

Climbing The Wall Of Worry With Optimism

Taking stock of the big scheme of things, and reflecting on the Capex that stirs the drink



MOSES STERNSTEIN

SEP 25, 2026  PAID

- first, the 10,000 foot view
- hard to call it a bubble
- market says “on the one hand, but on the other hand,” and lands in the middle, which may be the most wrong
- Jevons It! and emerging signs of the positive sum game
- reports of GPU obsolescence are greatly exaggerated and pondering the catbird seat

👍👍👍 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.

Climbing the Wall of Worry with Optimism

Public company earnings for the first half of 2026 are basically in the bag, so let's take a gander at where things sit.

Not all things, of course, no one has time for that, but the big things that are top of mind.

The 10,000 Foot View

Mostly, all the things that Random Walk has expected have come to pass.

Consoomer

The Almighty Consumer continues to be fine, somewhat pressured by higher gas prices at the low-end, but mostly just fine ([status wounds](#), notwithstanding).


- [K-shape? We hardly knew ye](#). Sure, higher-end services are in the pits for everyone but the price inelastic wealthiest, but that's because service shrinkflation makes higher-end services less-appealing.
- The bigger issue is that the [total number of shoppers isn't growing all that much](#)—that's what [aging](#), [no baby-making](#), and [deciding to enforce immigration laws](#) will do—and that can make it challenging to grow, if you're already a fully-scaled consumer business.

- [E-comm, especially for the cheaper stuff](#), continues to be a growth vertical (and if you're a grocer, you're selling high margin ads via edelivery ads to impress the ladies street).

But yeah, ew consumer discretionary:

Consumer discretionary washout

The lion's share of stocks in the sector are below 50 and 200 day moving averages



Source: Bloomberg, @LJKawa

Sherwood

[Told you to avoid that stuff.](#)

Labor market

Consumers are fine because the employment picture continues to be [no hire, no fire, no quit](#)—a secularly tight labor market will have that effect, especially when the biggest buyer of labor (the [National Nursing Home](#)) pays no nevermind to prices or profit (aka sustainability).

- We have “good” job prints, and then “soft” job prints, and then “good” ones, and then “soft” ones . . . mostly [reflecting the lower breakeven that is the new normal](#).
- Wage pressure, such that we have it (and we do still have some, but not much), [comes from the bottom](#), more so than the top. Have higher equity prices and tax refunds helped? Probably a little.

Inflation

[Inflation has stayed a bit high \(as expected\)](#), and wiggled up recently.

- Despite the fact that (a) AI Capex and (b) the IRGC’s ongoing war with civilization being the prime impulses, [the Fed raised rates for no reason that could be described as well-considered](#).
- As expected, the [long-end of the curve stays high](#), having nothing to do with monetary policy either, because its fiscal policy driving that train . . . and the only likely effect of raising short-term rates is to put Uncle Sam further in the hole. But hey, at least the Fed Board [showed Warsh who’s boss](#).

Rates & The Pass-Through

As for the rest of the economy, the pass-through mechanism of higher rates has already transpired.

- consumers [stopped buying big, financed durable goods](#), like cars and houses. There were some auto-defaults, as people walked away from underwater equity, but that’s already happened.
- The most meaningful “pass through” of higher rates (to consumers) are reflected in [unrealized losses in home equity](#) that [mostly sit on bank balance sheets](#) in the form of long-duration, low-interest loans—those nearly brought down the banking system in 2023, [when the Fed rug-pulled the banks](#), but the [BTFP saved us all](#), and it’s been a non-issue ever since.

On the commercial side of things, the other meaningful pass-through has also mostly transpired: *companies became lean, mean, profit-generating machines.*

- Higher rates have *hurt the Finance Bros*, as the *PE backlog continues to pile-up*—wounded by the *double-whammy re-pricing* of higher rates, and then AI disruption. Some of those shoes have dropped, but there are probably still more to come, *as PE sellers capitulate to the new-normal*, and looming refis will put more pressure on assets—and not just buyout assets, but real estate assets, as well.
- For all the talk of *Private Credit chicanery*, it's all held up pretty much fine, with the most pressure coming from the competition: banks have benefited from the un-inversion of the yield curve and got decidedly back in the lending game.
- In the broader universe of alts, as predicted, *concentration* has been the persistent theme. The *biggest, most experienced managers, are raising a bigger share of a smaller pie*, deploying larger checks into a smaller universe of companies. It's true in VC, PE, Credit, and Real Estate—fine times for the big dawgs, and a struggle for the rest.

