic growth will become a relic of the past. And so economists will be dragged back into the business of studying how the income pie is divided. In this post, we'll get ahead of the curve.

The national income pie

Jumping into the study of income distribution, today we have excellent data on how the income pie is divvied up. And that's somewhat ironic.

You see, the income data largely comes from the national accounts, which were designed as tools for measuring economic 'production'. Here's the irony. Because the national accounts are based on monetary transactions, they are a dubious measure of 'production'. But they are a good measure of income.[2]

To study the distribution of income, we do the opposite of what most economists do. Instead of using the national accounts to look at the *size* of the income pie, we measure the pie's *composition*. Typically, that means splitting the pie into five categories of income:

1. employee compensation
2. profit
3. proprietor income
4. rent
5. interest

To give you a sense for the size of each income category, Figure 2 shows the US income split in 2021.

Looking at these income categories, we naturally want to know where they come from. Economists have an answer. They claim that each class represents a 'factor of production'.

This jargon reflects economists' love of equating income with 'production'. I advise you to ignore it. The reality is that the various classes of income are *legal constructs*. In other words, the law defines different forms of property rights. And from these rights stem different categories of income.

For example, people who own corporations by definition earn 'profit'. People who own unincorporated businesses by definition earn 'proprietor income'. People who own physical and/or intellectual property earn 'rent'. People who own debt earn 'interest'. And people who own nothing earn 'wages'.

Are these different types of property rights arbitrary? Yes. But in some sense, rules are always arbitrary. What matters is that these property rules are heavily enforced, which means that they dictate the terms of class warfare.

[2]Income measures from national accounts still have problems. The main one is that because the national accounts ostensibly measure 'production', they exclude income that is deemed 'unproductive'. For this reason, the national accounts exclude capital gains, which are an important source of property income.

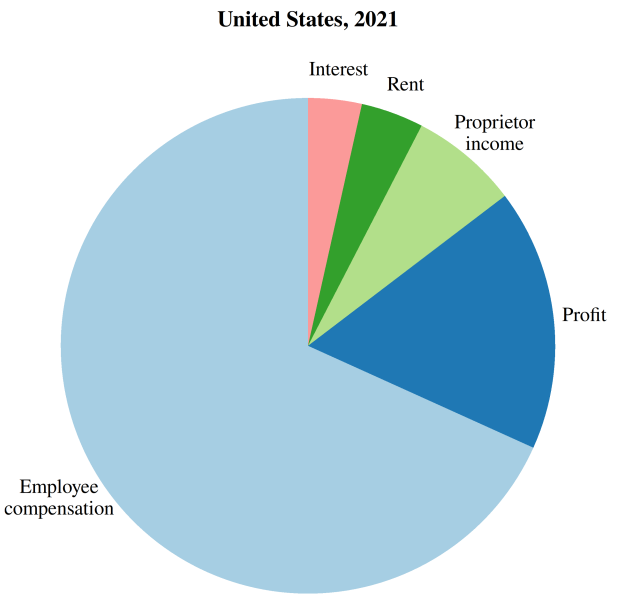

Figure 2: The US income pie in 2021
This chart shows the composition of US national income in 2021. Sources and methods

The interest share of income

Speaking of class warfare, let's talk about the rate of interest. As Rochon and Setterfield put it, interest rates are primarily a "distributional variable that affects the income shares of various social groups." That's a subtle way of saying that interest rates are a tool for waging class war.

On this fact, the evidence is quite clear. But before we get to the data, let's frame the battle. In the national accounts, interest income is tabulated as 'net interest'. (The 'net' part means that when adding up interest payments, statistical agencies subtract interest received from interest paid.)

Returning to Figure 2, we can see that net interest is currently the smallest slice of the US income pie. But don't be misled by this recent data. When we look at US history, we find that the interest share of income is extremely volatile.

Figure 3 tells the story. Here, the red curve shows the interest share of US national income. Over the last 120 years, it varied from a low of 1% (in 1946) to a high of nearly 10% (in 1982). What explains this variation? Well, the trough and peak of the interest share of income correspond to some important

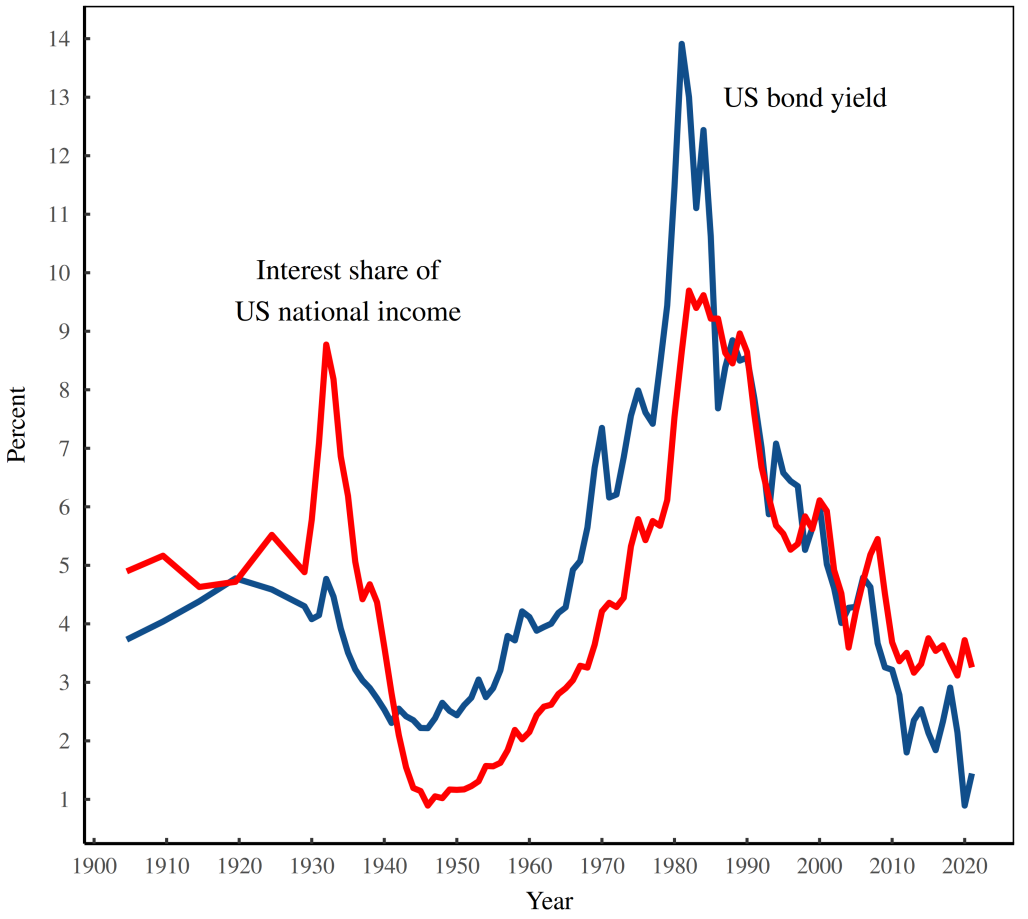

Figure 3: In the United States, the interest share of income rose and fell with the yield on bonds

The red curve shows how the interest share of US national income varied over the last century. The blue curve shows the corresponding change in bond yields, which proxy the rate of interest. Evidently creditors' cut of the pie is determined largely by the rate of interest. Sources and methods

---

geopolitical events (the end of WWII in 1945 and the stagflation crisis of the early 1980s, respectively). But in more abstract terms, the interest share of income is dictated by something simpler: the *rate of return on credit*.

In Figure 3, the blue curve testifies to the connection between the interest share of income and the interest rate. This curve shows the long-term history of US bond yields, which rise and fall with the interest share of income. And bond yields, if you're unfamiliar, are a way of measuring the return on credit — otherwise known as the rate of interest.

