Economics from the Top Down

new ideas in economics and the social sciences

Interest Rates and Inflation: Knives Out

Blair Fix

February 19, 2023

If you're just tuning in, I've spent the last few months debunking some common misconceptions about inflation:

- *Is inflation a uniform increase in prices?*
  No. Inflation is wildly differential.
- *Is inflation driven by an 'over-heated' economy?*
  No. Inflation tends to come with economic stagnation.
- *Do higher interest rates reduce inflation?*
  No. Higher interest rates are associated with higher inflation.

As expected, the last claim put mainstream economists into war mode. You see, the belief that interest rates down-regulate inflation has come to be sacred. So by scrutinizing this idea with evidence, I was effectively torching an effigy of the pope. (Novelist Cory Doctorow helped fan the flames by writing an incendiary essay about my research.) And so I spent a 'fun' week on Twitter being bombarded by econo-scorn.

Now back to science. When I published 'A Test of Monetary Faith', I had more evidence in the pipelines — evidence that debunks the idea that interest rates down-regulate inflation. In this post, I'll wade through the data.

The take-home message is clear: when we look at the World Bank database, there is no evidence that higher interest rates down-regulate inflation. If anything, the evidence suggests that rate hikes make inflation worse.

But before we get to the data, I'll respond to some of the more cogent criticisms that economist hurled my way.

Fisher, we call thy name

In 'A Test of Monetary Faith', I documented the fact the interest rates rise and fall with inflation. This covariation is not a new discovery.

My point, however, was to frame this evidence in terms of the science of regulation. If interest rates *down-regulate* inflation, then a bump in interest rates ought to come with a drop in inflation. And yet that's not what we find. Everywhere we look, we find that interest rates and inflation are positively coupled — just like the concentration of atmospheric CO_2 is tightly coupled with the Earth's temperature. From this evidence, there are two obvious conclusions. Either inflation up-regulates interest rates, or interest rates up-regulate inflation.

In response to these arguments, many economists resorted to namedropping. '*This is just the Fisher effect*,' they said.

What fun! I can also play the namedropping game. Economists are exhibiting the '*give-it-a-name-so-I-don't-have-to-think-about-it*' effect.

You see, the 'Fisher effect' is not some deep explanation for the empirical pattern that I highlighted. It's just a name given to the pattern itself. The 'Fisher effect' is the strong positive correlation between interest rates and inflation. It's named after the economist Irving Fisher who, in 1907, first discovered the relation.[1]

Tellingly, Fisher was interested in explaining interest rates (not inflation). And there, the pattern makes sense. When inflation rears its head, everyone fights to raise their income. Workers seek better wages. Firms pad their prices. And what do creditors do? Creditors hike the rate of interest. It's all rather obvious.

What does *not* make sense is what economists did with Fisher's observation. First, they canonized the 'Fisher effect' as a kind of 'natural law'. (They argued that interest rates automatically move with inflation, thus preserving

[1] A historical side note. Irving Fisher was an influential economist whose 1907 book *The Rate of Interest* helped consolidate the neoclassical understanding of capital. Fisher's task was to show how a stream of income could be traced back to a quantity of capital. His solution was to devise a feedback loop that seemed plausible, but was actually untestable, since the quantity of capital couldn't be measured independently of income. Jonathan Nitzan and Shimshon Bichler call the resulting logic "Fisher's house of mirrors". Given this history, it's fitting that economist continue to invoke Fisher's name to defend their dogma.

the 'real' rate of interest.)[2] Decades later, economists had the idea to flip the evidence on its head and claim that interest rates *down-regulate* inflation. The result was a kind of perpetual cognitive dissonance. Economists accepted the obvious evidence that inflation up-regulates interest rates. But they also insisted that lurking within the data was difficult-to-find evidence that interest rates down-regulate inflation.

The purpose of 'A Test of Monetary Faith' was to highlight this cognitive dissonance. So it's telling that economists think that invoking the 'Fisher effect' is a counter-argument. It's not. It's like staring at an asteroid hurtling towards Earth and dismissing calls for action as misplaced: 'Don't worry about the collision. It's just the Newton effect.'

A 'real' house of mirrors

Another criticism of 'A Test of Monetary Faith' was that I looked at the *actual* rate of interest — i.e. the number on your mortgage bill that determines your debt-servicing costs. Economist, however, have developed doublespeak for discussing reality. They call the *actual* rate of interest the 'nominal' rate. And they call the hypothetical rate of interest — one that has been adjusted for inflation — the 'real' rate.

Now, when it comes to explaining inflation, it's supposed to be the 'real' rate of interest that matters. However, I refuse to study this rate. Why? Because if we use the 'real' rate of interest to explain inflation, we're invoking circular reasoning. More on that in a moment.

First, some general principles of science. When you test a scientific theory, you start by interrogating its assumptions. If the assumptions are bullshit, the theory is bullshit. End of story. Now, when it comes to neoclassical economics, the assumptions are always bullshit. However, they also tend to be a sleight-of-hand that make the theory self-fulfilling.

Here's an example. Neoclassical economists claim that wages are proportional to productivity. What is their evidence? Well, it's the correlation between firms' average wage and their sales per worker. Now the assumption here is that sales per worker measures 'productivity'. But in reality, sales per worker is just another form of *income*. So the supposed 'test' of economic

[2] Wages also rise and fall with inflation. But for some reason, economists never thought to canonize this movement as a 'natural law'. How come the 'free market' automatically preserves the income of creditors, but lets workers fall on their face? It's almost as if economists have a class bias.

theory rests on correlating two forms of income — forms that we expect to correlate by virtue of an underlying accounting definition. In other words, if we accept that sales per worker measures 'productivity', we've already bought into the neoclassical theory of income.

So what does this productivity quagmire have to do with explaining inflation using 'real' interest rates? It's actually quite similar. Economists claim that inflation is a function of the 'real' interest rate. While this claim sounds reasonable, it comes with a hidden problem. You see, the 'real' interest rate *contains* the thing that it's supposed to *explain* — namely, inflation.

Let's illustrate the problem using a toy model. Imagine a society in which inflation has a simple cyclical pattern, as shown by the blue curve in Figure 1A. In this society, interest rates are constant. The correct conclusion is that inflation variation has nothing to do with interest rates.

Now let's get 'real' interest rates involved. To calculate the 'real' interest rate, we take the actual interest rate and subtract the rate of inflation. Doing so gives the red curve shown in Figure 1B. Comparing the red and blue curves, we find that real interest rates move counter-cyclically to inflation. In other words, it looks like 'real' rates down-regulate inflation.

Looking at the evidence, monetarists claim victory. When the 'real' price of money is low, monetarists explain, people borrow furiously. That causes the money supply to increases, pushing inflation higher. But when the 'real' price of money is high, people stop borrowing, the money supply contracts, and inflation slows. In other words, the data is consistent with monetarist theory. 'Real' interest rates down-regulate inflation.

It all sounds convincing until we realize that we're standing in a house of mirrors. Can you spot the reflection?

By construction, the 'real' interest rate is the *actual* interest rate minus the rate of inflation. If we compare this value back to the rate of inflation, we're introducing autocorrelation. We're correlating inflation with the negative version of itself:

$$\text{inflation} \sim \text{real interest rate}$$
$$\text{inflation} \sim \text{interest rates} - \text{inflation}$$

