

Economics from the Top Down

new ideas in economics and the social sciences

Interest Rates and Unemployment: An Underwhelming Relation

Blair Fix

April 30, 2023



For the last few months, I've been studying the distributional effects of interest-rate hikes. There's been no shortage of surprising results.

In this post, I'll discuss an effect that is surprising because it's underwhelming. Many economist claim that when interest rates rise, unemployment will increase. The idea is that higher rates make businesses tighten their belts, leading to less hiring and greater unemployment.

Looking at the evidence, I find that this claim is not particularly compelling.

Interest rates and unemployment in the United States

Although economists love to dazzle with fancy econometric models, I'm more of a here-is-the-data-let's-look-at-it person. If a simple scatter plot fails to show something interesting — and by 'interesting', I mean a relation that slaps you in the face — then I'm usually apathetic about further analysis. In other words, doing econometrics with a muddy scatter plot feels to me like polishing a turd.

To that end, when economists claim that higher interest rates worsen unemployment, my first impulse is to simply plot the data. And since the United States has the best historical data, I usually start there.

On that front, Figure 1 shows the long-term relation between the US unemployment rate and the US bond yield. (If you're unfamiliar, the bond yield is a common proxy for the rate of interest.) Staring at the two time series, you




Figure 1: US interest rates (proxied by the bond yield) fail to explain unemployment

This figure looks for a connection between US interest rates and unemployment. The main plot shows the last 130 years of US history. The red line plots the US unemployment rate. The blue line shows the US bond yield — a proxy for the rate of interest. The inset panel shows the scatter plot between the two series. With an R^2 of 0.01, the relation is statistically underwhelming.

can tell that there's not much of a connection. This intuition is confirmed by the inset scatter plot, which is a textbook example of statistical mud. Based on the R^2 value, we can say that the movement of bond yields explains about 1% of the rise and fall of unemployment. Impressive that is not.

A nexus is born

Since so much of economic theory is built on US observations, you have to wonder why economists insist that higher interest rates worsen unemployment. After all, a simple scatter plot tells us that this claim is underwhelming. Am I missing something?

First, let's make sure that I'm not making a straw man argument. Do economists actually claim that higher interest rates worsen unemployment? A quick internet search says yes. Here's a sample of such claims:

[H]igher interest rates tend to lower consumer spending and business investments, leading to a reduction in hiring and an increase in unemployment.

(CUPE)

When interest rates go up, it can have a negative impact on employment. Increased borrowing costs will likely lead to higher unemployment rates.

(Kathryn Underwood)

With interest rates going up, unemployment numbers also tend to rise.

(Eric Rosenberg)

So today, many economists think that interest-rate hikes lead to more unemployment. But have they always thought so? The answer seems to be no.

Now, if I was a judicious scholar, I'd read hundreds of papers and trace the evolution of the interest-rate-unemployment thesis. But to be honest, the prospect bores me. I'd rather let a bot do the work for me. And to be fair (to me), Google's text-scraping bots are far more thorough than I could ever be.

When we look at the text that Google has scraped, we find that the interest-rate-unemployment thesis seems to have emerged in the late 1970s. Figure 2A runs the numbers. Here, I've headed to the Google English corpus and measured the frequency of the phrase 'interest rates and unemployment'. Prior to the 1970s, it seems that virtually no one thought to connect these two phenomena. But by the mid-1980s, loads of people were making the connection.




Figure 2: The rise of the interest-rate-unemployment nexus

This figure traces the connection between interest rates and unemployment to a particular moment in US history. Panel A shows the frequency of the phrase ‘interest rates and unemployment’ in the Google English corpus. (I’ve included in this measurement the frequency for the conjugate phrase ‘unemployment and interest rates’.) The phrase exploded in popularity during the 1980s. Around the same time, there was a shift in how US interest rates related to unemployment. Panel B quantifies this shift by plotting the slope of the trailing 30-year regression between unemployment and the US bond yield. When this slope is zero, bond yields don’t respond to unemployment. But when this slope is one, bond yields show a one-to-one response to unemployment. [Sources and methods](#)

So if our modern internet writers were being forthright, they'd say that *today*, economist think that higher interest rates create more unemployment. But before 1970, hardly anyone made this connection.