Summarizing the evidence, it appears that Rochon and Setterfield are correct to call interest rates a 'distribution variable'. We can think of the rate of interest as a dial for setting the interest share of income.

That said, it's unwise to draw sweeping conclusions from studying a single country. So before we declare interest rates an income-share dial, let's widen our net.

To do that, we'd ideally look at data for a huge sample of countries. But the problem with this approach is that income-share data is not widely reported. (This data scarcity stems from the fact that economists mostly care about the *size* of the income pie, not its *composition*.) To date, I've been able calculate interest-share data for a handful of OECD countries covering the last two decades. This sample is not great. But it's enough for a consistency check.

On that front, Figure 4 shows that the OECD data fits with the pattern found in the United States. Across these countries, the interest share of income rises and falls with the rate of interest.

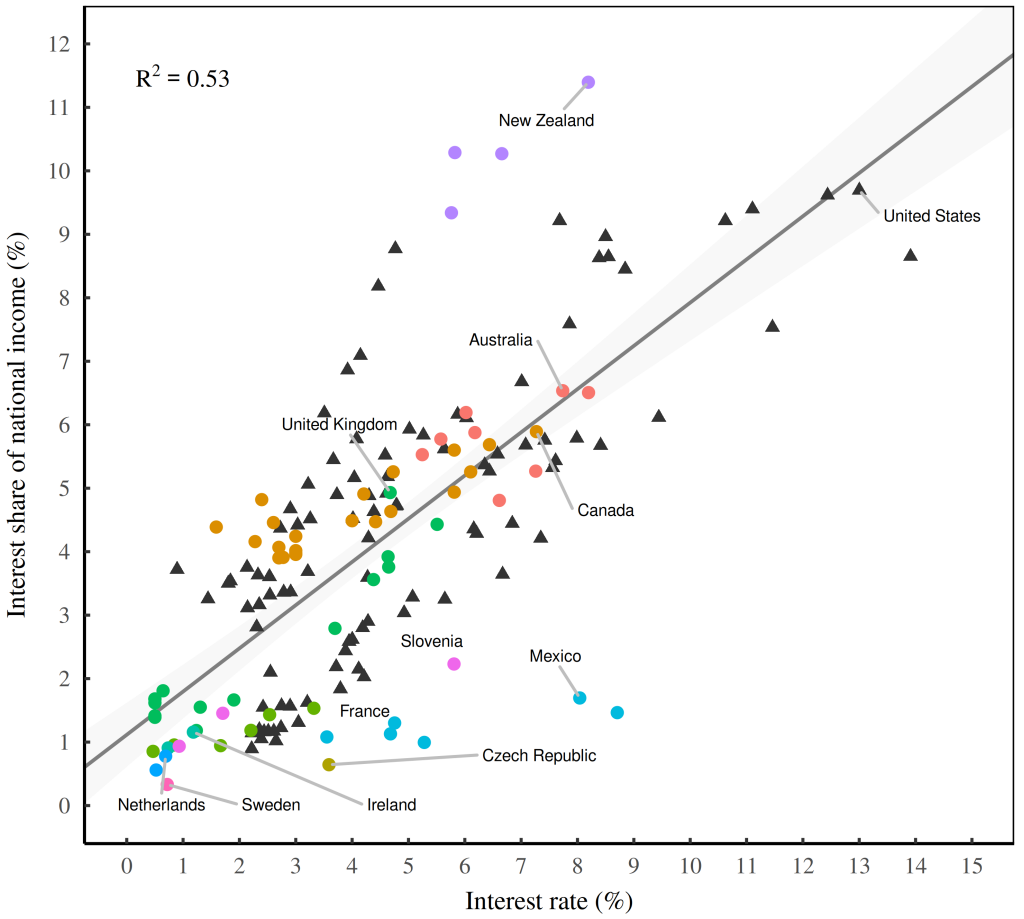

Figure 4: Interest rates control the interest share of income

This figure shows how the interest share of national income (vertical axis) relates to the rate of interest (horizontal axis). The black triangles show US data (plotted previously in Figure 3). Colored points show data from a handful of OECD countries, covering the years 1999 to 2019. Sources and methods

---

My debt, your income

So the interest share of income is controlled largely by the rate of interest. Is that surprising? In some sense, no. If we had to pick a rate of return that set the interest share of income, the interest rate is the obvious candidate. Yet when we reflect on how creditor income is determined, the primacy of interest rates is not obvious.

You see, creditor income is shaped jointly by the rate of interest *and* by the amount of outstanding debt. And taking a cue from individuals, we know that the level of debt can vary immensely. For example, my friend Alice owns $10 of debt. But her friend Bob owns $10 *billion* of debt. Who earns more interest income? Obviously Bob.

The point is that debt liabilities vary far more than the rate of interest. So it should be these liabilities (and not the rate of interest) that determine interest income. At least, that's how it works for individuals. But for some reason, when we switch to countries, the scale of debt liabilities somehow comes out in the wash, leaving the rate of interest as the main dial for setting the interest share of income.

How does that work?

The answer is that at the national level, debt liabilities can't be any old number. These liabilities are tightly coupled to aggregate income. Figure 5 shows the pattern in the US. Over the last two centuries, total US debt grew at about the same rate as nominal GDP (a measure of aggregate income).

So why are debt levels coupled to national income? Well, if you've familiar with how money is created, you'll know that there is a straightforward answer: debt is a *prerequisite* for income.

You see, to have income, there must be money in circulation. And to have money, there must be debt.[3] That's because most money is loaned into existence by commercial banks — created via the magic of double-entry book keeping.

Here's how it works. Starting with nothing, a bank lends you money by inserting digits into your bank account. That's your debt. On the other side of the ledger, the bank records a corresponding credit, causing the operation to look fiscally neutral. But the reality is that money has been created. When you spend your loan, your debt becomes someone else's income. Presto, debt creates income.

[3]Actually, the word 'must' is too strong. As MMT theorists point out, there's no fundamental reason that money must be created with debt. After all, money is an accounting convention. It's rules work however we define them. In other words, if we want money that is created without debt, we could make it so.

Then why don't we?

The catch is that debt-free money would deprive powerful people of prodigious income. You see, in capitalist societies, the power to create money has been largely ceded to private banks, who profit immensely from their power. When banks create money, they attach it to debt, which allows them to charge interest. Sure, we could socialize the power to create money. But bankers would put up a helluva fight.

