

# Economics from the Top Down

new ideas in economics and the social sciences

---

## Is Stagflation the Norm?

Blair Fix

January 17, 2023



[I]nflation in the midst of stagnation is not an anomaly.  
If anything, it is the general rule.

— [Nitzan and Bichler, 2009](#)

As much of the world grapples with post-Covid price gouging, it seems like a good time to revisit our understanding of inflation. In this post, I'm going to test Jonathan Nitzan and Shimshon Bichler's 'stagflation thesis'.

The idea is that 'stagflation' — economic stagnation combined with high inflation — is not some exogenous 'market shock'. According to Nitzan and Bichler, stagflation is a business strategy — one of two main routes to profit.

The first route to profit is for businesses to hold prices steady while they try to sell more stuff. The second route is to jack up prices. Since this latter option requires restricting the flow of resources (stuff that flows freely cannot be dear), Nitzan and Bichler reason that when inflation rears its head, it ought to come with economic stagnation. In other words, stagflation is the norm.

If this stagflation thesis is correct, then inflation ought to correlate negatively with economic growth. Looking at the United States, Nitzan and Bichler [find evidence](#) that it does. Here, I broaden their stagflation research by looking at all countries in the World Bank's [global development database](#).

I find that both within and across countries, economic growth (measured in terms of energy use) tends to decline as inflation increases. So Nitzan and Bichler appear to be onto something. Over the last half century, stagflation is the general rule.

## Follow the prices

To make sense of the stagflation thesis, we need to start with how Nitzan and Bichler understand inflation. Ironically, the best way to do that is to forget about the word ‘inflation’ and talk instead about prices.

The reason is that to many people, the word ‘inflation’ implies a decrease in the purchasing power of money. Although not wrong, the problem with this interpretation is that it is needlessly indirect. Framing inflation in terms of decreasing purchasing power is like discussing your child’s growth in terms of the ‘shrinking height capacity’ of your doors. Sure, it’s true in a sense. But it is also tediously circuitous. The fact remains that it is your child (not your doors) who changes.

The same is true of inflation. When inflation rears its head, money appears to lose its value. But the reality is that it is *prices* (not the nature of money) that change. So if we want to understand the phenomenon of ‘inflation’, we should study prices directly.

## Setting prices

When we focus on prices, it becomes clear that we are studying the behavior of business. Why? Because it is businesses who set prices.

Yes, I said *set*.

Contrary to the fantasy world of neoclassical economics — in which businesses ‘take’ prices from the market — in the real world, prices are always ‘set’. It’s a basic fact of property rights. If I walk into a grocery store, I cannot leave with food unless I agree to the price decreed by the store’s owner. That’s the way property rights work. Property gets transferred only if the owner agrees to the conditions.

Of course, I’m not claiming that businesses can set prices without fetter. No matter how powerful the company, the decision to raise prices will be met with friction. Still, the point is that raising prices is a business decision — a key part of every owner’s revenue strategy.

## Breadth vs. depth

Looking at business strategy, Nitzan and Bichler see two general paths to greater profit. To increase its income, a business can:

1. sell more stuff
2. raise prices

In principle, a business could pursue both strategies at the same time — raising prices while also selling more stuff. In practice, however, Nitzan and Bichler argue that the two strategies tend to be mutually exclusive.

Here's why.

If I want to sell more stuff — a strategy that Nitzan and Bichler call *breadth* — the easiest option is to lower the price of my property (i.e. put my stuff on sale). Conversely, if I want to raise my prices — a strategy that Nitzan and Bichler call *depth* — I'll need to restrict access to the stuff I'm selling. Why? Because stuff that flows freely cannot be expensive.<sup>1</sup>

So far, there should be nothing controversial about this thinking. That said, mainstream economist believe that the strategy of 'depth' should be nearly impossible. Here's their reasoning.

First, mainstream economists look at capitalist society and see 'perfect competition'. Next, they look at their theory of free markets and observe that it precludes coordinated price gouging. (In a competitive market, firms cannot raise prices without being undercut by rivals.) Putting the two assumptions together, economists conclude that coordinated price increases cannot exist unless they are driven by 'exogenous' factors (like the government printing money) or market 'distortions' (like oligopoly).

Of course, in the real world, oligopoly is the norm. So any theory that assumes cut-throat competition is dubious at best. But even without oligopoly, there's good reason to suspect that businesses might coordinate to raise prices.

## Coordination through conformity

As a social species, humans have an intense desire to conform to social norms. Sometimes this desire leads to stability — as with religious traditions that last for centuries. Other times, though, conformity leads to social change.

