

# Economics from the Top Down

new ideas in economics and the social sciences

---

## The Truth About Inflation

Blair Fix

November 2021



*Inflation is always and everywhere  
a phenomenon of structural change.*

— [Jonathan Nitzan, 1992](#)

Milton Friedman has been dead for more than a decade, but his ghost still haunts us. In the 1960s, Friedman [declared](#) that inflation is ‘always and everywhere a monetary phenomenon’ — a problem of printing too much money. Since then, whenever inflation rears its head, you can count on someone to reanimate Friedman’s ghost and blame the government for spending too much.

If only inflation were so simple.

Like much of economic theory, Friedman’s thinking appears plausible on first glance. Inflation is a general rise in prices. And since prices are nothing but the exchange of money, more circulating money means prices must increase. Hence, inflation is ‘always and everywhere a monetary phenomenon’.

Unfortunately, this thinking falls apart on further inspection. The problem is that it treats inflation as a *uniform* rise in prices. That’s theoretically convenient, but empirically false. In the real world, inflation is wildly divergent. At the same time that the price of apples rises by 5%, the price of cars could grow by 50%, and the price of clothing might fall by 20%.

To understand inflation as it actually exists, we must look not to economics textbooks, but to real-world data. That’s what political economist [Jonathan Nitzan](#) did during his PhD research in the early 1990s. His work culminated in a dissertation called [Inflation As Restructuring](#). In the real world, Nitzan

observed, price change is always ‘differential’, meaning there are winners and losers. The consequence is that inflation is not purely a ‘monetary phenomenon’, as Milton Friedman claimed. Inflation *restructures* the social order.

It is this real-world feature of inflation that is most important, because it means that inflation signals a change in society’s power structure. Predictably, it is this real-world feature that mainstream economists ignore — largely because it conflicts with their tidy theory of inflation as a ‘monetary phenomenon’. Fortunately, the evidence is clear. Inflation is (and has always been) overwhelmingly differential. Inflation is restructuring.

Today, as inflation fears [return](#) and Friedman’s ghost is [resurrected](#), it’s worth reminding ourselves of the real-world facts.

## The quantity theory of money

Let’s start with Milton Friedman’s ‘[quantity theory of money](#)’, which proposes that inflation is always caused by printing too much money.<sup>1</sup> Like so much of neoclassical economics, the theory is a mixture of two things:

1. Dubious assumptions about human behavior;
2. An accounting identity that makes the theory look good.

The potency of this mixture was solidified by Friedman’s famous ‘[F-twist](#)’, in which he argued that a theory’s assumptions are irrelevant. All that matters, Friedman claimed, is that the theory makes accurate predictions.

Friedman’s F-twist gets dubious assumptions off the hook. But there is still the problem of predictions. How do you ensure that your theory is consistent with the evidence? Here, neoclassical economists have hit upon a tidy trick: frame your theory in terms of an *accounting identity*. Since the identity is true by definition, any ‘test’ of the theory will come out in your favor.

### An old trick

Before we get to Friedman’s theory of inflation, let’s look at some other implementations of this accounting-identity trick. When neoclassical economists test their theory of income (the theory of marginal productivity), they [invoke](#)

---

<sup>1</sup>The ‘quantity theory of money’ is not so much Friedman’s theory as it was re-popularized by Friedman. The idea that prices relate to the quantity of money dates (at least) to a [1517 memorandum](#) written by the astronomer Nicolaus Copernicus.

an accounting identity. They correlate two related forms of income (usually sales and wages) and then call one of the incomes ‘productivity’. Since they always find a correlation, they always ‘confirm’ their theory of income. Nifty!

Then there’s the neoclassical theory of economic growth. The theory assumes that economic output is governed by a ‘production function’ that dictates how inputs of capital, labor and ‘technological progress’ are transformed into economic output. And guess what . . . this approach seems to have overwhelming empirical support. The problem, pointed out by Anwar Shaikh, is that the production function is actually a re-arrangement of a national-accounting identity. The production function ‘works’ because it is true by definition. Nice!

### Milton’s money

Back to Milton Friedman’s theory of inflation. Like a good neoclassical economist, Friedman grounds his theory in an accounting identity — one that relates the quantity of money  $M$  to the average price level  $P$ :

$$MV = PT$$

In this identity,  $V$  is the ‘velocity of money’ — the rate that money changes hands. And  $T$  is an index of the ‘real value’ of all transactions.

