Goldman Sachs Research

ISSUE 143 | October 22, 2025 | 4:00 PM EDT

TOP of MIND

AI: IN A BUBBLE?

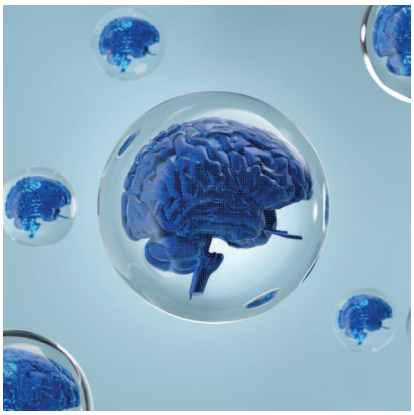

AI bubble concerns are back amid a rise in AI-exposed companies' valuations, ongoing massive AI spend, and the increasing circularity of the AI ecosystem. So, are bubble concerns warranted, or overblown? While GS' Eric Sheridan, Kash Rangan, Peter Oppenheimer, and Ryan Hammond all see some reasons for concern, they generally agree that the US tech sector is not in a bubble (at least not yet), with Sheridan more worried about the large gap between public and (higher) private market valuations. Sequoia's David Cahn sees it differently, arguing the only way to justify the large data center buildout forecasted by 2030 is AGI, while still seeing substantial opportunity in private AI application firms. But Bessemer's Byron Deeter is more optimistic on the AI capex boom, as is GS' Joseph Briggs. And NYU's Gary Marcus remains skeptical about the technology itself, at least in its current form. So, how should investors be positioned? We conclude that tech sector opportunities remain, but diversification makes sense.

"This isn't a 'hope-and-hype' cycle like the Dot-Com Era."
- Byron Deeter

"Some characteristics of the current period rhyme with past bubbles… [but] most of the Magnificent 7... generate outsized levels of free cash flow and engage in stock buybacks and pay dividends, which very few firms did in 1999."
- Eric Sheridan

"If you believe there's a data center bubble and there's going to be an overbuild of capacity, then you want to invest in consumers of compute."
- David Cahn

"Generative AI is still essentially autocomplete on steroids."
- Gary Marcus

WHAT'S INSIDE

INTERVIEWS WITH:

Byron Deeter, Partner, Bessemer Venture Partners

David Cahn, Partner, Sequoia Capital

Gary Marcus, Professor Emeritus, New York University

Eric Sheridan, US Internet Equity Research Analyst, Goldman Sachs;
Kash Rangan, US Software Equity Research Analyst, Goldman Sachs

Allison Nathan | allison.nathan@gs.com
Jenny Grimberg | jenny.grimberg@gs.com
Ashley Rhodes | ashley.rhodes@gs.com

Macro news and views

We provide a brief snapshot on the most important economies for the global markets

US

Latest GS proprietary datapoints/major changes in views

- We recently raised our 3Q/4Q25 US GDP growth forecasts to 3.3%/1.3% (QoQ, ann., from 2.8%/0.9%), partly to reflect a stronger-than-expected government contribution.
- We recently pulled forward our forecast for the end of Fed balance sheet runoff to Feb 2026 (from end-1Q26) following comments from Chair Powell.

Datapoints/trends we're focused on

- Government shutdown; while the shutdown continues to disrupt key data releases, we see value in alternative data.
- Tariffs; we believe President Trump's threat of an additional 100% tariff on China is more likely just an effort to gain negotiating leverage ahead of upcoming bilateral talks.

Japan

Latest GS proprietary datapoints/major changes in views

- No major changes in views.

Datapoints/trends we're focused on

- Japan politics; we think the election of LDP President Sanae Takaichi as prime minister won't have a significant near-term impact on fiscal or monetary policy.
- BoJ policy; we expect the next BoJ rate hike in January 2026, though we see potential for it to be delayed.
- US tariffs, which have led to a deterioration in the recurring profit outlook for Japanese firms, particularly those in industries with a high dependence on exports to the US.

Government shutdown: potentially longer to go

Kalshi odds: How long will the government shutdown last?, %

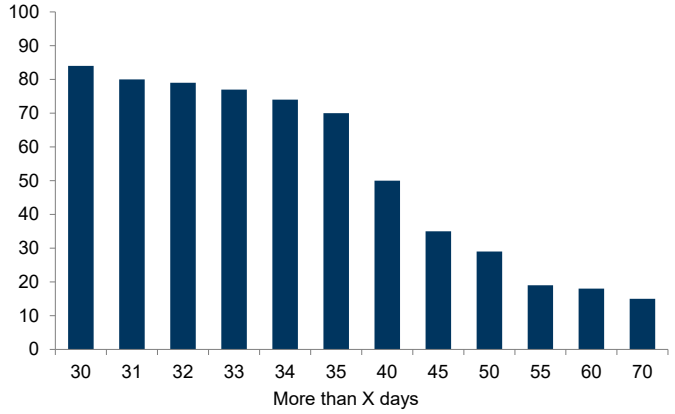

Source: Polymarket, Goldman Sachs GIR.

Japan: US tariffs cutting into corporate profitability

US export dependence (x-axis, % of US bound exports) vs. profit outlook for large manufacturers (y-axis, %, yoy)

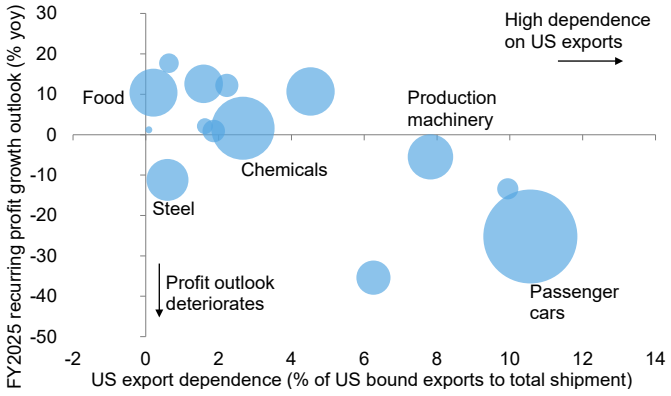

Source: BoJ, Ministry of Economy, Trade and Industry, Goldman Sachs GIR.

Europe

Latest GS proprietary datapoints/major changes in views

- We recently lowered our 2026 France GDP growth forecast to 0.8% (from 0.9%) to reflect renewed political uncertainty.

Datapoints/trends we're focused on

- ECB policy; we expect the ECB to remain on hold for the foreseeable future as Euro area wage growth and inflation expectations are consistent with 2% inflation.
- BoE policy; we expect the next BoE rate cut in February.
- European fiscal picture; while the UK and France face fiscal challenges, Germany's fiscal package should boost growth.
- EU trade diversification, which will likely have little macro impact, but may help some pressured industries recover.

Emerging Markets (EM)

Latest GS proprietary datapoints/major changes in views

- We recently raised our 2025/26 China headline PPI inflation forecasts to -2.6%/-0.7% yoy (vs. -2.8%/-1.0%) on stronger-than-expected PPI data, but reflation has a long way to go.

Datapoints/trends we're focused on

- China overcapacity, which could create headwinds to global manufacturing and lower DM GDP growth by 0.1-0.3pp/yr.
- EM growth, which has been resilient due to easier financial conditions, tariff front-running, cheaper oil, & AI investment.
- Argentina; the US' support represents a crucial lifeline for the Milei Administration ahead of midterm elections and could help anchor the extent of Peso depreciation.

Europe: in need of trade diversification

Market share in global exports, chg relative to 2015 average, pp

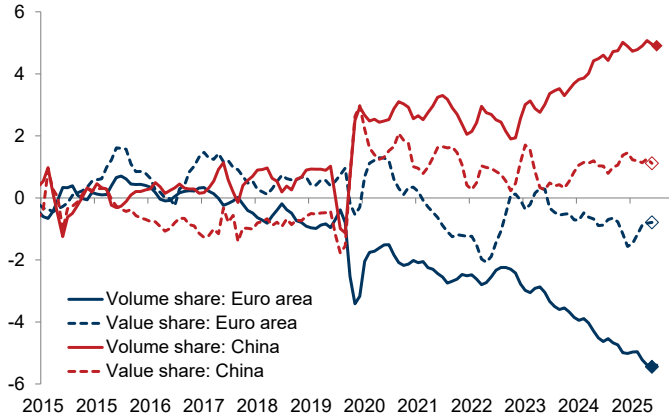

Source: CPB, IMF, Haver Analytics, Goldman Sachs GIR.

China capacity: exceeding global demand

China's capacity as a share of global demand, 2024, %

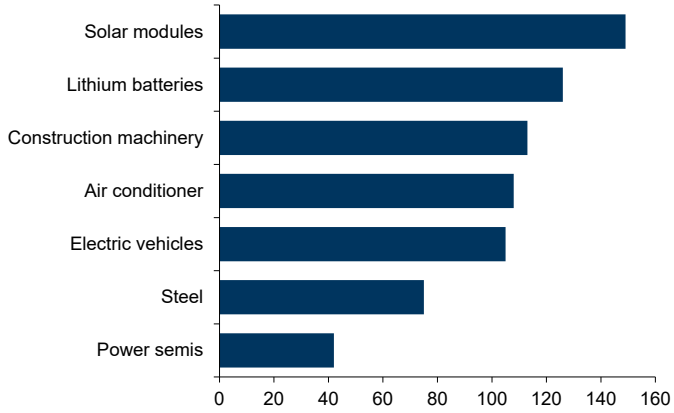

Source: Goldman Sachs GIR.

AI: in a bubble?

AI bubble concerns are back, and arguably more intense than ever, given a number of (worrying?) developments: a significant rise in the valuations of many AI-exposed companies, continued massive investments in the AI buildout, and the increasing circularity of the AI ecosystem, with model companies, infrastructure providers, and hyperscalers signing deals with each other that are blurring the boundaries between customers, suppliers, and capacity providers (see pg. 7). Amid these developments and growing worries that the megacap tech rally may be masking signs of weakness in the broader market, whether AI bubble concerns are warranted—or overblown—is Top of Mind.

We first ask GS US equity analysts Eric Sheridan (internet) and Kash Rangan (software). Sheridan notes that while some features of the current period rhyme with past bubbles, and the circularity of deals warrants caution, public market valuations and capital market activity levels remain below their Dot-Com peaks. He also points out that most of the Mag 7 continue to generate outsized free cash flows, engage in stock buybacks, and pay dividends—behavior seldom seen during the Dot-Com Bubble. So, he seems less inclined to call the situation in the public market a bubble today, though he concedes that "AI may just not be a bubble *yet*."

Rangan, for his part, sees few signs of a bubble in his coverage universe with many software valuations—if anything—too depressed. He more broadly worries about companies' increased reliance on debt to fund their AI ambitions (GS credit strategists Shamshad Ali and Ben Shumway quantify this debt buildup on pg. 20).

GS Chief Global Equity Strategist Peter Oppenheimer then digs deeper into the equity market parallels to past bubbles, also finding similarities but key differences—namely, the US tech leaders' current strong fundamentals, balance sheets, and AI market position—which lead him to agree with Sheridan that the US equity market is not in a bubble… yet. GS Senior US Equity Strategist Ryan Hammond also shares this view, finding little evidence of US tech sector froth outside of smaller pockets of the market like Quantum Computing.

All that said, Sheridan sees more reason to worry about the large gap between public and (higher) private market valuations of AI-exposed companies, with the latter largely based on revenues rather than profits and margins. He cautions that such a large and widening gap can indicate "risk in the system."

But David Cahn, Partner at Sequoia Capital, sees it differently. He argues that the only way to justify the massive data center buildout forecasted by 2030—which he estimates will cost several trillion dollars—is Artificial General Intelligence (AGI). So, in his words, "if you believe there's a data center bubble and there's going to be an overbuild of capacity, then you want to invest in consumers of compute. If you're a consumer of compute, having an overcapacity of compute means your gross margin goes up and your COGS goes down."

So, Cahn sees substantial opportunity in AI application companies that largely reside in the private market today: "Prices are elevated, but business model quality is high. These companies are doing interesting things. There's a lot of opportunity, and there's a lot of money to be made there."

Byron Deeter, Partner at Bessemer Venture Partners, is less concerned about the amount of capex pouring into the space, with AI "reshaping businesses in ways previously unimaginable." While he acknowledges that valuations are high today, he sees them as largely justified by AI firms' underlying fundamentals and revenue potential. And he views the ecosystem's circularity "less as artificial inflation and more as a reflection of strategic interdependence across the AI value chain." So, he believes that AI bubble concerns are overblown, and that "this isn't a 'hope-and-hype' cycle like the Dot-Com Era."

And GS Head of the Global Economics team Joseph Briggs is perhaps the most optimistic about the AI capex boom, arguing that the economic value generated by AI will ultimately justify the spend. Specifically, he estimates that generative AI will create $20tn in economic value (in present-discounted value terms), $8tn of which will flow to US companies, as it unlocks significant productivity gains. And although the range of outcomes remains wide, he notes that plausible revenue estimates generally exceed current cumulative AI investment forecasts even before factoring in AGI's potential emergence, though whether the firms making the investments today will be the ultimate beneficiaries remains less clear.

But Gary Marcus, Professor Emeritus at New York University, remains skeptical about the promise of the technology in its current form and the amount of capex being spent on it. Marcus explains that generative AI remains far from AGI today despite some incremental improvements over the past two years, with the technology "still essentially autocomplete on steroids." While he is encouraged by the industry's recent shift away from believing that large language models would be the solution to AGI and toward more promising approaches like neurosymbolic AI, he still sees significant challenges ahead on the road to AGI.

So, what does all this mean for how investors should be positioned? Within tech, Sheridan sees value in stocks well-positioned to benefit from potential AI disruption and underappreciated growth stories, while Cahn and Deeter see substantial investment opportunity in the mostly-private AI application layer.

But Oppenheimer also sees value in looking beyond the tech sector given the extreme level of market concentration today and ever-present risk of a market correction. He continues to recommend diversifying across regions, factors, and sectors—a strategy that has paid off this year.

Allison Nathan, Editor

Email: allison.nathan@gs.com
Tel: 212-357-7504
Goldman Sachs & Co. LLC

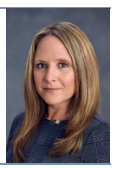

A discussion on artificial intelligence (AI)

Eric Sheridan and Kash Rangan are Senior Equity Research Analysts at Goldman Sachs covering US internet and software, respectively. Below, they argue that while some characteristics of the current period rhyme with past bubbles, we're not in an AI bubble today.

Allison Nathan: When we last spoke in June 2024, you explained that computing cycles typically follow an "IPA" progression—infrastructure first, platforms next, and applications last. Then, AI was very much in the infrastructure phase. Where does it stand today, and how does that compare to expectations?

Eric Sheridan: The AI infrastructure buildout is still very much ongoing, with the amount of capital and spend surprising to the upside as the demand for computational power, or compute, continues to outstrip Nvidia's ability to provide the chips fueling the compute. So, the model companies, including the hyperscalers, have been unable to scale inference—the process by which a trained AI model applies its knowledge to new data—and launch more products. At the platform layer, a few companies with the necessary scale of capital and talent have transitioned from just running foundational models to building API solutions or applications on top of them. And at the application layer, more consumer applications have emerged, mostly through the usage of OpenAI's ChatGPT and Google's Gemini, though significant monetization remains elusive as most consumers today are utilizing the free versions of these applications.

Kash Rangan: The infrastructure buildout has certainly gone on much longer than anybody expected as more foundational models have emerged. Three years ago, it was just ChatGPT. Today, six major foundational models are competing to be the model of the future: ChatGPT, Gemini, Anthropic's Claude, Microsoft's Phi, Mistral's models, and xAI's Grok. These models are very hungry because they require training, which is driving the significant demand for compute. But this activity is also beginning to percolate up to the platform layer. Databricks, Snowflake, and MongoDB—platform companies that provide solutions for data management and storage for enterprise use—are all experiencing an acceleration in revenue growth partly owing to AI. Salesforce's Data Cloud, another AI-enabled platform, is a nearly $1bn revenue business with a 100% growth rate. So, the platform layer is in a vastly better position today than when we last spoke. But the application layer for enterprise has been a disappointment. While enterprise AI applications have shown some signs of life, the revenue figures are well below where I expected them to be 1-2 years ago.

Allison Nathan: What's driving that disappointment?

Eric Sheridan: Most computing cycles are adopted first at the consumer level. The iPhone provides a relevant analogy; consumers owned iPhones well before iPhones were used in workplace settings. Blackberries remained ubiquitous in the corporate world for some time after consumers adopted the iPhone because the consumer chose their preferred device untethered from enterprise budget, compliance, and regulatory considerations. The AI application cycle is playing out the same way—enterprises are trying to work within their existing IT budgets, gain employee buy-in, and/or build products their customers want to adopt that may disrupt their own workflows, which has resulted in a relatively slow pace of enterprise adoption. Consumers encounter no such obstacles, so consumer adoption has proceeded more quickly.

Allison Nathan: Amid these disappointments, is the still massive amount of AI spend a reason to worry?

Kash Rangan: Some faith exists in the US capital system that if we get AI right, it will do wonders for productivity, scientific discovery, etc. So, the system feels that investing such a huge amount of capital is a risk worth taking. If that proves true, we all stand to benefit massively. That doesn't mean there won't be any pain along the way. The massive investment in fiber optic networks in the late 1990s resulted in a bubble. But that capacity allowed the internet to flourish, leaving everyone ultimately better off. Companies are once again investing heavily in capacity, this time in the form of AI chips, data centers, etc. Most of that capital is coming from hyperscalers' cash flows. And that may be what ultimately saves the day.

Eric Sheridan: The significant rise in AI spend has certainly led many investors to question the potential ROI. Most investors we speak to seem to struggle to justify the return profile on the $3-4 trillion of cumulative AI spend Nvidia estimates will occur by the end of the decade unless AI delivers on its enormous economic and societal promises. Admittedly, computing cycles sometimes play out in unpredictable ways. In the 1990s, people worried about an overbuilding of desktop computing as network and networking equipment companies expanded aggressively to capitalize on the PC boom. But then came Netflix and the browser and portal wars, which fueled a large rise in desktop usage. And when the buildout of spectrum and wireless towers began, nobody imagined that three billion people would one day own powerful smartphones. So, AI could surpass our wildest expectations, but visibility into the end state remains elusive.