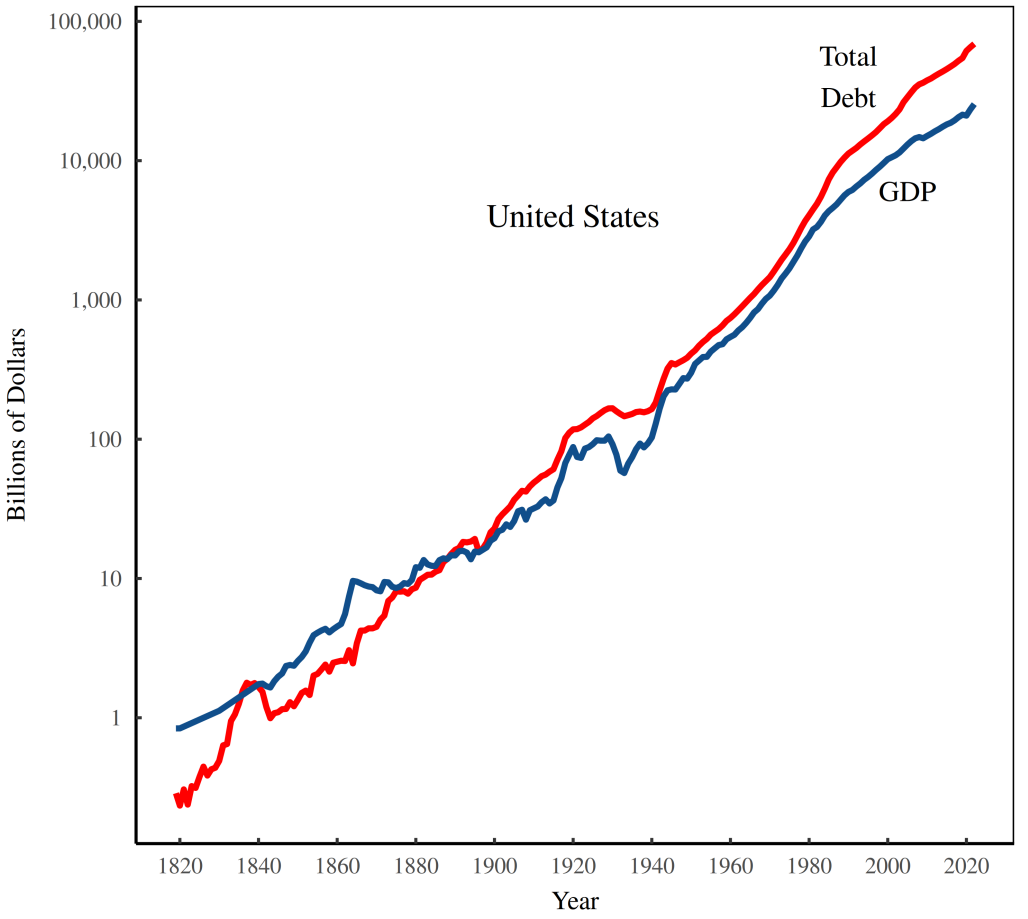

Figure 5: More debt, more income

This figure plots the growth of total US debt (red) and total GDP (blue), both measured in nominal dollars. (Note that the vertical axis uses a logarithmic scale, which makes constant exponential growth appear as a straight line.) The tight connection between the two monetary sums is easy to explain. To have more income (i.e. GDP), you must have more money. And to have more money, you need more debt. Sources and methods

---

With this income creation in mind, let's return to the interest share of income. Because total debt grows at roughly the same rate as total income, it follows that creditors' share of the pie is determined largely by the rate of interest. (For the math, see the appendix.)

Now let's take this thinking one step further. Since debt is a prerequisite for income, we can think of the interest rate as a kind of *royalty on income*. Like all royalties, interest stems from the enforcement of property rights. It's just that in the case of interest, the property concerned is perhaps the most important public good — the right to create money.

The labor share of income

Turning to other forms of income, we know that by definition, when the interest share of income rises, other income shares must *fall*. Let's have a look at these clawbacks.

I'm going to focus on the labor share of income. That's because (1) the labor share of income is an important income category, and (2) I can actually get widespread data for this share.

Labor foibles

So what is the 'labor share of income'. Well, it's obviously the income share paid to laborers. But then who is a laborer? I raise the question because mainstream economic theory taints the data that I'm going to use.

To start with, I think we can agree that wages and salaries are 'labor income'. But what should we do with proprietors — people who are self-employed? I'd prefer to either lump *all* of their income into the 'labor share', or *exclude* all of it.

Mainstream economists, however, do something that makes me cringe; they take proprietor income and *divide* it into a 'labor' component and a 'capitalist' component. The thinking is that a portion of proprietor income stems from working, and another portion stems from owning capital. In other words, economists attribute proprietor income to different 'factors of production'.

This division is ridiculous. Like all remuneration, proprietor income stems from a legal classification, nothing more. Anyway, I rant because the data that I use contains this dubious division of proprietor income. So keep that in mind as we review the results.[4]

Eroding the labor share of income

Returning to interest rates, we know that they are the main dial that controls the interest share of income (Figure 4). And since a growing interest share of income cuts into the other pieces of the pie, it seems plausible that interest rates might affect the labor share of income.

[4]The labor share data comes from ILOSTAT, a database run by the International Labour Organization. It's especially annoying that an organization dedicated to advancing workers rights is using accounting concepts cooked up by neoclassical economics.

Now, when I say 'affect', this implies causation. But in strict terms, I'm going to looking for an 'association' between the rate of interest and the labor share of income. We can sort out causation later.

Turning to the data, I've assembled an international dataset that compares the labor share of income (within each country) to the rate of interest (again, within each country). Figure 6 shows the resulting pattern across countries. Our hunch seems to be correct. Higher interest rates are associated with a lower labor share of income.[5]

Once more, it seems that Rochon and Setterfield are on the mark when they call interest rates a 'distributional variable'. Not only do interest rates affect the interest share of income, they also seem to affect the labor share of income.

In fact, the relation may be almost mechanical. If we assume that a growing interest share of income cuts directly into the labor share, then we automatically get the kind of pattern shown in Figure 6. (For details, see the appendix.)

[5]Readers sometimes complain that my high-level discussions of data are too abstract. Where is the *human* element? Well, the truth is that it's difficult to mix narratives about specific people and places with sweeping quantitative analysis that captures all places all a once. But to give you a glimpse of the human element behind the labor-share data, let's look at two extremes.

On the high end is Iceland. In 2006, Icelandic laborers received about 73% of national income. That makes sense. Iceland is well-known for its Nordic social democracy — an environment that is obviously a boon for workers.

On the low end is Qatar. In 2011, Qatari laborers received just 15% of national income. Again, this (pitiful) cut makes sense. Qatar is infamous for being a neo-feudal petrostate that relies on a huge army of migrant workers. And by 'huge', I mean that these migrant workers constitute about 95% of the Qatari workforce. In other words, we're talking about a country built on a foundation of indentured servitude. Unsurprisingly, this environment is bad for workers.

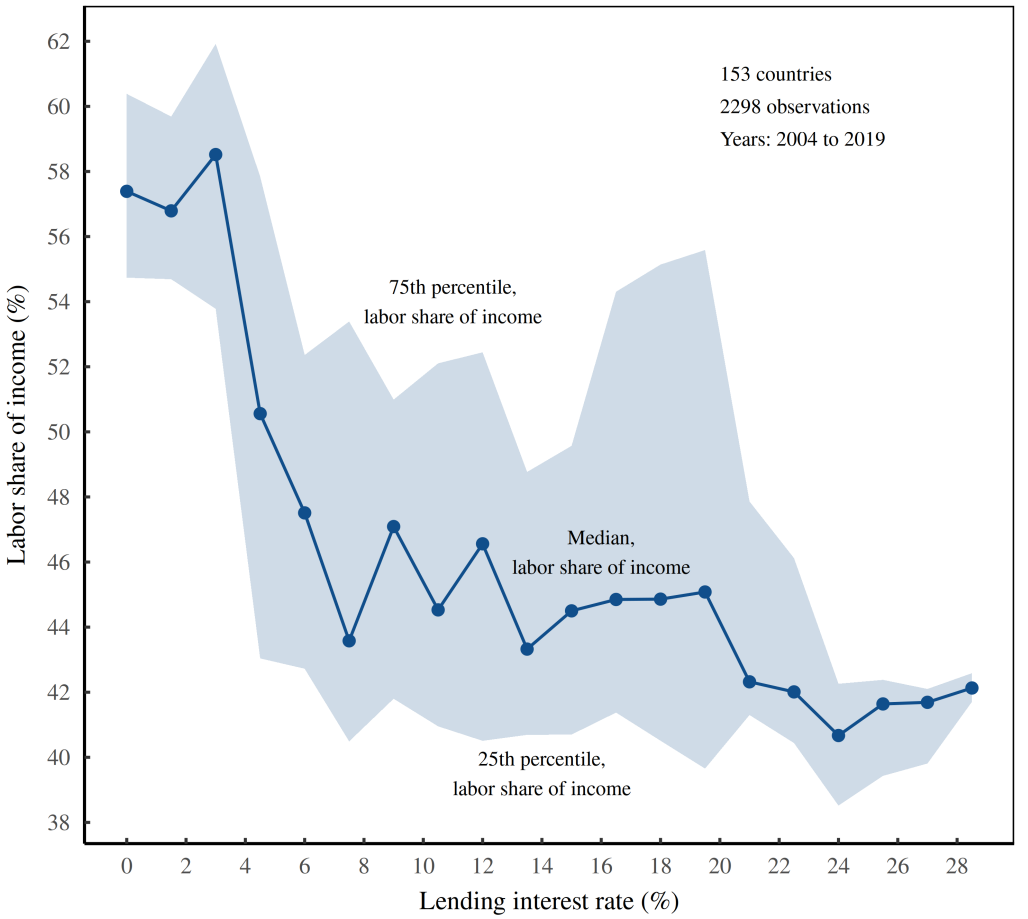

Figure 6: Across countries, higher interest rates are associated with a lower labor share of income

