Economic & Market Outlook

Executive Summary

~ The economy and market appear split, with AI-related capital spending supporting growth and profits while the "normal" economy slows. Since the ChatGPT moment, AI-linked names have contributed roughly 79% of S&P 500 earnings growth and have driven most of the equity returns.

~ Scale is the swing variable: multi-trillion dollar data-center investment acts like stimulus, and if the pace slows, the macro impulse will fade. Data-center capex alone is projected to reach about 1.7% of U.S. GDP in 2026.

~ The profits required to justify the build-out are significant and still rising, and whether they materialize remains uncertain.

~ AI datacenter funding structures continue to grow more complex, and circular commitments among model developers, chipmakers, and data-center operators increase market fragility.

~ We maintain exposure to the AI theme, but far less than US indices (like the S&P 500). Instead, we are leaning into global diversification and see opportunity in stocks that haven't benefited from the AI trade.

Introduction

Investing never seems to get easier. One issue is markets appear to grow more efficient with time, but that isn't the struggle we're alluding to. Rather the power of ignorance. Early in your investing career, everything seems obvious, making it easier to act, buy or sell with conviction. With experience, however, you learn the countless ways money can be lost. You learn that nuance matters, and once obvious investments become riddled with question marks. At times, we envy the naivety of youth.

Nowhere is that tension between conviction and hesitation more visible than in today's debate around artificial intelligence ("AI"). AI has consumed financial markets. To call it the most significant technological revolution since the dawn of the internet doesn't feel like an exaggeration, it might even be bigger. By most accounts, we're still in the early innings of adoption. Jensen Huang, CEO of Nvidia, has gone as far as to call this "the next industrial revolution," with AI's addressable market encompassing the knowledge economy itself, the overwhelming majority of global GDP.

Countless market veterans see the makings of a bubble, while others see a once-in-a-lifetime investment opportunity. For more than two years, we've struggled with this topic. We see the potential rewards but find it hard to shake our growing list of concerns. This commentary is our attempt to address the state of the market amidst the AI boom.

Michael Cembalest, JP Morgan

The tale of two economies

It's abnormal for us to devote an entire commentary to one theme. Our investments span multiple asset classes, both public and private, and are global in nature. There is always too much to talk about. The problem, however, is that AI has consumed both the economy and financial markets. It's difficult to invest in anything today without expressing a view on AI, implicitly or explicitly. The simple reality, AI is driving the economy and markets and little else seems to matter.

Consider the bifurcated state of our economy. On the one hand, economic growth is strong and appears to be accelerating. Last quarter, real GDP grew 2.1% year-over-year, and the Atlanta Fed's nowcast for third-quarter GDP sits at 3.8%. On the other hand, the Fed saw fit to cut rates due to "downside risks to the labor market." ADP's estimate of private payrolls has been trending lower since last October and suggests jobs were lost in June, August, and September. How do you square strong, real economic growth with a deteriorating labor market? The answer is AI.

The pace of hiring is down

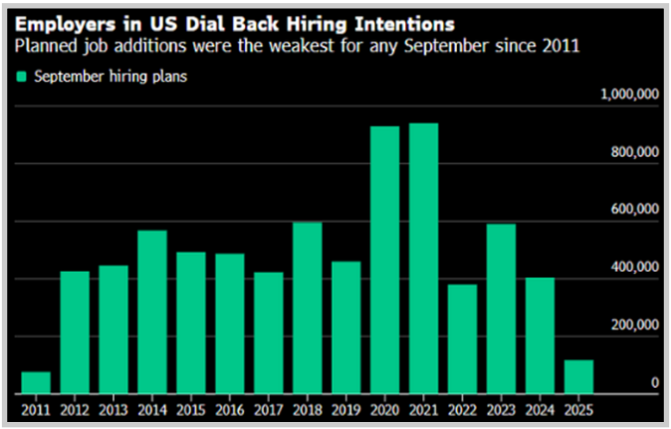

Bloomberg

ISM New Orders (Demand) is down, BUT prices are up

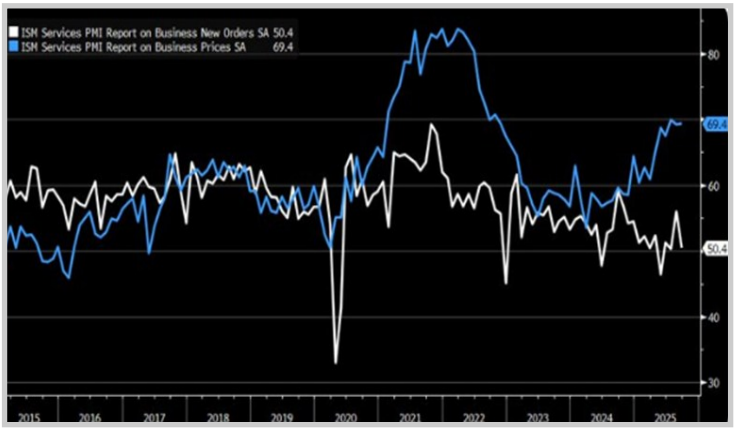

Bloomberg

Just read the comments from the September ISM survey. Seven of the nine quotations mention either high input costs (i.e., margin pressure) or weakening demand. One is neutral. The only positive commentary comes from the information technology sector, driven by AI demand.