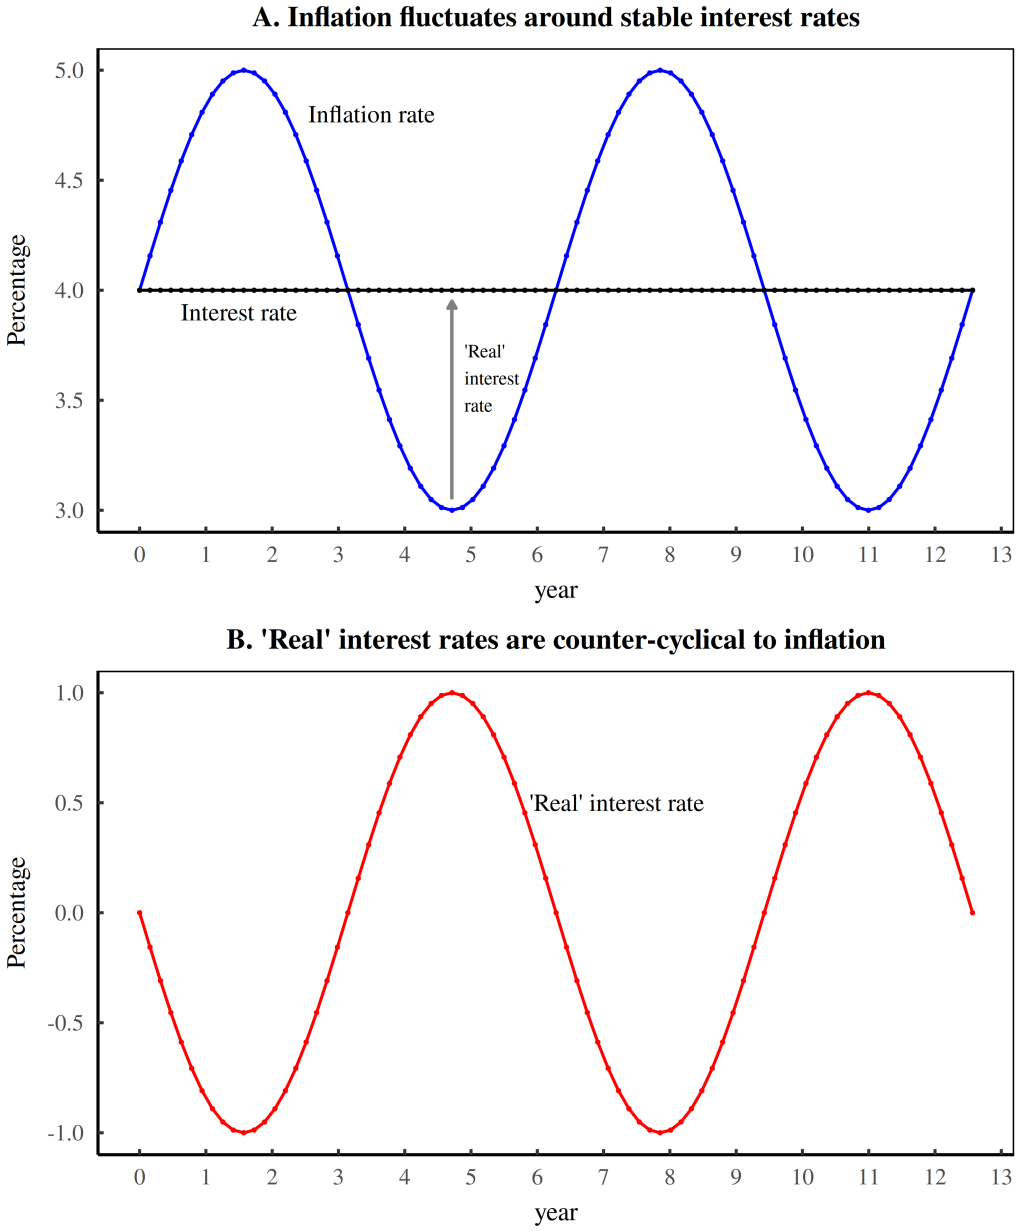

Figure 1: 'Real' interest rates — explaining inflation in terms of itself

This figure illustrates the problem with using 'real' interest rates to explain inflation. In Panel A, we imagine a society in which inflation is perfectly cyclical and interest rates are constant. The message is that interest rates don't influence inflation. In Panel B, we get the 'real' interest rate involved. This rate is the difference between the actual rate of interest and the rate of inflation. When we calculate the 'real' interest rate (shown in red) we find that it moves in the opposite direction of inflation (blue). In other words, the 'real' interest rate seems to down-regulate inflation. However, this finding is an illusion. What we're actually seeing is the effect of autocorrelation — we're correlating inflation with the negative version of itself.

So with 'real' interest rates, we have a recipe for defining our theory to be true. In a world in which *actual* interest rates have no relation to inflation rates, we can use autocorrelation to flip the conclusion. 'Real' interest rates appear to regulate inflation. Except that it's just inflation that is correlating with the negative version of itself.[3]

And that's why I don't play the 'real' interest game. You can't define your causal variable in terms of the thing you're trying to explain. In short, any theory that explains inflation in terms of 'real' interest rates is statistically illiterate.

Take your medicine

Let's move on to another favorite tool use by the econ tribe: 'proof' via analogy. It works as follows. You rescue your theory by imagining a situation in which analogous evidence leads to the wrong conclusion. *Ipso facto*, your theory is 'correct'.

Here's an example. When people get sick, they take medicine. So when we observe their behavior, we find a correlation between 'medicine taking' and 'getting sick'. We (wrongly) conclude that the medicine *causes* sickness. Dramatic pause ... did I mention that interest-rate hikes are the medicine and inflation is the disease? Therefore, rate hikes reduce inflation. *QED*.

Now, the trick here is that we're supposed to assume that the 'medicine' works. And so the correlation between getting sick and taking medicine leads us in the wrong direction. Unfortunately, this conclusion is a non sequitur. In reality, the medicine analogy simply restates what we already know: when people get sick, they do x. But on its own, this knowledge tells us nothing about whether x actually heals the patient.

Here are some examples that illustrate the fallacy. When people get sick, some of them swear by a placebo called 'homeopathy'. And in the Middle Ages, sick people had their blood drained. They called it 'bloodletting', and it killed them faster. The conclusion is obvious. The fact that sick people *use* a 'treatment' is not evidence that the 'treatment' *works*.

[3] Just to be clear, economists *do* compare the rate of inflation to the 'real' interest rate, seemingly unaware that they're introducing autocorrelation. For example, a 1981 paper by Frederic Mishkin finds a 'significant negative relationship between inflation and the real [interest] rate' ... a 'striking' correlation that has 'important implications'.

In more general terms, the problem is this: it's always possible to imagine scenarios in which a treatment 'works', despite the obvious evidence to the contrary. Fortunately, real-world medicine ignores these hypotheticals and instead relies on the precautionary principle. If there's a hint that a treatment causes problems, the drug doesn't get approved.[4] Put another way, doctors don't administer drugs based on their ability to construct hypothetical counter-arguments that the drug might actually work, despite the apparent evidence. Doctors go with the obvious conclusion.

Mainstream economists, however, throw caution to the wind. If they can imagine a scenario in which their interest-rate policy works (despite the evidence to the contrary), then the policy is sound. The 'medicine' should be administered in large doses.

Let's translate this thinking into the absurd. Do you see those people drinking cyanide-laced kool-aid? Sure, they all drop dead. But I can imagine a scenario in which it only *looks* like the kool-aid is killing them. Now here's a glass; drink up.

The dose of monetarism

Jonestown jabs aside, economists are correct that interest-rate hikes *could* conceivably lower inflation, even though they appear to do the opposite. On its own, this possibility doesn't justify monetary policy. But perhaps there's evidence that we're missing.

On that front, let's continue the medicine analogy. Suppose that interest-rate hikes are a drug, administered in response to inflation. The 'treatment' requires that we be vigilant with the dosage. Our aim is to tightly couple interest rates to inflation. If we manage to do so, the idea is that over the long term, we should achieve low inflation.

So does the medicine work? To find out, let's turn to Figure 2. Here, I've taken each country in the World Bank database (with relevant data) and done the following:

[4] OK, I'm being too kind to the pharmaceutical industry. They routinely engage in data shenanigans to get their drugs past government regulators.

1. **Measure the 'dose of monetarism'**.[5] To calculate this dose, I take each country and measure the correlation between interest-rate change and inflation change. The more positive the correlation, the more that a country is able to greet rising inflation with a dose of interest-rate hikes. The more negative the correlation, the more the country ignores the doctor's orders.

2. **Measure the 'treatment' outcome**. This is the average rate of inflation within each country. (I calculate the average using the geometric mean, applied over the years with available data). The lower the average rate of inflation, the more successful the 'treatment'.

In Figure 2, I compare the treatment outcome (vertical axis) to the dose of monetarism (horizontal axis). Now, if economists are correct, a larger dose of monetarism should lead to a decline in inflation. So in Figure 2, we should see a downward trend. Unfortunately, we do not.

Looking at the results, we see an *upward* trend. As the monetarist dose increases — as countries better match inflation jumps with interest-rate hikes — we find that the average rate of inflation *increases*. And where do we find the lowest rate of inflation? That would be among countries who flaunt the doctor's orders — countries that greet jumps in inflation with an interest-rate reduction.

In short, the evidence suggests that if the monetarist medicine 'works', it does so by killing the patient.

Lag it till you find it

To summarize, we still have no evidence that interest rates down-regulate inflation. However, economist won't be convinced until we torture the data with lags.

I say 'torture' because when it comes to inflation, lagging the data can be dangerous. You see, interest rates and inflation are both highly cyclical — they rise and fall over regular intervals. Given these cycles, the danger is that if we lag the data 'sufficiently', we can invert the obvious correlation.

[5] Monetarism (Milton Friedman's brainchild) is one of many theories claiming that raising interest rates lowers inflation. But it's the theory that has captured the public's imagination. So I'll use 'monetarism' as a catch-all term for orthodox economic policy — greeting rising inflation with rising interest rates.

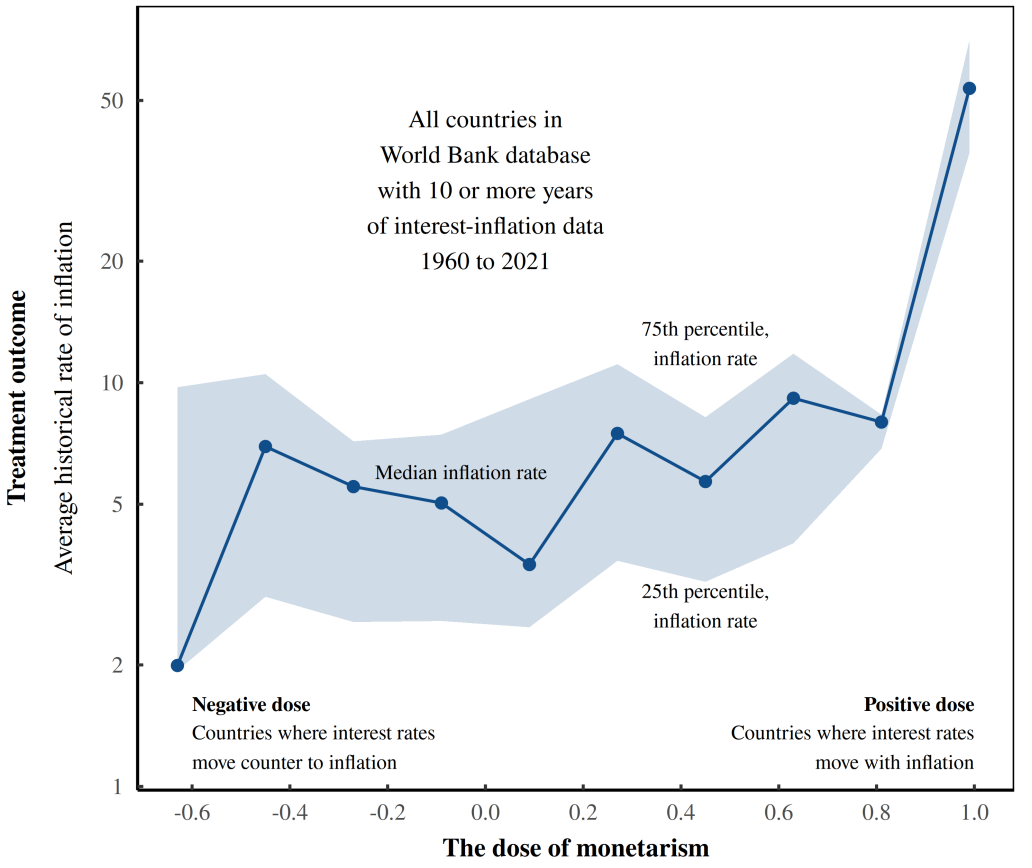

Figure 2: A dose of monetarism kills the patient?