The nice thing about this accounting identity is that it is true by definition. So if you tie a theory of inflation to it, your ‘predictions’ will always work. The problem, pointed out by critics, is that this identity tells us nothing about causation. It could be that printing too much money causes prices to rise. Or it could be that rising prices drive people to borrow (and hence ‘create’) more money.

### When identities mislead

A more subtle problem with our accounting identity is that it may mislead more than it guides. The identity tells us about the *average* price level,  $P$ . Economists then assume that this average is a useful measure of price behavior. But that may not be true.

Here's how things can go wrong. Consider a society that sells two commodities, apples and oranges. If the two commodities are sold in equal proportions, the price index  $P$  is the arithmetic average of the price of the two fruits. If each fruit costs \$1, then the price index is:

$$P_1 = \frac{\$1 \text{ (per apple)} + \$1 \text{ (per orange)}}{2}$$

$$= \$1 \text{ (per fruit)}$$

Suppose that after a few months, the price of apples *rises* by 50% while the price of oranges *falls* by 50%. This change represents a drastic restructuring of the price system — a mixture of hyper-inflation and hyper-deflation. Yet this instability doesn't register in our price index. The average price level remains unchanged:

$$P_2 = \frac{\$1.50 \text{ (per apple)} + \$0.50 \text{ (per orange)}}{2}$$

$$= \$1 \text{ (per fruit)}$$

So just because the *average* price is unchanged doesn't mean that *individual* prices are stable. If price change is divergent, then the idea of an 'average price level' is uninformative, if not outright misleading.

## The trouble with averages

In the hands of economists, the idea of an 'average price level' is, to echo Joan Robinson, "a powerful instrument of miseducation".<sup>2</sup>

---

<sup>2</sup>In this case, Joan Robinson was referring not to price indexes, but to another instrument of miseducation — the production function:

... the production function has been a powerful instrument of miseducation. The student of economic theory is taught to write  $O = f(L, C)$  where  $L$  is a quantity of labour,  $C$  a quantity of capital and  $O$  a rate of output of commodities. He is instructed to assume all workers alike, and to measure  $L$  in man-hours of labour; he is told something about the index-number problem involving choosing a unit of output; and then he is hurried on to the next question, in the hope that he will forget to ask in what units  $C$  is measured. Before ever he does ask, he has become a professor, and so sloppy habits of thought are handed on from one generation to the next.

(Joan Robinson, 1953)

The trouble is that averages are a mathematical identity — they are true by definition. I can calculate the average of any conceivable set of numbers. But that doesn't mean my calculation will be informative. That's because averages *define* a central tendency, yet do not indicate if this tendency actually *exists*.

Here's an example. Suppose two people have an average net worth of \$100 billion. Is this a central tendency? Perhaps ... if our two people are Warren Buffet (net worth \$104 billion) and Mukesh Ambani (net worth \$96 billion). Both have close to the same wealth. But what if the two people are Jeff Bezos (net worth \$200 billion) and me (net worth \$0 billion)? In this case, the average misleads more than it informs.

Of course, scientists are aware of this problem. That's why they are trained to report averages together with a measure of *variation*. Doing so gives a sense for the meaningfulness of the average.

Any measure of variation will do, but the most popular is the '[standard deviation](#)', which measures the average deviation away from the mean.<sup>3</sup> Returning to my wealth example, reporting the standard deviation of wealth tells us when the average measures a real central tendency, and when it does not.

For instance, Warren Buffet and Mukesh Ambani have an average net worth of \$100 billion, with a net-worth standard deviation of \$5.7 billion. The fact that the variation is small (about 0.06 times the average) indicates that there is a real central tendency. In contrast, Jeff Bezos and I have an average net worth of \$100 billion, with a standard deviation of \$141 billion. This enormous variation (1.4 times the average) indicates that there is no central tendency in the raw data. So the average is uninformative.

## Reported averages, unreported variation

The idea that averages should be reported together with a measure of variation is a basic part of empirical science. And yet when economists study inflation, this practice is conspicuously absent. Why?

Perhaps economists have a good reason for not reporting inflation variation. To consider this possibility, let's dig into how economists measure inflation. It starts with something called a 'price basket'. This is just a bunch of com-

---

<sup>3</sup>If you want to be technical, the standard deviation measures the [root mean square](#) of the deviation from the mean.

modities whose prices are tracked over time. The [consumer price index](#), for instance, tracks a basket of commodities typically purchased by consumers. The [wholesale price index](#), in contrast, tracks a basket of wholesale commodities. There are many different types of baskets, but here I'll focus on the consumer price index (CPI).

