Economics from the Top Down

new ideas in economics and the social sciences

From Commodity to Asset: The Truth Behind Rising House Prices

Blair Fix

August 22, 2024

Practical men who believe themselves to be quite exempt from any intellectual influence, are usually the slaves of some defunct economist.

— John Maynard Keynes, 1936

When it comes to rising house prices, nearly everyone has a theory about the cause. There's 'too much foreign money'. There are 'too many immigrants'. There's 'too little construction'. And so on.

What unites these explanations is that they appeal, in some way, to the idea that rising prices are caused by a mismatch between supply and demand. And surely that's true, right?

Yes, it *is* true . . . in the same way that death is caused by dying. But of course, that's circular logic. And so it goes with 'supply and demand'. Since prices are always caused by the interplay between what we *want* and what we can *get*, evoking 'supply and demand' leads us pretty much nowhere. Worse, it often puts the focus on short-term patterns, when the real scientific payoff lies in studying price trends over the long term.

Speaking of the long term, many people assume that rising house prices are a recent problem. But in the United States, the pattern dates to the early 1970s. For almost a century before that, US house prices had been *dropping* against income. And so Americans treated their house like a 'commodity' — a thing they bought to live in. But from 1972 onward, house prices began to slowly *appreciate* against income. And so Americans started to treat their house like an 'asset'.

It's this transformation — from commodity-like depreciation to asset-like appreciation — that is the real story of house prices. And the truth is that this story can't be understood using popular scapegoats. To see why US house prices headed south and then north, we need to forget about supply and demand and instead, peer into the belly of industrialism. We need to ground house prices in the use of energy.

Now, if going from *prices* to *energy* sounds like a non sequitur, I'll show you why it makes sense. And I'll show you how, when we bring debt into the energy fold, we can explain almost all of the historical variation in US house prices.

The lesson here is simple yet disturbing. When it comes to rising house prices, the trend has less to do with a 'supply crisis', and more to do with basic physical limits to industrial supply chains.

The price behavior of commodities and assets

Before we dive into the changing nature of house prices, it's helpful to start with a more general question: what's the difference between a 'commodity' and an 'asset'?

On this topic, there's much confusion. And that's largely because the distinction between the two concepts is ideological. It has less to do with the things themselves and more to do with the *goal* of ownership. People buy 'commodities' to *use* them. But they buy 'assets' as an *investment*.[1]

Crucially, these goals depend circularly on prices. In other words, what makes people treat property like a 'commodity' is that its value *depreciates* over time. And what makes people treat other property like an 'asset' is that its value *appreciates* over time.

Take cars as an example. Most cars lose value over time — a fact that makes people treat them like a commodity. If you buy a car, it's because you want to drive it. That said, some cars — old Ferraris, for example — get more valuable with time. When a collector buys one of these cars, they might never get behind the wheel. Instead, they treat the thing as an asset — a ledger line whose purpose is to appreciate in value.

[1] For an excellent philosophical dive into the distinction between commodities and assets, see Jesús Suaste Cherizola's paper 'From Commodities to Assets: Capital as Power and the Ontology of Finance'.

Now, when we talk about 'depreciation' and 'appreciation', we usually mean of a thing's nominal price. However, when we look at long-term trends, nominal prices aren't useful, because inflation makes everything appreciate in value. So to study long-term price patterns, we need a goalpost that moves with the times. In this post, I'll use *average income* as the benchmark.

When we judge prices against average income, we find a clear distinction between property that's considered a 'commodity' and property that's considered an 'asset'. The price of commodity-like property tends to *depreciate* against income, while the price of asset-like property tends to *oscillate*.

Let's visualize these patterns using historical data. We'll start with the price of commodity-like property, shown in Figure 1. Here, I've pegged the US consumer price index against the average American income. As you can see, the price of these consumer items depreciated drastically over the last two centuries.

Turning to asset-like property, we know that the same depreciation pattern can't hold. If it did, investors would always lose their shirts, and capitalism would have long ago collapsed. Since capitalism lives on (for now), we can deduce that asset-like property follows a different pricing pattern.

To see this pattern, let's look at the price of corporate stocks, shown in Figure 2. Here, I've pegged the value of the S&P 500 against the average US income.[2] The result is a long-term *oscillation*. Over the last century, stock prices kept pace with rising American income, albeit with some big swings in investors' fortunes. It's because of this price-holding pattern that corporate stocks are treated as an 'asset'. (That's not the only reason, of course.)

To summarize, long-term price patterns tells us about how people view their property. If property consistently *loses* value against income, it's treated as a 'commodity' — a thing you buy to use. But if property *maintains* its value against income, it gets treated as an 'asset' — a thing you buy as an investment.

[2] I should point out that the ratio between the S&P 500 and average US income is a variation on what Jonathan Nitzan and Shimshon Bichler call the 'power index'. For more details, see their paper 'A CasP model of the stock market'. Also see my post 'Stocks are up. Wages are down. What does it mean?'.

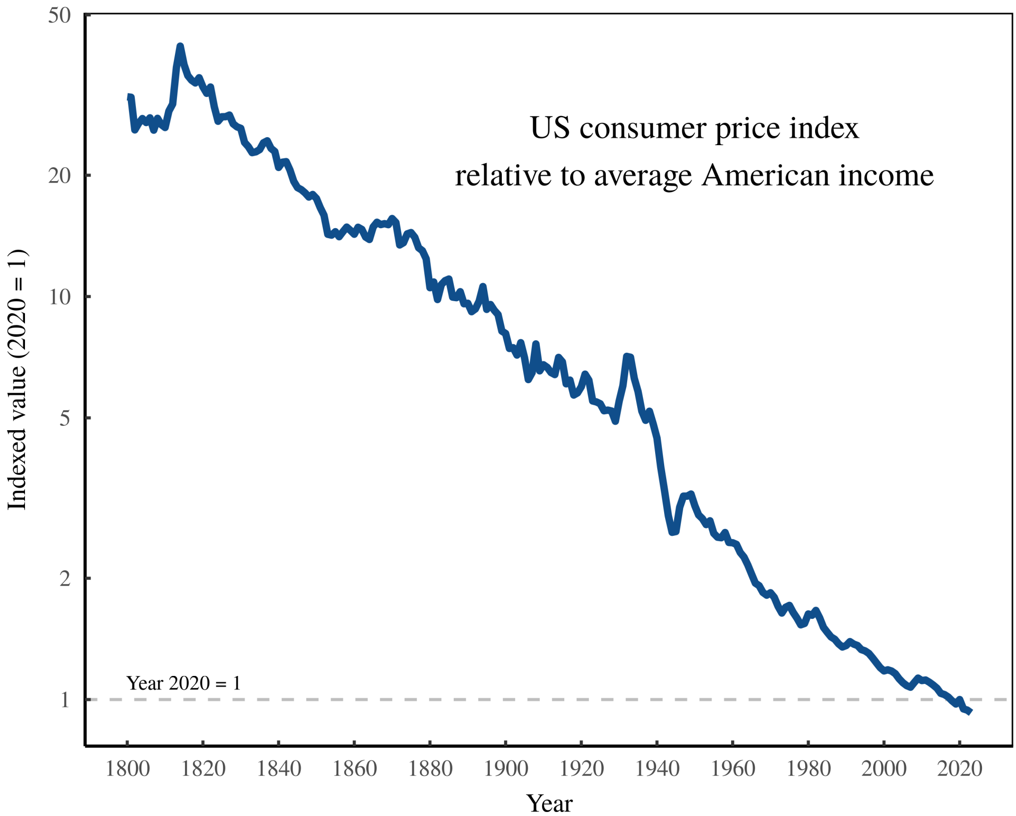

Figure 1: The depreciating relative value of US consumer items — typical price behavior of commodities

Since 1800, items tracked by the US consumer price index have seen their price drop over thirty-fold relative to the average American income. Note that the vertical axis uses a log scale. Sources and methods

US house prices: From commodity to asset

Now to housing. When we peg the price of US houses against the average American income, what pattern will we find? Will house prices depreciate like a commodity? Or will they maintain their value like an asset?

Turning to the historical data, the answer is ... both.

Figure 3 runs the numbers. From 1890 to 1972, US house prices behaved like a commodity, consistently declining against the average American income. But after 1972, US house prices behaved like an asset, (slowly) appreciating against income.