- *"We are beginning to see the impact of tariffs on our business… year-over-year cost increases are getting progressively greater."* [Accommodation & Food Services] [NEGATIVE]
- *"New residential construction continues to struggle in a tough market…"* [Construction] [NEGATIVE]
- *"Pharmacy costs continue to rise..."* [Health Care & Social Assistance] [NEGATIVE]
- *"Demand for artificial intelligence (AI) and cloud infrastructure remains very strong. Our primary focus this month was increasing production throughput to begin clearing the significant order backlog built up over the summer… the overall business outlook remains positive. We are still facing significant supply chain challenges… with lead times remaining extended."* [Information Technology] [POSITIVE]
- *"Client demand in professional services remains steady, though decision-making timelines are lengthening due to continued uncertainty… we are also seeing modest upward pressure on labor costs…"* [Professional, Scientific & Technical Services] [NEGATIVE]
- *"The overall housing market remains stagnant, which has forced our company to be hyper-vigilant about costs… Tariffs continue to inject an unnecessary level of uncertainty across the broader economy, and costs are now beginning to increase with the full effect of tariffs coming into play."* [Real Estate, Rental & Leasing] [NEGATIVE]
- *"Costs overall have stabilized, and we've not seen any interruption in sourcing or shipments."* [Retail Trade] [NEUTRAL]

~ "We've had more tariff charges last month than in previous months." [Utilities] [NEGATIVE]

~ "Business conditions continue to soften, even in markets that have historically been more resilient. Demand is simply weak. [Wholesale Trade] [NEGATIVE]

We don't think this bifurcation in our economy is the result of some dystopian scenario where AI has replaced white-collar workers, driving productivity gains at the expense of employment. That may come later. No, this is a good old fashion capital spending boom.

Capital expenditures on AI datacenters, spanning chips, land, equipment, and construction, are so large that they're boosting GDP. JP Morgan estimates that AI-related capital spending accounted for more than 40% of real GDP growth last quarter. In September, economists at Deutsche Bank noted, "AI machines – in a quite literal sense – appear to be saving the US economy right now… Nvidia, the key supplier of capital goods for the AI investment cycle, is currently carrying the weight of US economic growth… and that growth is not coming from AI itself, but from building the factories to generate AI capacity."

The estimates vary, but they all tell the same story: we have two economies, one in or near recession, and one booming on the back of AI data centers.

"Without tech the US economy would be close to recession"

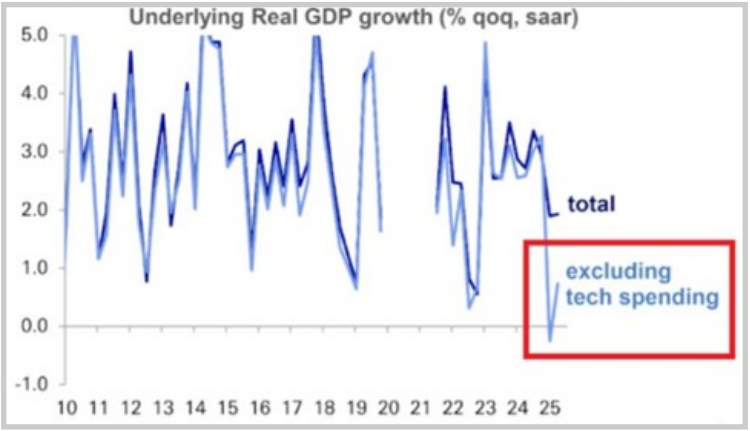

Deutsche Bank

"Without data centers, GDP growth was 0.1%"

Fortune

The tale of two stock markets

Unsurprisingly, the same story holds true for the capital markets. OpenAI's ChatGPT captured the world's attention in November 2022. Since that "ChatGPT moment," AI-related stocks have grown earnings 124%, versus just 9% for the S&P's non-AI constituents. Operating income is up 98% compared to only 16%. Put differently, AI stocks accounted for roughly 79% of the S&P 500's earnings growth. We could quibble over the numbers, but they approximate reality and align with the economic data outlined earlier: abnormal profit growth for companies supplying the AI revolution, and marginal growth for everyone else.

Growth & contribution to the S&P 500

	AI: Direct	AI: Utilities	AI: CapEquip	S&P 500 ex-AI
Performance since November 2022				
Price return	181%	65%	138%	25%
Earnings growth	124%	15%	58%	9%
EBIT growth	98%	11%	71%	16%
Capex + R&D growth	63%	21%	-14%	4%
Contributions to S&P 500 since November 2022				
Price return	75%	0.9%	0.9%	23%
Earnings growth	79%	0.5%	0.8%	20%
EBIT growth	62%	0.4%	0.9%	36%
Capex + R&D growth	90%	2%	-0.1%	8%

JP Morgan

As you'd expect, stocks benefiting from the AI trade have delivered exceptional performance. The S&P 500 Information Technology Index is up 23.8% year-to-date, compared to just 9.9% for the S&P excluding that sector. Keep in mind, Amazon, Google, and Meta aren't considered information technology stocks, nor are the utilities and industrials supporting the data center buildout. If the S&P published an S&P AI Index and an S&P ex-AI Index, the performance differential would be far more dramatic.