Using international data, this figure measures how interest rates relate to the labor share of income. Here's how the analysis works. I've cobbled together a dataset that covers 153 countries over the years 2004 to 2019. On the horizontal axis, I've plotted the interest rate within each country. The vertical axis shows the labor share of income (labor income as a percentage of GDP). To illustrate the trend across countries, I've binned the data by interest rate. Each point shows the midpoint of a bin. The blue line marks the median labor share of income. The shaded region shows the middle 50% of data. Note that each bin contains at least 10 country-year observations. Sources and methods

---

Income inequality

Chugging along, let's look at one last form of income distribution — the income share of the top 1% of earners.

At first glance, the top 1% share has little to do with the class-based income that I've discussed so far. After all, the legal category of your income tells us nothing about its size (large or small). And since income size and income class are seemingly unrelated, there's no reason to suspect that interest rates affect income inequality.

The catch is that income size and income class *are* related. And because they're related, the rate of interest *does* affect income inequality. Let's see how it works.

The interest-to-wage ratio

When it comes to interest income, the old adage is correct: it takes money to make money. And so it's a good bet that if you earn more income, you'll have more money to invest, hence you'll earn more interest.

To flesh this thinking out, let's look at the relation between income size and income composition. Specifically, I'm going to calculate something that I call the *interest-to-wage income ratio*. This ratio takes the income that an individual earns from interest and divides it by the income that they earn from wages/salaries:

$$\text{interest-to-wage ratio} = \frac{\text{interest income}}{\text{wage income}}$$

For most people, the interest-to-wage ratio is small, since their wages dwarf their interest income. But for a few individuals, interest is a significant source of money. As it turns out, these individuals also happen to be top earners.

Figure 7 shows the pattern in the United States. Here, the horizontal axis ranks Americans by their income percentile. The vertical axis plots their interest-to-wage ratio. You can see that as we approach the upper crust of earners, interest income explodes.

Now in direct terms, Figure 7 tells us about interest income. But indirectly, it tells us about the distribution of *credit*. You see, as a first approximation, the rate of interest doesn't vary much between investors. And so the trend in Figure 7 must be driven by the ownership of credit. Therefore, we can conclude that top earners own far more credit than the rest of us.

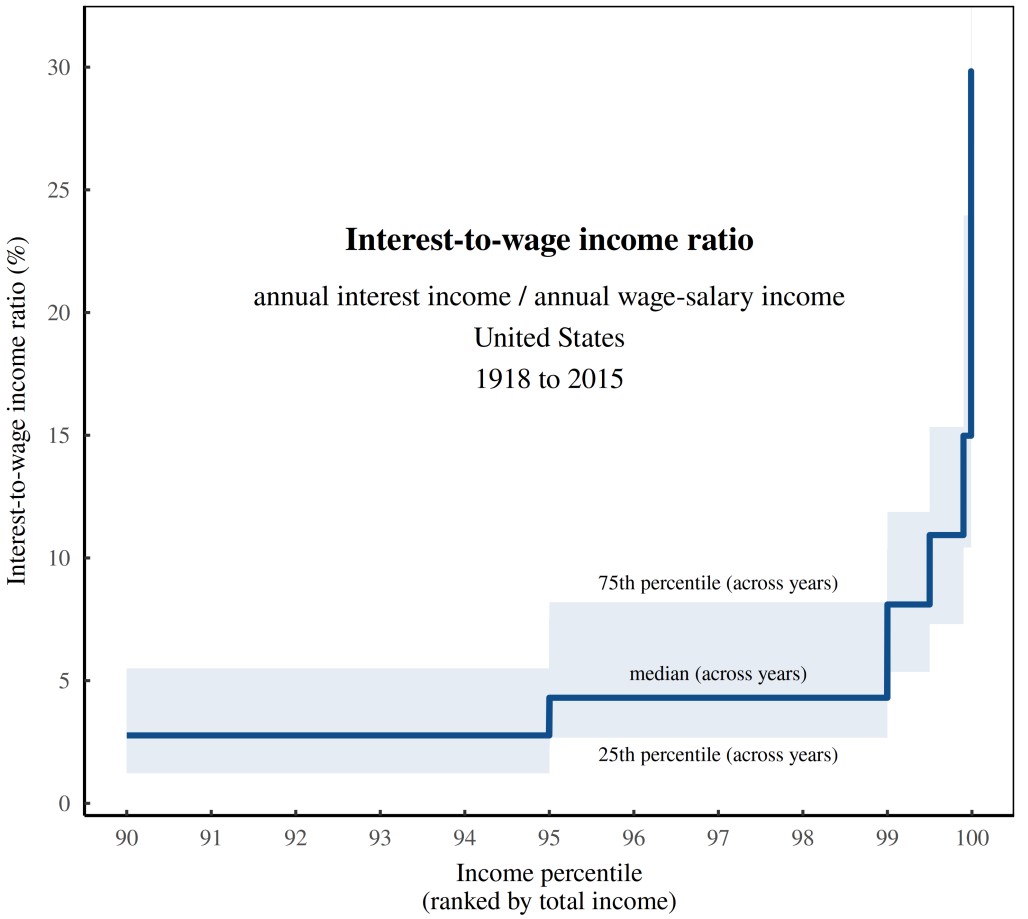

Figure 7: The interest-to-wage ratio in the United States

This figure measures the composition of Americans' income as a function of their income percentile. The horizontal axis shows income percentile, ranked by total income. The vertical axis shows the interest-to-wage ratio — the ratio between annual interest income and annual wage/salary income. The steps in the blue line indicate different percentile bins. The line indicates the median interest-to-wage ratio across the years 1918 to 2015. The shaded region indicates the middle 50% of data. Sources and methods

---

With this credit distribution in mind, let's think about what happens when we raise the rate of interest. Doing so sends money to people who own more credit. And the people who own more credit also happen to be top earners. And so what do interest-rate hikes do? They *send money to the rich*.

(Actually, the math is a bit more complicated than I'm making out. See the appendix for details.)

To summarize, there's good reason to suspect that higher interest rates might worsen inequality. So let's see if they do.