Having chosen a basket of commodities, economists then track the average price of the basket. If you read the technical literature, you'll see that economists have [various formula](#) for giving some commodities more weight than others when they average prices. The math looks fancy, but doesn't change the fact that the output is a type of 'average'. Watch out, though, because economists won't call it an 'average'. They'll call it a '[price index](#)'.

With their price index in hand, economist then judge the rate of inflation by seeing how fast the index rises over time. As an example, Figure 1 plots the percentage change in the US consumer price index since January 1, 2020. Over this period, the index rose by 7%. Inflation!

Having looked at the price-index trend, pundits will then discuss the cause and consequences of inflation. Here, for instance, is [commentary](#) from the New York Times:


Consumer prices surged at the fastest pace in more than three decades in October as fuel costs picked up, supply chains remained under pressure and rents moved higher — worrying news for economic policymakers at the Federal Reserve and for the Biden White House.

Given the trend in Figure 1, such commentary seems justified. And yet it is missing a crucial piece of information. I have shown you the movement of the *average* price. But I've told you nothing about price-change *variation*.

The reader is left to draw their own conclusions about how inflation varies by commodity. With this lack of information, most people will assume that the movement of the average price indicates a strong central tendency. In other words, they'll assume that inflation is *uniform*. And they might be right.

## Inflation, like gravity

Perhaps economists have a good reason for not reporting price-change variation. Maybe this variation is so small that it's not worth recording. In this case, inflation is like Earth's gravity: remarkably uniform.



**Figure 1: A typical inflation report**

I have plotted here the movement of the US consumer price index since January 1, 2020. [Sources and methods](#)

In high-school physics, I did hundreds of projectile-motion calculations. In each problem, we’d assume that the acceleration of gravity was a uniform  $9.81\text{ m/s}^2$ . I vaguely recall my teacher mentioning that this acceleration varied slightly around the world, depending on geography. But we never included this variation in our calculations. Nor did any of my textbooks report it. Were physicists hiding something from me?

To consider this possibility, I headed over to the [International Gravimetric Bureau](#) (I love the name) and downloaded [data](#) on the variation in Earth’s gravity. As with anything in science, there are an assortment of ways to estimate Earth’s ‘gravity anomalies’ (the differences in the acceleration of gravity

across the Earth's surface). But for the purposes of high-school physics, the different methods don't matter. Anyway you slice it, the geographic variation in Earth's gravity is incredibly small.

In the data that I downloaded, the standard deviation of Earth's gravity anomaly is about *40,000 times* smaller than the textbook (average) acceleration of  $9.81 \text{ m/s}^2$ . So unless you're doing an incredibly sensitive experiment, gravity variation is so tiny that you can ignore it.<sup>4</sup>

Perhaps economists don't report price-change variation because, like gravity variation, it is insignificantly small. To consider this possibility, let's return to the movement of the US consumer price index — a measure of the average price level. Suppose that alongside this price index, we plotted the price movement of *every* commodity in the CPI basket. If inflation was overwhelmingly uniform, the results might look like Figure 2.

In Figure 2, the colored lines show the indexed price of individual commodities. I should stress that this is *simulated* data. The actual data (which we'll see shortly) looks quite different. The point of this simulation is to show you what it would look like if inflation was overwhelmingly uniform.

Over the period shown, the official consumer price index increased by 7.2%. In my simulation, inflation is so uniform that across all commodities, the standard deviation of price change was 0.15%. At just 2% of the average, this variation is so small that we can ignore it.

So perhaps economists know that inflation varies slightly between commodities, but they also know that this variation is so small that it's not worth reporting.


## Inflation in the real world

Having shown you what price change would look like if inflation was overwhelmingly uniform, let's now look at inflation in the real world. It's rather different.

---

<sup>4</sup>The gravimetric dataset I downloaded is called the [Complete spherical Isostatic gravity anomaly \(Airy-Heiskanen, Tc=30 km\)](#). This is a model-based estimate of Earth's gravity anomaly, resolved to a grid size of 30km. The data is reported in units of mGal. A [Gal](#) is defined as an acceleration of  $1 \text{ cm/s}^2$ . A mGal is 1/1000 of a Gal. In this particular dataset, I calculated the standard deviation in Earth's gravity to be 25 mGal, or  $0.00025 \text{ m/s}^2$ . That's 39208 times smaller than the canonical (average) acceleration of  $9.80665 \text{ m/s}^2$ .