This figure analyzes the effectiveness of orthodox monetary policy — greeting inflation with interest rate hikes. Here's how the analysis works works. For each country in the World Bank database (that has ten-or-more years of relevant data), I measure the annual change in the inflation rate and the annual change in the lending interest rate. Then I measure the correlation between these changes. This correlation is the 'dose of monetarism', plotted on the horizontal axis. The more positive the correlation, the more that a country greets rising inflation with interest-rate hikes. The more negative the correlation, the more that the country flaunts the doctor's orders. After calculating the dose of monetarism, I then measure the treatment outcome — the average rate of inflation within the given country, plotted on the vertical axis. (I calculate the average using the geometric mean.) To display the results, I've binned the data. Each blue point shows the middle of a correlation bin. The blue line then shows the median inflation rate of countries within the bin. The shaded region indicates the middle 50% of data. Sources and methods

---

Figure 3 illustrates the problem. Here, we imagine a world in which inflation is perfectly cyclical, rising and falling over a 6-year interval. In this world, interest rates react to inflation — they mimic the inflation rate, but with a

slight delay. Interest rates, in turn, have no effect on inflation. In other words, if we look for evidence that interest rates down-regulate inflation, we should not find it.

Unfortunately, if we lag the data 'sufficiently', we *can* find (apparent) evidence for down-regulation. Figure 3B shows how to do it. When we lag the inflation data by about 2.5 years, we get a stunning negative correlation. Today's interest rates appear to down-regulate tomorrow's inflation.

The problem with this method should be self evident. When data is highly cyclical, we can always use lags to invert the obvious correlation. Here are some dubious examples. Want to 'show' that the sun *cools* the Earth? Take sunlight data and correlate it with temperature data 12 hours later. Want to 'show' that pressing your car's brakes speeds up the vehicle? Lag the speed data by the duration of a stop light.

In these scenarios, data lag is a recipe for confirmation bias: you lag the data until you find what you want. In more subtle terms, the problem is not the lag itself, but the freedom to apply it indiscriminately. It's part of a wider problem that statisticians call 'researcher degrees of freedom' — toying with your method until the data 'confesses'. It's a recipe for faulty results.

Fortunately, there's a simple method that solves the problem: *lock down your lag*.

Let's apply the lock-down method to the example of the sun seemingly cooling the Earth. The data 'confesses' when we turn night into day — i.e. when we let temperature data lag 12 hours behind the sunshine data. The trouble is, if we lock down this lag, it fails on other planets. On Mercury, the data would not confess. (Mercury has a 1400-hour day.) Nor would the data confess on Venus (which has a 5800-hour day). In fact, our 12-hour lag would fail everywhere but on Earth. So by locking down our lag, we turn our 'great results' into mud. That's good science.

Unfortunately, this falsification is not great if you want to push an ideology. To preserve your doctrines, it's best to change your methods as the need arises. You insist that the sun *does* cause the Earth's temperatures to fall. But the effect has lags that are 'long and variable'.

Does this language sound familiar? It should. It's exactly how economists describe the lags associated with monetary policy … they are 'long and variable'.

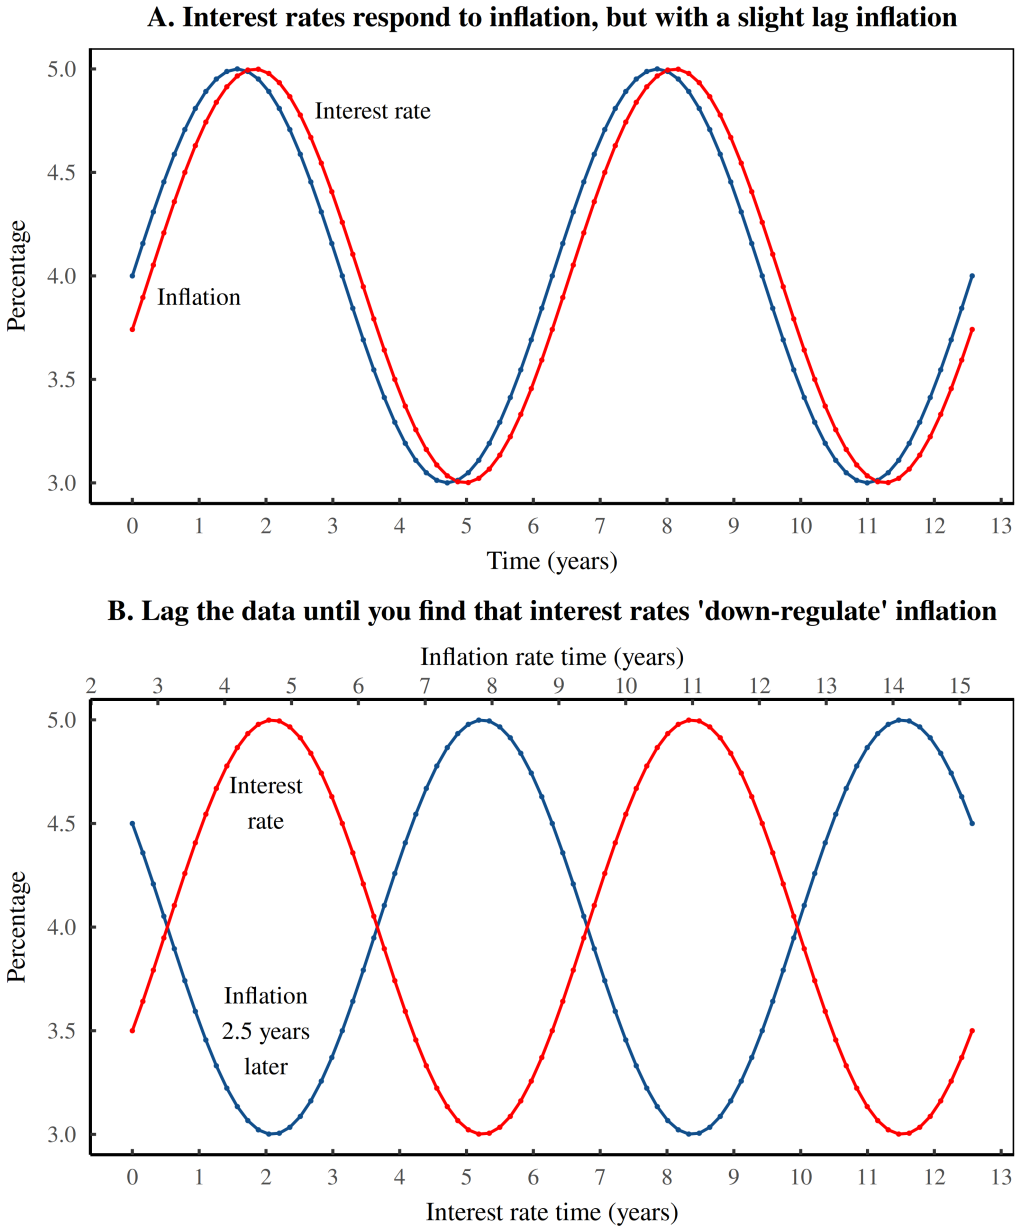

Figure 3: Time lag as a tool for confirmation bias

This figure illustrates the perils of using lagged analysis to infer causation. Here, we imagine a society in which inflation is perfectly cyclical. In response to inflation, bankers raise and lower interest rates, but with a slight delay. Interest rates, in turn, have no effect on inflation. Panel A shows the resulting pattern. The correct interpretation is that inflation up-regulates interest rates. As 'clever' economists, however, we want to show that interest rates down-regulate inflation. To do that, we lag the inflation data by about 2.5 years, as shown in Panel B. By introducing this lag, we get a perfect negative correlation, suggesting that today's interest rates down-regulate tomorrow's inflation. Unfortunately, this conclusion is spurious, since the underlying model contained no down-regulation. What actually happened is that we used time lag as a tool for confirmation bias. We lagged the data until we got the results we wanted.

This verbiage is a fascinating testament to faith-based policy. I suspect that ancient soothsayers would have used similar excuses. As they raised their knife for a ritual slaughter, they'd say to the audience: 'This sacrifice *will* bring rain. But beware; the delay can be long and variable.'[6]

Untangling up-regulation

Having ridiculed economists' use of lags, it's my turn to lag the inflation data. But I'll begin by discussing what lags can and cannot do.

Crucially, lags *cannot establish causation*. Instead, they can only demonstrate the *lack* of causation. For example, if the sun rises before a rooster's crow, it follows that the rooster did not cause the dawn. However, the dawn's precedence does not mean that it caused the rooster to crow. Precedence is not causation.

