

Things Are Cheaper (If You Know Where to Look)

5 Idea Wednesday: deflation abounds; 'cheapest' model isn't the cheapest model; private credit's other problem child; China's beggar thy neighbor; Still no housing shortage (nor renovation boom)



MOSES STERNSTEIN

AUG 19, 2026  PAID

5 Idea Friday Wednesday

- there's deflation, just not in the CPI (and other reflections on low prices that may not be low forever)
- open weight models may be cheaper, but that doesn't mean they're cheaper
- private credit's other problem child
- China continues to Beggar It's Neighbors, but maybe less-so (possibly)
- still no housing shortage, nor renovation boom, neither (and other tales from the housing market, where liquidity tells the story)

👉👉👉 Reminder to sign up for the [Weekly Recap only](#), if daily emails is too much.
Find me on [twitter](#), for more fun.

✓ Subscribed

👋👋👋 Random Walk has been piloting some other initiatives and now would like to hear from broader universe of you:

(1) 🛎 [Schedule a time to chat with me](#). I want to know what would be valuable to you.


(2) 💡 Find out more about [Random Walk Idea Dinners](#). High-Signal Serendipity.

1. Certain Things Are Cheaper (For Now), If You Know Where to Look

The recent CPI print has put some of the teeth-gnashing about raising rates to bed (even if it really shouldn't matter, one way or another).

Random Walk continues to maintain that [if monetary policy is not the problem, then monetary policy is not the solution](#). In this case, if the AI buildout is driving up the costs of technological components (and it is), and the impulse is driven by the spending decisions of some of the largest, most profitable companies the world has ever known (which it is), then the obvious cure to high prices is high prices.

For what it's worth, Goldman Sachs agrees with me:



The combined price effects of tariffs, energy prices and tech hardware demand are expected to fade over the next year.

Putting that aside, I have this half-baked theory about why self-correcting inflationary pressure is still raising the hackles of at least some members of the Fed (other than the fact that [letting everyone knows you've got hackles is the counter-Walsh tantrum de jour](#)).

It occurs to me that perhaps part of the disconnect may arise from a feature of how we measure inflation.

You see, ordinarily, the reason that non-monetary (and non-fiscal) inflation is self-correcting is that higher prices suppress demand—either for the thing itself, or for some other ‘nice to have’ that gets squeezed out by the more expensive ‘need to have,’ (aka, ‘the cure to high prices is high prices.’) In the latter case, one component of the CPI may rise, but then another should fall. One inflation is offset by another deflation.

There's deflation, it's just not measured by the CPI

In this case, however, while that's indeed what is happening—one price rises, while another one falls—it turns out that the deflation is not a measured component of the CPI (or PPI).

That's because the thing that's deflating are asset-prices:

Exhibit 7: Premium for Top 5 stocks has almost disappeared


S&P 500 12m forward P/E multiple



Source: Datastream, Goldman Sachs Global Investment Research

Exhibit 10: Industrials now trade at higher valuations than Technology

MSCI World sector/style valuations, 12-month forward P/Es relative to the last 20 years



Source: Factset, Datastream, Goldman Sachs Global Investment Research

Tech multiples, especially for those companies driving up the cost of memory etc. with their considerable outlays, have compressed to the middle of their historical ranges.

In other words, all that capex is driving the price of memory up, and the price of hyperscaler equity down. Put yet another way, the high price of “need to have” AI infrastructure has soaked up all the free cash flow, such that it’s squeezed out the “nice to have” of returning money to shareholders (mostly via buybacks).

One price goes up, and another one goes down. Self-correction working just fine.

But, here’s the thing: the price of equity is generally not a component of measured inflation—the component that comes closest is asset management fees (as a % of AUM), but they’re a relatively small piece

of the puzzle. Now, there are plenty of good reasons why we don't think of asset prices as part of the CPI, but in this case, it remains the case that the feedback loop of demand elasticity is functioning, just not within the usual framework for how we measure it.