**Figure 2: What price change would look like if inflation was nearly uniform across commodities**

The black line shows the change in the actual US consumer price index. The colored lines show what the price-change of individual commodities would look like if inflation was nearly uniform. [Sources and methods](#)

Figure 3 shows the price change of every commodity tracked by the consumer price index. Instead of clustering tightly around the average price level (the ‘official CPI’), real-world commodities have a mind of their own. Their prices head in all sorts of directions — often in ways that seem unrelated to the movement of the average price.

Notice how plotting the price-change of all CPI commodities alters the inflation story. Looking at Figure 3, no one would conclude that all prices are inflating uniformly. Yet when we looked at the consumer price index alone, this conclusion seemed plausible.



**Figure 3: Price change in the real world**

The black line shows the change in the US consumer price index since January 1, 2020. The colored lines show the indexed price of all the individual commodities tracked by the CPI. Many commodities are tracked in multiple locations. [Sources and methods](#)

When we study the whole range of price change, we see that inflation is a messy business. The numbers tell us as much. Since January 1, 2020, the consumer price index rose by 7.3%. That value seems significant ... until we measure price-change variation. Over the same period, the standard deviation of price change was 10.7%. So the variation in price change was about *1.5 times* larger than the price-change average.<sup>5</sup>

<sup>5</sup>When we divide the standard deviation by the mean, we get something called the ‘[coefficient of variation](#)’, or alternatively, the ‘relative standard deviation’. This measure allows us to quantify the relative size of variation within a data set. Like other measures of relative dispersion (i.e. the Gini index), the coefficient of variation has no units, and hence, can be used to compare dispersion across all sorts of data.

To put this variation in perspective, let's return to our wealth example. Jeff Bezos and I have an average wealth of \$100 billion. But this value does not indicate a real central tendency. Jeff Bezos is worth \$200 billion. I'm worth \$0 billion. We can tell that the average is misleading by measuring the standard deviation of our wealth, which happens to be \$141 billion. So the variation in our wealth is about 1.4 times the average.

This ratio of 1.4, you'll notice, is actually *less* than the ratio of 1.5 we found for price-change variation. So if we conclude that it's rather meaningless to average my wealth with Jeff Bezos's wealth, we should also conclude that the movement of the consumer price index is quite meaningless. Both averages mislead more than they inform.


## A story of inflation variation

The real story of inflation — the one that goes largely unreported — is of wildly divergent price change among different groups of commodities. Figure 4 shows how this inflation has played out across 12 major commodity groups tracked by the US consumer price index.

Let's break down the analysis. First, I've tracked the price change of every commodity in the US consumer-price-index basket (in every geographic location) between January 2020 and October 2021. Then I've put these commodities into major groups (as classified by the Bureau of Labor Statistics). Finally, I use box plots to show the variation in price-change within each commodity group.

The story told by this disaggregated analysis is dramatically different than the one told by the movement of the average price. We can see that inflation varies greatly between different commodity groups. Some groups, like 'men's apparel', have experienced little (if any) inflation. Other groups, like 'private transportation', have seen massive price hikes. Figure 4 also shows that inflation varies greatly *within* each commodity group. Inflation often coexists with deflation — a fact that's evident when the boxes cross the dashed red line.


So the real inflation story, which goes largely undiscussed, is that price change is remarkably non-uniform. In fact, it is so non-uniform that reporting the change in the average price borders on meaningless. So why does price-change variation go unreported?



**Figure 4: US price change by commodity group**

Box plots show the range of price change between Jan. 2020 and Oct. 2021, for US CPI commodities classified into 12 major groups. Here’s how to read the box plot. The thick vertical line indicates the median value. The ‘box’ shows the middle 50% of the data. And the line shows the range of the data, excluding outliers. [Sources and methods](#)

Perhaps we can excuse newspapers from not printing charts like Figures 3 and 4. These graphs are admittedly more challenging to interpret than simply reporting the percentage change of a price index. This difficulty, though, doesn’t get economists off the hook. Every trained economist ought to know how to interpret the type of evidence shown above. And yet, even in the technical literature, you’re unlikely to find a disaggregated analysis of inflation. Why?