AI stocks have accounted for the majority of market gains

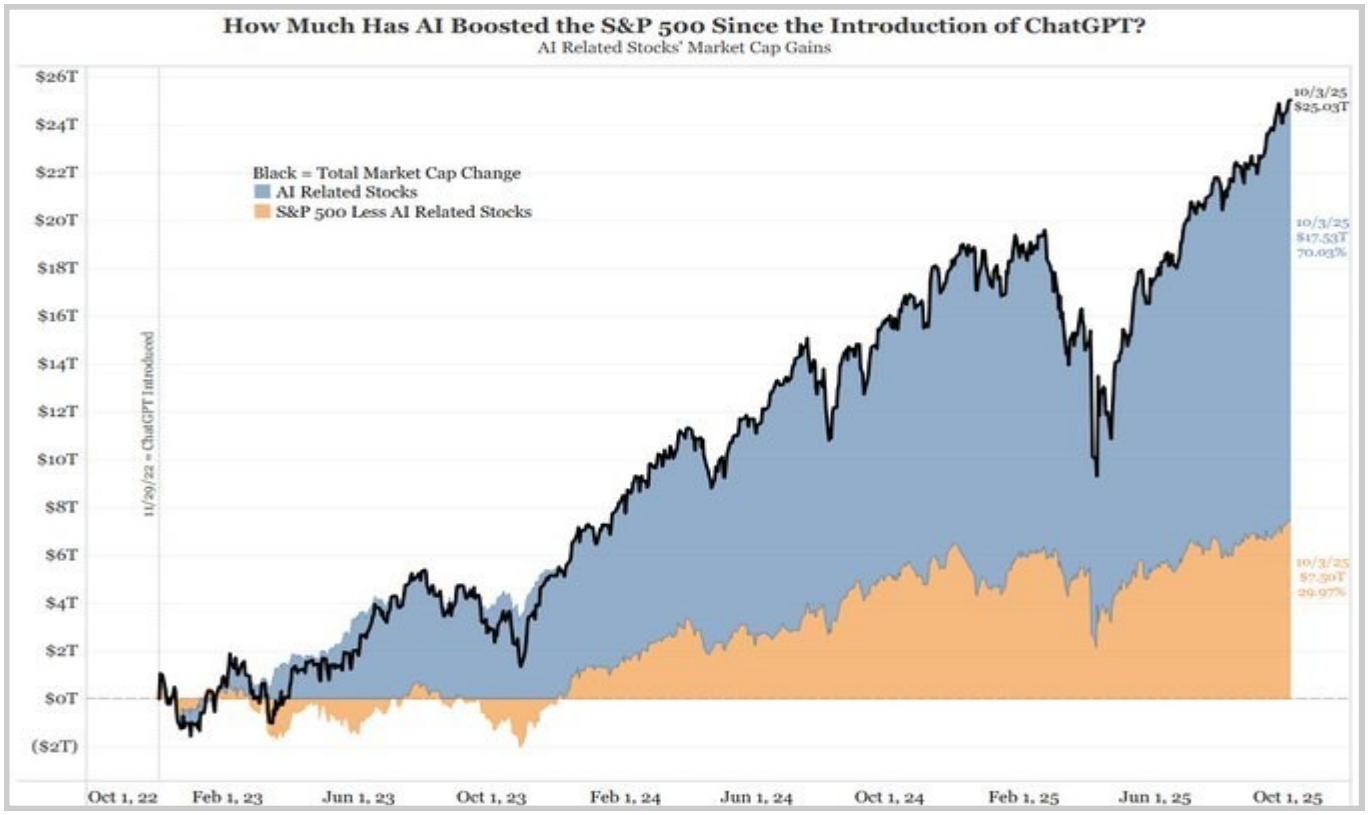

Bianco Research, JP Morgan

JP Morgan estimates, since the "ChatGPT moment," S&P constituents tied to the AI trade have returned 181% compared to 25% for those not associated with AI. The Goldman Sachs TMT AI Index is up 42.4% this year and now carries a total capitalization of $29.2 trillion. For context, the S&P is worth $56.7 trillion. The easiest way to grasp the US capital market's reliance on AI is to look at its largest companies: Nvidia (8.0%), Microsoft (6.7%), Apple (6.6%), Google (4.5%), Amazon (3.8%), Broadcom (2.8%), and Meta (2.8%). These seven companies, accounting for 35% of the S&P 500, are all zeroed in on the AI race. The eighth largest company in the world isn't US based, but it is Taiwan Semiconductor, another cornerstone of the AI supply chain. The information technology sector alone represents 35% of the S&P 500. If you added Google, Meta, and Amazon, it would exceed 46%. And if we included the industrials and power producers tied to the AI boom, at least half of US equities (by weight) would effectively be an AI trade.