[6]Of course, 'long and variable' delays can occur in established science. But what matters is that we can *predict* the lag. For example, if a star goes super nova — spewing radiation into the galaxy — the effect will have a long and variable lag. The reason is well understood.

Radiation propagates at the speed of light, which is fixed everywhere. So the effect of the super nova will be felt after time $t = d/c$, where d is your distance from the super nova and c is the speed of light. So before we even look at the data, we can predict the lag. That's hugely different than post-hoc smoke blowing.

If we insist that a theory make clear predictions, then looking for lags (or lack thereof) is a great way to do science. Here's a fascinating example. Some alternative theories of gravity predict that gravitational waves will propagate at a speed that is slightly different than the speed of light. This is a testable prediction that has been falsified. It's a stunning scientific achievement. Here's how it happened.

On August 17, 2017, the LIGO and VIRGO gravitational wave detectors picked up a strong gravitational signal and were able to triangulate its location to a small portion of the sky. Around that time, the Fermi gamma-ray satellite detected a signal from the same part of the sky. After-the-fact analysis suggests that the two detections came from the same event — thought to be the merger of two neutron stars.

Now, the light signal was detected about 2 seconds after the gravity signal — a lag that may sound like a significant difference. However, the signals came from a host galaxy that was pegged to be 130 million light years away. Over that distance, even the tiniest deviation in gravitational wave speed would lead to an enormous lag in the signal. (For instance, if gravitational waves travelled at 99.99% of the speed of light, they would arrive about 13,000 years after the light signal.)

After doing the math, physicists estimated that gravitational waves travel at the same speed as light, accurate to about one part in a quadrillion. This result then ruled out a host of alternative-gravity theories. For a nice summary of this research, see 'Reining in Alternative Gravity'. The scientific paper lives here.

Back to inflation. As I documented in 'A Test of Monetary Faith', interest rates rise and fall with inflation. (The evidence for this pattern is overwhelming.) Given this positive coupling, we want to know the cause. There are two obvious possibilities:

1. **Inflation up-regulates interest rates.**
2. **Interest rates up-regulate inflation.**

Before proceeding, I should note that these scenarios are not mutually exclusive. For example, it could be that creditors respond to inflation by hiking interest rates. And these rate hikes might then exacerbate inflation.[7]

Having proposed two causal directions, a lagged analysis can tell us which direction is more likely and/or more dominant. (Or to be rigorous, the lag points to the less-likely cause.) With that in mind, let's make some predictions.

To start with, we *know* that rate hikes are a response to inflation. There is no subtlety here. As Ryan Ortega observes, rate hikes come with much pomp and circumstance, broadcast via well-telegraphed meetings designed to pound the policy home. Rate-hike announcements are about as subtle as a declaration of war. And as with a declaration of war, the only confusing part is the official message.[8] Rate hikes are a response to inflation, yes. But in the long run, the claim is that higher interest rates reduce inflation. At least, that's what we're supposed to believe.

Moving on, let's discuss the opposite scenario — the idea that rate hikes drive inflation. Here, we can immediately rule out the most extreme version of the thesis. Suppose that rate hikes are the *sole* driver of inflation. And suppose that creditors always respond to inflation jumps with further rate

[7] Oddly, when it comes to *wages*, many economists think that mutual causation is self evident. They call it a 'wage-price spiral' — inflation leads to wage increases, which lead to more inflation, and so on. But when it comes to *interest* payments (which, like wages, are a cost of doing business), the thinking gets turned on its head. In response to inflation, interest rates rise. But the added debt-servicing costs don't stoke the inflationary fires. No, they somehow quash the flames.

[8] In war, it matters who attacks whom. If *we* attack *them*, we usually call it a 'liberation'. But if *they* attack *us*, that's an 'invasion'. And so it goes with responses to inflation. When *workers* demand a raise to shore up their standard of living, it's a cash-grab that will worsen inflation. But when *creditors* demand a raise (via rate hikes), it's an inflation-fighting sacrifice for the good of the many. Fortunately, as with war, we can investigate inflation outcomes without listening to the official propaganda.

hikes. The results would be catastrophic. The slightest rate bump would lead to a runaway feedback loop ending in hyperinflation. Clearly that doesn't happen. (Or if it does happen, it is rare.)

Having dismissed interest-rate hikes as the sole cause of inflation, it seems plausible that rate bumps might make inflation worse. The idea (articulated nicely by Tim Di Muzio) is that any form of cost creep puts upward pressure on prices. For example, if my suppliers raise prices, I try to pass those costs along to my customers. If my employees raise their wages, I pass those costs along as well. And if my *creditors* raise the price of debt (by hiking interest rates), I pass those costs along too.

So the idea is that any form of cost creep exacerbates inflation. That said, the various cost categories have different sizes. In particular, interest payments are usually a small portion of businesses' overall costs.[9] As such, interest-rate hikes should put a small upward pressure on prices.

Wrapping up, if I were to make predictions about the lag analysis that follows, I'd say that future interest rates will respond strongly to today's inflation. And if future inflation responds to today's interest rates, the effect will be relatively small.

Looking at the data, that's more or less what we find.

Do interest rates respond to inflation?

First up, we'll test the hypothesis that interest rates react to inflation. To do that, we'll return to the World Bank data, which contains interest-rate and inflation data for most countries. (The data covers roughly the last sixty years, although the exact coverage varies by country.) We'll measure how changes in inflation relate to changes in next year's interest rates.

Figure 4 shows the pattern across countries. Here, I take each country and compare the annual change in inflation to the annual change in interest rates *in the following year*. The blue line plots the average pattern across all countries. As expected, the trend is positive, suggesting that next-year's interest rates respond to today's inflation.

[9] Interest payments are not always small. When countries becomes heavily indebted, interest payments can become a substantial part of total expenses. In that case, rate hikes may become a major driver of inflation.

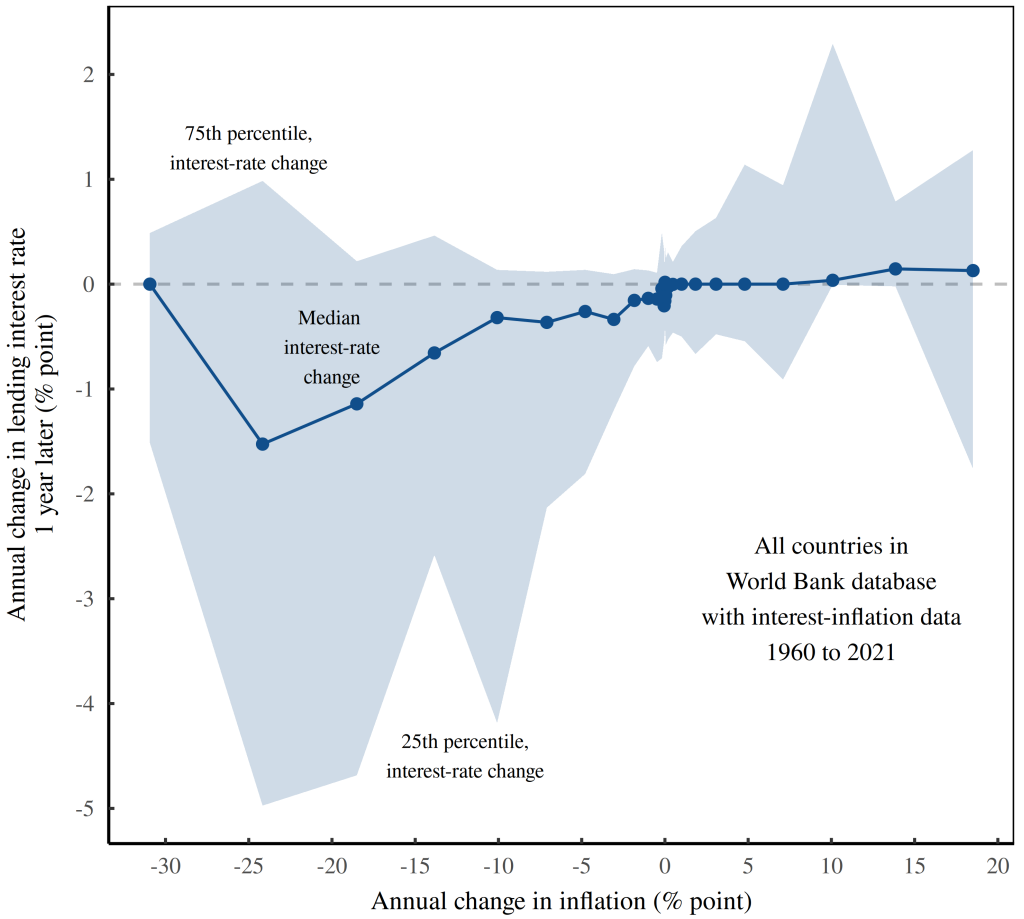

Figure 4: Changes in inflation precede changes in interest rates — the pattern across countries

This figure tests the hypothesis that interest rates respond to inflation. The chart analyzes the international relation between the annual change in inflation and the annual change in lending interest rates *in the following year*. To illustrate the trend, I place inflation changes into bins, plotted on the horizontal axis. The blue points indicate the midpoint of the bin. (Note that the size of the bins increases with the absolute value of inflation change. The goal here is to account for the spread in the data. Most inflation changes are small, but a few are quite large.) Within each bin, I then measure the range of interest-rate change. The blue line shows the median change. The shaded region indicates the middle 50% of the data. Sources and methods

---

When we switch to the correlation *within* countries, we find a similar pattern. Overall, changes in inflation correlate positively with next year's change in interest rates. Figure 5 shows the results. Here, the histogram illustrates the distribution of within-country correlations — the correlation between