But I will be brutally frank in saying that if the dollars keep rising, we will struggle to answer the ROI question based on what we know today. In every computing cycle I've ever analyzed, that has eventually led to a trough of disillusionment. I would be shocked if we avoided one this time. Beyond that, in any technology cycle, typically only 2-3 companies in the same vertical earn an excess return on their cost of capital. And here too, I see no reason why the AI cycle will prove any different.

Allison Nathan: So, especially given the high valuations of AI-exposed companies, are we in an AI bubble today?

Eric Sheridan: I am not trying to be flippant, but there has been more talk of an AI bubble over the last three years than of

the Dot-Com and Housing Bubbles in their midst. It's true that some characteristics of the current period rhyme with past bubbles. Private market valuations are well above public market valuations. Private companies are predominantly being valued on revenue and incremental revenue growth, with seemingly less focus on profits and margins. Public companies, by contrast, are still being evaluated on free cash flow, return on capital, margins, and how they trade relative to the market. So, the private and public markets are using two diametrically different valuation frameworks, and history shows that a widening between them can indicate risk in the system.

Of course, public market valuations are also above historical norms. But they are still below their 2000 peak. Capital market activity is also well below 1998-99 and 2007-08 levels, with the average deal today also much larger, indicating a more selective IPO market. Now, that's arguably a duration argument—AI may just not be a bubble *yet*. But, in 1999, firms that had no revenue were the ones with the most exuberant valuations. Today, most of the Magnificent 7—which trade at an aggregate P/E of 31x vs. 23x for the market, with Alphabet and Meta trading only slightly above the market multiple—generate outsized levels of free cash flow and engage in stock buybacks and pay dividends, which very few firms did in 1999.

Kash Rangan: I don't see a bubble in my coverage universe. Many software stocks are trading at depressed valuations given concerns that AI could disrupt their end-markets, whether through job dislocation or by enabling the writing of software in a much more cost-effective manner, which would increase competition in the application software space. At the platform layer, most companies are trading at more acceptable—if not frothy—valuations generally consistent with their strong financial results.

A larger story in the AI space that warrants watching is the emergence of a debt-fueled capital cycle. As we've discussed, most of the capital deployed to fund AI projects has so far come from hyperscalers utilizing cash flows from their core businesses. But now, entities are being funded with 80% debt and 20% equity, with the equity portion often backed by collateral from the sponsoring entity. Oracle recently completed an $18bn bond sale to fund its AI ambitions, and non-hyperscalers like CoreWeave have also secured significant debt financing. So, leverage is starting to emerge in the system, making it even more vital that the firms driving the need for capital hit their revenue and earnings targets.

Allison Nathan: Are you worried this is a house of cards with Nvidia investing in OpenAI, OpenAI pledging to spend on compute from Oracle, Oracle buying Nvidia chips, etc.?

Eric Sheridan: The circularity does make me nervous. This is another example of the current period rhyming with the Dot-Com Bubble. I started my career as a telecom analyst and investor in the late 1990s/early 2000s—the era of Global Crossing, Level 3 Communications, and Qwest. These companies traded capacity with each other using debt, with one company's revenue being another's capacity. The entire house of cards fell apart when the debt load became too high. When the revenue was ultimately untangled, it was nowhere near as much as people thought. While the capacity was eventually absorbed by the mid-2000s, returns for many investors were quite low for a very long time depending on the entry point. So, when companies start investing in other companies, suppliers start investing in capacity providers, and debt starts piling up, some caution is warranted.

Allison Nathan: So, how should investors be positioned?

Eric Sheridan: Investors should focus on industries where AI is either disrupting the industry structure or re-accelerating revenue growth. For example, AI will likely disrupt advertising profit pools in the coming years. Investors can also find value in underappreciated growth stories.

Allison Nathan: What are you watching that could make you less optimistic on the AI investment thesis?

Eric Sheridan: I continue to monitor utility, adoption, monetization, and free cash flows. Ultimately, it would be difficult to continue arguing the investment thesis if companies spend in a way that puts their free cash flow generation at real risk, and that could represent a tipping point in the market. So, I will be watching for any signs that companies are cutting dividends or buybacks or piling on debt.

Kash Rangan: As we've discussed, I will be closely monitoring how the debt-fueled credit cycle evolves. If the cycle doesn't cooperate, it would have a ripple effect throughout the tech ecosystem. And if something goes wrong with AI, application software companies would make good hedges because they have been under AI-related duress. If AI doesn't live up to its disruptive potential, these companies could benefit.

Q&A with Byron Deeter

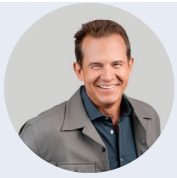

Byron Deeter is a Partner at Bessemer Venture Partners. Below, he argues that AI bubble concerns are overblown, with AI spend likely to eventually pay off.

Portions of this Q&A originally appeared as part of Goldman Sachs Exchanges' Great Investors series. The views stated herein are those of the interviewee and do not necessarily reflect those of Goldman Sachs.

Q: How would you characterize the AI investing landscape, and what—if anything—does it remind you of?

A: The AI investment opportunity is unprecedented. While we have experienced major transitions in software, with the first wave of cloud computing representing the closest analogy, the scale and scope of AI investments is unlike anything we've seen before. Billions of dollars are being poured into the infrastructure buildout. Companies like Anthropic and Perplexity, as well as sector-focused firms such as Abridge and MaintainX, are at the forefront alongside established players like Canva and Intercom, which are reinventing themselves through AI. And AI's impact extends beyond traditional technology layers, pulling in service revenues and reshaping businesses in ways previously unimaginable. So, AI truly represents the technology opportunity of our lifetimes. These developments will be discussed for generations, and our grandchildren will recount the early days of AI as a pivotal moment in history.

Q: Many question whether the return on AI investments will ever justify the huge amounts of AI capex. What's your view?

A: The AI infrastructure layer is already scaling and monetizing in many ways. You either believe that the trillions of dollars that are being monetized will carry over through software value capture, or this whole thing is a hoax and will collapse. I believe in the former. AI fundamentals are strong, and this value will inevitably roll through to the end-user through monetization, with applications capturing much of the value and the foundational model companies becoming the next generation of hyperscalers. Beyond the revenue story, AI's societal impact, in areas like education, healthcare, and scientific discovery, will arguably be some of the most enduring benefits.

Q: Are concerns about AI being a bubble warranted, or overblown?

A: Such concerns are overblown. While valuations are currently elevated, they are warranted by the underlying revenue potential of these companies. Looking back at the cloud computing cycle, just about every company that we brought in for a third meeting or a partner presentation in hindsight we should have invested in. Those thoughts sit heavily with me whenever people look at AI and say that the valuations are crazy and investors need to pull back. To me, this is the cloud opportunity all over again. AI is the next phase of cloud, and it is much bigger than the original.

Q: Some people worry about the wide gap between public and private valuations, pointing to this as rhyming with the Dot-Com Bubble. What's your response?

A: The comparison to the Dot-Com Bubble is an oversimplification. Unlike that era, many of today's leading AI companies—OpenAI, Anthropic, and even emerging players like Databricks and ClickHouse—are already generating meaningful, recurring revenue with enterprise-grade customers. We're seeing real demand and revenue pull-through from both consumer and large-scale enterprise adoption. The speed at which these companies are monetizing—particularly through APIs, licensing models, and partnerships—is unlike anything we've seen in prior tech cycles. So, again, while valuations are certainly ambitious, tangible revenue growth and extraordinary market momentum support them. This isn't a "hope-and-hype" cycle like the Dot-Com Era—it's a scale and monetization cycle. The size of the opportunity and real economic value being created justify a significant portion of the private valuation premium.

Q: Are you at all concerned that this is just a house of cards with Nvidia investing in OpenAI, OpenAI pledging to spend on compute from Oracle, Oracle buying Nvidia chips, etc.?

A: It's a fair concern. Some circularity exists in the ecosystem right now, with strategic investments often reinforcing each other's growth narratives. These dynamics can amplify short-term momentum and sometimes blur the picture of what's truly fundamental. That said, I see these relationships less as artificial inflation and more as a reflection of strategic interdependence across the AI value chain. Nvidia, for instance, has become a core supplier to nearly every player in the stack; Oracle's capital allocation is a bet on catching up in cloud; and OpenAI's commitments reflect genuine compute needs driven by user demand. The key question is whether these relationships are backed by real, recurring usage—and so far, the underlying demand for compute, model access, and AI-powered applications appears very real. While there's some reflexivity, I wouldn't characterize it as a house of cards.

Q: Given elevated valuations, what should investors look at to identify the most compelling investment opportunities?

A: Early on, investors can't look at business models and unit economics. In these especially dynamic markets, team matters much more than ever. Incumbent firms are recognizing that and joining the AI bidding wars to recruit top researchers and scientists. I have personally come to back the jockey more than the horse at this point, and that's probably the right way to go here. Valuations are high, but the opportunity set is also structurally larger than anything we've seen before. The total addressable market for AI—spanning chips, cloud infrastructure, model licensing, and downstream applications—is likely an order of magnitude greater than prior software cycles. Real opportunity also exists for AI to take over services revenue.

Q: How do you see the AI space evolving?

A: A progression similar to previous tech waves will likely occur, with value moving from the infrastructure layer to the application layer. The AI wave is playing out similarly to the cloud wave, with Nvidia, Broadcom, and the hyperscalers monetizing first and the foundational models following. It's been incredibly frustrating for public investors because the only way they can really take advantage is to gain exposure to the hardware companies, whereas the venture market is starting to see it working from the bottom up. And companies are staying private longer and going deeper into the private capital stack. So, public market investors have had much less exposure to the AI theme, while the private market opportunity has—and likely will—remain robust.

AI: an increasingly circular ecosystem

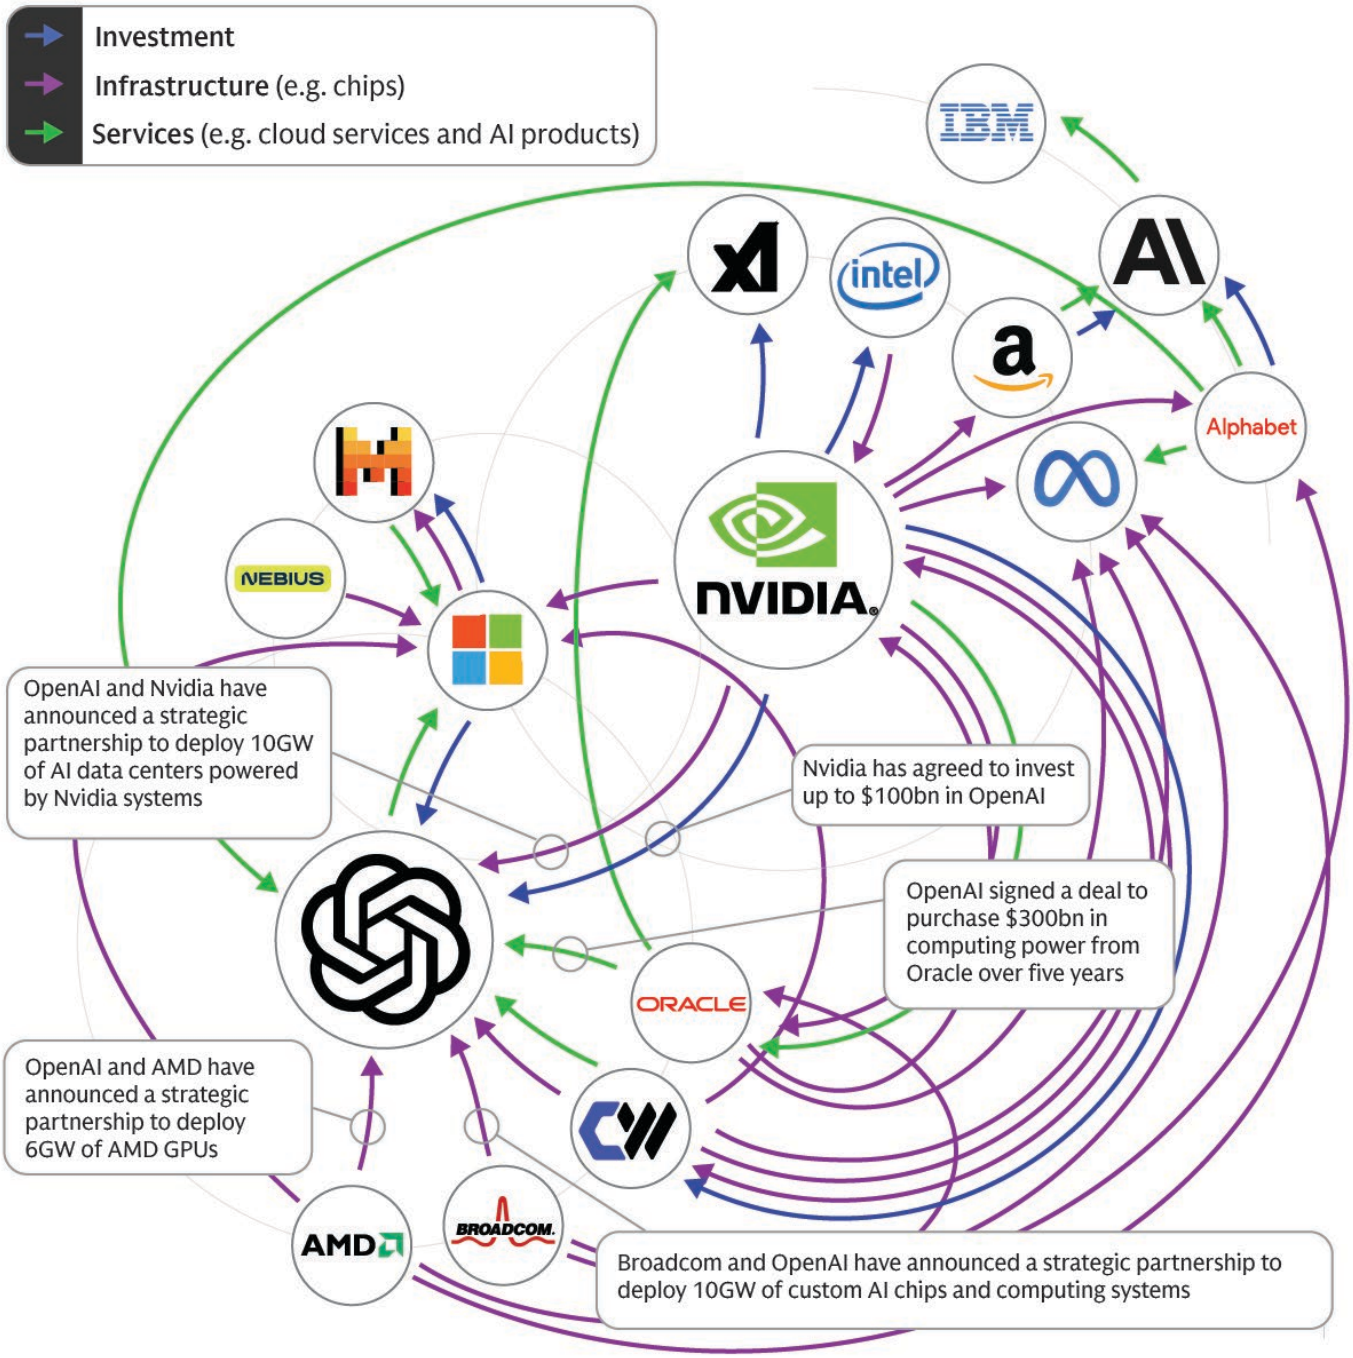

Note: Exhibit does not represent an exhaustive view of the AI ecosystem but rather the most prominent AI model companies, infrastructure providers, and hyperscalers.
Source: Company announcements, various news sources, compiled by Goldman Sachs GIR.

Interview with David Cahn

David Cahn is a Partner at Sequoia Capital. Below, he explains why the AI ecosystem has a long way to go to generate the revenues necessary to justify current data center spending. He argues that there is a capex bubble, but sees significant opportunity for AI long-term, especially at the application layer.

The views stated herein are those of the interviewee and do not necessarily reflect those of Goldman Sachs.

Allison Nathan: Why did you decide to explore AI's $600bn question?

David Cahn: I've been investing in AI for about eight years. I was on the board of Weights & Biases from the Series A, Runway before they invented Stable Diffusion, which is the technology behind all the AI image generators and video generators now, Replit, one of the AI coding companies, and backed HuggingFace. I'm very optimistic about AI. My core belief is that when you and I are 80, AI is going to have fundamentally changed the world. We're in for something that's going to be really interesting over a long time horizon.

Where my June 2024 piece, "AI's $600B Question" came from, originally AI's $200 billion question, is teasing out this long-term optimism about AI and then narrowing in on today's specific capex cycle. There's been this simplistic narrative of, either you're an AI bull or you're an AI bear. I've been trying to tease out nuance there, which is to say, you can believe that AI is going to be the greatest technology that humanity has ever invented. And you can also believe that we're overbuilding very specific data centers in a very specific time horizon. The resistance has been this idea that, no matter how much we spend on capex, no amount is too much. And it's been difficult to have that conversation without putting numbers around it.

So, the goal of my analysis was to put some numbers around it. If you take Nvidia's run rate revenue and you double it, roughly that is the total data center spending in any given year. And that's because half of data center capex is energy and generators and batteries and things that are not AI chips. You started with, in 2023, roughly $50 billion of run rate revenue in Q4. You got $100 billion of data center capex and $200 billion of implied revenue that you needed to generate for AI. In 2024, a lot of people reached out to me to redo the analysis. And that was the piece that went viral around AI's $600 billion question, which again was the same math, $150 billion of expected Q4 run rate for Nvidia. You double that to get $300 billion of expected data center spending, and then you double that to get the revenue needed to pay back that data center investment.

Allison Nathan: Will these investments generate high enough returns to justify the current spending?

