Economics from the Top Down

new ideas in economics and the social sciences

When Stocks Go Up, Who Benefits?

Blair Fix

October 28, 2023

Cui bono? For whose benefit?

Think of this question as a sword — a sharp piece of steel that cuts through bullshit. In this post, we'll use it to slice through business-press bullshit about the stock market. You know the stuff — the ubiquitous puff pieces that gush about rising stock prices, as though they benefit everyone.

When we ask *cui bono*, we carve through this BS. We discover that for most people, rising stocks are a tool not for gain, but for administering pain. Looking at the United States, I find that when stocks go up, the vast majority of people see their share of income (and wealth) *decline*.

So here's the truth about the stock market: it's a socially sanctioned way to take from the poor and give to the rich.

Number go up: a brief history of the US stock market

Before we dive into the details of how the stock market makes the rich richer, it's worth pausing for some history. Question: if you had to capture the history of the stock market with a catch phrase, what would it be?

Personally, I'd go with the aphorism '*number go up*'.

Figure 1 shows the number-go-up pattern in the United States. Here, I've plotted the century-long rise of the S&P 500 — a popular index of US stocks. To situate this history, I've labeled some of the major stock-market booms and busts. Note how these events add short-term froth to the mix. However, they don't disrupt the long-term trend, which is unrelentingly up.

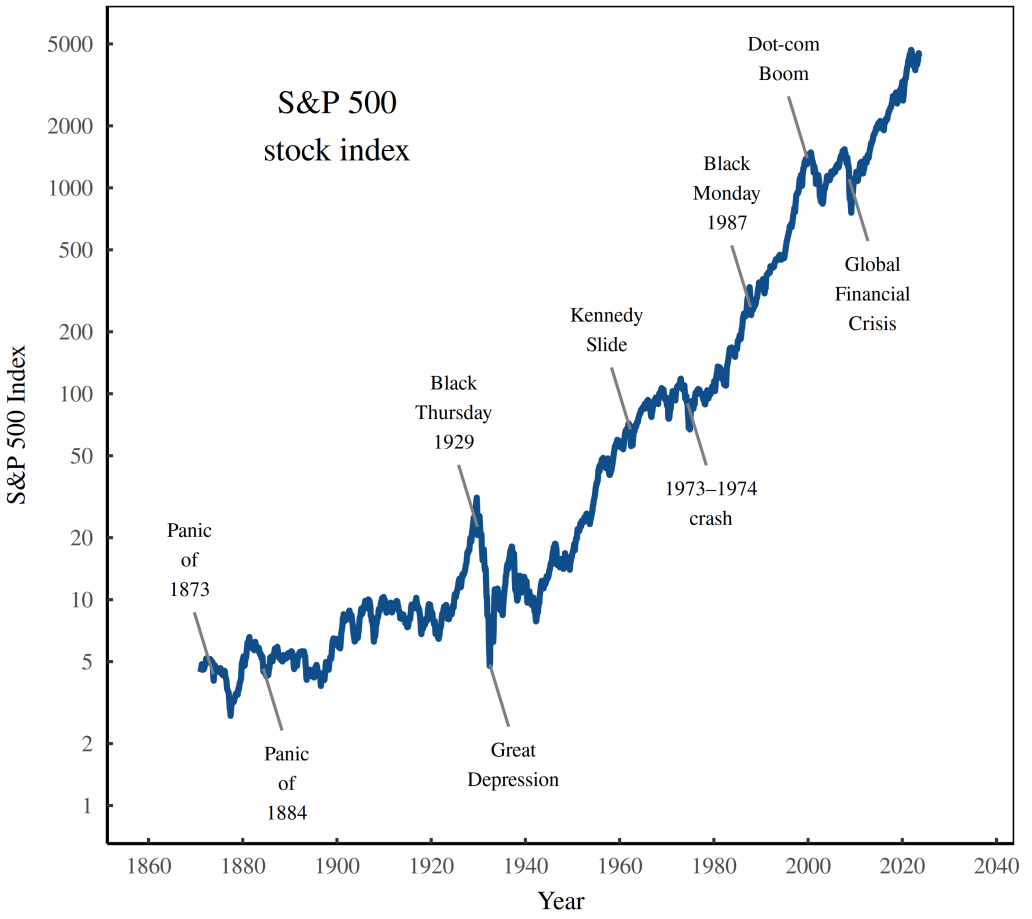

Figure 1: Number go up — the rise of the S&P 500

This figure plots Robert Shiller's data for the long-term history of the S&P 500 index — a popular measure of US stock prices. Note that the vertical axis uses a log scale. Sources and methods

---

Now that we've looked at the stock market's northward journey, realize that the upward trend is basically meaningless.[1] That's because like all financial quantities, stock-market returns don't mean anything until we've *compared* them to something.

[1] Speaking of meaningless, I once had a financial advisor show me a chart of the long-term rise of the S&P 500. "Look at these numbers go up," he said. "That's why you should invest in the stock market."

I almost laughed in his face. I wanted to tell him that you could make the same argument for why you should invest in being an *unskilled worker*. 'Look at how unskilled wages have grown over the last century. That's why it's always worthwhile remaining uneducated.'

Instead, I simply asked the guy if his chart was adjusted for inflation. He sheepishly replied "no". I smiled and changed the subject.

For example, suppose that last year, the S&P 500 rose by 15%. Is that a good rate of return? The answer is that without context, we have no idea. To judge this 15% return, we need to compare it to another rate of return. And that could be anything — the price of gold, the price of oil, the yield on bonds, the return on foreign stocks, and so on.

Typically, investors judge their stock returns against the price of other assets — other things they can own. But in terms of political economy, what's more interesting is to compare stock returns to other things you *can't* own … namely other people's income.

For example, political economists Jonathan Nitzan and Shimshon Bichler have done fascinating work studying how the stock market performs relative to average wages. They call this comparison the 'power index', and argue that it quantifies the class struggle between capitalists and workers. (For my take on their approach, see my piece 'Stocks are up. Wages are down. What does it mean?')

In this post, I compare stock returns to a broad measure of average income — GDP per capita.[2] Figure 2 shows the rise of the S&P 500 in this context. Over the long haul, stock prices rose at about the same rate as US GDP per capita. But over shorter periods, there's a dance between the two rates of return. Sometimes the stock market won. Other times GDP per capita took the lead.

A race into uncharted territory

When we compare stock-market gains to GDP per capita, we're effectively watching a financial race between two hypothetical people.