---

<sup>1</sup>Note that for exposition purposes, I'm simplifying Nitzan and Bichler's definition of breadth and depth. For a more nuanced discussion, see [Chapter 15](#) and [Chapter 16](#) of *Capital as Power*.

Fashion is a good example of both tendencies. As individuals, we like to dress the same as other people, leading to a (relative) uniformity in our attire. And yet over time, fashion changes — a herd behavior in which people conform to the *new* way of dressing. And so we get coordination (changing fashion) through conformity.

Something similar happens with businesses. Often, businesses compete by cutting costs and increasing the amount of stuff they sell. But this is not the only mode of competition. Sometimes, a business raises its prices and its competitor responds by *doing the same*. If enough businesses join in, suddenly we have a herd behavior in which every business is attempting to raise prices. Coordination through conformity. Inflation!

Looking at this herd behavior, Nitzan and Bichler argue that it tends to come in waves. When businesses compete by selling more stuff (breadth), they tend to hold prices steady. Conversely, when businesses compete by raising prices (depth), they tend to restrict how much they sell. The result, Nitzan and Bichler propose, is that breadth and depth tend to be mutually exclusive.

Translated into standard economic jargon, the idea is that when inflation rears its head, it should be accompanied by economic stagnation — a combination that economists call ‘stagflation’.

## Measuring stagflation

At first glance, testing the stagflation thesis seems straightforward. We simply measure inflation and see if it correlates negatively with economic growth. The problem, however, is that when we split open the sausage of how inflation and economic growth are measured, we find some big problems.

For starters, inflation is not one thing. When economists report a rate of inflation — say 5% — it implies that all prices have increased uniformly at this rate. But that’s not how inflation works. Inflation is always differential, meaning the prices of different commodities head in all sorts of directions. The consequence is that there is no single ‘rate of inflation’. Instead, there are many rates. True, we can average these rates to get a measure of central tendency. But doing so involves subjective decisions about which commodities to include, and how to weight their change in price. In short, the ‘rate of inflation’ is inherently ambiguous.

(For details about differential price change, see [‘The Truth About Inflation’](#) and [‘Inflation: Everywhere and Always Differential’](#).)

The situation is similar with measures of economic growth. The standard measure of economic growth, ‘real’ GDP, involves a [host of subjective decisions](#) — so many, in fact, that ‘real’ GDP is less of a measurement and more of a divination.

It gets worse.

To test the stagflation thesis, we need to correlate inflation with economic growth. But if we measure economic growth using ‘real’ GDP, we introduce a statistical dependence in our correlation. Why? Because ‘real’ GDP is calculated by taking nominal GDP and deflating it by a measure of average price change. The problem is that this measure of average price change is the same thing as the ‘rate of inflation’. So by construction, we are introducing an autocorrelation between the rate of inflation and the rate of economic growth.<sup>2</sup>

To deal with these issues, here’s what I’ll do. First, I’ll ignore the problems with measuring inflation and simply use the *consumer-price-index data* provided by the World Bank. (Sometimes doing science involves putting up with data that is less-than ideal.)

Second, I’ll use *energy use per capita* (rather ‘real’ GDP) to measure economic growth. Why energy use? Well, because energy is the ‘go of things’ — the resource that makes everything else that humans do possible. Energy is the ‘master resource’.

Of course, energy use correlates tightly with the growth of ‘real’ GDP. That said, measuring energy consumption is far simpler and more physically meaningful than measuring ‘real’ GDP. And unlike ‘real’ GDP, the measurement of energy use is not statistically connected to the rate of inflation. So when we test the stagflation thesis, it’s not contaminated by an autocorrelation.

---

<sup>2</sup>Here’s a simple way to think about the autocorrelation between economic growth and the rate of inflation. Let  $\dot{P}$  be the average rate of price change. And let  $\dot{Y}$  be the growth rate of nominal GDP. In simplified terms, the growth rate of real GDP is  $\dot{Y} - \dot{P}$ . And the rate of inflation is  $\dot{P}$ . So when we correlate inflation with economic growth, we correlate price change with the (negative version) of itself.

$$\dot{P} \sim \dot{Y} - \dot{P}$$

Of course, this equation doesn’t mean that we’ll automatically find a strong correlation between  $\dot{P}$  and  $\dot{Y} - \dot{P}$ . (If variation in  $\dot{Y}$  dwarfs variation in  $\dot{P}$ , then there will be little correlation.) Still, the circularity in our formula gives reason to worry.

On a sidenote, it turns out that some flashy results in the social sciences are due purely to autocorrelation. The [Dunning-Kruger effect](#) is the most famous example.