David Cahn: I think the question is returns for who. Let's talk about the supply chain a little bit. A lot of the AI debate and discussion tends to be myopically focused on specific companies. And that makes sense because Wall Street is looking at their quarterly results. But it's more interesting in my view to look at the ecosystem and to think of AI as a supply chain. Especially a hardware supply chain.

Any supply chain has a vertical stack top to bottom. At the top is a customer who pays you money. Then there's Microsoft Azure. And then under that, there's an OpenAI API. Under that, there's a data center and there's some real estate guys who build that data center. In that data center, there's generators, interconnect, and chips. Those chips are made by Nvidia. Nvidia fabs those chips at TSMC. TSMC buys from ASML.

In a healthy, mature supply chain, a dollar flows down the top from the customer and then some percentage of those dollars go to each player in the supply chain. One observation that a lot of people have made is, if a dollar comes in at the top, Nvidia keeps $1.20 today. So Nvidia is capturing a lot of the value in the supply chain today. And the second observation is that there's not enough dollars coming in at the top, regardless of how that dollar gets divided, so someone must be propping up the supply chain, by definition.

Now, this is something that investors do well. Think of food delivery as a business where the unit economics in the beginning were kind of backwards. Today, it's a fantastic business. It's highly profitable. The customer buys dinner, and my Dasher makes money, and the restaurant makes money, and DoorDash makes money, and everybody makes money. That's a healthy, mature supply chain at the end of the day. So, the question in AI is, how do we get to that kind of mature end state, who pays for the burn along the way, and how deep in the hole are they? Somebody is holding up the supply chain. It's not the customer. So, who is it?

A year ago, the answer was Microsoft, Amazon, and the cloud companies. Microsoft was calling the generator guy, and Microsoft's buying five years of generator capacity. That's propping up the supply chain. The supply chain wouldn't work without Microsoft doing that. Microsoft is guaranteeing a 20-year lease to the real estate developers. The real estate developer goes and builds a data center. Microsoft and the cloud giants were playing this pivotal role where they were holding up everybody else in the supply chain.

Fast forward a year, a lot has changed. If we ask ourselves, what is the biggest change in AI in the last 12 months? I think it is Microsoft and Amazon stepping back and a new set of players stepping forward to fill their shoes.

And this is where we're getting the circular deals, where the vendors themselves are now stepping forward. We've shifted from a hyperscaler funded buildout to a new world.

Allison Nathan: The players have changed, but what does that mean?

David Cahn: I think Microsoft and Amazon realized, okay, we've built a lot. We're in a good spot. Those companies are very well positioned. They have a lot of data center capacity.

They're the best in the world at building data centers. It's not like they've built a small amount of AI data center capacity.

The narrative on Wall Street is that these companies are not bullish enough on AI because they don't want to build more data center capacity and they're fundamentally losers, whereas Oracle and companies like this are winners. That's a very simplistic narrative. Those companies are behaving prudently. The intellectual rationale behind what they're doing is robust.

It's interesting to see that, when you're in a bubble, good behavior is penalized and bad behavior is rewarded. This is the vicious cycle of a bubble. And this influences the internal dynamics at companies, where there are always competing factions. It becomes self-reinforcing. I'm fascinated by these kinds of human dynamics and how they influence markets, which I try to tease out in my writing.

Allison Nathan: Are we in an AI bubble?

David Cahn: It's the tale of two AIs—maybe even the tale of three AIs. Let's get into it. There's three things that people talk about in AI. One is the data center bubble. That's fundamentally a public market phenomenon. That's a semiconductor and hardware bubble. And there are companies that have performed really well because of that.

Second, there are AI applications, and that is ChatGPT, which is an amazing business, that is Anthropic, that is Cursor, and that's what I'm trying to invest in. If you believe there's a data center bubble and there's going to be an overbuild of capacity, then you want to invest in consumers of compute. If you're a consumer of compute, having an overcapacity of compute means your gross margin goes up and your COGS goes down.

One new phenomenon is what we call the 0 to 100 million club. These are companies that overnight basically go 0 to 100 million of revenue. And it comes down to these companies building something useful. People are willing to adopt it. If you build something useful today, you can make a lot of revenue really quickly. This is a second element of AI, the application layer. Prices are elevated, but business model quality is high. These companies are doing interesting things. There's a lot of opportunity, and there's a lot of money to be made there.

The third AI story is AGI, which is a whole different topic.

Allison Nathan: The circularity of investments on the compute side have made the market nervous. How sustainable is all of that?

David Cahn: One thing that's happening on these circular deals is that we are substituting dollars for gigawatts. And it's obfuscating what's going on. So let me unpack that. We used to talk about run rate revenue and data center revenue. Now it's all gigawatts. We're now measuring data centers by the amount of energy they consume. There are so many billions of dollars that people don't want to use billions anymore. People throw around two forecasts. One is that we're going to have 100 gigawatts of energy built out for AI by 2030. The other is the bull case that 250 gigawatts will be built.

The dollar conversion of these numbers really makes you think. 100 gigawatts, let's say it's $40 billion per gigawatt. 40 billion times 100 gigawatts is $4 trillion of data center spend. Using my AI $600 billion math, you double that to get the implied revenue that you need to generate, and we're now at AI's $8 trillion question. And then you use the 250 gigawatt forecast and you do the same math, 40 times 250 gigawatts is $10 trillion. And then you double that and you get AI's $20 trillion question.

I don't think it's an accident that everybody is now quoting these numbers in gigawatts, because if you quoted them in dollars, you would have to do that as a percentage of US GDP. That ties into this third bucket around AGI that I've been writing about, which is to the extent that we're going to get 100 gigawatts of energy built out, the only way to justify that is AGI.

Allison Nathan: Do you think that investors, though, in the compute space are going to be in trouble?

David Cahn: The perspective that always stands out as most interesting is the real estate investor perspective. You ask real estate investors, "okay, you're financing the data center build out. So, at the end of the day, you should really believe in what the demand profile looks like in 15 years." And what they all say is, "I'm a credit investor. I'm buying Microsoft credit and I'm getting a spread. Microsoft is my guarantor. I don't actually have a view on 15-year AI demand." I think this is where credit funding can obfuscate the demand signal, because you would think, there's all this capital coming in. Well, no, it's just Microsoft and some nice yield on that. Credit investors don't see themselves as holding the bag on demand in 15 years.

It's the big companies and their shareholders who are funding the buildout. People argue that it is a positive thing that healthy companies with big balance sheets are funding this. I agree.

Allison Nathan: We covered a lot. Is there anything that we've missed that you wanted to cover?

David Cahn: My closing thought would be, AI is going to change the world. People who try to narrow this down into AI-good or AI-bad are incorrect. AI is probably the most important technology of the next 50 years. It is incumbent on all of us to figure out how it's going to change our lives. I recently wrote a piece about how young recent college grads should think about their career. I think it's crazy not to factor AI into your thinking about your future career and what type of business to work in.

For investors, though, we have this luxury or curse, that we are exposed to raw ground truth, and we don't get to delude ourselves. For us, time horizons do matter. These conversations about capex and overbuilding, these are things we need to grapple with. But for most people, the real question is, how do you lever your life to AI in some way?

This is the tension. On the one hand, we're overinvesting in data centers. On the other hand, most people are probably not changing their lives enough to factor in how AI is going to change their lives. We need to be thinking, what is our life going to be in 50 years? With a long enough time horizon, we'll avoid making investments that aren't going to pan out. And then, second, we'll make good career and life decisions, and we'll benefit from this technology. It's incumbent on all of us to think about that, and that's why these conversations are so important.

The AI spending boom isn't too big

Joseph Briggs argues that the economic value generated by generative AI will ultimately justify the AI capex boom

Since mid-2023, we have anticipated a large AI investment cycle driven by an initial surge in hardware spend to train AI models and run AI queries. But the size, speed, and circularity of recent AI investment announcements have raised questions around the sustainability of AI capex. Despite these concerns, we continue to see current investment levels as sustainable, though it is less clear whether companies making the investments today will emerge as AI winners.

AI investment has risen by around $200-300bn since 2023

Annualized AI hardware spending in US national accounts, chg vs. 2022, $bn

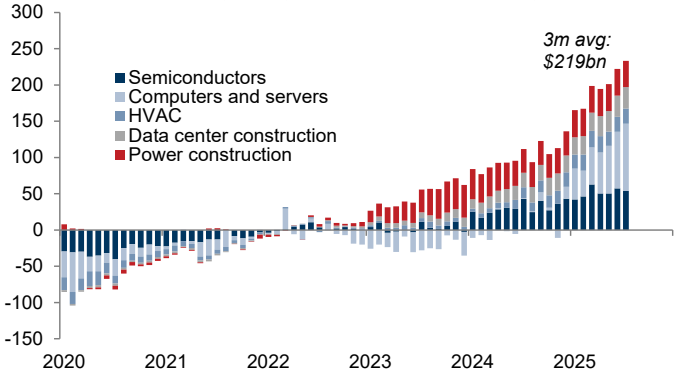

Source: Bureau of Economic Analysis, Goldman Sachs GIR.

Capex drivers remain supportive

The technological backdrop remains supportive of AI capex for two reasons.

First, generative AI still appears set to deliver rapid acceleration in task automation that will drive labor cost savings and boost productivity. We estimate a 15% gross uplift to economy-wide US labor productivity after full adoption, which we expect will realize over a 10-year period. Academic studies and company anecdotes support this, pointing to 25-30% average productivity gains following AI application deployment. While use cases remain fairly narrow, these early estimates highlight generative AI's potential to deliver transformative productivity uplifts.

Second, unlocking these productivity gains requires significant computational power and energy, with recent trends suggesting that demand for both will continue rising. The computational power necessary to train LLMs continues to grow more rapidly (floating-point operations (FLOPs); 400% per year) than computational costs are falling (FLOPs/$; 40% per year). Training queries (350%) and the number of frontier AI models (125%) are also growing rapidly, while energy efficiency is improving but at a relatively slower pace (40%). The punchline is that the differential between demand growth and computing cost declines remains wide.

Demand growth will likely continue outpacing technological cost declines in the near term, partly because frontier LLMs continue to improve with size. While predicting at what point the technological incentives to invest will diminish is difficult, recent evidence suggests that we are not at that point yet.

Sizing generative AI's macroeconomic value

We also still see the macroeconomic justification for AI investment as compelling and are less concerned about the dollar amount of AI capex.

Many have flagged the unprecedented amount of AI capex currently being deployed, with the cumulative AI-driven data center and infrastructure buildout likely to total multi-trillions of dollars. For example, Nvidia CEO Jensen Huang recently highlighted that AI infrastructure spend could total $3-$4tn by 2030. Our equity analysts similarly project sizable investments, particularly in data centers and power, with hyperscaler capex alone projected to total $1.4tn in 2025-2027.

While the AI investment buildout is admittedly larger than prior cycles in nominal dollar terms, it looks more benign when appropriately scaled. Historical infrastructure investment impulses generally peaked at 2-5% of GDP, while investment during the electrification of manufacturing in the 1920s and IT boom in the late 1990s peaked at ~1.5-2%. AI investment in the US over the last 12 months remains below 1% of GDP, a large—but not outsized—impulse by historical standards.

Large investment cycles have preceded prior general purpose technology productivity booms

Investment in general purpose technologies, share of GDP vs. pre-productivity boom level, pp

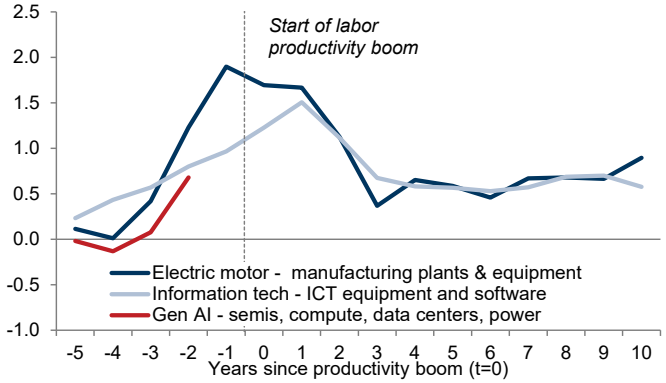

Source: Bureau of Economic Analysis, Goldman Sachs GIR.

More importantly, we believe generative AI's potential economic gains justify a multi trillion-dollar investment cycle.

To approximate future AI revenues, we estimate a present-discounted value (PDV) based on the following assumptions:

- *Productivity:* We assume a baseline 15% gross uplift to US labor productivity and GDP (equivalent to $4.5tn in economic value creation in today's dollars). We also consider "less powerful" (8% uplift) and "more powerful" (27% uplift) AI scenarios.
- *Timeline:* We assume that company-level adoption largely takes place over a 10-year period from 2027-2037, with a four-year intra-firm lag between adoption and full realization of productivity gains. We also consider scenarios where AI adoption takes place over five and 15 years.
- *Capital share of AI value-add:* We assume a 41% capital share of the incremental value add created by AI-driven productivity gains, in line with the economy-wide average. We also consider scenarios where the capital share of AI economic value creation corresponds to the 25th (28%) and 75th (60%) percentiles of the industry cross-section.
- *Discount rate:* We assume a discount rate of 15%, corresponding to the 75th percentile of the WACC for publicly traded AI companies. We also consider scenarios using higher (20%) and lower (10%) discount rates.

Under these assumptions, we estimate that generative AI in the US will create $20tn PDV of economic value, of which $8tn will flow to US companies as capital revenues. We estimate PDV of capital revenues ranging from $5-19tn under the alternative scenarios.

PDV of capital revenue from AI should exceed cumulative capex

Present discounted value of AI capital revenues, $tn

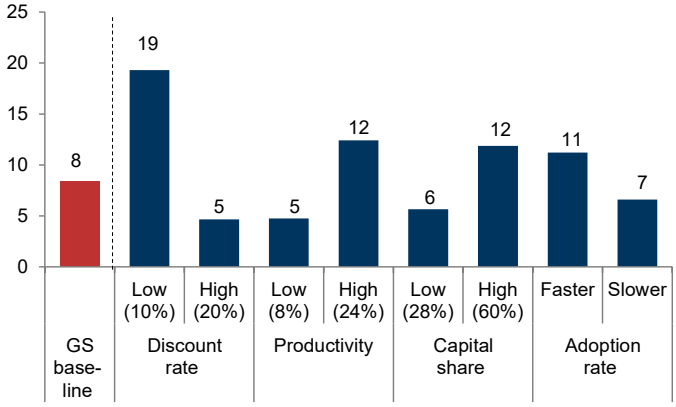

Source: Goldman Sachs GIR.

The key takeaway is that the PDV of generative AI capital revenue exceeds projected AI-related capex both under our baseline and alternative assumptions, suggesting that current and anticipated levels of AI capex are justified. We reach this conclusion without factoring in potential foreign profits, new profit pools, and the emergence of AGI, suggesting a meaningful cushion to the AI spending outlook.

Market structure and distribution of revenue

While the macro backdrop still looks solid, we see valid concerns around whether companies making AI investments today will benefit from this spending. This will depend critically on the timing and distribution of revenues across the AI stack.

On timing, investment in semiconductors and servers—which account for $112bn of the $240bn in AI spend in the US national accounts, respectively—will depreciate quickly given the rapid improvements in computing hardware. We estimate an 18% depreciation rate for current AI capex, raising potential for a mismatch between the timing of infrastructure build and revenue realization[1].

Timing considerations may be less important if AI investors can capture an outsized share of the long-run value, but "first-movers" have shown mixed performance in prior infrastructure builds. History suggests three patterns:

1. As was the case for UK canals in the late 1700s and early 1800s, US IT investment in the 1980s and 1990s, and (in some cases) for US railroads in the 1800s, first movers can command outsized returns due to high investment and switching costs.
2. In many cases—including the buildout of UK railroads in the 1800s and more recent buildout of fiber optic cables and US telecom—first-mover returns were quite poor, as an initial overbuild gave way to a subsequent bust that allowed "fast followers" to capture outsized returns by purchasing assets at low valuations.
3. In other cases—namely the development of UK turnpikes in the 1700s and US electricity in the early 1900s—first-mover returns were limited by regulation and the reorganization of capital as public utilities.

These historical precedents highlight that a complex set of factors, including timing, regulation, and market competition, determine the ultimate winners from infrastructure builds.

The current AI stack's market structure provides limited clarity into whether the companies leading AI investments today will be long-run AI winners. Competition is high at the application layer, reasonably high at the foundational model (despite OpenAI's lead) and data center layer, and more limited at the semiconductor layer (where Nvidia dominates design and TSMC dominates production). On the surface, this suggests outsized returns for AI hardware providers, consistent with the pricing of the AI trade so far.

The key question for investment sustainability, however, is whether incumbent advantages at the foundational model layer—i.e., the hyperscalers that are driving AI capex today—will lead to outsized returns over the longer-run that justify continued investment. History again provides mixed signals.

On the one hand, first movers can command outsized returns when high costs and limited access to key assets restrict competition, or when vertical integration allows first movers to control the entire production stack. Along these lines, the high costs of AI investment, OpenAI and other hyperscalers' recent moves to lock up key computing resources, and increasing signs of vertical integration suggest that hyperscalers may be successful in maintaining their first-mover advantage.

On the other hand, first-mover advantages have historically proven smaller when technology and market growth are happening rapidly, IP protections are limited, or when significant uncertainty around end-user applications exists. These patterns argue against the advantages for today's AI leaders that would justify sustained capex.

Other key determinants of first-mover advantages remain uncertain. High switching costs could lead to outsized returns for model and application providers, but most companies today are diversifying across foundational models, which could limit switching costs. Similarly, it is unclear whether first movers will benefit from network effects that yield advantages in determining product market standards and quality.

The case for continued spending

Ultimately, we think that the enormous economic value that generative AI promises justifies the current investment in AI infrastructure, and that overall levels of AI investment appear sustainable as long as companies expect that investment today will generate outsized returns over the long run.

So, we expect that the solid macro backdrop will support capex for as long as companies believe that 1) first-mover advantage will allow them to capture an outsized share of AI productivity-related revenues or 2) continued investment in compute capacity will drive improvement in model performance and potential AGI development—which could drive massive profits.