And I think that it may be obscuring an important part of the inflation picture. Over time, the rising cost of capital should itself soften the appetite for spending (so the theory goes), but probably not on any timeline that would make 2% benchmark-watchers satisfied.

Idk. It's a theory, and I'm not sure why it's wrong *per se* (even if it might be), so much as "that's just not how we measure inflation."

Discount on aisle hyperscaler

Anyways, there's another reason to bring this all up: it's once again [worth noticing that the largest, most-profitable companies the world has ever known are kinda cheap](#):


Current forward price/earnings multiple compared with distribution over the past decade



Past performance is not a guide to future performance and may not be repeated. The securities shown above are for illustrative purposes only and are not to be considered a recommendation to buy/sell. Blue bars show the interquartile range. 50% of data points lie within this range with 25% higher and 25% lower. The white lines across the blue bars highlight the median. Source: LSEG Datastream, Schroders. Based on month-end data 31 August 2016-31 July 2026. Please see relevant disclaimers on page 54


Other than Apple (which has largely opted out of the capex game), forward earnings multiples for 'big tech' are at or near the bottoms of their 10-year ranges.

On a monthly basis, the picture is largely the same, although Google is actually above its 10-year median and average:



Google [re-rated off the Deepseek sell-off towards the end of '25](#), and while it's come down from its peak, it's still somewhat elevated relative to the historical norm. But Meta, Microsoft, and Amazon share prices are definitely feeling the cashflow burn.

The reason, of course, that these companies are so discounted is the point above:




Source: FactSet, Goldman Sachs Global Investment Research

Capex is expected to exceed 100% of cash flow this year and the next. That's a lot of cash flow that's (a) not being returned to shareholders; and (b) flowing into the AI infra coffers . . . and so investors are putting their dollars elsewhere.

It's either short-termism, opportunity cost, or [fundamental skepticism](#) about [capex payoff](#), but either way, hyperscaler equity is trading for a relative song. (Another possibility is that investors are concerned that AI cloud is just a less-profitable business than non-AI cloud and therefore a re-rating is the new-normal.)

That's part of why this is perhaps one of the most consequential charts in equity pricing right now:




Source: FactSet, Goldman Sachs Global Investment Research

Hyperscaler free cash flow is expected to tank in 2027, and then come roaring back with a fury in '28-'29.

If the estimates are right (and that's a big if), then the current price of the hyperscalers is very wrong. For Random Walk's part, I'm still long hyperscaler ([for all the heads-I-win-tails-you-lose reasons laid out before](#)), but you do you.

ICYMI




Oil Shock or Nah (and the Machinations of Iran's Biggest Customer)

MOSES STERNSTEIN • JUN 5

[Read full story](#)

Retail on a Heater



MOSES STERNSTEIN • MAY 30

[Read full story](#)



Why Isn't the On-Ramp Working?

MOSES STERNSTEIN • MAY 29

[Read full story](#)





Chart Dive on Private Credit, PE/VC

MOSES STERNSTEIN • MAY 21

[Read full story](#)



That Escalated Quickly

MOSES STERNSTEIN • MAY 15

[Read full story](#)

2. Open Weight Open For Business (cont): When Cheapest Isn't Always Cheapest

When the Kimi K3 Open Weight model kerfuffle exploded a few weeks back, Random Walk was the first (and perhaps only) person to state the obvious: [if Chinese labs are going to freeload on lab capex to release distilled "same but cheaper" knockoffs, then it's inevitable that frontier model progress would slow.](#)

Lots of people were gleefully celebrating what they perceived to be the demise of lab margins (and/or were angrily protesting any lab efforts to protect those margins), without considering the fact that labs without margins could not keep pouring billions of dollars into training a new frontier, if the pirated copy would be released a few weeks later.