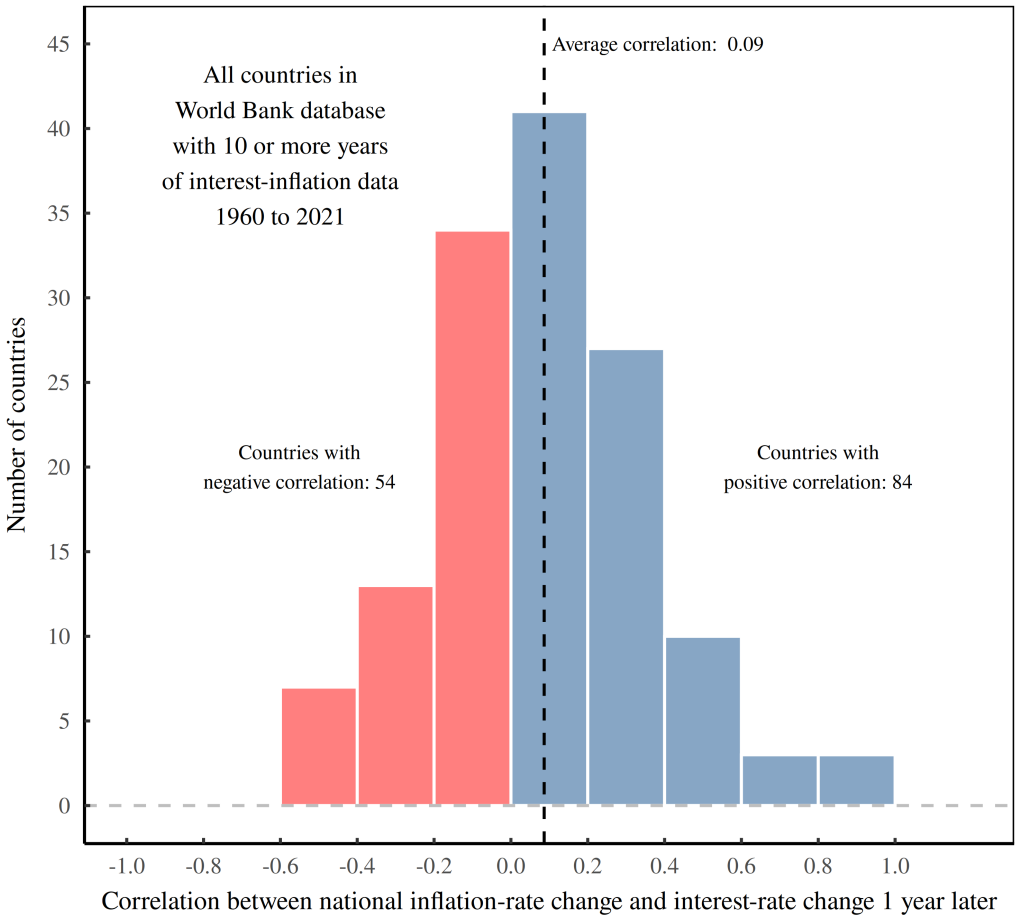

Figure 5: Changes in inflation precede changes in interest rates — the correlation within countries

This figure tests the hypothesis that interest rates respond to inflation. The chart analyzes the within-country connection between inflation changes and interest-rate changes *in the following year*. I start with all countries in the World Bank database with ten or more years of interest rate and inflation data. Within each country, I then measure the correlation between the annual change in the inflation rate and the annual change in interest rates in the following year. The histogram then shows the distribution of correlations. Countries with a negative correlation are shown in red. Those with a positive correlation are shown in blue. Sources and methods

---

the inflation-rate change and the interest-rate change *in the following year*. The correlations are largely positive, suggesting that rate hikes respond to inflation.[10]

[10]Note that in Figures 4 and 5, the lagged effect is weaker than the non-lagged effect. (For the non-lagged effect, see Figure 7 and Figure 8 in 'A Test of Monetary Faith'.) This weaker effect suggests that if interest rates respond to inflation, they mostly do so within a year.

Does inflation respond to interest rates?

So far, the evidence suggests that interest rates respond to inflation. Now let's test the reverse scenario — the idea that inflation responds to interest rates.

We'll repeat the lag analysis from above, but this time flip the delay. We'll correlate annual interest-rate changes to annual inflation changes *in the following year*. Figure 7 shows the resulting pattern across countries.

Immediately, we see that the trend is muddier than in Figure 5. Oddly, it seems that large interest-rate drops are associated with *decreases* in next-year's inflation. (That's not something we'd expect if interest-rates down-regulated inflation.) But in the rest of the plot, there's no clear pattern. (If you squint, you might see a downward trend. More on that in a moment.)

Turning to the pattern *within* countries, we find much the same thing. Figure 7 shows the results. Here, the histogram plots the distribution of within-country correlations between the annual interest-rate change and the inflation-rate change *in the following year*. Again, the relation is muddy. The average correlation is trivially small — about -0.04. (Yes, the average correlation is negative. I'll address that shortly.)

To summarize, the evidence in Figures 4–7 suggests that interest rates largely respond to inflation. Or in more strict terms, the evidence falsifies the idea the interest rates are the sole driver of inflation. Of course, that's what we expected. Central bankers shout at us that they're raising rates in response to inflation. The evidence suggests that we should believe them.

Now to the elephant in the room. When we let inflation lag behind interest rates (Figures 6 and 7), we get a slight downward trend. Looking at this pattern, I can hear economists proclaiming victory. So let me put words in their mouth:

> A negative correlation! In your lagged analysis, you found a negative correlation between interest-rate changes and changes in next-year's inflation. Interest rates down-regulate inflation! Our theory is vindicated!

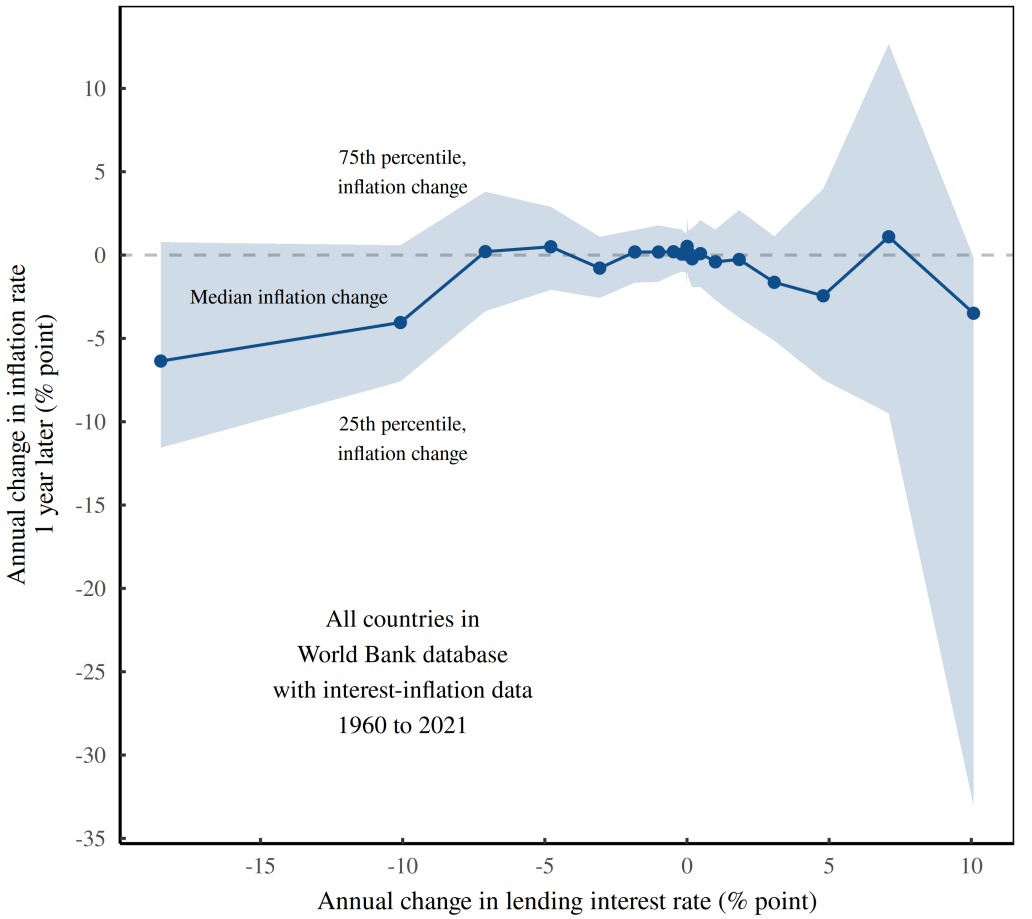

Figure 6: Change in interest rates precede changes in inflation — the pattern across countries

This figure tests the hypothesis that inflation responds to interest rates. The chart analyzes the international relation between the annual change in lending interest rates and the annual change in inflation *in the following year*. To illustrate the trend, I place interest-rate changes into bins, plotted on the horizontal axis. The blue points indicate the midpoint of the bin. (Note that the size of the bins increases with the absolute value of interest-rate change. The goal here is to account for the spread in the data. Most interest-rate changes are small, but a few are quite large.) Within each bin, I then measure the range of inflation-rate change. The blue line shows the median change. The shaded region indicates the middle 50% of the data. Sources and methods

---

Let me now respond. On the face of it, the evidence *does* seem to justify monetary orthodoxy. But that's only because economists are accustomed to doing junk science. As you'll see, when we do rigorous science, we find that the small negative correlation in Figure 7 is not evidence for down-regulation. If anything, it indicates that hiking interest rates makes inflation worse.