Joseph Briggs, Head of the Global Economics team

Email: joseph.briggs@gs.com *Goldman Sachs & Co. LLC*
Tel: 212-902-2163

[1] We assume a 5-year lifespan for tech equipment and use the BEA's depreciation estimates for HVAC, data center structures, and power investment.

AI: still strong hardware investment...

AI investment growth has remained strong...

Actual and forecasted revenues by AI-exposed sectors of Russell 3000, index, 4Q19=100

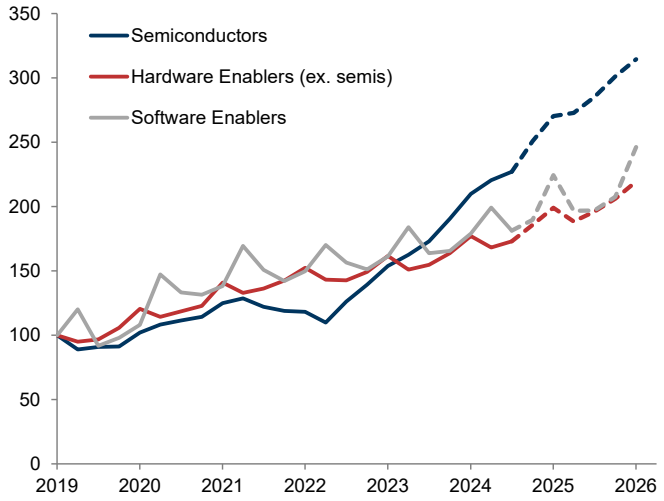

Note: Dashed lines in this chart indicate consensus forecasts.
Source: FactSet, Goldman Sachs GIR.

...with the hyperscalers continuing to make significant investments in AI infrastructure

AI infrastructure capex spend by company, $bn

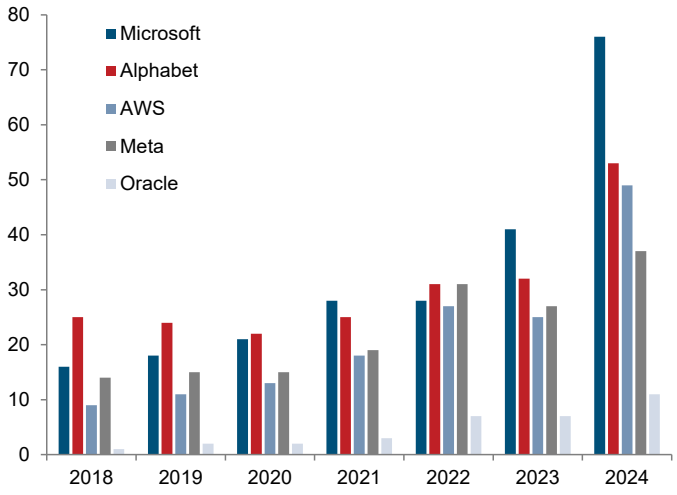

Source: Company data, Goldman Sachs GIR.

AI investment growth has been concentrated in the US...

AI-related investment in hardware*: national accounts, log index**, 3Q22=100

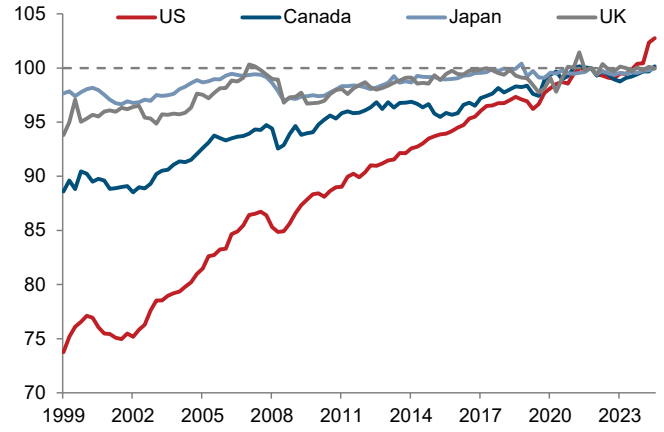

**Categories related to information and communication technologies (ICT) and communications equipment.*
Source: Haver Analytics, Goldman Sachs GIR.

...as US AI-related hardware and software investment has risen

AI-related investment in software: national accounts, log index**, 3Q22=100

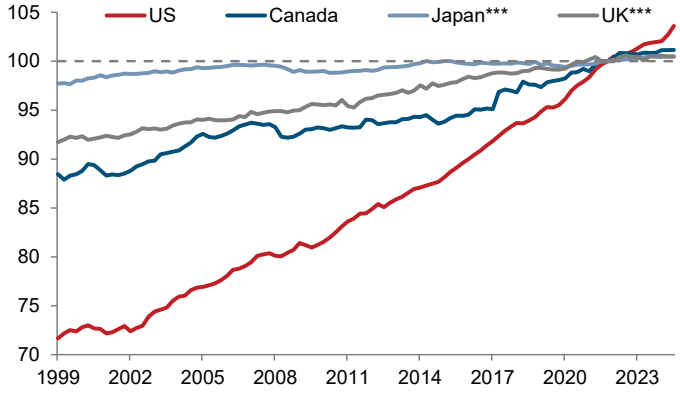

***Shown as log index because software investment grows by different exponential rates across countries. Steady growth in investment would appear as a line with a constant slope, while accelerating growth would appear as a line with an increasing slope.*
****Intellectual property products category, which includes software.*
Source: Haver Analytics, Goldman Sachs GIR.

The market has significantly upgraded its revenue expectations across the AI hardware stack, but not for software...

Change in Russell 3000 consensus revenue forecasts since March 2023, $bn, annualized

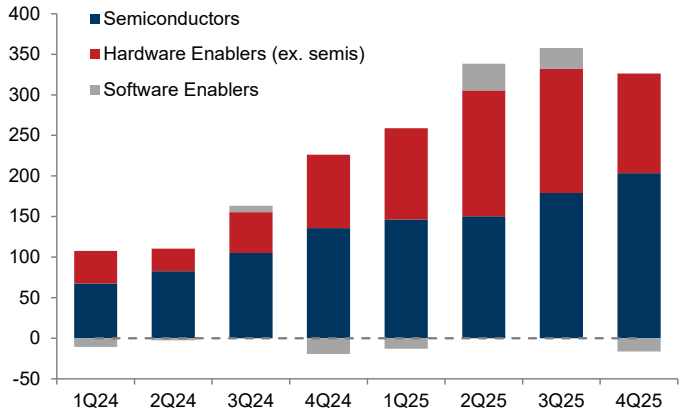

Source: FactSet, Goldman Sachs GIR.

...while the broader AI space has seen a relatively small upgrade

Change in Russell 3000 consensus revenue forecasts since March 2023, $bn, annualized

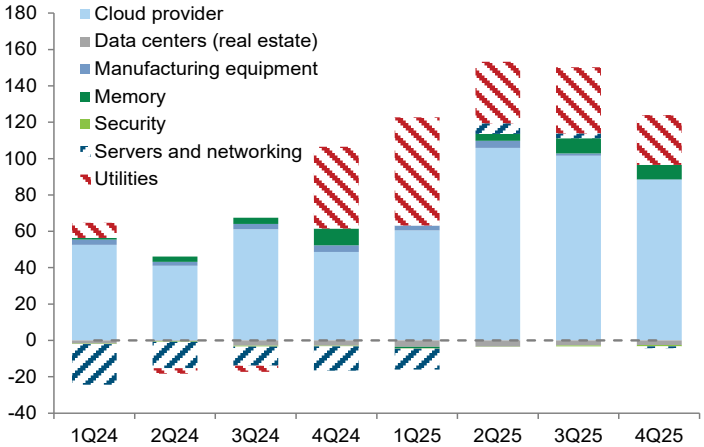

Source: FactSet, Goldman Sachs GIR.

...slowing adoption, but signs of impact

The pace of corporate AI adoption slowed in the recent quarter...

Economy-wide firm AI adoption rate, %

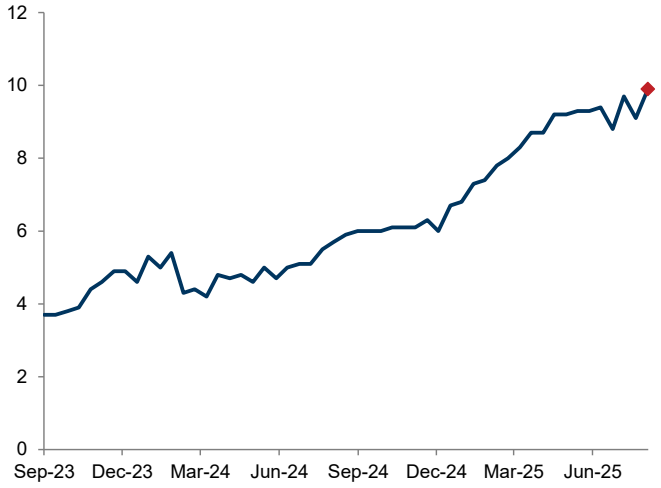

Source: Census Bureau, Goldman Sachs GIR.

...with 9.9% of companies across all industries having adopted AI, up only slightly from 9.3% in Q2...

Change in share of US firms using AI since quarter ago, pp

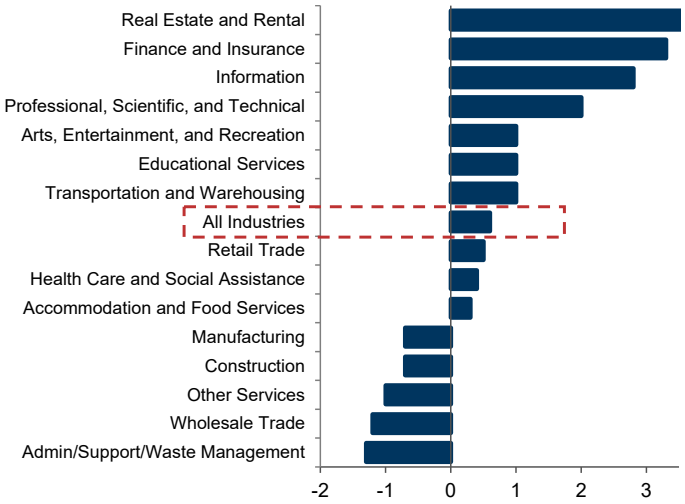

Source: Census Bureau, Goldman Sachs GIR.

...as many companies cite security and data quality concerns as barriers to adoption

Source	Date	Future Adoption
Cisco	October 2025	AI personal assistants and robotics control agents were among the top planned future agentic AI use cases. Overall, most companies expect the next 12-36m to be the critical window to drive ROI on AI projects.
Anthropic	September 2025	The share of Claude conversations geared around task automation (vs. augmentation) has continued to edge up, suggesting a growing confidence in delegating complete tasks to AI. This could raise risks of displacement of workers who were previously performing these tasks.
Gartner	June 2025	Security threats, finding the right use cases, and data availability/ quality were among the top barriers for AI implementation reported.
Paypal and Reimagine Main Street	June 2025	Respondents who had not yet fully adopted AI cited security concerns, resource constraints, and uncertain value add as the top barriers faced.
Bain & Company	May 2025	Data security and privacy concerns have grown substantially since initial surveys from early 2023. Infrastructure buildout and scaling remains a challenge to full adoption, with 75% of companies reporting difficulty in finding in-house expertise.
EPAM	April 2025	Improving productivity and greater operational efficiency were the top goals for firms' AI-related initiatives. Surveyed firms also report a 14% yoy planned increase in AI spending in 2025. 43% of firms plan to hire more AI-related roles to support these initiatives, and machine learning engineers and AI researchers are the most in-demand positions. Data security, insufficient cloud infrastructure, and outdated tech stacks were the most commonly cited concerns for adoption.

Source: Goldman Sachs GIR.

But adoption rates remain high in tech and other digitally-enabled fields...

Share of US firms using AI, top 15 subsectors, %

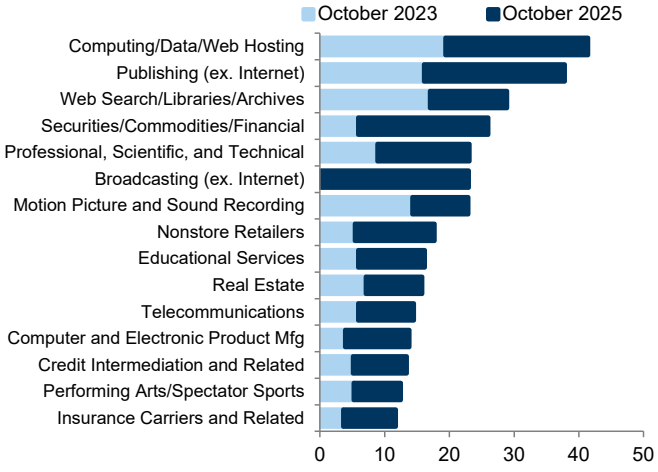

Source: Census Bureau, Goldman Sachs GIR.

...and clear signs of AI's impact are evident in the tech sector, as tech's share of total employment has fallen below trend...

Tech sector* share of total employment, %

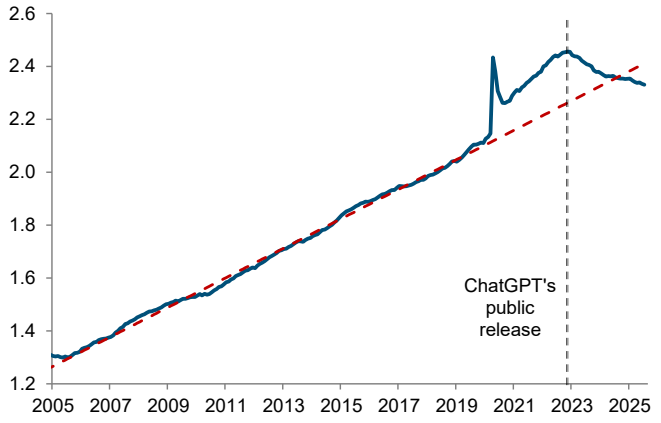

**Tech is the software publishers, data processing and related, web search and related, and computer systems design subsectors. Red line is 2007-2019 trend.*
Source: Bureau of Labor Statistics, Haver Analytics, Goldman Sachs GIR.

...and youth unemployment in tech-exposed sectors has risen

Unemployment rate, %

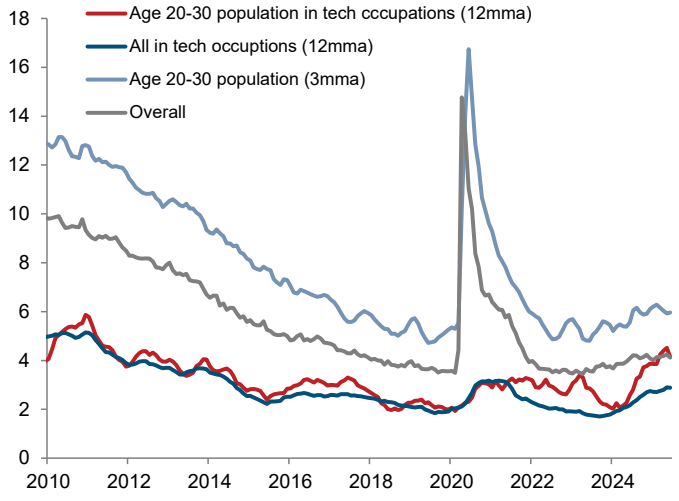

Source: Bureau of Labor Statistics, iPUMS, Goldman Sachs GIR.

Special thanks to GS GIR global economists Joseph Briggs and Sarah Dong and equity analyst Jim Schneider and team for charts.

Interview with Gary Marcus

Gary Marcus is Professor Emeritus at New York University and author of *Taming Silicon Valley*. Below, he argues that, despite some incremental improvements in generative AI tools, the core challenges persist, highlighting the need for an alternative approach to AI development.

The views stated herein are those of the interviewee and do not necessarily reflect those of Goldman Sachs.

Jenny Grimberg: When we last spoke in July 2023, you were quite critical of the capabilities of generative AI tools. Given the technological developments since then, are you still skeptical?

Gary Marcus: Yes. While generative AI has been billed as tantamount to Artificial General Intelligence (AGI), it remains far from it. Some incremental technical improvements have been made, and these tools are useful for tasks like brainstorming and coding, but generative AI is still essentially autocomplete on steroids, trained to predict the next word in a sequence. Current AI tools lack a genuine understanding of the world and struggle with reliably executing even basic tasks as issues like hallucinations and the inability to fact-check their own output persist. So, these systems still cannot be trusted. Every day people send me new examples of these models hallucinating in ways they really shouldn't be, bringing to mind the famous line from Forrest Gump about life being like a box of chocolates. Well, life is like an LLM—you never know what you're going to get.

Many anticipated that subsequent iterations of AI models, such as ChatGPT-5, would miraculously solve these problems. But, as I predicted, they have not. These issues are inherent to the architecture. So, AI's widely promised productivity gains and transformative economic impact remain largely unrealized, and the predicted billion-dollar businesses run by a single employee have not materialized. In short, despite some improvements, the fundamental challenges remain.

Jenny Grimberg: Since we last spoke, several companies' AI models have reportedly achieved medal-level performance in the International Mathematical Olympiad (IMO), a difficult milestone that many people thought was years away. Isn't that a sign of meaningful progress?

Gary Marcus: I don't yet put great stock in these achievements. Google's DeepMind has claimed notable progress with its model, and to their credit, they used actual IMO judges to evaluate it. But the model has yet to be released, and its broader capabilities remain unknown. I expect the model will excel in the math and coding, where it's possible to produce synthetic data tailored for training and verification. This approach works for well-defined domains, such as solving quadratic equations, but it doesn't translate well to open-ended real-world problems. For example, in business forecasting, synthetic data that accurately mimics the real world can't easily be created. So, the applicability is limited.

Technological developments are generally marked by a recurring pattern of people claiming that "the next one" will solve this or that, but these promises don't often materialize. While I am almost certain that more advanced AI will eventually emerge, I don't see sufficient evidence in the technical literature to suggest that we're close to that breakthrough.

Jenny Grimberg: What has gone wrong on a technical level that has led generative AI to underdeliver on its proponents' promises of where the tech would be today?