While I can't say for sure that that's what's happening, [there was some indication at the time that it was](#), and now there's this:



Greg Brockman  
@gdb



we temporarily slowed scaling of our frontier training, including our largest planned frontier RL, to strengthen security and monitoring. we believe confidence in safety will increasingly set the pace of AI development:



OpenAI  @OpenAI · 36m

As models become more capable, the risks associated with developing and testing them internally also grow.

We temporarily paused reinforcement learning (RL) training on our latest models intended for deployment for two weeks while we hardened and red-...

2:36 PM · Aug 18, 2026 · **15.4K** Views

OAI announces that its slowing frontier training . . . because of security and monitoring. I mean, sure, that's possible . . . or it could just be a euphemism for "we're not going to release a new frontier model until we figure out a better way to prevent distillation."

Anyways, that's not really the point, here. The point is really upstream of frontier progress, and it's the question of whether open models are actually "better."


The thing is that models are developing so quickly on so many different dimensions that pure apples-to-apples comparisons are difficult to impossible. Among other things, there are different modes to consider (e.g. text, code, video, voice, etc.), different capabilities (e.g. context window size, caching, latency, etc.), and different success rates, depending on both the task and the degree of difficulty.

Random Walk had promised a chartapalooza to make the point, but things have evolved so much in such a brief period of time, that it became kind of silly, so I'll offer just three anecdotal data instead.

The first comes from Pitchbook, which ran an experiment to illustrate precisely the point that the cheapest token does not necessarily mean the cheapest model. The reason is pretty straightforward: a cheaper model may cost less per token, but if it burns more tokens on the same task (or really, string of tasks to accomplish a goal), then it's functionally more expensive:

The priciest tier runs cheapest; the cheapest sticker runs dearest

Top AI model output cost per MTok



True cost per completed enterprise task (lower is better)

Sources: PitchBook, Claude, OpenAI, xAI, Google, and Meta • Geography: US • As of July 17, 2026

Mythos 5, the priciest sticker on the sheet at \$50, runs near cheapest, while the cheapest sticker, Muse Spark, runs dearest. Cost per completed enterprise task at a fixed quality: The near-top sticker finishes the work for less than half what the cheapest model on the sheet costs to run.

[Pitchbook](#)

According to pitchbook's experiment, Opus 4.8 and Gemini 3.1 Pro were the cheapest options, even though Opus 4.8 cost \$25/token and Gemini

just \$12/token.

The difference comes down to what Pitchbook calls the “reliability premium.”

That is, every time a model reads and re-reads (and writes and re-writes) context to generate an output, it burns tokens (and consumes memory). An agent is basically doing that over and over again—it starts with the token-consumptive “prefill,” (i.e. the core context behind the assigned task), writes new tokens (which get added to the context), evaluates the growing corpus of context, and writes again (and again) until the task is done.

The more mistakes it makes and/or the less efficient the caching is, the more token consumptive (and costly) the overall task becomes. The net result is that two models with very different \$/token prices can have roughly equivalent \$/tasks.

Doordash ran a similar experiment, in this case routing different levels of code review to different models:





Figure 4: Recall by severity.

[Doordash](#)

For Doordash, the “best” model varied by the complexity of the problem, where different models had different success rates, leading to different optimal model routing.

3Fourteen Research ran an experiment of their own, once again demonstrating that price/task is not the same as price/token:




Kimi K3 turned out to be ~30% more expensive than Opus 5 on a task basis.

You get the idea: just because open models (or any model) has a cheaper sticker price, does not mean it's in fact cheaper. There's a substantial amount of nuance involved, and it depends on the nature of

the task, the hardware, the software running in the background, and I strongly suspect, the skills of the agent-makers.

Don't get me wrong: the cost advantages of open models are real, but it doesn't obviously follow that they are as-good or cost-competitive in all cases. Not by a long shot.

In the big scheme of things, though, one thing is clear: computational demand continues to increase, likely driven in large part by “agentic” workflows:



Silicon Data

Spot prices for GPU rentals continue to climb across the board, even for the oldest GPUs.