To summarize, we have two economies, the AI economy and the "normal" economy," and two stock markets: the AI haves and the AI have-nots. You can't discuss the US economy or financial markets without crossing this topic. As Dylan Patel, founder of SemiAnalysis, put it, the AI race represents "the highest stakes game of capitalism of all time," and if AI models stopped improving, the US economy would be "absolutely screwed."

Jensen's TAM Math

Now that the stage is set, let's address the bull case from the view of an AI maximalist, Nvidia CEO Jensen Huang. Huang sees AI as the next industrial revolution, and he offers two primary reasons why.

First, he believes the era of general-purpose computing (CPUs) is over; the future is accelerated computing. Over time,

workloads handled by traditional data centers will migrate to AI data centers. You can already see this shift in "recommendation engines." Accelerated computing is better suited for showing you a Facebook ad, suggesting an alternative in your Amazon cart, or recommending a movie on Netflix. That said, it's unclear whether it's needed for simpler, linear tasks, like calculating your total bill (taxes included) when you check out of that same cart.

Second, and this is his bigger point, Huang argues that new applications will emerge from artificial intelligence. AI models, unlike humans, can run 24 hours a day, seven days a week. While we're sleeping or cooking dinner, they can perform work on our behalf. Already, we use ChatGPT daily for work-related questions and let it "think" while we move on to other tasks. Eventually, we may have AI agents capable of managing our personal and professional lives. For example, an AI assistant with access to my calendar might alert me to a flight delay, propose alternate routes, and automatically reschedule affected meetings. The potential applications are nearly limitless. All of this thinking and activity will be executed at "AI factories," Huang's term for AI data centers.

Defining the total addressable market (TAM) for AI is difficult because, if you let your imagination run, the use cases seem almost unlimited. Initially (in 2023 and 2024), Huang cited roughly $1 trillion in cumulative capital spending, his estimate of what it would take to replace general-purpose computing with accelerated computing. Roughly $1 trillion had already been spent on legacy data centers, so another trillion would be needed to modernize the infrastructure. On an annualized basis, that implies $200–$250 billion in capital expenditures over the next four to five years.

By 2025, Huang broadened his estimates: "We see $3 to $4 trillion in AI infrastructure spend by the end of the decade," and "$1 trillion annually by 2028." In a September interview, that figure grew again, to $5 trillion by 2030. His rationale extends well beyond the replacement cycle. Knowledge work, physicians, lawyers, analysts, scientists, and others, accounts for 55–65% of global GDP, roughly $50 trillion of output. If AI can make these workers "two or three times" more productive, Huang argues, why wouldn't we pursue it? He posits $10 trillion of augmented output at a 50% gross margin, which implies roughly $5 trillion in cumulative capital investment.

To Huang's credit, the hyperscalers (Google, Amazon, Microsoft, Meta, and now Oracle) have consistently surprised investors with their capital expenditure guidance. Last year, the big four hyperscalers spent $231 billion. This year, Citi projects $393 billion in AI-related capital spending, a 70% year-over-year increase. Citi's estimates for the hyperscalers alone achieve 63% of Huang's $1 trillion 2028 run-rate target, and probably closer to 70% if Oracle is included. The remaining growth would come from sovereigns and "neo-clouds" such as CoreWeave and Nebius.

It's also worth noting that Citi's long-term forecast, $5.5 trillion in incremental spend and a $2.2 trillion annual run-rate, closely mirrors Huang's projections. Whether that alignment reflects shared conviction or simple deference is an open question.

Citi's AI datacenter capital spending projections (per annum)