Gary Marcus: The core issue is the reliance on black box statistical AI with its opaque internal workings—data is fed into the model and results are generated, but what happens inside can't be seen or traced. Unlike classical AI, which allows for the examination of each step and debugging of errors directly, with LLMs, no clear way exists to pinpoint what went wrong in the event of an error. The only real options are retraining the entire model or utilizing reinforcement learning to patch errors. But these fixes are often temporary and can create new problems elsewhere—essentially creating a game of whack-a-mole. So, these systems have come to resemble the original Mechanical Turk, with teams of contractors working behind the scenes to manually correct AI errors who never foresee everything that requires fixing. Ultimately, this approach lacks the reliability required for critical fields like medicine or finance, and I don't believe it will ever overcome these limitations. So, committing fully to black box AI was a mistake, in my view.

Jenny Grimberg: When we last spoke, you were the lone voice in the wilderness espousing these concerns. Has the broader population now caught up?

Gary Marcus: The mood has shifted significantly. Two years ago, my skepticism was met with open hostility. Most dismissed my concerns and invested billions of dollars based on hope and hype, believing these systems would inevitably improve. For years, the industry put their faith in so-called scaling laws—the idea that bigger models would always yield better results—which I have long cautioned against. However, with recent disappointments surrounding releases like Llama 4, Grok 4, and GPT-5, it has become clear that scaling isn't yielding the hoped-for results, with the phrase "diminishing returns" entering mainstream discourse. Even industry leaders have begun publicly acknowledging the limitations, with Satya Nadella recently saying that these scaling laws might not be physical laws of the universe. That was the first time I heard someone in the industry echo what I have been saying for quite some time. So, the industry has begun to rethink its assumptions and become more open to exploring alternatives.

Jenny Grimberg: What alternatives do you see promise in?

Gary Marcus: One alternative that I have long advocated for, including in my book *The Algebraic Mind*, is neurosymbolic AI. Neurosymbolic AI combines two traditions in AI: neural networks, which are fast, statistical, and automatic—akin to Daniel Kahneman's "System 1" thinking—and classical AI, which uses symbolic reasoning, logic, and rule-based systems, much like Kahneman's "System 2" thinking. The aim is to integrate both approaches to create more robust, interpretable,

and flexible AI systems. For years, the industry dismissed this method, insisting that the brain doesn't use symbols despite clear evidence to the contrary from logicians, computer programmers, and even studies on infant cognition. The dominant AI approach became single feed-forward neural networks: input in, output out, with pattern recognition but little deliberative reasoning. However, recent developments show a quiet shift—all the major LLM companies have begun embedding symbolic tools like code interpreters within their LLMs. AlphaFold, an AI program developed by DeepMind that predicts protein structures, is the best example of neurosymbolic AI come to life. So, I am glad to see steps in this direction. But it is not enough.

A core challenge remains: building AI models that can truly understand and represent the real world. Humans constantly build mental models—of their environments, social interactions, or even fictional worlds like Harry Potter—enabling us to reason, adapt, and make predictions. LLMs try to shortcut this, and they've suffered for it.

Jenny Grimberg: In 2023, you argued that AGI could be achieved this century. So, are you now less optimistic?

Gary Marcus: I'm actually slightly more optimistic. The last five years have been characterized by an intellectual monoculture as companies repeatedly pursued similar approaches and achieved similar, flawed results—a cycle that epitomizes the definition of insanity. Despite the introduction of dozens of new models since ChatGPT-4, each striving to eliminate hallucinations and errors, none have meaningfully succeeded. But, as we've discussed, the industry is finally beginning to recognize that LLMs aren't the solution.

I continue to believe that achieving AGI this century is entirely plausible. While I'm often labeled as an AI critic, I am instead an AI realist—I want AI to succeed, but I remain unconvinced that LLMs are the path forward. The field currently lacks a robust theoretical foundation, and I believe that real progress will require a fundamental change in approach. I liken it to the scientific race in the 1920s to uncover the molecular basis of genes: for years, researchers fixated on proteins, until a paradigm shift led to the groundbreaking discovery of DNA. In much the same way, I expect that, eventually, a radically new approach to AI will emerge, sparking rapid and transformative advances once the right paradigm is found.

Jenny Grimberg: Some AI model companies have generated large revenues. Has that impressed you at all?

Gary Marcus: The top model companies are undeniably generating significant revenue, particularly from coding applications as developers readily pay for tools to automate routine tasks. However, much of this revenue still stems from companies running proof-of-concept pilots rather than full-scale deployments, leaving the sustainability of these revenues in question. The high and rising costs associated with generative AI—especially as the industry shifts toward more resource-intensive domains like video generation—further complicate the picture. OpenAI, for example, has seen its revenue increase significantly but is simultaneously burning through billions of dollars, setting new records for losses every quarter. So, the real question is not just about revenue, but profitability, long-term viability, and how much companies and investors are willing to gamble that those will materialize.

Society is essentially making a trillion-dollar bet that these systems will become reliable and broadly useful, but persistent challenges make this a high-stakes gamble based on very abstract optimism that these models will improve. Unless future AI models deliver dramatic breakthroughs—rather than the incremental improvements we've seen so far—the massive expenditures are difficult to justify.

Jenny Grimberg: So, is AI in a bubble today?

Gary Marcus: AI is certainly in a financial bubble. While LLMs offer genuine utility and are undoubtedly here to stay, the economics don't add up. As more companies build AI models, the technology becomes commoditized, eliminating meaningful competitive advantages and driving down prices. Most revenue still comes from developers, which may reach a few billion dollars annually, but this is dwarfed by the trillion-dollar infrastructure investments that companies like OpenAI, Microsoft, and Nvidia are making. So, their lofty valuations are detached from actual profits, and even OpenAI's own optimistic forecasts don't anticipate profitability before 2030. The business model hinges on hope for discovering new, vastly larger revenue streams and for future models to deliver substantial breakthroughs. Not to mention all the circular economic shenanigans where these companies are passing a dollar back and forth, except it's $100 billion. So, this is undoubtedly a Wile E. Coyote moment.

Ultimately, returns drive markets, and current AI ventures aren't generating enough profit to justify their lofty valuations. People are starting to get that message. And if enough people get that message on the same day, it will start to look like a bank run. Just as no one could predict the exact moment the Dutch Tulip Mania would end, it's impossible to know when this bubble will burst—but it will burst.

> "While I'm often labeled as an AI critic, I am instead an AI realist."

Jenny Grimberg: What else concerns you about AI today?

Gary Marcus: I remain deeply concerned about generative AI's broader societal impact. The rise of generative AI technology has enabled widespread automated disinformation that threatens democracies worldwide, a proliferation of non-consensual deepfakes, and a rise in cybercrime—particularly through automated phishing attacks, which have become much easier to execute. Environmental costs are also mounting owing to AI data centers' immense energy demands. And with the benefits so far largely limited to modest productivity gains for programmers, generative AI currently appears to be a net negative for society. Yet, we continue to allow AI companies to operate without meaningful regulation or accountability, concentrating substantial power in the hands of a few. Ultimately, regulation is needed to balance the costs and benefits of generative AI technology, though such measures remain unpopular in Washington. At the very least, these issues demand careful consideration.

Why we are not in a bubble… yet

Peter Oppenheimer argues that the US equity market is not in a bubble… yet

The present market environment exhibits several characteristics that rhyme with previous bubbles, including a notable rise in absolute valuations, a high degree of market concentration, increased capital intensity among leading companies, and the emergence of vendor financing. The sheer scale of market dominance of leading technology companies is also striking: the top five US technology companies collectively hold a value exceeding the combined size of the Eurostoxx 50, the UK, India, Japan, and Canada, representing approximately 16% of the entire global public equity market. And the ten largest US stocks, eight of which are technology-related, account for nearly 25% of the global equity market with a valuation near $25 trillion.

The top 10 US companies dominate the world equity market

2024 GDP vs. market capitalization, $tn

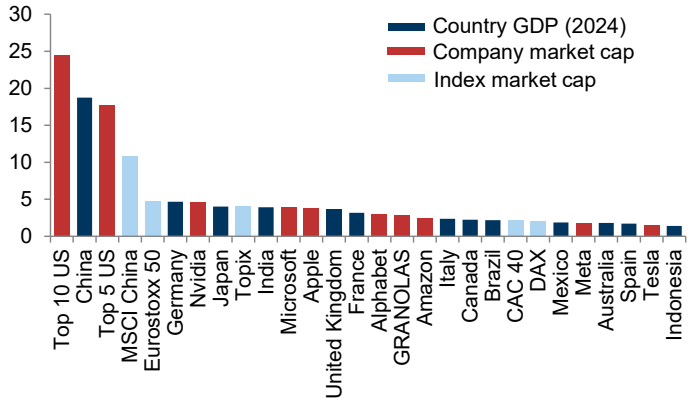

Source: IMF, FactSet, Datastream, Goldman Sachs GIR.

The rise in technology stocks and the excitement surrounding AI share similarities with past bubbles—including the UK's Canal Mania in the 1790s, the Railway Bubbles of the 1840s (UK) and 1873 (US), and the global Technology Bubble of the 1990s—which centered on a new technology or innovation that generated significant investor excitement, leading to substantial investment from both existing and new companies. One study found that 73% of 51 major innovations between 1825 and 2000 exhibited equity price bubbles, with their magnitude increasing alongside the radicalness, potential for indirect network effects, and public visibility of the innovations[1].

The current enthusiasm for AI resonates with these historical patterns, particularly the Dot-Com Bubble of the late 1990s. A technological sea change appears to be at a critical commercialization point, like the early internet, promising higher future growth. However, this presents a threefold challenge: accurately valuing the future market size, determining the timeline for commercialization and scaling, and identifying the primary beneficiaries of future growth. The inherent uncertainty often leads investors to purchase options in numerous companies, hoping to identify the next dominant player. Ultimately, bubbles form when the aggregate value of companies associated with an innovation surpasses the future potential cash flows they are likely to generate.

Despite these similarities, crucial differences suggest that the current environment may not yet constitute a full-blown bubble. These include:

1. **Fundamental growth vs. speculation:** Fundamental growth and robust earnings have thus far, been the primary drivers of technology sector appreciation, rather than irrational speculation about future potential. This contrasts with many past bubbles when expectations of future growth and market dominance drove companies at the epicenter of the fervor rather than proven success. The extraordinary rise in the earnings per share (EPS) of the technology sector, particularly since the Global Financial Crisis, highlights this sustained fundamental performance.

Tech earnings have outstripped those of the global market

12m trailing EPS (USD), index, January 2009=100

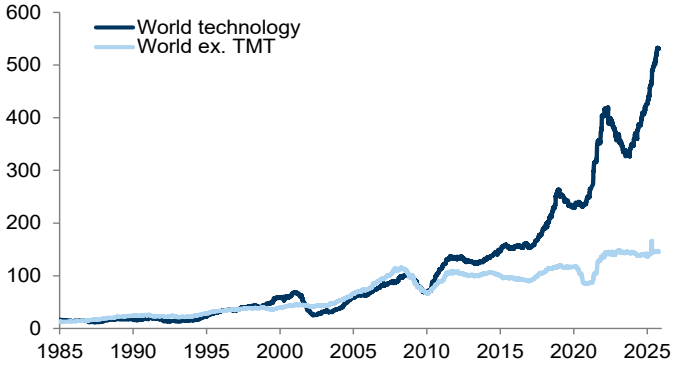

Source: Datastream, Worldscope, Goldman Sachs GIR.

2. **Strong balance sheets:** The companies that have experienced the strongest returns possess unusually strong balance sheets and staggering profit growth. This financial strength provides a buffer that was often absent in companies at the heart of previous bubbles.
3. **Incumbent dominance in AI:** A few established incumbents have dominated the AI space to date. Most historical bubbles, conversely, formed during periods of intense competition as both investors and new entrants flocked into the space. While competition in AI is increasing, established players have led the initial charge.

Valuations: stretched but not at bubble levels

We also examine several valuation metrics and find that while big tech valuations are stretched, they are not at bubble levels:

1. **Past bubble comparison:** The median 24-month forward P/E ratio for the Magnificent 7 tech stocks is 25x (24.5x excluding Tesla) (based on consensus forecasts). This is roughly half the equivalent valuation of the biggest seven companies during the late 1990s Dot-Com Bubble. Enterprise value to sales (EV to sales) ratios are also considerably lower than those of dominant companies in the late 1990s. So, while high, current valuations are generally not at the extreme levels typically observed at the peak of a financial bubble.
2. **Price-to-earnings growth (PEG) ratio comparison:** The PEG ratio for technology stocks, which compares valuation to earnings growth, remains comparable to other stocks and well below late 1990s levels. A more conservative PEG ratio based on 12-month forward P/E versus trailing 3-year EPS growth puts the current ratio at 1.6x, significantly lower than the 3.7x observed at the Dot-Com Bubble's peak.

[1] Chancellor, E., and Kramer, C. (2000). Devil Take the Hindmost: A History of Financial Speculation. Finance and Development, 37.

3. **ROE vs. price-to-book (P/B) comparison:** The US tech sector boasts historically high ROEs, which partially justifies a high P/B ratio. While the sector appears somewhat stretched on this metric, it is not as extreme as in 1999/2000, when P/B was similar, but ROE was much lower.
4. **Dividend discount model (DDM):** A one-stage DDM, assuming an equity risk premium (ERP) of 4.5%, implies an 8% per annum nominal perpetual growth for the US Technology, Media, and Telecom (TMT) sector. This is high compared to recent years (4-5%) but lower than the 10% implied in 1999/2000. We find similar results using a three-stage model.

Capex spending and financing

While valuation metrics are generally less extreme than those typically observed during previous bubbles, a surge in investment and capex is occurring, which has been a typical feature of past technology bubbles. Companies at the forefront of major technological innovations often fail to achieve the returns implied by their high valuations as marginal costs fall and capacity increases. Meanwhile, new entrants may leverage existing capex by providing new products and services. While current incumbents have generated extraordinary profit growth with a relatively light capital intensity model, their capex spend has rapidly increased since the emergence of ChatGPT in 2022. This raises the risk that future returns on this capital may be less than the market currently implies.

That said, we have yet to see the boom in equity or debt financing that has typically occurred in prior bubbles. While capex-to-sales ratios for the technology sector in the US are rising sharply, capex to free cash flow remains relatively low.

Capex-to-sales ratios have risen for US tech but capex to free cash flow is still relatively low

US technology sector capex to sales (lhs, %) vs. capex to free cash flow (rhs, %)

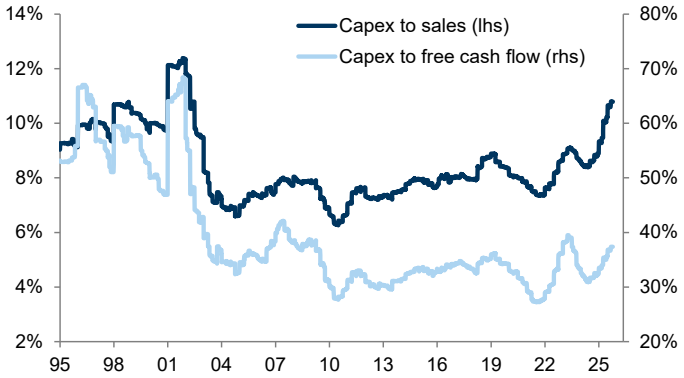

Source: Datastream, Goldman Sachs GIR.

This doesn't mean that leading companies cannot correct or be compromised by new entrants, but rather that their funding models, at least until now, have been more sustainable. Accordingly, the fallout of any market correction is likely to be less systemic, particularly given the strength of bank balance sheets. That said, rising examples of credit financing and vendor financing models imply that risks are rising.

Market concentration

While valuations and funding models suggest we are not yet in a bubble, the significant outperformance of technology has indeed led to a dangerous degree of market concentration across geographies, sectors, and stocks. The US market's consistent outperformance over the past 15 years—which has resulted in it accounting for over 60% of the global stock market—is largely attributable to the technology sector's record share of the US index. Such concentration is extreme, with the top 10 US companies alone making up nearly a quarter of the global public equity market. However, high market concentration alone does not equate to a bubble. Historically, dominant sectors have maintained their leadership for extended periods, reflecting prevailing economic drivers.

The largest firms are usually in the dominant sector of the time

% Market capitalization, % of net income before 1974

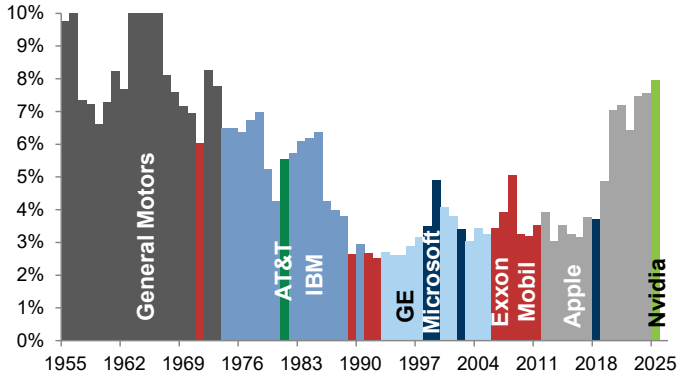

Source: Fortune 500, Datastream, Goldman Sachs GIR.

Investment implications: focus on diversification

All told, while the risk of a bubble forming persists, our current assessment suggests we are not yet in one. The healthy state of private sector balance sheets, reduced leverage financing the current spending boom, and strong bank balance sheets should also mitigate the risk of economy-wide repercussions if investor confidence in AI wanes. Nevertheless, a market correction remains possible if technology and AI growth prospects are de-rated.

Given these risks, we recommend investors focus on several diversification strategies:

1. **Regional diversification:** Despite lower tech exposure outside the US, regional performances have been similar year-to-date. Non-US equity markets have outperformed in US dollar terms, with some regions like Germany, Italy, Spain, and Korea experiencing significant gains.
2. **Factor and sector diversification:** The investment landscape is broadening beyond the "growth" versus "value" bifurcation that dominated the post-financial crisis decade, which provides an expanding opportunity set across various styles and sectors.
3. **Infrastructure interdependence:** The fortunes of leading tech stocks are increasingly tied to physical infrastructure. Surging demand for electricity, for instance, necessitates real spending in energy generation and distribution, broadening growth prospects for industries such as capital goods, energy, resources, real estate, and transport.
4. **Intra-technology diversification:** Within the technology sector itself, returns are likely to broaden. While current leaders may remain dominant, rapid innovation, particularly in machine learning and AI, should create a new wave of tech superstars capable of generating new products and services, leveraging the current capex boom.