The longer-term unit economics for GPUs are still very much a moving target, but the possibility that compute is “overbuilt” is becoming more and more remote.

3. Private Credit has a Software Healthcare Problem?

[Private Credit continues to not-apocalypse](#), and everyone seems to have mostly lost interest for now.

Things never really got that bad, and most of the badness more or less stabilized. It certainly could get worse, and it's definitely the case that the great software repricing is making life uncomfortable for some loans, but for the most part, it's fine.

The worst deals were from the ZIRP vintages, and those shoes have been dropping over the past year and a half:




The vast majority of recent foreclosures came from the peak-zirp vintages.

Again, [mistakes were definitely made](#), but no apocalypse has ensued. To be sure, there is still some white-knuckling left to go, as '28 and '29

are banner years for software term-loans coming due:

Exhibit 19: Software & Services Maturity Wall by Credit Quality



Data as of March 2026.
Sources: JPMorgan, Bloomberg


[Apollo](#)

~\$100B+ isn't a huge deal in the big scheme of things, but these were junk-rated to begin with, and likely were issued before software re-rated at the beginning of 2026.

So, yeah, there will be some painful conversations. But melt-down? Doesn't seem like it's in the cards.

One interesting aside, however, is that software is arguably not the biggest problem-child in private-credit land:


Software Defaults Are Slightly Higher



Random Walk has made this observation before, [but healthcare and life-sciences are pacing default rates](#), with a substantial uptick since the end of '25.

Healthcare also has some of the widest dispersion in spreads, of any category of BDC loans:

Figure 2: BDC Lending Spreads by Industry, 2026:Q1



Richmond Fed

Median spreads are pretty much the same everywhere (but Utilities), but the inter-quartile range for healthcare and social assistance is the largest (and much larger than software's).

In other words, there's far more turmoil and uncertainty running ripshit through healthcare than there is in software.

Why that is, I'm not entirely sure, but it could have something to do with the [fundamental unsustainability of providing "healthcare services."](#) With downward pressure on reimbursements, I do wonder if some business models are simply upside down. It's pure speculation, however.

4. China's Beggar Thy Neighbor Trade (cont.)


Random Walk periodically riffs on [China's efforts to be the maker of all things to everyone everywhere, while buying nothing from nobody.](#)

- It seems to reflect a strategic approach to identify what's selling, and then make that, better and/or cheaper than whomever was making it before (at least partially by dint of subsidies).
- Random Walk would posit that the impetus for the approach is demographics: China (and the OECD) is getting older, and perhaps smaller soon, and so if neither domestic markets, nor global markets are getting bigger, then share-capture is the logical next move. In other words, if we've decided to weather our democratic winter by selling our debt to subsidize our consumption, then China has decided to sell their stuff, by subsidizing their production.
- the observation matters for understanding the state of play more broadly, but also with respect to trade, and it reflects something of a challenge to trade econ 101, insofar as the theory says e.g. Germany is supposed to enjoy the consumer surplus, while its manufacturing base is brought to ruin, which may indeed prove net beneficial to Germany in the longer-run, but it's obvious to see why the claim is a harder sell.

Anyways, some charts to add to the theme.

First, yes, China is in fact selling more stuff everywhere, because it's own domestic market has no juice left to squeeze:

Exhibit 8: Chinese Export Volumes Surge, But Imports and Retail Sales Stagnate




Source: China General Administration of Customs, China National Bureau of Statistics, Haver Analytics, Goldman Sachs Global Investment Research

Chinese exports have surged past its imports, all the while Chinese retail sales continue to fall.

And yes, once China decided to swallow Germany's industry, that's exactly what it did:

Exhibit 10: China Is Undoubtedly Moving Up the Value Curve and Now Competes Directly With Germany

Germany: Investment Goods Trade Balance with China (€Millions, 12-Month Moving Average)





Data as at April 30, 2026. Source: Statistisches Bundesamt.

