J.P.Morgan

EYE ON THE MARKET | OUTLOOK 2026

Smothering Heights

The market cap of four hyperscalers and the semiconductor ecosystem companies they rely on has grown from $3 trillion to $18 trillion in just a few years, and a broader group of 42 AI related companies has generated 65%-75% of S&P 500 earnings, profits and capital spending since ChatGPT's launch in November 2022. Is this moat really indestructible? In this year's Outlook we focus on four medium term risks: US power generation constraints, China scaling the moat on its own, Taiwan and a "metaverse moment" for hyperscaler profits after $1.3 trillion of capex and R&D. We conclude with a few pages on everything else, including a history of populism for investors.

By **Michael Cembalest** | Chairman of Market and Investment Strategy for J.P. Morgan Asset & Wealth Management

MARY CALLAHAN ERDOES
Chief Executive Officer
J.P. Morgan Asset & Wealth Management

As we begin a new year, I am delighted to share with you our much-anticipated 2026 Eye on the Market Outlook, written by my longtime investment partner, Michael Cembalest. This year's edition, "Smothering Heights," explores the extraordinary impact of generative AI on global markets, the unprecedented concentration of returns and the critical questions that will shape the future of investing.

Michael's research dives deep into the forces driving today's market—spectacular technological progress, record capital spending and the risks that come with rapid innovation. As always, Michael brings clarity and candor to complex topics, helping us look beyond the headlines to understand both the opportunities and the risks. I encourage you to pay special attention to the cover and the key themes highlighted throughout the report, as they will be central to investment decisions in the year to come.

On behalf of all of us at J.P. Morgan, thank you for your continued trust and partnership. We remain committed to delivering thoughtful insights and working tirelessly to help you achieve your goals.

Wishing you a healthy, peaceful and prosperous new year.

Smothering Heights: is the largest moat in market history indestructible?

After a rally comprised in equal parts of technological progress, a surge in tech capital spending and frenzied speculation[1], we've arrived at the following destination: 65%-75% of S&P 500 returns, profits and capital spending since the launch of ChatGPT in 2022 have been derived from 42 companies linked to generative AI. In other words, **the generative AI investment theme has smothered the rest of the US equity market**. Without the benefit of these 42 AI stocks, the S&P 500 would have underperformed Europe, Japan and China. To reinforce the point: tech sector capital spending contributed 40%-45% of US GDP growth over the last three quarters, up from less than 5% in the first three quarters of 2023.

Returns, earnings and capex/R&D of AI-related stocks in the S&P 500 since ChatGPT launch in Q4 2022

	Direct AI 28 stocks	AI Utilities 8 stocks	AI Cap Equip 6 stocks	Total AI 42 stocks	S&P 500 ex-AI	MSCI Europe	MSCI Japan	MSCI China
Since November 2022								
Price return	195%	66%	174%	190%	26%	33%	59%	50%
Earnings growth	159%	64%	155%	153%	19%	4%	52%	15%
Capex / R&D growth	72%	13%	20%	68%	19%	28%	37%	27%
Share of changes since November 2022								
Price return	76%	0.8%	1.3%	78%	22%			
Earnings growth	63%	1.6%	1.5%	66%	34%			
Capex / R&D growth	70%	1.0%	0.2%	71%	29%			

Source: Bloomberg, JPMAM, December 22, 2025

One lesson I've learned on investing: after a correction, ask "what could go right" rather than obsessing over factors that led to the selloff. And when markets are highly concentrated and near all-time highs, the right question to ask is "what could go wrong".

In this year's Outlook, we focus on four major moat risks: US power generation, China's ability to scale the technology moat on its own, China's approach to Taiwan and the ultimate profits earned on $1.3 trillion in hyperscaler capital spending and R&D since 2022. This year looks to be another version of 2025: a 10%-15% correction at some point due to profit-taking and a growth scare, but then equity markets end the year higher than where they began. Even so, it's the right time to start focusing on these questions. We conclude with a few pages on everything else, including a history of populism for investors.

Michael Cembalest, JP Morgan Asset Management

Tech capital spending in 2025 vs spending on major US infrastructure projects
Peak annual project percent of GDP

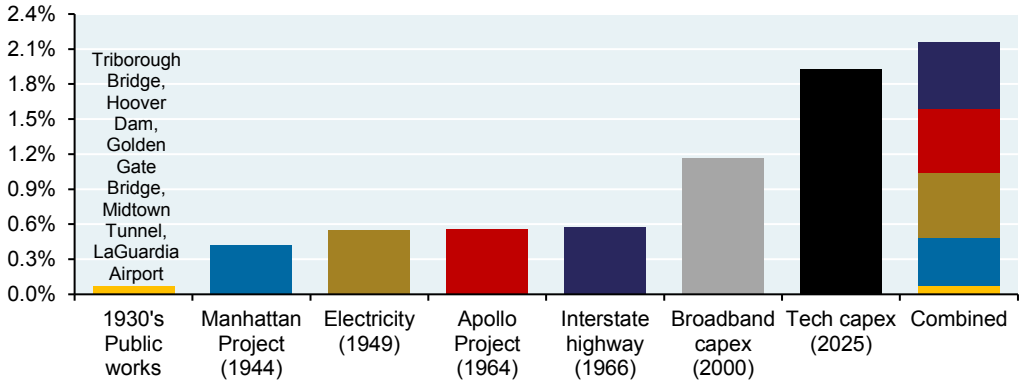

Source: Manhattan District History, BEA, Planetary Society, Eno Center for Transportation, San Francisco Fed, Hoover archives, Baruch, GoldenGate.org, New York Times, JPMAM, 2025

[1] An equal weighted basket of bitcoin-sensitive stocks, quantum computing, nuclear SMRs and unprofitable tech was up +120% by mid-October and still ended the year at +70% after a fall correction

Smothering Heights: is the largest moat in market history indestructible?

The foundation of the moat on the cover is comprised of NVIDIA chip designs, TSMC chip manufacturing and ASML lithography machines[2]. The market cap of just four semiconductor companies and four hyperscalers has grown from $3 trillion seven years ago to $18 trillion today. These stocks make up ~20% of developed world equity markets and ~16% of global equity markets (including emerging markets).

8 largest Hyperscaler and Semiscaler market caps
US$, trillions

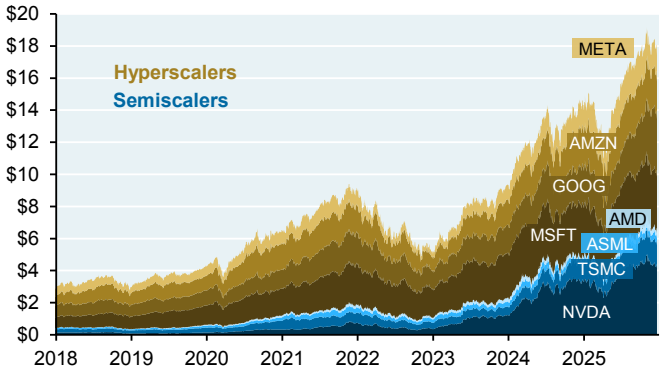

Source: Bloomberg, JPMAM, December 21, 2025

8 largest Hyperscaler and Semiscaler market caps
Percent of global market cap

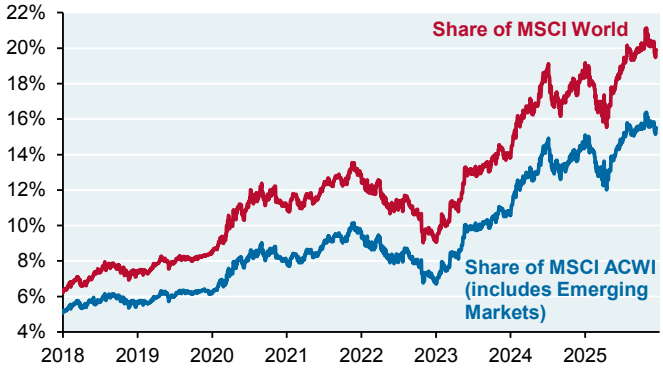

Source: Bloomberg, JPMAM, December 21, 2025

The revolution will be imported. The US semiconductor ecosystem has so far been spared many of Trump's new tariffs, allowing for a lot of imports. As shown on the right, moat-related industries benefit from larger tariff exemptions, defined as the share of US imports currently exempt from reciprocal, fentanyl and product-specific Section 232 tariffs. A Section 232 investigation on semiconductors is ongoing but may provide substantial exemptions to companies that make commitments to produce in the US.

US semiconductor, large computer and computer parts imports, US$, billions, monthly annualized

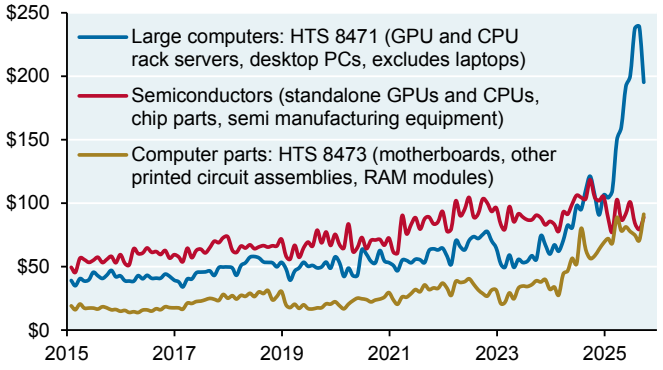

Source: USITC, JPMAM, September 2025

Trump tariff exclusions by product
Percent of 2024 US product imports excluded from Trump tariffs

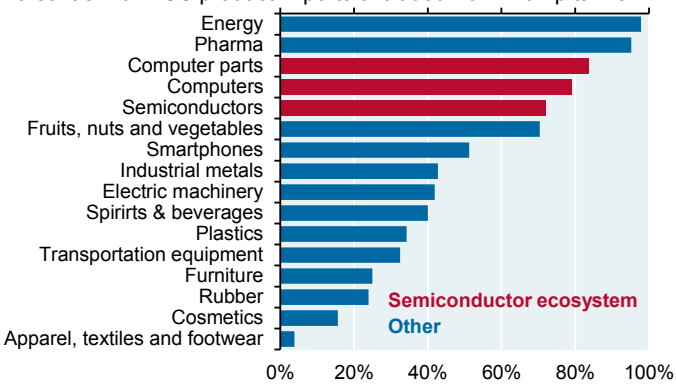

Source: USITC, White House, JPMAM, 2025

Our S&P 500 AI universe is composed of 42 stocks from three categories. **Direct AI (28)**: NVIDIA, Microsoft, Apple, Alphabet, Amazon, Meta, Broadcom, Tesla, Oracle, Palantir, AMD, Salesforce, IBM, Uber, ServiceNow, Qualcomm, Arista, Adobe, Micron, Palo Alto, Intel, Crowdstrike, Cadence Design, Dell, NXP, Fortinet, HP and Super Micro Computer. **AI utilities (8)**: NRG, Vistra, NextEra, Southern, Constellation, Public Service Enterprise, Entergy, NiSource. **AI capital equipment (6)**: Eaton, Trane, Johnson Controls, Quanta, GE Vernova, EMCOR

[2] ASML extreme ultraviolet light (EUV) machinery is essential to production of artificial intelligence technology. In machines the size of large buses, ASML lasers fire 13.5 nanometer plasma pulses at the rate of thousands per second at microscopic pieces of ultra-pure tin. Tiny mirrors capture the refracted light and etch features smaller than 10 nanometers onto thinly sliced wafers of polysilicon. With a market cap over $300 billion, ASML is the largest technology company in Europe

The four hyperscalers are spending increasingly large shares of revenue on capital spending and R&D. Meta's capex and R&D share of revenue is hitting new highs at 70%, and the company mentioned on its Q3 earnings call that it would "aggressively ramp up spending to stay competitive in the AI arms race"[3]. For context, the ratio of capex and R&D to revenues for the median S&P 500 company is 10%. But unlike capital spending booms of the past (casinos, airlines, fiber optic cable, gas turbines), the latest one has been mostly financed via internally generated cash flow…until very recently.

- Our universe of 28 direct AI stocks represents 50% of S&P 500 market cap and just 5% of S&P 500 net debt. The second chart shows a proxy for the share of capex and dividends for tech/communications companies that are financed via debt as opposed to cash flow from operations. **Despite the recent rise in capex, the amount financed by debt was still very low as of Q3 2025 in contrast to the late 1990's**
- Even after accounting for the Q4 surge in AI-related debt, many AI companies have net debt to cash flow ratios less than zero since cash and securities exceed total debt (third chart). That said, credit markets are picking up on the shift from cash flow to debt financing as described on the next page. Since we highlighted Oracle's challenges in September, its stock is down ~35% and its credit spreads have risen by 90 bps
- Free cash flow to revenue for most direct AI companies exceeds the S&P 500 by a wide margin, although as shown on page 15, hyperscaler cash balances are falling as a share of assets

Hyperscaler capex and R&D as a share of revenues
Percent

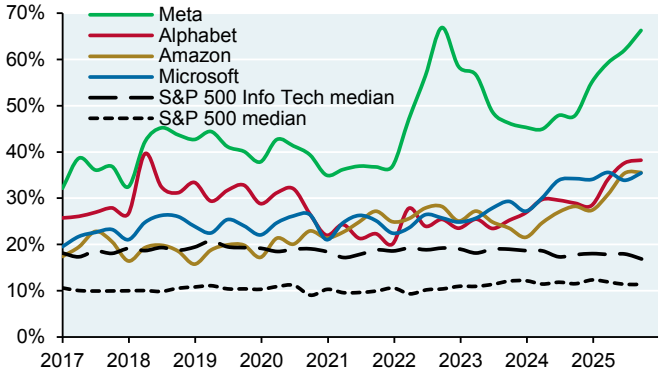

Source: Bloomberg, JPMAM, Q3 2025

Capex financing vs capex cycle
Russell 3000 universe: US tech and communications companies

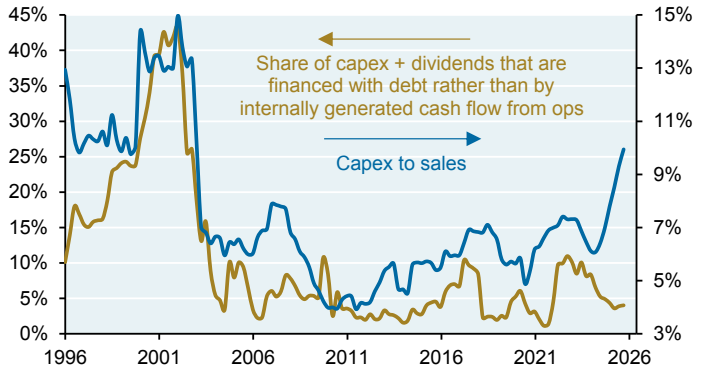

Source: Bloomberg, JPMAM, September 2025

Net debt to EBITDA ratios of Direct AI stocks
Multiple; includes bonds, loans and SPV triple net leases

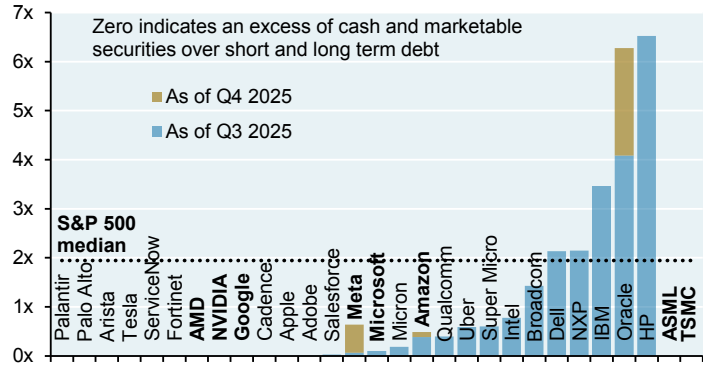

Source: Bloomberg, JPMAM, December 17, 2025

Free cash flow to revenue ratios of Direct AI stocks

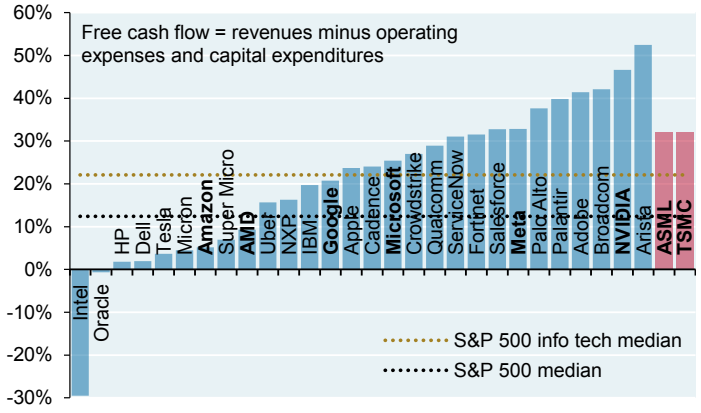

Source: Bloomberg, JPMAM, December 27, 2025

[3] **It would be great if Meta spent some of this cash flow hunting down and delisting fraudulent WhatsApp stock scams that have my name attached to them**. According to Senators Hawley and Blumenthal, Meta earns up to $16 billion per year hosting advertising for scams and banned goods on its platform and shows users 15 billion higher risk ads every day that show clear signs of being fraudulent (illicit gambling, payment scams, crypto scams, AI deepfake sex services and fake offers of federal benefits). "Meta's central, facilitating role in scams against consumers is unprecedented: by its own employees' assessment, Meta was involved in one-third of all successful scams in the US and was unmatched by other Big Tech platforms". Source: Blumenthal/Hawley Press release November 24, 2025

Two debt outliers: Oracle and Meta start the process of financing data centers via debt markets

Oracle. OpenAI committed to pay Oracle $60 bn per year, an amount OpenAI doesn't earn yet, to provide cloud computing facilities that Oracle hasn't built yet, and which will require 4.5 GW of power (the equivalent of 2.25 Hoover Dams or four nuclear plants). Since Oracle free cash flow ratios pale in comparison to the hyperscalers, the company started borrowing heavily in loan, bond and project finance markets last year.

Meta partnered with Blue Owl to build the Hyperion data center complex in Louisiana. Hyperion is financed by a joint venture SPV owned 20% by Meta and 80% by Blue Owl; the Blue Owl entity borrowed $27 bn of investment grade debt to finance construction, debt which apparently won't affect Meta's credit rating. Let's get into it:

- Meta entered into 4-year renewable operating lease agreements which begin upon Hyperion's completion in 2029. Meta provided the SPV with a residual value guarantee (RVG) that declines over time, that is always greater than outstanding debt and that declines to zero after 20 years. From the perspective of debt investors, Meta's periodic payment and RVG obligation effectively function as a triple net lease in most scenarios*
- S&P will not currently consolidate this obligation for purposes of analyzing Meta credit risk; it will instead add the present value of an assumed eight years of lease payments onto Meta's adjusted debt only in 2029 when lease payments commence (assuming for some reason that Meta walks away after 8 years). If this adjustment were made today, S&P states that Meta's leverage would only increase by 0.2x, well below S&P's threshold for downgrading Meta of 1.0x. As a result, S&P affirmed Meta's AA- rating
- S&P doesn't include the RVG in its leverage calculation since it's out of the money; in other words, the RVG is contingent and only applies if (a) Meta terminates the lease and (b) after Blue Owl sells the project, there aren't enough proceeds to pay SPV bondholders. Since S&P believes the RVG will decline faster than the value of the project, they do not assume Meta would ever have to make the RVG payment even if it terminated the lease (although they will evaluate this periodically). Et Voila! Off balance sheet…
- Let's focus on real economics rather than financial engineering and consolidate the SPV debt**. Meta's net debt to EBITDA ratio was negative starting the year; it rose to 6.6% by Q3; it rose further to 37% when including its October 2025 $30 bn bond issue; and rises to 63% when consolidating its Hyperion obligation. Still well below the S&P 500 median of 2x, but a very different picture than when the year began

Annual change in hyperscaler long term debt (bonds, loans and leases), US$, billions

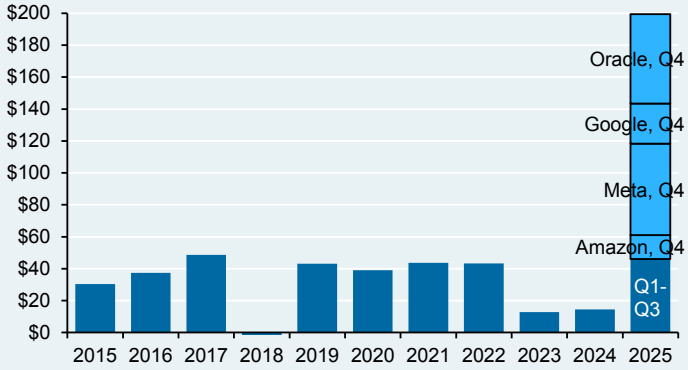

Source: Bloomberg, Company sources, JPMAM, 2025

Credit default swap comparison, select AI stocks
Spread, five year CDS, bps

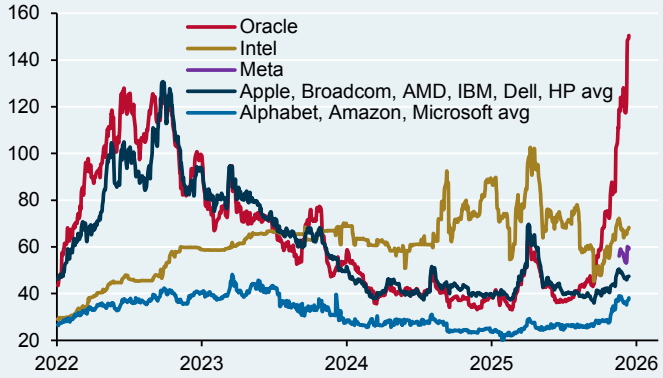

Source: Bloomberg, JPMAM, December 16, 2025

* While most permutations of the RVG involve a sale of the data center should Meta terminate the lease, if the SPV Board does not approve a sale, a third-party appraiser would value the project to determine Meta's RVG obligation. Should this appraisal exceed the project's true value, bondholders would be exposed to potential loss. Finally, **bondholders do not hold a first mortgage**; as a result, the SPV can incur more debt that is senior to debt already raised as long as incremental pari-passu or senior debt would not cause a ratings downgrade

** The same financing approach has been used in other infrastructure projects by Intel, EQT, Rogers, Sempra and KDP. **While these structures take advantage of accounting rules ASC 810/842, they're de facto triple net leases and investors should look past the "creative" accounting/rating agency treatment**. Meta is borrowing $27 bn, plus an extra $3.5 bn in present value terms for the presumed benefit of not having to consolidate (i.e., a 1% higher rate on project finance debt vs straight debt). Whether these structures would survive substantive consolidation tests in bankruptcy court is unknown at this time

Valuations: high but not as high as you might think

Many analysts see valuations as the primary risk to the current AI cycle. But as shown below, there's an internal consistency in the way US equity markets are valued: the higher the profitability (x-axis), the higher the valuation markets ascribe (y-axis). This is the case across sectors as shown on the left, and also when adding in the individual stocks on the right. **You might think the *slope* of the curve is too steep, and maybe it is; but just looking at higher valuations in isolation misses the extraordinarily high profit margins of the tech sector vs the rest of the market.** Another analysis that gets to the same place: the PEG ratio (P/E multiples divided by projected earnings growth) for the MSCI World Technology Index is not that different than the same ratio for the MSCI World Equity market as a whole.

Last figure to keep in mind: in 2025 the forward P/E of the S&P 500 only rose by ~3%; the rest of the market's 18% YTD gain was attributable to earnings growth.

S&P 500 price to book and ROE
Consensus forward price to book ratio

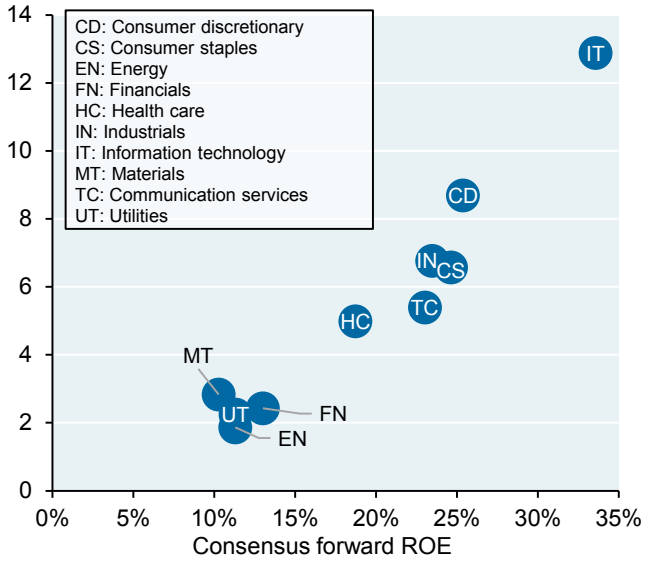

Source: Bloomberg, JPMAM, December 17, 2025

S&P 500 vs Hyperscalers & Semiscalers
Consensus forward price to book ratio

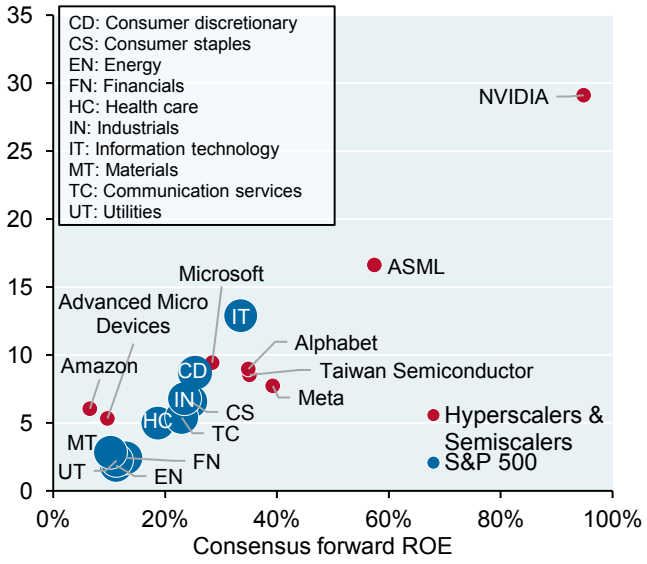

Source: Bloomberg, JPMAM, December 17, 2025

S&P 500 free cash flow margins: tech & interactive media vs the rest, Percent

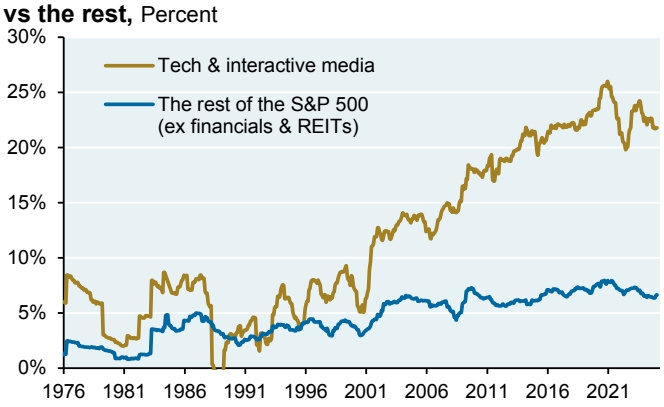

Source: Empirical Research Partners, December 2025

PEG ratios for global tech stocks vs the market
Forward price/earnings ratio divided by year 2 forward EPS growth

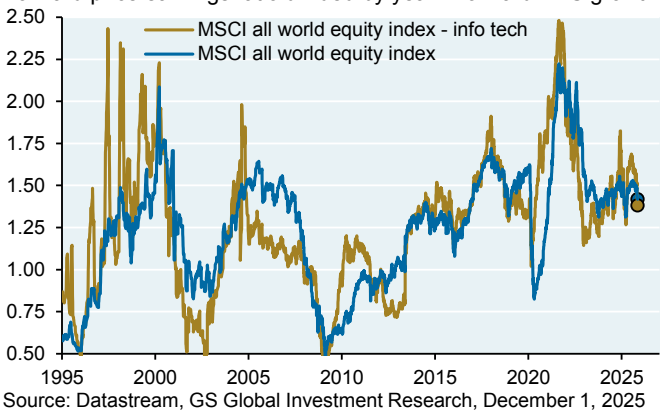

Source: Datastream, GS Global Investment Research, December 1, 2025

While the current tech capital spending surge as a share of GDP has now exceeded the dotcom era, technology valuations are roughly half the levels seen at the dotcom peak, and US tech PEG ratios are only 1x-3x in recent years compared to 4x-8x in the dotcom era. Just as importantly, **the equity market is less reliant on the market cap and spending by the "YUCs"**: young unprofitable companies the likes of which flooded the equity market during the dotcom boom and again during the SPAC era.