Peter Oppenheimer, Chief Global Equity Strategist

Email: peter.oppenheimer@gs.com *Goldman Sachs International*
Tel: 44-20-7552-5782

A recent history of AI developments

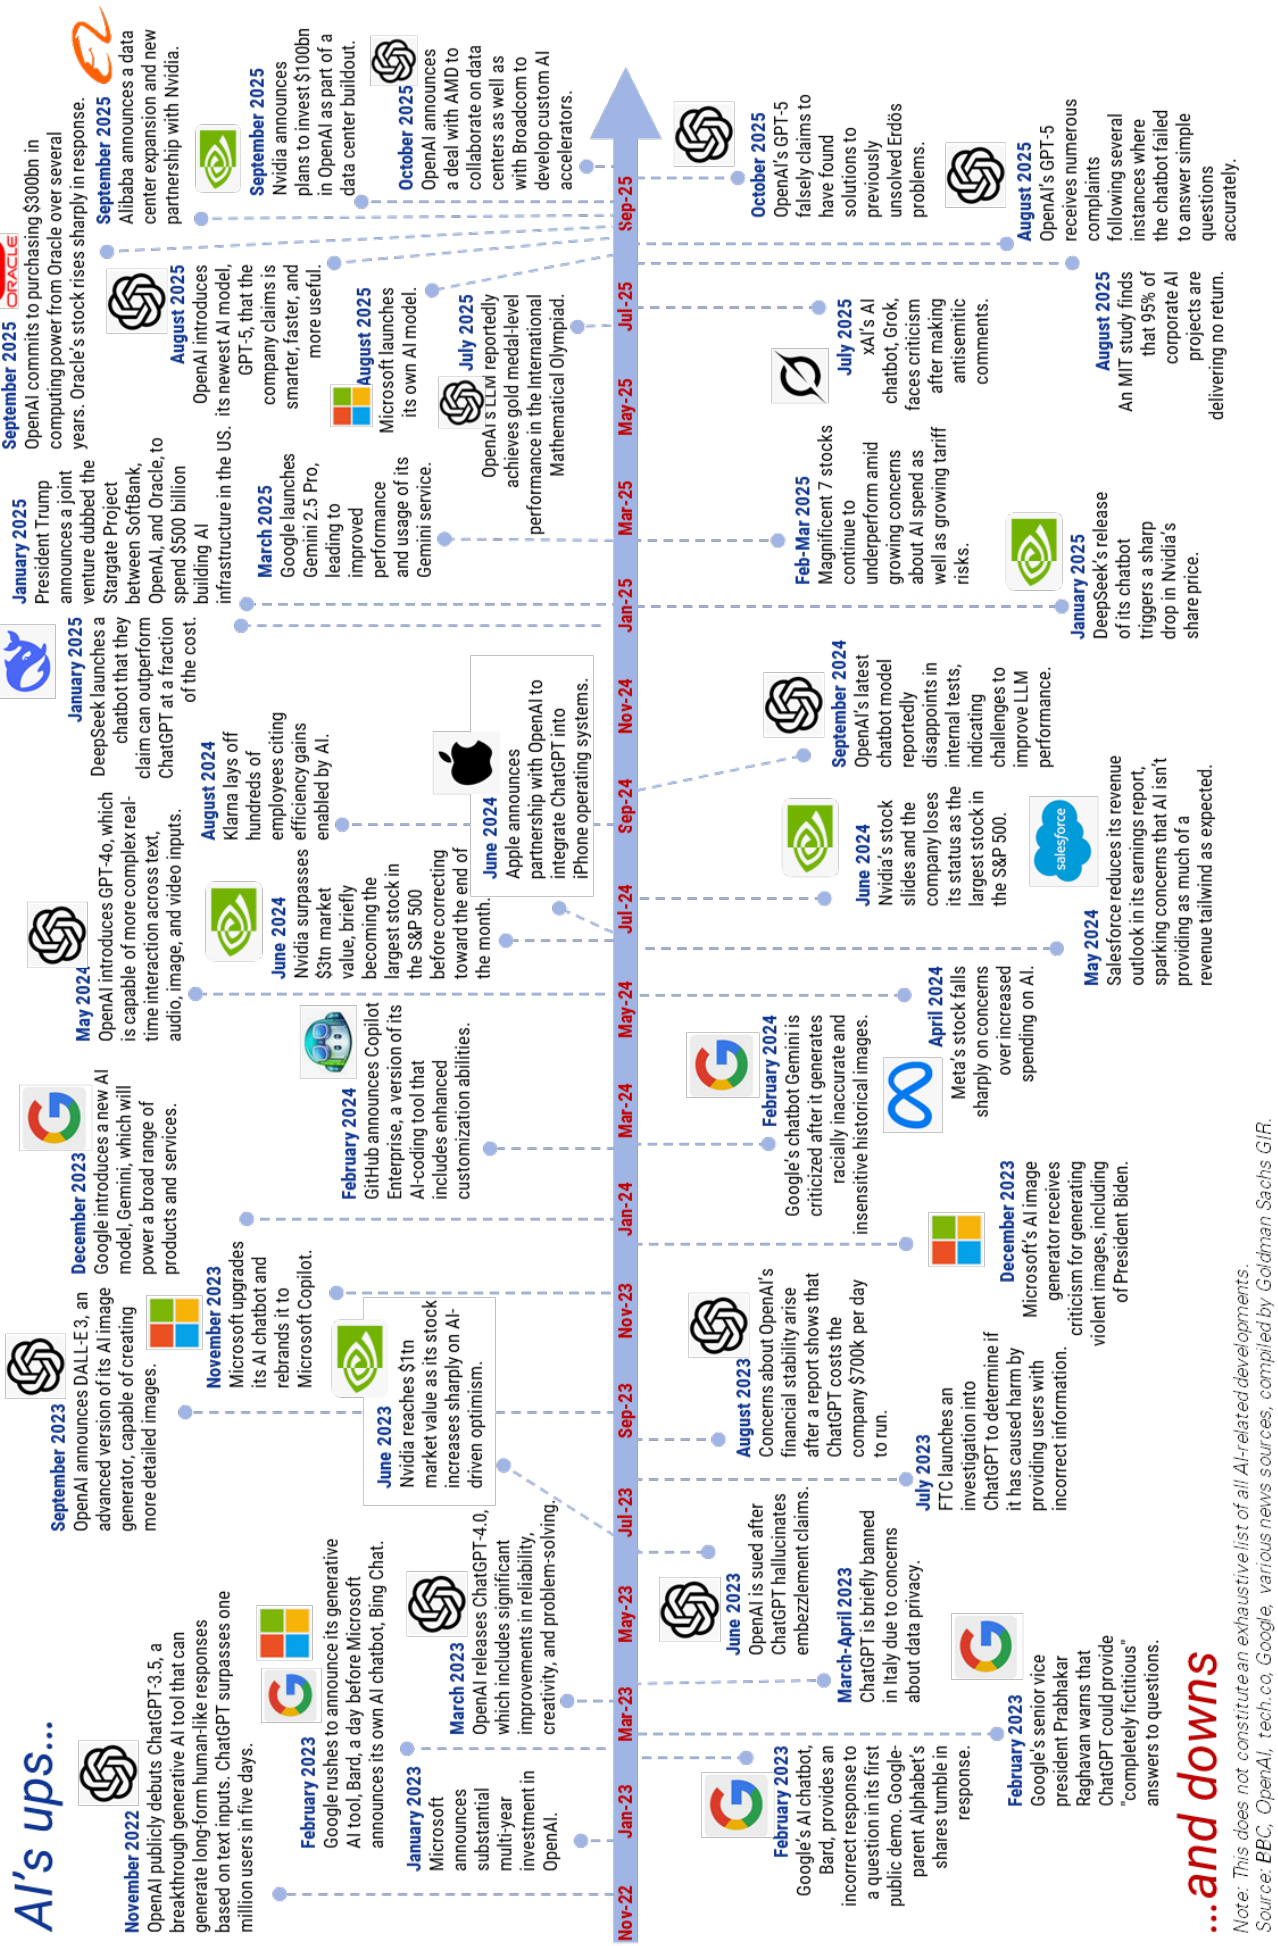

Note: This does not constitute an exhaustive list of all AI-related developments.
Source: BBC, OpenAI, tech.co, Google, various news sources, compiled by Goldman Sachs GIR.

US equities: no bubble but signs of froth

Ryan Hammond argues that, in aggregate, the US stock market is not in a bubble today

After several years of AI enthusiasm and a stock market that continues to reach new highs, some investors are now expressing concern about a US equity market bubble, drawing parallels to the capex boom and subsequent bust of the Dot-Com Bubble. We think such concerns are overblown. While we find that the level of US equity optimism is modestly above average, it is not at exuberant levels. Three key factors support this view: market-implied long-term earnings growth, valuations, and IPO activity all remain below levels observed during prior bubble periods. That said, emerging signs of froth in select areas of the market merit monitoring.

Factor 1: Near average long-term earnings growth expectations

While equity prices have continued to reach record highs, S&P 500 long-term growth expectations remain near historical averages. We estimate that market-implied long-term earnings growth stands at 10% today, only slightly higher than the long-term average of 9% and well below levels reached at the height of the Tech Bubble in 2000 (16%) and in 2021 (13%)[1].

Market-implied long-term growth expectations only modestly above average

Market-implied long-term S&P 500 growth expectations, %

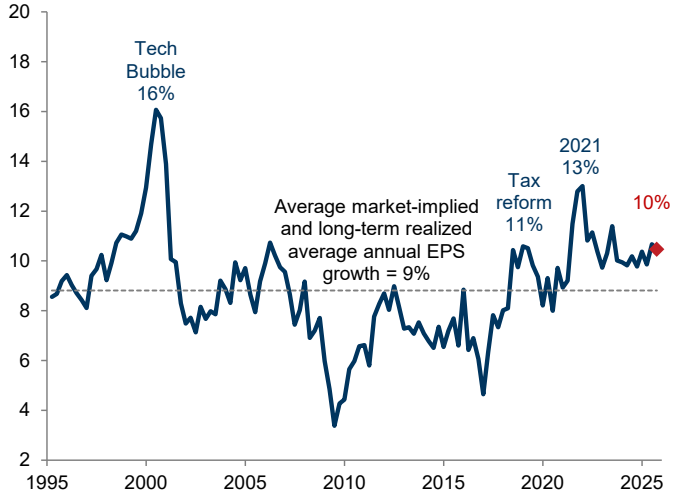

Source: Goldman Sachs GIR.

Factor 2: Largest stock valuations below historical peaks

Valuations for today's largest stocks also remain well below those of the largest stocks at the peaks 2000 and 2021. The five largest stocks in the S&P 500 (Nvidia, Apple, Microsoft, Alphabet, Amazon) currently trade at a P/E multiple of 29x, compared with 40x for the largest stocks at the peak in 2021 and 50x at the peak of the Tech Bubble. And with real yields over 200bp lower today, the gap between these stocks' earnings yield and the real 10-year Treasury yield is 4pp wider than at the height of the Tech Bubble. For most of the largest stocks, earnings growth rather than valuations has been the primary driver of returns. For example, Nvidia's stock price has increased by 13x since December 2022 alongside a 13x earnings increase, leaving its valuation effectively unchanged. That said, some large stocks, such as Tesla and Palantir, trade at much more elevated valuations (190x and 230x, respectively).

The largest stocks trade at a valuation of 29x, below Tech Bubble and 2021 levels

S&P 500 forward P/E multiple

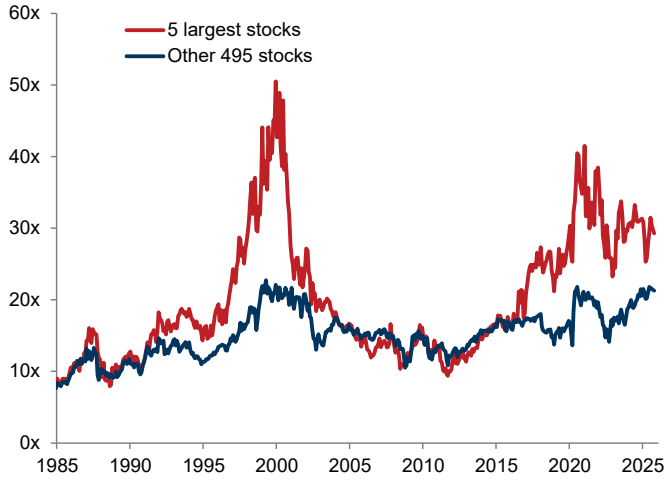

Source: Goldman Sachs GIR.

Factor 3: Still muted IPO activity

2025 will likely represent the busiest year for IPOs since 2021, but activity remains well below historical averages. So far this year, 51 IPOs over $25 million have occurred in the US. While this represents a 6% increase compared with the same point last year, it is substantially below the number of IPOs in 1999 (388) and 2021 (261). The average first-day return of IPOs this year (30%) is also below that of 1999 (65%) and 2020 (40%), though it is somewhat above the historical average (16%).

Select areas of froth

Despite few signs of a bubble in aggregate, some smaller pockets of the market show more evidence of froth. Our Speculative Trading Indicator shows a pickup in trading activity among expensive, unprofitable, and penny stocks, although this activity too remains below levels observed in the Tech Bubble and 2021. Other slices of the equity market have also risen sharply since the start of September, including a Quantum Computing basket (GSXUQNTM, +83%), Bitcoin Sensitive Equities basket (GSCBBTC1, +62%), and Retail Favorites basket (GSXURFAV, +17%), though these generally comprise relatively small stocks[2]. All told, while these pockets merit watching and we continue to see value in diversification partly given high levels of market concentration (see pgs. 16-17), we don't see a bubble in the US equity market in aggregate today.

Ryan Hammond, Senior US Equity Strategist

Email: ryan.hammond@gs.com *Goldman Sachs & Co. LLC*
Tel: 212-902-5625

[1] To estimate this implied growth rate, we use a cross-sectional regression of company return on equity and price to book. See the March 2024 US Equity Views for details.
[2] Baskets mentioned here were developed by the GS Global Banking & Markets division.

AI capex turns to credit

Shamshad Ali and Ben Shumway discuss the growing reliance on debt to fund AI ambitions

Since the rise of the generative AI theme in 2022, the five largest AI hyperscalers (Amazon, Alphabet, Microsoft, Meta, Oracle) have financed the majority of their AI-focused capex out of the large cash reserves on their balance sheets. But over the last few quarters, their cash balances have declined notably. The cash-to-total assets ratio among hyperscalers declined to 15% as of end-Q2 from 29% at end- 2021. And while the absolute level of balance sheet liquidity remains substantial relative to other industries, the downward trend marks a clear shift in how these firms are managing their liquidity positions.

Growing reliance on debt financing

The drawdown in cash balances has coincided with a growing reliance on debt financing—a reflection of both shifting capital management priorities and the scale of ongoing investment needs tied to AI and cloud infrastructure. This trend has been visible across many corners of the credit ecosystem, spanning traditional bank lending, the asset-backed securities (ABS) market, and public and private debt markets.

The cash-to-total assets ratio for the five largest AI hyperscalers has notably declined over recent years

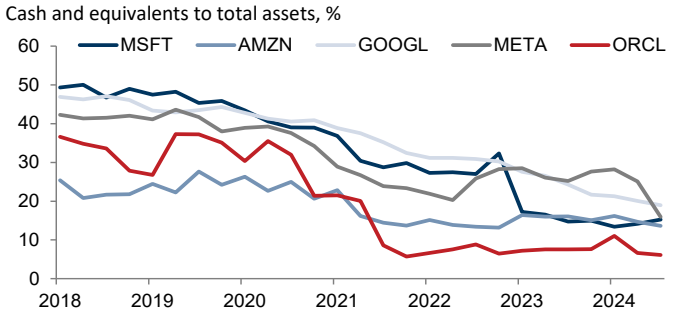

Source: Bloomberg, Goldman Sachs GIR.

AI-exposed companies have been more active in the primary corporate bond market this year

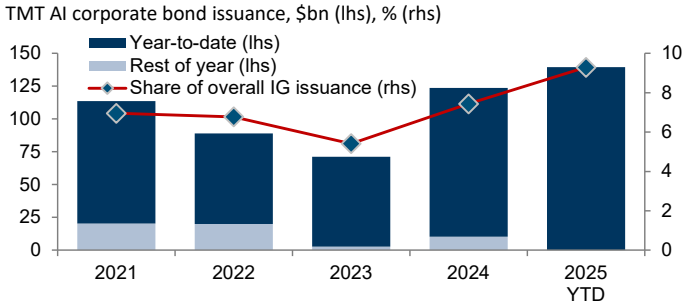

Source: Bloomberg, Goldman Sachs GIR.

Higher bond issuance…

In public markets, bond issuance by AI-exposed companies has increased sharply. Year-to-date, the constituents of the US TMT AI Basket (ticker: GSTMTAIP, developed by the GS Global Banking & Markets division), which consists of companies that are pursuing AI or can help enable new technologies, have collectively issued $139bn worth of corporate bonds, representing ~9.3% of total investment grade supply—a 23% increase relative to the same period last year, when issuance from these firms accounted for ~7.4% of overall investment grade issuance volumes. For context, as of end-Q2, these firms collectively spent $493bn on capex in the prior 12 months, including $313bn from hyperscalers. Disclosure around the use of proceeds in the investment grade market has typically limited granularity, but we suspect a meaningful share of this year's issuance has been directed toward capex. Just last week, Meta completed a financing deal for a Louisiana data center with a $27.2bn investment grade-rated private placement debt issue, the single largest bond issue recorded in the USD IG market.