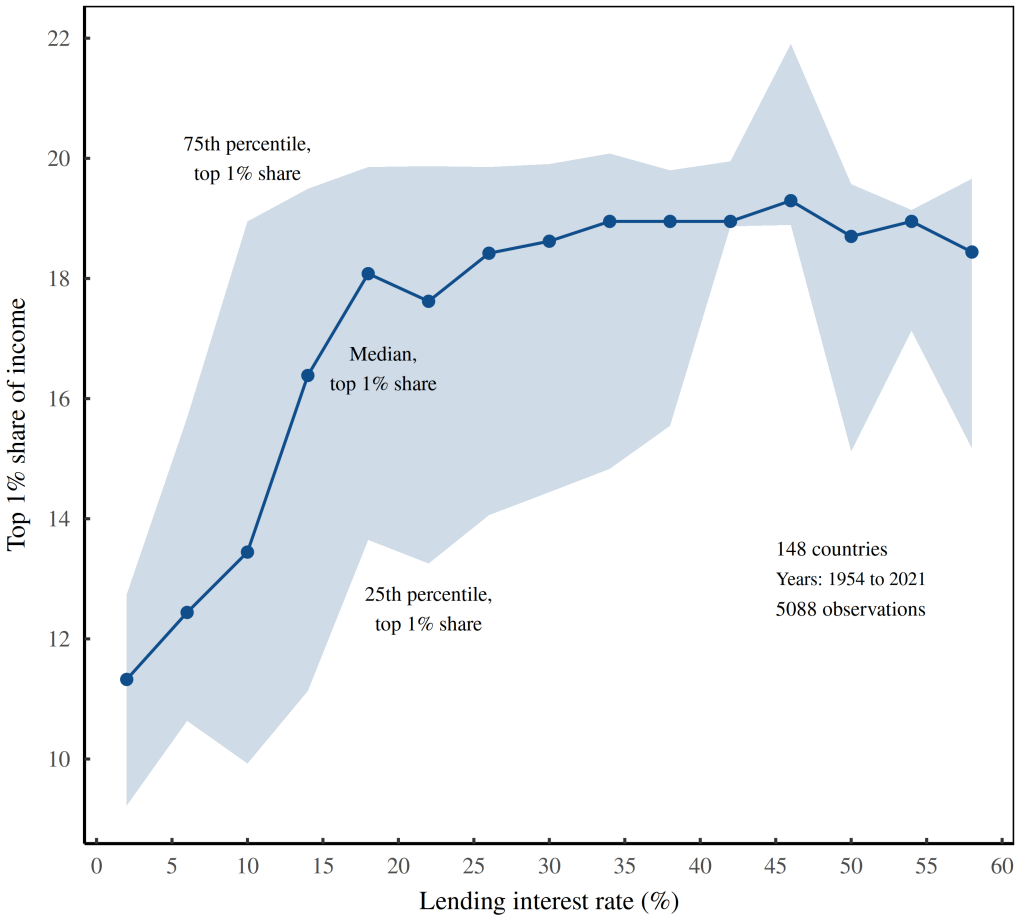

Figure 8: Higher interest rates are associated with greater income inequality

This figure looks at the international relation between the lending interest rate (horizontal axis) and the top 1% share of income (vertical axis). The dataset covers 148 countries, observed over the years 1954 to 201. To illustrate the trend, I've binned the data by interest rate. Each point shows the midpoint of an interest-rate bin. Blue line shows the median top 1% share of income within each bin. The shaded region shows the middle 50% of data. Note that each bin contains at least 10 country-year observations. Sources and methods

---

Interest rates and the top 1% share of income

With inequality in mind, let's look at Figure 8. Here, I've gone to the World Inequality Database and downloaded all of their data for the top 1% share of income within countries. Then I've merged this data with interest-rate time series, and plotted the result. As you can see, the pattern is fairly obvious. Across countries, higher interest rates are associated with greater income inequality.

Now, the caveat here is that the inequality pattern is non-linear. As interest rates increase, the top 1% share grows — but only to a point. Why? Well, if we assume that changing interest rates operate on a fixed distribution of wages and credit, then this type of non-linear pattern is exactly what we expect. (For details, see the appendix.)

Specifics aside, it's clear that income inequality is related to the rate of interest. So again, Rochon and Setterfield are correct to label interest rates a 'distributional variable'.

The average sabotage

All in all, the evidence shouts at us that the rate of interest affects the distribution of income. Of course it's nice to have this empirical confirmation. But then again, the 'distributional' nature of interest rates was never in doubt. As a rate of return, interest rates are by definition a 'distributional variable'. Or in more incendiary terms, interest rates are a weapon of class warfare.

Now, this language might sound hyperbolic, but I think it's accurate. Contrary to what neoclassical economists claim, there are no neutral market forces that allocate income in proportion to productivity. Instead there are only *ideas* and the *power* to implement them. In other words, people have ideas about what their income should be (and also what other people's income should be). And they have the power (or lack thereof) to make these ideas a reality. That's it.

So viewed through the lens of power, rates of return are, by definition, outcomes of social conflict. Still, this vantage point doesn't get us very far in understanding real-world outcomes. It's like saying that animal behavior is an outcome of evolution. It's true. But it's only the starting point for a scientific explanation.

On that front, how should we study the class struggle involved in interest income? For Marx, the answer was that interest is about inter-capitalist competition. Interest payments, Marx argued, transfer to the 'money-capitalist' some of the profits received by the 'industrial capitalist'.

Although I admire the simplicity of this approach, it implies that we can cleanly divide between 'industrial' and 'money' capitalists. And by extension, Marx's view implies that profit is earned from 'productive' activities, whereas interest stems from the unproductive ownership of money.

Faced with this Marxist division, capital-as-power theorists Jonathan Nitzan and Shimshon Bichler think that it misses the point. Looking at capital, they argue that *all of it* is unproductive. That's because capital is nothing but the quantification of property rights. And property rights, in turn, are inherently negative; they are an institutional act of exclusion. So in that sense, profit and interest both stem from enforced exclusion — what Nitzan and Bichler call *strategic sabotage*.

Viewed this way, profit and interest represent different tactics for inflicting sabotage:

> [I]nterest on debt represents *average* sabotage, while profit on equity denotes *differential* sabotage.
>
> (Nitzan and Bichler, 2009; emphasis added)

Unpacking this claim, the idea is that when you purchase equity in a specific company, you are investing in that company's ability to wield property rights for your benefit. And because you hope that this company will beat the average rate of return, you are investing in 'differential' sabotage.

With debt, however, you're not investing in any specific set of property rights. Instead, you're buying access to an *average* return on all property rights — the average sabotage.

Clearly this thinking is a world away from mainstream macroeconomics, which views the rate of interest as a variable for bringing 'financial markets into equilibrium' (Gregory Mankiw's words).[6] But we should expect as much. The core of neoclassical economics has always been to paper over class conflict.

By using language like 'sabotage', Nitzan and Bichler emphasize the conflict involved in setting rates of return. Now whether the rate of interest represents the 'average sabotage' is something that we can debate. But one thing seems clear: when it comes to distributing income, interest rates are not 'neutral'.

[6]There's a delightful irony to Mankiw's interest-rate description, found in the 7th edition of his tome *Macroeconomics*. After claiming that interest rates equilibrate financial markets, he mounts a case study which he calls 'war and interest rates'. In it, he shows that in the UK, interest-rate spikes correlate with war-driven spikes in military spending (measured as a portion of GDP). It takes a true neoclassical economists to describe war in terms of financial equilibrium.

Support this blog

Economics from the Top Down is where I share my ideas for how to create a better economics. If you liked this post, consider becoming a patron. You'll help me continue my research, and continue to share it with readers like you.

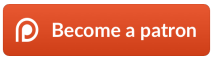

Debt and the interest share of income

Assuming that total debt and total income grow together, here's why the interest share of income should be largely determined by the rate of interest.

Let r be the annual rate of interest. Let D be the amount of outstanding debt within a country. The interest earned on this debt will be $r \cdot D$. Letting Y be national income, the interest share of national income is:

$$\text{interest share of income} = \frac{r \cdot D}{Y}$$

Now, if total debt is roughly the same as national income ($D \approx Y$), it follows that the interest share of income is determined solely by the rate of interest, r:

$$\text{interest share of income} = \frac{r \cdot D}{Y} \approx \frac{r \cdot Y}{Y} = r$$

Of course, this equivalence is an approximation. In the real-world, debt levels are rarely identical to national income — much to the chagrin of debt hawks.

Interest rates and labor-share clawbacks

Here's a simple model of how higher interest rates cause the labor share of income to decline.