Private fixed investment in information processing equipment & software as a share of potential GDP, %

Source: BEA, CBO, JPMAM, Q3 2025

US tech valuations
P/E ratio

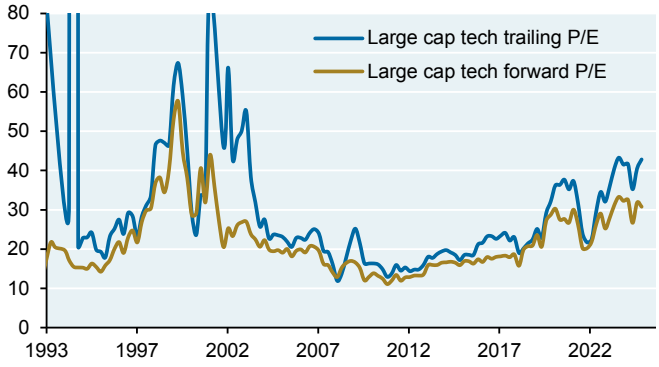

Source: Bloomberg, JPMAM, Q3 2025

Market cap of young unprofitable technology companies
% of total technology market cap

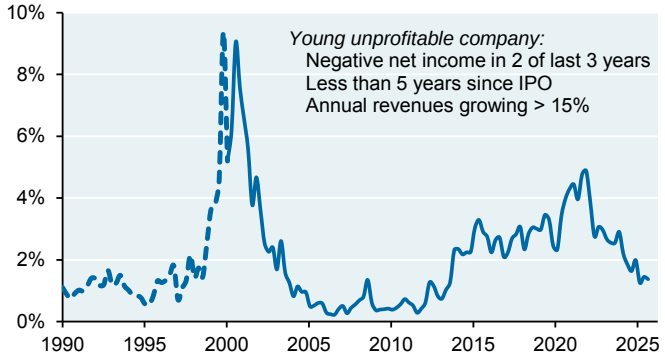

Source: Factset, JPMAM, Q3 2025

Spending by young unprofitable technology companies
% of total corporate spending on Selling, General & Administrative Expenses, Capital Spending and Research & Development

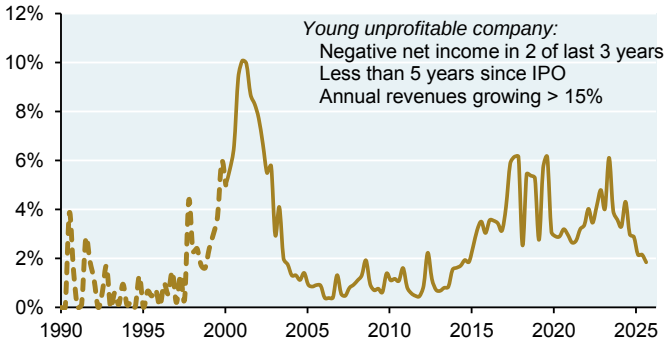

Source: Factset, JPMAM, Q3 2025

If valuations, tariffs, borrowing capacity and YUCs don't represent immediate threats to the moat companies, then what does? This year's Outlook walks through **What Could Go Wrong** (WCGW) over the medium term to the existing moat. After that, a few pages on everything else.

WCGW #1: A Metaverse moment for the hyperscalers 7

WCGW #2: US power generation constraints 20

WCGW #3: China scales the moat with its own lithography-semiconductor technology…when, not if 27

WCGW #4: China and Taiwan 33

The Rest: tech spending contributions to GDP, the Fed's inflation/employment dilemma, tariffs, immigration and the labor supply, the OBBBA, the US$, capital markets, US equity outperformance, healthcare stocks, China, Japan, a history of populism for investors and a Presidential cabinet time capsule 38

WCGW #1: A Metaverse moment for the hyperscalers

Of the four risks to the moat, this is the most immediate. Since Q4 2022, the four hyperscalers have spent $1.3 trillion on capital spending and R&D, much of it related to generative AI. Will all this investment eventually result in commensurate profits? If not, moat companies could face a version of the Metaverse reckoning that took place in 2022 when the Mag7 stocks fell in half.

Ancient history: Mag 7 stocks fell by 50% in 2022

Total return index

Source: Bloomberg, JPMAM, December 17, 2025

This first WCGW section[4] simplifies the avalanche of information on generative AI by bucketing everything into six categories. If you're exhausted from all the LLM benchmark scoring jargon and AI adoption surveys and just want to focus on actual revenue and profit consequences, **jump to page 11**; and if you're just interested in the bottom line on hyperscaler revenues and profits, **jump to page 15**.

[a] The progress of genAI models according to data science benchmarks

[b] Generative AI corporate adoption rates

[c] Generative AI impacts on corporate profits and employment, including the MIT/Wharton AI gap

[d] Hyperscaler revenues, free cash flow margins, cash balances and genAI commentary

[e] Questions around GPU and networking equipment depreciation

[f] Why might the genAI capital spending juggernaut continue, and the market risk from OpenAI

[4] **TPUs vs GPUs.** I did not include Google TPUs displacing NVIDIA GPUs as one of the four main risks since both companies are part of the moat on the cover. Some background: for most developers, Nvidia's software makes it easy to run applications on GPUs. Anthropic, Apple and Meta have fewer challenges using TPUs since they're capable of writing their own software. GPUs can generally handle a wider variety of tasks (rendering video graphics, model training, repetitive math operations, matrix multiplication); Google TPUs are more specialized to handle matrix multiplication but can only run certain models faster than GPUs. The software tools and the nature of the task heavily affect which option is cheaper/ faster. Some analysts see TPUs as more efficient users of high bandwidth memory, which is often a binding constraint for GPU systems. For example: as with crypto, hardware eventually moved from more general solutions to application-only products like ASICs. TPUs and Groq LPUs are examples of this applied to AI.

To compete with NVIDIA, Google would need to overhaul its supply chain to secure enough chips from foundries and make sure customers can use TPUs reliably, which would require investment in sales distribution networks, server designers and customer support engineers

[a] The progress of genAI models according to data science benchmarks

Examples of improved capabilities are everywhere, including the first chart below combining performance scores across 10 different benchmarks. **Don't get too carried away with these results since all benchmarks can be gamed with enough trial and error**.

- AI models match or exceed humans in image recognition, medium reading comprehension, visual reasoning, language understanding and competition-level math when they are trained on such materials in advance
- OpenAI's o3 achieved a rating better than 99% of human competitors in Codeforces coding competitions[5]
- Otto-SR, a language model for synthesizing medical research (Cochrane's Systematic Review), took just 2 days to reproduce a review that would have taken humans 12 work-years; it more accurately identified relevant research than human counterparts and more accurately extracted data from research sources[6]
- Look how rapidly most models improved in 2025 in solving real world GitHub issues, particularly Google's

AI still struggles in some areas, particularly on tasks where models aren't explicitly trained:

- The bottom left chart shows hallucination rates by model from Artificial Analysis, defined as the share of incorrect responses out of all questions asked; some of these hallucination rates are shockingly high
- Only some models appear to be experiencing real improvements in opaque reasoning ability that result from pretraining efforts (Google Gemini and Anthropic Opus); other models rely more on increased inference-time compute to improve performance scores. The bottom right chart is a proxy for this: models solve math problems without the benefit of chain-of-thought

AI model performance tracker
Artificial Analysis Intelligence Index v3.0*

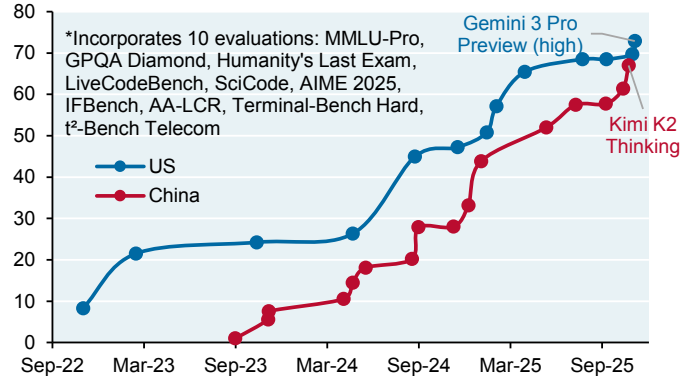

Source: Artificial Analysis, JPM Private Bank, December 23, 2025

Ability to solve real world GitHub issues
Percent resolved, highest scoring model by developer per quarter

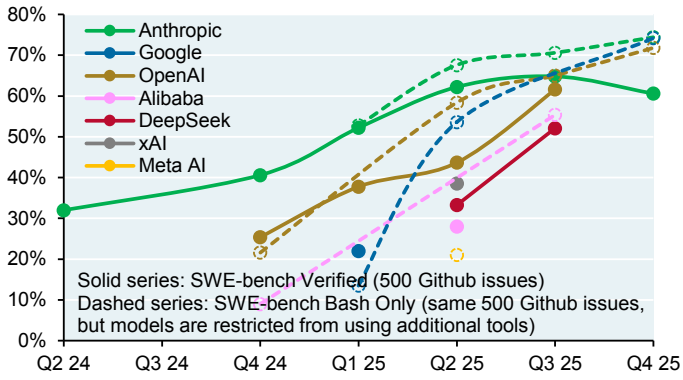

Source: Epoch AI, SWE-bench, JPMAM, December 17, 2025

AI model hallucination rates
Incorrect responses as a share of total questions

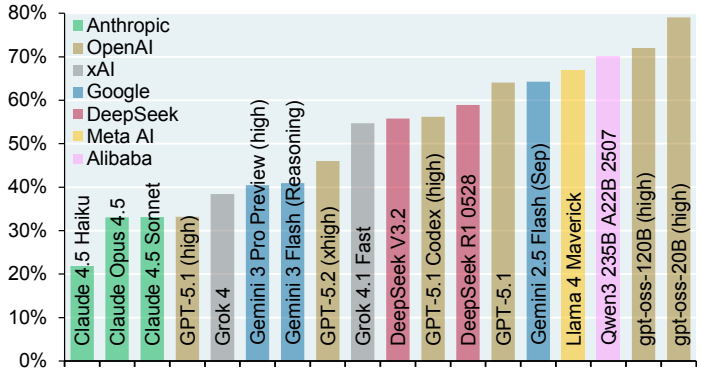

Source: Artificial Analysis, December 16, 2025

A proxy for opaque reasoning ability: the time it takes for humans to solve math problems that models can solve with no chain-of-thought (single pass), in minutes

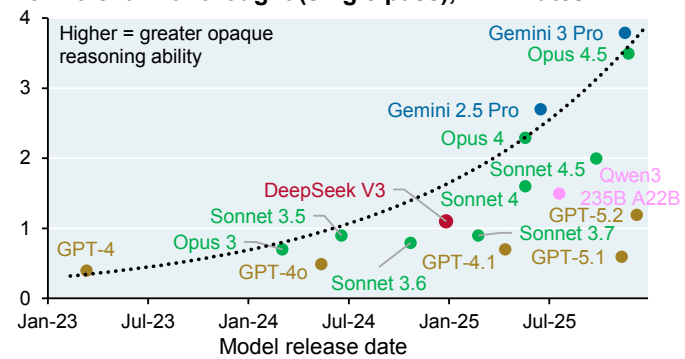

Source: Ryan Greenblatt, Redwood Research, December 2025

[5] Noam Brown (OpenAI), May 2025

[6] *"Automation of Systematic Reviews with Large Language Models"*, Cao, Arora and Cento, June 13, 2025

Wary as I am of self-serving promotional materials from AI companies, I did find a September 2025 piece from OpenAI to be interesting. Its analysts assessed the ability of AI models to complete realistic work tasks at levels approaching expert humans:[7]

- OpenAI released GDPval, a benchmark that evaluates AI capabilities on real-world, economically valuable work. The benchmark tasks draw from 44 professions such as law, finance and engineering (i.e. creating a business review PowerPoint, drafting a will and drawing a blueprint)
- Human experts with ~14 years of industry experience designed tasks that would typically take people seven hours to complete. AI models and humans attempted the tasks and human experts graded the work
- **While graders still preferred human work overall**, judges preferred some AI tools in industries like software engineering. Note on the bottom right how AI software engineering task complexity has been rising
- If human experts partner with AI tools (asking AI to attempt the task first, then just doing manual work to fix any issues), both cost and speed improved by ~1.5x

AI model overall performance on realistic work tasks
% of time AI work product is preferred vs work from human experts

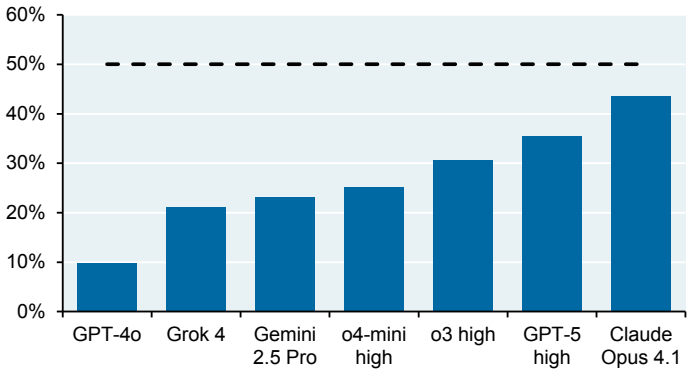

Source: "GDPval", Patwardhan et al (OpenAI), September 2025

AI model performance on software developer work
% of time AI work product is preferred vs work from human experts

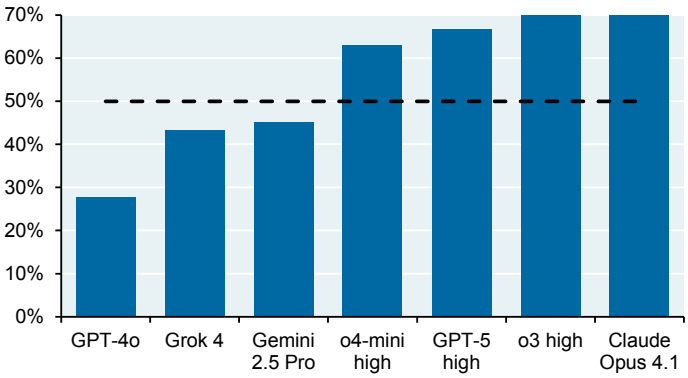

Source: "GDPval", Patwardhan et al (OpenAI), September 2025

Efficiency gains from AI assistance to human experts
Cost improvement

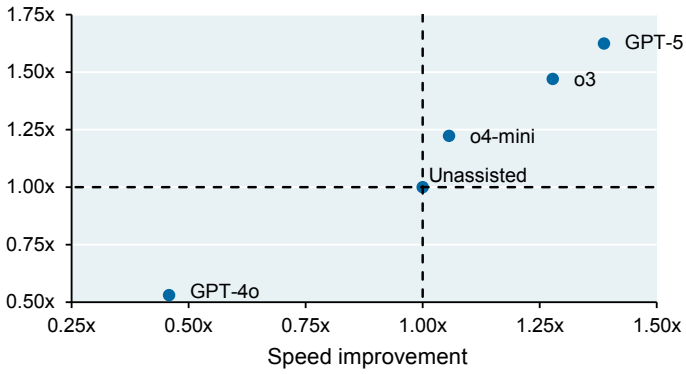

Source: "GDPval", Patwardhan et al (OpenAI), September 2025

AI software engineering task complexity, Time humans take to complete tasks AI models can complete at 50% accuracy, hours

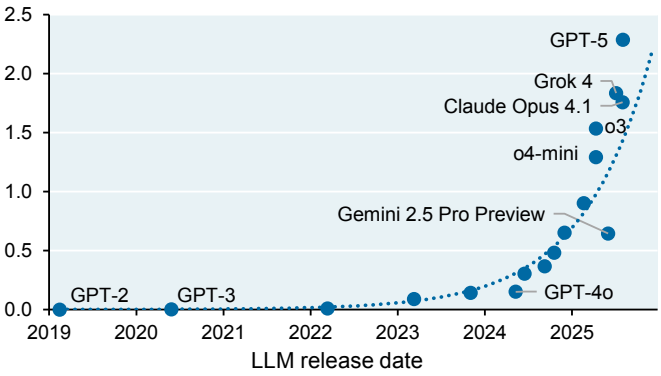

Source: METR, August 2025

Will models run out of data to train on? In its 2025 AI Index Report, Stanford concluded that this is unlikely before 2030. Common Crawl, an open repository of web crawl data frequently used in AI training, is estimated to contain a median of 130 trillion tokens. The indexed web holds approximately 510 trillion tokens, while the entire web contains around 3,100 trillion. Additionally, the total stock of images is estimated at 300 trillion tokens, and video at 1,350 trillion tokens.

[7] *"GDPval: Evaluating AI Model Performance on Real-World Economically Valuable Tasks"*, Patwardhan et al (OpenAI), September 2025. Note: Gemini 2.5 was used by OpenAI since Gemini 3 was not yet available

[b] Generative AI adoption rates

The usual suspects are shown in the first three charts on rising AI adoption according to surveys. The fourth chart from Anthropic shows the fraction of occupations (y-axis) that have some tasks executed by or with AI (x-axis). For example, 36% of occupations use AI on 25% of their tasks[8]. From a December 2025 OpenAI report: "Enterprise usage is scaling, with deeper workflow integration. ChatGPT message volume grew 8x, and API reasoning token consumption per organization increased 320x year-over-year, demonstrating that more enterprises are using AI, and their intensity of usage has increased"[9].

To be clear, these kind of surveys tell us nothing about how much companies or individual users are willing to pay for these models, or how much can be earned by advertising while they use them.

AI adoption rate by sector, Share of US businesses with paid subscriptions to AI models, platforms & tools

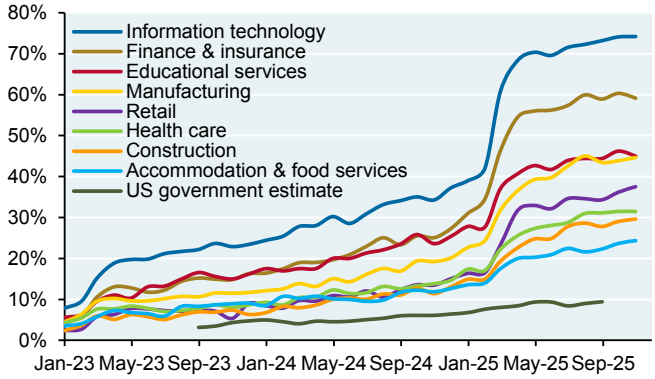

Source: Ramp, November 2025

Census: AI adoption rates by sector and date
Share of firms using AI, n=200,000 businesses

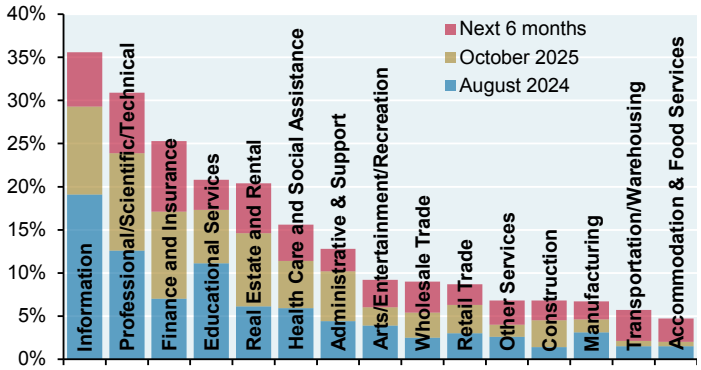

Source: US Census Bureau, JPMAM, October 2025

Generative AI adoption survey
Percent of adults using generative AI, n=25,000 respondents

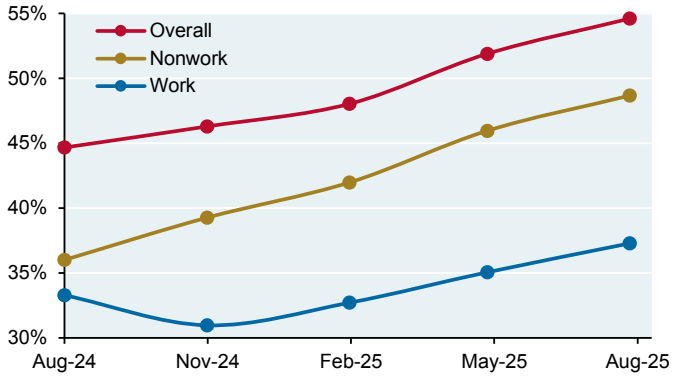

Source: St. Louis Fed, November 2025

Depth of AI usage across organizations
Fraction of occupations, n=1 million Claude conversations

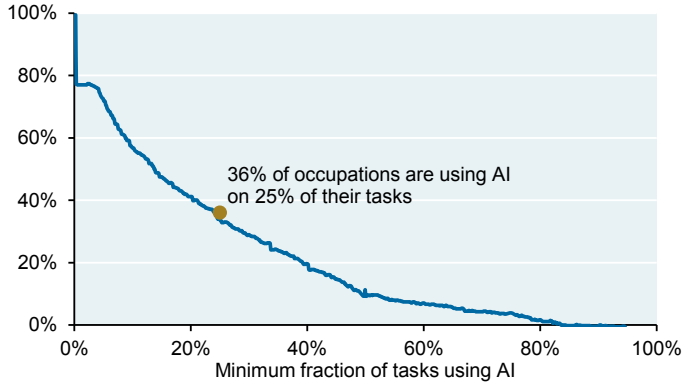

Source: Anthropic, Stanford AI Index Report, 2025

[8] Another use case: generative AI text-to-video tools in the entertainment industry, a topic we covered in the **November 2025 Eye on the Market**

[9] "*The state of Enterprise AI*", OpenAI, December 7, 2025

[c] Generative AI impacts on profits and employment

The first chart tracks the number of companies reporting quantifiable AI benefits and the second chart breaks down these benefits by type, with the most frequent being labor productivity gains and other cost savings.

Companies experiencing quantifiable AI benefits
Percent of companies

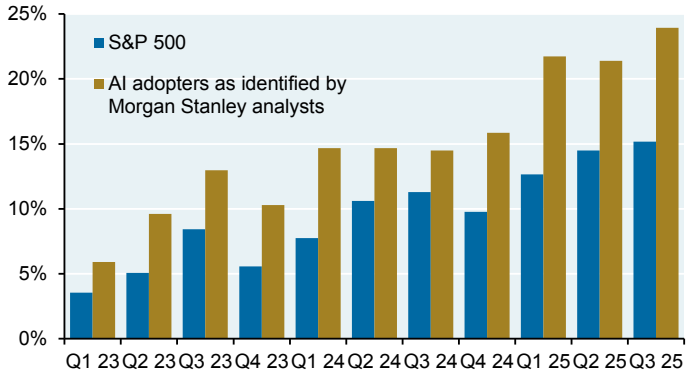

Source: Morgan Stanley Research, November 2025

Mentions of AI benefits in earnings/conference transcripts
Number of mentions per quarter by category, n=13,500 transcripts

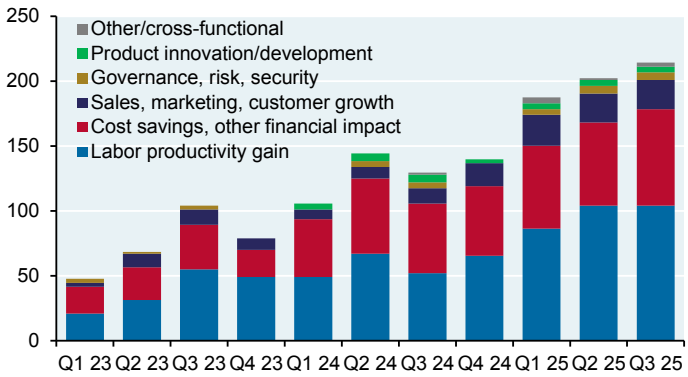

Source: Morgan Stanley Research, November 2025

The next chart gets closer to measuring the AI impact we're looking for: revenue increases and cost declines resulting from generative AI adoption. Note that the largest number of respondents found cost and revenue benefits from AI to be rather small, at least so far (the blue and grey bars).

Cost decreases and revenue increases from generative AI use
Percent of survey repondents, n=2,854 respondents

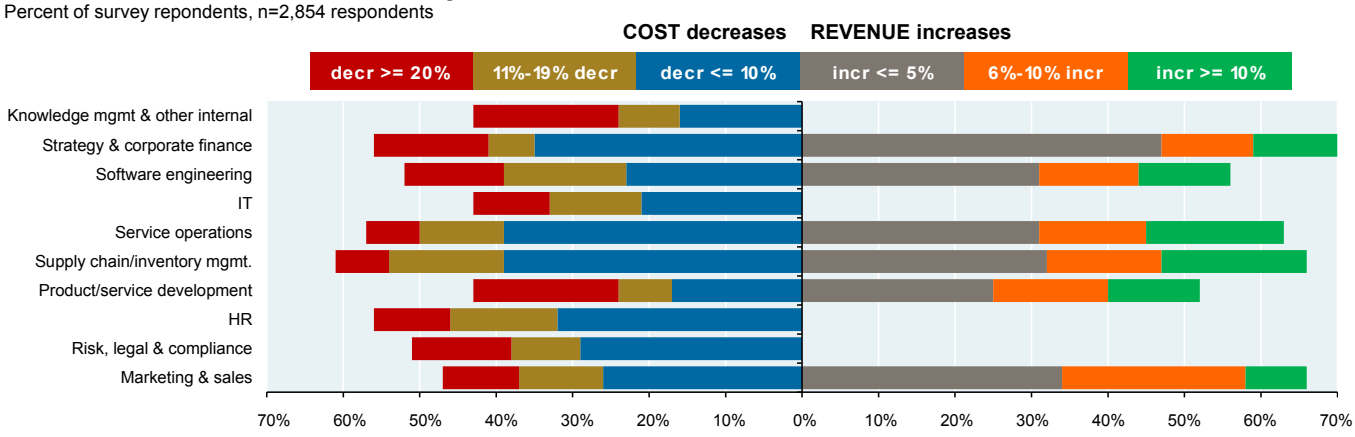

Source: Stanford AI Index Report, 2025

Goldman's latest piece concluded that AI applications are already boosting labor productivity and showed the two charts below as supporting evidence[10]. When we look at the data, overall non-farm productivity and tech productivity are higher since ChatGPT's launch when compared to the pre-COVID period (see tables), although these results could be entirely coincidental.

While use cases underlying these charts are mostly limited to coding, customer service and consulting, Goldman projects that adoption will spread to other industries. They also estimate a 15% boost to US productivity and an $8 trillion net present value gain (!!) associated with revenues unlocked by generative AI. Given the projected decline in US labor force growth (from an aging population, low birth rates and Trump immigration policies), inflationary pressures from tariffs, a weaker dollar and high budget deficits, a large AI productivity boost may be needed to sustain US GDP growth and keep inflation low.

Company estimates of the effect of AI on labor productivity, Percent

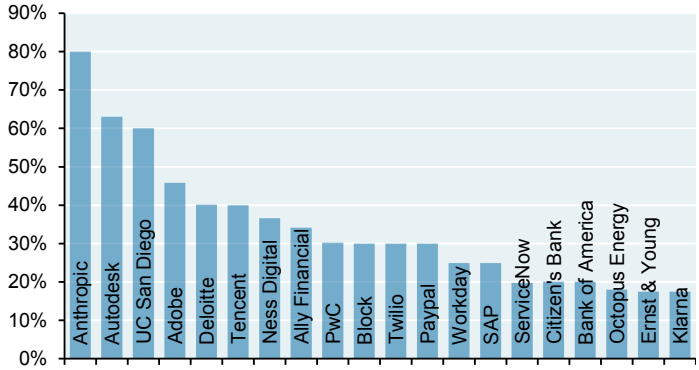

Source: GS Global Economics Research, December 2025

Academic estimates of the effect of AI on labor productivity, Percent

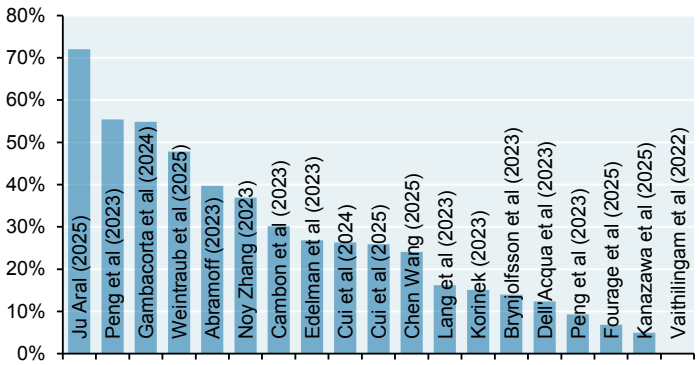

Source: GS Global Economics Research, December 2025

Productivity gains, pre-COVID vs post-GPT

Real GDP method (annualized)	From To	Q1 2016 Q1 2020	Q4 2022 Q2 2025	Q3 2023 Q2 2025
Non-financial corporate		1.1%	3.1%	3.9%
Information sector		5.0%	8.9%	6.0%
Data processing		8.0%	12.1%	9.9%

Source: BLS, Bloomberg, JPMAM, Q2 2025

Productivity gains, pre-COVID vs post-GPT

Real gross output method (annualized)	From To	Q1 2016 Q1 2020	Q4 2022 Q2 2025	Q3 2023 Q2 2025
Non-financial corporate		1.1%	3.1%	3.9%
Information sector		3.7%	7.3%	5.7%
Data processing		7.8%	10.2%	10.6%

Source: BLS, Bloomberg, JPMAM, Q2 2025

[10] "*The AI spending boom is not too big*", Goldman Sachs US Economic Research, October 15, 2025

Point-Counterpoint: MIT and Wharton are miles apart, literally and figuratively

A July 2025 MIT[11] paper came to a subdued conclusion on gen AI: disruption has been confined so far to media & telecom and professional services, and the rate of successful genAI deployment is a fraction of genAI investigation and pilot rates. Here are some excerpts from the MIT paper:

Generative AI disruption index by industry
AI market disruption index

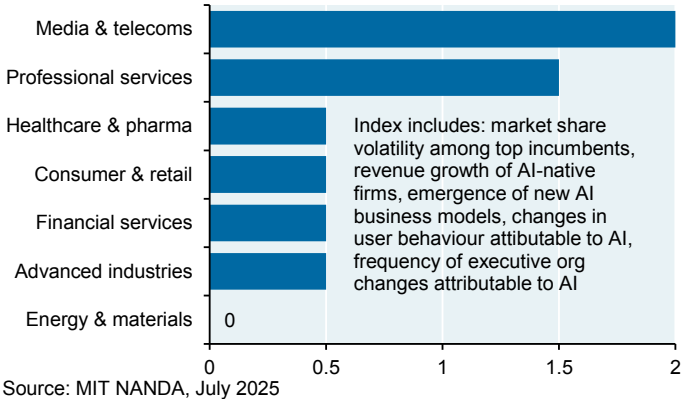

Source: MIT NANDA, July 2025

Generative AI project deployment rates
Percent of organizations

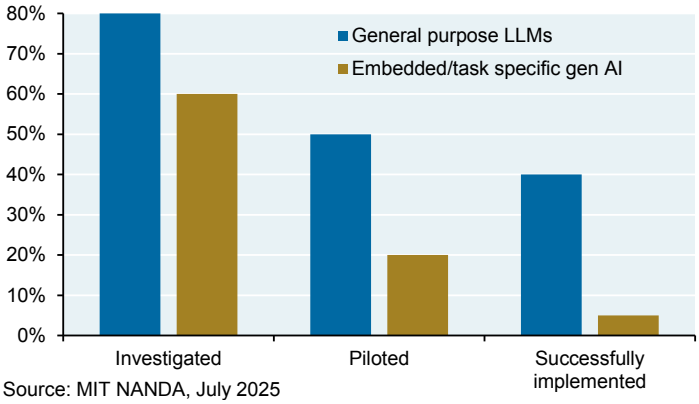

Source: MIT NANDA, July 2025

Survey results on the barriers to core workflow integration of generative AI, Percent

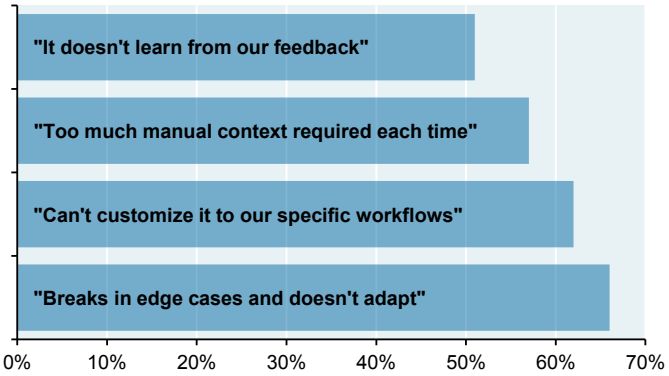

Source: MIT NANDA, July 2025

Some quotes from the MIT report:

- Despite $30-$40 bn in enterprise investment in genAI, 95% of organizations are getting zero return
- The share of CEOs that are "very confident" in their AI strategy has fallen from 82% in 2024 to 49% in 2025
- Indicative feedback from respondents: "It's excellent for brainstorming and first drafts but doesn't retain knowledge of client preferences or learn from previous edits. It repeats the same mistakes and requires extensive context input for each session. For high stakes work, I need a system that accumulates knowledge and improves over time"

[11] "*The State of AI in Business 2025*", MIT NANDA, July 2025. The report is based on reviews of 300 publicly disclosed AI initiatives, interviews with 52 organizations and survey responses from 153 company leaders

Wharton is 300 miles from MIT and a million miles away regarding perceptions of genAI's impact. In October 2025, Wharton published "*Accountable Acceleration: Gen AI Fast-Tracks into the Enterprise*". Wharton surveys of 801 senior decision makers at US commercial organizations with over 1,000 employees and $50 million in revenue found a more optimistic view on corporate genAI ROI and future AI spending than MIT (see the blue and gold bars in the charts).