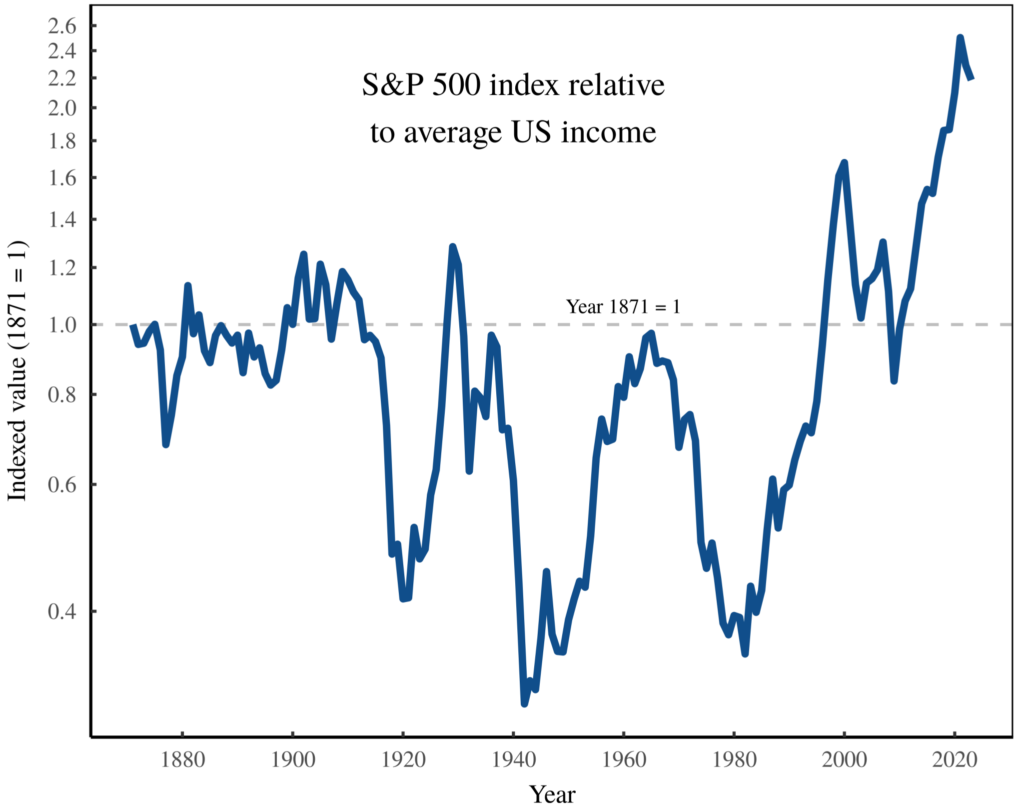

Figure 2: The oscillating relative value of US corporate stocks — typical price behavior of assets

Since the late 19th century, the price of the S&P 500 has kept pace with average US income. Of course, over the short term, there are periods of stock booms and busts. But this volatility is what motivates speculation. Note that the vertical axis uses a log scale. Sources and methods

---

Before we search for the cause of this house-price transformation, it's worth commenting on the difference between the two price environments. In today's asset-like environment, the game is to buy a house before prices outpace your income. For many people, a house purchase requires borrowing (or inheriting) money from parents, as well as taking on Herculean levels of debt.

In contrast, two generations ago, American house buyers pursued a different strategy. In the commodity-like environment of the 1950s, purchasing a house was similar to buying consumer electronics today. The expectation (and the reality) was that each year, houses got cheaper relative to income.

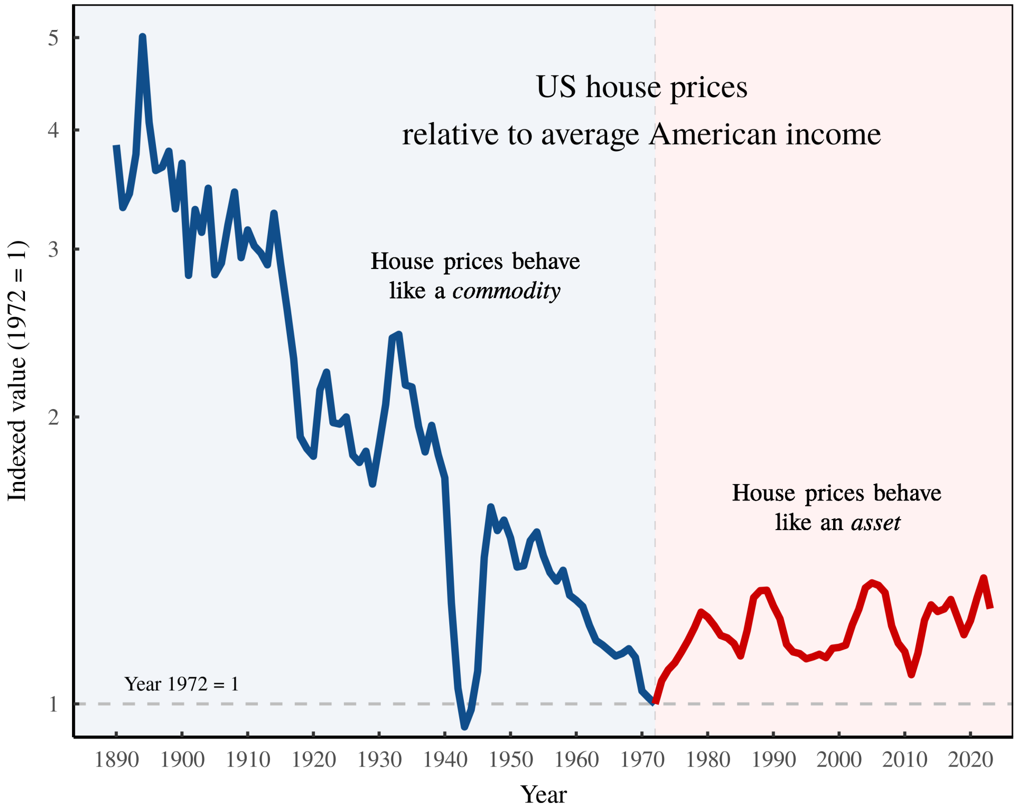

Figure 3: From commodity to asset — the long-term history of US house prices

This chart plots the (indexed) value of US house prices, pegged against the average American income. Prior to 1972, houses behaved like a commodity, depreciating with time. But after 1972, houses started to behave like an asset, with prices that oscillated (and slightly rose) against income. Note that the vertical axis uses a log scale. Sources and methods

---

This depreciation meant that real-estate was not a game of 'investment', and there was no pressure to 'buy now'. If you wanted a house, you simply waited for prices to drop such that you could afford one.

My how things have changed.

Case (temporarily) closed

Let's turn now to what's driving the trend in US house prices. For their part, mainstream economists already know the answer. Prices, they say, are everywhere and always caused by the interplay of supply and demand. As such, the task is to take this pre-determined hypothesis and find data to support it.

For the moment, I'll play the game.

To quantify the 'supply' of US houses, I'll look at the pace of construction, measured in terms of the number of 'housing starts'. And to quantify the 'demand' for (new) houses, I'll use the annual change in US population. The idea is that when the population grows, so does the demand for housing. So if the supply of new houses stagnates while the population expands, people will bid up prices and house will get more expensive. It's simple economics.

And it's supported by evidence.

Figure 4 runs the numbers. Here, the grey curve shows the number of US housing starts, measured relative to the annual change in US population. The red curve illustrates the trend in this data — a trend which conspicuously mirrors the long-term pattern in house prices. As construction rates increased, house prices dropped. And when construction stagnated, house prices gradually rose.

And there we have it folks, the case is closed. The problem of rising house prices is caused by a mismatch between supply and demand.

I kid, of course. For (at least) three reasons, I find this standard explanation dissatisfying. First, the correlation in Figure 4 is not compelling. Notice that I had to *smooth* the construction data to see that it resembles the house-price data. If prices were completely determined by the interplay of supply and demand, data smoothing should be unrequired.

Second, the appeal to supply and demand raises more questions than it answers. For example, if an expanding population requires more houses, why doesn't the construction industry supply them automatically? After all, that's what we'd expect in a fairy-tale free market. Confusingly, there are moments in US history when housing construction far outpaced population growth. And there are other moments when the construction industry dropped the ball. Why? If the investigation begins and ends with 'supply and demand', we'll never know the answer.

Third, it seems plausible that in Figure 4, cause and effect might be reversed. Instead of construction rates driving prices, it could be that house prices drive the pace of construction. In other words, it's people's willingness to buy new houses that sets the pace of new house construction. So when prices rise, people become less willing to buy, hence fewer new houses get built.