KKR

China went from importing from Germany to exporting to Germany in a heartbeat. Perhaps Chinese cars really are better, and perhaps this makes Germany better off. I suppose we shall see.

Bigger picture, none of this should come as any surprise, nor should the tepid growth of durable goods sales everywhere be a surprise:



China is getting older (and even more childless) very quickly.


It used to be a thing where 'experts' would periodically remark that 'Chinese consumers just need to consume more because they're under-consuming' or whatever. Maybe so, but plainly 'under-consuming' is not the main problem here for aggregate retail demand.

What China does do is make a lot of energy, **and a lot of robots:**

DISPLAY 2

China Is Well Ahead in Humanoid Patents

Monthly Patent Applications Mentioning "Humanoid" by Country



Source: Google Patents, Morgan Stanley Research. As of 3/31/2026. The data may not capture 100% of patents.


China has (ironically) ramped up its robot patent game a lot. Robots are a key ingredient when it comes to doing more with less, and it seems like something China prioritized well before the West. We're playing catch-up now, and while we'll get there, a lot of the components and technology come from China, so it will take some time.

Another thing that China has done has moved its manufacturing to other countries, like [Vietnam](#), but also, Mexico:

Mexican exports of server-related hardware have surged past exports of cars

Monthly exports (\$bn)

— Vehicles* — Hardware*



FT


Mexico got very good at manufacturing dishwashers, and more recently, has gotten very good at manufacturing the hardware critical to the AI infrastructure buildout.

Here's another look at how our neighbors have fared in relative import growth to the US:

Chart 4

Artificial intelligence-related demand drives Mexico's nominal import growth

Contributions to U.S. import growth, percent



NOTES: Imports reflect the total nominal value of merchandise cleared through Customs. Merchandise either entered consumption channels immediately or after withdrawal from bonded warehouses or Foreign Trade Zones under Customs and Border Protection custody. Import growth is measured year over year for December–February (2025/2024; 2026/2025). Corresponding North America Industry Classification Codes are 334111 (electronic computers), 211120 (crude petroleum) and 336110 (automobile/light-duty vehicle manufacturing).

SOURCES: Census Bureau, accessed via the U.S. International Trade Commission's DataWeb; authors' calculations.

Federal Reserve Bank of Dallas

[*FRB Dallas*](#)

Mexico has been doing great on the AI trade, while Canada, not so much.

In this case, however, China is probably not the prime mover. It's Taiwan (not China) that's been setting up the AI infra-shop in Mexico.

Why would Taiwan build in Mexico rather than the US of A? Presumably labor costs are a big part of the story, but I do not know that for sure. Claude seems to agree:

Fully loaded assembly labor in Mexico runs roughly \$5.70–7.80/hr depending on the region, against about \$31.60/hr in Texas. That gap would matter less if this were an automated process, but it isn't — an NVL72-class rack is thousands of hand-routed cables, connectors and manifolds, then burn-in and test. It's one of the most manual stages left in the whole chain, which is exactly the stage where a 5x wage differential dominates.

A 5x gap in labor costs is definitely a worthwhile consideration.

But, back to China for a moment, the great Michael Cembalest just published a big chart pack titled [The Year of the Trojan Fire Horse: China's imbalanced economy and unrelenting mercantilism](#).


It's interesting throughout, but the gist of it is that China's (subsidized) export machine is running roughshod over the rest of the world, including its "trading partners," of which there are basically none.

- "Overcapacity is essentially the hallmark of China's impact on the global economy"

Europe is not the only region struggling with Chinese excess production. As shown on the left, China runs massive trade surpluses with developing countries as well. When sourcing critical raw materials, chemicals and electronic component inputs, firms in developing countries are increasingly left with few options other than Chinese suppliers. Overcapacity is essentially the hallmark of China's impact on the global economy. As shown on the right, China's manufacturing *capacity* is often 2x or more than its current *production*, setting the stage for years of low-cost exports whether the rest of the world is prepared for it or not.