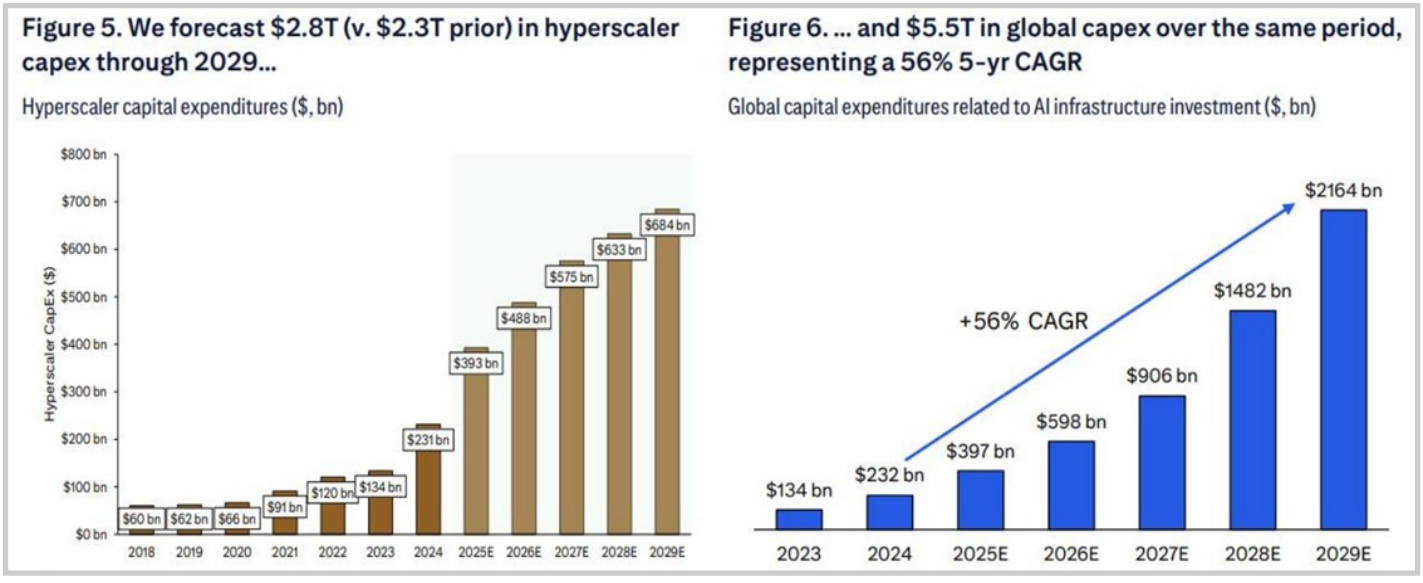

Capital Spending in Context

We have no clue what the actual numbers will be, 2030 is a long way off, and there are countless embedded assumptions in both Jensen's and Wall Street's estimates. What we can say is that the scale of the numbers being tossed around is mind-boggling.

For example, the combined operating earnings of all 500 S&P constituents totaled $1.76 trillion over the past 12 months. The estimates from Jensen and Citi suggest that annual AI capital expenditures could reach 1.23 times that amount by 2029. If spending hits $2 trillion annually, that would represent roughly 7% of current U.S. GDP. Keep in mind, this figure refers only to data center capital spending, it doesn't include the additional expenditures required by power providers to expand capacity, or the multiplier effects as this spending ripples through the economy.

US data center capacity (GW)

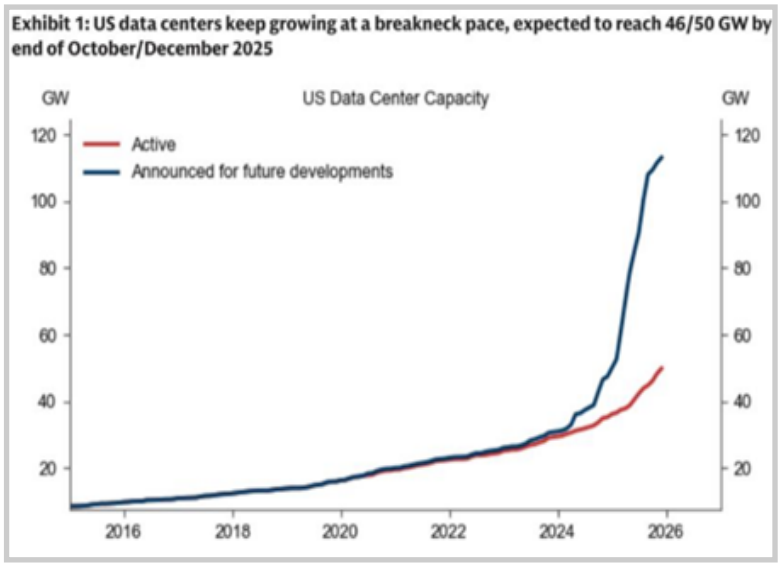

Goldman Sachs

Another way to grasp the scale is through power consumption, the primary bottleneck in bringing AI data centers online. A 1-gigawatt AI data center costs roughly $50 billion to build, so $5.5 trillion in cumulative spending implies around 110 gigawatts of planned capacity. That aligns with McKinsey's estimate of 112 gigawatts of incremental demand between now and 2030. For context, Texas all-time peak power consumption was 86 gigawatts. From the perspective of power generation, we're talking about 50–55 Hoover Dams or 22–24 Plant Vogles, the nation's largest nuclear power plant.

Perhaps the easiest way to comprehend the data center buildout is visually. Below are images of the Stargate Data Center outside Abilene, Texas, taken by Citrini Research. The site spans 875 acres, enough to fit Central Park or roughly 662 football fields. The facility is scheduled to come online in mid-2026 with 1.2 gigawatts of capacity and 400,000 Nvidia GPUs. This campus represents just one of five Stargate sites and only a fraction of the total capacity under construction. It isn't even the largest data center being built; that title belongs to Meta's Hyperion campus, with a footprint two and a half times larger.

Arial photo of the Stargate Abilene Data Center

Citrini Research

And if you listen to Sam Altman, co-founder of OpenAI, everything just described would sound almost trivial. In August, he told reporters, "You should expect OpenAI to spend trillions of dollars on data center construction in the not very distant future." His target is 250 gigawatts of capacity by 2033 for OpenAI alone, about 20% of total U.S. generating capacity. That said, Altman is wearing a sales hat, in a constant race to raise funds for OpenAI.