Putting on our sleuth hat, we have to wonder why the theoretical zeitgeist changed.

One possibility is that in the 1970s, the 'truth' was discovered. In other words, higher interest rates had always led to greater unemployment, but this fact wasn't discovered until 1970(ish). While this scenario is conceivable, the statistical mud shown in Figure 1 makes me discount it.

Another possibility is that there is no general relation between interest rates and unemployment. Instead, there are periods when interest rates tend to move with unemployment, making it *look* like rate hikes cause unemployment. Conversely, there are periods where interest rates have nothing to do with unemployment (and so no one thinks to connect the two phenomena).

In Figure 2B I test this second possibility. Here's what I've done. The blue line shows the slope of a trailing 30-year regression between US unemployment and the US bond yield. In simple terms, this slope tells us how bond yields relate to unemployment over the preceding thirty years. When the slope is *zero*, bond yields don't respond to changes in unemployment. But when the slope is *one*, bond yields have a one-to-one reaction to unemployment.

Looking at Figure 2B, we can see that the slope of our trailing regression bears an eerie resemblance to the word frequency data plotted above. Prior to 1970, there was essentially no connection between unemployment and bond yields. But from 1980 to the late 2000s, there was a one-to-one connection. In other words, just as bond yields started to move with unemployment, economists began to connect unemployment with the rate of interest.

Given economists' penchant for reactionary fads, it seems plausible that the interest-rate-unemployment nexus is not a general truth. Instead, it may have been a theoretical reaction to a transient period in US history.

Interest rates and unemployment across countries

If my 'reaction thesis' is correct, then when we look at the broadest scale, there should be no pattern between interest rates and unemployment. Speaking of broad scales, the World Bank has extensive cross-country data for both the rate of interest and the rate of unemployment.

The nice thing about this international data is that it reaches extremes that are unheard of in the US. For example, in the World Bank database, lending interest rates cover a range of over 100 percentage points. And the spread in unemployment is similarly large, ranging from a low of 0.1% in Qatar (in 2019) to a high of 38% in Lesotho (in 1997). In short, if higher interest rates create greater unemployment, this database ought to reveal the effect.

Yet when we stare at the cross-country data, it is shockingly unimpressive. Figure 3 tells the story. Here, I've plotted the cross-country trend between unemployment and the lending rate of interest, compiled using data from 132 countries observed over the last thirty years. Looking at the blue line, you can see a slight uptick in unemployment as interest rates climb from 0.5% to 100%(!). But this uptick is statistical mud.¹

¹Playing with the raw interest-rate-unemployment data in Figure 3, a log-log regression returns a p -value of 0.13. Other forms of regression fair worse. While I don't think much of p -values in general, I do think that big ones tell you to put down the data you're playing with and move on to something else.




Figure 3: Across countries, interest rates fail to explain unemployment

This figure shows the international relation between interest rates and unemployment. (Note that the horizontal axis uses a log scale.) To show the trend across countries, I've binned the data by inflation rate. Each blue point shows the midpoint of a bin. The blue line shows the median unemployment rate. The shaded region indicates the middle 50% of data. [Sources and methods](#)

A lagged effect?

At this point, I'm ready to quit. But from experience, I know that if I stop here, I'll regret it.

You see, when it comes to monetary policy, economists have been taught that the effects come with lags that are 'long and variable'. So if I *don't* do a lag analysis, I'll get an endless stream of requests to 'lag the data'. Let me preempt that torture.