(If you're interested, I wrote a [whole book](#) in which I tested economic growth theory using energy consumption rather than 'real' GDP.)

## A tour of international inflation

Before I test the stagflation thesis, I think it's worth taking a tour of the World Bank inflation data.

As individuals, we're often focused on the inflation rate in our own corner of the world. But as scientists, it's best to avoid this kind of myopia and look at the big picture. On that front, the World Bank reports [inflation rates](#) (measured using the consumer price index) for a wide variety of countries over the last 60 years. Let's take a tour of the data.


The general pattern is that inflation is always differential. Just as the rate of price change [varies significantly between commodities](#) (within the US), the rate of inflation (the average price change) varies wildly between countries.

Figure 1 shows the scale of this variation in 2021 — the year that post-Covid inflation reared its head. Here, each blue point shows the inflation rate in the corresponding country. To put the numbers in context, note that I've plotted the inflation rate on a logarithmic scale. So each tick mark on the horizontal axis indicates a *tenfold* increase in the rate of inflation.

Behold the enormous variation in price change! In 2021, countries like Fiji and Ecuador saw virtually no inflation, while countries like Sudan and Lebanon experienced hyperinflation. And some countries, like Chad and Japan, actually saw prices *fall* (inset box).


The take-home message from Figure 1 is that we should be wary of any simple explanation of inflation. For example, one of the narratives about post-Covid price gouging is that it's been driven by supply-chain shortages. True, global supply chains are a mess. But the point is that you can't look at the scale of inflation variation (as shown in Figure 1) and claim it all boils down to 'supply-chain problems'.

Take the example of Japan and Haiti — two island nations that depend heavily on imports. In 2021, Japan experienced mild deflation while Haiti experienced double-digit inflation. Why? I don't know the answer, but it seems tenuous to chalk it up solely to 'supply-chain problems'.



**Figure 1: Inflation in 2021 by country**

Each blue point shows the 2021 inflation rate within the corresponding country, measured using the consumer price index. Note that these rates are plotted on the horizontal axis using a logarithmic scale. The inset panel shows data for countries that experienced deflation. (These negative values cannot be plotted on the main log scale.) [Sources and methods](#)



**Figure 2: The median inflation rate across all countries**

Using data from the World Bank, this figure shows the median inflation rate across all countries that have inflation data. Note, however, that the country sample is not consistent over time. See Figures 7 and 8 for details. [Sources and methods](#)

Having looked at inflation variation, it's important to recognize that there is a pattern to the madness. Countries tend to experience bouts of inflation together. Figure 2 testifies to this fact. Here, I've plotted the median inflation rate over time for all countries in the World Bank database. Across the globe, inflation peaked in the 1970s and declined thereafter.

(The caveat here is that the composition of the inflation data changes over time. Overall, the World Bank data is biased towards poor countries, although this bias lessens with time. For details, see Figures 7 and 8 in the [Sources and methods](#).)



Looking at the global inflation trend in Figure 2, it's important that we not fool ourselves. Yes, we can take international inflation rates and measure their median. And yes, this median has a pattern. But no, the median doesn't tell us much about what's happening in any specific country. Figure 3 illustrates why.

Here I've replotted (from Figure 2) the median inflation rate across countries. But now I've placed this median within the context of inflation variation — the range capturing 90% of the inflation data in the World Bank dataset. Looking at Figure 3, we can see that the inflation range dwarfs the trends in the central tendency. (Note that the vertical axis uses a compressed scale to show the range of inflation variation.) It's another reminder that inflation is always differential.

In Figure 3, you can see that there was a conspicuous spike in inflation variation in the 1990s. What happened then?


If you lived in the West during the 90s, you'll remember it as a period of prosperity. Elsewhere, however, it was a period of crisis. Most conspicuously, the Soviet Union collapsed in 1991, sending many of the newly sovereign states into bouts of severe hyperinflation. (For example, in 1993, the Ukrainian inflation rate was close to 5,000%.) During the same period, hyperinflation also reared its head in many parts of Latin America. (In 1990, Peru's inflation rate was about 7,500%.)

To summarize, yes, international inflation has a central tendency. But no, this tendency doesn't tell us much about what's going on in any particular country.

## **The rich avoid inflation**

One more detour before I test the stagflation thesis. The World Bank data clearly demonstrates that inflation is related to the global distribution of income. To predict whether a country will have high (or low) inflation, we need only rank its per capita income. If the country is poor, inflation will be high. But if the country is rich, inflation will be low.