Imagine that your friend Alice puts all her money into a stock fund that tracks the S&P 500. And imagine that your friend Bob manages to index his salary to US GDP per capita. With their 'investments' in hand, Alice and Bob meet each year to see who came out on top. To their surprise, they find that the race has a cyclical pattern. For a few decades, Alice wins. But then she loses ground, and Bob takes the lead for another few decades.

[2]Note that the term 'gross domestic *product*' is a misnomer. GDP does not measure 'production'. It measures *income* using double-entry book keeping. On one side, there is the sum of firm's 'value added' — their sales less their non-labor costs. And on the other side, there is the sum of personal income. Neither side has anything to do with the quantity of production.

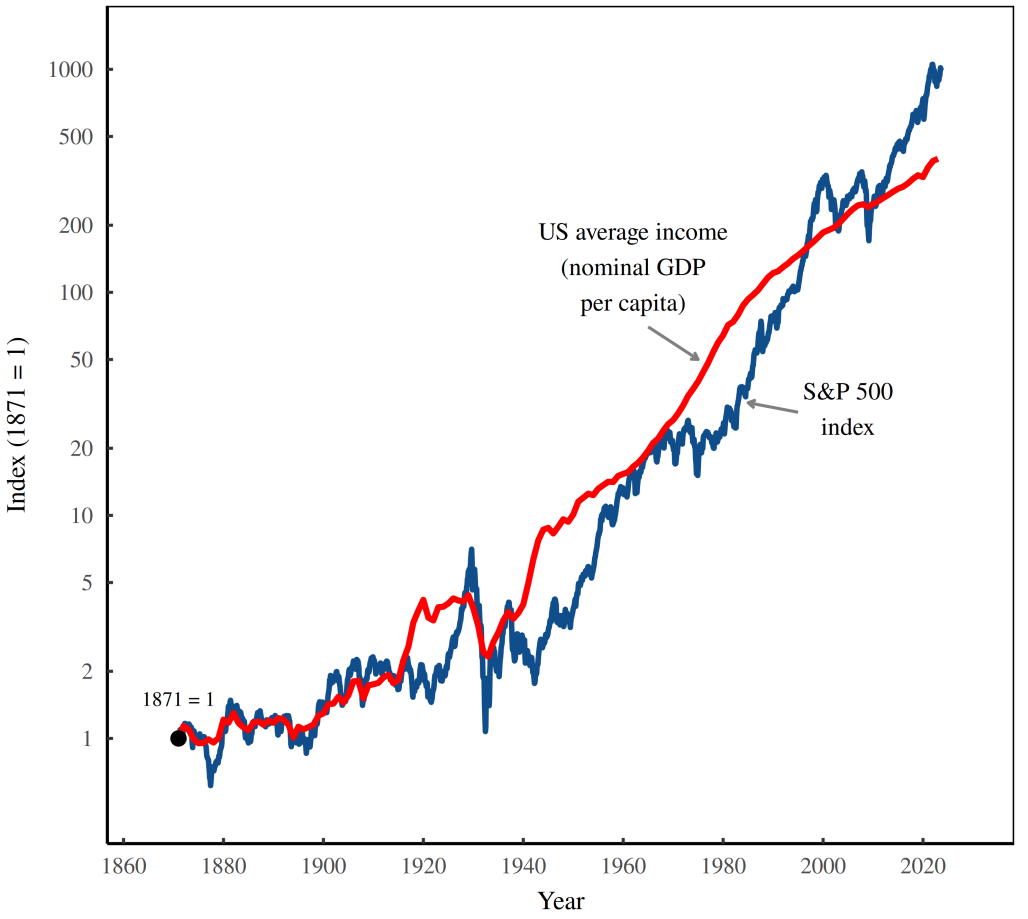

Figure 2: US stock returns in context

This figure shows how the S&P 500 index (a measure of US stock prices) has risen relative to US nominal GDP per capita. (Note the log scale on the vertical axis.) Over the long haul, the race is quite evenly matched. But during the short term, the competition goes in cycles. Sometimes the stock market wins. Other times GDP wins. Sources and methods

---

If the race began in 1871 and lasted until today, it would look like Figure 3. Here, the blue curve takes the S&P 500 index and divides it by US nominal GDP. This ratio indicates Alice's lead over Bob.

As the twentieth century unfolds, we see a fairly competitive race, with Alice sometimes gaining ground but then later losing it. Notice, though, that the race has recently become one-sided.

During the dot-com boom of the late 1990s, Alice's stock investment took a commanding lead over Bob's investment in GDP. True, Alice got pummeled during the 2008 financial crisis. But during the bull market of the 2010s,