Figure 4: Changes in the US bond yield predict jumps in next-year’s unemployment

Taking data originally plotted in Figure 1, this chart looks at the relation between changes in the US bond yield (horizontal axis) and changes in next-year’s unemployment (vertical axis). The red line shows the smoothed trend. The red shaded region shows the associated uncertainty in the trend, which is significant. [Sources and methods](#)

In Figure 4, I’ve returned to the US data and done a lagged analysis of bond yields and unemployment. The horizontal axis shows the annual change in the US bond yield. The vertical axis shows the annual change in unemployment in the following year. Surprisingly, we get a positive trend. In other words, today’s interest-rate hikes seem to *increase* next-year’s unemployment.

The caveat here is that the lagged trend is produced in large part by a few outlier years, all of which are in the late 1970s and the early 1980s. Of course, what constitutes an ‘outlier’ is mostly a matter of taste. Still, from Figure 2,

we know that the 1980s stand out as a unique period when US interest rates moved consistently with unemployment. In other words, it's dubious to take a pattern from that decade and pronounce it a 'general tendency'.

Things get worse when we realize that a lagged effect doesn't mean much on its own. That's because when we're dealing with cyclical data, we'll inevitably find that an observation today predicts an observation later. For example, tonight's sunset predicts tomorrow's dawn. Why? Because the Earth has a reliable rotation.

Matters are similar (although less reliable) with unemployment. Looking at Figure 1, you can see that US unemployment rises and falls with a fairly regular frequency, corresponding to what economists call 'business cycles'. Because of these cycles, today's unemployment will predict unemployment in the future.

In the case of a one year lag, it's fairly easy to understand what we'll observe. In the US, unemployment oscillates with a roughly 8-year cycle. In that context, a one year lag represents an eighth of a cycle. If we do the math, we find that an uptick in this year's unemployment should be followed by *another* uptick next year. In other words, changes in unemployment this year ought to correlate positively with unemployment changes next year. And indeed they do.

Figure 4-self shows the pattern. Here, the blue curve shows the self-correlation between US unemployment changes today and those same changes next year. As expected, the pattern is positive — a fact that doesn't bode well for the idea that interest-rate hikes worsen unemployment. As the red curve shows, the bond-yield 'treatment effect' is pretty much the same as the relation between unemployment and itself.²

(For a longer explanation of why lagged effects should exceed self-correlation, see [this post](#).)

When we leave the US behind and look at the lagged pattern found across countries, the picture grows even muddier. Figure 6 shows what happens.

²In Figure 4-self, the bond-yield effect is slightly stronger than the unemployment self-correlation. But keep in mind that for both effects, the uncertainties are large — as indicated by the red and blue shaded regions. The fact that the red curve barely ever leaves the blue shaded regions says, in a statistical sense, that the bond-yield 'treatment effect' is consistent with no effect at all.




Figure 5: US bond yields fail to affect US unemployment more than unemployment ‘affects’ itself

This figure measures the lagged relation between the change in interest rates (proxied by the bond yield) and the change in unemployment in the United States. The red curve shows the average trend when we take the annual percentage-point change in the US bond yield and compare it to next-year’s percentage-point change in unemployment. (For the raw data, see Figure 4.) The blue curve shows the autocorrelation between unemployment and itself — the annual change in unemployment compared to next-year’s annual change in unemployment. [Sources and methods](#)

Here, the red curve shows the effect that a rise in interest rates has on next-year’s unemployment. Even on its own terms, the effect is lack-luster. And when we compare this ‘treatment effect’ to the self-correlation effect — the relation between changes in unemployment today and the same changes next year — things really fall apart. The self-correlation is significantly stronger than the ‘treatment’ effect. In short, when we look at the lagged trend across countries, there’s no evidence that interest-rate hikes worsen unemployment.




Figure 6: Across countries, interest rates fail to affect unemployment more than unemployment affects itself

Looking across countries, this figure measure the lagged relation between the change in interest rates and the change in unemployment in the United States. The red curve shows the average trend when we take the annual percentage-point change in interest rates (within a country) and compare them to next-year's percentage-point change in unemployment. The blue curves shows the autocorrelation between unemployment and itself — the annual change in unemployment (within a country) compared to next-year's annual change in unemployment. Note that I've binned the data by values on the horizontal axis. The points indicate median change within each bin. The shaded regions indicate the middle 50% of data. [Sources and methods](#)

Is full employment always good?