We start with the observation (plotted in Figure 4) that higher interest rates are associated with a greater interest share of income. Fitting a linear regression to the trend, we can say that for every percentage point increase in the rate of interest, the interest share of income increases by about 0.7 percentage points.

Next, we assume that a growing interest share of income is directly clawed out of labor income. In that case, for every percentage point increase in the rate of interest, the labor share of income will *decrease* by 0.7 percentage points:

$$\text{labor share of income} = c - 0.7 \times \text{(rate of interest)}$$

Yes, this model is almost naively simple. And yet it fits the international trend fairly well. Figure 6-mod runs the numbers. Of course, the caveat is that I've chosen the parameter c to fit the data. But the main point is that the slope of the model is on par with the empirical trend.

Now this clawback model has obvious limitations. For example, if the interest rate exceeds 80%, the model predicts that the labor share of income becomes negative — an impossibility. Still, it's fascinating that such a simple model says *anything* about the labor share of income.

How interest-rate hikes redistribute income

Looking at Figure 7 (the relation between income percentile and the interest-to-wage ratio), I think it's intuitive that raising interest rates sends income to the rich. But since intuition sometimes misleads, let's do some formal math.

To dive into the mathematics of income redistribution, we'll start with a common misconception. If Alice owns more credit than Bob, it seems obvious that she will preferentially benefit from an interest-rate hike. If that's your intuition, know that it is wrong.

You see, when it comes to income inequality, it is income *ratios* that matter. And these ratios, in turn, are not affected by changes in the rate of interest. For example, suppose that Alice owns 10 times more credit than Bob. As long as Bob and Alice earn the same rate of interest, Alice will receive 10 times as much interest. In other words, (uniform) rate hikes don't redistribute interest income.

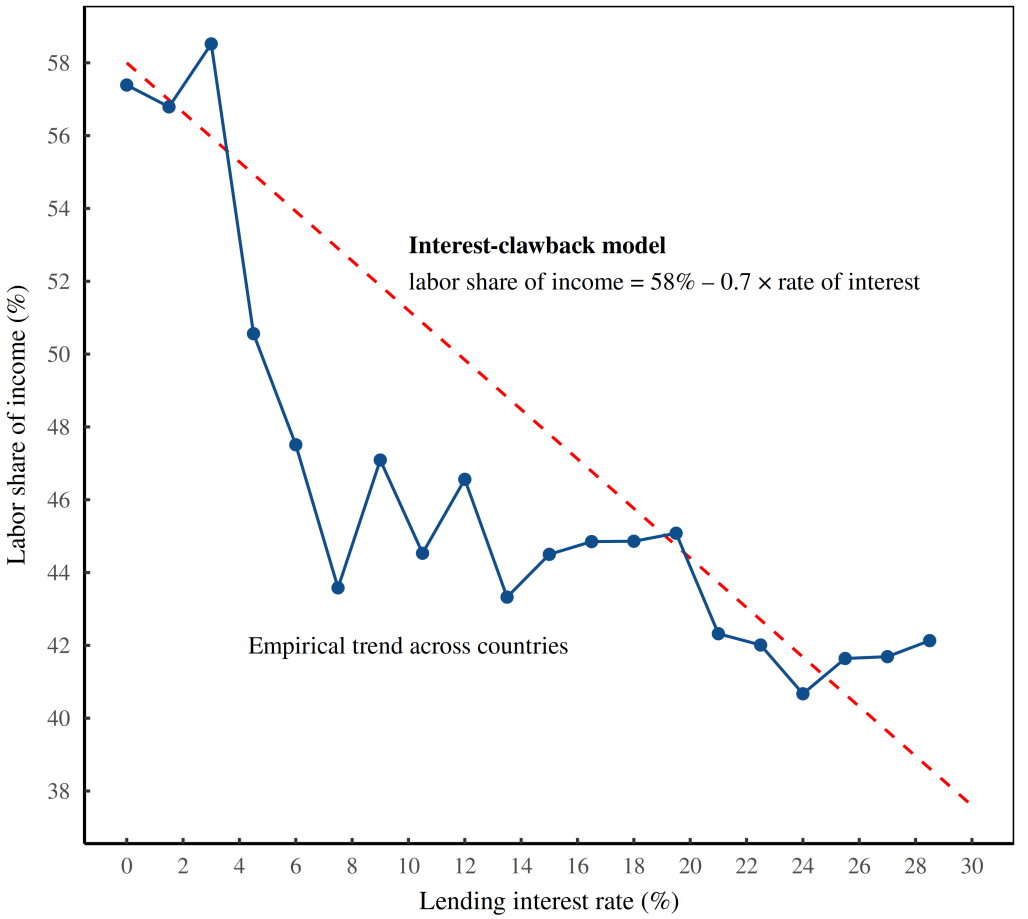

Figure 9: A clawback model of the labor share of income

The blue line replots (from Figure 6) the cross-country trend between interest rates and the labor share of income. The red line shows a simple model of this pattern. Using the regression from Figure 4, I assume that the interest share of income is proportional to the rate of interest. Next, I assume that this interest income gets clawed out of labor income. From there, it follows that the labor share of income will decline linearly with rising interest rates. Sources and methods

---

And now you are confused. You see, I just claimed that interest-rate hikes send money to the rich. But now I'm now arguing that they *don't* redistribute interest income. What gives?

To unconfuse ourselves, we need to bring other forms of income into the equation. To do that, let's imagine a world in which there are two types of income: *wages* and *interest*. Next, imagine that we hike the rate of interest while keeping wages unchanged. What happens to the distribution of income? Obviously creditors get a raise relative to workers.

That's nice and simple. But here's the problem: in the real world, there's no dividing line between 'workers' and 'creditors'. In other words, almost everyone earns some combination of wage and interest income. And so when it comes to the distribution effect of interest-rate hikes, we're back to being confused.

To unconfuse ourselves for a second time, we need to do some math. We'll start by writing an equation that describes individual income. In our imaginary world, each person's income consists of their annual wage (w) plus any interest (i) they earn on credit investments:

$$\text{income} = w + i$$

Interest income, in turn, is set by the size of the credit investment (K) and the rate of interest (r). In other words, $i = K \cdot r$. Putting that into our income equation, we get:

$$\text{income} = w + K \cdot r$$

Next, we want to know how income will change if we hike the rate of interest but leave wages the same. Let's suppose that the rate of interest rises from r_1 to r_2. In that case, income growth is:

$$\text{income growth} = \frac{w + K \cdot r_2}{w + K \cdot r_1}$$

One last step. In most cases, we don't care about the absolute value of your wage or your credit investment. What matters is the *ratio* of these two quantities — the credit-to-wage ratio, $\frac{K}{w}$. So let's rewrite our income-growth equation in terms of this ratio. Factoring out w in both the numerator and the denominator gives:

$$\text{income growth} = \frac{1 + \frac{K}{w} \cdot r_2}{1 + \frac{K}{w} \cdot r_1}$$

This equation tells us how your income grows in response to an interest-rate hike. The larger your credit-to-wage ratio, the more you benefit.