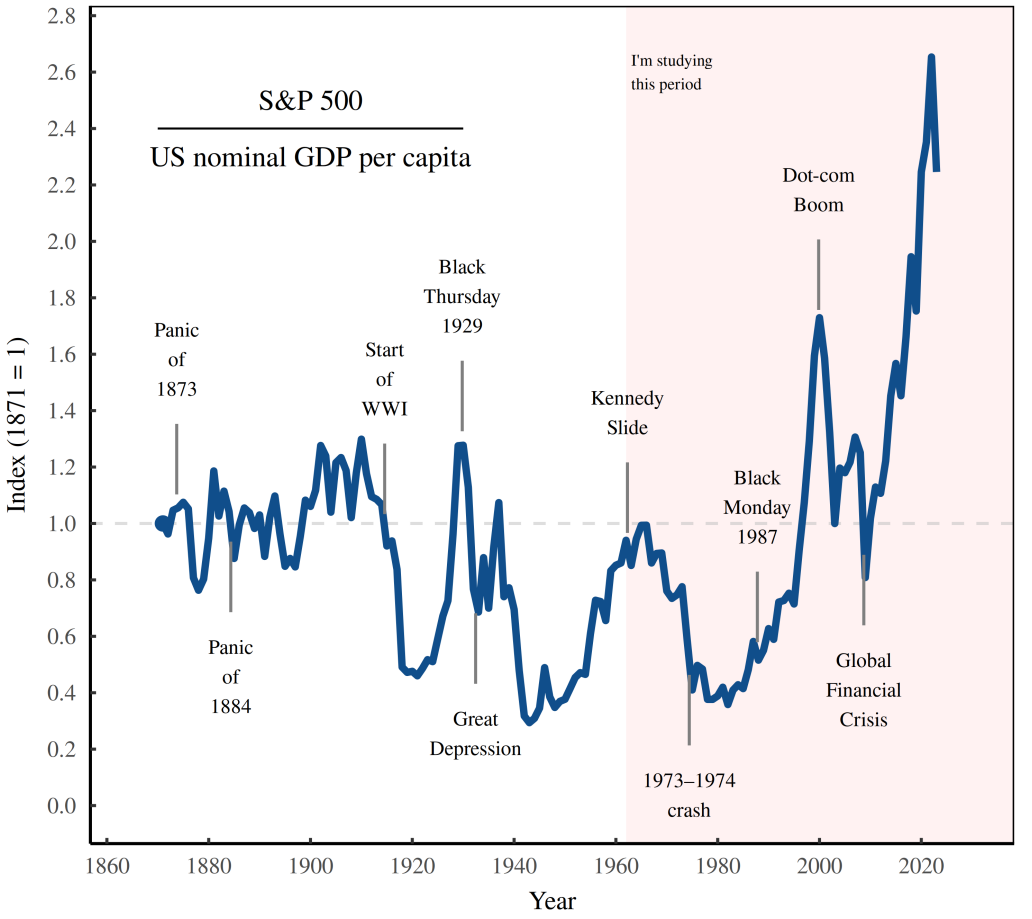

Figure 3: A race into uncharted territory

This figure plots the ratio between the S&P 500 index and US nominal GDP per capita. For much of the last century and a half, the race was fairly equal. But in the 21st century, stocks have taken a commanding lead over GDP. Sources and methods

she regained the lead. Or more accurately, she got *catapulted* into uncharted territory. By 2020, Alice's stock investments gained an unprecedented lead over Bob's GDP-indexed income.

Looking at the US stock-market's presently uncharted territory, I'd warn Alice not to get cocky. Sure, her investment is at an all-time high relative to GDP. But if the past is any indication, there's nowhere to go but down.

But I digress. This post isn't about investment advice. It's about *cui bono*. When stocks rise relative to GDP per capita, who benefits?

The race to divide the financial pie

To answer our *cui bono* question, we need to leave Alice and Bob behind and turn our attention to a different race: the race to divide the financial pie.

First, some spoilers. By definition, the distribution race is zero sum, which means that if I enlarge my share of the financial pie, someone else has their share reduced. In this race, win-win is not an option.

With zero-sum competition in mind, let's turn to Figure 4. Here, I've plotted the race to distribute US income. It's a contest with 100 participants — one for each income percentile. The colored lines show how the income share of each percentile has changed since 1962. (I'll explain why I chose this date in a moment.)

If the distribution race were a draw, then all the colored lines in Figure 4 would travel horizontally, indicating that each income percentile preserved its share of income. But in the modern US, that's not what happened. Instead, top percentiles saw their income share rise. And everyone else saw their income share decline. The visual result is a pretty rainbow that gradually fans outward.

Notice that in Figure 4, I'm focusing on the period since 1962. That's partly because the most dramatic swings in the stock market have occurred in the last half century. (See Figure 3.) But it's also because I want to expand my analysis to include the distribution of *wealth*. And this wealth data only goes back to 1962.

Figure 5 shows the US wealth competition. Each colored line plots the wealth share of a corresponding wealth percentile. As with income, the distribution of US wealth has grown more unequal, with top wealth brackets increasing their share, and everyone else making do with less.

Measuring stock-market gain and pain

Now that we've assembled our income and wealth data, we're ready to see who benefits from the rise and fall of the stock market. Figure 6 illustrates my method.

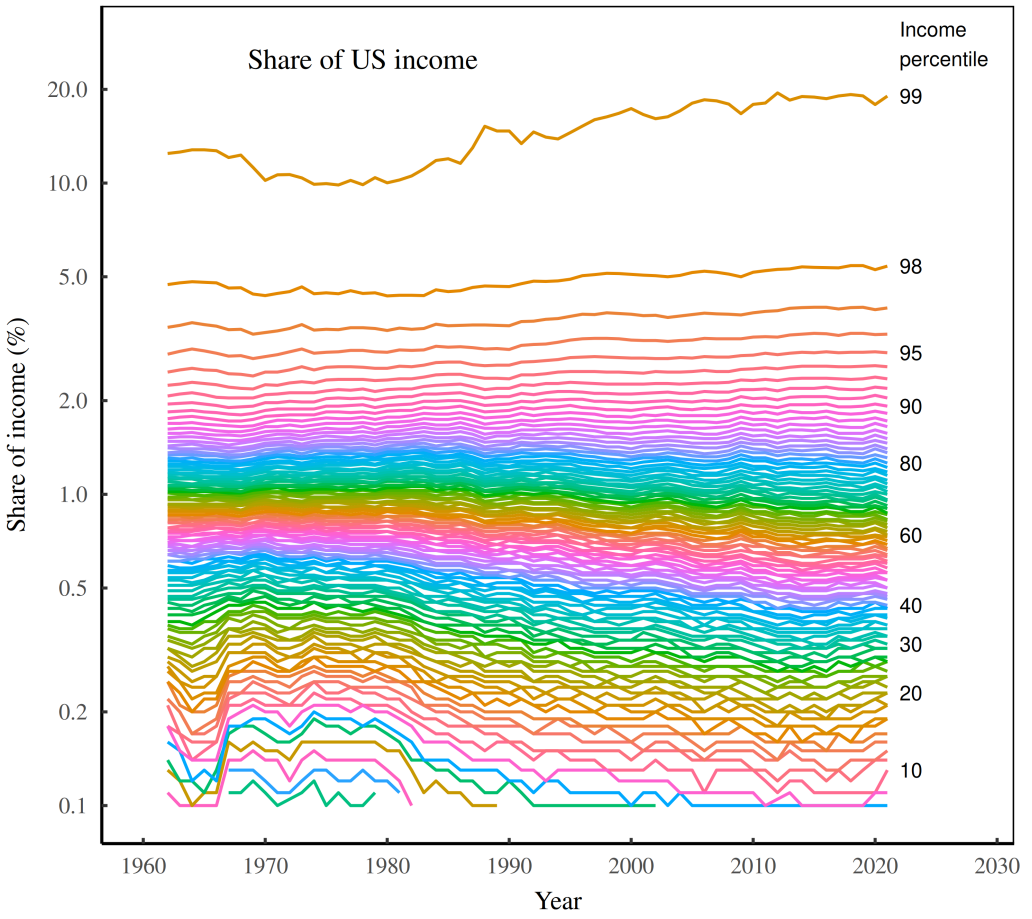

Figure 4: The race to divide the US income pie

Each colored line shows the income share of a particular income percentile. For example the line labeled '99' shows the income share of individuals in the 99th to 100th percentile. And the line labeled '98' plots the income share of individuals in the 98th to 99th percentile. And so on. Note that the vertical axis uses a log scale. Also note that below the 10th percentile, the income share is typically zero, which isn't plottable on a log scale. Sources and methods

---

Basically, it's a game of correlation. We start by selecting a specific US income percentile (or later on, a wealth percentile). For example, in Figure 6A I've selected the 50th income percentile. Next, we see how the *income share* of this percentile relates to the motion of *stock prices*, as captured by the stock-market-to-GDP ratio.