OpenAI wants to grow capacity 125x

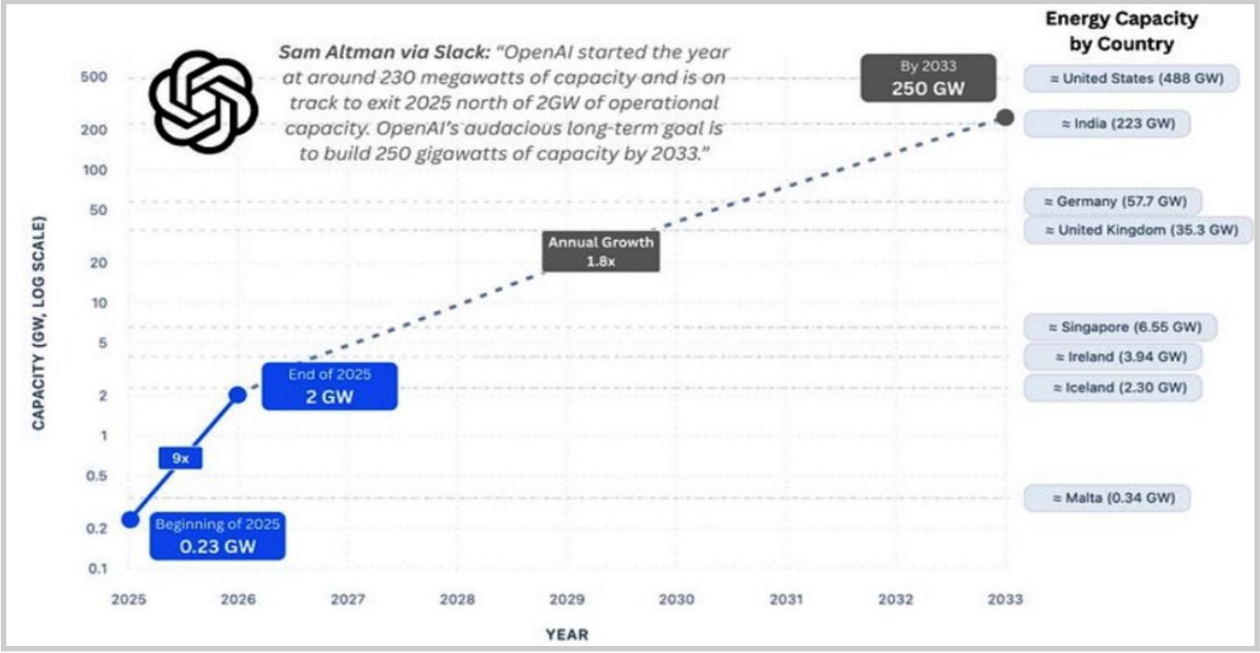

Alex Heath

Financing the Boom

Funding a construction boom of this scale isn't easy, but no one is better positioned than the hyperscalers racing to dominate AI. Google, Amazon, Microsoft, and Meta all have the cash flow and balance sheets to finance massive investments without relying on outside capital. They also have their own internal use cases for AI computing beyond renting chips to customers. Still, even these giants are approaching the limits of what they can fund through operating cash flow. Meta, for instance, started the year with $30 billion of cash net of operating leases and debt, but as of last quarter end carried net cash of negative $2.5 billion. As a result, many are seeking external financing and strategic partners, and Meta has already tapped the debt market to the tune of $39 billion.