There has been no *“back to (ab)normal,”* and no broader collapse.


Demographics and the fruits born of decisions made long ago—both policy and culture—are doing their thing, and since the course was set ~75 years ago, there's no easy way to unwind all the decisions made along the way, let alone set a new course.

There is some hope, though. And that's why we're mostly in a uni(pro)cycle economy right now: *AI uber alles*, the *deus ex machina solution to all that ails us.*

OK, so that's where we are. It feels good to get that all recapped and summarized.

But now let's swerve from the macro to the more micro, and get to the point of today's passing missive: AI Capex is the straw that stirs the drink, and everyone (rightfully) keeps asking: how long can it last, and relatedly, will it pay off (and by how much)?


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](#)


It's hard to call it a “bubble”

Let's get one thing out of the way: while there's [certainly plenty of downside risk](#)—nothing ventured, nothing gained, after all—it seems [wrong to call this a bubble](#).

Skepticism abounds, and [it's most definitely reflected in prices](#):

Exhibit 10: Industrials now trade at higher valuations than Technology

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




Source: FactSet, Datastream, Goldman Sachs Global Investment Research

Tech is trading below its 20-year median, and near the bottom of its inter-quartile range.

That's right. In the midst of an historic capex boom [pouring trillions of dollars into computing infrastructure](#), the builders of all that compute and the direct recipients of all those trillions are relatively . . . cheap?


It's true. The price of earnings for semis is now cheaper than the rest of the index:



The discount on semis earnings—again, with a multi-trillion dollar build out in full swing—is now bigger than the longer-term average discount typical of the industry’s historically sleepy, cyclical nature.

Does this look sleepy to you, thought?

Exhibit 7: Earnings revisions have unusually been positive
Consensus estimates in USD. MSCI AC World EPS



Source: Datastream, Goldman Sachs Global Investment Research


Earnings revisions keep bending upwards *off their already high base*. It's pretty darn unusual.

The most straightforward explanation for why that's the case is, as per above, skepticism abounds.

Investors are concerned that we're at or near the cycle's peak, i.e. *"it's pretty darn unusual' is the problem, dummy!"* They simply do not believe the forward earnings forecasts. They apparently do not believe the hyperscaler capex commitments will come to pass (even if they are [still pricing the hyperscalers as though their money will be spent](#)).


Capex from the major AI hyperscalers*

USD billions; Alphabet, Amazon, Meta, Microsoft, Oracle




Hyperscalers' cash flow and capex

% of sales



The AI infrastructure buildout is expected to continue for years to come

Hyperscaler capital expenditures (actual and consensus)



Notes: Hyperscalers include Alphabet, Amazon, Meta, Microsoft, and Oracle. Forecasts for 2026 onward reflect Bloomberg consensus estimates.

Sources: Vanguard calculations, based on data from Bloomberg, as of July 31, 2026.


No one has given any indication that they intend to taper their investments in chips any time soon. Not even close.

And yet, some of the biggest chipmakers are trading as though back-to-normal for earnings will arrive any moment now.

I mean, look at estimates for semi margins:

Exhibit 12: Semiconductor gross margins have expanded dramatically during the last few years

reflects current S&P 500 constituents



Source: FactSet, Goldman Sachs Global Investment Research

Semi margins, and especially memory, have ripped higher as demand has far out-paced supply . . . but then going forward, margins are forecast to just stand still. Not fatter, not thinner—but just stay as they are.


To Random Walk, that's just a symptom of the uncertainty that abounds.

The one hand

On the one hand, there's the trillions of dollars being spent and a persistent supply backlog that seems to keep pricing power in favor of the semi-makers.

Figure 1: HBM S/D ratio (CoWoS adjusted)

1GB equiv. Mn, % - RHS



Source: Company data, J.P. Morgan estimates.

JPM's estimates for high bandwidth memory put demand ahead of supply comfortably through 2028.

There is also the hundreds of billions of dollars being spent on all the compute that's already been bought and installed (with [a parabolic revenue backlog](#)), which provides the impetus for that ongoing trillion dollar outlay.