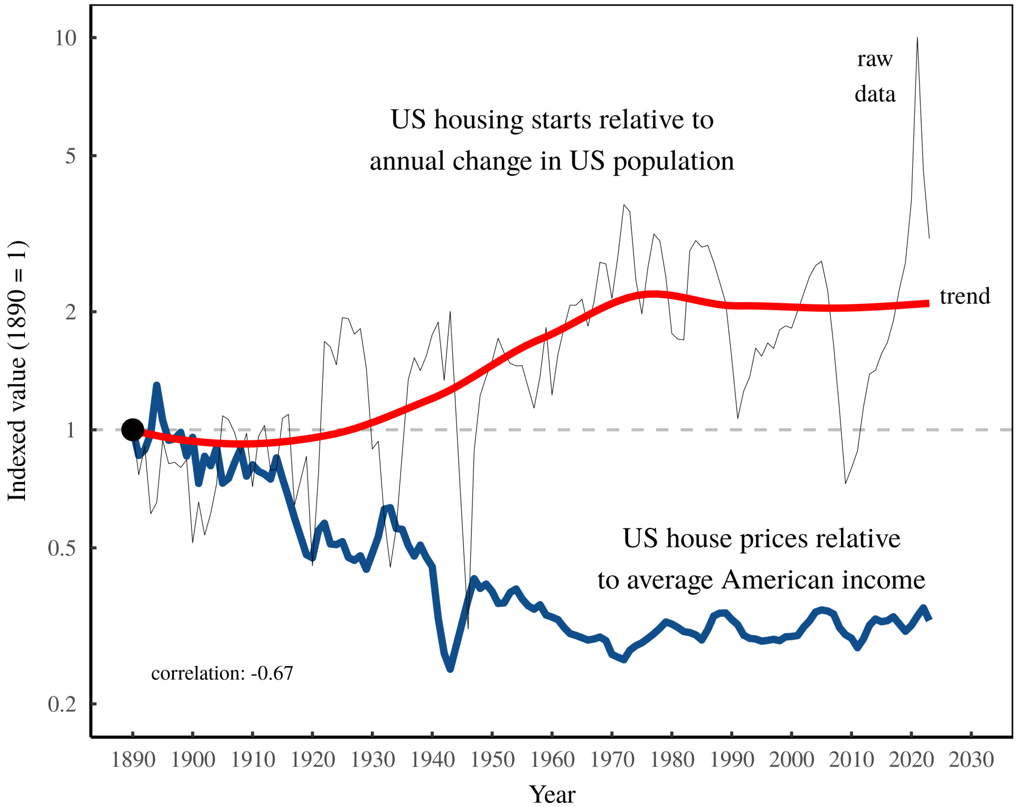

Figure 4: Are house prices driven by a mismatch between supply and demand?

The grey curve plots the number of US housing starts per new American resident. The red curve shows the smoothed trend in this ratio. The blue line shows the price of US houses, pegged against average American income. Both series are indexed to be one in 1890. Note that the vertical axis uses a logarithmic scale. Sources and methods

In short, when I see an explanation based on supply and demand, it's like finding a sealed box labelled "*interesting thing inside . . . do not open*". I am, of course, going to open it.

How stuff gets cheaper

To peer inside the sealed box of house prices, we need to leave supply and demand behind and instead, start with a basic question: how does stuff get cheaper?

The answer is that there are two main routes. To reduce prices, businesses can either increase workers' productivity or they can slash workers' wages. If these two routes seem obvious to you, feel free to skip ahead. But for detail-oriented readers, here's the math.

***

In economics, nearly all obvious truths start from accounting definitions. Equations for a commodity's price are no exception.

To define a price, we'll begin with the basics of double-entry book keeping. When you buy something, your *expense* becomes someone else's *income*. As such, we can define a commodity's price by adding up this income and then dividing by the number of commodities sold:

$$\text{price} = \frac{\text{gross income}}{\text{number of commodities sold}}$$

Let's assume that it's a firm who receives this income. From its gross income, the firm then sends money to its owners, its employees, and its suppliers. Of these expenses, payments to workers typically dominate. So to simplify the math, we'll pretend that the worker payroll is the only expense.[3] Hence in rough terms, a commodity's price is set by the size of the firm's payroll per item sold:

$$\text{price} \approx \frac{\text{worker payroll}}{\text{number of commodities sold}}$$

Next, we note that a firm's payroll is set by the number of workers it employs, multiplied by their average pay. So we can update our price equation as follows:

$$\text{price} \approx \frac{(\text{number of workers}) \times (\text{pay per worker})}{\text{number of commodities sold}}$$

[3] On the aggregate scale, money sent to suppliers eventually gets paid to other workers and other owners. So on this scale, the 'supplier' term vanishes from our price equation. And for most firms, profit margins are small relative to the wage bill. So as an order-of-magnitude approximation, we can define commodity prices in terms of the wage bill per item sold. And no, that doesn't mean the labor theory of value holds.

Finally, note that if we divide the number of commodities sold by the number of workers employed, we're calculating 'worker productivity'.[4] Looking at our price equation, we've got the *inverse* of this quantity. So with a quick substitution, we get our final formula:

$$\text{price} \approx \frac{\text{pay per worker}}{\text{worker productivity}}$$

Now let's return to our original question: how do commodities get cheaper? According to our equation, there are two options. To reduce prices, we can either:

a) *decrease* workers' pay;

b) *increase* workers' productivity.

Price reductions through productivity

Given the two routes to cutting prices, we want to know the path that dominates. Historically, the answer is that *productivity* has been king. And it's easy to understand why. Cutting prices with wage clawbacks elicits pushback from workers. But productivity gains? *Everyone* likes those.

In a moment, I'll show you the tight connection between worker productivity and prices. But first, a word about measurement.

[4] An important caveat is that for many activities, 'selling' has nothing to do with 'producing'. For example, when toll workers sell access to a highway, nothing is 'produced'. The lesson is that productivist language comes with all kinds of problems which I'm ignoring here.

Given how frequently economists blather about productivity, you'd think they would have a sound way to measure it. But they *don't*. Instead, economists dubiously equate 'productivity' with income, leading to official data that's laced with toxic ideology.[5] Since I refuse to use this sludge, I'm left devising my own measures of productivity — ones that are reasonably objective.

The key to these objective measures is that they look for activities where we can use basic science to aggregate output. The energy sector is a good example. When the goal is to produce energy, the science of energetics leaves no wiggle room for ideology — it tells us exactly how to compare different resources. So it's in the energy sector where we'll measure productivity and connect it with prices.

To make the connection, let's look at the price of *oil*. Today, many people assume that oil prices are driven by speculation. And to some extent, that's true. But over the long-term, the price of oil is set by something more mechanical: the rising productivity of oil-and-gas workers.

Figure 5 shows the relation. Here, the red line shows the (indexed) productivity of workers in the US oil-and-gas sector. From 1865 to the early 1970s, this productivity rose by a factor of roughly fifty. In contrast, the blue line

[5] The reason that productivity is difficult to measure (objectively) is deceptively straightforward: different people do different tasks.

To see the problem, ask yourself who is more 'productive': a doctor who treats four patients per hour or a farmer how grows four pounds of potatoes per work hour? The correct answer is that the question is *ill-defined*. To judge 'doctoring' against 'potato farming', we need to supply a criteria for comparing the two tasks. And since this criteria is subjective, so is any comparison of 'productivity'.

Not to worry, say mainstream economists. The goal of 'production', they say, is to create 'utility'. And since prices 'reveal' utility, we can use prices to sum the quantity of utility produced, and the resulting levels of 'productivity'.

Not so fast. There are two big problems with this logic. First, it is deeply ideological, because it ends up equating productivity with income — an assumption that justifies the bloated incomes paid to the rich and the pittance paid to the poor. Second, the reliance on prices leads to fundamental ambiguity in the measurement of output — ambiguity caused by the simple fact that prices are an *unstable* unit of measure.

For more detailed criticism of how mainstream economist measure productivity, here's a collection of essays:

- No, Productivity Does Not Explain Income
- Productivity Does Not Explain Wages
- Debunking the 'Productivity-Pay Gap'
- Productivity and Income ... Again
- Real GDP: The Flawed Metric at the Heart of Macroeconomics
- The Aggregation Problem: Implications for Ecological and Biophysical Economics

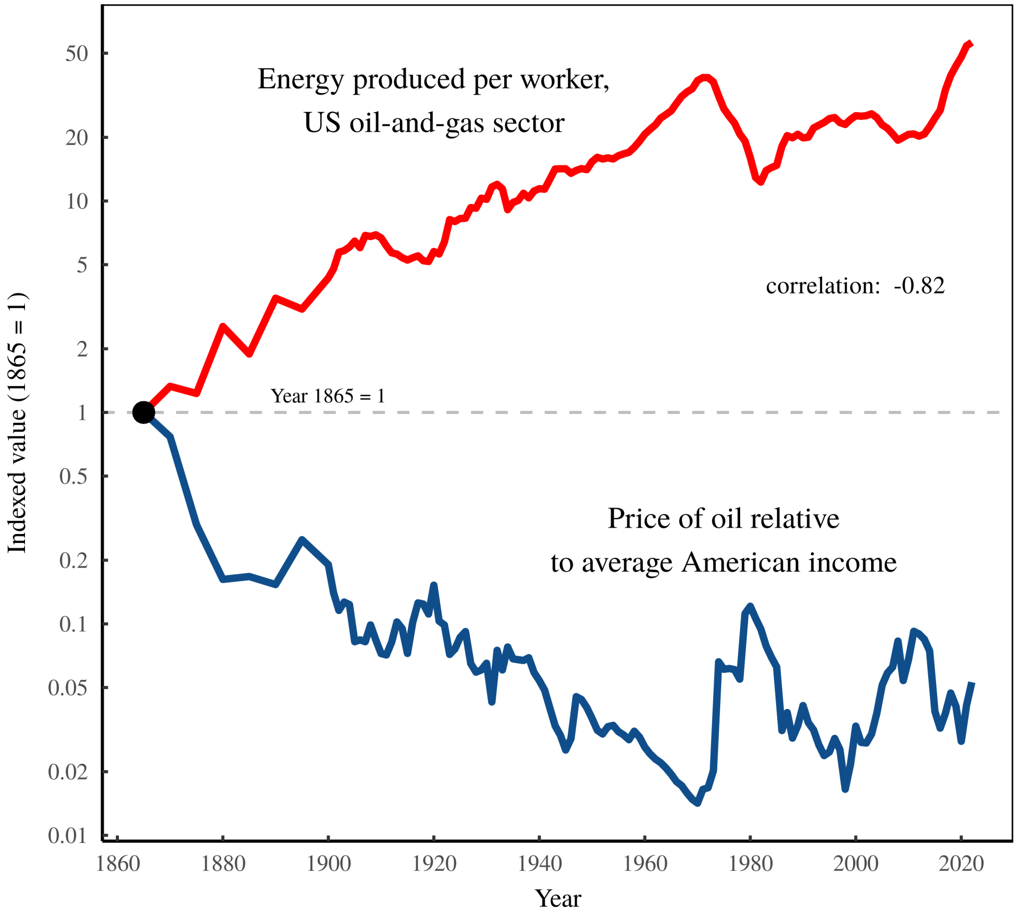

Figure 5: Cheaper oil, underwritten by rising productivity

The red curve shows the indexed productivity of workers in the US oil-and-gas sector (measured in terms if energy production). The blue curve shows the price of oil pegged against the average American income. Note that the vertical axis uses a log scale. Sources and methods