Figure 4 illustrates the pattern. Here, I've taken countries and ranked them (annually) by their nominal GDP per capita. Then I've put countries into income bins (plotted on the horizontal axis) and calculated the range of




**Figure 3: Inflation variation across all countries in the World Bank database**

Here, the blue line shows the median inflation rate across countries — the same data as in Figure 2. The shaded region shows the scale of inflation variation between countries (the range covering 90% of the data). Note that the vertical axis is compressed to show the scale of variation. (It's a log scale, where percentage growth rates are converted to decimal growth ratios, as in 5%  $\rightarrow$  1.05  $\rightarrow$   $\log(1.05)$ ). [Sources and methods](#)

inflation rates within each bin (plotted on the vertical axis). The resulting trend is easy to spot. As countries get richer, their inflation rates tend to decrease.

Since this data foray is mostly a detour from my main goal of testing the stagflation thesis, I won't dive too deeply into the income-inflation pattern shown in Figure 4. But I will say that it begs for a systems analysis.



**Figure 4: Richer countries tend to have lower rates of inflation**

This figure analyzes national inflation rates as a function of nominal GDP per capita. Here's how it works. In each year, I rank countries in the World Bank database by their (nominal) GDP per capita. Based on this rank, I then sort countries into bins by GDP-per-capita percentile. In the chart, each point indicates a GDP bin, with the percentile plotted on the horizontal axis. Within each GDP bin, I then measure the range of inflation experienced by the corresponding countries. The blue line indicates the median inflation rate within each GDP bin. The shaded region shows the range for the middle 50% of inflation data. [Sources and methods](#)

Yes, poor countries tend to have less stable political-economic regimes, which may lead to price instability. But poor countries are also peons in the global halls of power. In other words, it seems plausible that the policies that benefit rich countries may also destabilize prices in poor countries. (In particular, I'm thinking of the International Monetary Fund's penchant for imposing 'structural adjustment' — aka brutal austerity — on indebted countries.) At any rate, the relation between international income rank and inflation is worth analyzing in greater detail.

## The stagflation evidence

Now to the main task — testing the stagflation thesis. Using World Bank data, I'm going to measure the correlation between the annual inflation rate and the annual growth rate of per capita energy use. If the stagflation thesis is correct, this correlation should be negative.

### The energy-inflation pattern within countries

I'll begin by looking at the energy-inflation trend within countries. Figure 5 summarizes the evidence.

The analysis works as follows. First, I search for all countries in the World Bank database with 20 or more years of energy-inflation data. (The purpose of this time cutoff is to isolate the countries where we can study a long(ish)-term trend.) Within each country, I then measure the correlation between the annual inflation rate and the energy-per-capita growth rate. (Because both series are quite volatile, I measure the correlation on the log of the fractional rate of change.) Finally, the histogram in Figure 5 plots the distribution of within-country correlations.


Looking at Figure 5, we can see that the evidence is not slap-yourself-in-the-face obvious. In some countries, inflation correlates negatively with the growth of energy use. But in other countries, the reverse is true. In other words, stagflation is by no means universal.

Still, the evidence leans in favor of the stagflation thesis. Among the 114 countries studied here, energy growth correlates negatively with inflation among 74 of them — about two thirds of the sample. And overall, the average energy-inflation correlation is negative, at roughly -0.1.

### The energy-inflation pattern across countries

Let's turn now to the energy-inflation trend across countries. Figure 6 shows the pattern.

Let me break down the chart, starting with some general principles of data visualization. When we visualize a correlation, the simplest option is our old friend the [scatter plot](#). To make a scatter plot with the current data, we'd plot international inflation rates against international energy growth rates. Looking at Figure 6, you can see that these are indeed the axes that I've used. But there's no scatter plot. Why?




### Figure 5: Inflation vs. energy growth — the correlation within countries

This figure analyzes how inflation correlates with the growth rate of energy use per capita within countries. Looking only at countries with 20 or more years of data, I measure the log-log correlation (within each country) between annual inflation rates and energy-per-capita growth. The histogram shows the resulting distribution of correlation. [Sources and methods](#)

The answer is that scatter plots work well when the correlation leaps off the page. However, when the data is ‘messy’ (i.e. has lots of outliers), scatter plots can be hard to interpret. In the current case, the inflation data is wildly noisy, while the inflation trend with energy growth is quite subtle. To see the energy-inflation trend, I’ve binned the raw data to show the average pattern.

Here’s what that means.