Hyperscalers are spending the majority of cash flows on data center CAPEX

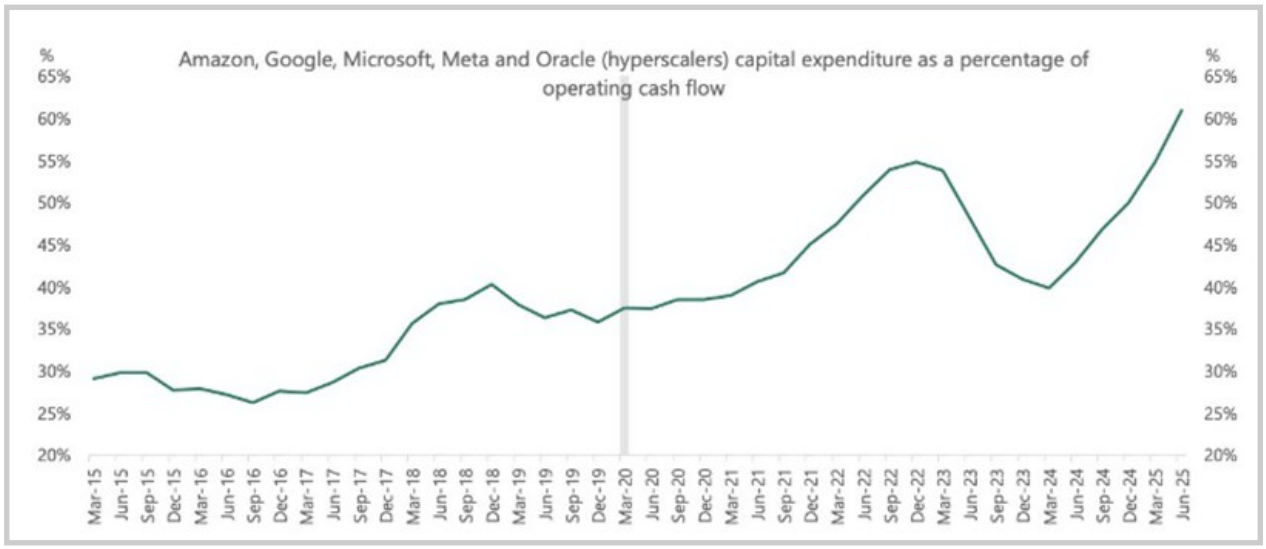

Apollo

More recently, new entrants with weaker balance sheets and thinner cash flow profiles have joined the fray. The most obvious example is Oracle, which now boasts a $50 billion run-rate data center business. On September 9, Oracle's market capitalization jumped 36% in a single day after reporting $455 billion in remaining performance obligations (RPO), a 359% year-over-year increase. RPO represents contracted future revenue, akin to bookings. During its quarterly earnings call, Oracle revealed it had signed four multibillion-dollar contracts with three different customers and expected to add several more, likely pushing RPO beyond $500 billion. Overnight, Oracle went from a major data center player to a hyperscale-level competitor. We later learned that one of those contracts was a $300 billion cloud computing deal with OpenAI.

Oracle's remaining performance obligations by quarter

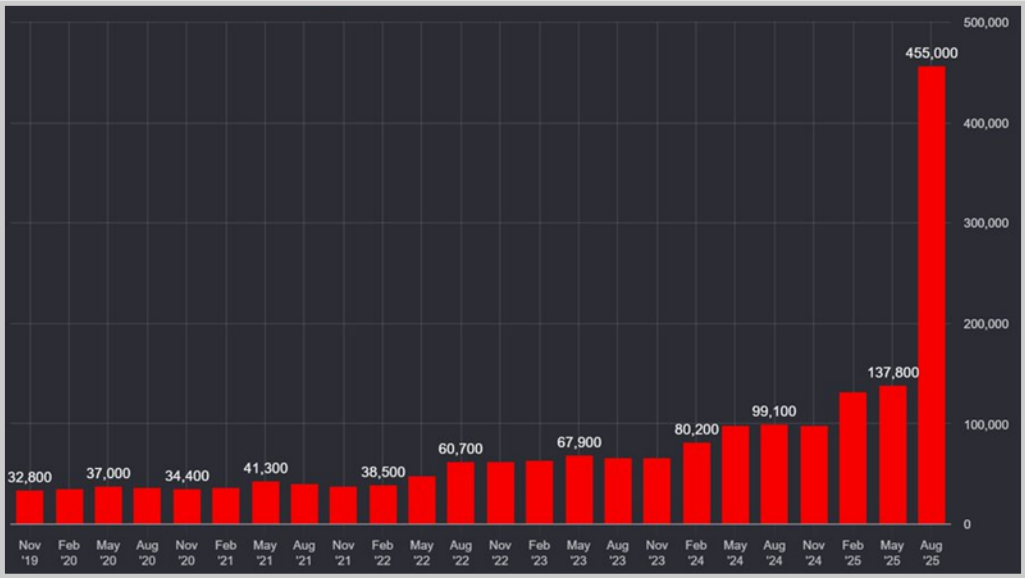

Fiscal AI

We've also seen the rapid rise of neo-clouds such as CoreWeave, Lambda Labs, and Nebius. Many of these firms are former crypto miners who repurposed their infrastructure to capitalize on the AI boom. For perspective, CoreWeave now carries a $76 billion enterprise valuation, and consensus estimates have ballooned from $5 billion in 2025 run-rate revenue to $44 billion by 2029.

The broader point is that the data center market appears to be fragmenting. Consolidation may come once growth slows, but for now, new entrants are pouring in, and this fragmentation seems partly by design. Nvidia, the dominant GPU manufacturer, effectively dictates who can purchase its chips. Its two largest customers account for 39% of total sales, giving Nvidia incentive to diversify and fragment the customer base. There's also another reality: the primary tenants of AI data centers, model developers like OpenAI and Anthropic, may not want the bundled services hyperscalers provide, preferring instead the bare-metal environments offered by the neo-clouds. Will AI datacenters drive margin by offering value added services or will they be commoditized as we move forward?


How Jensen Huang is Using Nvidia Cash to Rule the AI Economy

By Anissa Gardizy, Nick Wingfield, Wayne Ma and Qianer Liu

The Information

Bubble Vibes: Massive Scale

The argument that we're entering a bubble stems from a few concerns.