When we crunch the numbers, we find that the resulting pattern depends on the income percentile we've selected. For example, if we select the 50th income percentile (Figure 6A), we find that stock-market gains come with an

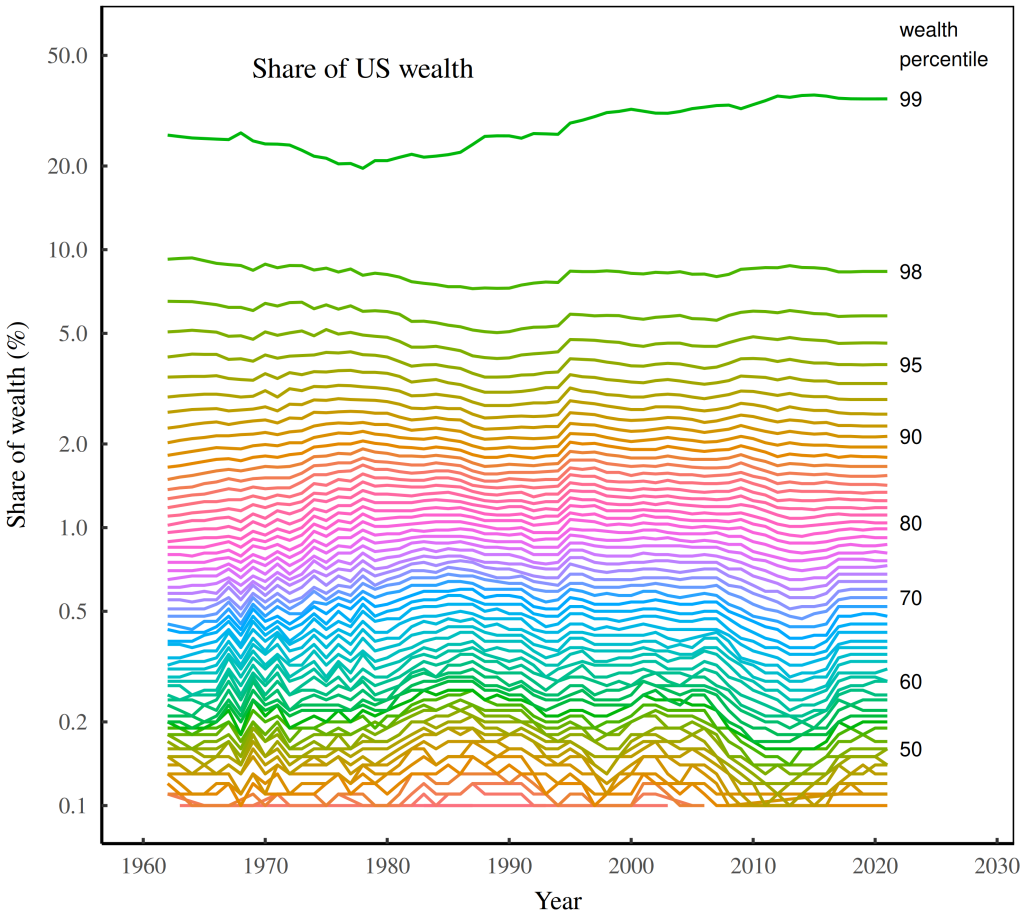

Figure 5: The race to divide the US wealth pie

Each colored line shows the wealth share of a particular US wealth percentile. Note that the vertical axis uses a log scale. Also note that below the 50th percentile, the wealth share is typically zero, which isn't plottable on a log scale. Sources and methods

---

income-share *decline*. However, when we select the 98th income percentile (Figure 6B), we get the opposite trend; stock-market gains come with an income-share *increase*.

Looking at Figure 6, the message is clear: a rising stock market doesn't benefit everybody. In reality, one person's gain is another person's pain.

Now, if you're a fan of simple analysis, we could close the case here. In the US, rising stock prices appear to harm the income share of the 50th percentile, while they bolster the income share of the 98th percentile.