The other hand

On the other hand, as seasoned cyclical investors know, P/E is always lowest before the cycle's peak.

- Every shortage becomes a glut. Efficiency gains and supply competition must be around the corner.
- The cycle is already at historic levels, and how much bigger can it possibly get?

- Plus, both compute and AI demand are highly concentrated, and if those big dawgs were ever to dial things back—[perhaps because their own unit economics didn't pencil](#)—the whole thing could collapse like a house of chips.
- And where are the miracle AI productivity gains? [If this thing is world-changing, then where are the changes](#) (other than the trillions of dollars trading hands)?!

[People have doubts, and rightfully so.](#)

To return to the point about forecasts for memory margins, it's almost as though the one thing that *shouldn't* be possible is that margins flatten out. Either they get fatter or they get thinner, but “just right” seems like the least plausible outcome. (Of course, that's not really true, either, because if capex stops growing, and merely holds at the expected ~\$1T+/year for the next few years, then memory pricing power could also conceivably just hold for the next few years, but I digress.)

So, yeah, “unfettered exuberance” is hardly the setup.

There are clearly pretty divergent views of how this all plays out, with some very bullish commitments in one corner, very bearish positioning in another, and everything in-between, and it all adds up to an aggregate that sits somewhere, and perhaps most implausibly, in the middle.

Jevons It and the Positive Sum Game

Where does Random Walk land?

Well, if you follow all the backlinks above, Random Walk is [no stranger to the trials and travails](#) that may be looming . . . but [I've increasingly warmed to the upside case](#).

I think we're in the still early innings of a massive positive sum game with a demand frontier (and profit pool) that's going to substantially push the boundaries of what's come before.


To oversimplify, [it's a Jevons It! story](#). AI is getting exponentially cheaper, and rather than simply serve extant demand at a lower cost, it's [pushing the demand curve wayyy up and to the right](#).

Intelligence is getting much cheaper, very quickly

Consider this, and note the logarithmic scale:

The price of artificial thought may have fallen faster than for any other transformative technology in history

Relative cost (cheaper)



All prices adjusted for inflation. Compute series stops merely for lack of newer data. Sources: Electricity: Hersh (1999), Figure 2.1, and EEI (1995), Table 63. Compute: Nordhaus (2007), Appendix Table 2. Batteries: Fouquet (2026) via Our World in Data. DNA sequencing: National Human Genome Research Institute. AI line is 47%/quarter (12.7x/year), per estimates reported here for 2023–26; the extrapolation of that rate back to 2021 is broadly validated by Appenzeller (2024)'s finding that per-token prices for models clearing a certain capability threshold fell 10x/year in 2021–24.


EPOCH AI | CC-BY

epoch.ai

Through a combination of technological efficiency, pricing wars, and model piracy, the “cost of intelligence” is collapsing.


Again, it's not just the models making things cheaper. Data centers are getting more expensive, but unit costs are falling dramatically:

Exhibit 33: Overall Data Center Costs are Rising, While Costs on a Unit Basis are Declining




Source: Morgan Stanley Research

Exhibit 27: Rate of Change Accelerates as Throughput Scales



Source: Morgan Stanley Research

Exhibit 28: Rate of Change Accelerates as Throughput Scales



Source: Morgan Stanley Research

Whichever way you measure it, whether cost per cycle, tokens per power, or tokens per chip, the cost of running models is falling rapidly.


... Leading to even more demand

And again, as AI has gotten cheaper, the most proximate effect has been even more spending:

Top AI Models

Volumes Shares

Total token spend by businesses, indexed ①



Source: Ramp AI Index, data from Ramp AI Token Spend Management. Includes model-attributed API spend from connected AI providers. [Get the data.](#)

ramp ↗



Ara Kharazian ✓
@arakharazian



AI spend September 14-20 versus September 7-13:

Token spend: +19.6% WoW

Token volume: +19.5% WoW

Price per million tokens: +0.1% WoW

Notably the decline in token spend we observed in August reversed last week, but token spend is still below its July peak.


12:28 PM · Sep 25, 2026 · 90 Views

According to Ramp's data, which is directionally useful, but has some sample bias, August reflected a huge step-up in token spending, which tapered a bit in September, and then snapped nearly all the way back.