**Figure 5: The history of US inflation**

The blue line shows the annual rate of change of the US consumer price index. Dashed red lines show the threshold for double-digit inflation. [Sources and methods](#)

Is this time different?

Before we chastise economists for not reporting price-change variation, let’s give them the benefit of the doubt. Perhaps inflation is *normally* uniform, but our current circumstances are unusual. It could be that the pandemic has wrought chaos on an otherwise stable price system. In this case, if we look at the deeper history of inflation, we ought to see a pattern of uniformity in which all prices move together.

To investigate this possibility, let’s look at the long-term history of US inflation. We’ll start with the standard picture, shown in Figure 5. I’ve measured the inflation rate in terms of the annual change in the consumer price index. The red lines show the threshold for double-digit price change — an arbitrary but frequently referenced threshold for ‘major’ inflation.

I could write a book about the inflation ups and downs shown in Figure 5. But I won't, since many such books exist. Instead, I am interested in what is *absent* from this picture — namely, price-change variation.

Perhaps economists ignore price-change variation because, although it is large today, historically it has been negligible. That would make our current situation what economists call a 'distortion' — an event that has forced prices out of 'equilibrium'. In less 'distorted' times, we ought to find price-change uniformity.

To see if this is true, let's look at the long-term history of US price-change variation. Figure 6 shows the data. I start by replotting the average inflation rate (blue line). But then I add some much-needed information — the range of price change across *all* CPI commodities. That's the blue region, which plots the 95% range for the annual price change of all commodities (in all locations) tracked by the CPI. The price-change range is ... rather large.


The evidence in Figure 6 suggests that our current situation is not unusual. Since the CPI data began in 1913, the US inflation rate averaged 2.8%. But over the same period, the standard deviation of annual price change averaged 5.2%. So the inflation variation was historically about *1.8 times* larger than the inflation average. To remind you, the variation between Jeff Bezos's wealth and my wealth was only 1.4 times our average wealth. So looking at the average US inflation rate is even less meaningful than averaging Bezos's wealth with my own.

To summarize, the data is pretty clear: the historical norm has been for price-change variation to trump the average rate of inflation. So why does this variation go unreported?

## **An unreported pattern in the unreported data**

Before we call the science police, let's give economists one last chance to redeem themselves. Although it certainly seems dubious that price-change variation is rarely reported, perhaps we can think of a good reason not to worry about it.

The only reason I can come up with is that price-change variation is large, but otherwise *constant*. If price-change variation is steady over time, then perhaps we can ignore it when we measure the movement of the average price.




**Figure 6: The history of US price-change variation**

The blue line shows the annual change in the consumer price index, replotted from Figure 5. The shaded region shows the 95% range (the range for the middle 95% of the data) in the annual price change of all commodities (in all locations) tracked by the CPI. [Sources and methods](#)

To consider this possibility, let’s imagine a counterfactual US in which price-change variation is constant over time. In this US, annual price change varies by commodity, but this variation is steady year to year. It is fixed at the average range found in the real-world US.

The red lines in Figure 7 show what this counterfactual US would look like. We see that the range of price-change variation does not fluctuate. Instead, it has a constant ‘width’. (The red lines are a constant distance apart.) We also see that this counterfactual US is rather different than the real world.

Our counterfactual thought experiment demonstrates that in the real world, price-change variation (blue shaded region) is *not* constant. Sometimes this variation was far *larger* than the historical average (indicated in Figure 7 when the blue shaded region extends outside the red lines). Such was the case during the stagflation crisis of the 1970s, and during the inflationary



**Figure 7: Is price-change variation constant over time?**

The blue lines and the shaded region reproduce (from Figure 6) the real-world trends in US prices. The red lines show what price-change variation would look like if it was constant over time at the US historical average. [Sources and methods](#)


bouts of both world wars. But sometimes price-change variation was much *smaller* than the historical average, (indicated in Figure 7 when the blue shaded region is inside the red lines). Such was the case during the roaring '20s, and again during the boom years of the 1960s.

Given these observations, it appears that we have a problem. Price-change variation is large, yet goes unreported. Moreover, price-change variation itself *varies* over time. So we have an unreported pattern in the unreported data. That's bad. It means economists have been ignoring an important aspect of inflation.

But it gets worse.

Had economists bothered to measure and report price-change variation, they'd have discovered an interesting correlation. Price-change variation rises and falls with the average rate of inflation.