---

shows the price of oil, pegged against the average American income. Over the same period, this price *dropped* fifty-fold.[6] In short, oil got cheaper through rising productivity.

[6]Detail-oriented readers might protest that if I'm going to look at *oil* prices, I should measure the productivity of *oil* production only — not the aggregate of *oil and gas*. Well, I would if I could. The problem is that US employment statistics don't separate the two tasks, and for good reasons. Wells that produce oil often also produce gas. In fact, in the early days of oil drilling, the gas was deemed waste and burnt off on site. Even today, gas flaring remains common, and is a major source of methane emissions.

Productivity through energy gluttony

It's at this point that energy enters the equation. That's because, despite what mainstream economists preach, productivity gains don't flow magically from the workings of the free market.[7] Instead, rising productivity is driven almost exclusively by the gluttonous use of energy.

The reason is straightforward. Basic physics tells us that making stuff takes energy. And so making *more* stuff takes *more* energy. Put another way, when people get more productive, it's largely because they've subcontracted the tough work to energy slaves. For example, a locomotive driver can haul millions of times more stuff than someone hauling a handcart. But this productivity boost isn't magic; it's underwritten by the locomotive's profligate use of diesel fuel.

Returning to the US oil sector, we know that over the 20th century, oil-and-gas workers got vastly more productive. Unsurprisingly, this productivity was no miracle; it was underwritten by energy gluttony. Figure 6 runs the numbers. As US oil-and-gas *output* per worker grew (blue curve) so did the energy *consumed* per worker (red curve). The message is that when it comes to rising productivity, there are no shortcuts. It's energy all the way down.

The energetics of house manufacturing

Back to our main story — the long-term trend in US house prices. The point of my sidetrack into energetics was to make the connection between energy and prices. In the US oil business, cheaper prices were driven by rising productivity, which in turn, was driven by energy gluttony. We expect the same to be true with house prices. And it is.

The catch is that with housing, I'll cut out the productivity middleman because I don't think construction productivity can be measured objectively. So I'll jump straight from energy use to house prices.

Actually, there's another catch, which is that I'll compare US house prices to the energy used in US *industry*. ('Industry' consists of the sum of agriculture, mining, construction, and manufacturing). Here's why this comparison makes sense.

[7] Sadly, 'magic' is the right word. According to the neoclassical theory of economic growth, productivity gains are mostly attributed to an unknown residual term opaquely dubbed 'total factor productivity'. The term is a fudge factor that says 'insert magic here'.

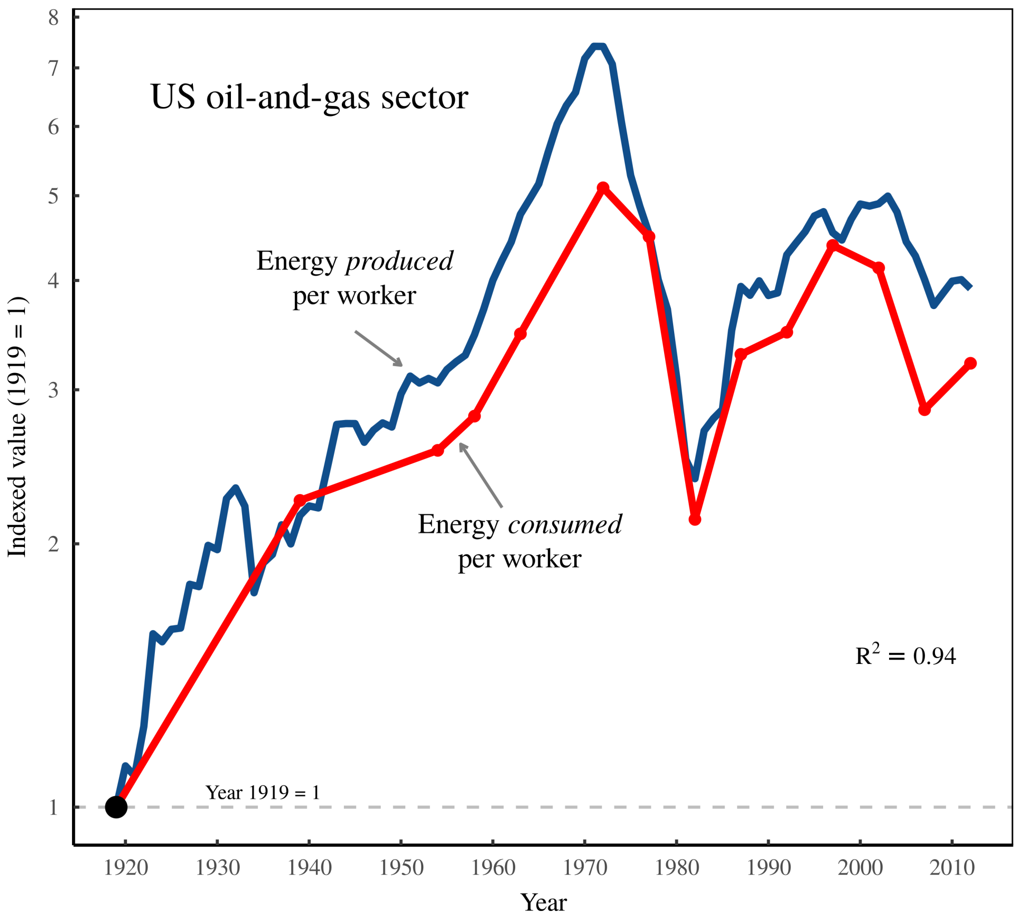

Figure 6: Rising oil-and-gas productivity, paid for with energy gluttony
The blue curve plots the energy produced per worker in the US oil-and-gas sector. The red curve shows the energy consumed per worker in the US oil-and-gas sector. Note that the vertical axis uses a log scale. Sources and methods

---

Although we typically say that houses are 'constructed', it's more apt to say that they are *manufactured* and then *assembled*. In other words, if you look at a modern house, you'll find a wide array of industrially-made components — things like lumber, concrete, pipes, hose, drywall, paint, windows, shingles, wiring, not to mention thousands of screws and nails.

To make houses cheaper, the construction business must get more productive at assembling housing components. But in addition, the whole industrial supply-chain must get more productive at pumping out the required parts and materials. And since productivity is underwritten by physical work, we expect that the affair will be visible in the book of energy.[8]

Let's read the story. As Figure 7 demonstrates, there's a tight connection between cheapening US house prices and rising energy use per worker in US industry. The message is that Americans made houses more affordable by throwing energy at the problem. And when this energy throwing stagnated, houses stopped getting cheaper. It's that simple.

The debt antagonist

Of course, the story of US house prices doesn't end here. That's because if energy is the *protagonist* of our tale, then debt is the *antagonist*.

Let me explain.

When consumers are given access to credit, they have the option to use debt to bid up prices. For example, I could walk into Walmart and, with my credit card in hand, offer to buy bananas for double their listed price. True, the strategy sounds insane when applied to bananas. But when applied to *houses*, debt-driven bidding is standard fare. In other words, house buyers routinely make bids based on the amount of credit they can procure. As such, we expect that house prices will be connected to levels of household debt.

[8]There's a common perception that the price of housing is dominated by the price of land. But that's only true in dense urban centers. For the average new American home, it's construction costs that dominate the price.

Engineer Brian Potter has an insightful article that breaks down these costs. For the average new build, land makes up about 18% of the final price. In contrast, construction costs constitute about 56% of the finished price. And of these construction costs, about half consists of materials alone.

Looking at these materials, Potter worries that there's "no obvious path for making [them] cheaper". But actually, that's not true. There *is* an obvious path for making materials cheaper: throw energy at the whole industrial sector. Of course, the problem with this obvious path is that it's unlikely to happen.