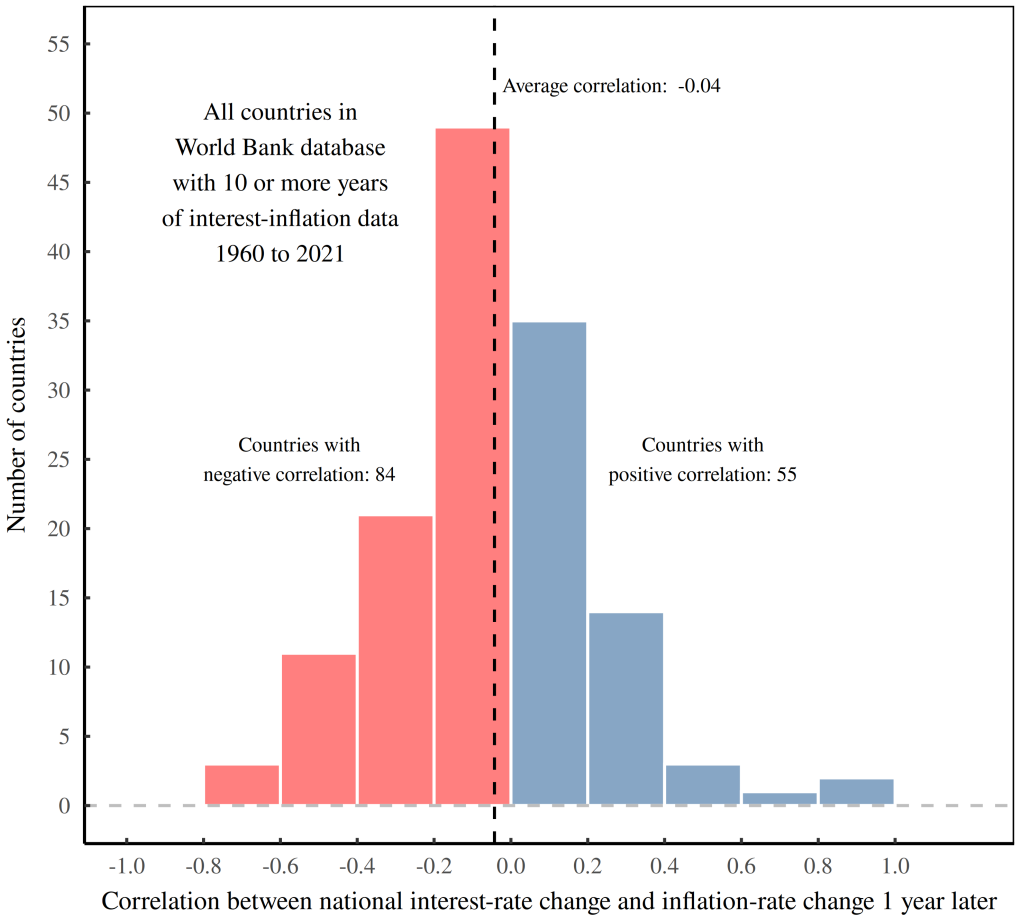

Figure 7: Change in interest rates precede changes in inflation — the correlation within countries

This figure tests the hypothesis that inflation responds to interest rates. The chart analyzes the within-country connection between interest-rate changes and changes in inflation *in the following year*. I start with all countries in the World Bank database with ten or more years of interest rate and inflation data. Within each country, I then measure the correlation between the annual change in the lending interest rate, and the annual change in inflation in the following year. The histogram then shows the distribution of correlations. Countries with a negative correlation are shown in red. Those with a positive correlation are shown in blue. Sources and methods

---

Pass the cigarettes

At this point, you think I'm crazy. So let me explain my reasoning.

When we lag our inflation data, the purpose is to test the direction of *up-regulation*. (Does inflation up-regulate interest rates? Or do interest rates up-regulate inflation?) Given the evidence, it seems clear that interest rates largely respond to inflation.

Now here's where things get confusing. While we're testing the direction of up-regulation, the inflation lag returns a (tiny) negative correlation. Mainstream economists will likely jump on this negative correlation and claim that it is evidence for *down-regulation.* But if they do, they're revealing their scientific ignorance. That's because you cannot use the results of a test to *change the thing that you are testing*. Our lag analysis tests the direction of up-regulation. As such, it says nothing about down-regulation. If we want to actually test for down-regulation, we have to use a different method.

And that brings me to cigarettes. You see, every time I get a cold, I smoke cigarettes like fiend.[11] And do you know what? After about ten days of this treatment, my cold disappears. So clearly, cigarettes cure colds.

Or not.

As you can guess, my conclusion is dubious because it invokes the fallacy of precedence. Yes, my dose of cigarettes precedes my cold recovery. But that doesn't mean my 'treatment' worked.

The problem is obvious. The human body comes equipped with an immune system that fights colds all by itself. So regardless of the 'treatment' I administer, a cold will typically resolve itself in about a week. In other words, if we want to test the cold-healing efficacy of cigarettes, we can't take the recovery as evidence in itself. We have to compare the cigarette-treatment to the baseline recovery pattern.

There are two basic ways to do make this comparison. The gold standard is the clinical trial, in which we give one group of people the treatment, and the other group a placebo. If the treatment doesn't beat the placebo, it's garbage.

When we cannot run trials, the other option is to study observational data. Here, it's not clear what a 'placebo' would be, since people always do *something* when they get sick. As such, a good option is to use the general population as the baseline. We take a large sample of people and measure how long it takes (on average) for them to heal from a cold. Whatever the number, that's the baseline that a proposed treatment must beat. So in my example, we'd test if a cigarette 'treatment' out-performs the baseline for a cold recovery. (I'd guess that it doesn't.)

[11] No, I don't treat my colds with cigarettes. But short of bloodletting, it's the stupidest 'treatment' I could think of.

With this method in mind, let's return to inflation. To test for down-regulation, we can't just look at how inflation behaves after the interest-rate 'treatment'. That's junk science. Instead, we must compare the interest-rate 'treatment' to the baseline effect. If the treatment doesn't beat the baseline, it's garbage.

Inflation: A cancer or a cold?

So what is the baseline inflation effect? To frame the question, let's return to the big picture. In broad terms, there are two ways of understanding inflation: it's either a *cancer* or a *cold*.

In mainstream economics, inflation is portrayed as a cancer. The fear is that capitalist societies are always on the knife edge of hyperinflation. If inflation is not constantly exterminated (via interest-rate medicine) the cancer will take hold and prices will head to the stratosphere. So in this scenario, the baseline effect is runaway inflation.

An alternative view is that inflation is more like a cold. Sure, in rare cases it can lead to severe illness. Unusually, though, the cold resolves itself. Why? According to political economists Jonathan Nitzan and Shimshon Bichler, it's because inflation is self defeating. The idea here is that inflation is not about 'too much money chasing too few goods'. Instead, inflation is a business strategy in which businesses expand their income by raising prices.

Now the key is that this greater income is not shared equally. When businesses raise prices, they tend to keep the profits for themselves. Yes, some of the added income trickles down to workers. But not enough to offset rising prices. As a result, when inflation rears its head, workers see their purchasing power erode, and so they're forced to buy less stuff. When workers can no longer afford the items whose prices are being hiked, the inflation cycle ends.

In this second scenario, the baseline effect is for inflation to by cyclical. Just like a cold, inflation resolves itself, regardless of what we do.

And that brings me back to cigarettes.

What if greeting inflation with a dose of rate hikes it is like greeting a cold with a dose of cigarettes? The appropriate test is not whether the symptoms resolve. The real test is whether the outcome is better than the baseline effect — the results when we do *anything* to treat inflation.

Here we get to the crux of the problem. When it comes to inflation, the 'anything-goes' effect is huge.

Measuring the anything-goes effect

Having played the analogy game, let's now talk statistics. When we're dealing with time-series data, the 'anything-goes' effect is the data itself. Or more precisely, it is the relation between the data *now*, and the data at some point in the *future*. The strength of this relation tells us about the cyclicity of the data.[12]

Figure 8 illustrates the basic principle. Here, the left-hand panels show hypothetical time series for inflation. The right-hand panels show the 'anything-goes' effect — the correlation between inflation now and inflation one year later. The more cyclical the inflation data, the tighter this self correlation.

In example A, inflation is highly cyclical, giving rise to a strong self correlation. In example B, the inflation data is less cyclical, so the self correlation is weaker. And in panel C, the inflation data is acyclical, causing the self correlation to vanish.

Noting these examples of self correlation, let's continue our exposition. When we study cyclical data (like inflation), we must imagine that it covaries with many other phenomena. Here's an extreme analogy. Life on Earth exists on a rotating sphere. Unsurprisingly, many biological cycles tend to mirror the sphere's rotation. But it would be foolish to use lagged data to conclude that biological cycles *cause* the Earth to spin.

So how can we avoid fooling ourselves?

That's where the 'anything-goes' effect comes in; it helps us avoid drawing dubious conclusions from cyclical data. In the case of the Earth's spin, regardless of how life scurries about, the planet's rotation strongly predicts its future state. So if we want to claim that a biological pattern influences the Earth's rotation, that biological pattern must predict future rotation *better* than the Earth itself. It's a tremendous hurdle to pass.

[12] Statisticians often call this self prediction 'autocorrelation' — the ability of a time series to predict a delayed copy of itself. This nomenclature is useful in the sense that 'auto' is a synonym for 'self'. However, 'auto' also connotes things that are 'automatic' or 'inevitable', and I find that misleading. There nothing inevitable about 'autocorrelation' in the sense used here. It is a function of data cyclicity. The more cyclical the data, the greater the self correlation.