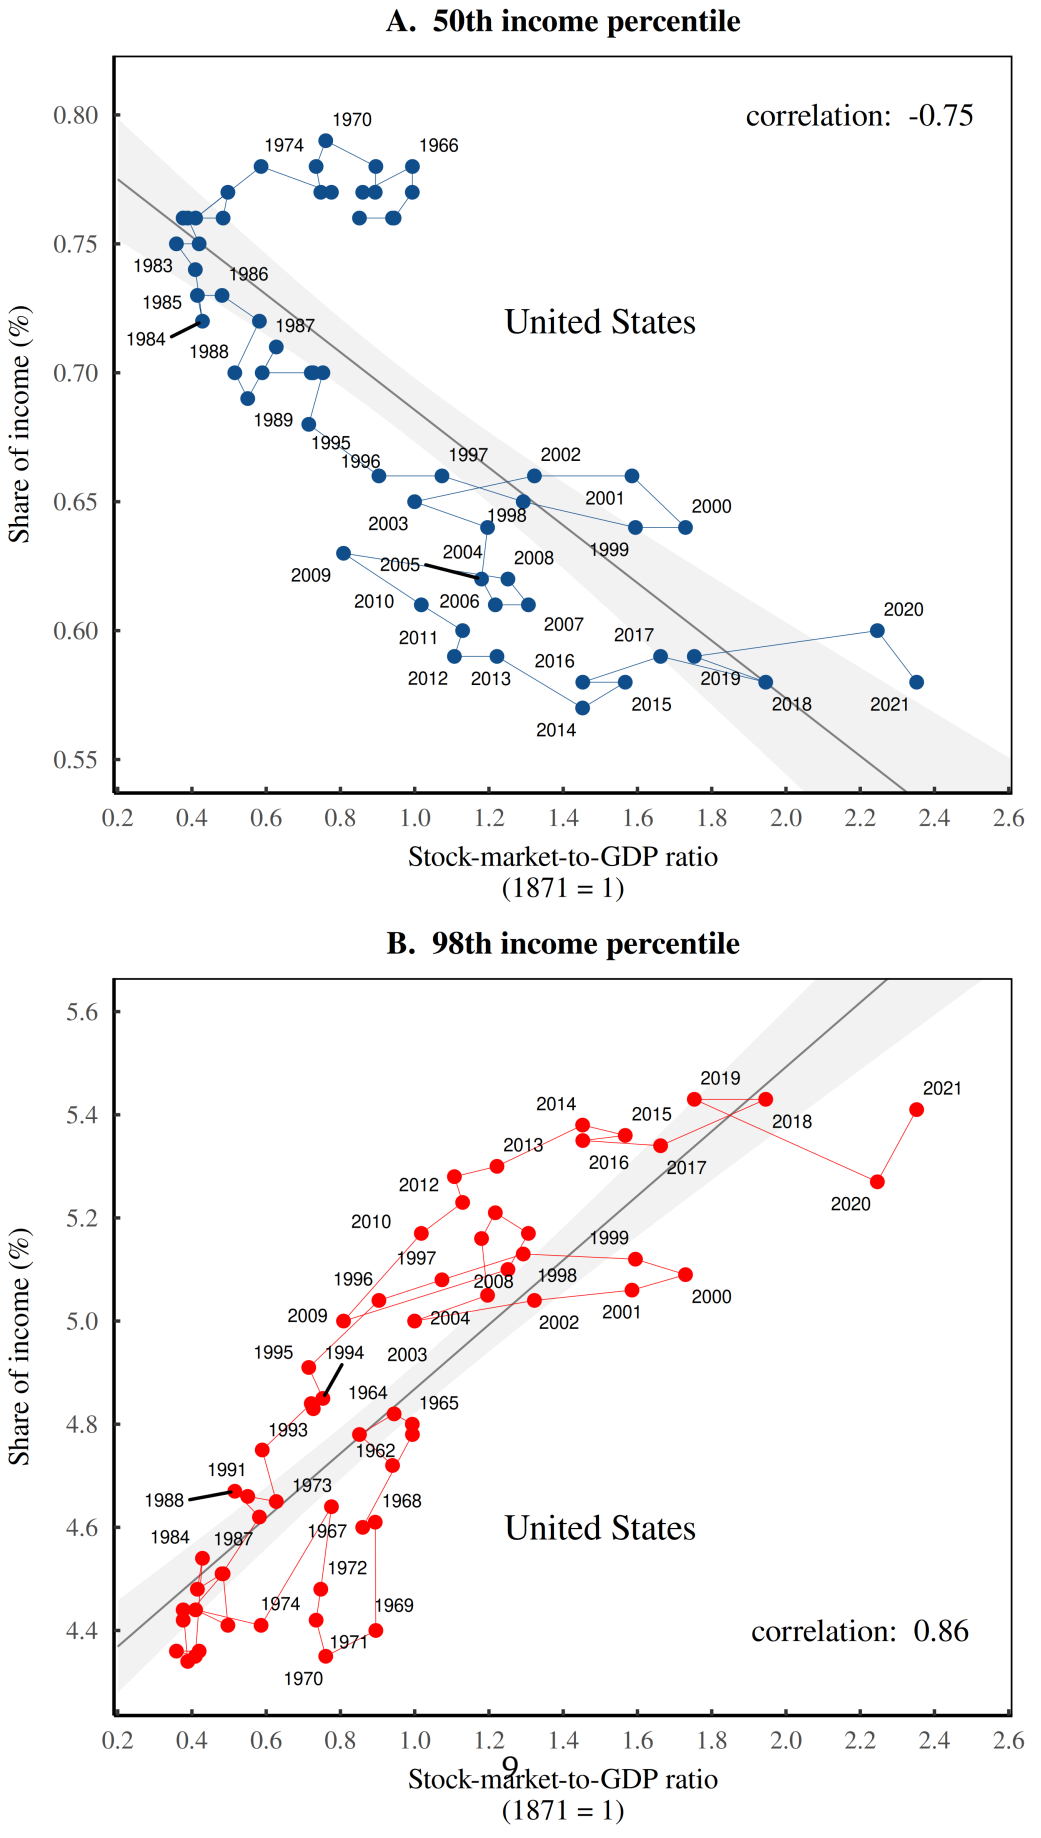

Figure 6: Stock market losers and winners

But as you can probably guess, I'm not going to stop here. To me, Figure 6 feels like a photo with many missing pixels. By selecting two income percentiles, we get a hint of the whole picture. But I want more than a hint. I want high-resolution glory. I want to fill in every pixel by applying the same analysis to *every* US income percentile.

To do that, I'm going to take each US income percentile and see how its share of income relates to the movement of the stock market. But rather than visualize the raw data (which would result in dozens of scatter plots) I'm going to reduce the data to a *correlation*.[3]

In other words, we take the scatter plot in Figure 6A and reduce it to a correlation of -0.75. (The negative value indicates that as stocks go up, the income share held by the 50th percentile declines.) And we take the scatter plot in Figure 6B and reduce it to a correlation of +0.86. (The positive value indicates that as stocks go up, the income share held by the 98th percentile rises.)

This reduction gives us two pixels. But if we repeat the analysis for every US income percentile, and we'll fill in the whole picture.

Lifting all boats

Before we get to our hi-res picture, it's worth setting some (naive) expectations. If the stock market actually lifted all boats, what would it look like?

Well, it would look something like Figure 7 — a delightfully dull flatline.

Now at first, this flatline seems counter-intuitive. If the stock market is lifting all boats, shouldn't we see some sort of upward trend? Actually, no. The key here is that we're measuring how the stock market relates to income *distribution* (not income itself). And if the stock market lifts all boats equally, that means it has *no* effect on the distribution of income. Hence our flatline in Figure 7.

For every income percentile (horizontal axis), the correlation between income share and the stock-market-to-gdp ratio (vertical axis) hovers around zero. In short, the race to distribute income bares no relation to the movement of the stock market.

[3]Fun fact: the World Inequality Database reports income shares for 127 different income/wealth percentiles. There are 99 observations for the bottom 99 percentiles (one for each percentile). Then there are 28 observations that split up the top 1% into fine-grain detail.