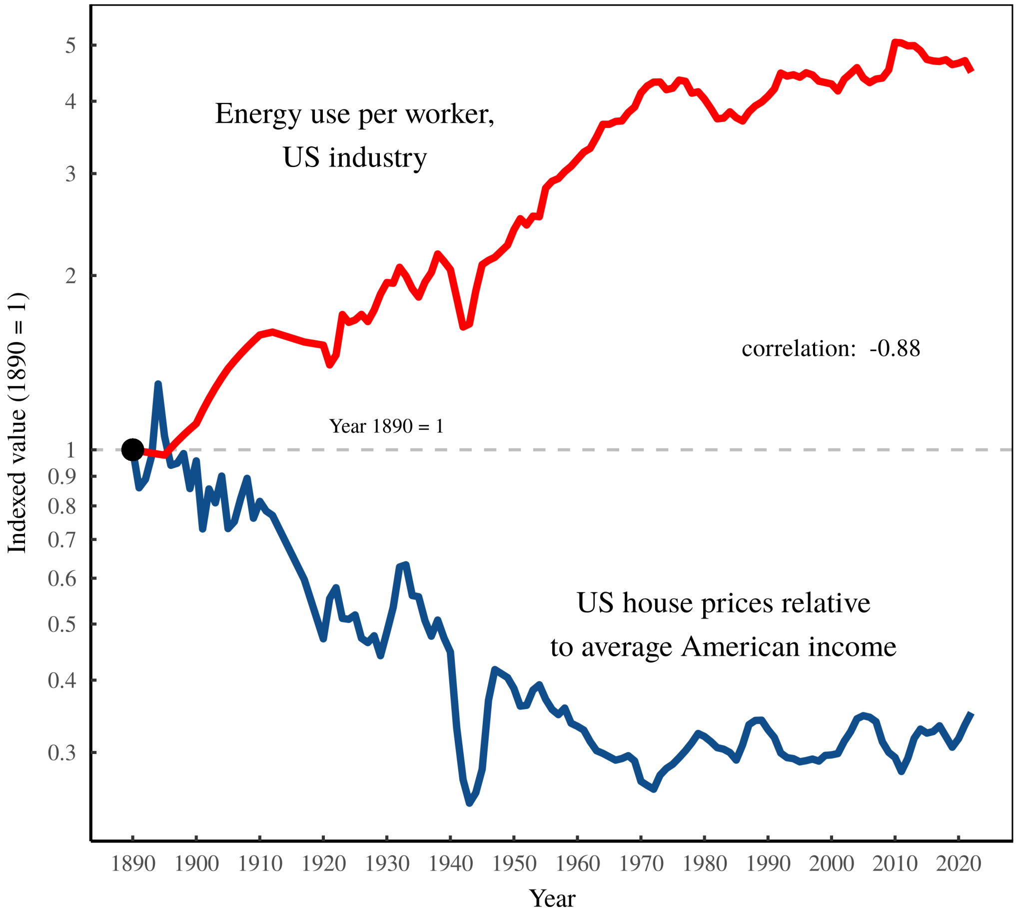

Figure 7: Cheaper houses, paid for with profligate energy use

The red curve shows the energy consumed per worker in US industry. ('Industry' consists of agriculture, mining, construction and manufacturing.) The blue curves shows US house prices, pegged against average US income. Both series are indexed to equal 1 in 1890. Note that the vertical axis uses a log scale. Sources and methods

---

Perhaps more than any other thinker, it's renegade economist Steve Keen who's made the connection between prices and debt. For example, in a piece called 'Your Margin and Your Life', Keen shows how the recent rise of US house prices corresponds to a conspicuous build-up of household debt.[9]

To continue the story of US house prices, I'm going to extend Keen's analysis back in time. Figure 8 shows the results. Over the past century, changing US house prices were tightly coupled to changing levels of American household debt.

[9] I was recently reminded of Steve Keen's work on debt when reading Hilliard MacBeth's excellent book on the Canadian housing market, *When the Bubble Bursts*. Full disclosure: Hilliard is a long-time patron of my research.

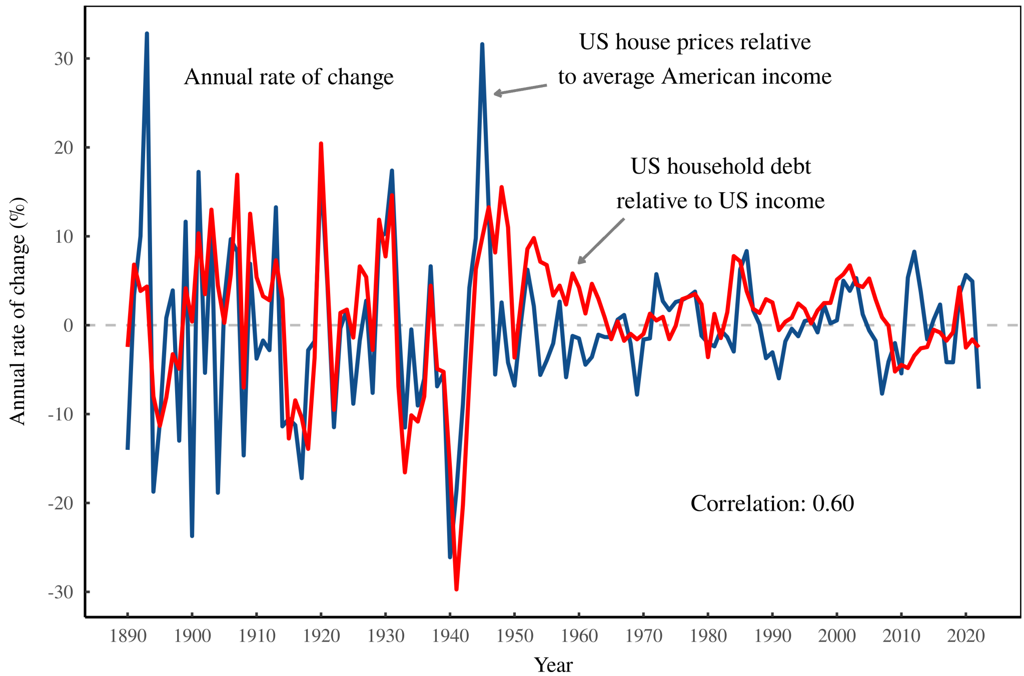

Figure 8: Using debt to bid up house prices

The blue curve shows the annual percentage change in US house prices, pegged against average American income. The red curve shows the annual percentage change in household debt, measured relative to total US income. Sources and methods

---

A tale of energy and debt

At this point, our house-price story has two main characters. The protagonist is industrial energy use, which drives prices *down*. And the antagonist is household debt, which drives prices *up*. Now, it would be naive to think that these two characters are the only factors affecting US house prices.[10] That said, the evidence suggests that energy and debt dominate the story.

Figure 9 makes the case. Here, the red curve is a model of US house prices based solely on historical levels of industrial energy use and household debt. The model explains about 90% of the variation in US house prices. In short, energy and debt may not be the only factors affecting US house prices, but they seem to be the most important players.

[10]Journalist Matt Stoller argues that in the US, a home-builder cartel is driving up house prices, largely by buying land and then *not* building houses. I find the idea intriguing. That said, Stoller mostly demonstrates that in major US markets, big financiers dominate the house-building industry. What he doesn't do — and what someone *should* do — is test how this concentration relates to prices.

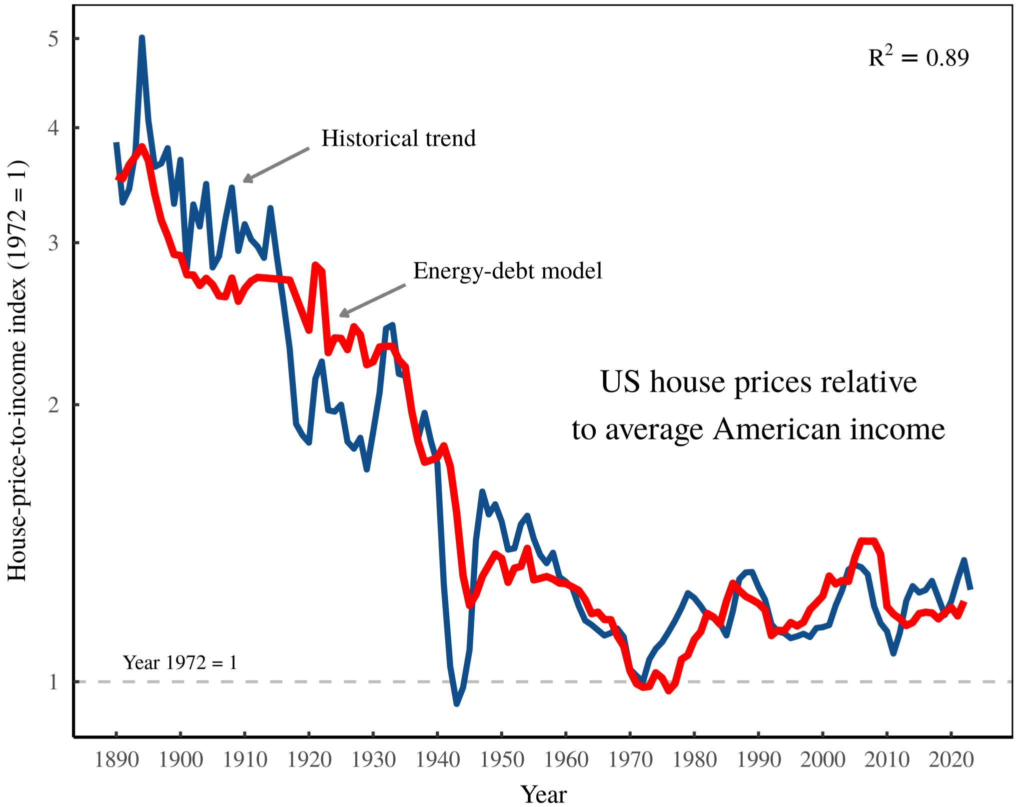

Figure 9: US house prices — a story of energy and debt

Here, the blue curve shows the historical trend in US house prices, pegged against average American income. The red curve shows a model of US house prices based solely on energy and debt. For the math behind the energy-debt model, see the appendix. Note that the vertical axis uses a log scale. Sources and methods

---

(For the math behind the energy-debt model, see the appendix.)