…and rising private credit and structured finance channels

In parallel, private credit and structured finance channels have also become increasingly important sources of funding. Nowhere has this been more evident than in the acceleration of data center-related debt financing, which has expanded rapidly to meet surging demand for digital infrastructure. As of end-September, banks provided $73bn worth of financing for data center operators, up 31% versus full-year 2024 (according to data from Project Finance International). The same trend has prevailed in structured finance markets, where certain vehicles provide data center operators with the flexibility to raise debt incrementally, as long as they remain within pre-defined limits on leverage and debt servicing capacity. Commercial mortgage-backed securities (CMBS) and ABS linked to data center financing have grown their total balance outstanding from a meager $13bn at end-2021 to $49bn today.

Where to from here?

Consensus forecasts call for a 20% increase in capex among the five largest hyperscalers next year, marking a sharp slowdown from this year's 74% pace. However, our portfolio strategy team sees significant upside risk to this estimate, implying that the recent shift toward greater reliance on debt financing will likely persist and potentially intensify ahead. While not yet a cause for alarm given both the high cash flow generation and low leverage among large tech companies, the shifting funding mix of capex beyond cash suggests balance sheet leverage will likely grow, which is negative for credit quality in the corporate bond market, albeit only marginally.

Another way to invest in AI

For structured credit investors, the growth of data center ABS and CMBS markets offers a viable way to express a positive view on AI. However, it also comes with risk management challenges. The first is the risk of a supply/demand imbalance given the rapid expansion of the past few years. Such an outcome could raise the risk of tenants either breaking or not renewing leases, which may negatively impact investors' cash flows. The second is the risk of obsolescence, as hardware in the ever-advancing technology industry will naturally need replacing. All told, the AI opportunity for credit investors should rise alongside the industry's growing reliance on debt financing, though it will come with risks.

Shamshad Ali, Credit Strategist

Email: shamshad.ali@gs.com *Goldman Sachs & Co. LLC*
Tel: 212-902-6712

Ben Shumway, Credit Strategist

Email: ben.shumway@gs.com *Goldman Sachs & Co. LLC*
Tel: 801-578-2553

AI talk, in pics

The outsized performance of the Magnificent 7 stocks on the back of AI enthusiasm has raised concerns about an AI bubble

Magnificent 7 and S&P 500, 12m trailing EPS, Jan 2005=100

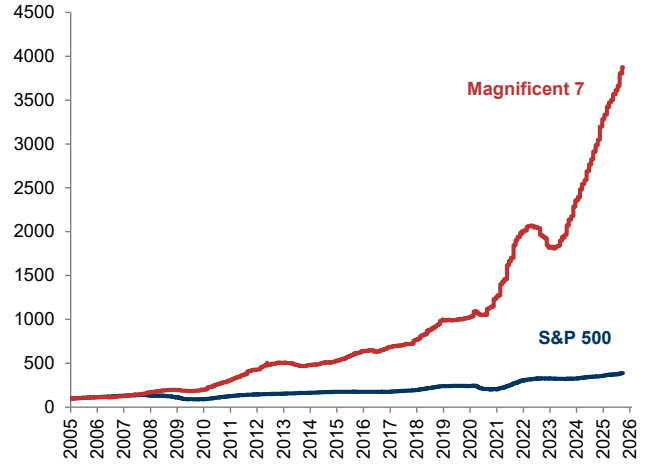

Source: FactSet, Goldman Sachs GIR.

Google searches for "AI bubble" have declined from their August peak, but have ticked up in recent weeks

Google searches of "AI bubble"

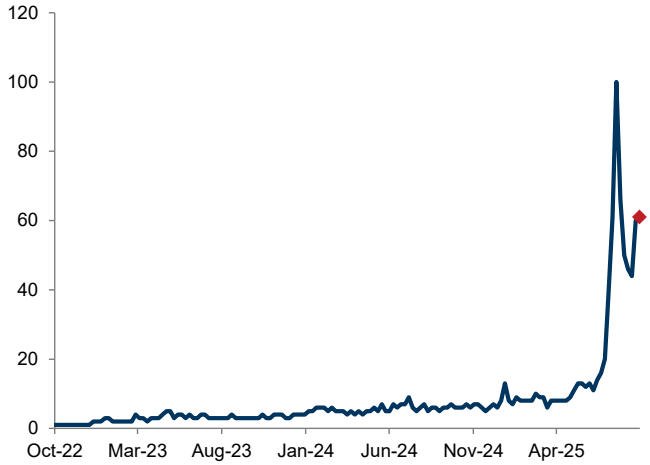

Source: Google, Goldman Sachs GIR.

News stories about AI bubbles show waves of concern, with the peak coming after the DeepSeek announcement in January

"AI" and "AI bubble" mentions in news on Bloomberg terminals, count

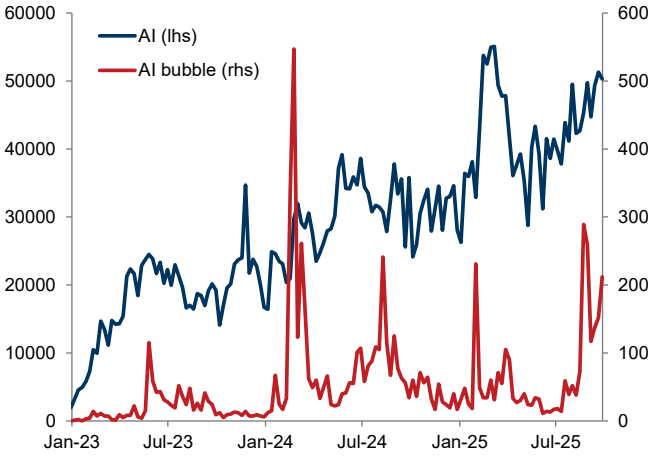

Source: Bloomberg, Goldman Sachs GIR.

More broadly, AI mentions in earnings calls continue to grow, with over 50% of S&P 500 companies mentioning AI in Q2

Share of S&P 500 companies mentioning AI in earnings calls, %

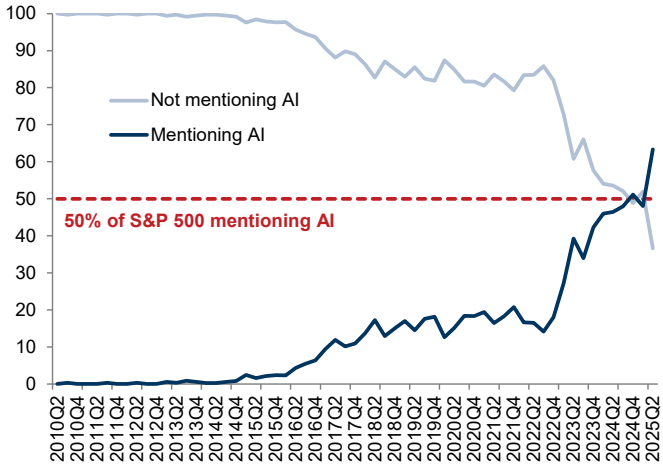

Source: LSEG, Goldman Sachs GIR.

AI focus remains broad, with companies continuing to mention AI investment and opportunity with a focus on their customers

AI-related mentions in 2Q2025 S&P 500 earnings calls, scaled by number of mentions

Source: LSEG, Goldman Sachs GIR.

Tech companies have increasingly shifted their AI focus from infrastructure to the application phase, with rising emphasis on AI agents

Mentions in S&P 500 TMT/Telco earnings calls*

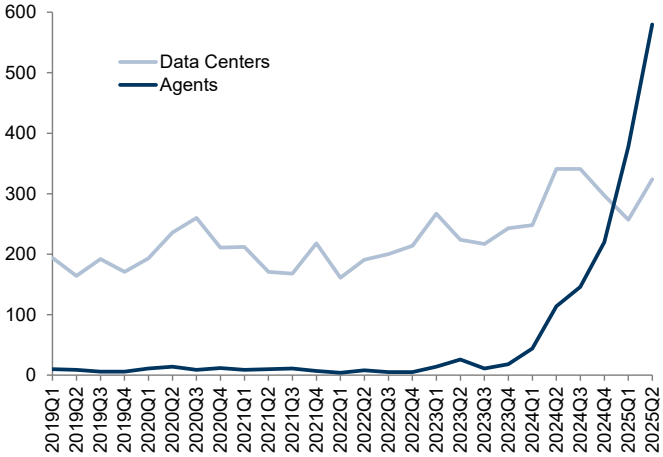

**Keyword counts from S&P 500 TMT and Telecommunications companies.*

Source: LSEG, Goldman Sachs GIR.

Special thanks to Dan Duggan and GS Data Works for its AI-related analysis.

Summary of our key forecasts

GS GIR: Macro at a glance

Watching

- **Globally,** we expect real GDP growth to slow to 2.7% yoy in 2025, reflecting headwinds from higher US tariffs. We expect global core inflation to remain relatively steady this year and end the year at around 2.8% as the tariff-driven boost to inflation in the US is largely offset by disinflationary impulses from declines in shelter inflation and wage inflation.
- **In the US,** we expect GDP growth to slow to 1.3% in 4Q25 (QoQ, ann.), reflecting drags from higher tariffs, slower job gains, and the federal government shutdown for full-year 2025 GDP growth of 1.9% on a Q4/Q4 basis. We expect the economy to modestly reaccelerate toward potential next year as the tariff drag fades, fiscal policy turns more expansionary, and the Fed eases. We expect core PCE inflation to rise slightly to 3.0% yoy by end-2025, reflecting a boost from higher tariffs. We expect the unemployment rate to rise to 4.5% by end-2025.
- **We expect the Fed** to deliver two more 25bp rate cuts this year in October and December followed by two 25bp cuts in 2026 for a terminal rate range of 3-3.25%.
- **In the Euro area,** we expect real GDP growth of 1.3% yoy in 2025 amid higher US tariffs, though the trade-related growth headwind should recede next year. We expect core inflation to fall to 2.2% by end-2025 and remain around 2% in coming years as a resilient labor market will likely result in wage growth running in line with target and inflation expectations remain well anchored at 2%.
- **We believe the ECB's** cutting cycle is now finished, although the risks around our baseline remain slightly skewed to the downside.
- **In China,** we expect growth to slow further over the remainder of the year for full-year real GDP growth of 4.9% yoy in 2025 amid higher US tariffs. On the inflation front, we believe China's overcapacity problem will take time and effort to solve and expect CPI/PPI inflation of 0%/-2.6% this year.
- **WATCH US POLICY, GLOBAL FISCAL DYNAMICS, AND GEOPOLITICAL DEVELOPMENTS.** Uncertainty about US policy remains elevated, presenting risk to the US and global economies. Fiscal pressures—including in the US, UK, France, and Japan—could have important macro and market implications. Geopolitical developments also remain important to watch as US-China relations remain fraught and the outlook for both the Russia-Ukraine war and the situation in the Middle East remains uncertain.

Goldman Sachs Global Investment Research.

Growth

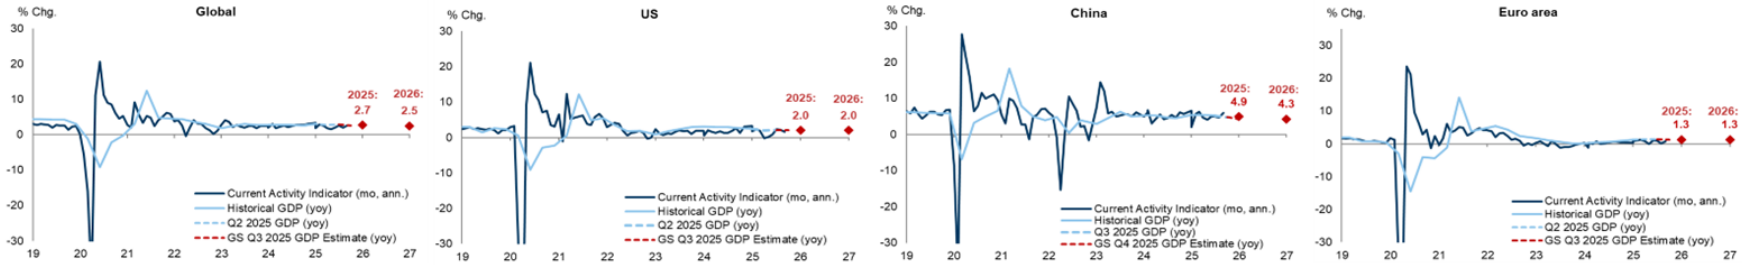

Source: Haver Analytics, Goldman Sachs Global Investment Research.
Note: GS CAI is a measure of current growth. For more information on the methodology of the CAI please see "Technical Updates to Our Global CAIs," Global Economics Comment, Sep. 01, 2025.

Forecasts

Economics							Markets								Equities								
GDP growth (%)	2025				2026		Interest rates 10Yr (%)	Last	E2025	E2026	FX		Last	3m	12m	S&P 500	E2025		E2026		Returns (%)	12m	YTD	E2025 P/E
	GS (Q4/Q4)	Cons. (Q4/Q4)	GS (CY)	Cons. (CY)	GS (CY)	Cons. (CY)											GS	Cons.	GS	Cons.				
Global	2.4	--	2.7	2.5	2.5	2.4	US	3.98	4.20	4.20	EUR/$		1.16	1.20	1.25	Price	6,800				S&P 500	6.9	14.5	25.7x
US	1.9	1.4	2.0	1.8	2.0	1.8	Germany	2.55	3.00	3.25	GBP/$		1.34	1.38	1.39	EPS	$262	$267	$280	$304	MXAPJ	5.2	26.4	17.8x
China	4.2	4.1	4.9	4.8	4.3	4.2	Japan	1.66	1.70	1.90	$/JPY		152	150	145	Growth	7%	9%	7%	14%	Topix	4.7	16.7	17.2x
Euro area	1.0	1.0	1.3	1.3	1.3	1.1	UK	4.42	4.40	4.25	$/CNY		7.10	7.00	6.90						STOXX 600	1.4	12.9	16.3x
Policy rates (%)	2025				2026		Commodities	Last	3m	12m	Credit (bp)		Last	2Q25	4Q25	Consumer	2025		2026		Wage Tracker 2025 (%)			
	GS	Mkt.			GS	Mkt.											CPI (%, yoy)	Unemp. Rate	CPI (%, yoy)	Unemp. Rate	Q1	Q2	Q3	Q4
US	3.63	3.62			3.13	2.87	Nat Gas, NYMEX ($/mmBtu)	3.47	4.00	4.50	USD	IG	78	88	85	US	2.7	4.5	2.7	4.4	3.9	3.9	3.9	--
Euro area	2.00	1.89			2.00	1.75	Nat Gas, TTF (EUR/MWh)	32.14	39	29		HY	288	303	300	Euro area	2.1	6.4	1.8	6.5	--	--	--	--
China	1.30	1.46			1.10	--	Copper ($/mt)	10,593	10,390	10,580	EUR	IG	91	102	100	China	0.0	--	0.8	--	--	--	--	--
Japan	0.50	0.63			1.00	1.03	Gold ($/troy oz)	4,170	4,160	4,745		HY	299	310	307									

Source: Bloomberg, Goldman Sachs Global Investment Research. For important disclosures, see the Disclosure Appendix or go to www.gs.com/research/hedge.html.

Market pricing as of October 21, 2025

Glossary of GS proprietary indices

Current Activity Indicator (CAI)

GS CAIs measure the growth signal in a broad range of weekly and monthly indicators, offering an alternative to Gross Domestic Product (GDP). GDP is an imperfect guide to current activity: In most countries, it is only available quarterly and is released with a substantial delay, and its initial estimates are often heavily revised. GDP also ignores important measures of real activity, such as employment and the purchasing managers' indexes (PMIs). All of these problems reduce the effectiveness of GDP for investment and policy decisions. Our CAIs aim to address GDP's shortcomings and provide a timelier read on the pace of growth.

For more, see our CAI page and Global Economics Comment: Technical Updates to Our Global CAIs.

Dynamic Equilibrium Exchange Rates (DEER)

The GSDEER framework establishes an equilibrium (or "fair") value of the real exchange rate based on relative productivity and terms-of-trade differentials.

For more, see our GSDEER page, Global Economics Paper No. 227: Finding Fair Value in EM FX, 26 January 2016, and Global Markets Analyst: A Look at Valuation Across G10 FX, 29 June 2017.

Financial Conditions Index (FCI)

GS FCIs gauge the "looseness" or "tightness" of financial conditions across the world's major economies, incorporating variables that directly affect spending on domestically produced goods and services. FCIs can provide valuable information about the economic growth outlook and the direct and indirect effects of monetary policy on real economic activity.

FCIs for the G10 economies are calculated as a weighted average of a policy rate, a long-term risk-free bond yield, a corporate credit spread, an equity price variable, and a trade-weighted exchange rate; the Euro area FCI also includes a sovereign credit spread. The weights mirror the effects of the financial variables on real GDP growth in our models over a one-year horizon. FCIs for emerging markets are calculated as a weighted average of a short-term interest rate, a long-term swap rate, a CDS spread, an equity price variable, a trade-weighted exchange rate, and—in economies with large foreign-currency-denominated debt stocks—a debt-weighted exchange rate index.

For more, see our FCI page, Global Economics Analyst: Our New G10 Financial Conditions Indices, 20 April 2017, and Global Economics Analyst: Tracking EM Financial Conditions – Our New FCIs, 6 October 2017.

Goldman Sachs Analyst Index (GSAI)

The US GSAI is based on a monthly survey of GS equity analysts to obtain their assessments of business conditions in the industries they follow. The results provide timely "bottom-up" information about US economic activity to supplement and cross-check our analysis of "top-down" data. Based on analysts' responses, we create a diffusion index for economic activity comparable to the ISM's indexes for activity in the manufacturing and nonmanufacturing sectors.

Macro-Data Assessment Platform (MAP)

GS MAP scores facilitate rapid interpretation of new data releases for economic indicators worldwide. MAP summarizes the importance of a specific data release (i.e., its historical correlation with GDP) and the degree of surprise relative to the consensus forecast. The sign on the degree of surprise characterizes underperformance with a negative number and outperformance with a positive number. Each of these two components is ranked on a scale from 0 to 5, with the MAP score being the product of the two, i.e., from -25 to +25. For example, a MAP score of +20 (5;+4) would indicate that the data has a very high correlation to GDP (5) and that it came out well above consensus expectations (+4), for a total MAP value of +20.