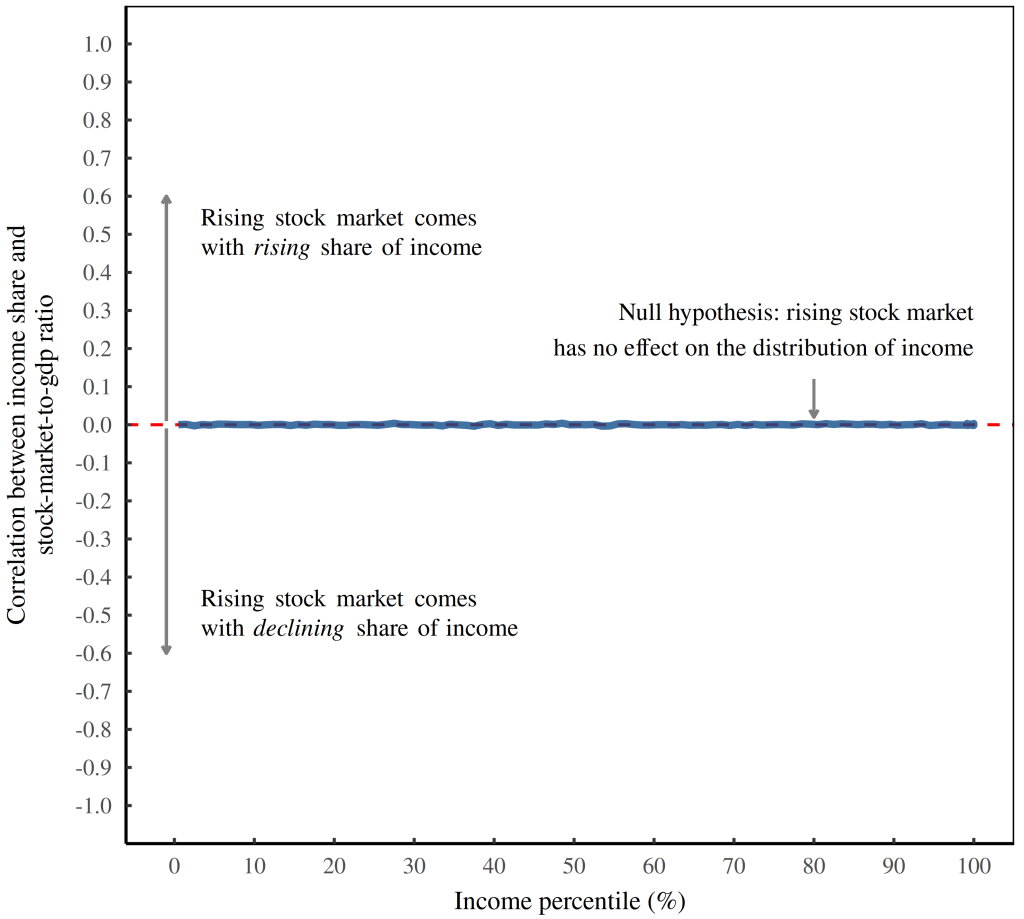

Figure 7: If the stock market lifted all boats, it would look like this

This figure shows what would happen if stock-market returns had no affect on the distribution of income (i.e. the null hypothesis). In this case, for all income percentiles (horizontal axis) the correlation between income share and the stock-market-to-gdp-ratio (vertical axis) is essentially zero. Sources and methods

Lifting the rich boats, sinking the rest

So does the stock market actually lift all boats? Of course not! In reality, stock-market gains are a recipe for lifting a few rich boats, and sinking the rest.

Figure 8 tells the US story. The key result is that there's no flatline to be found. Instead, we get an L-shaped pattern. Here's what it means.

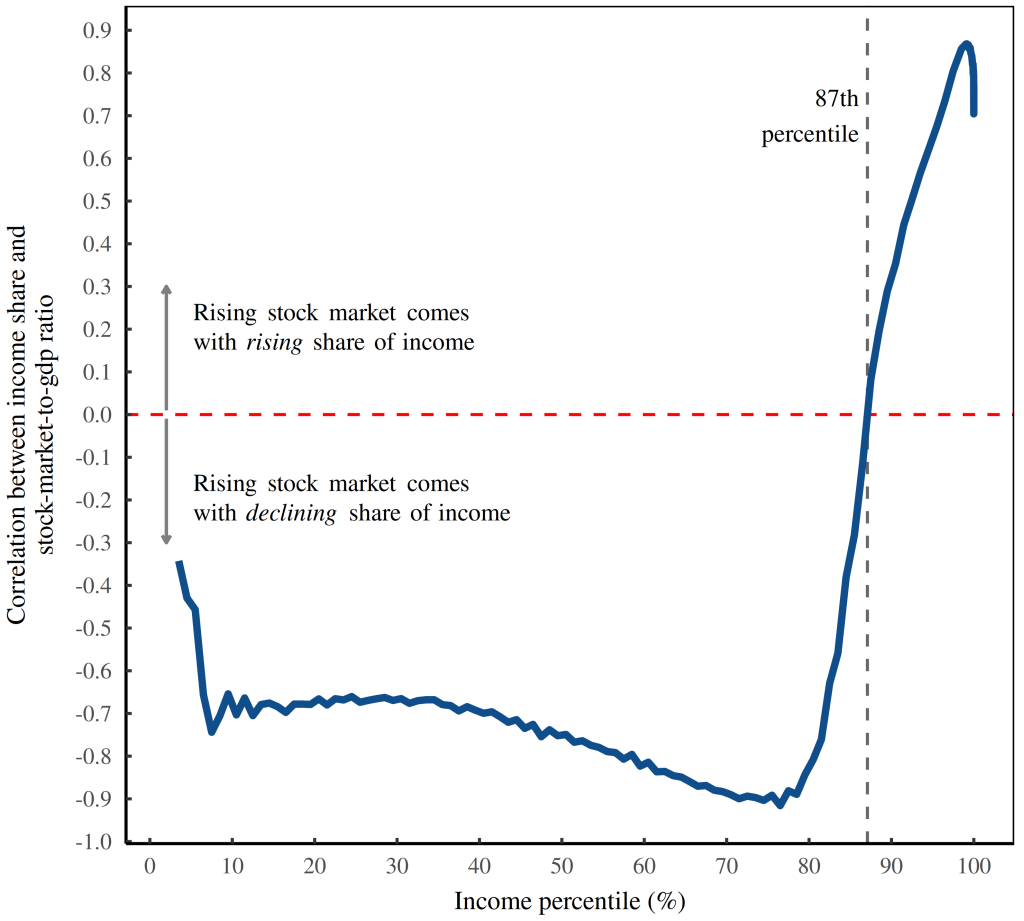

Figure 8: Stock-market pain and gain as a function of income percentile

This figure illustrates how income gets redistributed as stocks go up. For each US income percentile (plotted on the horizontal axis) I measure the correlation between income share and the stock-market-to-GDP ratio. The blue curve shows how this correlation varies as a function of income percentile. For the vast majority of Americans (the bottom 87%) the correlation is negative, meaning stock-market gains *harm* their share of income. It's only among the top decile where things turn positive. Sources and methods

---

The blue line shows the correlation between income share and the stock-market-to-gdp ratio, measured as a function of income percentile. Notice that for the bottom 87% of people, this correlation is negative. That means stock-market gains came with a *declining* share of income.