Reopening the case of short supply

With energy and debt in hand, let's return to the case of a short housing supply. Back in Figure 4, I showed you that construction rates correlate with the trend in US house prices. But I remained skeptical that the house-price story began and ended with a mismatch between supply and demand.

The truth is that part of my skepticism was driven by an ace up my statistical sleeve. You see, I'd already run the numbers (below) and found that in the context of energy and debt, rates of US construction have virtually no ability to explain house prices.

Let's have a look at this evidence.

It's based on something called 'partial correlation' — the correlation between two variables once the modelled effects of other factors has been removed.[11] Looking at US house prices, I've identified three possible causes:

1. levels of industrial energy use;
2. levels of household debt;
3. levels of construction relative to population growth.

To run a partial correlation, we take one of these factors and measure its correlation with US house prices, after having removed the modelled effects of the other two factors. The purpose of this type of analysis is less to establish causation and more to *disprove* it. You see, if a correlation vanishes when put in the context of other variables, the factor in question is unlikely to be a cause. And as it turns out, that's exactly what happens with the construction data.

Figure 10 does the math. Here, each column shows the correlation between US house prices and the variable on the horizontal axis. The blue columns show the standard *univariate* correlation. And the red columns show the *partial* correlation — the correlation once the modelled effects of the other two variables has been removed.

What interests us is the *difference* between the blue and red columns — the change in correlation once we introduce context. On that front, the correlation for energy use hardly changes, while the correlation for household debt dramatically switches from negative to positive.[12] But what's most important is the data on the right — the collapsing evidence that house prices are driven by a supply shortage. When we put construction rates in the context

[11] It's a pet peeve of mine when social scientists claim to have 'adjusted' for factor x. No you haven't. What you did is build a *model* of x, and then use the *model* to make a statistical adjustment. The distinction is important, because if the model is wrong, then so is the 'adjustment'.

[12] In Figure 10, you might be wondering why the debt correlation flips signs when switching from a univariate to a partial correlation. Here's why. Over the long term, household debt levels have little to do with the trend in US house prices. (The univariate correlation is positive, but I don't think it means anything.) Instead, debt levels are strongly related to *short-term* variations in house prices. (See Figure 8.) And it's this short-term variation that appears once the modelled effects of industrial energy use have been removed. (From Figure 7, we know that energy explains most of the long-term trend in US house prices.)

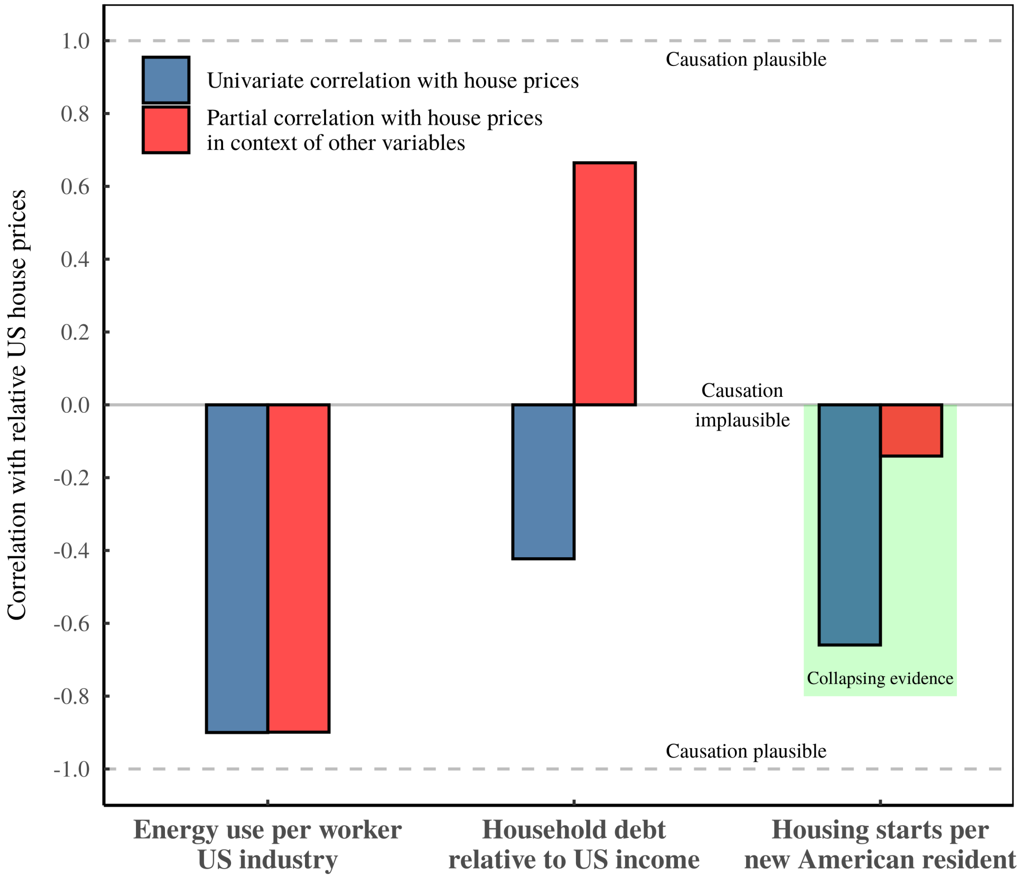

Figure 10: A case of vanishing evidence

This figure shows how three factors — (1) energy use per worker in US industry, (2) household debt relative to US income, (3) housing starts per new American resident — correlate with the price of US houses (pegged against average American income). The blue columns show the standard univariate correlation. Red columns show the partial correlation — a variable's correlation with house prices once the modelled effects of the other two factors has been removed. The results signal trouble for the supply-shortage thesis. In the context of energy and debt, construction rates show virtually no correlation with US house prices. Sources and methods

of energy and debt, their ability to explain house prices all but vanishes. The message is that stagnating construction is almost surely *not* what's driving the long-term trend in US house prices.

A supply-chain problem

In many ways, the key to a good brain worm is its ability to evoke circular logic. On that front, appeals to a housing 'supply shortage' are consistently circular. When people refer to a 'housing shortage', they typically mean that there's a shortage of *affordable* homes. But that's just another way of saying that houses are *expensive*.

Since this circularity goes largely unnoticed, it leads to numerous oxymorons. For example, a recent CBC article complains that "in the midst of a housing supply crisis, thousands of condos in Canada's largest cities are sitting empty". Tortured rhetoric aside, what the authors mean is that despite a glut of empty condos, prices are still unaffordable.

And so we're back to square one. If it's not a supply shortage, then what's making houses so expensive? The evidence suggests that the problem is less about 'supply' and more about the limits of *supply chains*.

Unlike how we conjure money from nothing, humans cannot 'will' housing into existence. Instead, houses are built by tirelessly wrenching resources from the Earth and transforming them, bit by bit, into the thing we call home. The process takes a tremendous amount of work, most of which is done cheaply by the energy from fossil fuels.

For much of the 20th century, Americans were able to make houses cheaper by ramping up the energy subsidy that went into their creation. But since the early 1970s, that subsidy has plateaued. And in my mind, it's unlikely to resume rising. So today, we sit at an impasse. If we can't throw energy at the problem, how else can we make houses cheaper?

Why, by cutting wages of course!

The offshored house?

A quick trip to Walmart shows that American businesses are adept at cutting prices by slashing wages. And no, I don't mean the crappy wages paid to Walmart's staff. I mean the rock-bottom wages paid to the army of *offshore* workers who make the stuff that Walmart sells. It's a nifty trick. If a pesky minimum wage stops you from cutting domestic pay, just move your whole supply chain elsewhere.

Turning to houses, IKEA has successfully offshored the production of home furnishings. Perhaps the same principle could be applied to whole homes? Well, shipping houses across land and sea would create obvious logistical problems. But let's ignore these issues and focus on more fundamental limits to offshoring.