I have to admit that I find these results disappointing. Although I've learned to take economists' pontifications with a boulder of salt, my intuition was that interest rates would connect with unemployment And yet the evidence suggests otherwise.

That said, there are ways connect interest income to unemployment — ways that are better supported by evidence. In my next post, I'll discuss Jonathan Nitzan and Shimshon Bichler's concept of the '[maturity of capitalism](#)', which pits interest income against profit income. It turns out that unlike interest rates, the interest-to-profit income *ratio* is related to unemployment.

But for now, back to interest rates and unemployment.

Karl Marx famously argued that unemployment is a tool of class warfare. The unemployed, Marx noted, are a 'reserve army' that bids down the price of labor, to the obvious benefit of capitalists. Viewed in this light, it seems like the unemployment rate should relate to the distribution of income.

On that front, my [last post](#) made a big deal about how interest rates are a 'distributional variable' that ratchets up or down the income shares of different groups. If unemployment is also a 'distributional variable', then it seems plausible that interest rates might related to unemployment. And yet they seemingly don't.

So what went wrong?

To some extent, I think the problem lies with the concept of unemployment itself. Obviously, mass unemployment is bad (as during the Great Depression). But the flip side is that low unemployment isn't necessarily 'good'.

Take the example of Qatar. In 2019, it had an unemployment rate of just 0.1% — a rate so low that there was effectively 'full employment'. In other words, everyone who wanted a job had one. So Qatar sounds like a good place to work, right?

Wrong.

In 2019, Qatari workers took home just 25% of the income pie.³ (The rest presumably went to property owners.) So despite having full employment, Qatar had an extremely despotic distribution of (class-based) income. Why?

The answer is found not in the unemployment rate, but in the nature of Qatari 'employment'. You see, Qatar is a nation built on indentured servitude. Exceedingly rich from oil money, few Qatari citizens work for a living. Instead, they import an army of foreign laborers to do the dirty work. (Migrant workers constitute about [95%](#) of the Qatari labor force.)

³Data for the Qatari labor share of income is from ILOSTAT series [SDG_1041_NOC_RT_A](#), Labor income share as a percent of GDP.

The effect of this arrangement is that unemployment is basically impossible. If a foreign worker loses their job, they lose the right to stay in Qatar. So yes, Qatar has ‘full employment’. But by that, we mean that Qatar has a workforce that *cannot* be unemployed.

The lesson here is that the scale of ‘unemployment’ doesn’t necessarily tell us anything about the welfare of the working class. Hell, if you’re willing to bend categories, you could say that slave states have ‘full employment’. But that’s because ‘unemployment’ is impossible. In a slave state, you can either be a *slaver*, a *slave*, or a *fugitive*. There is no unemployment.

The point is that measurements must take into account the social order that they are quantifying. On that front, if we divide society into the ‘employed’ and the ‘unemployed’, we’re excluding a third category: the *shittily* employed. As big corporations increasingly turn to ‘flexible’ labor to do their bidding, we’d best pay attention to this third category.

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.



Sources and methods

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 US unemployment

- Unemployment from 1890 to 1946: Historical Statistics of the United States, [Table Ba475](#).
- Unemployment from 1947 to 2020: Bureau of Labor Statistics, series [LNU04000000](#)

Text frequency

Word frequency data is from the 2019 Google English corpus, downloaded with the excellent R package [ngramr](#).

International interest rates

International 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.

International unemployment

- International unemployment data is from the World Bank, series [SL.UEM.TOTL.ZS](#)

Interestingly, when describing this unemployment data, the World Bank warns that low unemployment isn't necessarily 'good':

Paradoxically, low unemployment rates can disguise substantial poverty in a country, while high unemployment rates can occur in countries with a high level of economic development and low rates of poverty. In countries without unemployment or welfare benefits people eke out a living in vulnerable employment.

Further reading

Standing, G. (2011). *The precariat: The new dangerous class*. New York: Bloomsbury.