**Figure 6: Inflation vs. energy growth — the trend across countries**

This figure analyzes the international relation between annual rates of inflation and the growth rate of energy use per capita. On the horizontal axis, I take energy use growth rates (within each country) and then place the data into bins. (Each blue point represents the bin midpoint.) Within each energy-growth-rate bin, I then measure the spread of inflation rates. The line shows the median inflation rate across countries. The shaded region plots the middle 50% of the data. [Sources and methods](#)

I start by lumping (national) energy growth rates into different bins based on the size of the observation. In Figure 6, each blue point represents the midpoint of an energy-growth-rate bin. Within each bin, I then measure the spread of international inflation rates. The blue line shows the median inflation rate among countries. The shaded region shows the middle 50% of the data.

Despite being subtle, the pattern in Figure 6 is easy to see. As energy growth rates increase, inflation tends to decline. Or put another way, energy stagnation is associated with greater inflation. Stagflation!

## The stagflation future

According to the World Bank data, higher inflation tends to come with stagnating energy consumption. So Nitzan and Bichler's stagflation thesis appears to be supported by the evidence, at least over the last half century.

But so what?

Well, the 'so what' depends on our vision for the future. If you think that perpetual economic growth is good, then the stagflation evidence gives a clear policy goal. To limit inflation, you should strive for economic growth. So let the fossil fuels flow!

Or not.

If, like me, you think that perpetual economic growth is both impossible and insane, then the 'so what' is more complicated. Over the long term, we *want* energy consumption to decrease. But we also want this 'degrowth' to be a non-traumatic experience.

And here we have a problem. At present, the lack of economic growth represents a crisis. It comes, among other things, with high inflation — something that nobody likes. And so degrowth skeptics have low-hanging fruit to pick. "Look at what happens when economic growth stops," they say. "It's a disaster!"

The critics are not wrong. But they also lack imagination.

If perpetual economic growth is both impossible and insane, then we have no choice but to build a social system in which the lack of growth is not a crisis. ([Slower by design, not disaster](#) is how ecological economist Peter Victor puts it.) In this world, stagflation cannot be the norm.

How do we get there? I have no idea. But the point is that we should start thinking about it . . . now.

## 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.

 Become a patron

## Sources and methods

All data comes from the World Bank, as follows:

- Inflation data: series [FP.CPI.TOTL.ZG](#)
- Energy use per capita: [EG.USE.PCAP.KG.OE](#).
- Nominal GDP per capita (current US\$): [NY.GDP.PCAP.CD](#)

### Sample bias in the World Bank data


It's important to recognize that the World Bank dataset is not an exhaustive sample of international historical data. When it comes to inflation, the World Bank sample size varies with time. Figure 7 shows the pattern. In 1960, the World Bank reports inflation data for just over 60 countries. That number increases with time until the late 2000s, after which it decreases.

The composition of this country sample is not random. Instead, it is biased towards low-income countries. Figure 8 runs the number.

Let's break down what I've done. In each year, I take World Bank data for (nominal) GDP per capita and calculate each country's corresponding GDP-per-capita percentile. (Because economists are obsessed with GDP, this data happens to be quite exhaustive, meaning few if any countries have missing data.) Next, I remove countries that lack inflation data. Finally, I calculate the average GDP-per-capita percentile of the remaining sample.

If this sample was non-biased, the GDP-per-capita-percentile would average 50%. Any deviation from this value indicates that the inflation sample is biased. Figure 8 shows that the World Bank inflation data is biased towards






**Figure 7: Number of countries with inflation data in the World Bank database**

The World Bank does not have exhaustive international inflation data. Instead, the country sample increases with time.

---

low-income countries. We know this because our country sample has an average GDP-per-capita-percentile that is well below 50%. Note, though, that this bias lessens with time.



**Figure 8: Measuring income bias in the World Bank inflation data**

This figure illustrates how the World Bank inflation data is biased towards low-income countries. Here's how the analysis works. First, I take each country and calculate their global GDP-per-capita percentile (a value from 0 to 100.) in each year. Next, I remove countries that lack inflation data. Finally, I calculate the average GDP-per-capita percentile of the remaining sample. If the inflation data had no income bias, the results would hover around the 50% mark (dashed red line). They do not. Instead, the World Bank inflation data is biased towards low-income countries (those below the median GDP per capita).

## Further reading

Fix, B. (2015). *Rethinking economic growth theory from a biophysical perspective*. New York: Springer.

Fix, B., Nitzan, J., & Bichler, S. (2019). Real GDP: The flawed metric at the heart of macroeconomics. *Real-World Economics Review*, (88), 51–59.

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

Victor, P. (2008). *Managing without growth: Slower by design, not disaster*. Northampton, Mass.: Edward Elgar.