That's the challenge with surveys; they're not substitutes for tracking actual profit and spending data, and often have small and unrepresentative sample sizes.

GenAI return on investment by company size

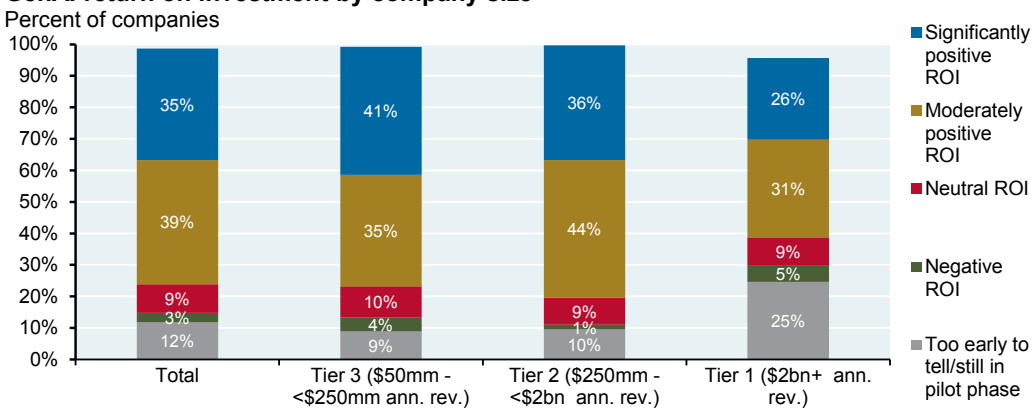

Source: Wharton Human-AI Research, October 2025

GenAI return on investment by industry

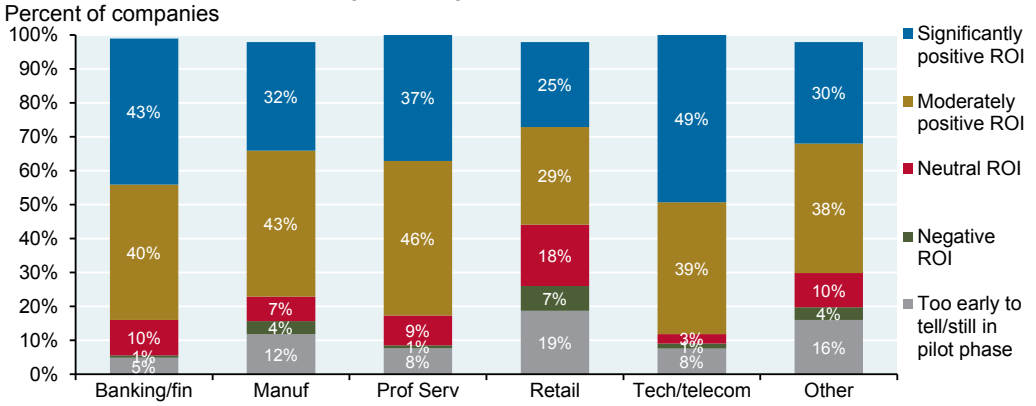

Source: Wharton Human-AI Research, October 2025

GenAI budget investment for the next 12 months by functional area

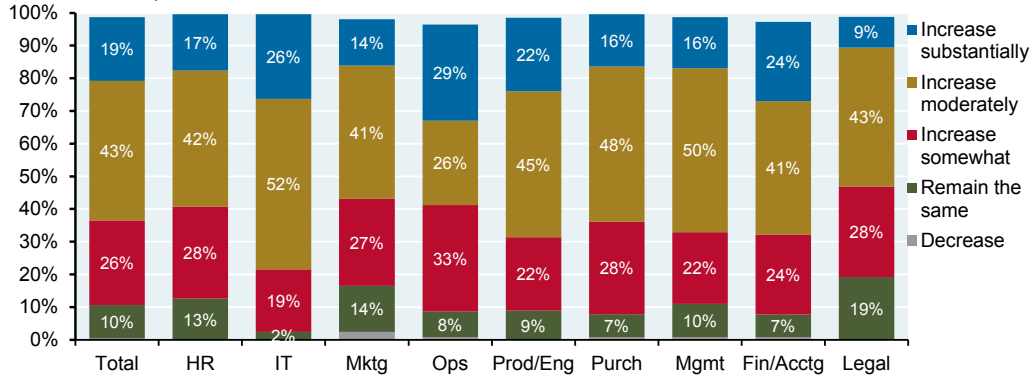

Source: Wharton Human-AI Research, October 2025

[d] Hyperscaler revenues, free cash flow margins, cash balances and genAI commentary

Despite the enormous amount of capital deployed so far, tracking hypescaler AI profitability is difficult. As far as we can tell, Microsoft is the only hyperscaler that discloses AI specific revenue, although there are a few other public and private AI-specific revenue observations shown below in the fourth chart. **Given gradually declining hyperscaler free cash flow margins and falling cash balances, a clearer path to profitability on AI investments may be needed in 2026 for current valuations to be sustained**. That path is made murkier by questions on GPU and networking equipment depreciation which we discuss on the next page. As for vendor financing, NVIDIA sent a note in November to Wall Street analysts to try and put this issue in context[12].

Comments from earnings calls

- **Microsoft's** inference business: fastest line in the business history to reach a US$10 bn revenue run rate, with growth constrained by a lack of available compute capacity [Q1 2025]… "Revenue from our AI business was above expectations" [Q3 2025]
- **Amazon's** AWS AI business is a "once-in-a-lifetime opportunity", achieved billions in annualized revenue but requires investment in infrastructure "in advance of when we can monetize it with customers" [Q3 2024]… Rufus monthly users were up 140% y/y, on track to deliver over $10 bn in incremental annualized sales [Q3 2025]
- **Google's** revenue from products built on its generative AI models grew more than 200% year-over-year. AI Max for Search is used by hundreds of thousands of advertisers, currently making it the fastest-growing AI-powered Search Ads product. In Q3 alone, AI Max unlocked billions of net new queries [Q3 2025]

Hyperscaler quarterly revenue from cloud services
US$, billions

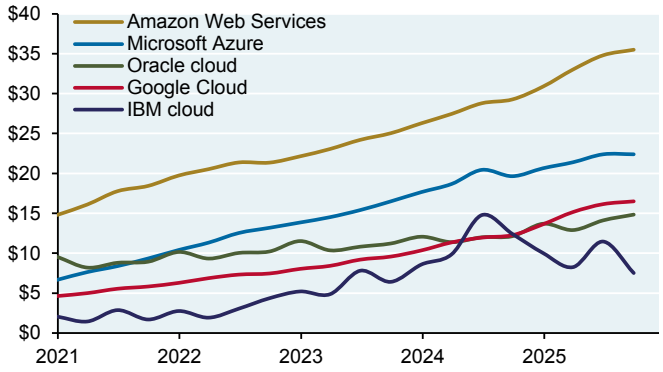

Source: Bloomberg, JPMAM, Q3 2025

Hyperscaler free cash flow margins (net of capex)
Percent, trailing 12 months

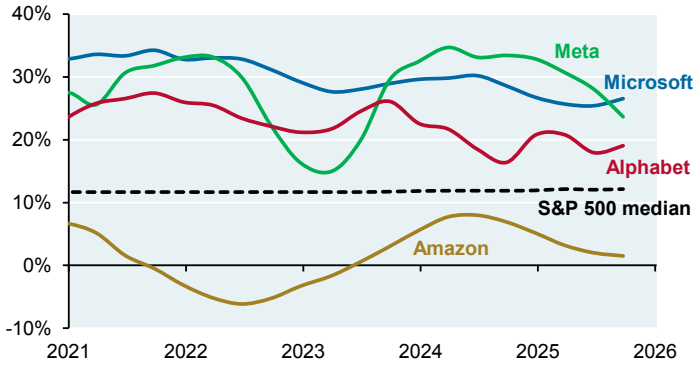

Source: Bloomberg, JPMAM, Q3 2025

Hyperscaler cash and cash equivalents
Share of total assets

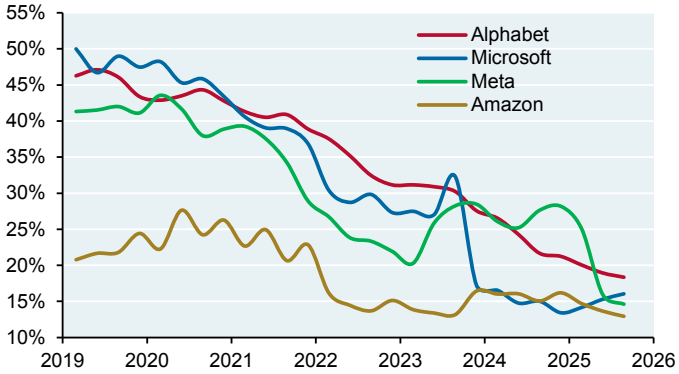

Source: Bloomberg, JPMAM, Q3 2025

Tracking annualized revenue from AI
US$ billions, quarterly/periodic revenues annualized

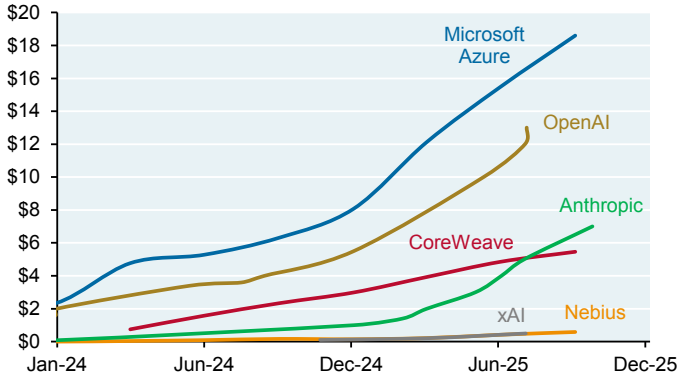

Source: Bloomberg, Epoch AI, JPMAM Q3 2025

[12] According to NVIDIA: in 2025, it invested $4.7 bn in private companies which represented just 3% of its revenues; these companies primarily raise capital from third party sources other than NVIDIA; and they predominantly generate revenue from third-party customers rather than NVIDIA. The company also does not rely on extensive vendor financing arrangements and cites a stable DSO (days sales outstanding) of 53

[e] Questions around GPU and networking equipment depreciation

Since 2020 hyperscalers ramped up depreciation lives of GPU and networking assets. Rationale: older chips remain in use for several years; even NVIDIA A100 chips still run at high utilization and generate positive margins beyond 2 - 3 years. GPUs don't become useless when new versions arrive; processors can be repurposed for inference or resold to emerging countries. Also: it's expensive to replace GPUs since power/HVAC equipment may need to change, so some users will not swap out for every upgrade. That said, most observations for GPU hourly rental rates paid by neocloud and hyperscaler customers have declined by 20% - 25% over the last year, trends which may partially explain the 50%+ decline in CoreWeave stock since its June 2025 peak (the stock is still up 90% from its March IPO price).

NVIDIA GPU hourly rental rates

GPU	Rate customers pay to:	US$/hr rental cost: Nov-2024	Dec-2025	Percent change
A100	Neoclouds	$1.8	$1.4	-22%
A100	Hyperscalers	$3.1	$3.6	15%
H100	Neoclouds	$2.7	$2.0	-26%
H100	Hyperscalers	$8.9	$7.0	-21%
H100	AWS	$4.5	$3.4	-24%
H200	AWS	$3.0 [a]	$3.3	11%
A100	AWS	$2.7	$2.0	-26%

Source: Silicon Data, 3Fourteen Research, JPMAM, December 1, 2025.
[a] Time series begins in February 2025

GPU/networking equipment depreciation assumptions

Years

Company	2020	2021	2022	2023	2024	2025
META	3.0	4.0	5.0	5.0	5.0	5.5
Google	3.0	4.0	4.0	6.0	6.0	6.0
Oracle	5.0	5.0	5.0	5.0	6.0	6.0
Microsoft	3.0	4.0	6.0	6.0	6.0	6.0
Amazon	4.0	4.0	5.0	5.0	6.0	5.0

Source: Company SEC filings, November 2025

In any case, what would happen if depreciation lives declined to where they used to be, a change possibly triggered by (a) hyperscalers refurbishing existing data centers with new GPUs while discarding the older ones, (b) evidence confirming that data center GPU failure rates are ~9% annually or (c) users opting for Blackwell GPUs even for inference, which could reduce the value of older generation chips?

The table estimates the impacts on depreciation, margins and EPS using 3-year depreciation for GPUs and networking equipment added since GPT's launch in 2022. Given incomplete disclosures, we had to make some assumptions which are listed below the table. The results: EPS and operating margin revisions would range from -6% to -8%, other than for Oracle whose declines would be larger. The next stage might entail more disclosures from hyperscalers clarifying the basis for GPU and networking depreciation assumptions which could partly defuse this issue. Trump's decision to allow NVIDIA H200 chips to be sold to China could support long term values and support existing depreciation assumptions.

Income statement impact of reducing GPU/networking depreciation assumption to 3 years on capex since Jan 2023

	Amazon	Meta	Microsoft	Google	Oracle
Current GPU/NW depreciation assumption, years	5.0	5.5	6.0	6.0	6.0
Current GPU/NW depreciation, $ million (estimated)	10,315	7,394	8,365	8,667	2,509
Current operating margin, percent	11.0	43.2	46.3	32.2	30.4
Current pro-forma EPS, $ per share	7.17	23.22	14.11	10.29	4.38
Revised GPU/NW depreciation assumption, years	3.0	3.0	3.0	3.0	3.0
Revised GPU/NW depreciation, $ million (estimated)	17,192	13,556	16,731	17,335	5,017
Revised operating margin, percent	10.0	40.0	43.4	29.9	26.2
Revised pro-forma EPS, $ per share	6.64	21.52	13.19	9.70	3.62
Change in operating margin, percent	(1.0)	(3.3)	(2.8)	(2.2)	(4.3)
% change in operating margin	-9%	-8%	-6%	-7%	-14%
Change in pro-forma EPS, $ per share	(0.53)	(1.70)	(0.92)	(0.59)	(0.76)
% change in pro-forma EPS	-7%	-7%	-7%	-6%	-17%
Assumptions. AI share of capex: 2023 40%, 2024 55%, 2025 70% (for Amazon, 20%, 30% and 50%). GPU share of AI spend 45%, networking share of AI spend 15%; GPU/NW salvage value of 10%					

Source: Bloomberg, JPMAM, 2025.

[f] Why might the genAI capital spending juggernaut continue and the market risk from OpenAI (warning: tech jargon overload ahead; skip to page 20 if you want to get to the power discussion)

There are examples of companies whose earnings and stock prices rose after adopting genAI to manage supply chains and logistics (CH Robinson), to automate mobile advertising and ad-bidding (AppLovin) and to improve speed/efficiency of customer service (Mr. Cooper Group). Palantir and Databricks are also examples of AI embedded into workflows and contributing to earnings growth. But at a market-wide level, **so far only the AI infrastructure providers have generated substantial excess returns**. Baskets of companies presumed to benefit from selling AI products and/or lower labor costs have been flat to the equal weighted S&P 500. If so, other than the infrastructure plays, genAI is still a stock-pickers story rather than a secular one.

So far, AI infrastructure company benefits only
Indexed return vs equal weighted S&P 500, Dec 31 2023 = 100

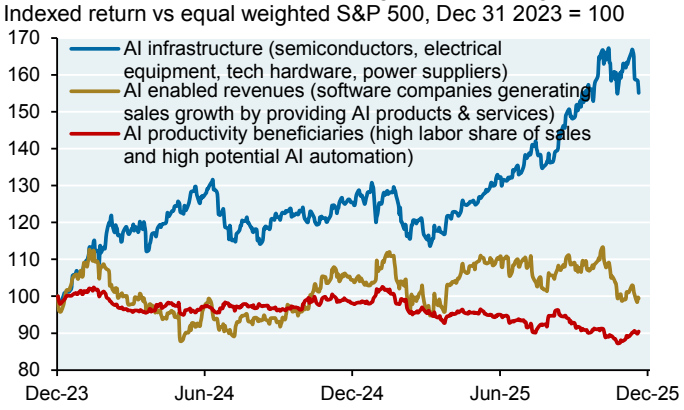

Source: Goldman Sachs, JPMAM, December 17, 2025

Token consumption through OpenRouter API calls
Trillions of tokens

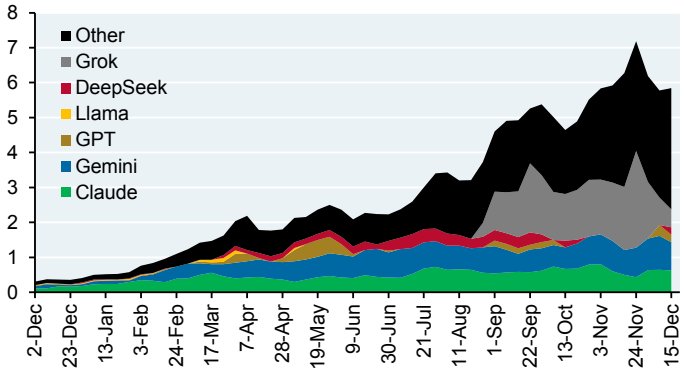

Source: OpenRouter LLM Ranking Leaderboard, December 15, 2025

In the interest of equal time, what's the bull case on the genAI capex juggernaut? I apologize in advance since this is about to get very technical. The chart on the right shows developer token consumption via API calls to OpenRouter which started 2025 at less than 0.5 trillion per week and rose to 6-7 trillion as the year ended. API calls are a small subset of overall token consumption, but it's an interesting high-frequency proxy for AI activity. Similarly, Google monthly token consumption rose 100x from 9.7 trillion in April 2024 to 980 trillion by July 2025.

I'm not a futurist and generally avoid this kind of thing, but here's a thought experiment: **how do future computing needs for consumer inference stack up against current AI workload capacity?**[13]

- Assume that 5 billion people end up using AI by the 2030's (compared to 6 billion current internet users)
- Each user consumes ~1.6 mm tokens per day for search, coding assistance, other agents, background assistants and creative purposes. That would be **a LOT** of demand vs current levels and assumes a paradigm shift in how AI is used in daily life[14]. Across all users, that would be 8×10^{15} tokens per day
- If a model needs 2*N flops per token for inference, a 5 trillion parameter mixture of experts model with 150 billion activated parameters might need roughly 0.3 teraflops per token. If we multiply this amount by 3x to incorporate the need for long-context/agent overhead, we end up needing 0.9 teraflops per token
- **Last step (I promise):** 5 billion users, 1.6 mm tokens per day per user and 0.9 teraflops per token yields ~7.2 billion exaflops per day to meet future consumer inference demand. However, that appears to be well above total available AI workloads which we explain next

[13] I would like to thank Matt Dratch at ExodusPoint Capital for most of these figures and hypotheticals

[14] EpochAI uses a high-side estimate of 500 tokens per query, so 25 queries per day would only require 12,500 tokens. Other tasks require more; Vamsi Chemitiganti at AWS describes the following token intensity by task:

- Basic chatbot interactions 50-500 tokens; document summarization 200-6,000 tokens; code assistance 500-2,000 tokens; complex code generation: 20,000-100,000+ tokens; legal document analysis: 75,000-250,000+ tokens; multi-step agent workflows: 100,000-1,000,000+ tokens
- 1.6 mm tokens per day would be very expensive at current GPT 5.1 rates (~$300 per month). The implicit assumption is that token prices continue to sharply decline (they have declined by at least 240x so far)

How much AI workload compute exists today[15]? For 2026, SemiAnalysis estimates Western AI workload compute of ~36,000 exaflops at FP8 precision, a large increase over 2025 levels of 11,500 exaflops. Assuming a 50% utilization factor, that would imply ~1.7 bn exaflops per day compared to the 7.2 bn exaflops of consumer inference demand cited on the prior page. Using a broader range of GPU utilization rates from 33% to 75%, 2026 AI compute would only meet 14% - 66% of future demand just for consumer inference, **before adding any demand** from enterprises, sovereigns/defense or low-latency inference (robotics, autonomous cars, warehouse automation, delivery robots, farming and commercial drones).

Current hyperscaler compute as a share of future consumer inference demand, 5 bn users, based on GPU utilization and intensity of consumer token demand

	GPU utilization rate		
	75%	50%	33%
800,000	66%	44%	29%
1,200,000	44%	29%	19%
1,600,000	33%	22%	14%

Cloud tokens per day of future inference demand
Source: JPMAM

On flops and floppers
GPU discussions often involve flops, or floating point operations per second. Such measurements are used to assess computational power, training/inference time and model efficiency
The other major context for flops is in **soccer**. In 2018, Bleacher Report rated the 15 worst floppers in the sport based on a diving score derived from frequency of dives, diving success rate and the degree of dive difficulty. **The five worst floppers** according to their scoring: Angel di Maria (Real Madrid) 104.9, Sergio Busquets (Barcelona) 102.5, Cristiano Ronaldo (Real Madrid) 92.7, Arjen Robben (Bayern Munich) 87.9 and Neymar (Santos) 82.2. Honorable Mention: Jurgen Klinsmann

This is not meant as a projection; I just wanted to illustrate what some optimistic arguments I see and hear are based on. AI professionals at JP Morgan helped me outline below the factors that may affect future demand for AI workload compute; the hyperscalers are obviously betting the ranch on the top part.

Increased demand for cloud computing resources
Reasoning models use 20x - 40x more tokens than non reasoning ones. More and more agents will use reasoning models, and most coding tasks will use reasoning models
We are still in very early stages and many of the domain specific models (e.g., legal, finance..) are still in their infancy
Most large companies have still not scaled up AI use cases (especially the high-volume ones)
Creative work is very token intensive, and most studios, ad agencies and citizen content generators are still in early stages of using this technology
Agents will be used extensively for constant monitoring of software services in production for dev-ops, security ops, anomaly detection …. and these tend to be very input token intensive
AI for healthcare/drug discovery is in very early stages and could be big wildcards in terms of compute required
More and more data will become AI ready and AI will generate more and more data (flywheel)
Fraud Detection, recommender systems and personalization is already moving to transformer based architecture that requires far more compute than traditional models
Decreased demand for cloud computing resources
Ensemble of models, mixture of experts, GPU advances, inference optimization and specialized accelerators will continue to bring the cost down by reducing the compute required for a given workload
Big companies will likely move from generalized models to post trained smaller purpose-built models
AI compute will move to edge devices, PCs and mobile devices as cost of adding accelerators is cheap

[15] While there are 125 GW of data centers around the world, only 10-13 GW are currently capable of handling AI workloads (as opposed to traditional enterprise workloads)

A simpler way to think about the bull case: the latest forecasts from OpenAI[16]. By some accounts OpenAI is 6-12 months ahead of competitors based on user growth[17], monetization and model/product capabilities. The first table shows how OpenAI's forecasts for 2027 that the company made in mid-2025 were already eclipsed by forecasts for 2027 it made just three months later.

If OpenAI stays on its current trajectory, the company projects it will still need more compute for training than for inference through to 2030. But look at the second table; for these forecasts to materialize, OpenAI would need 30 GW of new generation capacity by 2030. That's a tough hurdle; as we explain in the next section, the entire US added just 25 GW in 2024 after adjusting for the intermittency/reliability of new capacity. **Given combined cycle turbine delays, grid interconnection bottlenecks, wildly overly optimistic timetables for small modular reactors and nuclear fusion and limited opportunities to recommission shuttered but still viable nuclear fission plants, energy may be a significant constraint on AI, a topic we discuss next**.

Before covering the energy topic, I want to add that **OpenAI is arguably the biggest individual corporate risk to the whole AI story, even more than NVIDIA and despite OpenAI being a private company**. OpenAI is on track to make $10-$20 bn in revenue and has commitments of $1.4 trillion to its corporate partners[18], and currently survives on subscription/developer AI fees with limited or zero search/advertising revenue, cloud computing revenue or hardware sales. When pressed on this issue on Brad Gerstner's BG2 podcast, Altman replied "If you want to sell your shares, I'll find you a buyer. I just.....enough"[19]. Meanwhile, OpenAI's own chief economist disclosed that 72% of GPT queries are not business related. The next 12-18 months will be very interesting.

OpenAI forecasts keep rising

	2027 forecast		
	Mid-2025	10/31/2025	Change
Revenues, US$ bn	$60	$90	50%
Inference compute cost, US$ bn	$21	$30	43%
Weekly average users, bn	1.4	1.8	29%
Free user average rev per user	$7	$12	71%
Paid subscribers, mm	58	76	31%
Paid user average rev per user	$740	$880	19%
Billions of API tokens/minute	11	17	55%

Source: Barclays Research, November 2025

Current and future GPU and power requirements

	2024	2025E	2030E
New B200 GPU equivalents added			
For inference	72,098	260,354	1,974,686
For training	120,163	380,518	2,231,034
Implied power requirements assuming 400,000 B200 equiv. GPUs per GW			
New capacity per year, GW	0.5	1.3	5.6
Cumulative capacity, GW	0.5	1.8	30.2

Source: Barclays Research, November 2025

OpenAI projected compute budget

US$, billions

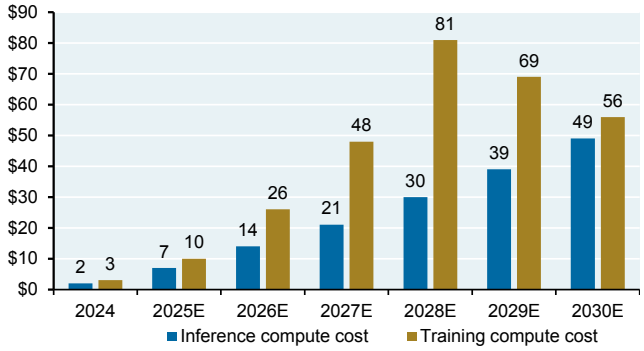

Source: Barclays Research, November 2025

[16] "*Unpacking AI demand from OpenAI*", November 18, 2025, Barclays Research

[17] According to Sensor Tower, OpenAI's non-business Chat GPT users spend the equivalent of 80,000 years per week using it compared to an aggregate of 7,000 years per week for DeepSeek, Gemini, Perplexity, Grok, Claude, Kimi and Tencent non-business use combined

[18] Fortune, November 28, 2025; CNBC December 16, 2025

[19] "*Risks Facing OpenAI and its $1.4T in Spending Commitments*", Carnegie Investment Counsel, Nov 20, 2025

WCGW #2: US power generation constraints

"*Eventually the cost of intelligence, the cost of AI, will converge to the cost of energy. How much you can have – the abundance of it - will be limited by the abundance of energy*", Sam Altman Senate testimony, May 2025

"*The increase in peak load forecasts for 2026 and 2027 were almost entirely due to existing and projected large data center load additions to the grid*", Monitoring Analytics, PJM Independent Market Monitor, October 2025

Data centers only represent 4% - 8% of US power demand but are projected to account for two thirds of load growth, even after accounting for ~40% annual energy efficiency improvement in chip design. Are we getting closer to a wall of power constraints? Let's start here: why would Microsoft and Constellation Energy restart a decommissioned nuclear reactor at Three Mile Island at a cost of $110-$130 per MWh for 20 years (~2x the national average for a 20-year power purchase agreement) if they could acquire baseload power more cheaply?