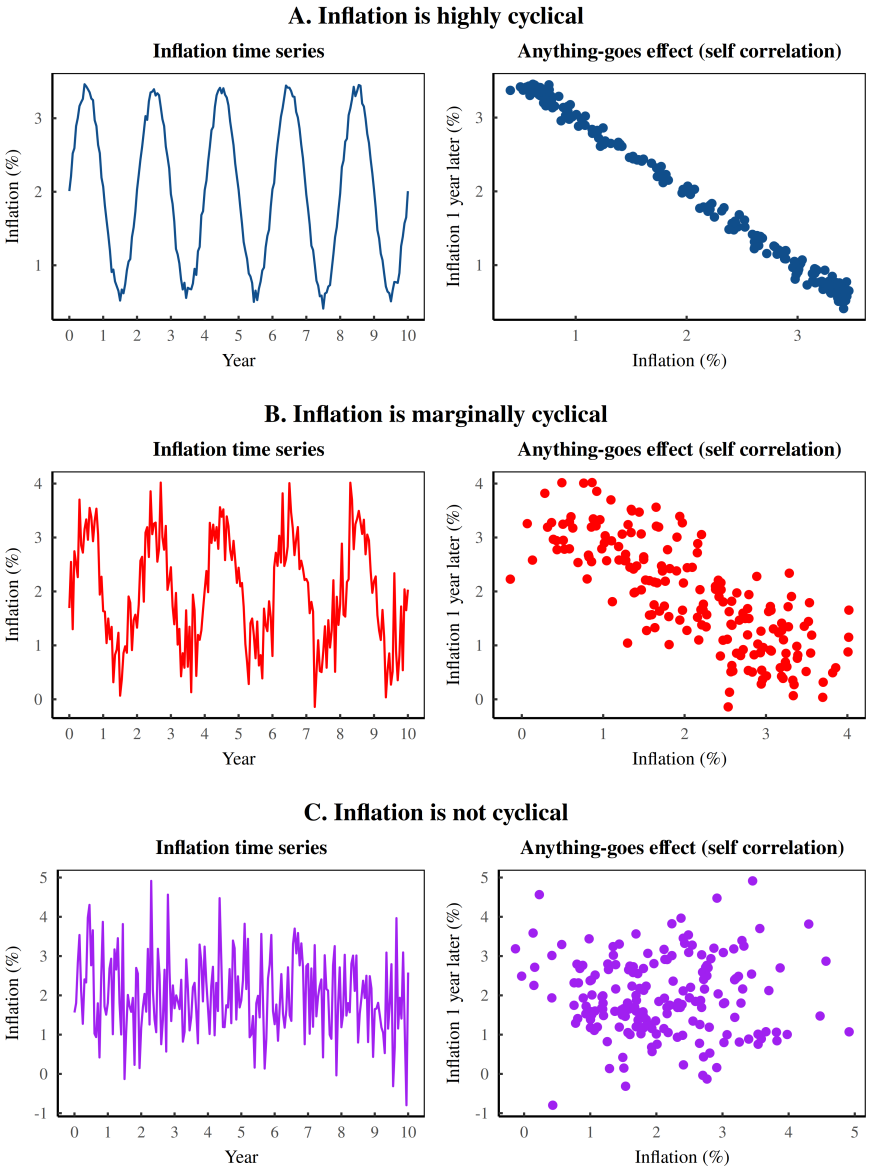

Figure 8: Measuring self correlation — the 'anything-goes' effect

This figure illustrates how the cyclical nature of inflation data affects its ability to predict itself. I call this self-prediction the 'anything-goes' effect, because it assumes that many different phenomena covary with (and hence, potentially cause) the observed inflation cycles. What interests us not whether these phenomena can predict the inflation cycles, but whether they can do so better than the *inflation data itself*. Crucially, this prediction hurdle depends on the cyclicity of the observed data. When inflation data is highly cyclical (example A), inflation accurately predicts its future state, one year later. However, as the inflation data becomes noisier (example B), the self correlation declines. In the limit that the data has no cyclical pattern (example C) the self correlation vanishes. When we do a lagged analysis, this anything-goes effect is the threshold to beat. Put another way, it is the expected outcome of a null hypothesis — the pattern we expect if a variable co-varies with inflation but has no effect on it. If a lagged effect does not surpass the anything-goes effect, we can rule out causation.

When we turn to inflation, the same thinking holds. Many phenomena covary with inflation. If we're to believe that one of these phenomena is actually *driving* future inflation, its lagged effect must exceed the anything-goes effect. Importantly, this anything-goes threshold is not fixed. It depends on the cyclicity of our data.[13]

When data is highly cyclical, as in Figure 8A, a lagged effect must pass a huge threshold to be convincing. That's because anything that co-varies with inflation will automatically predict future inflation. In contrast, when the data is acyclical — as in Figure 8C — any kind of prediction about future behavior is impressive. That's because the inflation data has virtually no ability to predict its future self. So the hurdle for a lagged effect is small.

Testing the interest-rate treatment

With the anything-goes effect in hand, we're ready to test the efficacy of treating inflation with interest-rate hikes.

Let's begin by reviewing the results so far. When we lag the inflation data, we find that a rise in interest rates is associated with a small reduction in inflation one year later (Figures 6 and 7). Now we're going to ask — does this pattern exceed the anything-goes effect? If so, it indicates that interest rates actually down-regulate inflation. If not, it suggests that the treatment effect is spurious — an artifact of the covariation of cyclical data.

Looking at the evidence across countries, we find that (with a few exceptions) the treatment effect is *identical* to the anything-goes effect. Figure 9 shows the pattern. Here, the red curve is the treatment effect. This curve is what happens to the change in inflation one year after a change in interest rates. (The data is replotted from Figure 6.)

Now, it's tempting to interpret the treatment curve in isolation. But that's a mistake. Interest rates covary with inflation, and both data series are highly cyclical. As such, we need to account for the anything-goes effect — the effect of inflation on the lagged version of itself. When we include this effect (the blue line), we find that the treatment outcome is (mostly) consistent with a

[13] Here's another way to think about the 'anything-goes' effect. It's the maximum *apparent* effect of a null hypothesis. Suppose that A and B are perfectly correlated over time. Suppose also that B has no effect on A. Despite the lack of real effect, we may find that B predicts future states of A. However, this result is an apparent effect only — it is created by the ability of A to predict its future self. The more cyclical the data, the greater this apparent, or 'null' effect. In other words, the more cyclical the data, the more easily we can fool ourselves with lagged analysis. Buyers of interest-rate analysis beware.

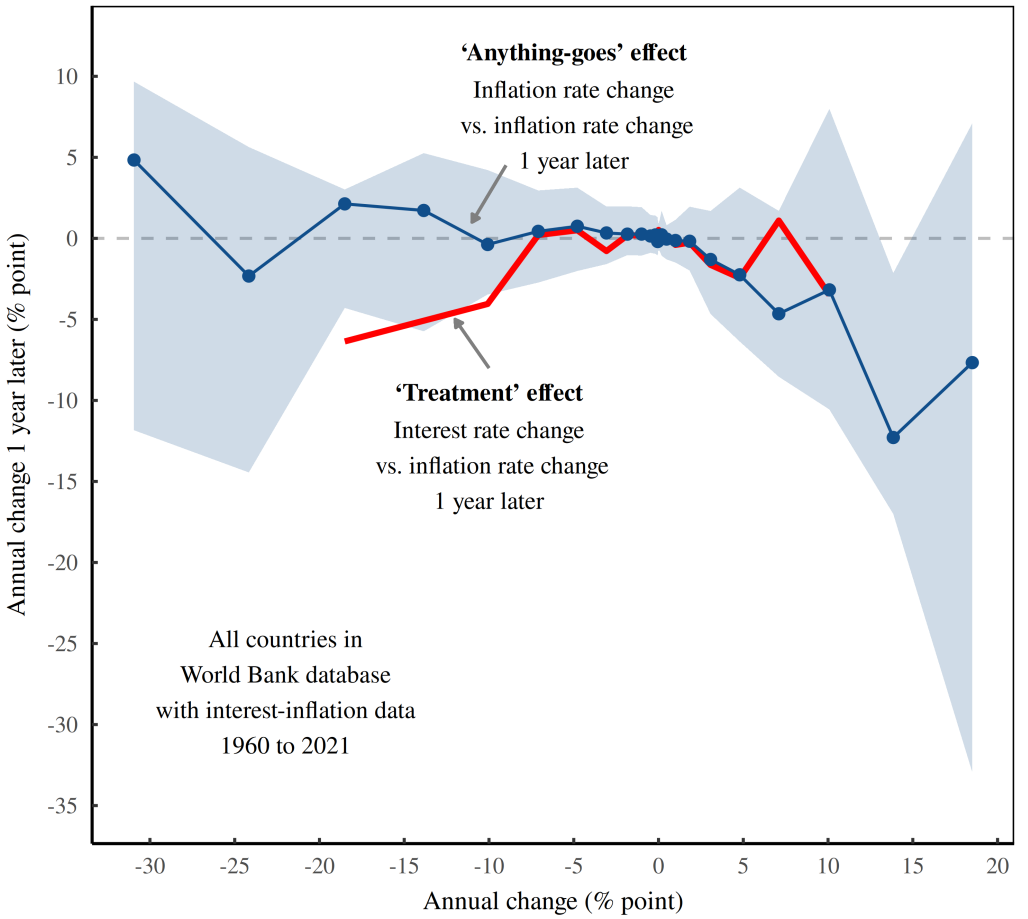

Figure 9: Comparing the interest-rate 'treatment' to the 'anything-goes' effect — the pattern across countries

This figure analyzes the effectiveness of 'treating' inflation with interest rate hikes. The red curve is the 'treatment' effect — the results of treating inflation with interest-rate hikes. (It replots the data from Figure 6 — the cross-country relation between interest rate change and inflation change one year later.) To put the treatment effect in context, I compare it to the 'anything-goes' effect (blue). The anything-goes effect measures the ability of the inflation data to predict itself. To calculate it, I measure how annual inflation change relates to the lagged version of itself, one year later. With a few exceptions, the results suggest that the apparent treatment effect *is* the anything-goes effect (i.e. the red and blue lines mostly overlap.) Note that all data is binned by annual change on the horizontal axis. The blue line shows the median anything-goes effect. The shaded region indicates the middle 50% of the data. Sources and methods

---

null result. With a few exceptions, the red and blue lines overlap. In other words, the treatment outcome is almost exactly what we expect if interest rates covary with inflation but have *no effect on it*.