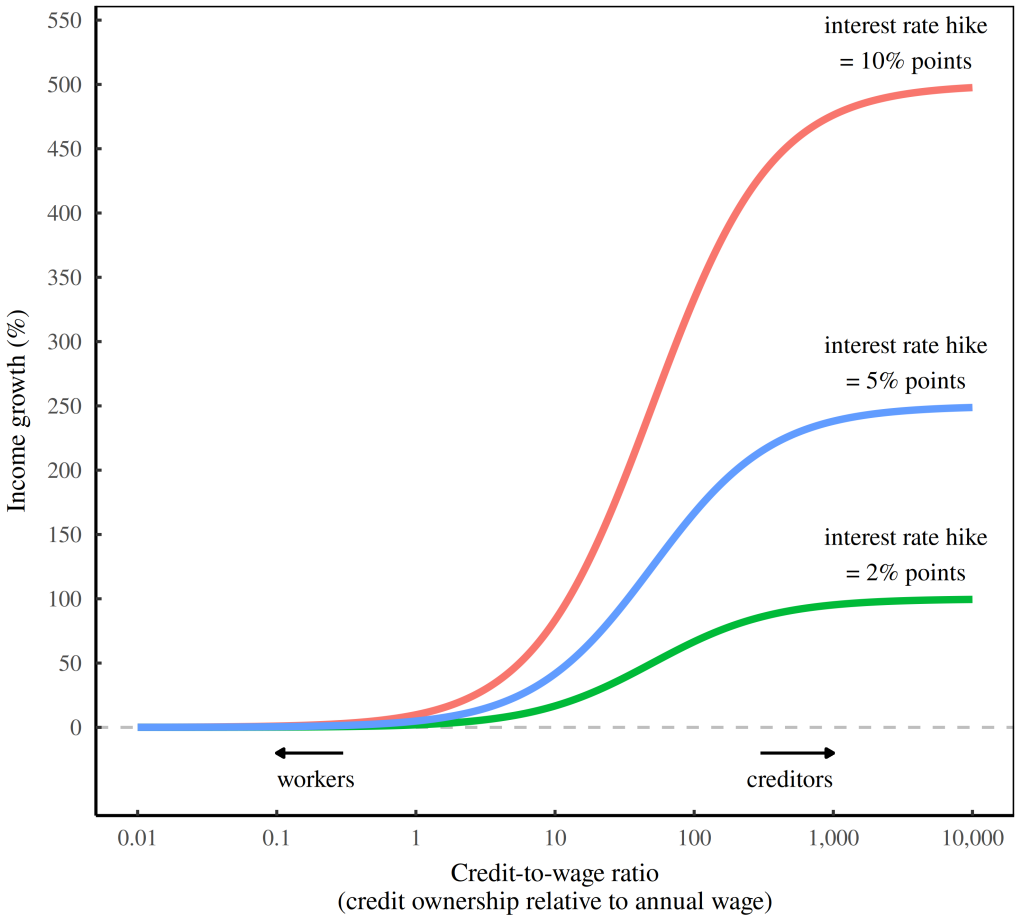

Figure 10: How interest-rate hikes redistribute income

This figure illustrates how the income affect of interest-rate hikes depends on the amount of credit that you own. To that end, the horizontal plots the credit-to-wage ratio — the ratio of an individual's credit ownership to their annual wage. (The larger this ratio, the more you are a 'creditor'. The smaller this ratio, the more you are a 'worker'.) The vertical axis shows how income grows in response to an interest-rate hike (while wages are kept constant). Each colored curve shows the effect of a different rate hike, which starts from a baseline interest rate of 3%.

To get a sense for how this equation works, let's look at Figure 10. Here, the horizontal axis shows the credit-to-wage ratio, which basically indicates the degree to which you are a worker (left) or a creditor (right). The vertical axis then shows how income grows in response to a given rate hike. Obviously, the larger the rate hike, the bigger the pay raise for creditors. But more importantly, the colored curves show how the degree of benefit changes *continuously* as a function of your class status.

Having twice confused ourselves and then unconfused ourselves, we can now go back to our initial intuition. Raising the rate of interest sends money to the rich.

How do we know that? Well, from Figure 7, we know that as incomes get larger, people tend to earn more money from interest and less money from wages. In other words, their interest-to-wage ratio increases with total income. Now, assuming that everyone earns roughly the same rate of interest, the interest-to-wage ratio is a proxy for the *credit*-to-wage ratio — the ratio of credit ownership to the annual wage. So as incomes get larger, the credit-to-wage ratio increases.

Now let's put it all together. From Figure 10, we know that interest-rate hikes preferentially benefit people with a large credit-to-wage ratio. And since this credit-to-wage ratio grows with (total) income, we're back to our original intuition: raising the rate of interest sends money to the rich.

A fixed-wages fixed-credit model of how interest rates affect income inequality

Building on the mathematics above, let's create a simple model of how interest rates affect income inequality.

We'll start by supposing that among individuals, wage income follows a log-normal distribution with scale parameter σ:

$$w \sim \text{lnorm}(\sigma)$$

Next, let p be each individual's income percentile, ranked by their wage w. Now suppose that for each individual, their interest-to-wage ratio, $\frac{i}{w}$, is some function of their income percentile. Let's notate that function as $(\frac{i}{w})_p$.

Next, let's assume that the interest-to-wage ratio varies with income percentile as found in Figure 7 (the average relation in the United States). I fit this pattern with a power-law regression, giving the following equation:

$$\log\left(\frac{i}{w}\right)_p = 0.015 - 0.31 \cdot \log(1-p)$$

Given an individual's interest-to-wage ratio, we can then calculate the size of their credit investment, K. To do that, we multiply their interest-to-wage ratio by their wage, w, and then divide by some hypothetical rate of interest, r_0:

$$K_p = w_p \cdot \frac{(\frac{i}{w})_p}{r_0}$$

Applying this operation to every individual, we know both their wage w and their credit investment K.

Next, let's assume that these wages and credit investments stay the same while the rate of interest changes. If the variable rate of interest is r, the total income of a person in wage percentile p is:

$$I_p = w_p + K_p \cdot r$$

After we've calculated income for every individual, we can see how the distribution of income relates to the rate of interest r.

I'm calling this model the 'fixed-wages fixed-credit model', because it assumes that the distribution of wages and credit are fixed over time. The model has two tunable parameters. The lognormal scale parameter σ determines the inequality of wage income (which is the baseline inequality when interest rates are zero). And the interest rate r_0 determines the average size of credit investments (relative to the average size of wages). When combined, the two parameters determine how income inequality changes as interest rates rise.

While the specific results depend on the exact parameters, a general feature of this model is that the top 1% share of income grows non-linearly with the rate of interest. Figure 8-mod shows an example. Here, the model parameters are $\sigma = 0.85$ and $r_0 = 0.4\%$.

So what produces the plateau in inequality as interest rates continue to rise? The answer is that in this model, the interest rate is essentially a weighting variable that determines the average size of interest income relative to wages. As this rate heads north, interest income starts to dominate total income. And in that limit, the distribution of income reflects the distribution of credit, which is assumed to be constant.