Power required to train frontier models has doubled each year[20]; data center and power spending are on the rise after years of low growth; and in data center-heavy PJM competitive power markets, capacity prices are soaring[21]. **The 4th chart is the most important: growth in new generation and storage capacity is a lot more gradual after adjusting for its reliability/intermittency (i.e., derating of solar, wind and hydro)[22].**

Power used to train frontier AI models
MW

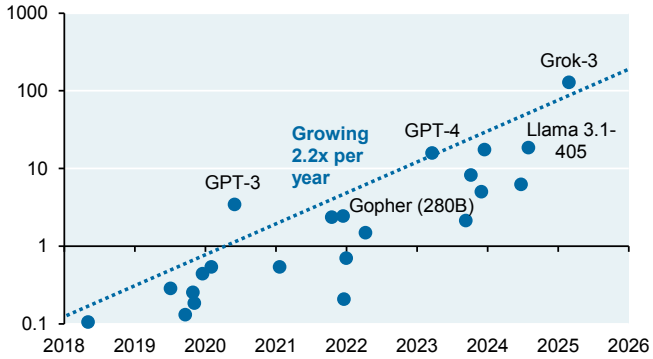

Source: Epoch AI, August 2025

US construction spending
US$, billions, seasonally adjusted annual rate

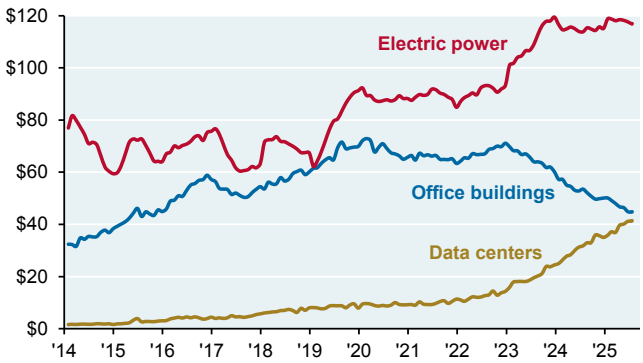

Source: Census Bureau, JPMAM, August 2025

PJM capacity prices are soaring
Base Residual Auction rate, US$ per MW by delivery year

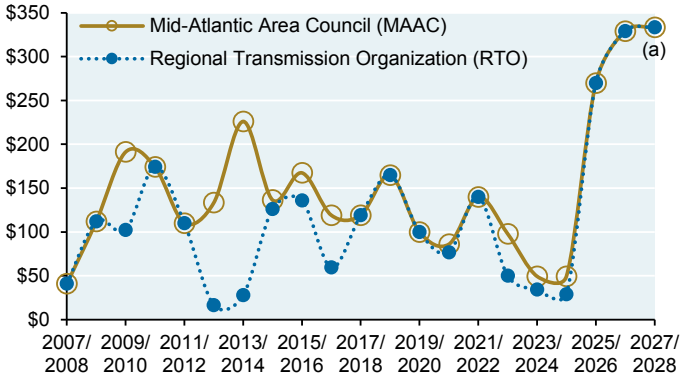

Source: NRG Energy, July 23, 2025. (a) uncapped would have been $530/MW

US electricity generation and storage capacity additions: nameplate and ELCC-weighted, Gigawatts, annual

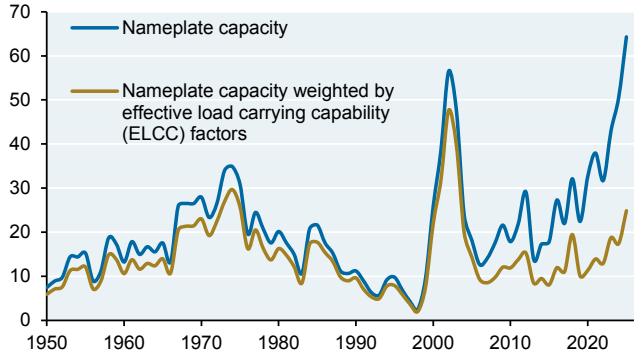

Source: EIA, PJM, MISO, ERCOT, CAISO, Thundersaid Energy, 2025

[20] According to Epoch AI, power demand would grow 4x per year without mitigating factors. Increased duration of training runs could reduce power demand growth to 3x, and improved chip efficiency could further reduce growth to 2.2x. That happens to be roughly the pace at which power required to train frontier models has grown so far

[21] **PJM** region (data center alley: VA, PA, MD, OH) has 67 GW of existing and planned data center capacity, followed by **ERCOT** (TX) with 20 GW and **CAISO** (CA, NV) with 8 GW. Factors driving the spike in PJM capacity payments include retirement of thermal assets, payments to keep generators running past planned retirement, data center growth and a transition to intermittency-weighted capacity formulas. **Capacity payments** are insurance premiums paid to generators to commit future supply or demand response during periods of peak demand

[22] We use effective load carrying capability estimates from PJM, MISO, ERCOT and CAISO for the adjusted series. **ELCC** estimates anticipated firm capacity (as a % of nameplate capacity) a specific resource is projected to supply during the highest net-load hours and reflects its intermittence and reliability

While many utilities apply higher tariffs to data centers than to other customers, such higher tariffs are generally still not high enough to cover the cost of new generation capacity[23]…

Specialized data center power rates still trail cost of new generation

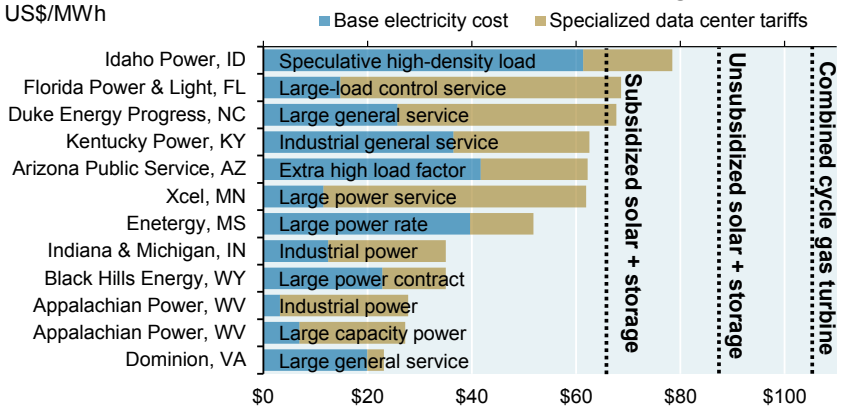

Source: Wood Mackenzie, Lazard, JPMAM, 2025

The food fight for new gas turbines

- Over the last 24 months, the cost of a new combined cycle turbine has risen from ~$1,200 per kW to $2,500 per kW with delivery times of 3 - 7 years
- Demand has impacted single cycle turbines as well whose costs have risen from $800 per kW to $1,500 per kW with delivery times of 3 -4 years
- Some companies are buying **Bloom Energy** solid oxide fuel cells as an alternative, which we will discuss in the March 2026 energy paper

…while at the same time, there's a shortage of gas turbines (see box) just as PJM and ERCOT project a surge of large load demand from data centers. The logjam isn't just turbines: transmission equipment (transformers, capacitor banks, medium and high voltage switches) now have lead times of 3- 5 years…

Global gas turbine orders

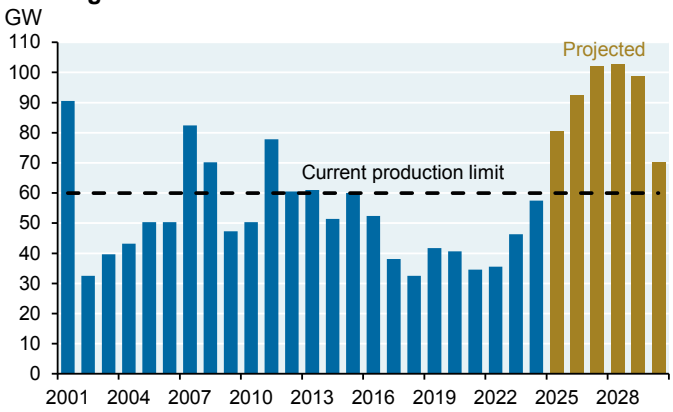

Source: Bloomberg News, October 2, 2025

ERCOT large load interconnection requests

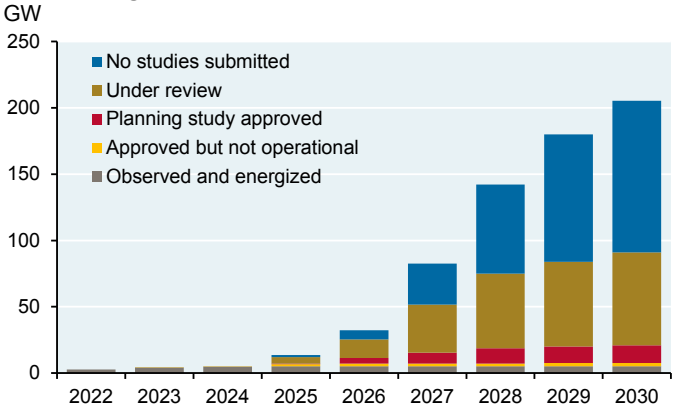

Source: ERCOT, October 2025

…and while the semiconductor ecosystem benefits from substantial tariff exemptions, that's not the case for US electrical equipment imports used to build generation capacity

Power generation vs semiconductor tariff exclusions

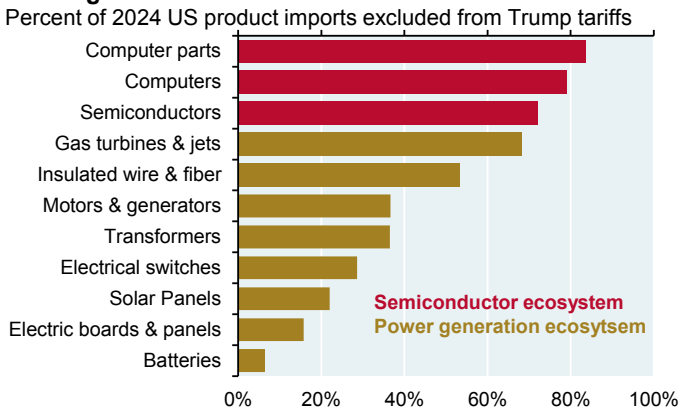

Source: USITC, White House, JPMAM, 2025

US power generation imports

Source: USITC, JPMAM, September 2025

[23] There are proposals in CA, GA, IN, MD, MI, OH and OR that require data centers to bear more of the costs associated with their additions to the grid

Some projections of data center power demand look daunting. OpenAI alone announced four partnerships that require 30.5 GW of new power[24]. To put this in context, that's around 75% of the peak level of nuclear GW completed over 5 years during the US nuclear era…just for one company in the AI ecosystem. And as shown earlier, the entire US only added 25 GW in 2024 after adjusting for intermittency/reliability of capacity. A power draw of 30.5 GW would also require ~300% of NVIDIA 2025 GPU shipments and more than 100% of current global high bandwidth memory supply[25].

Data center power demand from announced OpenAI partnerships, GW

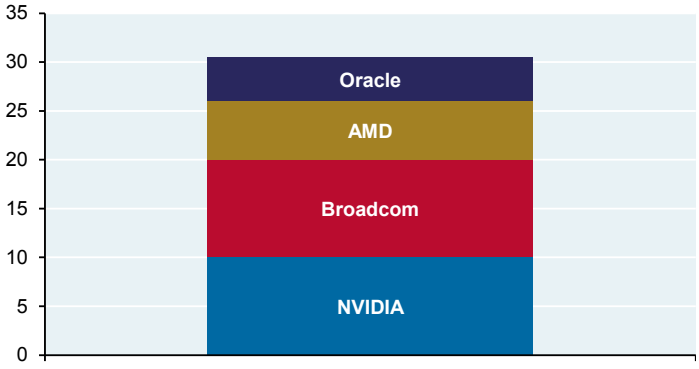

Source: OpenAI, 2025

US nuclear plants built by year of completion
GW, 5-year rolling sum

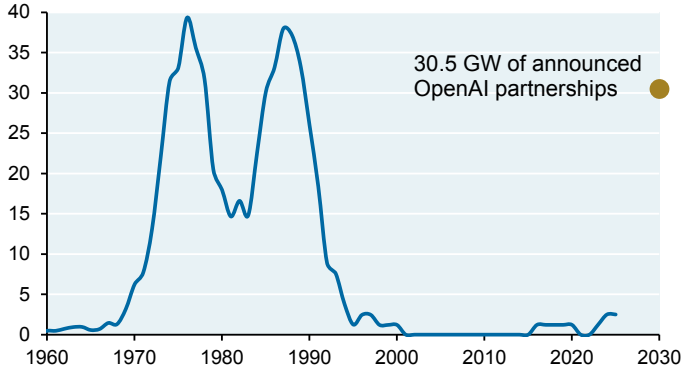

Source: Power Reactor System Database, JPMAM, September 2025

US live data center capacity
GW

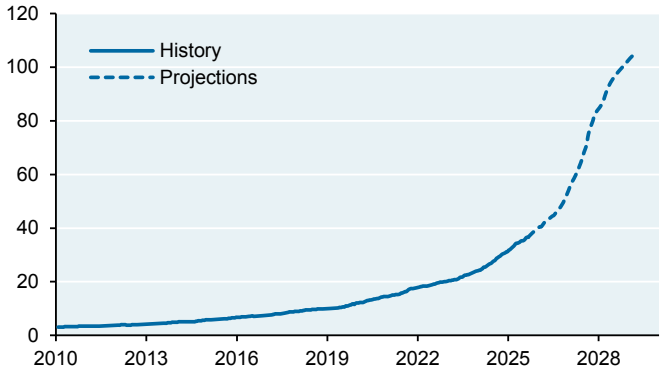

Source: Bridgewater AIA Labs, November 9, 2025

US data center additions by size and year
GW, including hyperscaler and retail/wholesale colocation facilities

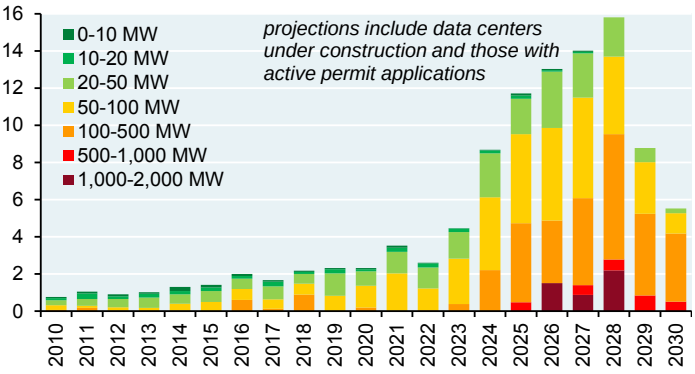

Source: Hiatt and Ryu (USC Zage Business of Energy Initiative), Q2 2025

Signs of power constraints

- In California, data center developers Digital Realty Trust and Blue Owl have projects in NVIDIA's hometown of Santa Clara that sit empty since utilities aren't ready to supply power
- Dominion Energy recently cited 5-7 years as completion timeframes for some data centers, and Amazon is in a legal dispute with a Berkshire Hathaway-owned utility regarding data center access to power
- Most data center construction is taking place in locations that NERC describes as already being at "elevated risk" regarding foundational grid reliability
- According to CBRE, data center demand is so high that developers are often able to lease projects years before completion with 74% of the current US construction pipeline already committed to tenants[26]

[24] A September 2025 article in *The Information* cited Altman as stating in an internal memo that the ultimate goal is 250 GW (!!!) by 2033, an amount equal to one third of peak US power consumption

[25] Bridgewater Daily Observations, November 10, 2025

[26] "*Data centers in NVIDIA's hometown stand empty awaiting power*", Bloomberg, November 10, 2025

As shown in the bar chart, BNEF projects that data centers may boost US gas demand by 3-4 bn cubic feet per day by 2030. This is a meaningful addition in the context of total US dry natural gas production of ~107 bn cubic feet per day in 2025. BNEF projects that new data centers added to the grid will be ~60% reliant on natural gas on a look-through basis through to the end of the decade.

Data center natural gas demand by 2030
Billion cubic feet per day

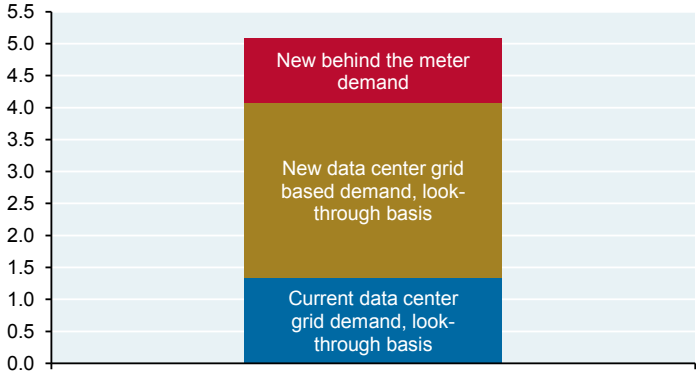

Source: BloombergNEF, JPMAM, November 2025

Incremental grid-connected data center fuel mix
Percent of generation

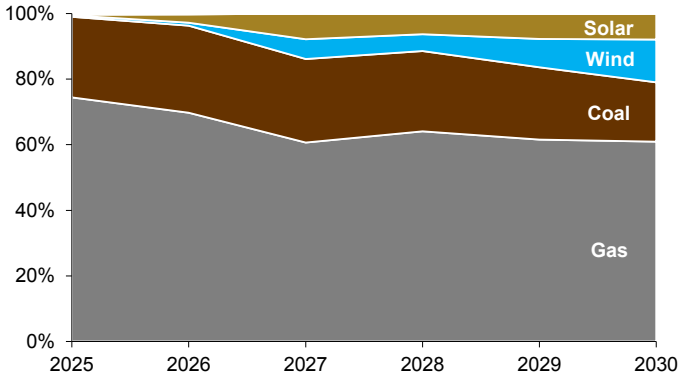

Source: BloombergNEF, JPMAM, November 2025

What about BYOG: bring your own generation? While only 1% of data centers currently rely on their own onsite generation, a survey from Bloom Energy indicates that 38% expect to rely *partially* on their own power by 2030, with 27% expecting to *fully* rely on their own onsite power. Long lead times for grid interconnection shown on the left may explain why. There are several options for onsite generation shown in the table, all with their own costs, lead times and performance characteristics.

Time elapsed from interconnection request to commercial operation date, Months

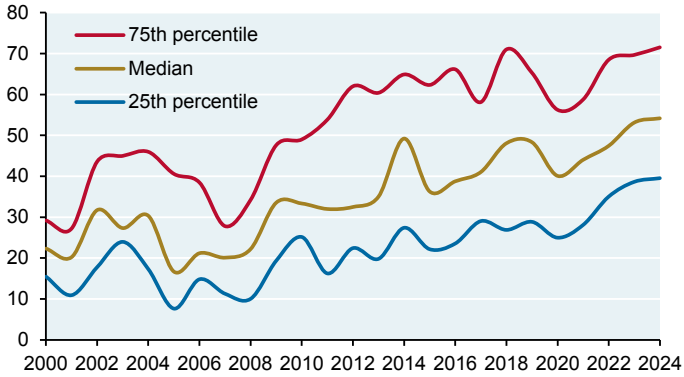

Source: LBNL, 2025

Data centers with onsite power, current & proposed

- xAI Colossus 1 and 2, OpenAI/Oracle Stargate, Microsoft Three Mile Island, Equinix, CoreWeave, Meta Socrates South

New entrants in gas fired power generation

While GE Vernova, Mitsubishi and Siemens dominate this industry, new entrants provide other options:

- Doosan Enerbility (S Korea) H class turbines
- Wartsila (Finland) modified ship engines
- Boom Supersonic (US) jet engines modified into aeroderivative gas turbines
- Caterpillar (US) gas turbines and reciprocating engines

Generation type	Size per unit MW	All-in capex $/kW	Lead time Months	Ramp rate Minutes	Land use MW/acre	Efficiency %	Est LCOE $/MWh
Aeroderivative gas turbines	30-60	1,500-1,800	18-36	10	30-50	35-40%	80-130
Industrial gas turbines	5-50	1,000-1,300	12-36	20-30	20-40	35-40%	65-110
Small combined-cycle gas turbines	40-100	1,800-2,500	18-36	30-60	20-30	40-55%	85-160
Medium-speed reciprocating engines	7-20	1,700-2,200	15-24	5-10	8-15	40-50%	80-150
High-speed reciprocating engines	3-5	2,200-2,800	15-24	5-10	5-10	40-50%	120-200
Fuel cells	0.325	3,000-4,000	3-4	Baseload	30-100	50-55%	100-200
H-Class combined-cycle gas turbines	600-1000	2,200-2,800	36-60	30-60	20-30	50-60%	100-200

Source: SemiAnalysis, December 30, 2025

BNEF also cited the risk that LNG demand will tighten US natural gas markets, and projects that data center gas demand on its own could flip the US from a supply surplus in 2030 to a supply deficit. But it looks like there's a looming supply glut in the global LNG market which could leave more gas for US domestic consumption.

Global LNG capacity utilization

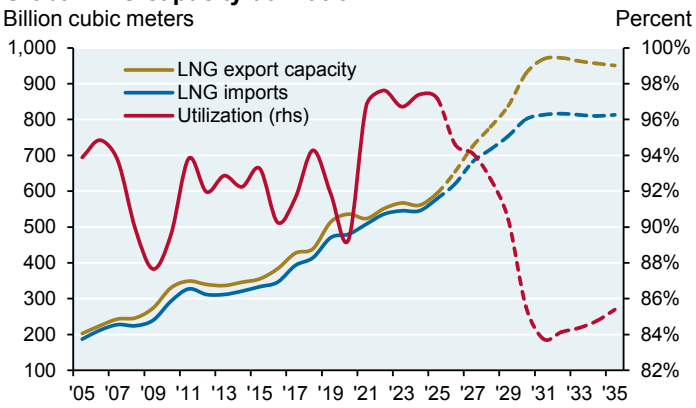

Source: OIES, October 2025

Data center behind the meter configuration, $ per MWh

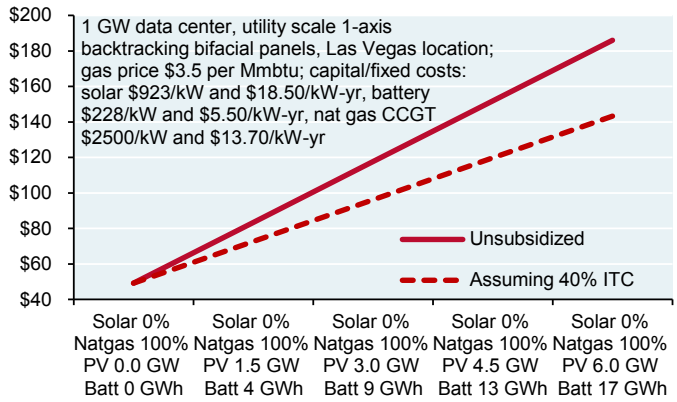

Source: NREL, JPMAM, EIA, December 2025

Could behind the meter solar + battery configurations provide baseload power substitutes for gas turbines? As we will discuss in our energy piece later this year, even when including a 40% investment tax credit, solar + battery installations are much more expensive per MWh than combined cycle turbines. The chart above on the right is an example of this analysis. **But for operators that are less price-sensitive on power and more focused on getting up and running**, solar + storage *might* work; a delay of even a few months to secure power can mean billions in lost revenue for AI infrastructure developers. Google's December 2025 acquisition of Intersect Power will be interesting to watch given Intersect's 6 GW of solar and 12 GWh of storage (existing + planned).

What about **nuclear**? I will discuss in the annual energy paper in March, but there are only 2-3 decommissioned nuclear plants that haven't been dismantled yet and which may be restarted (Three Mile Island-PA, Palisades-MI and Duane Arnold-IA). As for small modular reactors, **I remain skeptical since they're a reverse learning curve in motion**. When nuclear power was first commercialized, the learning curve experience from 1960 to 1980 led developers and engineers to increase their size to ~1 GW to spread large fixed capital costs over more MWh. There's not enough evidence suggesting that modularization can be done cost-effectively, and by that I mean $4-$6 mm per MW for SMRs or SMR costs per MWh that are less than 2x grid power.

The original learning curve led to larger nuclear plant sizes

Average capacity of global nuclear plants, gross MW

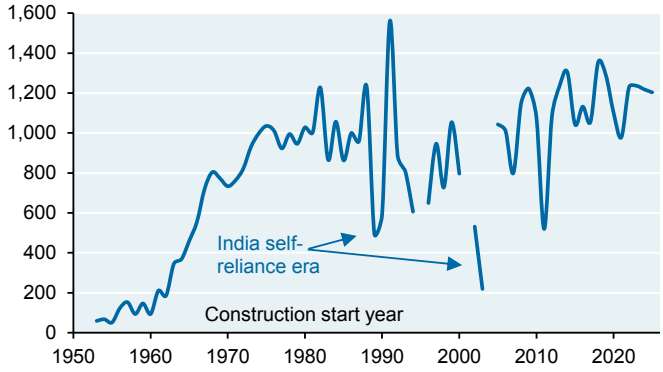

Source: Power Reactor System Database, JPMAM, September 2025

What about quantum computing as a means of reducing data center power demand, in which quantum computing is used as a replacement for GPUs for complex mathematical operations? This appears to be several years away at the earliest. I may cover this topic as well in the 2026 energy piece.

Some analysts point out slack capacity in US gas-fired power plants that in theory could allow for data centers to be added if they're willing to be curtailed. The chart on the left shows our aggregation of all combined cycle plants in the US (roughly 310 GW) and what their utilization rates were in 2024; most capacity factors clustered between 50% and 75%. The chart on the right shows combined cycle capacity factors from the EIA by region for 2016-2020 with similar results. In other words, perhaps more data centers could be added to the grid if they're willing to be curtailed at moments of peak demand, at which point they would reduce consumption or draw power from backup behind-the-meter diesel generators.

US natural gas combined cycle capacity factor in 2024
Share of total US natural gas combined cycle capacity, percent

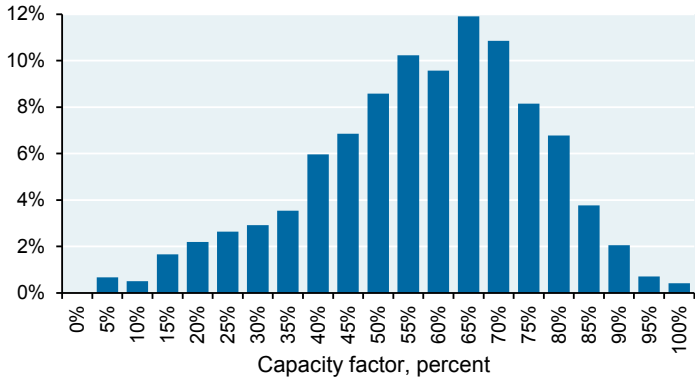

Source: EIA forms 923 and 860, JPMAM, 2024

Natural gas combined cycle capacity factor by region, plants built from 2008-2020, Average capacity factor, percent

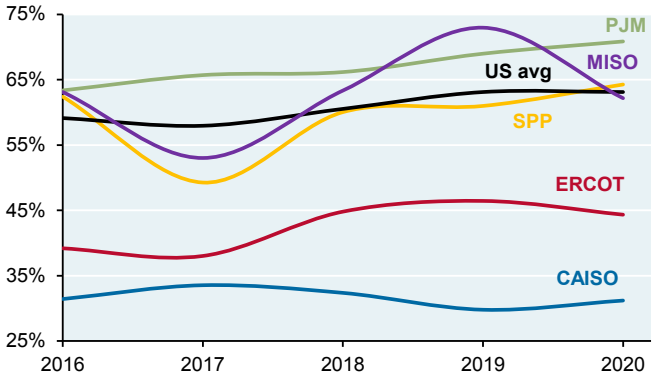

Source: EIA, May 20, 2021

There are two "buts". First, combined cycle plants generally need 7-8 weeks of downtime per year for maintenance and repair, so most plants already operate closer to max capacity. Second, some grid professionals don't believe that slack gas capacity or demand response is enough. Monitoring Analytics, the independent market monitor for PJM, wrote in November 2025 in very stark terms that new data center loads require new dedicated capacity and that demand response/flexibility is not enough[27]:

- "Large data center load additions have already had a significant impact on other customers as a result of higher transmission costs, higher energy market prices and higher capacity market prices…Continuing to simply accept interconnection of large data center loads that cannot be served reliably because there is not adequate dispatchable capacity is not a reasonable path forward and is not a solution of any kind"
- "Current capacity in PJM is not adequate to meet demand from large data center loads and will not be adequate in the foreseeable future. This is a simple factual issue. **The market solution is to establish a queue for large new data center loads which would not be interconnected until there is adequate capacity to serve them**. This solution to the issues created by the addition of unprecedented amounts of large data center load does not require a massive wealth transfer"
- "The assertion that large new data center loads can be demand side resources and do not require new capacity is a regulatory fiction…PJM does not have the authority to enforce reductions in load from emergency demand resources. Proposals that include the demand side option in place of adding actual generation capacity do not make the demand side option a mandatory condition for interconnection"
- "Implementing a load queue for large new data center loads is the only enforceable way to address the impacts of such loads and to require large new data center loads to pay for a significant part of the costs and risks that they would otherwise impose on other customers"

[27] Monitoring Analytics LLC, "*State of the Market Report for PJM*", November 13, 2025. On a related note, last October the Secretary of Energy proposed expanding FERC's authority over large loads to accelerate data center grid connection. The proposals included the stipulation that large loads would need to be curtailable for system operators to be able to conduct system adequacy planning

Some academics and consulting firms disagree with Monitoring Analytics. They argue that by combining flexible grid connections with behind the meter capacity, "data centers can reach full operation years sooner while maintaining reliability and improving affordability for all customers"[28]. I'm not sold on this argument. Google-funded analyses like these can be influenced by (a) hyperscaler and data center companies simply seeking faster data center grid interconnections, and (b) renewable energy advocates fundamentally opposed to development of natural gas resources under any circumstances. The ultimate arbiters will be the utilities and the Independent System Operators who run the grids and who are accountable to customers.