The third chart is also very telling: at the beginning of the 21st century, China's share of value added in global apparel, textiles, footwear and leather exports was ~35% of the low- and middle-income (LMI) country total. At the time, that figure was around the same as China's LMI population share. Today China's value added share is two thirds of the LMI total even as its population share has gone down.

China's manufactured goods trade surplus by region




Source: World Bank WITS, UN Comtrade, JPMAM, 2024

China manufacturing capacity vs production, 2024



Source: IEA, 2025

Apparel, textiles, footwear and leather exports from low- and middle-income countries, Percent




Source: OECD, Foreign Affairs, June 18, 2026

- “~30% of Chinese industrial companies are unprofitable zombies that survive due to a broad-based subsidy approach that dwarfs the rest of the world.”

While China strenuously objects⁶ to such descriptions of its mercantile approach, the story is pretty clear. Normally, a prolonged period of negative industrial producer prices would lead to competition and reduced capacity. Not in China, however; instead, **~30% of Chinese industrial companies are unprofitable zombies that survive due to a broad-based subsidy approach that dwarfs the rest of the world as shown below.**


To complete the picture, countries that experienced large increases in manufactured imports from China also experienced slowdowns in domestic manufacturing. From 2021 to 2024, as Chinese imports to the ASEAN region increased, every ASEAN country except for Brunei, Cambodia and Laos experienced a decline in their manufacturing share of GDP. Some consequences were outlined in the US-China Economic and Security Review Commission 2025 report to Congress: dozens of Indonesian textile and garment firms shuttered with 250,000 jobs lost in 2023 and 2024 with another 280,000 jobs at risk; and the closure of 4,000 small household appliance, furniture, electronics, garment, auto and steel factories in Thailand in 2023 and 2024.

China industrial producer prices & entities operating at a loss, y/y % change




Source: China National Bureau of Statistics, Haver, JPMAM, June 2026

Manufacturing value added by ASEAN countries
Percent of GDP




Source: World Bank, 2024

Industrial policy spending in key economies, 2022
Percent of GDP



Source: CSIS, 2022

Industrial subsidies, 2005 - 2024
Percent of annual firm revenue




Source: OECD MAGIC database, June 1, 2026

“China is pulling up the ladder behind it”, Foreign Affairs, June 18, 2026

Many poorer countries fear that China’s economic rise will not leave room for them to industrialize. Unlike the US and Europe, which in the past helped facilitate the industrialization of poorer countries (including China), Beijing is currently on track to have the opposite effect. China is not merely climbing the technological ladder; it is pulling up the ladder behind it. It dominates the new commanding heights of manufacturing (electric vehicles, solar panels, batteries, drones) but it has not vacated the older, labor-intensive sectors through which it and other rich and middle-income countries escaped poverty. In effect, China is trying to do what economic theory says no country should be able to do: retain comparative advantage in almost everything. [Shoumitro Chatterjee, Johns Hopkins and Arvind Subramanian, Peterson Institute]

Anyways, you get the idea. China can produce way more than the world can consume, much to the chagrin of anyone else aspiring to make much of anything. It appears to be a matter of policy, as well, as latent capacity is kept afloat by state largesse.

One interesting aside, is that there is some evidence that the subsidies may be winding down:



Chinese state investment in fixed assets is flirting with negative yoy growth (and private investment has been negative for about 2 years).

If subsidies really are winding down, that certainly could change things quite a bit. It would force prices to rise, and perhaps give non-Chinese manufacturers more of a fighting chance. Now, don't kid yourself too much, because China may well win with a level playing field, as well, but at least we'll know.

5. No, “restrictive permitting” did not cause a renovation boom

A few weeks back, the great Arnold Kling posted something [a little silly about home renovation spending](#).

Joseph Politano writes,

In 2025, Americans spent more money on home improvements and renovations than on building new single-family homes — a historical first, excepting a brief period during the Great Recession. Total renovation spending now exceeds \$430 billion per year, an increase of 36% in inflation-adjusted dollars over just the past decade.