Top of Mind archive

Issue 142
Data (un)reliability
September 11, 2025

Issue 141
Stablecoin summer
August 19, 2025

Issue 140
US fiscal worries: Is this time different?
June 12, 2025

Issue 139
Fed independence: how concerning?
May 15, 2025

Issue 138
Tariff-induced recession risk
April 17, 2025

Issue 137
Europe's shifting security landscape
March 13, 2025

Issue 136
Trump tariffs: mostly talk, or big action?
February 26, 2025

Issue 135
High bond yields: here to stay?
February 5, 2025

Special Issue
2024: 4 themes in charts
December 18, 2024

Issue 134
Will China's policy stimulus be enough?
December 11, 2024

Issue 133
Market concentration: how big a worry?
November 25, 2024

Issue 132
Post-election economic policies
October 21, 2024

Issue 131
Is the Fed behind the curve?
September 3, 2024

Issue 130
How investable is Europe?
August 1, 2024

Issue 129
Gen AI: too much spend, too little benefit?
June 25, 2024

Issue 128
Central bank divergence room to run?
May 21, 2024

Issue 127
Weighing the GLP-1 market
April 12, 2024

Issue 126
Global transit & trade: in rough waters
March 12, 2024

Issue 125
2024: The year of elections
February 1, 2024

Issue 124
Middle East risks
December 5, 2023

Issue 123
US outperformance: at a turning point?
October 30, 2023

Issue 122
Commercial real estate risks
October 9, 2023

Issue 121
Corporate credit concerns
August 10, 2023

Issue 120
Generative AI: hype, or truly transformative?
July 5, 2023

Issue 119
Daunting debt limit dynamics
May 22, 2023

Issue 118
US-China: more decoupling ahead?
May 1, 2023

Issue 117
All about bank(panic)s
April 3, 2023

Issue 116
(Japanese) Bonds, Bonds, Bonds
February 23, 2023

Issue 115
The Bigger Worry: Growth or Inflation?
January 27, 2023

Special Issue
2022: 3 themes in charts
December 15, 2022

Issue 114
The Winter of Crypto's Discontents
December 9, 2022

Issue 113
Central Bank Tightening: what could break?
November 11, 2022

Issue 112
China's Congress: an inflection point?
October 11, 2022

Issue 111
Will slaying inflation require recession?
September 13, 2022

Source of photos: www.gettyimages.com.

Disclosure Appendix

Reg AC

We, Allison Nathan, Jenny Grimberg, Ashley Rhodes, Peter Oppenheimer, Joseph Briggs, Ryan Hammond, Ben Shumway and Shamshad Ali, hereby certify that all of the views expressed in this report accurately reflect our personal views, which have not been influenced by considerations of the firm's business or client relationships.

We, Eric Sheridan and Kash Rangan, hereby certify that all of the views expressed in this report accurately reflect our personal views about the subject company or companies and its or their securities. We also certify that no part of our compensation was, is or will be, directly or indirectly, related to the specific recommendations or views expressed in this report.

Unless otherwise stated, the individuals listed on the cover page of this report are analysts in Goldman Sachs' Global Investment Research division.

Logo disclosure

Please note: Third party brands used in this report are the property of their respective owners, and are used here for informational purposes only. The use of such brands should not be viewed as an endorsement, affiliation or sponsorship by or for Goldman Sachs or any of its products/services.

Other disclosures

Kash Rangan has announced his intention to retire from Goldman Sachs as of January 31, 2026. His current coverage will be transferred or suspended by that date.

DataWorks

GS DataWorks leverages alternative data sources and advanced analysis techniques to create unique data-driven insights across Global Investment Research.

GS DataWorks analysis provided by Dan Duggan, Ph.D.

Regulatory disclosures

Disclosures required by United States laws and regulations

See company-specific regulatory disclosures above for any of the following disclosures required as to companies referred to in this report: manager or co-manager in a pending transaction; 1% or other ownership; compensation for certain services; types of client relationships; managed/co-managed public offerings in prior periods; directorships; for equity securities, market making and/or specialist role. Goldman Sachs trades or may trade as a principal in debt securities (or in related derivatives) of issuers discussed in this report.

The following are additional required disclosures: **Ownership and material conflicts of interest:** Goldman Sachs policy prohibits its analysts, professionals reporting to analysts and members of their households from owning securities of any company in the analyst's area of coverage. **Analyst compensation:** Analysts are paid in part based on the profitability of Goldman Sachs, which includes investment banking revenues. **Analyst as officer or director:** Goldman Sachs policy generally prohibits its analysts, persons reporting to analysts or members of their households from serving as an officer, director or advisor of any company in the analyst's area of coverage. **Non-U.S. Analysts:** Non-U.S. analysts may not be associated persons of Goldman Sachs & Co. LLC and therefore may not be subject to FINRA Rule 2241 or FINRA Rule 2242 restrictions on communications with a subject company, public appearances and trading in securities covered by the analysts.

Additional disclosures required under the laws and regulations of jurisdictions other than the United States

The following disclosures are those required by the jurisdiction indicated, except to the extent already made above pursuant to United States laws and regulations. **Australia:** Goldman Sachs Australia Pty Ltd and its affiliates are not authorised deposit-taking institutions (as that term is defined in the Banking Act 1959 (Cth)) in Australia and do not provide banking services, nor carry on a banking business, in Australia. This research, and any access to it, is intended only for "wholesale clients" within the meaning of the Australian Corporations Act, unless otherwise agreed by Goldman Sachs. In producing research reports, members of Global Investment Research of Goldman Sachs Australia may attend site visits and other meetings hosted by the companies and other entities which are the subject of its research reports. In some instances the costs of such site visits or meetings may be met in part or in whole by the issuers concerned if Goldman Sachs Australia considers it is appropriate and reasonable in the specific circumstances relating to the site visit or meeting. To the extent that the contents of this document contains any financial product advice, it is general advice only and has been prepared by Goldman Sachs without taking into account a client's objectives, financial situation or needs. A client should, before acting on any such advice, consider the appropriateness of the advice having regard to the client's own objectives, financial situation and needs.

A copy of certain Goldman Sachs Australia and New Zealand disclosure of interests and a copy of Goldman Sachs' Australian Sell-Side Research Independence Policy Statement are available at: https://www.goldmansachs.com/disclosures/australia-new-zealand/index.html. **Brazil:** Disclosure information in relation to CVM Resolution n. 20 is available at https://www.gs.com/worldwide/brazil/area/gir/index.html. Where applicable, the Brazil-registered analyst primarily responsible for the content of this research report, as defined in Article 20 of CVM Resolution n. 20, is the first author named at the beginning of this report, unless indicated otherwise at the end of the text. **Canada:** This information is being provided to you for information purposes only and is not, and under no circumstances should be construed as, an advertisement, offering or solicitation by Goldman Sachs & Co. LLC for purchasers of securities in Canada to trade in any Canadian security. Goldman Sachs & Co. LLC is not registered as a dealer in any jurisdiction in Canada under applicable Canadian securities laws and generally is not permitted to trade in Canadian securities and may be prohibited from selling certain securities and products in certain jurisdictions in Canada. If you wish to trade in any Canadian securities or other products in Canada please contact Goldman Sachs Canada Inc., an affiliate of The Goldman Sachs Group Inc., or another registered Canadian dealer. **Hong Kong:** Further information on the securities of covered companies referred to in this research may be obtained on request from Goldman Sachs (Asia) L.L.C. **India:** Further information on the subject company or companies referred to in this research may be obtained from Goldman Sachs (India) Securities Private Limited, Research Analyst - SEBI Registration Number INH000001493, 951-A, Rational House, Appasaheb Marathe Marg, Prabhadevi, Mumbai 400 025, India, Corporate Identity Number U74140MH2006FTC160634, Phone +91 22 6616 9000, Fax +91 22 6616 9001. Goldman Sachs may beneficially own 1% or more of the securities (as such term is defined in clause 2 (h) the Indian Securities Contracts (Regulation) Act, 1956) of the subject company or companies referred to in this research report. Investment in securities market are subject to market risks. Read all the related documents carefully before investing. Registration granted by SEBI and certification from NISM in no way guarantee performance of the intermediary or provide any assurance of returns to investors. Goldman Sachs (India) Securities Private Limited compliance officer and investor grievance contact details can be found at this link: https://publishing.gs.com/disclosures/hedge.html - /general/equity. **Japan:** See below. **Korea:** This research, and any access to it, is intended only for "professional investors" within the meaning of the Financial Services and Capital Markets Act, unless otherwise agreed by Goldman Sachs. Further information on the subject company or companies referred to in this research may be obtained from Goldman Sachs (Asia) L.L.C., Seoul Branch. **New Zealand:** Goldman Sachs New Zealand Limited and its affiliates are neither "registered banks" nor "deposit takers" (as defined in the Reserve Bank of New Zealand Act 1989) in New Zealand. This research, and any access to it, is intended for "wholesale clients" (as defined in the Financial Advisers Act 2008) unless otherwise agreed by Goldman Sachs. A copy of certain Goldman Sachs Australia and New Zealand disclosure of interests is available at: https://www.goldmansachs.com/disclosures/australia-new-zealand/index.html. **Russia:** Research reports distributed in the Russian Federation are not advertising as defined in the Russian legislation, but are information and analysis not having product promotion as their main purpose and do not provide appraisal within the meaning of the Russian legislation on appraisal activity. Research reports do not constitute a personalized investment recommendation as defined in Russian laws and regulations, are not addressed to a specific client, and are prepared without analyzing the financial circumstances, investment profiles or risk profiles of clients. Goldman Sachs assumes no responsibility for any investment decisions that may be taken by a client or any other person based on this research report. **Singapore:** Goldman Sachs (Singapore) Pte. (Company Number: 198602165W), which is regulated by the Monetary Authority of Singapore, accepts legal responsibility for this research, and should be contacted with respect to any matters arising from, or in connection with, this research. **Taiwan:** This material is for reference only and must not be reprinted without permission. Investors should carefully consider their own investment risk. Investment results are the responsibility of the individual investor. **United Kingdom:** Persons who would be categorized as retail clients in the United Kingdom, as such term is defined in the rules of the Financial Conduct Authority, should read this research in conjunction with prior Goldman Sachs research on the covered companies referred to herein and should refer to the risk warnings that have been sent to them by Goldman Sachs International. A copy of these risks warnings, and a glossary of certain financial terms used in this report, are available from Goldman Sachs International on request.

European Union and United Kingdom: Disclosure information in relation to Article 6 (2) of the European Commission Delegated Regulation (EU) (2016/958) supplementing Regulation (EU) No 596/2014 of the European Parliament and of the Council (including as that Delegated Regulation is implemented into United Kingdom domestic law and regulation following the United Kingdom's departure from the European Union and the European Economic Area) with regard to regulatory technical standards for the technical arrangements for objective presentation of investment recommendations or other information recommending or suggesting an investment strategy and for disclosure of particular interests or indications of conflicts of interest is available at https://www.gs.com/disclosures/europeanpolicy.html which states the European Policy for Managing Conflicts of Interest in Connection with Investment Research.

Japan: Goldman Sachs Japan Co., Ltd. is a Financial Instrument Dealer registered with the Kanto Financial Bureau under registration number Kinsho 69, and a member of Japan Securities Dealers Association, Financial Futures Association of Japan Type II Financial Instruments Firms Association, The Investment Trusts Association, Japan, and Japan Investment Advisers Association. Sales and purchase of equities are subject to commission pre-determined with clients plus consumption tax. See company-specific disclosures as to any applicable disclosures required by Japanese stock exchanges, the Japanese Securities Dealers Association or the Japanese Securities Finance Company.

Global product; distributing entities

Goldman Sachs Global Investment Research produces and distributes research products for clients of Goldman Sachs on a global basis. Analysts based in Goldman Sachs offices around the world produce research on industries and companies, and research on macroeconomics, currencies, commodities and portfolio strategy. This research is disseminated in Australia by Goldman Sachs Australia Pty Ltd (ABN 21 006 797 897); in Brazil by Goldman Sachs do Brasil Corretora de Títulos e Valores Mobiliários S.A.; Public Communication Channel Goldman Sachs Brazil: 0800 727 5764 and / or contatogoldmanbrasil@gs.com. Available Weekdays (except holidays), from 9am to 6pm. Canal de Comunicação com o Público Goldman Sachs Brasil: 0800 727 5764 e/ou contatogoldmanbrasil@gs.com. Horário de funcionamento: segunda-feira à sexta-feira (exceto feriados), das 9h às 18h; in Canada by Goldman Sachs & Co. LLC; in Hong Kong by Goldman Sachs (Asia) L.L.C.; in India by Goldman Sachs (India) Securities Private Ltd.; in Japan by Goldman Sachs Japan Co., Ltd.; in the Republic of Korea by Goldman Sachs (Asia) L.L.C., Seoul Branch; in New Zealand by Goldman Sachs New Zealand Limited; in Russia by OOO Goldman Sachs; in Singapore by Goldman Sachs (Singapore) Pte. (Company Number: 198602165W); and in the United States of America by Goldman Sachs & Co. LLC. Goldman Sachs International has approved this research in connection with its distribution in the United Kingdom.

Goldman Sachs International ("GSI"), authorised by the Prudential Regulation Authority ("PRA") and regulated by the Financial Conduct Authority ("FCA") and the PRA, has approved this research in connection with its distribution in the United Kingdom.

European Economic Area: Goldman Sachs Bank Europe SE ("GSBE") is a credit institution incorporated in Germany and, within the Single Supervisory Mechanism, subject to direct prudential supervision by the European Central Bank and in other respects supervised by German Federal Financial Supervisory Authority (Bundesanstalt für Finanzdienstleistungsaufsicht, BaFin) and Deutsche Bundesbank and disseminates research within the European Economic Area.

General disclosures

This research is for our clients only. Other than disclosures relating to Goldman Sachs, this research is based on current public information that we consider reliable, but we do not represent it is accurate or complete, and it should not be relied on as such. The information, opinions, estimates and forecasts contained herein are as of the date hereof and are subject to change without prior notification. We seek to update our research as appropriate, but various regulations may prevent us from doing so. Other than certain industry reports published on a periodic basis, the large majority of reports are published at irregular intervals as appropriate in the analyst's judgment.

Goldman Sachs conducts a global full-service, integrated investment banking, investment management, and brokerage business. We have investment banking and other business relationships with a substantial percentage of the companies covered by Global Investment Research. Goldman Sachs & Co. LLC, the United States broker dealer, is a member of SIPC (https://www.sipc.org).

Our salespeople, traders, and other professionals may provide oral or written market commentary or trading strategies to our clients and principal trading desks that reflect opinions that are contrary to the opinions expressed in this research. Our asset management area, principal trading desks and investing businesses may make investment decisions that are inconsistent with the recommendations or views expressed in this research.

The analysts named in this report may have from time to time discussed with our clients, including Goldman Sachs salespersons and traders, or may discuss in this report, trading strategies that reference catalysts or events that may have a near-term impact on the market price of the equity securities discussed in this report, which impact may be directionally counter to the analyst's published price target expectations for such stocks. Any such trading strategies are distinct from and do not affect the analyst's fundamental equity rating for such stocks, which rating reflects a stock's return potential relative to its coverage universe as described herein.

We and our affiliates, officers, directors, and employees will from time to time have long or short positions in, act as principal in, and buy or sell, the securities or derivatives, if any, referred to in this research, unless otherwise prohibited by regulation or Goldman Sachs policy.

The views attributed to third party presenters at Goldman Sachs arranged conferences, including individuals from other parts of Goldman Sachs, do not necessarily reflect those of Global Investment Research and are not an official view of Goldman Sachs.

Any third party referenced herein, including any salespeople, traders and other professionals or members of their household, may have positions in the products mentioned that are inconsistent with the views expressed by analysts named in this report.

This research is not an offer to sell or the solicitation of an offer to buy any security in any jurisdiction where such an offer or solicitation would be illegal. It does not constitute a personal recommendation or take into account the particular investment objectives, financial situations, or needs of individual clients. Clients should consider whether any advice or recommendation in this research is suitable for their particular circumstances and, if appropriate, seek professional advice, including tax advice. The price and value of investments referred to in this research and the income from them may fluctuate. Past performance is not a guide to future performance, future returns are not guaranteed, and a loss of original capital may occur. Fluctuations in exchange rates could have adverse effects on the value or price of, or income derived from, certain investments.

Certain transactions, including those involving futures, options, and other derivatives, give rise to substantial risk and are not suitable for all investors. Investors should review current options and futures disclosure documents which are available from Goldman Sachs sales representatives or at https://www.theocc.com/about/publications/character-

risks.jsp and https://www.goldmansachs.com/disclosures/cftc_fcm_disclosures. Transaction costs may be significant in option strategies calling for multiple purchase and sales of options such as spreads. Supporting documentation will be supplied upon request.

Differing Levels of Service provided by Global Investment Research: The level and types of services provided to you by Goldman Sachs Global Investment Research may vary as compared to that provided to internal and other external clients of GS, depending on various factors including your individual preferences as to the frequency and manner of receiving communication, your risk profile and investment focus and perspective (e.g., marketwide, sector specific, long term, short term), the size and scope of your overall client relationship with GS, and legal and regulatory constraints. As an example, certain clients may request to receive notifications when research on specific securities is published, and certain clients may request that specific data underlying analysts' fundamental analysis available on our internal client websites be delivered to them electronically through data feeds or otherwise. No change to an analyst's fundamental research views (e.g., ratings, price targets, or material changes to earnings estimates for equity securities), will be communicated to any client prior to inclusion of such information in a research report broadly disseminated through electronic publication to our internal client websites or through other means, as necessary, to all clients who are entitled to receive such reports.

All research reports are disseminated and available to all clients simultaneously through electronic publication to our internal client websites. Not all research content is redistributed to our clients or available to third-party aggregators, nor is Goldman Sachs responsible for the redistribution of our research by third party aggregators. For research, models or other data related to one or more securities, markets or asset classes (including related services) that may be available to you, please contact your GS representative or go to https://research.gs.com.

Disclosure information is also available at https://www.gs.com/research/hedge.html or from Research Compliance, 200 West Street, New York, NY 10282.