More to come in our 2026 energy paper but after 20 years of flat electricity demand, the US now needs to meet new demand from data centers, EVs and electrification of commercial, industrial and residential heating. While US electricity generation grew by 2%-4% each year from 1950 to 2006, during that era 70%-90% of new capacity additions came from a small number of large coal, nuclear and gas plants. Whether the same pace can be maintained today is another question, particularly given associated shortages of skilled energy labor.

One country that is not having much trouble adding more power generation: **China**, which we discuss next.

Generation and capacity additions since 1950

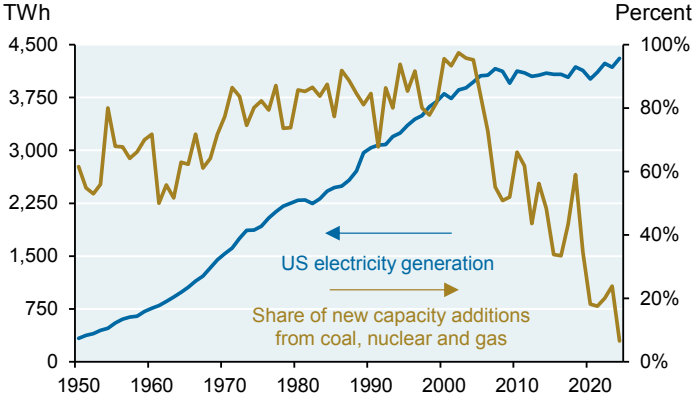

Soruce: EIA, JPMAM, 2025

Projections of a US power supply-demand gap

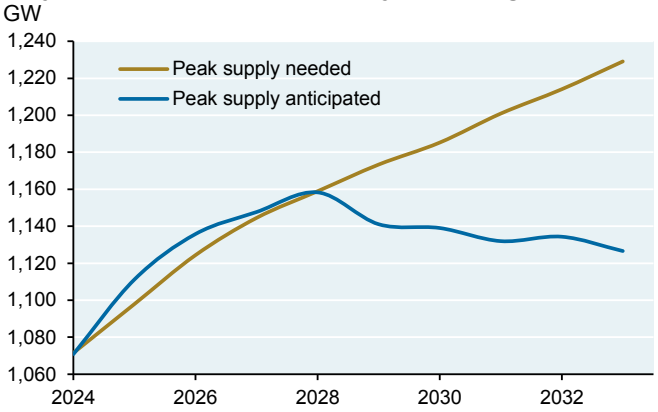

Source: NERC, Schneider Electric, 2025

Difficulties in finding energy employees

Percent of survey respondents

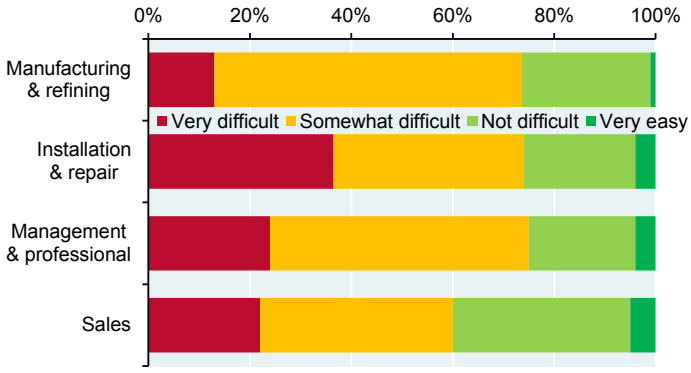

Source: Goldman Sachs, July 23, 2025

Electricity generation

TWh

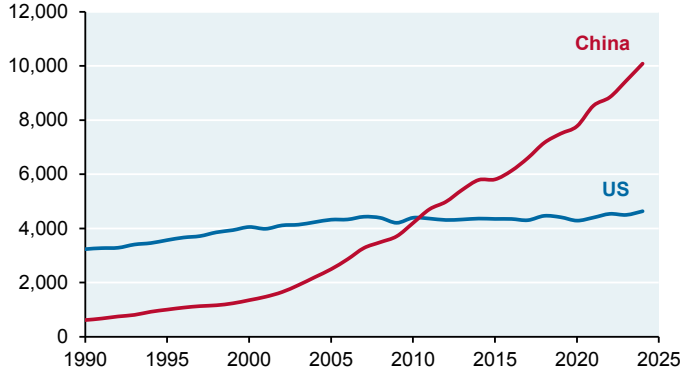

Source: Energy Institute, JPMAM, 2025

[28] "*Flexible Data Centers: A Faster, More Affordable Path to Power*", Camus Energy and Princeton University Zero Lab, December 2025

WCGW #3: China scales the moat with its own lithography-semiconductor technology…when, not if

Let's start with a basic premise: **China has been steadily climbing the innovation food chain across a range of industries**. The first chart shows China breaking into the top ten on innovation, while the second shows China and the US converging with respect to the complexity of exported products. China has also been outpacing the rest of the world on clean energy patents, AI patents and nuclear engineering research.

China breaks into top 10 most innovative economies
Global Innovation Index, country rank, (1 = highest innovation)

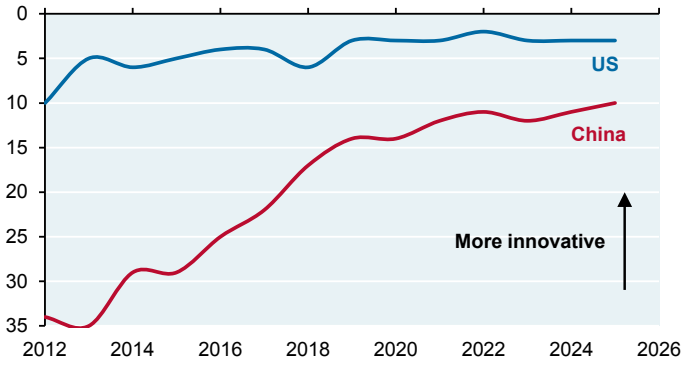

Source: World Intellectual Property Organization, 2025

US and China export complexity
Country rank, (1 = highest complexity)

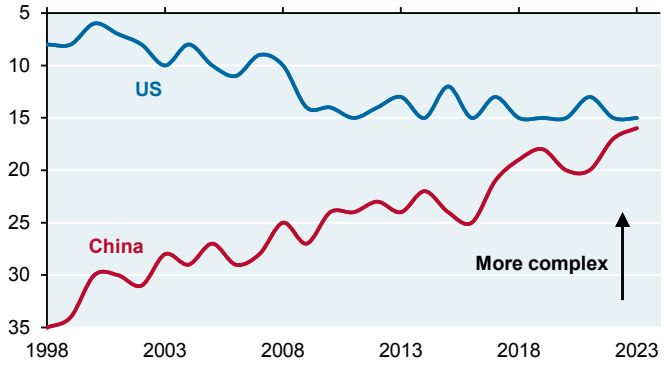

Source: Harvard Growth Lab, JPMAM, 2025

Annual clean energy patent applications by country
Thousands

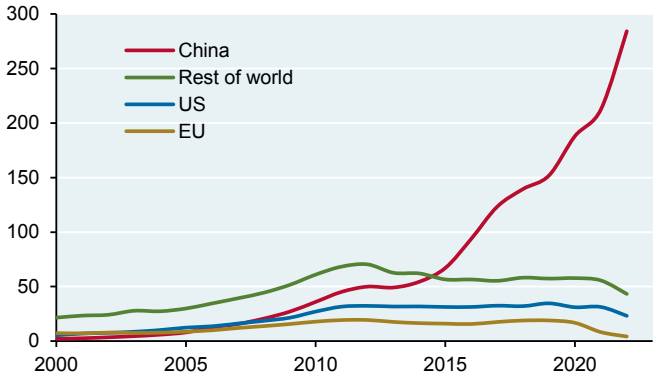

Source: IRENA, JPMAM, 2025

AI patents by geographic area
Percent of world total AI patents

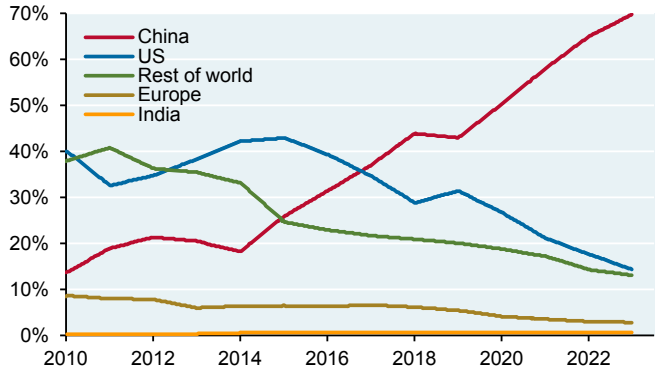

Source: Stanford AI Index Report, 2025

Nuclear science and engineering research by country
Count of publications in top 10% of most-cited publications

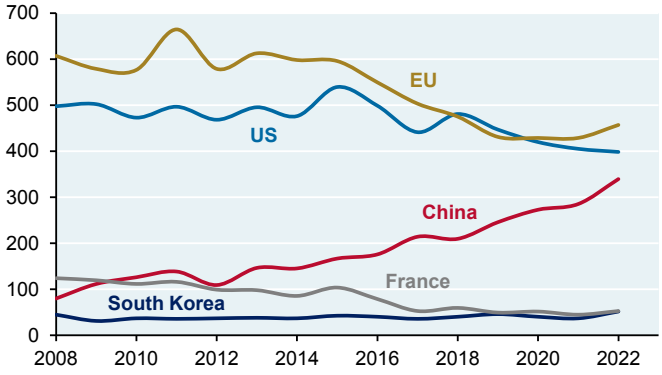

Source: "How innovative is China in nuclear power", Stephen Ezell, ITIF, 2024

China's willing innovation partners: US and European antitrust regulators

EU regulators and Biden's FTC (Lina Khan) signaled that they would block Amazon's acquisition of iRobot, after which the deal was scrapped. Since then the company has laid off half of its staff and just filed for bankruptcy. Even the NYT lamented that a Chinese supplier will now take control of the company. In contrast to the US and Europe, Chinese robotics companies receive enormous government support: China's robotics hubs receive 20x to 40x more funding than US counterparts. iRobot's heavily subsidized Chinese rivals Ecovacs and Roborock will benefit from iRobot's demise

China's ability to scale the moat may rely in equal parts on innovation and brute force power capacity. China is adding every kind of electricity generation, electrifying its energy consumption at the fastest pace in the world, and is also the world leader in deployment of high voltage direct current lines. Power generation changes since 2019: **China +2,500 TWh, the US +221 TWh and Europe -110 TWh**.

China electricity generation by source
Terawatt-hours

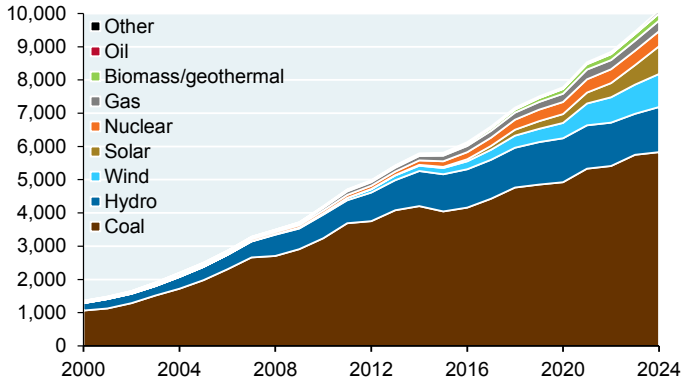

Source: Energy Institute, JPMAM, 2025

Electrification: China and Japan power ahead
Electricity share of useful final energy

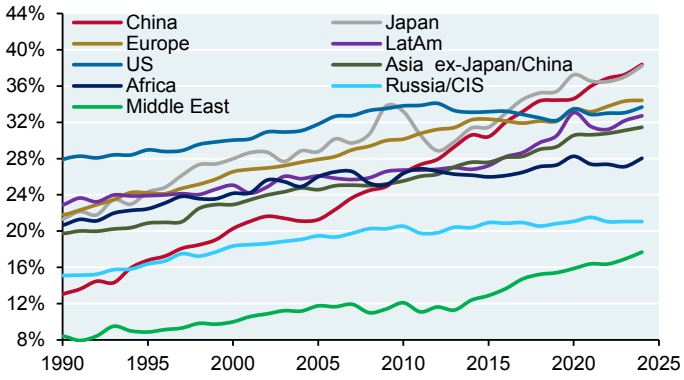

Source: Energy Institute, EIA, JPMAM, 2025

Many of these achievements are driven by China's support for its national champion industries. While the US government's investment in MP Materials is a belated effort to narrow this gap (see box), China is by most accounts the most mercantile country in the world, including a #1 ranking in the ITIF Mercantilist Index and in an analysis by CSIS regarding industrial policy spending by country. Look at the enormous share of Chinese corporate profits that government subsidies represent on the right.

Industrial policy spending in key economies
Percent of GDP

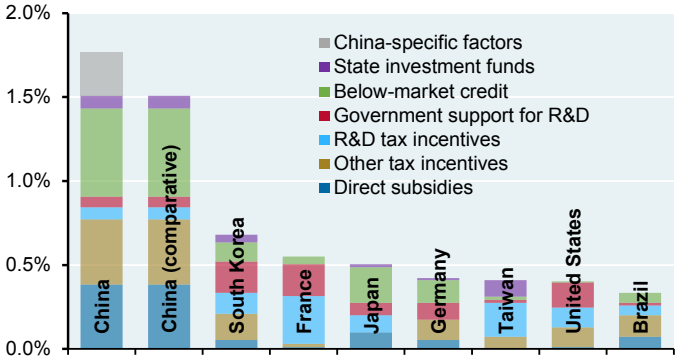

Source: CSIS, 2022

China direct subsidies for listed firms by subsector
Share of profits, 2017-2020

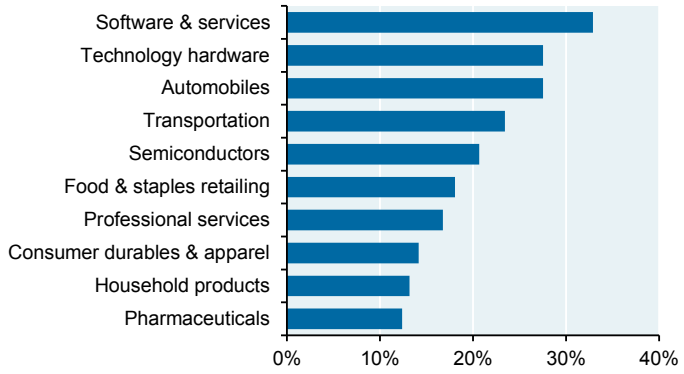

Source: CSIS, 2022

MP Materials vs Intel: two different approaches to US industrial policy

US gov't support for the critical minerals company **MP Materials** entails an anchor of demand: 10 years of price support for neodymium and praseodymium at $110 per kg. In other words, should market prices decline below $110 per kg, the US government would make up for the rest. Furthermore, the Department of Defense has committed to purchase 100% of the output from MP's second magnet facility which is still in the pre-construction phase.

US gov't support for **Intel** is different: there's no anchor of gov't demand. Of $12 bn in grants and subsidies provided to Intel by the Chips Act, only $3 bn were deployed. The remaining $9 bn was converted by the Trump Administration into a cash investment in Intel with warrants to purchase more shares if Intel sells > 50% of its manufacturing arm. As a result, the deal entails no new demand creation for Intel's leading edge manufacturing capacity. The implicit presumption is that Intel's primary problem is inadequate liquidity and access to capital…but its challenges may run deeper than that. NVIDIA and Softbank also pledged capital but without committing to using Intel manufacturing given its prior defect issues with 13th and 14th Gen Core CPUs. Will the US gov't take the next step and compel US companies like Apple, NVIDIA, Qualcomm, Broadcom, etc to use Intel manufacturing? That remains to be seen.

China, sanctions and the NVIDIA-TSMC-ASML moat

China's AI development still relies heavily on Western chips acquired via official and black market sources. China also lags the US and its allies in production of cutting-edge chips, particularly those less than 14 nanometers (nm). The key question: how quickly might this change?

Origin of chips used to train Chinese LLMs
Number of models

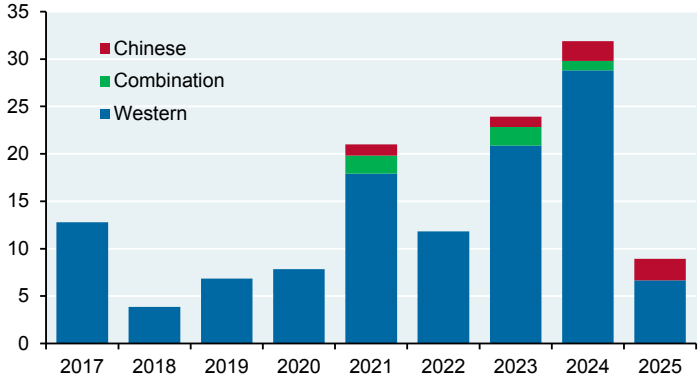

Source: EpochAI, July 2025

Installed semiconductor capacity by market and node
Millions of 200 mm equivalent wafer starts per month

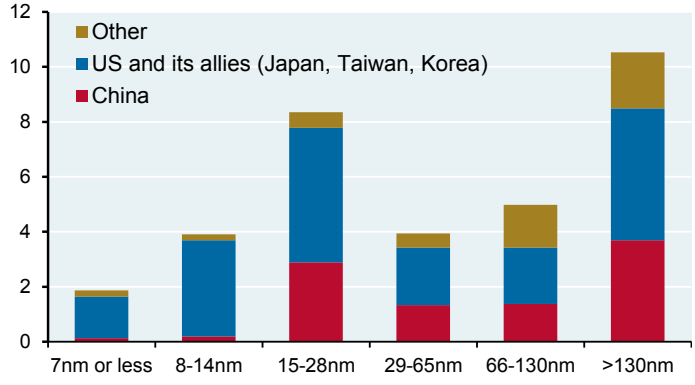

Source: Bernstein Societe Generale Group, November 2025

The impact of sanctions and the long-term Huawei cluster challenge to NVIDIA

There are over 200 Chinese individuals and 620 Chinese entities in the US Treasury sanctions database. Biden's Department of Commerce described these sanctions as weakening China's domestic semiconductor ecosystem and its ability to manufacture advanced chips, and stated they were the strongest controls ever enacted by the US. As for lithography, ASML has been prohibited for many years from selling its extreme ultraviolet light (EUV) machines and even some less advanced DUV (deep ultraviolet technology) technology to China.

Sanctions have accelerated China's push to develop its own semiconductor ecosystem. Its progress can be inferred from a variety of statements, actions and consequences:

- The Chinese government has pledged unprecedented and nearly unlimited financial support for its domestic semiconductor industry. China launched its 3rd National Integrated Circuit Industry Investment Fund with $48 billion, more than the first two phases combined and with a 15-year time horizon rather than 10. The reported goal: resolve "chokepoint technologies" like semiconductor ecosystems
- In April 2025, Huawei announced the launch of its 910C chip to fill the void of NVIDIA's H20 chip, which the US had restricted at the time. The US has since unbanned exports of the H20, a modified version of NVIDIA's H100, but Beijing has begun restricting Chinese firms from using the modified chip. In December 2025, the US approved NVIDIA's H200 chip for sale in China again
- Huawei is committed to becoming a fully verticalized company via its various affiliates and partners, from wafer fab equipment (SiCarrier) to chip design (HiSilicon) to manufacturing/advanced packaging, including proprietary high bandwidth memory (Peng Chip, Fujian Jinhua). Fabs built by SiCarrier could exceed SMIC manufacturing capacity as early as next year
- Huawei's Ascend 910C has reportedly improved its foundry yield to ~40% from 20% in September 2024, and now accounts for 75%+ of China's total AI chip production. Huawei aims to improve its new node foundry yields to 60% to match industry norms[29]; for a mature node, foundry yields would need to be 75%-85%
- Chinese scientists are working on methods to generate EUV laser light at wavelengths of 13.5 nanometers, and a December article in Reuters confirmed that a team of former ASML engineers in Shenzhen completed a working prototype of an EUV machine with production targeted for 2028-2030[30]

[29] Digitimes Asia, February 25, 2025

[30] "*How China's award-winning EUV breakthrough sidesteps US chip ban*", SCMP, January 19, 2025; "*How China built its Manhattan Project to rival the West in AI chips*", Reuters, December 17, 2025; "*ASML employee who stole chip secrets went to work at Huawei*", Asia Financial, October 24, 2023

NVIDIA still leads Huawei on a number of metrics on a per-chip basis. The NVIDIA advantage in the last column would rise further if we were to compare Huawei's 910C NPU with NVIDIA's forthcoming R100 GPU.

CHIP BASIS	Huawei 910 C	NVIDIA B300	NVIDIA vs Huawei	Comments
Computing power, teraflops, BF16 dense	780	2,250	2.9x	1 teraflop = 1 trillion floating pt operations
High bandwitdh memory (HBM), gigabytes	128	288	2.3x	High speed stacked memory chips
Memory speeds				
HBM bandwidth, terabytes per second	3.2	8.0	2.5x	
Intranode (between GPUs/NPUs in a server), GB/s	700	1,800	2.6x	
Power, watts, thermal design power only	310	1,100	3.5x	
Power, kW, all-in-power at the server level *	1.455	2.014	1.4x	
All-in power efficiency, watts per teraflop	1.87	0.90	0.5x	This is the key metric on power draws

* All-in power calculated using total power draw of Huawei CloudMatrix and NVIDIA Blackwell Ultra servers (Schneider Electric reference design), incorporating networking/cooling power draws in addition to thermal design power. Intranode bandwidth is also known as scale-up bandwidth

That's why Huawei aims to compete at the cluster level rather than at the chip level. The next table shows Huawei's roadmap, with its clusters made up of 384 to over 15,000 individual chips. Huawei asserts that its product offerings are already on par with NVIDIA at the cluster level; that its Atlas 950 SuperPod slated for 2026 will deliver 6.7x more computing power, 15x more memory and 62x more intranode bandwidth than NVIDIA offerings planned for the same year (Rubin NVL144); and that Huawei will be ahead on all fronts in 2027.

CLUSTER BASIS	Huawei			NVIDIA		
	Atlas 900 A3 SuperPoD*	Atlas 950 SuperPoD	Atlas 960 SuperPoD	GB300 NVL72 Blackwell Ultra	NVL144 Vera Rubin	NVL576 Rubin Ultra
Release	March 2025	Q4 2026	Q4 2027	Current	2026	2027
Chip cluster (number and type of chips)	384 Ascend 910C	8,192 Ascend 950DT	15,488 Ascend 960	72 B300s	144 Rubin GPUs	576 Rubin GPUs
Computing power, dense exaflops with mathematical precision level	**0.30 BF16**	**8 FP8, 16 FP4**	**30 FP8, 60 FP4**	**0.18 FP16, 0.36 FP8, 1.08 FP4**	**1.2 FP8, 3.6 FP4**	**5 FP8, 15 FP4**
High bandwitdh memory (HBM), terabytes	49	1,152	4,460	21	21	365**
Memory speeds						
Aggregate peak HBM bandwidth, petabytes/second	1.2	32.7	148.9	0.576	1.4	4.6
Intranode (between GPUs/NPUs in a server), PB/s	0.268	16.3	34.1	0.13	0.26	1.5
Internode (between servers), terabytes/second	38.4	819 [e]	NA	14.4	28.8	115.2
All-in power, kW	599	12,000 [e]	NA	145	225 [e]	600 ***
All-in power efficiency, watts per teraflop	**2.00 BF16**	**1.50 FP8, 0.75 FP4**	**NA**	**0.81 FP16, 0.40 FP8 0.13 FP4**	**0.19 FP8 0.06 FP4**	**0.12 FP8 0.04 FP4**

Source: Huawei press releases, Gavekal Technologies, SemiAnalysis, TheWhiteBox consulting, JPMAM. December 2025. *Also known as CloudMatrix. **Includes 144 TB of GPU memory and 221 TB of CPU memory. *** Based on 4.16 kW per package of 4 chiplets each. NVIDIA processors are GPUs (graphics), Huawei processors are NPUs (neural). Intranode and internode speeds are quoted on a bidirectional basis. 1 exaflop = 1,000 petaflops = 1,000,000 teraflops. [e] = Rough estimate of future capabilities. BF16, FP16, FP8 and FP4 refer to mathematical precision in the number of decimals; lower precision entails faster computing

OK, but: system performance is more than just the amount of exaflops. For Huawei to succeed, its entire ecosystem (low-latency interconnect via optical networking, HBM bandwidth/capacity and training software) would have to perform exceptionally well if hundreds or even thousands (!!) of chips are joined together. Unlike NVIDIA and other Western servers in which copper wiring is the predominant scale-up gear, Huawei's scale-up is entirely optics-driven. This would be prohibitively expensive in the West but China is less concerned about today's profits and is forced to rely on optical networking due to the physical size of its servers.

All the extra hardware and power is expensive. Huawei's CloudMatrix 384 product is estimated to cost 2x more, and its 2026 cluster products may **require 8x-12x more power per teraflop** than NVIDIA for comparable levels of precision. While we cannot yet make informed power efficiency estimates for the Atlas 960, I suspect that Huawei will still trail NVIDIA materially on that front as well. Then there's the real estate: ~7 racks of Ascend 910C NPUs are needed to match the power of one rack of NVIDIA B300 GPUs. And while we mostly focus on NVIDIA here, Google's Ironwood Pod is also much more efficient than Huawei's: Ironwood offers 42 exaflops at FP8 precision, comprised of 9,216 TPUs with cluster-level power efficiency of just 0.23 watts per teraflop. **As a result, Huawei compute clusters may only be used for now by Chinese state-owned companies who receive subsidies, and by Chinese companies pressured or explicitly required to purchase them.**

While ByteDance bought more NVIDIA chips than any other Chinese company in 2025, **China has now blocked ByteDance and other Chinese companies from using them in new state funded data centers for inference purposes[31].** The oversupply of inference chips shown below is a byproduct of China stockpiling NVIDIA chips via official and black markets, ***and*** local chipmakers catching up on quality. Chinese chips have higher performance scores than NVIDIA H20 chips (but not NVIDIA H200 chips approved for sale in December 2025)[32].

As another sign of declining China reliance on the West, ASML announced that its China share of revenues would fall to 20% from 25%-40% in 2024/2025.

Milestones to watch: how quickly will Chinese companies adopt chips made by Cambricon, one of China's NVIDIA chasers which is backed by the Chinese Academy of Sciences?

China: a large surplus of inference chips
Exaflops

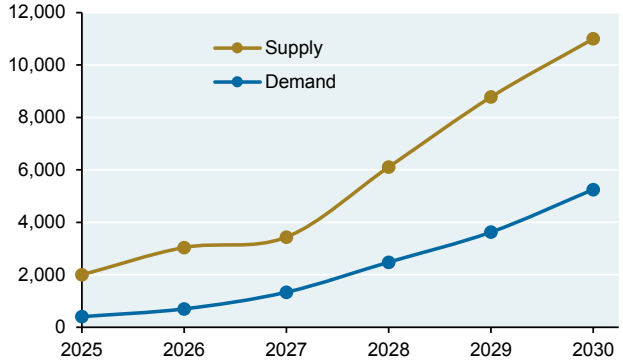

Source: Jefferies, November 2025

AI chip performance
Total processing performance (TPP) score, operations/second, trn

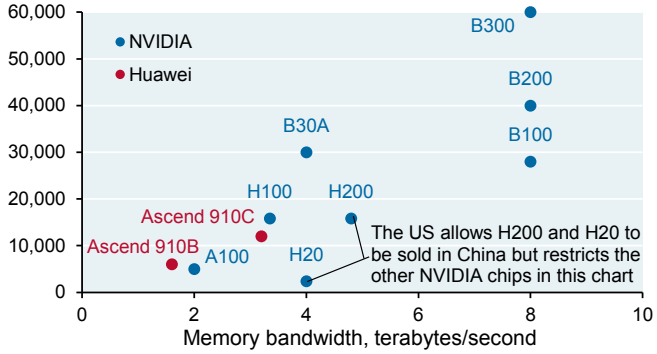

Source: Institute for Progress, Rand, The Information, JPMAM, Dec 9, 2025

ASML revenue share by country
Percent of total net sales

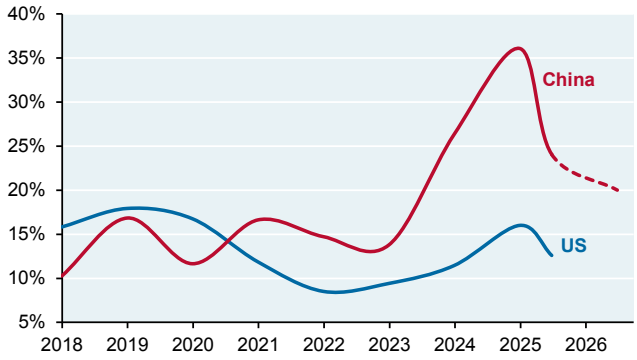

Source: Bloomberg, ASML, JPMAM, June 2025

Market cap of China's NVIDIA-chaser, 2023-2025
Index (100 = January 1, 2023)

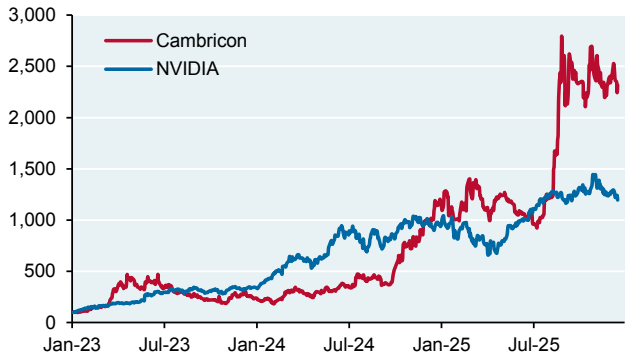

Source: Bloomberg, JPMAM, December 17, 2025

[31] "*Signaling Confidence in Its Domestic Industry, China Bans Foreign AI Chips in State-Funded Data Centers*", FDD, Nov 10, 2025; and "*China Is Slowly but Surely Breaking Free From Nvidia*", The Information, Nov 26, 2025

[32] In the XY scatterplot, TPP scores on operations per second (y axis) are typically used to assess model training capabilities while memory bandwidth in bytes per second (x axis) are used to assess inference capabilities

The bottom line. The US still has an enormous advantage regarding the actual user experience since its moat is based on training budgets and inference-time compute rather than data science benchmark scores, a metric on which China is catching up[33]. Consider the gap in actual 2023-2025 hyperscaler spending (US $694 bn vs China $124 bn), or the gap in compute capacity added in 2025 (NVIDIA on its own at 18 zettaflops vs China at < 1 zettaflop[34]). Or consider the chart on the right: using the US Commerce Dep't formula that compares chips normalized across precisions that are used to define export control thresholds for China, the amount of compute the US will deploy is much larger than China, not accounting for US chips shipped to China like the H200.