Now let's discuss the points where the treatment outcomes differ from the anything-goes effect. Intriguingly, it seems that large drops in interest rates might stimulate a *drop* in inflation — something we don't see in the anything-goes effect. Conversely, there's some evidence that steep interest-rate hikes might make inflation worse (evident where the treatment curve lies above the anything-goes curve).

These exceptions are certainly worth studying. But let's be clear about the entirety of the results: there is no evidence that rate hikes actually reduce inflation. If there were such evidence, we'd see the treatment curve drop below the anything-goes curve as interest rates increased. Yet there is no hint of such a pattern.

Turning to the pattern *within* countries, we find a similar outcome. The interest-rate treatment fails to show its face.

Figure 10 illustrates the results. Here, the red curve replots the treatment outcomes from Figure 7 — the distribution of within-country correlations between interest rate changes and next-year's change in inflation. (Note that instead of using a histogram, I've switched to a density function to illustrate the distribution.)

Beside the treatment effect, I've plotted the anything-goes effect. This is the distribution of within-country correlations between inflation rate changes now and inflation rate changes one year later. The gap between the two outcomes is palpable. The anything-goes effect is strongly negative, indicating that a rise in inflation today is typically followed by a drop in inflation next year. In other words, regardless of how it is 'treated' inflation tends to resolve itself. In contrast, the treatment outcome is much weaker, meaning it is consistent with no effect at all.

An effect that is not an effect

Let's summarizing our results. Looking at Figures 9 and 10, we can conclude that interest-rate hikes *do not reduce future inflation*. There is no evidence that the effect of an interest-rate treatment exceeds the effect of doing 'anything'. Quite the contrary. It seems that the apparent rate-hike effect can be chalked up entirely to self correlation.

Let's review how it works. Because inflation is highly cyclical, it strongly predicts its future state. (If inflation jumps this year, it will likely fall next year.) And since interest rates covary with inflation, they make the same

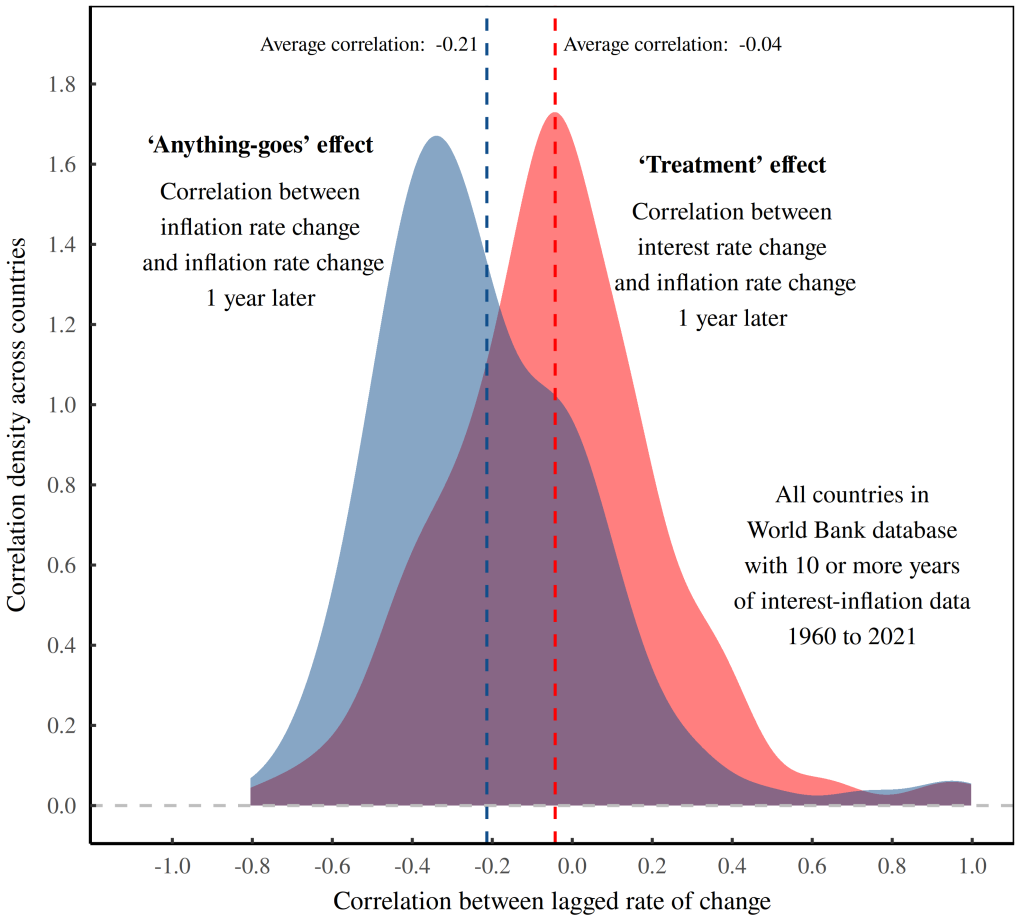

Figure 10: Treating inflation with 'anything goes' — the correlation within countries

This figure analyzes the effectiveness of 'treating' inflation with interest rate hikes. The red distribution replots the treatment outcome from Figure 7. Let's review the method. Within each country that has relevant data, I measure the correlation between interest-rate change and the inflation-rate change *one year later*. The density curve plots the distribution of correlations. The blue curve uses the same method, but measures the 'anything-goes' effect — the correlation between inflation change and the lagged version of itself. To measure this effect, I take each country and measure the correlation between inflation change and lagged inflation change *one year later*. The blue curve illustrates the distribution of correlations. The gap between the treatment and anything-goes effect is easy to spot. Changes in inflation have a strong, negative correlation with themselves. In other words, if inflation rises this year, next year it will likely fall. This self-correlation is far stronger than the treatment effect. As such, the evidence hints that treating inflation with interest-rate hikes may be worse than doing 'anything'. Sources and methods

---

prediction. However, there is no evidence for actual down-regulation. Rather, the interest-rate effect is either *equal to* or *worse* than the anything-goes effect.

Returning to our cigarette analogy, our results are similar to finding that a cigarette treatment fails to outperform the baseline for a cold recovery. Instead, we find that the treatment outcomes may be *worse*; cigarette patients recover in ten days rather than the usual seven.

And so it goes with the interest-rate medicine.

Looking at the data, not only is there no evidence for down-regulation, there is some indication that rate hikes actually make inflation worse. That is the conclusion suggested by the 'dosage' analysis in Figure 2. It's also hinted at by the treatment gaps in Figure 9. And then there's the within-country results in Figure 10, which scream at us that rate-hike outcomes are worse than the anything-thing goes effect.[14]

In short, we can conclude that the interest-rate medicine does not work. It is certainly no better than a placebo. And in some cases, it may actually be a poison.

Economists will continue to prescribe the rate-hike medicine

Having exhaustively examined the data, I find no evidence that interest rates down-regulate inflation. Given this result, we're left with a paradox. Economist vigorously prescribe a 'medicine' that does not work. Why?

In scientific terms, the behavior makes no sense. But in ideological terms, it is easy to understand. By and large, mainstream economics exists to justify the behavior of the powerful. On that front, the idea that interest rates 'reduce' inflation is a convenient veneer for an otherwise self-serving behavior.

Faced with inflation, everyone tries to raise their income. But not everyone has an official ideology backing them up. Tellingly, when workers try to boost their wages, economists scream that it will make inflation worse. Yet when *creditors* do the same thing, economists pass it off as an act of altruism. 'It just looks like a money grab,' economists declare. 'But it's actually a noble medicine for everyone.'

[14]You might be wondering, *how much* do interest-rate hikes exacerbate inflation? It's a difficult question to answer. It's also not very interesting. The main reason to care about inflation is that it affects the distribution of income. And that's why we should care about rate hikes. They're a way to bolster the income of creditors, with obvious distributional consequences.

This verbiage is a new variation on an old trick. The tithes may come in the name of God, but they always seem to benefit the church.

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.

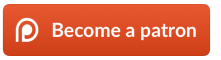

Sources and methods

All monetary data is from the World Bank as follows:

- inflation rates: series FP.CPI.TOTL.ZG — inflation consumer prices (annual %)
- interest rates: series FR.INR.LEND — lending interest rate (%)

Some comments on methods. When putting the cross-country data into bins, I keep only the bins containing 10 or more observations. (The exception is Figure 2, which keeps all the binned data, because there is far less of it.)

For time-series analysis, I keep only the countries with 10 or more years of interest-rate and inflation data.

Further reading

Di Muzio, T. (2022). Do interest rate hikes worsen inflation? *Strange Matters*. https://strangematters.coop/interest-rate-hikes-worsen-inflation-volcker-shock/

Fisher, I. (1907). *The rate of interest: Its nature, determination and relation to economic phenomena*. Macmillan.

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