And again, all of this AI spending is happening while the frontier labs face increasing competition from [open-weight alternatives](#):

Share of weekly spend and tokens for Anthropic+OpenAI+Google models on OpenRouter
Data from Sept 11, 2024 through Sept 10, 2026

Combined share of Anthropic + OpenAI + Google models in dollars spent and tokens



openrouter.ai/data

Per OpenRouter (which also captures only a slice of the picture), frontier lab token share is only 30% . . . but spend share is near 70%.


In other words, it seems like users [are getting more sophisticated at spend optimization, as well](#). They're using cheaper models for the cheaper problems, and using more expensive models for more expensive problems (because [cheapest isn't always the cheapest](#)) . . . but the overall pie keeps getting bigger.

Rather than pure negative sum substitution, we're seeing positive sum value accruing up and down the value chain.

And perhaps the [biggest price-driven unlock for all this demand has been agents](#), where cheaper cached tokens have made the unit economics of agent-loops much more viable (and memory that much more important). And agents are barely off the ground . . . [probably less than 1% of users](#) are actually using agents with any meaningful scale, and they're already pushing the limits of what existing compute can reasonably serve.

A win-win-win for the hyperscalers (reprise)

That cheaper intelligence only increases demand shows up in other data, as well—with some [pretty straightforward implications for the potential catbird seat](#) in this ecosystem may sit:




Silicon Data has consistently shown that falling price/token has coincided with rising prices for compute, particularly the most advanced models.

In other words, once again, cheaper intelligence = more demand.


GPUs were supposed to be obsolete a few years after their install, but the opposite is happening: if anything, GPUs are appreciating as rental prices *rise* over time:

GPU Compute Prices Sep 2026

Year to date · Nvidia Accelerator One-Year Contract Rate




Last 30 days · Nvidia Accelerator One-Year Contract Rate



Even older models of GPUs are not only keeping pricing stable—because you'd actually expect them to get at least somewhat cheaper—they've actually increased rates since the year began.

These are merely estimates based on purchase price and rental income, again via Silicon Data, but this looks nothing like 3-4 year depreciation to obsolescence:




The oldest A100s are still holding ~25% of their original value through year 6, and the curve isn't overly steep, so there may be another ~2 years left. That's not bad, at all.

The point being is that the impetus for all this capex [is real, live paying demand](#), that appears to be accelerating.

The ultimate question, is of course, the return on capital, and for the compute layer, the prospects are promising:

We Present 3 Frameworks Detailing GenAI ROIC of 25-50%

We estimate ROIC across 3 different monetization vectors including: **1) Hyperscaler renting out GPU/compute capacity**, **2) Model builders serving their API through their own, internal infrastructure** and **3) Model builders serving their API through 3P infrastructure**




Source: Company data, Morgan Stanley Research Estimates

23

Morgan Stanley

Depending on where you sit—from being a pure rental, to serving your own models on your own stack, to serving your own models on someone else's stack—estimates for ROICs run from ~25% to 46%, and that dog will definitely hunt.


Fully mature cloud ROIC ran in the high ~30s:



Hyperscalers include AMZN, GOOGL, META, MSFT, ORCL. ROIC = (EBIT less 1 debt less cash).

Obviously ROIC has come down as the Capex buildout is in the steep part of the j-curve, but all indication is that it's not going to stay that way for long.

Behold, another look, courtesy of Warren Pies, showing the relationship between rental rates and cloud revenues:



Imagine that. More demand for compute leads to more expensive rates leads to [more revenue for the firms who own all that compute.](#)

Idk, man. Obviously do your own homework, and there are plenty of ways that this can go sideways, but there's also plenty of reason to think that this cycle is not like all the others.


"AI Adoption" isn't anywhere near mature, even with the model capabilities that we have. [Jevons Paradox appears to be in full effect](#), as cheaper intelligence pushes the frontier of still nascent demand higher and higher. And demand is nonetheless already in excess of supply such that prices are rising for compute.

I [hate to repeat those awful words](#), but just looking at the pieces where they lie, it increasingly feels like I have no choice, so here it goes: it's so early . . . and people don't seem to know what's coming.

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



five charts on the rise of private credit in life insurance

[Read full story](#)

3 odd observations about AI that are probably misleading somehow




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?




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](#)