These metrics are better reflections of the real moat since most AI progress is driven by two factors: model size and the ability to deploy more computing resources on inference tasks (scaling laws). In other words, China's ability to scale inference-time compute is nowhere near current US levels. And remember, NVIDIA is not standing still. The table on page 30 projects that new NVIDIA Rubin GPUs may offer another 70% improvement in power efficiency. In that regard, Huawei servers may not have many current takers outside China itself given their size and power demands.

All things considered, China has a long way to go to scale the moat since it's one to three generations behind depending on how you measure it, but may be further along than investors realize. **But**…given China's fundamentally different approach than the West and its Asian allies (industry subsidies, a willingness to absorb higher costs in the interest of national security and domestic supply chains, a no-holds-barred buildout of power generation, transmission and distribution, and plenty of industrial espionage[35]), **China's gradually declining technological reliance on Taiwan and TSMC is important to consider, a topic which we discuss next**.

US vs China AI model performance
LMSYS Chatbot Arena score

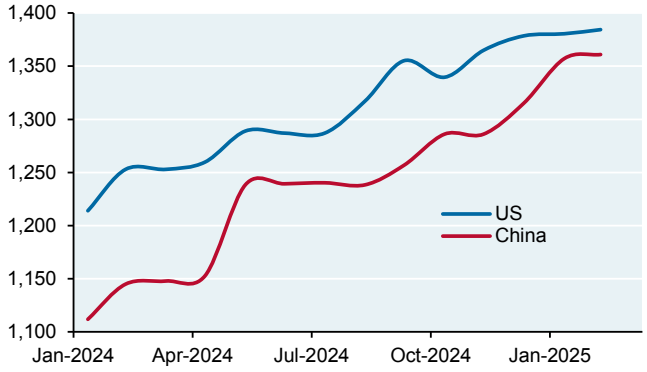

Source: "Artificial Intelligence Index Report", Stanford HAI, 2025

Forecasted annual AI compute production
Total processing performance (TPP), terabit operations/second, bn

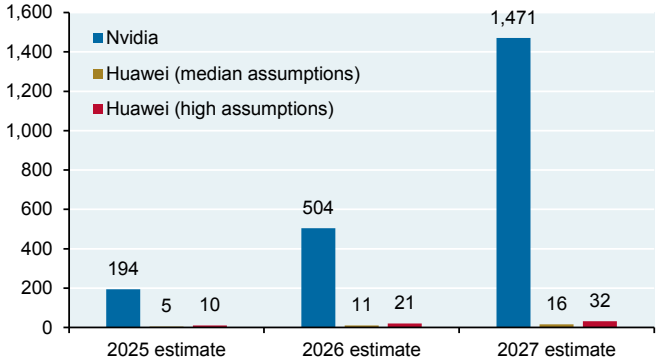

Source: Council on Foreign Relations, December 15, 2025

[33] **How has China narrowed the performance gap in the first chart?** Some analysts believe that given smaller training runs and model sizes and yet achieving close benchmark results, China is heavily relying on distillation of US model outputs. In other words, China trains their models by using US models in parent-teacher mode. According to the Trump Administration AI czar, OpenAI has evidence that DeepSeek violated its terms of service by querying OpenAI's model at scale and using its responses as training data to improve DeepSeek's own model

[34] "*The AI Value Chain: the US has chips but no power, China has power but no chips*", Bernstein Research, Stacy Ragson, November 24, 2025

[35] ASML won an $845 million judgment in 2019 against a former Chinese engineer accused of stealing trade secrets, but the defendant filed for bankruptcy and continues to operate in Beijing with Chinese government support. **The October 17, 2024 Eye on the Market included a detailed analysis of Chinese industrial espionage.** The main conclusions: China is conducting industrial espionage on an unprecedented scale; its businesses are now seen by Western intelligence agencies as inseparable from the Chinese gov't; its espionage efforts include infiltration of US transport, telecom, water and electricity networks; the breadth of China's cyberattacks are broadening; and its espionage efforts include hacking of private defense companies working with the Pentagon

WCGW #4: China and Taiwan

It's no secret that the world is highly reliant on Taiwan for advanced semiconductor chips; the nuanced part is estimating the degree. The chart on the left uses a simple approach by volume, while the chart on the right focuses on the US and includes the value of imports from Taiwan directly and through third party intermediaries. In the third and fourth charts we show US imports of semiconductors and computer parts by country, but the Taiwan figure in the third chart is understated since USITC data does not track the flow of Taiwanese chips to other countries for assembly and packaging.

The bottom line: global and US reliance on Taiwanese chips dwarf European reliance on Russian energy before Russia's invasion of Ukraine, as one example for geopolitical comparison.

Regional reliance for critical goods

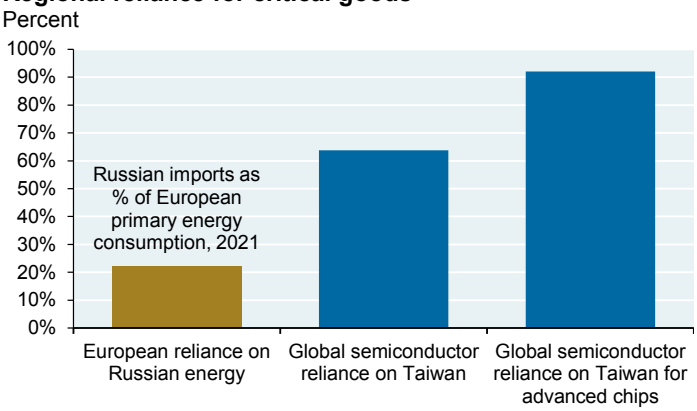

Source: BP Statistical Review, ROC Taiwan, Global Guardian, 2024

US semiconductor reliance on Taiwan

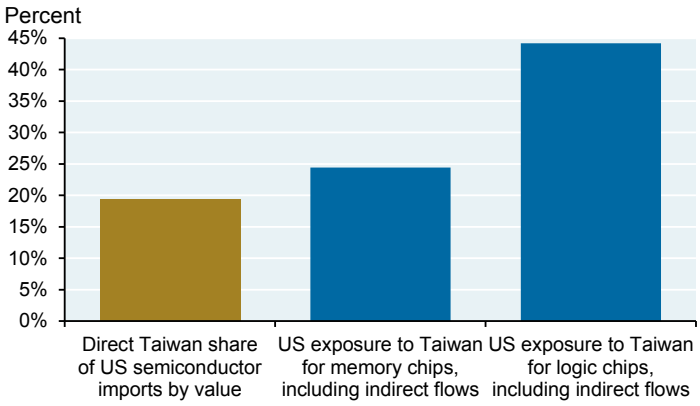

Source: USITC, November 2023

US semiconductor related imports

Billions, US$, monthly annualized

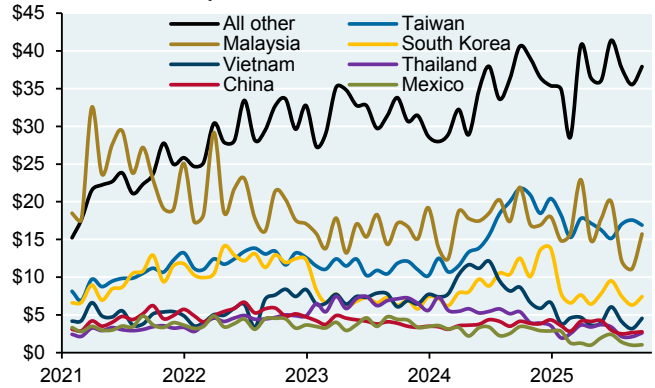

Source: USITC, JPMAM, September 2025

US AI related computer imports

Billions, US$, monthly annualized, includes large comp, GPUs & parts

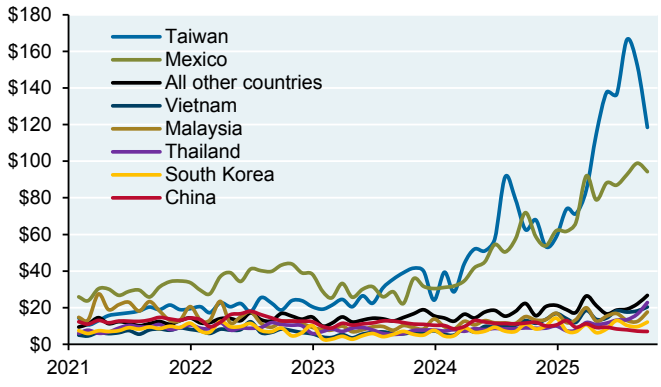

Source: USITC, JPMAM, September 2025

Biden's Taiwan moment, predictions and the 2024 election

In 1979 the US recognized the People's Republic of China and ended diplomatic relations with the Republic of Taiwan, ushering in the era of "strategic ambiguity" regarding the US position on Taiwan's independence. Biden was not the first US president to depart from the strategic ambiguity doctrine regarding whether the US would defend Taiwan if attacked; George W Bush did the same in an April 2001 interview. But the White House and the State Department immediately walked back Bush's comments by claiming that US policy had not changed, and later that same day, Bush reiterated a commitment to the One-China policy and advised Taiwan against declaring independence. In contrast, Biden repeatedly stated that the US would militarily defend Taiwan and left poor Jake Sullivan to try and clean up the mess. These repeated departures from the strategic ambiguity doctrine were part of the reasoning for my prediction in our January 2024 Outlook that Biden would drop out of the race before the Democratic convention due to health reasons.

How unique is TSMC? Very. How reliant is the US on TSMC? Very, at least until the end of the decade

Eight of the 10 largest companies in the world by market cap depend in large part on TSMC supply (the only exceptions are TSMC itself and Saudi Aramco). For these companies, more than one third of their combined $2 trillion in revenue comes from hardware that uses TSMC products[36]. The digital economy, automotive economy and industrial production economy rely on TSMC, and without it the world economy would sputter.

TSMC's ecosystem is anchored by three major science parks in Hsinchu, Central Taiwan and Southern Taiwan. Hsinchu accounts for ~60% of TSMC output value and hosts more than 150 semiconductor-related companies, including integrated circuit design firms, foundries and testing/packaging firms. The maximum distance between any two TSMC fabs in Taiwan is less than three hours by car, enabling coordination and resource sharing. As we discussed in the September 2025 EoTM, Taiwan's National Development Fund is still the largest shareholder in TSMC with 6%-7% and a seat on the company's board, and the government supports TSMC via subsidized water and electricity prices, tax credits and creation of industrial parks and research centers.

As TSMC founder Morris Chang suspected, **TSMC's Arizona fab is much less efficient (so far) than Taiwan based facilities.** Also: all TSMC chips made in Arizona **are still sent back to Taiwan** for packaging, dicing and testing until Amkor builds its US facility to support TSMC Arizona fabs, targeted for 2028.

TSMC 5 nm wafer economics, 2025

Item	Unit	Taiwan Fab 18 Phase 1-3	US Fab 21, Phase 1
Fab capex	US$	27,000,000,000	14,380,000,000
Wafer fab equipment	US$	22,950,000,000	8,600,000,000
Non-wafer fab equipment	US$	4,050,000,000	5,690,600,000
Installed wafer starts per month	Thousand	90	24
Raw materials	$/wafer	1,520	3,040
Labor cost	**$/wafer**	**1,800**	**3,600**
Utilities	$/wafer	630	630
Consumables	$/wafer	760	760
Overhead	$/wafer	471	803
Depreciation per wafer	**US$**	**1,500**	**7,289**
Variable costs per wafer	US$	5,181	8,833
Total cost per wafer	US$	6,681	16,123
Gross margin per wafer	US$	10,819	1,377
Gross margin per wafer	**%**	**62%**	**8%**

Source: Semianalysis, November 30, 2025

Will Taiwan's N-1 rule become an N-2 rule?

TSMC Arizona Phase 1 fab will initially produce 4/5 nm chips. Taiwan's gov't has been applying an N-1 rule, allowing TSMC to only export technology that's at least one generation behind its leading-edge fabrication process. Taiwan is reportedly considering a shift to an N-2 rule instead. In Taiwan, TSMC is about to begin production of 2 nm chips, while its Arizona Phase 2 fab is set to produce 3 nm chips in 2027

The Arizona project's core purpose is to rebuild part of the supply chain that the US surrendered decades ago, even at higher cost. US Commerce Secretary Lutnick set a goal for the US to onshore 40% of semiconductor demand by the end of Trump's term in 2028. By that time, global advanced node demand is expected to grow from 850k wafer starts per month (wspm) in 2024 to 1.4 mm. For the US share, we estimate that ~75% of TSMC leading edge wafer revenues come from US customers, implying US advanced node demand of ~1 mm wspm in 2028. If the estimated production levels in the chart are achieved, the US might be 30%-35% self-sufficient in advance node production by the end of the decade but would still be highly reliant on Taiwan.

US might reach 30%-35% of advanced node (< 5 nm) production by 2028-2030

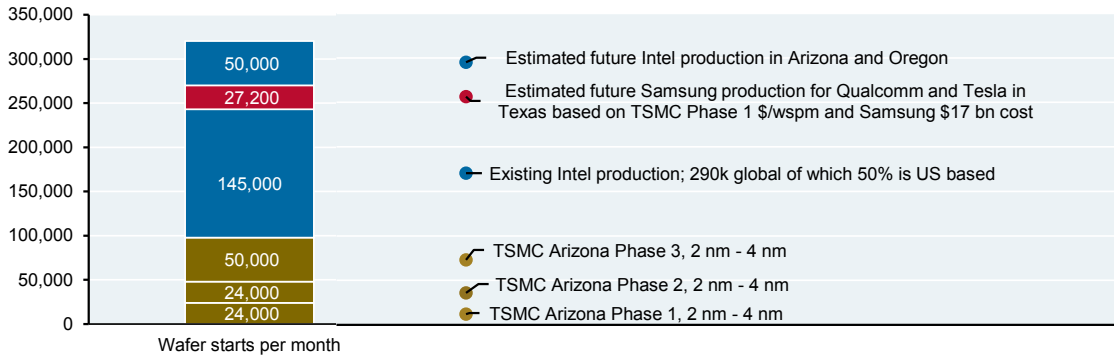

Source: JPMAM, 2025

[36] "*TSMC Overseas Fabs*", Semianalysis, November 30, 2025

How vulnerable is Taiwan to a naval blockade or quarantine? Very

Taiwan is one of the most import-dependent countries in the world: **90% of its primary energy consumption relies on fossil fuel imports**. In what might be the most ill-advised energy policy shift in history, Taiwan's share of electricity generation from home-grown nuclear power fell from 50% in the 1980's to just 5% today.

Instead, imported LNG now accounts for 40% of power generation and Taiwan only has 10-11 days of domestic natural gas storage; I almost didn't believe these figures when I read them[37]. Taiwan plans to ratchet up LNG reliance further by expanding existing LNG terminals and building new ones. The risks are so apparent that LNG suppliers in Singapore's trading hub require force majeure and act of war provisions in Taiwan delivery contracts in case they're unable to deliver cargoes and have to seek alternative buyers.

Taiwan also imports ~60 percent of its food supply, and according to Taiwan's Council on Agriculture, the country relies on imports for 67% of its caloric intake. As shown below, **Taiwan is the only country outside the Middle East in the top ten list of food import dependency, other than the island of Hong Kong**. This is all the more unusual given how much higher Taiwan's arable land share is compared to the rest of the countries in the top ten of food import dependency.

Bottom line: Taiwan may be the most blockade-sensitive advanced economy in the world.

Net imports of fossil fuels as a share of primary energy consumption

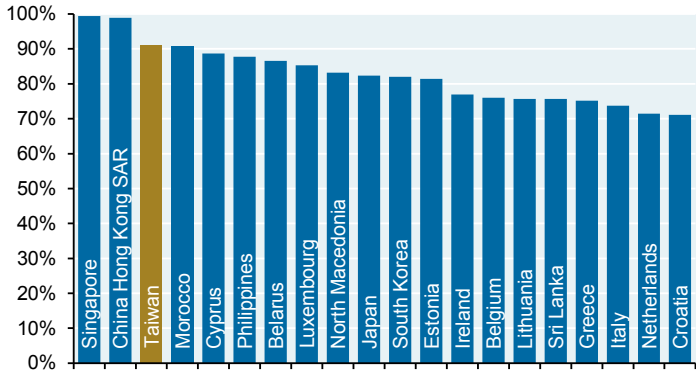

Source: Energy Institute, JPMAM, 2025

Taiwan nuclear generation

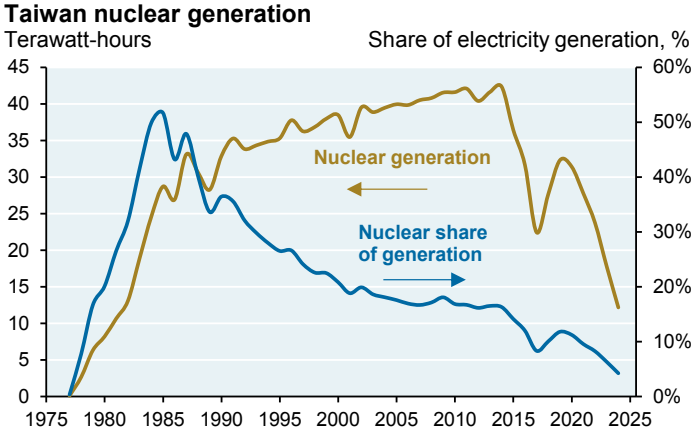

Source: Energy Institute, Taiwan Power, JPMAM, 2025

Value of food imports as a % of domestic food supply
Countries with at least $50 bn in GDP

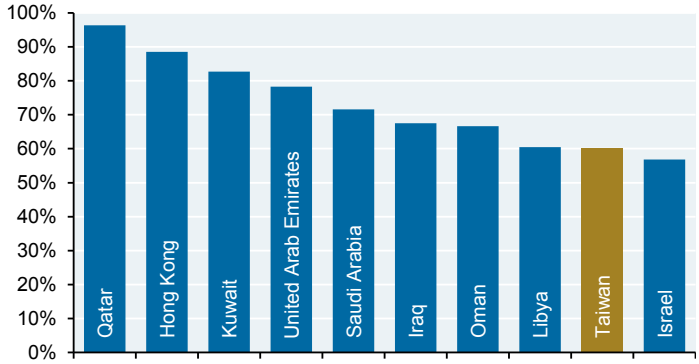

Source: UN FAO, 2025

Arable share of total land
Countries with at least $50 bn in GDP

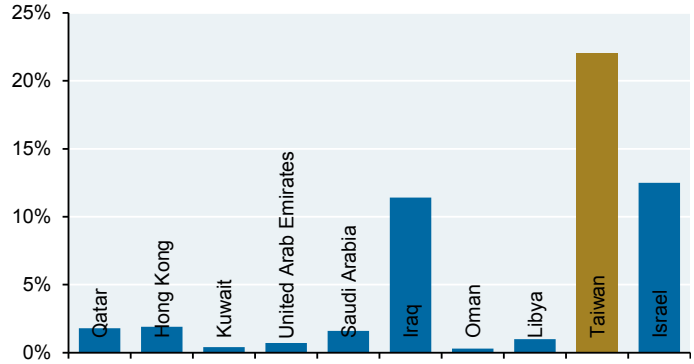

Source: World Bank, 2025

[37] "*Taiwan vulnerable to LNG supply risks in the event of a maritime blockade*", S&P Global, May 2024

The first map shows the small number of major ports through which imports flow into Taiwan; Kao-hsiung alone accounts for 57% of Taiwan's imports including containerized goods, oil and LNG. In the second map, the red shaded areas show where China's "Joint Sword 2024-A" military exercises occurred following the inauguration address of Taiwan's president William Lai. The exhibits on the next page illustrate the extreme power imbalance between China and Taiwan, China's rapid military advances and the share of Chinese military assets that are deployed in or near the Taiwan Strait.

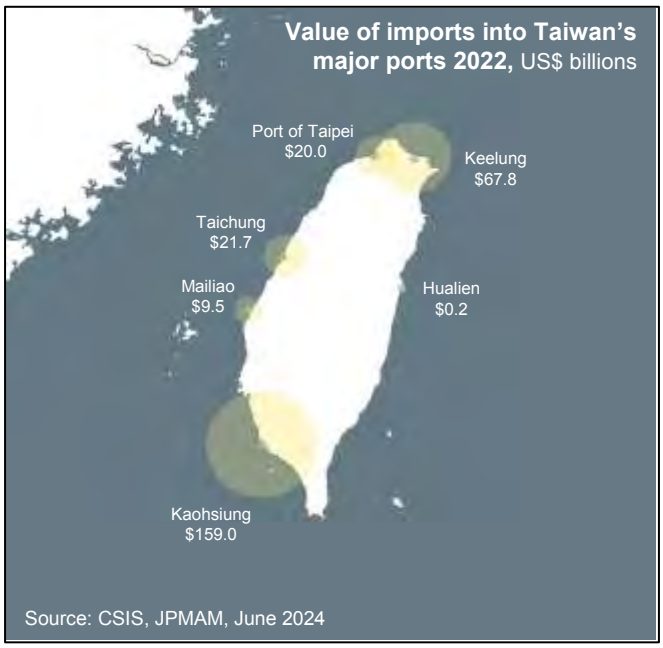

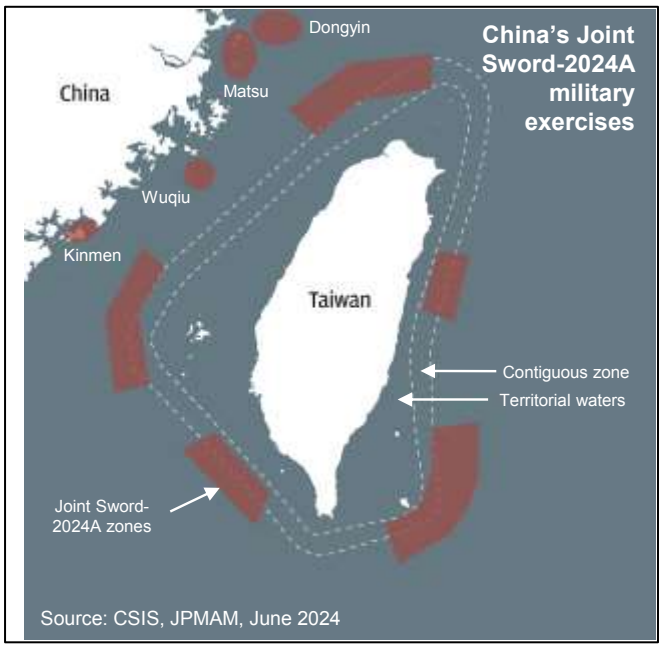

Heating up

"*No one can sever our family bonds, and no one can stop the historical trend of reunification of the motherland*" "*Taiwan's return to China is a central element of the post-World War II international order*"…*President Xi*

In 2024, China massed dozens of navy and coast guard vessels in the largest maritime exercise targeting Taiwan and the Western Pacific since 1996. This was the first time such a large-scale maritime operation involved multiple Chinese commands[38], and in 2024 Chinese probing of Taiwanese air and maritime boundaries increased by 60% vs 2023. In September 2025, hacked documents revealed that Russia is using battlefield experience to train Chinese airborne units[39] in areas such as dropping heavily armored vehicles from high altitudes using special purpose parachute systems, and is selling China heavy-lift transport aircraft, light amphibious vehicles, self-propelled anti-tank guns and airborne armored personnel carriers. In a conflict over Taiwan, Russia's supply of oil, gas and other natural resources could become strategic backup for China.

It's impossible to know what President Xi's timetable is, or to what lengths China would go to achieve Xi's goals. It's also possible that reunification would not disrupt Western access to TSMC chips, and that China would not exercise leverage should it control Taiwan's output. But given how China has been flexing its muscles regarding critical minerals, I wouldn't bet on it. I would also not be surprised to learn that China sees Trump's second term as a window during which the US is inclined towards non-intervention in a Taiwan conflict given the "Trump Corollary to the Monroe Doctrine" outlined in a December 2025 US National Security Strategy document. In other words, the US may focus militarily on the Western Hemisphere instead; the record $11 bn US arms sales package to Taiwan announced in December may be more of an economic arrangement than a military one.

[38] Wall Street Journal, December 10, 2024

[39] Washington Post, September 26, 2025 and the British Royal United Services Institute

China/Taiwan military exhibits: every picture tells a story

The balance of naval and air power in the Taiwan Strait

Assets	China	China (near Taiwan)	Taiwan
Naval			
Aircraft carriers	2	1	0
Amphibious assault ships	3	3	0
Cruisers	8	4	0
Destroyers	42	30	4
Frigates	47	30	22
Corvettes	50	40	0
Landing ships	57	50	50
Attack submarines	47	31	4
Nuclear-powered attack submarines	6	2	0
Nuclear-powered ballistic submarines	6	6	0
Coastal patrol (missile)	60	60	43
Air force & naval aviation			
Fighters	1900	750	300
Bombers/attack	500	300	0
Transport	500	40	50
Special mission aircraft	250	150	20

Source: CSIS, 2023

January 1, 2026 update: Reunification is unstoppable

In a New Year's Day speech, President XI announced that "the reunification of our motherland, a trend of the times, is unstoppable". This speech came after China conducted "Justice Mission" military exercises in waters north and southwest of Taiwan that involved long-range artillery rockets, destroyers, bombers and fighter jets within 24 miles of Taiwan's coast. These miliary exercises appear designed to demonstrate the PLA's ability to cut off Taiwan from any outside efforts to come to its aid in a conflict, and to demonstrate the PLA's ability to enforce a blockade. The exercises included amphibious assault ships and drones whose drills focused on precision strike operations and seizing/ controlling key ports, according to the Chinese military.

China and US Navy size comparison
Number of battle force ships

Source: CSIS, CRS, DOD, September 2025

China fighter aircraft by generation
Count

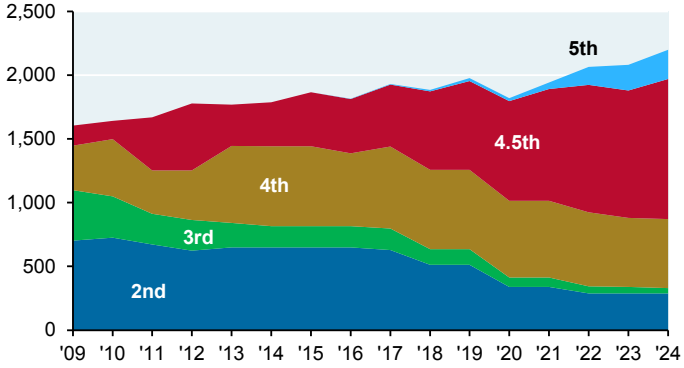

Source: CSIS, IISS, September 2025

Share of Chinese military assets deployed in the Taiwan Strait, Percent

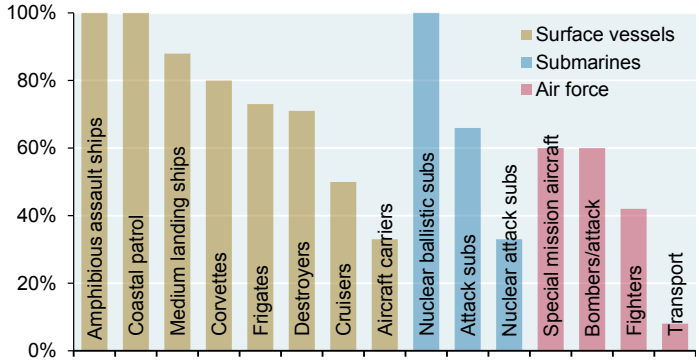

Source: CSIS, DOD, September 2025

China military and dual-use satellites
Number of satellites launched

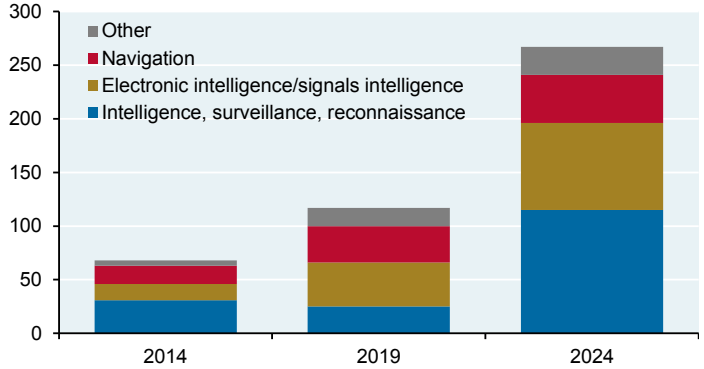

Source: CSIS, IISS, September 2025

The Rest: tech spending contributions to GDP, the Fed's inflation/employment dilemma, tariffs, immigration and the labor supply, the OBBBA, the US$, capital markets, US equity outperformance, healthcare stocks, China, Japan, a history of populism for investors and a Presidential cabinet time capsule

If it were not for a once-in-a-generation surge in AI-related tech capital spending in Q1 and Q2 of 2025, the US economy would have been in modestly worse shape. Annual real GDP growth was 2.3% from Q3 2024 to Q3 2025 even when including the benefit of Q3's 4.3% print; this annual figure would have been 1.8% without the benefit of above-trend tech capital spending. In Q1 2025, slowing growth in personal consumption expenditures actually dipped below tech related equipment & software spending, which is a pretty rare occurrence.