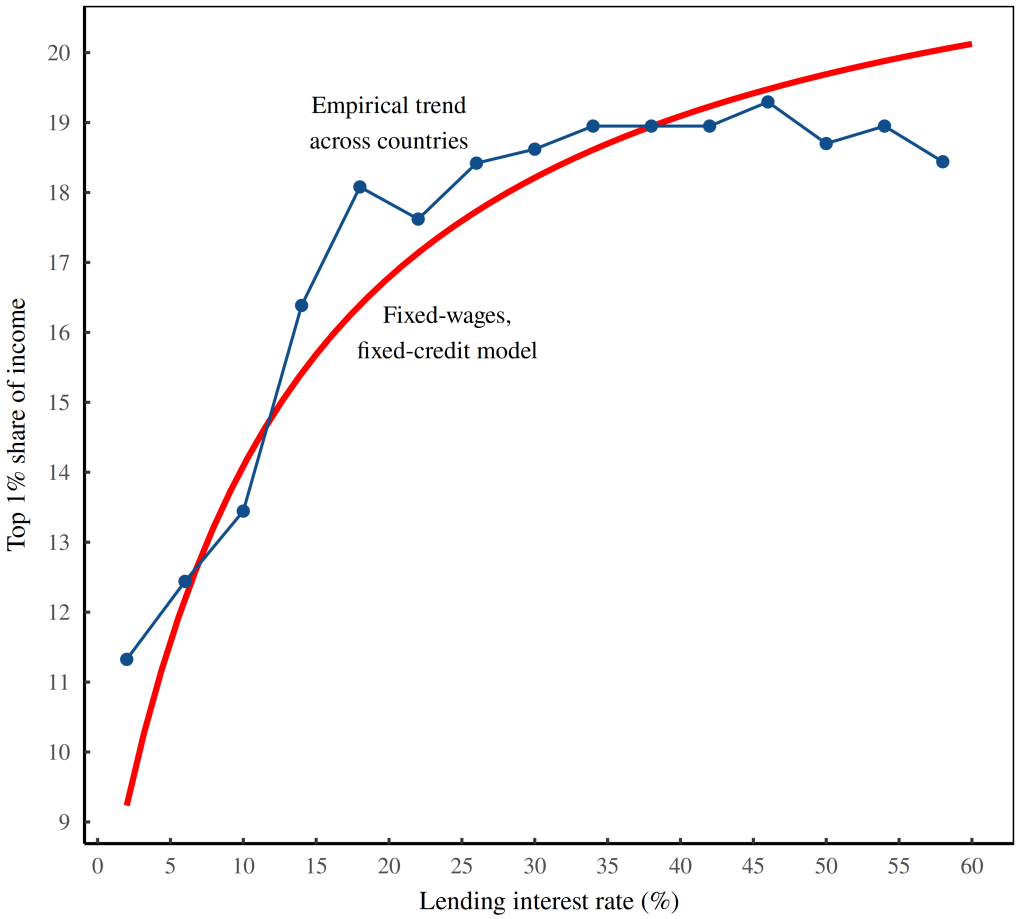

Figure 11: The fixed-wages fixed-credit model of how interest rates affect income inequality

This figure shows how a simple model of credit income can explain why income inequality increases non-linearly with the rate of interest. The crux of the model is that the size of credit investments (measured relative to wages) tend to grow as incomes get larger. From this assumption, we can derive a static distribution of wages and credit investments. When we then take this arrangement and hike the rate of interest, we send money to the rich, increasing inequality. The important feature of the model is that that the interest-inequality pattern is non-linear, as shown by the red curve. The blue curve replots the empirical data from Figure 8.

Sources and methods

Figure 1

Word frequency data is from the 2019 Google English corpus, downloaded with the excellent R package ngramr.

Figure 2

- Data is from the US Bureau of Economic Analysis, Table 1.12. National Income by Type of Income. I've used data before taxes, without inventory adjustment and without capital cost adjustment.
- Note that for illustration purposes, the pie chart defines national income as the sum of employee compensation, profit, proprietor income, rent, and interest. This sum excludes some terms that the BEA uses to make the income side of the national accounts consistent with the expense side.

Figure 3

Data for US bond yields:

- Bond yields from 1798 to 1959: Historical Statistics of the United States, Table Cj1192-1197 (long-term bond yields). This table contains several series for bond yields, each of which covers a different period of time. To construct the long-term index, I calculate the average (the unweighted mean) of the reported data in each year.
- Bond yields from 1960 to 2022: FRED series IRLTLT01USM156N, long-term government bond yields, 10-year.

Data for the interest share of income:

- 1929 to 2021: US Bureau of Economic Analysis, Table 1.12, National Income by Type of Income. (I divide net interest by national income.)
- 1900 to 1929: Table 1 in Gale Johnson's paper 'The Functional Distribution of Income in the United States, 1850-1952'.
  - Note that Johnson's national income data does not include the various adjustment terms that are part of modern national income data. To adjust for this discrepancy, I first take the BEA data and sum the 5 types of income and calculate their average share of national income. (It's about 90%.) Then I multiply Johnson's income share data by this adjustment.
  - Also note that Johnson's data is reported as averages over roughly 5-year intervals. To make this data comparable to the bond yield data, I average the latter over the same time intervals.

Figure 4

For US data, see the sources for Figure 3. Non-US data is from the following sources:

- Interest share of national income; data is from the OECD, Experimental Statistics: Distributional information on household income, consumption and saving.
  - National income: series TOTRESPIA
  - interest income: D41R
- interest rate data comes from two sources:
  - World Bank, series FR.INR.LEND, lending interest rate
  - OECD, series LTINT, long-term interest rates
  - Note: when merging the World Bank and OECD data, if/when I found duplicate country-year observations, I used the World Bank data.

Figure 5

Data for nominal US GDP:

- 1790 to 1928: Historical Statistics of the United States, series Ca10.
- 1929 to 2021: Bureau of Economic Analysis, Table 1.1.5 Gross Domestic Product

Data for total US debt:

- 1819 to 1833: Historical Statistics of the United States, series Cj148 (state banks, loans and discounts)
- 1834 to 1915: Historical Statistics of the United States, series Cj252 (total assets or liabilities of commercial banks)
- 1916 to 1945: Historical Statistics of the United States, series Cj870 (Net public and private debt)
- 1946 to 2021: sum of FRED series DODFS (Domestic Financial Sectors; Debt Securities and Loans; Liability, Level) and FRED series BOGZ1LA384104005A (Domestic Nonfinancial Sectors; Debt Securities and Loans)

Notes:

- I've spliced the various debt and GDP series together by indexing to the most modern data. In other words, the GDP series is consistent with the BEA data going backwards in time. And the debt series is consistent with the FRED data going backwards in time.
- Prior to 1916, the debt data is based on bank liabilities only, so it should be treated with appropriate uncertainty.

Figure 6

Interest rate data comes from two sources:

- World Bank, series FR.INR.LEND, lending interest rate
- OECD, series LTINT, long-term interest rates
- Note: when merging the World Bank and OECD data, if/when I found duplicate country-year observations, I used the World Bank data.

Data for the labor share of income:

- Ilostat, series SDG_1041_NOC_RT_A, Labor income share as a percent of GDP

Figure 7

Data is from the World Inequality Database, as follows:

- US interest income is from series afiint992t (interest income threshold by income percentile)
- US wage income is from series afiwag992t (wage and pension income threshold by income percentile)
- Figure 7 plots the ratio of the two series (afiint992t / afiwag992t).

Note that according to the World Inequality database, these series are considered 'legacy' measurements, and are not being updated. Why?

To understand what's going on, you have to understand the history of World Inequality Database. It started life as the 'World Top Incomes Database', which was built to house the data generated by Thomas Piketty and his collaborators. And Piketty, in turn, pioneered the study of income inequality based on *tax records*.

Now, tax records come with some problems — the simplest being that taxes are typically reported by households, not individuals. But the other 'problem' is that tax records are not consistent with the accounting definitions used by the national accounts.

I've used scare quotes here, because I don't think this inconsistency is a problem. If anything, tax records offer the best way to account for income. You see, income is fundamentally a legal concept. And so what better way to measure it than by using the records generated by the tax code — the legal system that codifies and classifies income.

But for better or worse, Piketty and collaborators have decided to abandon the tax approach, and instead, make their methods consistent with the national accounts. (They got the ball rolling in this 2016 paper.)

In one sense, the move is a good idea, since it's nice to have accounting techniques that are consistent across databases. But in another sense, I think the change is a mistake. You see, the national accounts were never designed to give a comprehensive measure of income. They were designed to measure *production*. That's why capital gains — which are obviously income — are not part of the national accounts. They were deemed 'unproductive'.

At any rate, the legacy tax data still sits in the World Inequality Database (if you know where to find it). But it's not being updated. And it's available for only a handful of rich countries.

Figure 8

- For interest rate sources, see the sources for Figure 6.
- Data for the top 1% of income is from the World Inequality Database, using the following income share series: sptinc992j, sptinc992i, sptinc992t, sptinc999i, sptinc996i

Further reading

Nitzan, J., & Bichler, S. (2009). *Capital as power: A study of order and creorder*. New York: Routledge.

Rochon, L.-P., & Setterfield, M. (2007). Interest rates, income distribution, and monetary policy dominance: Post Keynesians and the "fair rate" of interest. *Journal of Post Keynesian Economics*, *30*(1), 13–42.