Let's say that again. For nearly 9 out of 10 Americans, the stock market is a tool for *clawing back* their share of the pie. But don't worry. Their loss is someone else's gain. Among the top 10% of earners, rising stock prices are wildly beneficial, upping their share of the pie.

Switching from income to wealth makes the story even more scandalous. Figure 9 relays the illicit details.[4]

Again, the L-shaped pattern indicates that when the stock market rises, most boats get sunk, while a few luxury yachts float even higher. But compared to income, the sinkage of wealth is even more extreme. When the stock market rises, a whopping *96%* of Americans see their share of wealth decline. But don't worry, for the top 4%, everyone else's pain is their gain.

From the many to the few

Returning to our starting point, we set out to look at the stock market and ask *cui bono*: who benefits? We now have our answer. In the United States, the stock market takes wealth (and income) from the many and hands it to the few.

Now, I'm personally not surprised by this pattern. But I suspect that for many Americans, the detrimental nature of stock-price gains might be shocking. In particular, I'm thinking of members of the professional class — the folks who are not rich, but who still devoutly read Bloomberg. My guess is that when stocks go up, these folks cheer.[5]

Funny. Unless these professionals are in the top 4% wealth bracket, the evidence suggests that they're celebrating on a sinking ship.

[4]Looking at Figure 9, you might wonder why the blue line doesn't extend much below the 40th wealth percentile. The reason is that below the 40th, these folks have a share of the wealth pie that is unchanging — it is constantly *zero*. When you run a correlation on this stream of zeros, you get a value that's undefined. Hence below the 40th percentile, there's no correlation to plot.

[5]In his book *Disciplined Minds*, Jeff Schmidt argues that the professional class are among the most of indoctrinated of groups. I think there's something to this argument. Professionals wield a fair amount of power, but for the most part, are left out of the windfall that flows to society's real owners. So it's best if these professionals believe in ruling-class doctrines.

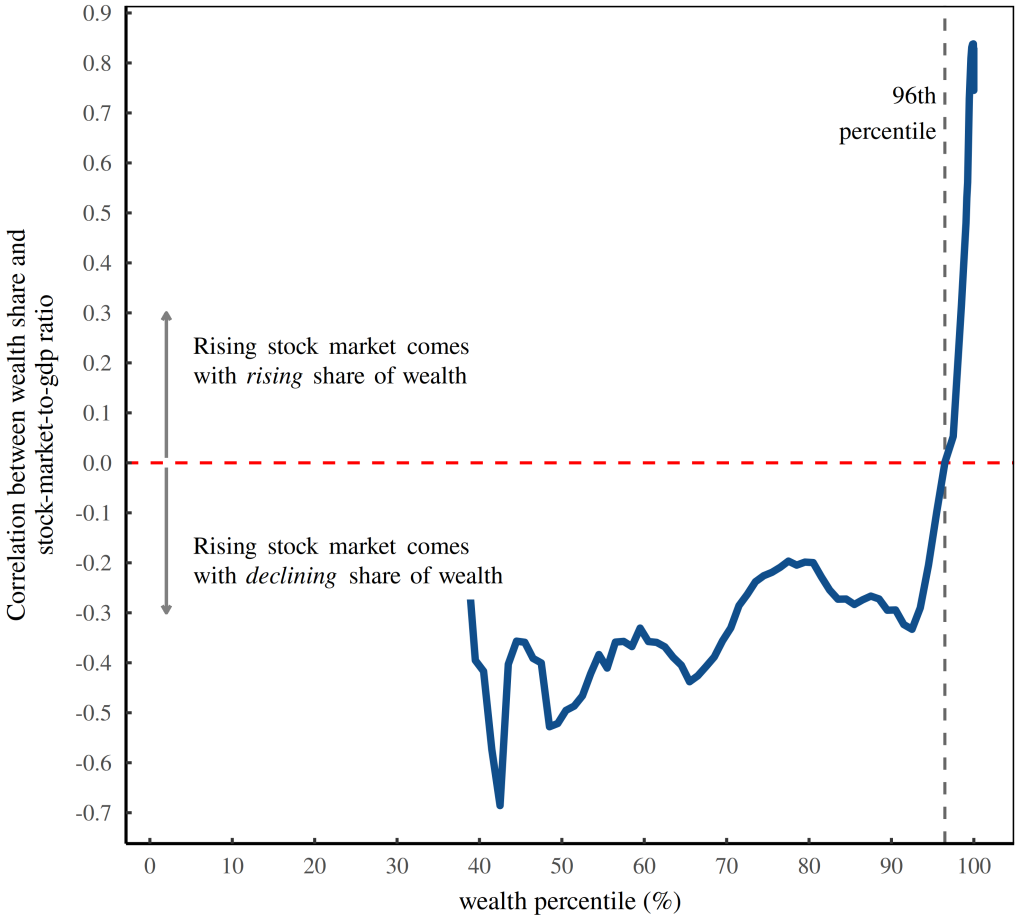

Figure 9: Stock-market pain and gain as a function of wealth percentile

This figure illustrates how US wealth gets redistributed as stocks go up. For each US wealth percentile (plotted on the horizontal axis) I measure the correlation between wealth share and the stock-market-to-GDP ratio. The blue curve shows how this correlation varies as a function of wealth percentile. Here's the message: for the vast majority of Americans (the bottom 96%) the correlation is negative, meaning stock-market gains *harm* their share of income. It's only among the top top 4% where things turn positive. Sources and methods

---

Support this blog

Hi folks. I'm a crowdfunded scientist who shares all of his (painstaking) research for free. If you think my work has value, consider becoming a supporter.