...Restrictive permitting processes make replacing many old homes a nonstarter, so an entire cottage industry has erupted around renovating and flipping existing homes to meet the needs of new buyers.

Ah yes, the great zoning barrier is keeping people from spending on new houses, so they're renovating instead.

Nonsense. New construction spending has regressed because demand has softened, and higher interest rates have depressed purchasing power.

Like you, I have no recollection of YIMBY victories in 2020, nor the NIMBY retrenchment in 2023, that apparently drove construction spending up and then down, respectively:

Private Residential Construction Spending Index




Source: Construction Spending, Census Bureau, NAHB Analysis
Index Level 01/01/2000 = 100

I do however recall that the non-existent cost of borrowing caused a building surge, and then a dramatic increase in the cost of borrowing, caused an equally dramatic building flatline.

There [is no housing shortage and/or affordability crisis](#), despite what the discourse wants you to believe. Rates (i.e. the demand subsidy) are the story here—and population non-growth—which shouldn't be some deep contrarian insight, and yet the commitment to the 'housing shortage' narrative is indefatigable:

Why is housing so expensive?

Major housing papers, colored by ideology, sized by citations



Full methodology: <https://homeeconomics.us/the-great-debate>
Source: Citations: OpenAlex.




[Home Economics](https://homeeconomics.us)

My friend Aziz Sunderji made this neat chart showing the weight of attention that zoning receives as a driver of home values, relative to the actual prime-mover, demand.

The reality is that **all public builders are offering historically high concessions to move inventory, operating at GFC-level margins**, while sellers continue to outnumber buyers, and inventory builds. You can holler “shortage” all you want, but the data is very much not on your side.

Once again, take a gander at one of Random Walk's favorite series on home values:



[Opendoor](#)

Opendoor shows ~1.5+ of steady home value depreciation. Since they're in the business of actually selling homes (albeit somewhat regionally biased), it's data worth paying attention to.

In general, new homes are built and sold at the prices that people are able to pay, such that there are net-new people to pay them. For existing homes, [there is a bid-ask spread](#) driven by sticky-prices and un-motivated sellers (who offload their "[unrealized losses](#)" to their mortgage lenders).

And yes, when homeowners don't sell and/or trade-up, they may renovate a bit, but that too is effected by rates, and the data does not support a booming renovation industry:

Exhibit 18: Expected business activity indicator
Construction sector




Source: Houzz

Exhibit 19: Backlog business activity indicator
Construction sector



Source: Houzz

Exhibit 1: Sales-weighted vendor index 1-year and 2-year growth history



Vendors include: SWK, MAS, WHR, FBIN, MHK, SMG, SHW, and PPG

Source: Company data, Goldman Sachs Global Investment Research


Both “business activity” and the renovation backlog (as measured by Houzz) have been depressed since rates went up, and sales for renovation goods and materials tell the same story: higher rates are keeping renovation activity (often funded by HELs) on backburner.

Anyways, in terms of when things go “back to normal,” the answer is that with time, and rising incomes, the bid-ask spread will close (and it

is, ever so slowly):

Income Needed to Buy a Home Flattens, While Typical U.S. Income Rises

Junes only, 2012-2026




Source: Redfin analysis • Created with [Datawrapper](#)

Redfin

Home values will hold steady and/or decline moderately, [while some combination of higher incomes and lower interest rates](#) gradually closes the gap.

For now, existing homes are mostly just not selling. And, to repeat the point above, they won't sell, until lower rates and/or higher incomes move prices up and to the right, which is indeed, the historical pattern:



[First American](#)

Existing home sales tend to rise when values go up (or sellers otherwise become motivated, e.g. during a massive recession).

In other words, **low-liquidity is hiding the fact that prices would be lower, if sellers were actually motivated** to meet the bids where they are. Where there is liquidity, as with the new builds, prices are indeed lower.