Quarterly contributions to real GDP growth
US$, billions

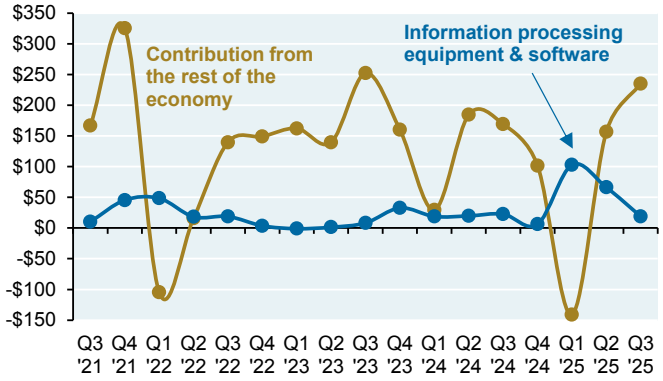

Source: Bloomberg, BEA, JPMAM, Q3 2025

Quarterly contributions to real GDP growth
Percent

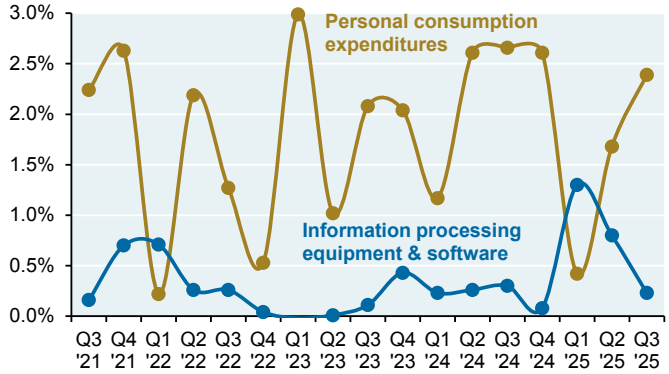

Source: BEA, Bloomberg, JPMAM, Q3 2025

Another way to visualize the magnitude of tech capital spending: compare it to some of the largest capital outlays of the 20th century. Tech capital spending in 2025 was roughly equal to the Manhattan Project, farm electrification, the Moon Landing, the Interstate Highway system and several FDR-era public works projects combined, measured as a share of GDP.

Tech capital spending in 2025 vs spending on major US infrastructure projects
Peak annual project percent of GDP

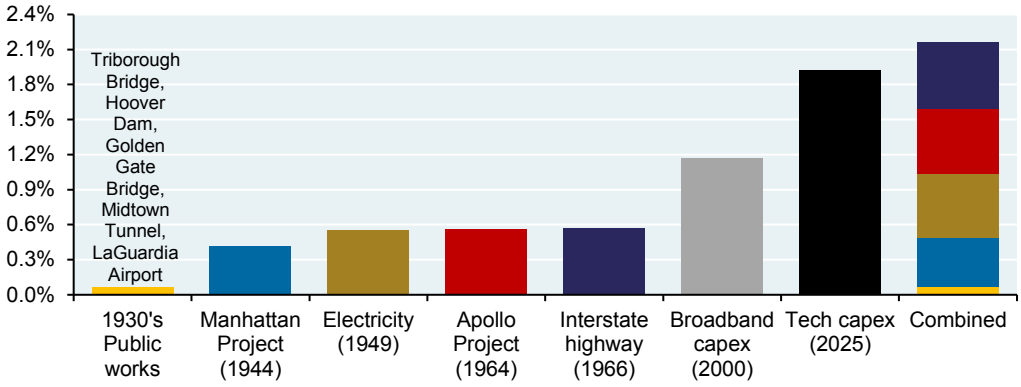

Source: Manhattan District History, BEA, Planetary Society, Eno Center for Transportation, San Francisco Fed, Hoover archives, Baruch, GoldenGate.org, New York Times, JPMAM, 2025

The Fed's pig-in-the-python bet. With some labor market indicators weakening and with prices paid surveys in positive territory, **Fed easing is a bet that any tariff-driven inflation spike will subside next year**. The challenge for the Fed[40]: today's employment-inflation divergence is in the 70th percentile for manufacturing and the 83rd percentile for services (4th chart).

Monthly net change in nonfarm payrolls
Thousands, 3 month average, seasonally adjusted

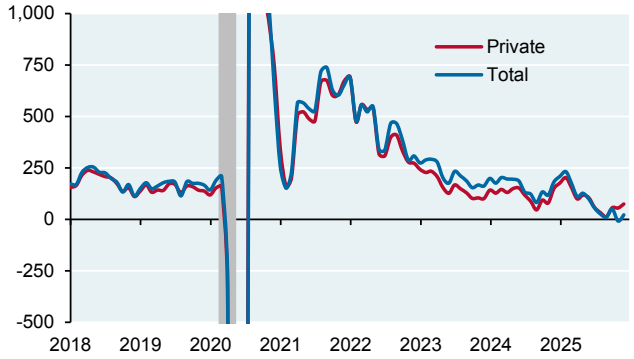

Source: BLS, Bloomberg, JPMAM, November 2025

ISM surveys: prices paid
Index (50+ = increasing), 3 month moving avg

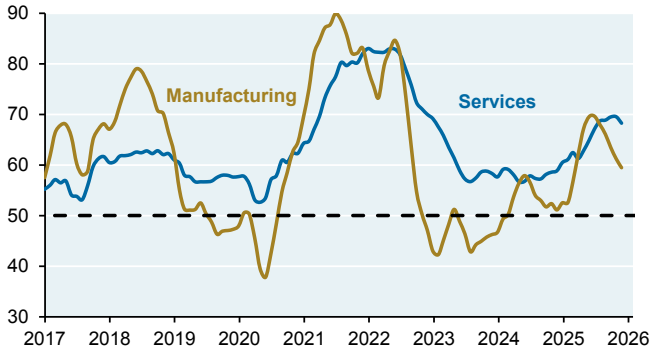

Source: ISM, Bloomberg, JPMAM, November 2025

ISM surveys: employment
Index (50+ = increasing), 3 month moving avg

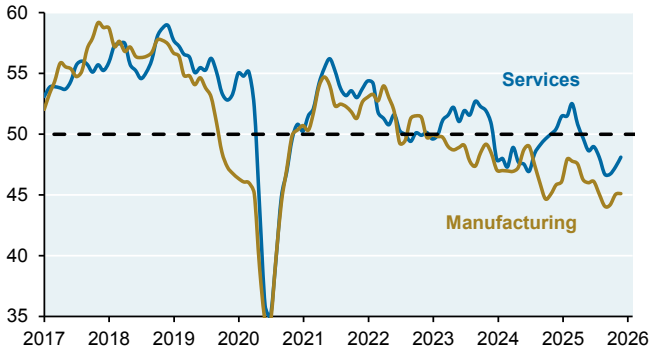

Source: ISM, Bloomberg, JPMAM, November 2025

Fed's dilemma: rising price surveys, falling employment surveys, ISM prices paid minus ISM employment

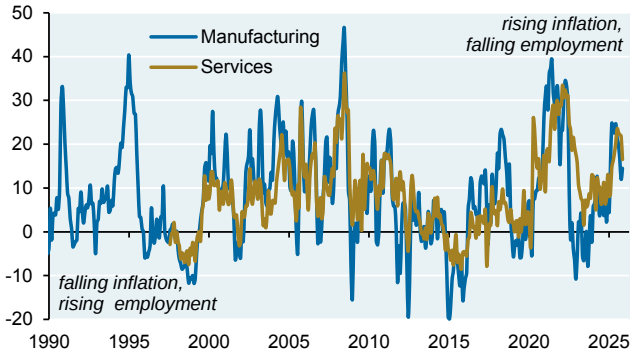

Source: ISM, Bloomberg, JPMAM, November 2025

Labor market indicators

Indicator	Dec-22	Dec-24	Latest	Indicator	Dec-22	Dec-24	Latest
Atlanta Fed wage tracker, y/y % change	6.1%	4.2%	3.8%	Premium of job switchers vs stayers, y/y % change	16%	-8%	12%
Challenger US monthly announced job cuts	43,651	38,792	71,321	Share of industries with employment growth in last mo	64%	61%	57%
Employment cost index, y/y % change	5.1%	3.8%	3.6%	Unemployment less new college grad unemploym.	-0.6%	-0.8%	-0.8%
Employment to population ratio (age 25-54)	80%	81%	81%	Unit labor cost index for nonfarm business sector	116	121	123
ISM employment indicator	50	45	44	US labor force participation rate: 25-54	83%	83%	83%
Jobs plentiful minus jobs hard to get	35%	22%	10%	US labor force participation rate: 55-64	66%	67%	66%
JOLTS layoffs and discharges, % of total employment	1.0%	1.1%	1.2%	US labor force participation rate: over 64	19%	19%	19%
JOLTS voluntary quit rate, % of total employment	2.7%	1.9%	1.8%	US labor force participation rate: women 25-54	77%	78%	78%
Labor market pressure gauge	1.0%	-1.1%	-0.9%	US small businesses with plans to raise wages	27%	24%	24%
Labor market surprise index (negative=downside surprise)	0.3%	0.0%	-0.3%	US small businesses with unfilled job openings	41%	35%	33%
Net probability of losing minus finding job, $50-100k income	-45%	-36%	-37%	US unemployment U3	3.5%	4.1%	4.6%
Net probability of losing minus finding a job, >$100k income	-52%	-46%	-39%	US unemployment U6	6.6%	7.5%	8.7%
Nonfarm payrolls m/m change, 3 mo avg, thousands	274	209	22	Wage growth: immigration sensitive services	6.3%	3.2%	2.7%
Nonfarm private payrolls m/m change, 3 mo avg, thousands	243	177	75	Wage growth: overall economy	6.3%	3.2%	2.6%
Part time employees unable to find full time work, mm	925	1,195	1,759				

Source: Atlanta Fed, Conference Board, FRB of Cleveland, NY Fed, Fed Consumer Expectations Survey, BLS, US Department of State, Indeed, Bloomberg, Haver, JPMAM, December 19, 2025

[40] *What about Fed independence?* I get asked this question a lot. It's too soon to tell since we can only analyze a new Fed chair over a business cycle. Presidents have a long history of putting pressure on the Fed dating back to Nixon's dirty tricks campaign against Fed chair Arthur Burns. If you're interested in Presidential decisions regarding the Fed, here's an article from Peter Conti-Brown, law/ethics professor at Wharton, critiquing Biden's Fed appointments as having a "stunning anonymity within central bank circles" and a lack of relevant academic contributions or previous service at the Fed

The impact of the reconciliation bill on US growth. The chart below summarizes the bill's sources and uses which restructure $5 trillion in taxes and spending. Economic drags include a gutting of the safety net by $1.5 trillion and slashed energy bill subsidies. Offsets are important to focus on too: the transfer of tax benefits from the green economy to the broader corporate sector, and new tax cuts for working families with children. The bill is an expression of White House policies: if you don't have a job, get one; and if you don't have children, start producing them. OBBBA may also be a long-term fiscal disaster if its policies are sustained for long[41].

Some tax cut bills provide the largest tax benefits to the wealthiest families. The OBBBA (One Big Beautiful Bill Act) is a bit different; as shown below, benefits when measured by changes in after tax income are more equally distributed and begin with families whose incomes are in the 40th percentile. And as shown in the last chart, US consumer spending appears increasingly reliant on the wealthy. Bottom line: consumer spending may hold up in 2026 despite the regressive features of the OBBBA, boosted in the first part of the year by rising tax refunds.

OBBBA sources & uses, Current Policy basis, Senate version
US$, trillions, 2025-2034

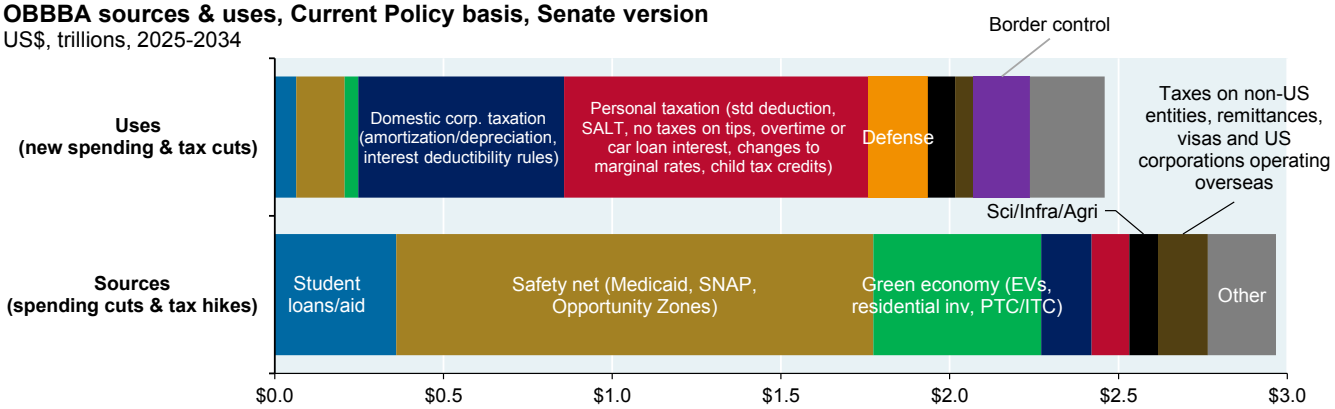

Source: JPM Global Markets Strategy, CBO, JPMAM. August 2025. Current Policy basis means that all figures are compared to 2024 fiscal conditions rather than comparing them to fiscal conditions assuming a sunset of all TCJA provisions (Current Law basis).

Change in after-tax income by income group for OBBBA tax changes, Percent, excludes impact of TCJA extension

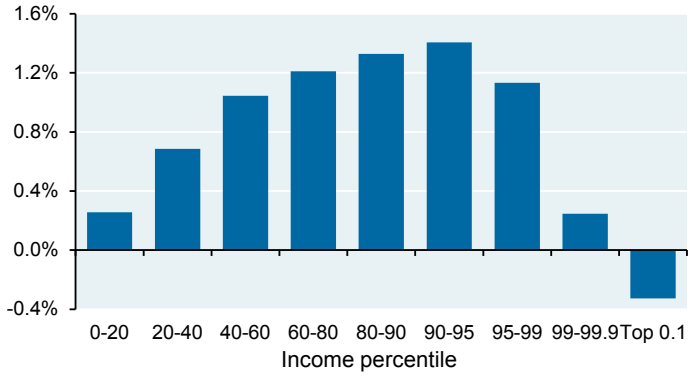

Source: Tax Policy Center, October 2025

Share of spending by consumers in top 10% of income distribution, Percent

Source: Moodys Analytics, Federal Reserve, Q2 2025

[41] Before the OBBBA, the CBO projected a debt to GDP ratio of 154% in thirty years compared to today's 100%. Alan Auerbach (Berkeley) and William Gale (Brookings) found that OBBBA provisions could increase the US debt to GDP ratio to 183% in thirty years, and to 199% if OBBBA temporary tax and spending provisions are made permanent. "*Then and Now: A Look Back and Ahead at the Federal Budget*", October 2025

I'm more worried about immigration than tariffs as risks to growth. As we wrote in October, after accounting for tariff exclusions granted to certain products, importer country-switching, importer tax avoidance and high margins on imported goods which dilute the impact of tariffs, selling prices on imported goods have risen by ~5% rather than by 15%. As for the muted impact of tariffs on overall inflation so far, see the chart on the right: services are an enormous share of overall consumption and only 6% of this amount is imported. Even for non-durable goods, the import share is only 20%. For durable goods, the impact on inflation is a bit larger since imports are 26% of consumption. Overall, the import share of US consumption is 10%-11%; this includes spending on final goods produced abroad, less the local content embedded in prices of these goods, plus the import content in domestically produced goods and services[42].

A bigger concern: slowing labor force growth and the reduction in the breakeven level of payroll growth (above which job growth may cause wage inflation).

US tariff rate estimates, tariff revenues and imported goods selling prices, Percent

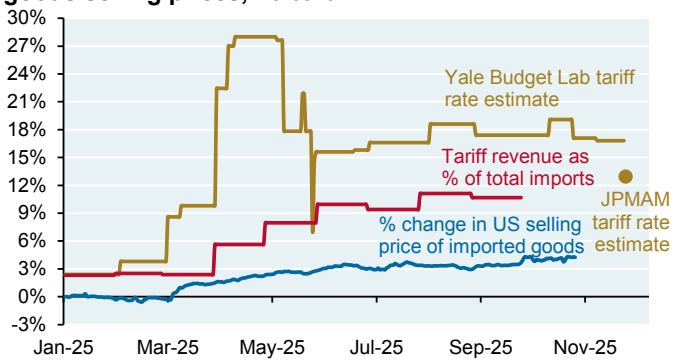

Source:Yale Budget Lab, US Treasury, Census Bureau, Digital Data Design Institute Pricing Lab, JPMAM, December 2025

Import composition of US consumption
Share of US personal consumption expenditures

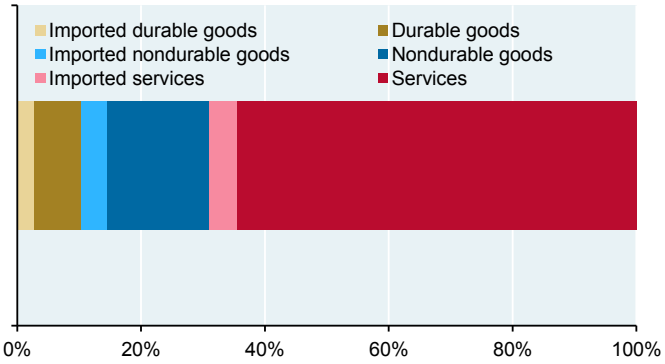

Source: Federal Reserve Bank of San Francisco, JPMAM, 2019

Estimates of the breakeven level of payroll gains vs actual payroll gains, Thousands, monthly

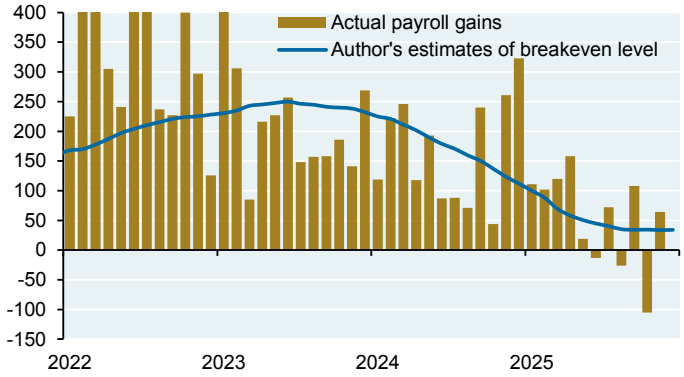

Source: Bloomberg, BLS, Dallas Fed, JPMAM, November 2025

[42] "*How Much Do We Spend on Imports?*", Federal Reserve Bank of San Francisco, January 2019

Ultimately, capital markets will be an important barometer of administration policies. At the end of 2025 we began to see signs of recovery in primary and secondary IPO activity and announced M&A activity. As we discussed in the December alternatives piece, we expect buyout and venture managers to take advantage of these improved conditions and monetize part of their enormous backlog of unsold positions.

US IPO activity, 3 month rolling avg

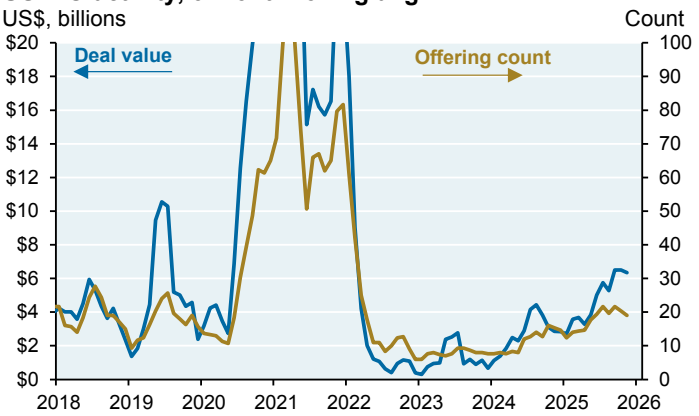

Source: Bloomberg, JPMAM, November 2025

US average IPO return on first trading day

Percent return, IPOs > $25 million

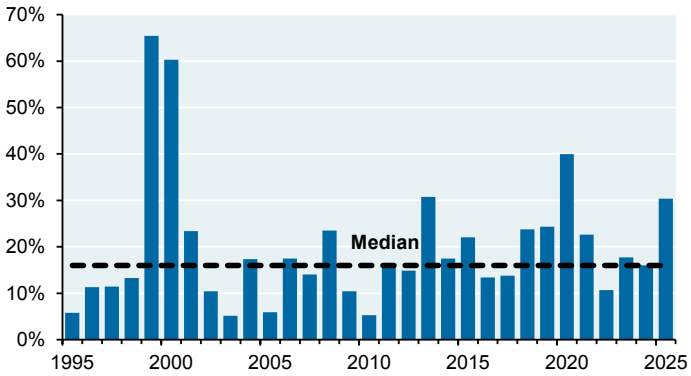

Source: GS Global Investment Research, September 2025

US secondaries activity, 3 month rolling avg

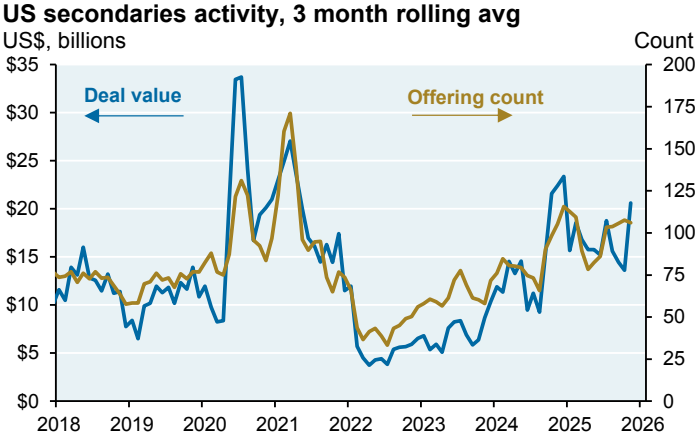

Source: Bloomberg, JPMAM, November 2025

Total announced M&A volumes by quarter

US$, billions

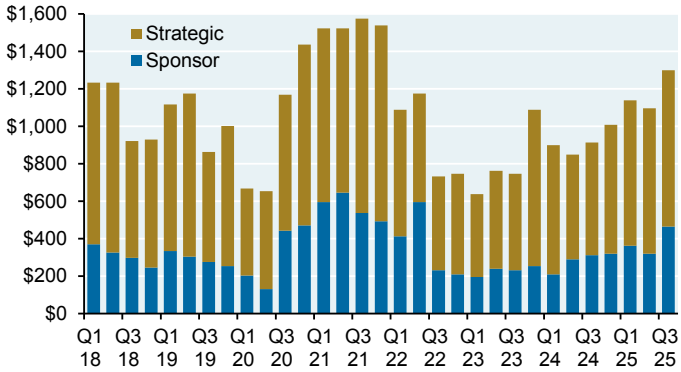

Source: Morgan Stanley Research, Q3 2025

Waiting for Godot: unmonetized venture and buyout shares, Median remaining value as a % of fund value + distributions

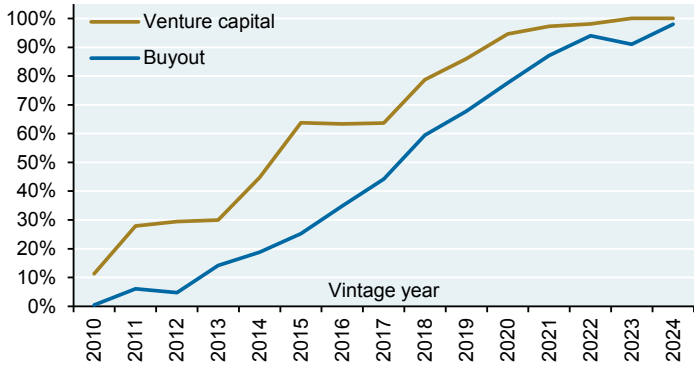

Source: Steve Kaplan (U Chicago), MSCI/Burgiss, JPMAM, Q2 2025

The dollar should be stable in 2026. In 2025, the US dollar declined ~10% from its post financial crisis high and led to the re-emergence of dollar Armageddonists. While US debt and deficit projections continue to deteriorate, we're not seeing material declines in the six measures we use to assess the dollar's role in the world economy, and I am not surprised to see the dollar having stabilized after its initial decline.

2025: 5%-10% decline in the US dollar from elevated levels
Index (100 = January 31, 2006)

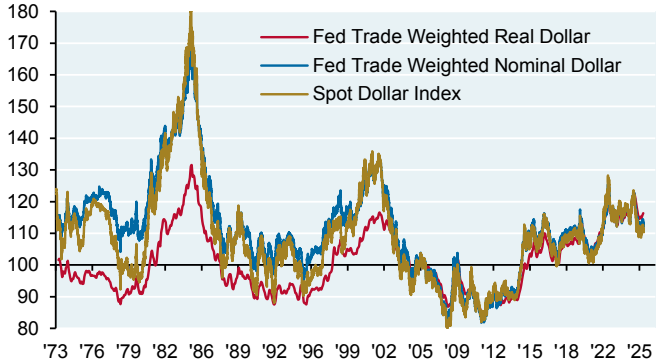

Source: Bloomberg, JPMAM, December 17, 2025

US$ share of debt, FX, reserves and trade

	2022	Latest	Latest as of date
Cross-border loans	50%	52%	Q4 2024
Intl. debt securities	49%	53%	Q4 2024
FX transaction volume	88%	89%	April 2025
Official FX reserves	60%	56%	Q2 2025
Trade invoicing	50%	42%	June 2025
SWIFT payments	42%	47%	August 2025

Source: BIS, ECB, IMF, SWIFT, JPMAM

Healthcare. In our August 12, 2025 Eye on the Market, we highlighted the extreme underperformance of US healthcare stocks relative to technology stocks since 2020 and how there was a lot of very bad news priced in. Since that time, healthcare has rebounded. I expect healthcare stocks performance to improve in 2026 for many of the reasons cited in the August piece: low valuations cure a lot of wounds.

Healthcare vs information technology returns
Total return index (08/12/2025 = 100)

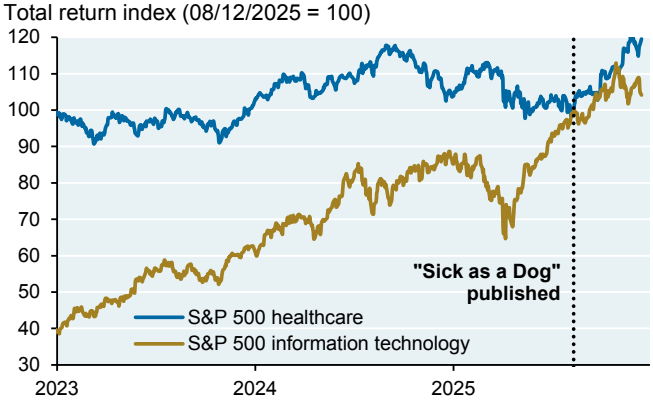

Source: Bloomberg, JPMAM, December 15, 2025

Overweight US & EM, underweight Europe & Japan
3-year rolling out (under) performance vs MSCI All World Index

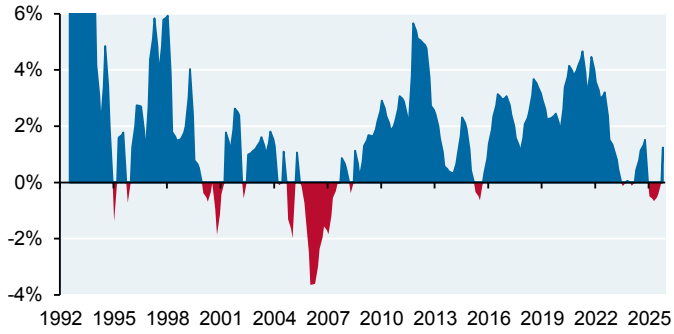

Source: Bloomberg, JPMAM. December 2025. All equity portfolio, rebalanced quarterly. O/W US by 10%; U/W EUR by -10%; U/W JPN by -5%; O/W EM by 5%. Assumes no currency hedging. Past performance is not indicative of future results

The US-EM barbell keeps chugging along. The rolling three-year excess return of the overweight US-Emerging Market, underweight Europe-Japan barbell yielded a small negative figure for a few months during 2025 but ended the year in positive territory again. With the exception of the 2006-2007 period, this barbell has been one of the most consistently beneficial asset allocation strategies since I joined JP Morgan in 1987.

Is the era of US outperformance over? From 2009 to 2024, the S&P 500 generated 3.6x the return of the MSCI World ex-US Index. In 2025 the US finally underperformed most equity markets, leaving the S&P 500 vs rest of world multiple since 2009 at 3.0x. To be clear, **the right move has generally been to ignore annual strategist calls since 2010 to bail on the US given cheaper P/E multiples elsewhere in the world.** The second chart shows cumulative returns since 2009 based on when you bailed on the S&P 500 and switched to non-US stocks instead. If you switched in 2025, that was ideal; any year before that left money on the table vs never switching at all.

Trump policies on trade, the dollar and the deficit argue for a less US-centric portfolio than in the prior decade, particularly since non-US equities still trade at a large P/E discount to the US (60% when the year began, now ~70%). But most US sectors delivered higher earnings growth in 2025 and also generate higher RoE and RoA than non-US counterparts. For that reason, I don't think parity is the right expectation for relative P/E ratios; I would be surprised if the MSCI World ex-US Index P/E exceeded 80% of US levels.