Become an ETD supporter

Sources and methods

S&P 500

Historical data for the S&P 500 is from Robert Shiller, and can be downloaded here.

If you're planning on using this data, here's a warning. It has a weird YYYY.MM notation for dates. For example, January 2001 would be 2001.01. And December 2001, would be 2001.12

Do you see the problem? If you read this date data into any stats program (in my case, R), it's going to think that the month value is a *decimal*. But of course, it's not.

Fun fact: I learned this lesson the hard way. A few years ago, I wrote a piece call 'Stocks are up, wages are down: What does it mean?'. In Figure 1 of that post, I plotted Robert Shiller's data *without* knowing that the month data was getting misinterpreted. The resulting plot actual has a cool 'digital' look. But the monthly data is plain wrong. (The long-term trend is still correct.)

I only realized my mistake when I reran my old code for this post. Live and learn.

US nominal GDP per capita

Data for US nominal GDP is from:

- 1871–1928: Historical Statistics of the United States, Series Ca10
- 1929–2021: Bureau of Economic Analysis, Table 1.1.5
- 2021–2023: quarterly GDP per capita data from FRED, series A939RC0Q052SBEA.

Data for US population is from:

- 1871–1959: Historical Statistics of the United States, Series Aa7
- 1960–2021: World Bank, series SP.POP.TOTL

US share of income/wealth by percentile

All data is from the World Inequality Database. Income data is from series `sptincj992`. Wealth data is from series `shwealj992`.

Stock market null hypothesis

I generated the null-hypothesis data (plotted Figure 7) using the bootstrap method. I randomly sampled from my US stock-market and income-distribution data, and then computed the correlation.

The reason we get a flatline correlation (across all income percentiles) is that I allow dates to be mismatched. For example, the stock-market-to-gdp ratio in 1997 might be randomly matched to income-share data in 1973. The effect of this random mismatching is to remove any time-based correlation, ensuring that the movement of the stock market has no relation to the distribution of income.

Expanding the evidence

In Figures 8 and 9, I used data from 1962 onwards. That's because data for the US distribution of wealth only goes back to 1962. However, data for the distribution of *income* goes back to 1913. So what happens if we extend the analysis backwards another 50 year?

The answer is that we get much the same result. Figure 10 shows the pattern. Here, the blue curve plots, for each income percentile, the correlation between income share and the stock-market-to-gdp ratio for the period 1913 to 2021. The pattern is quite similar to what we saw in Figure 8, in that stock-market gains harm most people's share of income, while helping a small minority. But notice the weird 'elephant curve' — the right portion of the curve that rises, then drops, then rises again. What's the reason for this pattern?

I'm not being rhetorical. I seriously have no idea why this elephant-curve pattern exists.

Further reading

Bichler, S., & Nitzan, J. (2016). A CasP model of the stock market. *Real-World Economics Review*, (77), 119–154.

Schmidt, J. (2001). *Disciplined minds: A critical look at salaried professionals and the soul-battering system that shapes their lives*. New York: Rowman & Littlefield.

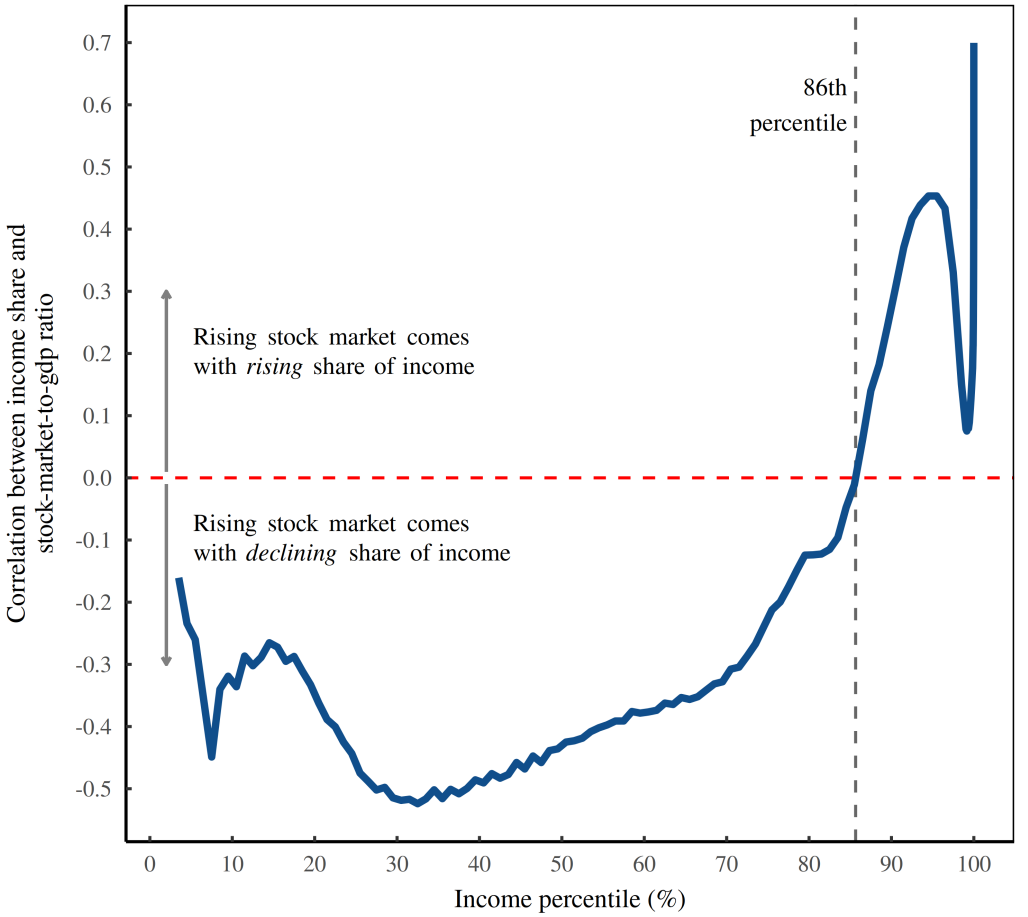

Figure 10: Stock-market pain and gain as a function of income percentile — extending the data back to 1913

This figure illustrates how income gets redistributed as stocks go up. For each US income percentile (plotted on the horizontal axis) I measure the correlation between income share and the stock-market-to-GDP ratio. The blue curve shows how this correlation varies as a function of income percentile. The analysis is conceptually identically to Figure 8, except that I've extended the time-frame back to 1913 (instead of to 1962).