**Figure 8: As prices change more rapidly on average, inflation becomes more differential**

The horizontal axis shows the annual change in the US consumer price index. The vertical axis shows the standard deviation of annual price change among all commodities tracked by the CPI. The red curve shows the smoothed trend. [Sources and methods](#)

Figure 8 shows the trend. On the horizontal axis, I’ve plotted the annual change in the consumer price index — a measure of inflation. On the vertical axis, I’ve plotted the *variation* in the annual price change of all CPI commodities, as measured by the standard deviation. The red line shows the smoothed trend, which has a U shape. The data suggests that the more rapidly prices change on average, the more *differential* inflation becomes.

What we see, in Figure 8, is a manifestation of Jonathan Nitzan’s discovery: inflation *restructures* the price system. It’s this feature that makes inflation socially traumatic.

## An insignificant thing?

In the 19th century, political economist John Stuart Mill famously declared that money was of little interest:

There cannot, in short, be intrinsically a more insignificant thing, in the economy of society than money.

(John Stuart Mill, 1871)

Although this claim has always struck me as dubious, it's easy to imagine a world in which it were true. Consider a world in which inflation is utterly uniform. All prices (including wages) rise and fall together. In this world, Mill's dictum holds: price change is meaningless. Today the price of an apple is \$1. Tomorrow it doubles to \$2. But since the price of everything else also doubles (including your income), nothing changes. Inflation is an 'insignificant thing'.

Unfortunately (for Mill's dictum), the real world works differently. In it, inflation is never uniform ... it is always differential. And that makes it highly significant. Inflation restructures the social order, producing winners and losers. It is this restructuring, Jonathan Nitzan realized, that is the most important aspect of inflation. And yet it is this feature that economists almost completely ignore.

## But it works in theory ...

For some reason, economists didn't listen to the memo given to all other scientists — the one that said 'thou shalt report an average alongside a measure of variation'. So why the miseducation? Is it an accident? Or is it by design?

Sadly, I think it's the latter. Economists ignore the variation memo because their notion of inflation assumes that price change is uniform. It's much like the old joke:

An economist says to a physicist: "Sure, this equation works in *practice*. But does it work in *theory*?"

Applied to inflation, the joke is:

Sure, inflation is wildly divergent in *practice*. But is it wildly divergent in *theory*?

The answer is overwhelmingly no. In economic theory, inflation is assumed to be uniform. But why would economists assume something so at odds with reality?

Here's what I think is going on. I treat the 'does-it-work-in-theory' joke as a litmus test for *ideology*. It's a test to see if someone elevates ideas above evidence. The more they do so, the less they are doing science and the more they are promulgating ideology. Apply this litmus test to mainstream economics, and you see that it is a secular priesthood masquerading as science.

### The monetarist monastery

As an example of this priesthood, take [monetarism](#) — the school of thought popularized by Milton Friedman. According to monetarists, most social ills are caused by the government printing/spending too much money. Unsurprisingly, monetarists think that these problems have a simple solution: government austerity.

Back to inflation. Faced with rising prices, most monetarists will quickly call for government belt tightening. Their logic works like this:

1. Inflation is linked to the money supply, via the formula  $MV = PT$  (where  $M$  is the quantity of money and  $P$  is the average price level);
2. The government controls the money supply;
3. The government needs to spend less.

Like most good ideologies, this argument contains a devious trick. What monetarists won't tell you is that the money supply gives meaningful insight into inflation only if price change is uniform. If price change varies wildly by commodity (as it does in the real world), then the movement of the average price tells you little (if anything) about the movement of individual prices. And that means the money supply tells you little (if anything) about real-world inflation.

Faced with this problem, the monetarist solution is to make their ideas 'work in theory'. Assume inflation is uniform. Call for austerity. Repeat.

## Solow's snake oil

Although popular in the late-20th century, monetarism was always a controversial sub-sect of neoclassical economics. Many 'moderate' economists thought monetarism was quackery. Hence Robert Solow's famous dig at Milton Friedman. "Everything," Solow [quipped](#), "reminds Milton of the money supply. Well, everything reminds me of sex, but I keep it out of the paper."

Monetarist digs aside, Robert Solow and his fellow macroeconomists maintained their own brand of snake oil. For instance, Solow's famous paper describing his model of economic growth began with this zinger:

There is only one commodity, output as a whole. . . Thus we can speak unambiguously of the community's real income.