Non-US P/E discount vs US

MSCI World ex-US fwd P/E divided by US fwd P/E

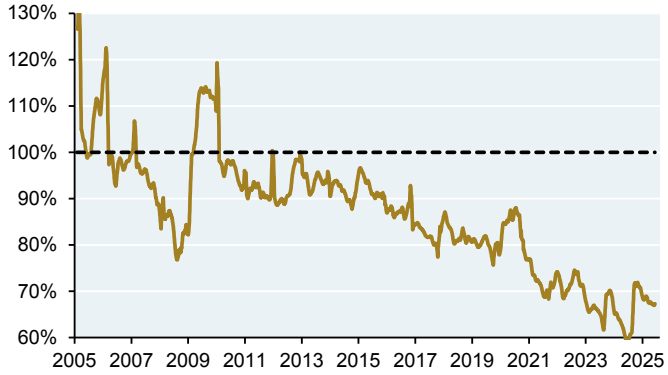

Source: Bloomberg, JPMAM, December 26, 2025

Investors reaped large gains by sticking with US stocks

Cumulative total returns, January 2009 to December 2025

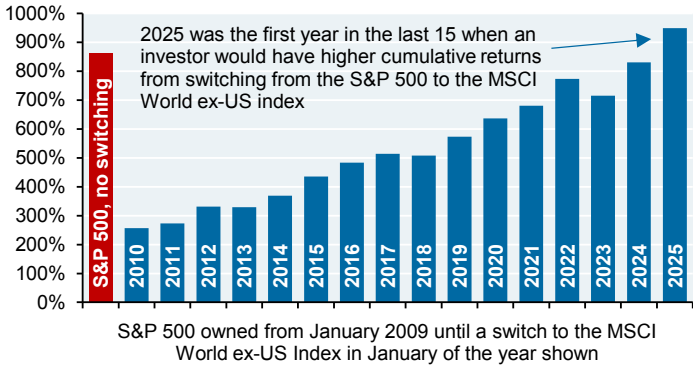

S&P 500 owned from January 2009 until a switch to the MSCI World ex-US Index in January of the year shown

Source: Bloomberg, JPMAM, December 24, 2025

2025 ytd total returns for major equity markets

Percent, US$ returns

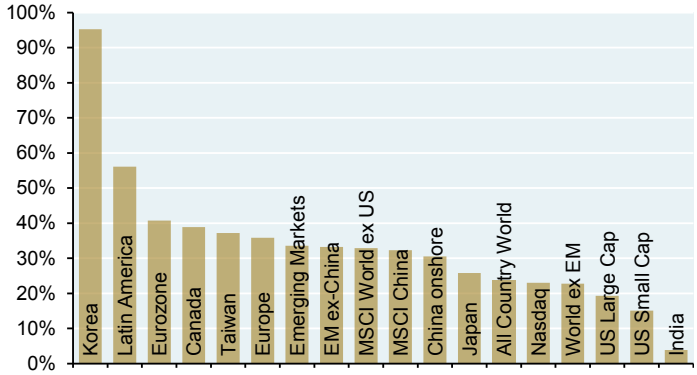

Source: Bloomberg, JPMAM, December 27, 2025

2025 year over year EBIT growth by sector and country

	US	Europe	China	Japan
Communication Services	19%	-7%	18%	42%
Consumer Discretionary	-3%	-21%	-23%	-14%
Consumer Staples	0%	-3%	2%	-1%
Energy	-10%	-16%	-17%	-22%
Financials	23%	9%	11%	20%
Real Estate	5%	54%	-16%	17%
Health Care	12%	4%	22%	11%
Industrials	7%	6%	-24%	10%
Information Technology	27%	13%	40%	11%
Materials	-6%	-11%	29%	-29%
Utilities	7%	-2%	-3%	9%
Total	12%	-1%	0%	8%

Source: Bloomberg, JPMAM, December 18, 2025

Return on Assets: Higher in the US						
Country	Staple	Con Disc	Tech	Hcare	Comm Serv	Finan
US	6.9	7.6	13.1	6.4	10.4	1.5
Europe	5.0	2.0	8.8	7.3	2.1	0.7
Japan	2.3	3.2	7.5	4.8	5.0	0.5
China	4.2	1.3	9.1	3.8	3.0	0.8

Source: Bloomberg, JPMAM, December 17, 2025

Return on Equity: Higher in the US						
Country	Staple	Con Disc	Tech	Hcare	Comm Serv	Finan
US	23.5	26.0	29.3	18.8	24.1	13.5
Europe	15.6	5.5	17.1	17.4	7.7	12.8
Japan	6.4	9.2	11.7	8.4	16.7	9.7
China	15.6	13.6	10.7	12.7	15.3	11.2

Source: Bloomberg, JPMAM, December 17, 2025

China and Japan. While China property investment and fixed asset investment remain weak, broader surveys of activity are stable[43] and JP Morgan economists are upgrading prospects for 2026. More importantly for investors, Chinese companies have substantially increased distributions to shareholders in excess of cash rates.

On Japan, ongoing corporate reforms and economic stimulus are positive signs. Japan still has plenty of wood to chop on the former, as the Topix cash to asset ratio is more than 2x US and European levels. What about rising Japanese bond yields? In August 2024, rising JGB yields and a rising Yen spooked markets since they were seen as signs of repatriation of capital to Japan and unwinding of global Yen-funded assets. At the current time, the Yen is falling rather than rising vs the dollar and while Yen yields are rising, they're still below US and German yields. As a result, the impetus for capital repatriation back to Japan may be muted for now. **But this will remain one of the market's greatest potential fissures since Japan has the largest Net International Investment Position in the world**, a measure of its assets invested abroad (close to $3.5 trillion).

China economic activity monitors
Percent, y/y change
Index

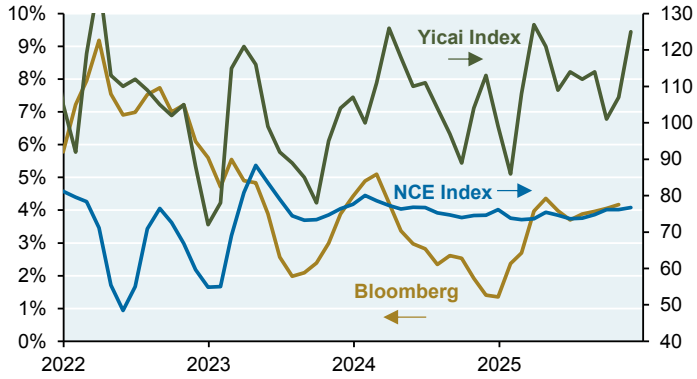

Source: Bloomberg, Now-Casting Economics, Yicai Research, November 2025

Capital returned to China shareholders vs cash rates
Percentage of market cap, 12 month moving avg

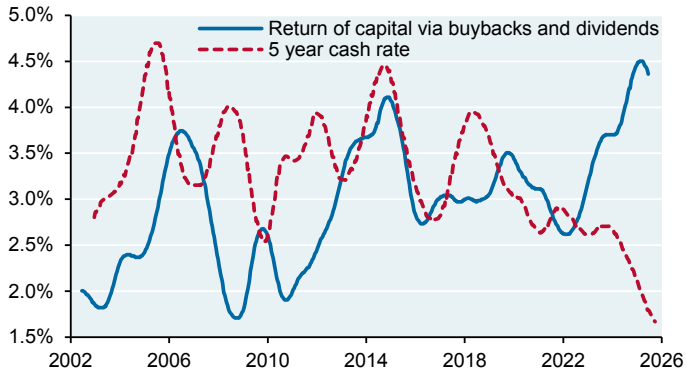

Source: Bridgewater Daily Observations, December 1, 2025

Median cash to asset ratio
Percent

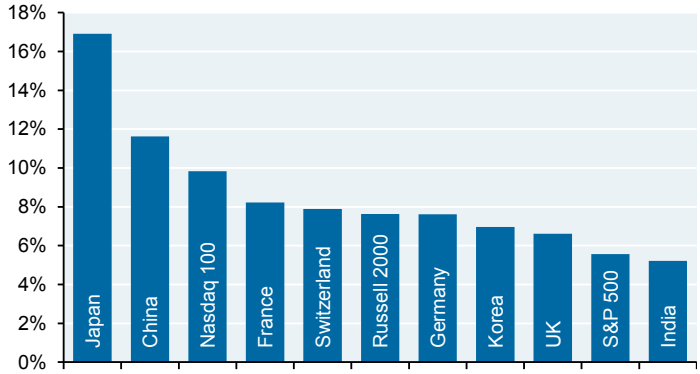

Source: Bloomberg, JPMAM, December 27, 2025

US, German and Japanese 10 year bond yields
Percent

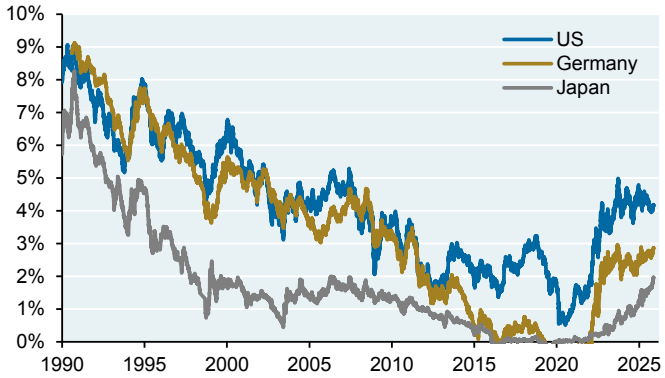

Source: Bloomberg, JPMAM, December 17, 2025

[43] **Bloomberg China index**: exports, coal/electricity consumption, steel production, motor vehicle sales, real estate investment, medicine consumption, telecom equipment production, SOE output, private output. **Yicai Index**: coal/electricity consumption, employment, traffic congestion, commercial real estate sales, subway flow, air pollution, dry bulk freight shipping, bankruptcies. **NCE Index**: retail sales, industrial production, exports, imports, new construction starts, coal/electricity consumption, motor vehicle sales, purchasing manager surveys

Some clients believe that something has fundamentally changed regarding US governance, institutions and rule of law. I'm not going to opine on that, but I will share the following: populism is generally not good for investors. In each chart below the vertical line indicates when 20th and 21st century populists took power[44]. The timeline: first, judicial constraints, free elections and media freedoms are curtailed; populists follow with higher tariffs, less trade and less financial openness; debt and inflation then rise; and finally, real per capita GDP falls.

Judicial constraints, free elections and media freedom

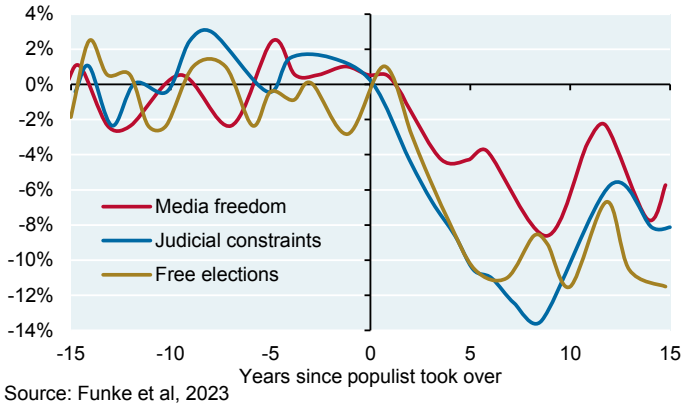

Source: Funke et al, 2023

Tariff rates, trade/GDP and financial openness

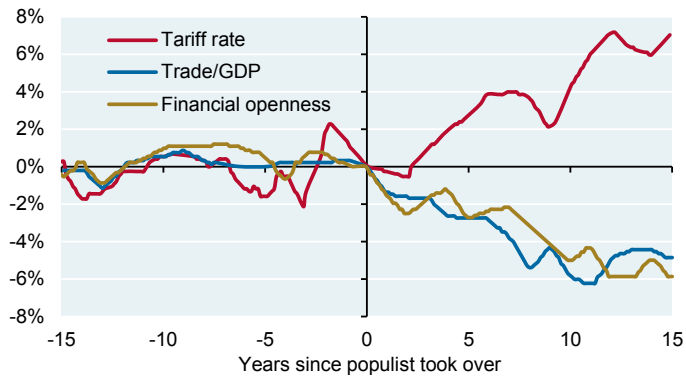

Source: Funke et al, 2023

Debt/GDP and inflation

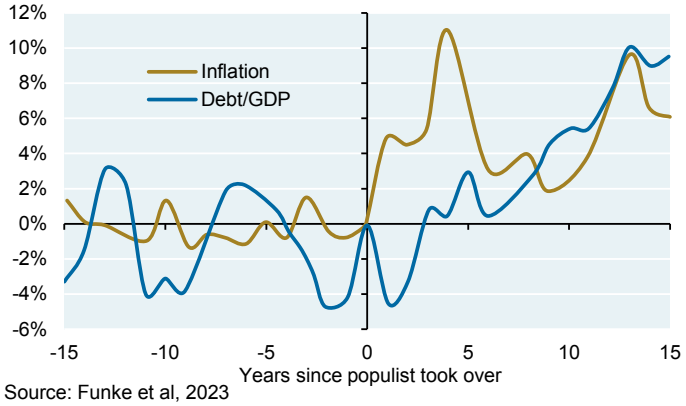

Source: Funke et al, 2023

Real per capita GDP

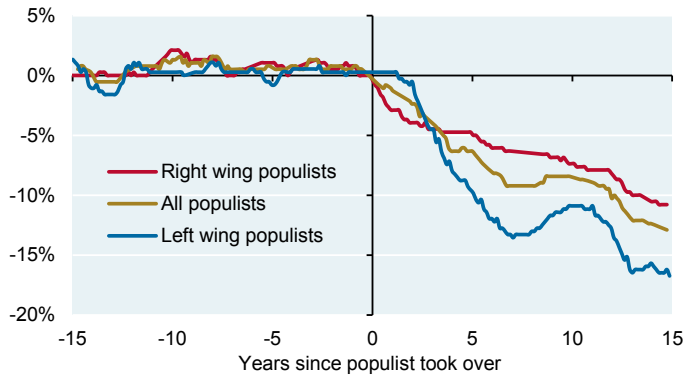

Source: Funke et al, 2023

[44] "*Populist Leaders and the Economy*", Funke, Schularick and Christoph Trebesch, American Economic Review, 2023. The authors analyzed 20,000 pages of populist leader case studies to identify 51 presidents and prime ministers from 1900 to 2020 that fit the modern definition of a populist politician: someone with a political strategy that pits the people against the elite and positions themselves as the sole representative of the people. The authors isolate the impact of populist leaders by taking the economic characteristics of each country 15 years before and after a populist leader takes power and comparing it against a synthetic economy comprised of a weighted combination of other economies constructed to most closely track the path of each country in the 15 years leading up to the populist leader taking power.

A run for their money

In my lifetime, the two most controversial Presidential cabinet members might have been JFK/LBJ Defense Secretary Robert McNamara (1961-1968) and Nixon Attorney General John Mitchell (1969-1972):

- **McNamara**, a primary architect of the Vietnam War, wrote a 1995 mea culpa conceding that the war should have been avoided; that it could have been halted at several points after it started; that he and other advisers to President Johnson suffered from ignorance, inattention, flawed thinking and political expediency; that their strategy had little chance of success; and that communism would not have prevailed in Asia with the strategic position of the United States no worse than it is today. From an editorial at the time of McNamara's 1995 mea culpa: "McNamara must not escape the lasting moral condemnation of his countrymen. Surely he must in every quiet and prosperous moment hear the ceaseless whispers of those poor boys in the infantry, dying in the tall grass, platoon by platoon, for no purpose. What he took from them cannot be repaid by prime-time apology and stale tears, three decades late"

- **Mitchell's** policy decisions included expansion of warrantless wiretapping and "no-knock" warrants, attempts to stop publication of the Pentagon Papers (the conservative Warren Burger Supreme Court allowed publication), approving the Watergate break-in, suppression of anti-war protestors, sabotage of 1968 Paris Peace Accords, foul language regarding former Washington Post publisher Katherine Graham and participation in a bizarre kidnapping and sedation of his own wife to prevent her from talking to the press. Mitchell was sentenced to prison in 1977 and served 19 months in jail

There was an October 7, 2025 Washington Post OpEd entitled "***Six surgeons general: It's our duty to warn the nation about RFK Jr***". The authors all served as surgeon general under George HW Bush, Clinton, George W Bush, Obama, Trump and Biden[45]. Their concern: RFK's "profound, immediate and unprecedented threat to public health". They describe the current period as one during which science and expertise take a back seat to ideology and misinformation. They cited the following policy issues:

- RFK's claims about vaccines and autism, the HPV vaccine and mRNA technology, all of which they see as lacking in scientific foundations
- RFK's de-emphasis of vaccines and promotion of vitamins during the worst measles outbreak in 30 years. [Note: a recent outbreak in South Carolina puts the US on the precipice of losing its measles elimination status. Over 90% of cases have been in unvaccinated children aged 5 to 17. The measles outbreak is mostly a byproduct of falling vaccination rates, now below 93% vs the 95% level needed to prevent contagion]
- RFK's termination of all 17 members of the Advisory Committee on Immunization Practices and replacement with unqualified replacements
- How the American Academy of Pediatrics and the American College of Obstetricians issued public guidance urging doctors and patients to ignore RFK's warnings about prenatal acetaminophen use and autism
- How RFK sidelined hundreds of scientists, public health officials and medical professionals and created what they describe as a culture of intimidation where scientific findings are censored, evidence is disregarded and career officials are pressured to rubber stamp conclusions not backed by science

Rest easy, McNamara and Mitchell. At the rate things are going, in terms of controversial policy decisions, it looks like someone may give you a run for your money.

Michael Cembalest
JP Morgan Asset Management

[45] Op-Ed authors: Antonia Novello (1990-1993), Joycelyn Elders (1993-1994), David Satcher (1998-2002), Richard Carmona (2002-2006), Vivek Murthy (2014-2017, 2021-2025) and Jerome Adams (2017-2021). The authors include all living Senate-confirmed surgeon generals with the exception of Regina Benjamin (2009-2013)

MICHAEL CEMBALEST is the Chairman of Market and Investment Strategy for J.P. Morgan Asset & Wealth Management, a global leader in investment management and private banking with $6 trillion of client assets under supervision as of 2025. He is responsible for leading the strategic market and investment insights across the firm's Institutional, Funds and Private Banking businesses.

Mr. Cembalest is also a member of the J.P. Morgan Asset & Wealth Management Investment Committee and previously served on the Investment Committee for the J.P. Morgan Retirement Plan for the firm's more than 256,000 employees.

Mr. Cembalest was most recently Chief Investment Officer for the firm's Global Private Bank, a role he held for eight years. He was previously head of a fixed income division of Investment Management, with responsibility for high grade, high yield, emerging markets and municipal bonds.

Before joining Asset Management, Mr. Cembalest served as Head Strategist for Emerging Markets Fixed Income at J.P. Morgan Securities. Mr. Cembalest joined J.P. Morgan in 1987 as a member of the firm's Corporate Finance division. Mr. Cembalest earned an M.A. from the Columbia School of International and Public Affairs in 1986 and a B.A. from Tufts University in 1984.

IMPORTANT INFORMATION

This material is for information purposes only. The views, opinions, estimates and strategies expressed herein constitutes Michael Cembalest's judgment based on current market conditions and are subject to change without notice, and may differ from those expressed by other areas of JPMorgan Chase & Co. ("JPM"). **This information in no way constitutes J.P. Morgan Research and should not be treated as such.** Any companies referenced are shown for illustrative purposes only, and are not intended as a recommendation or endorsement by J.P. Morgan in this context.

GENERAL RISKS & CONSIDERATIONS Any views, strategies or products discussed in this material may not be appropriate for all individuals and are subject to risks. Investors may get back less than they invested, and **past performance is not a reliable indicator of future result**s. Asset allocation/diversification does not guarantee a profit or protect against loss. Nothing in this material should be relied upon in isolation for the purpose of making an investment decision.

NON-RELIANCE Certain information contained in this material is believed to be reliable; however, JPM does not represent or warrant its accuracy, reliability or completeness, or accept any liability for any loss or damage (whether direct or indirect) arising out of the use of all or any part of this material. No representation or warranty should be made with regard to any computations, graphs, tables, diagrams or commentary in this material, which are provided for illustration/ reference purposes only. Any projected results and risks are based solely on hypothetical examples cited, and actual results and risks will vary depending on specific circumstances. Forward-looking statements should not be considered as guarantees or predictions of future events. Nothing in this document shall be construed as giving rise to any duty of care owed to, or advisory relationship with, you or any third party. Nothing in this document shall be regarded as an offer, solicitation, recommendation or advice (whether financial, accounting, legal, tax or other) given by J.P. Morgan and/or its officers or employees. J.P. Morgan and its affiliates and employees do not provide tax, legal or accounting advice. You should consult your own tax, legal and accounting advisors before engaging in any financial transactions.

For J.P. Morgan Asset Management Clients:

J.P. Morgan Asset Management is the brand for the asset management business of JPMorgan Chase & Co. and its affiliates worldwide.

To the extent permitted by applicable law, we may record telephone calls and monitor electronic communications to comply with our legal and regulatory obligations and internal policies. Personal data will be collected, stored and processed by J.P. Morgan Asset Management in accordance with our privacy policies at https://am.jpmorgan.com/global/privacy.

ACCESSIBILITY
For U.S. only: If you are a person with a disability and need additional support in viewing the material, please call us at 1-800-343-1113 for assistance.

This communication is issued by the following entities: In the United States, by J.P. Morgan Investment Management Inc. or J.P. Morgan Alternative Asset Management, Inc., both regulated by the Securities and Exchange Commission; in Latin America, for intended recipients' use only, by local J.P. Morgan entities, as the case may be.; in Canada, for institutional clients' use only, by JPMorgan Asset Management (Canada) Inc., which is a registered Portfolio Manager and Exempt Market Dealer in all Canadian provinces and territories except the Yukon and is also registered as an Investment Fund Manager in British Columbia, Ontario, Quebec and Newfoundland and Labrador. In the United Kingdom, by JPMorgan Asset Management (UK) Limited, which is authorized and regulated by the Financial Conduct Authority; in other European jurisdictions, by JPMorgan Asset Management (Europe) S.à r.l. In Asia Pacific ("APAC"), by the following issuing entities and in the respective jurisdictions in which they are primarily regulated: JPMorgan Asset Management (Asia Pacific) Limited, or JPMorgan Funds (Asia) Limited, or JPMorgan Asset Management Real Assets (Asia) Limited, each of which is regulated by the Securities and Futures Commission of Hong Kong; JPMorgan Asset Management (Singapore) Limited (Co. Reg. No. 197601586K), which this advertisement or publication has not been reviewed by the Monetary Authority of Singapore; JPMorgan Asset Management (Taiwan) Limited; JPMorgan Asset Management (Japan) Limited, which is a member of the Investment Trusts Association, Japan, the Japan Investment Advisers Association, Type II Financial Instruments Firms Association and the Japan Securities Dealers Association and is regulated by the Financial Services Agency (registration number "Kanto Local Finance Bureau (Financial Instruments Firm) No. 330"); in Australia, to wholesale clients only as defined in section 761A and 761G of the Corporations Act 2001 (Commonwealth), by JPMorgan Asset Management (Australia) Limited (ABN 55143832080) (AFSL 376919). For all other markets in APAC, to intended recipients only.

For J.P. Morgan Private Bank Clients:

ACCESSIBILITY
J.P. Morgan is committed to making our products and services accessible to meet the financial services needs of all our clients. Please direct any accessibility issues to the Private Bank Client Service Center at 1-866-265-1727

LEGAL ENTITY, BRAND & REGULATORY INFORMATION
In the **United States**, **JPMorgan Chase Bank, N.A.** and its affiliates (collectively "**JPMCB**") offer investment products, which may include bank managed investment accounts and custody, as part of its trust and fiduciary services. Other investment products and services, such as brokerage and advisory accounts, are offered through **J.P. Morgan Securities LLC ("JPMS")**, a member of FINRA and SIPC. JPMCB and JPMS are affiliated companies under the common control of JPM.
In **Germany**, this material is issued by **J.P. Morgan SE**, with its registered office at Taunustor 1 (TaunusTurm), 60310 Frankfurt am Main, Germany, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB). In **Luxembourg**, this material is issued by **J.P. Morgan SE – Luxembourg Branch**, with registered office at European Bank and Business Centre, 6 route de Treves, L-2633, Senningerberg, Luxembourg, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Luxembourg Branch is also supervised by the Commission de Surveillance du Secteur Financier (CSSF); registered under R.C.S Luxembourg B255938. In the **United Kingdom**, this material is issued by **J.P. Morgan SE – London Branch**, registered office at 25 Bank Street, Canary Wharf, London E14 5JP, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – London Branch is also supervised by the Financial Conduct Authority and Prudential Regulation Authority. In **Spain**, this material is distributed by **J.P. Morgan SE, Sucursal en España**, with registered office at Paseo de la Castellana, 31, 28046 Madrid, Spain, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE, Sucursal en España is also supervised by the Spanish Securities Market Commission (CNMV); registered with Bank of Spain as a branch of J.P. Morgan SE under code 1567. In **Italy**, this material is distributed by **J.P. Morgan SE – Milan Branch**, with its registered office at Via Cordusio, n.3, Milan 20123, Italy, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Milan Branch is also supervised by Bank of Italy and the Commissione Nazionale per le Società e la Borsa (CONSOB); registered with Bank of Italy as a branch of J.P. Morgan SE under code 8076; Milan Chamber of Commerce Registered Number: REA MI 2536325. In the **Netherlands**, this material is distributed by **J.P. Morgan SE – Amsterdam Branch**, with registered office at World Trade Centre, Tower B, Strawinskylaan

1135, 1077 XX, Amsterdam, The Netherlands, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Amsterdam Branch is also supervised by De Nederlandsche Bank (DNB) and the Autoriteit Financiële Markten (AFM) in the Netherlands. Registered with the Kamer van Koophandel as a branch of J.P. Morgan SE under registration number 72610220. In **Denmark**, this material is distributed by **J.P. Morgan SE – Copenhagen Branch, filial af J.P. Morgan SE, Tyskland**, with registered office at Kalvebod Brygge 39-41, 1560 København V, Denmark, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Copenhagen Branch, filial af J.P. Morgan SE, Tyskland is also supervised by Finanstilsynet (Danish FSA) and is registered with Finanstilsynet as a branch of J.P. Morgan SE under code 29010. In **Sweden**, this material is distributed by **J.P. Morgan SE – Stockholm Bankfilial**, with registered office at Hamngatan 15, Stockholm, 11147, Sweden, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Stockholm Bankfilial is also supervised by Finansinspektionen (Swedish FSA); registered with Finansinspektionen as a branch of J.P. Morgan SE. In **Belgium**, this material is distributed by **J.P. Morgan SE – Brussels Branch** with registered office at 35 Boulevard du Régent, 1000, Brussels, Belgium, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE Brussels Branch is also supervised by the National Bank of Belgium (NBB) and the Financial Services and Markets Authority (FSMA) in Belgium; registered with the NBB under registration number 0715.622.844. In **Greece**, this material is distributed by **J.P. Morgan SE – Athens Branch**, with its registered office at 3 Haritos Street, Athens, 10675, Greece, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Athens Branch is also supervised by Bank of Greece; registered with Bank of Greece as a branch of J.P. Morgan SE under code 124; Athens Chamber of Commerce Registered Number 158683760001; VAT Number 99676577. In **France**, this material is distributed by **J.P. Morgan SE – Paris Branch**, with its registered office at 14, Place Vendôme 75001 Paris, France, authorized by the Bundesanstaltfür Finanzdienstleistungsaufsicht(BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB) under code 842 422 972; J.P. Morgan SE – Paris Branch is also supervised by the French banking authorities the Autorité de Contrôle Prudentiel et de Résolution (ACPR) and the Autorité des Marchés Financiers (AMF). In **Switzerland**, this material is distributed by **J.P. Morgan (Suisse) SA**, with registered address at rue du Rhône, 35, 1204, Geneva, Switzerland, which is authorised and supervised by the Swiss Financial Market Supervisory Authority (FINMA) as a bank and a securities dealer in Switzerland.

In **Hong Kong**, this material is distributed by **JPMCB, Hong Kong branch**. JPMCB, Hong Kong branch is regulated by the Hong Kong Monetary Authority and the Securities and Futures Commission of Hong Kong. In Hong Ko.ng, we will cease to use your personal data for our marketing purposes without charge if you so request. In **Singapore**, this material is distributed by **JPMCB, Singapore branch**. JPMCB, Singapore branch is regulated by the Monetary Authority of Singapore. Dealing and advisory services and discretionary investment management services are provided to you by JPMCB, Hong Kong/Singapore branch (as notified to you). Banking and custody services are provided to you by JPMCB Singapore Branch. The contents of this document have not been reviewed by any regulatory authority in Hong Kong, Singapore or any other jurisdictions. You are advised to exercise caution in relation to this document. If you are in any doubt about any of the contents of this document, you should obtain independent professional advice. For materials which constitute product advertisement under the Securities and Futures Act and the Financial Advisers Act, this advertisement has not been reviewed by the Monetary Authority of Singapore. JPMorgan Chase Bank, N.A., a national banking association chartered under the laws of the United States, and as a body corporate, its shareholder's liability is limited.

With respect to countries in **Latin America**, the distribution of this material may be restricted in certain jurisdictions.

*Issued in **Australia** by JPMorgan Chase Bank, N.A. (ABN 43 074 112 011/AFS Licence No: 238367) and J.P. Morgan Securities LLC (ARBN 109293610).*

References to "J.P. Morgan" are to JPM, its subsidiaries and affiliates worldwide. "J.P. Morgan Private Bank" is the brand name for the private banking business conducted by JPM. This material is intended for your personal use and should not be circulated to or used by any other person, or duplicated for non-personal use, without our permission. If you have any questions or no longer wish to receive these communications, please contact your J.P. Morgan team.

© 2026 JPMorgan Chase & Co. All rights reserved.

J.P.Morgan