The dirty secret behind offshoring is that it makes commodities cheaper only if *domestic wages remain high*. In other words, Walmart is able to hawk 'cheap' imported stuff only to the extent that Americans earn much more than the offshore workers who made the goods. But if American wages decline, then the offshore-made stuff will no longer appear 'cheap'.

As it turns out, this domestic wage decline is already happening. The consequence is that, for the poorest Americans, it's not just houses that are getting more expensive. The price of *everything* is going up.

Figure 11 makes the case. Here, the line shows the price of consumer items, benchmarked against the income of the *tenth* American percentile. For these folks, nearly everything has been getting more expensive since the 1980s. In short, for the poorest Americans, offshoring didn't 'work'.

Requiem for an American Dream

It's at this point that tech optimists would feel obliged to highlight an easy technological solution to our housing woes. But I won't offer one, because I don't think it exists.

If anything, the house-price problem is more severe than we might think. You see, while house prices are rising slowly against the average American income, they're *exploding* relative to the income at the *bottom*. For the poorest tenth of Americans, houses are now as expensive as they were at the turn of the 20th century. Figure 12 runs the (alarming) numbers.

In short, if the goal is ever-cheaper housing, then in my mind, there is no solution. This version of the American dream is dead and we should let it rest in peace. But that doesn't mean all is lost.

Whether cheap or expensive, having a place to live is a basic human right. And long before people wielded the power of fossil fuels, we found ways to house our fellow humans. I'm sure we can do it again. We just need to stop

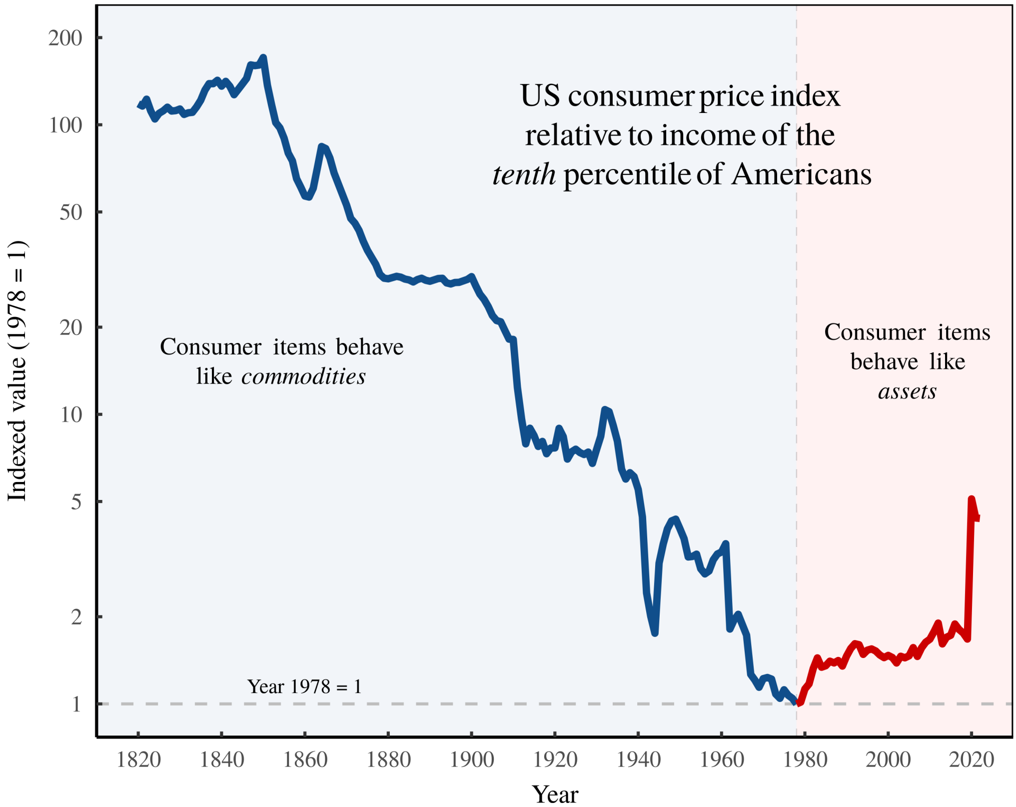

Figure 11: For the poorest Americans, consumer items are getting more expensive

This chart shows the US consumer price index, pegged to the income the tenth US percentile. Note that the vertical axis uses a log scale. Sources and methods

believing that the 'market' will do it automatically. Left alone, the 'market' will do what it always does: give the rich what they want, and take from the poor as much as they can suffer.

Speaking of taking from the poor, have you noticed how governments are responding to the housing crisis by penalizing the unhoused? It's as if the ruling class is clueless to the problem.

Well, in a real sense, they *are* clueless. You see, from the vantage point of the American top 1%, there *is no housing affordability crisis*. For these wealthy folks, house prices continue to decline relative to income, as Figure 13 shows. So while the ruling class basks in cheap housing, they ignore the affordability problem by outlawing its victims.

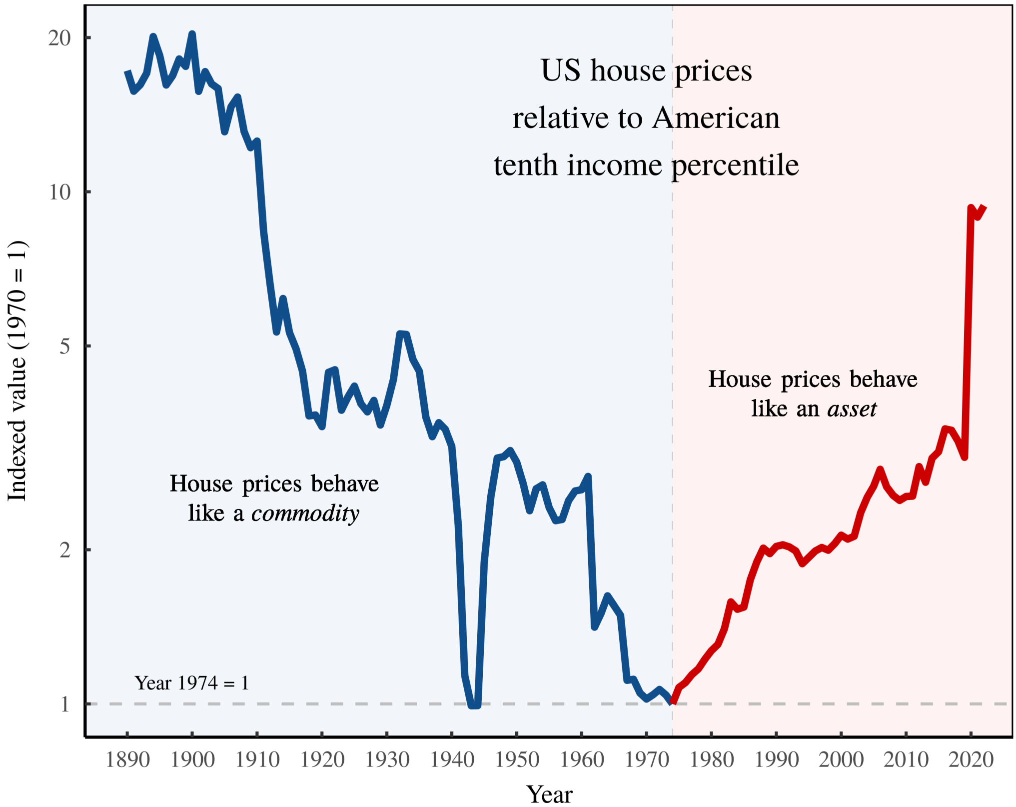

Figure 12: For the poorest Americans, the dream of home ownership is dying

The line shows the indexed price of US houses, pegged against the income of the tenth US percentile. Note that the vertical axis uses a log scale. Sources and methods

And that brings me to a reasonable solution to the housing crisis. While we can't build ever-cheaper homes (the world doesn't have enough energy), we can make homes more affordable to the working class.

To do it, we need to (paradoxically) ignore house prices and instead, focus on *income*. The root driver of the housing crisis is that poor people can't afford to buy houses or pay rent. And yet the rich have their cake and get to eat it too. Maybe ... just maybe ... if we took some of this money and gave it to the poor, then these folks could afford a place to live.