([Robert Solow, 1956](#))

Among critics, the 'one-commodity' assumption gets heaps of scorn. And rightfully so. But I think we should give Solow credit for following it up with a clever insight. If 'real income' is to be unambiguous, we must be able to treat society as though it produces only one commodity.

What Solow is hinting at (but not acknowledging) is a severe problem in macroeconomics. The field is built on the principle that you can take the monetary value of production,  $Y$ , and separate it into two components — 'real production'  $Q$  and the nominal price level  $P$ :

$$Y = Q \cdot P$$

If there is only one commodity, then  $Q$  is unambiguous. (Hence Solow's comment.) But the formula also works if there are multiple (unchanging) commodities whose prices move as one. In that case,  $P$  is an index of inflation that unambiguously describes the movement of all prices.

The problem for Solow and his fellow macroeconomists is that neither of these assumptions holds in the real world. Actually-existing societies produce many commodities whose prices do not change uniformly. And that creates a problem. It means that the quantity of production,  $Q$ , is hopelessly ambiguous. (See [this paper](#) for a detailed discussion.)

Like monetarists, macroeconomists solve the problem by making their ideas ‘work in theory’. They simply *define* inflation as a uniform change in prices, and then assume that ‘real income’ is unambiguous. [Procrustes](#) would be proud.

## Why differential inflation matters

Having seen that price change varies greatly between commodities, you might wonder why this matters. Well, it matters because it means that inflation is not purely a ‘monetary’ phenomenon. Inflation *redistributes* income.

Since doing his doctoral work in the 1990s, Jonathan Nitzan has gone on to publish (together with Shimshon Bichler) some eye-opening research on the distributional effects of inflation. Nitzan and Bichler have discovered, for instance, that inflation systematically benefits big business. Figure 9 shows the most recent version of their research.

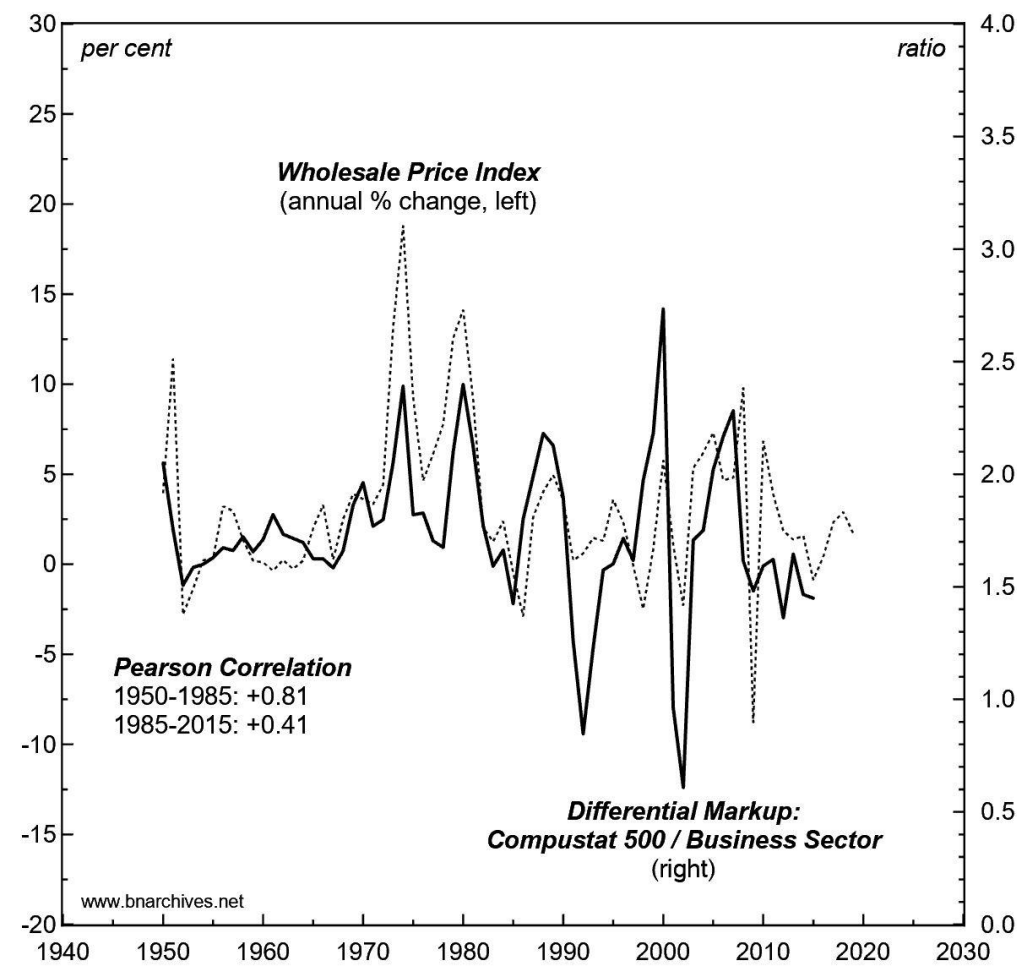
Let’s break it down. We start with the ‘markup’ — the size of a company’s profits as a portion of sales. Nitzan and Bichler then measure the markup for two groups of companies:

1. The ‘Compustat 500’ — the 500 largest US companies (ranked by capitalization) in the Compustat database;
2. All US corporations;

Next, Nitzan and Bichler take the ratio of these two markups. The resulting ‘differential markup’ measures the relative profitability of big business compared to the profitability of all US businesses. What Figure 9 shows is that this differential markup rises and falls with the wholesale price index, a measure of inflation. In other words, inflation benefits big business’s bottom line.

This update is from the [Capital as Power forum](#).

Notice how this evidence changes your view of inflation. It makes it hard to blame government for the problem. You see, if big business is systematically benefiting from inflation, it implies that these big corporations are raising prices *faster* than everyone else. In other words, it is oligopolies that are driving inflation.



**Figure 9: Inflation benefits big business**

This is Nitzan and Bichler’s chart illustrating how inflation relates to the profitability of US big business. The thick line shows what Nitzan and Bichler call ‘differential markup’ — the markup of the 500 largest US public firms (ranked by market capitalization) relative to the markup of all US corporations. The dashed line shows the wholesale price index, a measure of inflation that focuses on wholesale commodities. An earlier version of this research can be found in [Chapter 16](#) of *Capital as Power*.

So it seems that in the real world, inflation looks nothing like it does in economics textbooks. Yes, inflation is a ‘monetary phenomenon’ — as is anything to do with prices. But more importantly, inflation is a *power struggle* over who can raise prices the fastest.

## Support this blog

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

Become an **ETD** supporter

## Sources and methods

Data for commodities on the US consumer price index comes from the Bureau of Labor Statistics. I downloaded every unique series from their [bulk-download facility](#), and then aggregated it into a single database. For the analysis of individual commodities, I used the lowest level of aggregation. That is, I kept all commodity groups that did *not* have a subgroup.

For the  $i$ th commodity, I define price change,  $\Delta P_t^i$ , as the percentage change in price between year  $t$  and year  $t + 1$ :

$$\Delta P_t^i = \left( \frac{P_{t+1}^i}{P_t^i} - 1 \right) \times 100$$

In Figure 6, the shaded region shows the 95% confidence interval, in each year, of  $\Delta P_t^i$ . In Figure 8, the vertical axis plots the standard deviation, in each year, of  $\Delta P_t^i$ .

## Constant price-change variation

Here's how I calculate what price-change variation would look like if it were constant (dashed red lines in Figure 7). In the empirical data, I calculate, in each year  $t$ , the 2.5%, 50% and 97.5% percentiles for price change among all commodities (notated as  $P_t^{2.5}$ ,  $P_t^{50}$ , and  $P_t^{97.5}$ ). Then I calculate the average distance (using the geometric mean) between the upper/lower percentiles and the median price:

$$\bar{P}^{97.5} = \text{geometric mean}(|P_t^{97.5} - P_t^{50}|)$$

$$\bar{P}^{2.5} = \text{geometric mean}(|P_t^{2.5} - P_t^{50}|)$$

Finally, I take the official inflation rate in each year, measured using the CPI, and add to it  $\bar{P}^{97.5}$  and subtract  $\bar{P}^{2.5}$ . The result is a 95% confidence interval on price change ... if price-change variation was uniform over time (which it is not).

## Further reading

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

Friedman, M. (1963). *Inflation: Causes and consequences*. Asia Publishing House.

Nitzan, J. (1992). *Inflation as restructuring. A theoretical and empirical account of the US experience* (PhD thesis). McGill University.

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