It shows up elsewhere in the data, as well. If you look at the historical discount for forced foreclosure sales, the discounts in the Northeast are rather high—some of that is consistent with relative decrepitude of some foreclosures, but the other part of that are market prices held up by low-liquidity. A foreclosure sale, though, is a forced sale that doesn't get to pick its spots.

The high relative discount to the Northeast is still true, but look at where discounts are way *above* trend:

Florida and Texas which account for 21% of active foreclosures and are seeing historically wide REO discounts

REO discount by home price tier — average across 100 largest markets, Jun 2026




Regional REO discounts — historical median vs current

Region	Hist. median	Current	vs median	Pctile
Florida	19.4%	27.8%	+8.5 pts	97th
California	14.1%	22.3%	+8.2 pts	98th
Mountain West / Pacific NW	15.8%	21.9%	+6.1 pts	90th
Texas	19.1%	24.2%	+5.2 pts	95th
Southeast / South	25.0%	29.2%	+4.2 pts	80th
Midwest	25.9%	29.9%	+4.0 pts	81st
Northeast / Mid-Atlantic	23.8%	25.5%	+1.6 pts	60th
Rust Belt	31.7%	35.4%	+3.7 pts	64th


Florida + Texas combined account for 21% of active foreclosure inventory and 18% of recent REO sales nationally.

Source: ICE Home Price Index – June 2026

REO Discount by Market



REO Discounts by Market – June 2026



Source: ICE Home Price Index – June 2026

ICE

Northeastern foreclosure discounts are high, but well within their normal ranges . . . Florida and CA, on the other hand, are historically lower, but are both currently running very hot (at nearly ~97-98th percentile of the normal range).

It makes sense that more liquid markets would have less of a forced-seller discount (because the presence of a forced-seller is less of



anomaly). It also makes sense that currently, even those more liquid markets are facing steep discounts . . . because Sun Belt demand has peaked (as all the builders will be the first to tell you).

For the Northeast, it's closer to business as usual. Homes don't sell frequently, because people are aging in place, often without a mortgage, and they won't sell, unless they get a price that meets their fancy. Foreclosure sales, however, take what the market gives, and that's considerably less than the usual low-volume bid.


Previously, on Random Walk

Private Credit and Insurance, two peas in a pod (reprise), and a chart dump on default rates


MOSES STERNSTEIN • JULY 9, 2025



e) all takeover



f) 2014-16 takeovers only




five charts on the rise of private credit in life insurance

[Read full story](#)

3 odd observations about AI that are probably misleading somehow

MOSES STERNSTEIN • JULY 8, 2025



paid adoption is going down?

[Read full story](#)


Energy in 1776

MOSES STERNSTEIN • JULY 4, 2025

How has U.S. energy use changed since 1776:

U.S. energy consumption (1776–2024)

quadrillion British thermal units



Data source: U.S. Energy Information Administration, *Monthly Energy Review*


Data values: Primary Energy Consumption by Source and Estimated primary energy consumption in the United States, selected years, 1635–1945

It's July 4th, so Happy Birthday America, and we're going to keep it light and only semi-topical.

[Read full story](#)

Maybe the OBBB is actually fiscally responsible?

MOSES STERNSTEIN • JULY 3, 2025



frontloaded stimulus < backloaded cuts?

[Read full story](#)

LPs want their liquidity and they want it now

MOSES STERNSTEIN • JULY 1, 2025



an anecdote from the world of cap-intro

[Read full story](#)

*Random Walk is an idea company dedicated to the [discovery of idea alpha](#). **It's not investment advice, conflicts of interest abound, it's totally unreliable, and it reflects the thoughts and ideas of Random Walk, only.** Find differentiated data, perspectives and people, and keep your information mix lively. A foolish consistency is the hobgoblin of small minds. Fight the Great Idea Stagnation. Join Random Walk. Follow [me on twitter](#). Follow me on substack:*



Get more from Moses Sternstein in the Substack app

Available for iOS and Android

[Get the app](#)