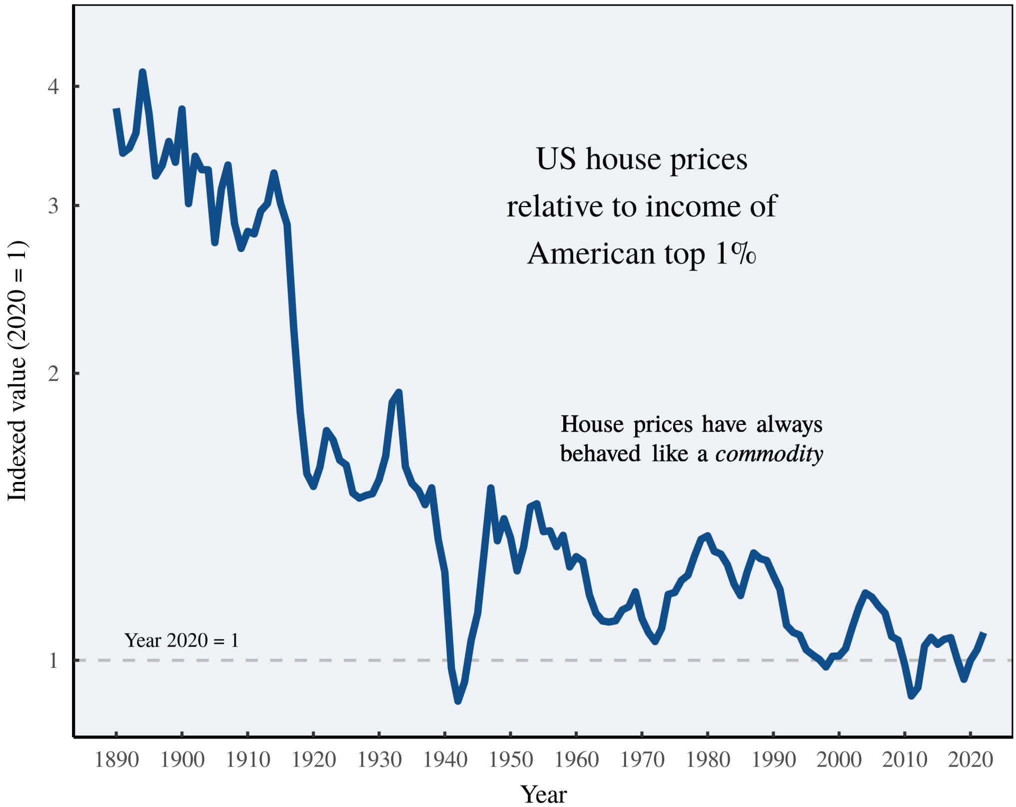

Figure 13: For the American top 1%, there is no housing affordability crisis

The blue line shows the indexed price of US houses, pegged against the income threshold to be part of the American top 1%. Note that the vertical axis uses a log scale. Sources and methods

---

Support this blog

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

Become an ETD supporter

The energy-debt model of US house prices

Here's the math behind the energy-debt model. It's based on the following multivariate regression:

$$\log\left(\frac{P_h}{Y_{pc}}\right) = a \cdot \log(E_{pw}) + b \cdot \log\left(\frac{D_h}{Y}\right) + c$$

Let's break down the math. On the left side, P_h/Y_{pc} is the indexed price of US houses, divided by US GDP per capita. On the right side, E_{pw} is the energy used per worker in US industry. D_h/Y is US household debt divided by US GDP. And a, b, and c are regression parameters, fit to empirical data.

When fed US data, the regression model has an R^2 value of 0.893 and an F-statistic of 525.

To test the importance of house supply, we can added housing starts to the regression model as follows:

$$\log\left(\frac{P_h}{Y_{pc}}\right) = a \cdot \log(E_{pw}) + b \cdot \log\left(\frac{D_h}{Y}\right) + c \cdot \log\left(\frac{S_h}{\Delta P}\right) + d$$

Here $S_h/\Delta P$ is the number of housing starts divided by the annual change in US population.

Tellingly, adding housing starts to the model barely improves the R^2 value (bumping it up from 0.893 to 0.897). But the addition worsens the F-statistic, reducing it from 525 to 360. The statistical message is clear: in the context of energy and debt, housing starts *don't* explain house prices.

Sources and methods

US population

- 1959 to 2023: FRED series `POPTHM`
- 1800 to 1958: Historical Statistics of the United States, series `Aa7` (US Resident population total)

US income

I use (nominal) GDP data to calculate US income. Data is from the following sources;

Data for nominal GDP:

- 1947 to 2023: FRED series `GDP`

- 1800 to 1946: Historical Statistics of the United States, series Ca10 (nominal GDP)

Data for nominal GDP per capita:

- 1947 to 2023: FRED series A939RC0Q052SBureau of Economic Analysis
- 1800 to 1946: Historical Statistics of the United States, series Ca10 (GDP) and series Aa7 (population)

US consumer price index

Data for the consumer price index is from:

- 1913 to 2023: Bureau of Labor Statistics series CUUR0000SA0 (consumer price index for all urban consumers)
- 1800 to 1912: Historical Statistics of the United States, series Cc1 (consumer price index, all items)

US housing starts

Data for US housing starts is from:

- 1959 to 2023: FRED series HOUST (new privately-owned housing units started: total units)
- 1889 to 1958: Historical Statistics of the United States, series Dc510 (non-farm housing starts, units started total)

US household debt

Data for US household debt is from:

- 1946 to 2023: FRED series CMDEBT (households and nonprofit organizations; debt securities and loans; liability, level)
- 1916 to 1945: Historical Statistics of the United States, series Cj879 (total private debt, individual and non-corporate)
- 1896 to 1915: Historical Statistics of the United States, series Cj254 (commercial bank assets; real-estate loans)

Energy *produced* per worker, US oil-and-gas sector

Data for US oil-and-gas production is from the following sources:

- 1949 to 2023: Energy Information Agency, table T01_02
- 1860 to 1948: Historical Statistics of the United States, table Db155-163

Data for US oil-and-gas employment is from:

- 1929 to 2023: Bureau of Economic Analysis, persons engaged in production, tables 6.8A-D
- 1860 to 1928: Historical Statistics of the United States, Colonial Times to 1970, table M5-6

Energy *consumed* per worker, US oil-and-gas sector

Data for the energy consumed by the oil-and-gas sector is from:

- 1919 to 2007: Table 6 in Guilford et al (2011), 'A New Long Term Assessment of Energy Return on Investment (EROI) for U.S. Oil and Gas Discovery and Production'
- 2012: My own estimate.
  - Data for fuel used and/or purchased by the oil-and-gas industry is from the US Census, table EC1221SM1. To convert fuel expenses into energy, I used price data from the Energy Information Agency, and fuel energy content from a variety of sources.
  - To estimate the amount of electricity used by the oil business, I used input-output tables from the Bureau of Economic Analysis and electricity price data from the EIA. I also used EIA data to convert electricity use into primary energy equivalent. (For details about this calculation, see the appendix here.)
  - To make my estimate consistent with the data from Guilford, I also used census data to estimate oil-and-gas sector energy consumption in 2007. (Sources are the same as above, but the 2007 census file is called EC0721SM11.) Then I used this 2007 estimate to index the 2012 estimate.
  - As it happens, I also used census data (EC1721MATFUEL) to estimate oil-and-gas energy consumption in 2017. But I'm not convinced the data is sound, as there are a number of missing entries.

Data for US oil-and-gas employment is from:

- 1929 to 2023: Bureau of Economic Analysis, persons engaged in production, tables `6.8A-D`
- 1860 to 1928: Historical Statistics of the United States, Colonial Times to 1970, table `M5-6`

Energy use per worker, US industry

Data for US industry energy consumption is from:

- 1949 to 2023: Energy Information Agency series `TEICBUS` (total energy consumed by the industrial sector)
- 1912 to 1948: I estimate industry energy use as follows; I start with data for total US energy use from Benjamin Warr's REXS database. Then I estimate the portion of energy used by industry by looking at industry's share of electricity use. This electricity-share data is from Historical Statistics of the United States, Colonial Times to 1970, series `S120`
- 1890 to 1911: I assume that industry energy use is proportional to total US energy consumption, derived from Energy Information Agency Appendix D1

Data for the number of workers in US industry is from:

- 1929 to 2023: persons engaged in production, from Bureau of Economic Analysis tables `6.8A-D`
- 1890 to 1928: Historical Statistics of the United Dates, series `Ba814-830` (labor force, by industry)
- Note: 'Industry' consists of the following sectors: Agriculture, Mining, Construction, and Manufacturing.

US house prices

Data for US house prices is from:

- 1963 to 2023: FRED series `ASPUS` (average sales price of houses sold for the United States)
- 1890 to 1962: data is from Robert Shiller (US National Case Shiller Home Price Data, nominal house price index)

S&P 500

Data for S&P 500 price is from Robert Shiller (US Stock Price, Earnings and Dividends as well as Interest Rates and Cyclically Adjusted Price Earnings Ratio (CAPE) since 1871)

Income of US 10th and 99th percentile

Data is from the World Inequality Database, series `sptincj992` (income shares) and series `tptincj992` (income thresholds). To make the estimate consistent with my other measures of income, I do the following:

1. Get the income of the 10th/99th percentile using WID income threshold data.
2. Get the average US income by taking the mean of WID income threshold data, weighted by the population share of each data point.
3. Calculate the income of the 10th/99th percentile relative to the mean income, and index this value to US GPD per capita.

Further reading

Keen, S. (2009). Household debt: The final stage in an artificially extended Ponzi bubble. *Australian Economic Review*, *42*(3).

MacBeth, H. (2018). *When the bubble bursts: Surviving the Canadian real estate crash*. Dundurn.

Suaste Cherizola, J. (2021). From commodities to assets: Capital as power and the ontology of finance. *Review of Capital as Power*, *2*(1), 1–29.