First is the sheer scale of everything. As famed investor David Einhorn put it, "The numbers that are being thrown around are so extreme that it's really, really hard to understand them." Hopefully, we've done an adequate job illustrating that scale earlier in this commentary. Whenever spending reaches such magnitude, caution flags should go up, because it carries serious implications for the broader economy.

If today's AI boom were to stall, a recession would almost certainly follow. In 2026, the hyperscalers alone are projected to spend 1.7% of U.S. GDP on capital projects. Imagine if the spending and the multiplier effects it drives throughout the economy vanished overnight. The current level of economic and financial concentration creates fragility.

Bubble Vibes: Big Valuations

Second, stock market valuations are high, and valuations of AI-related enterprises, particularly chipmakers, are even higher. This holds true in both public and private markets.

Let's start with the private side, since that's the simpler of the two. Earlier this month, BlackRock-backed Global Infrastructure Partners acquired Aligned Data Centers for $40 billion, implying a 40x revenue multiple. On the venture capital front, nearly all new funding is flowing into AI deals. To be fair, most of that capital is going to Anthropic and OpenAI, but smaller rounds are being struck at equally eye-watering valuations. Thinking Machine Labs, for instance, raised $2 billion at a $10 billion valuation, with no product, no revenue, and no customers, based purely on the promise that OpenAI's former CTO, Mira Murati, would eventually build something of value.

On the public equity front, valuations are more hotly debated. Many observers are quick to point out that this isn't the same as the dot-com bubble, at least not yet. We concede that point; we aren't at the extremes of 1999 or early 2000. But we'd push back on the "all-clear" narrative. The benchmark for "froth" shouldn't be the absolute peak. If you instead compare valuations to an earlier point, say, the start of 1999, the Nasdaq still fell roughly 50% from that level.

At the start of 1999, Cisco carried a $105 billion market capitalization, far below its eventual $530 billion peak. Its valuation was 15x sales, 74x earnings, and 56x free cash flow, all on a trailing basis. Nvidia, by comparison, trades at 27x sales, 54x earnings, and 62x free cash flow. The two are in the same valuation ballpark, with some trade-offs: Nvidia is more expensive relative to sales, cheaper relative to earnings, and roughly equivalent on free cash flow, which makes sense, given Nvidia's current pricing power. The same conclusion holds using forward multiples.

AI funding crowds out the venture market

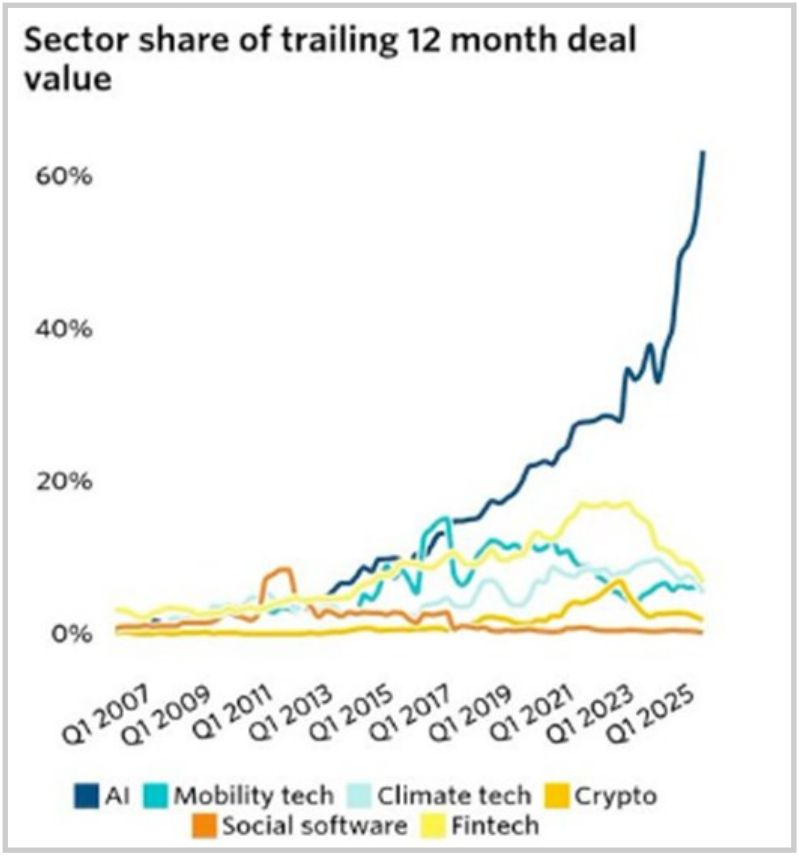

Quantitative Perspectives: A Fork in the Road

Had you purchased Cisco on January 1, 1999, your peak-to-trough decline would have been -84.1%. Even holding for the next decade, your cumulative return would have been –29.8%, or –3.5% per annum. Yes, there was a 246% rally from early 1999 to the peak, but those gains evaporated, and you ultimately lost most of your principal. Our point is that, if a dot-com comparison exists, it's closer to 1997 or 1998, not the late-1999 or early-2000 mania, although that's far from reassuring.

Not yet 1999/2000 bubble valuations

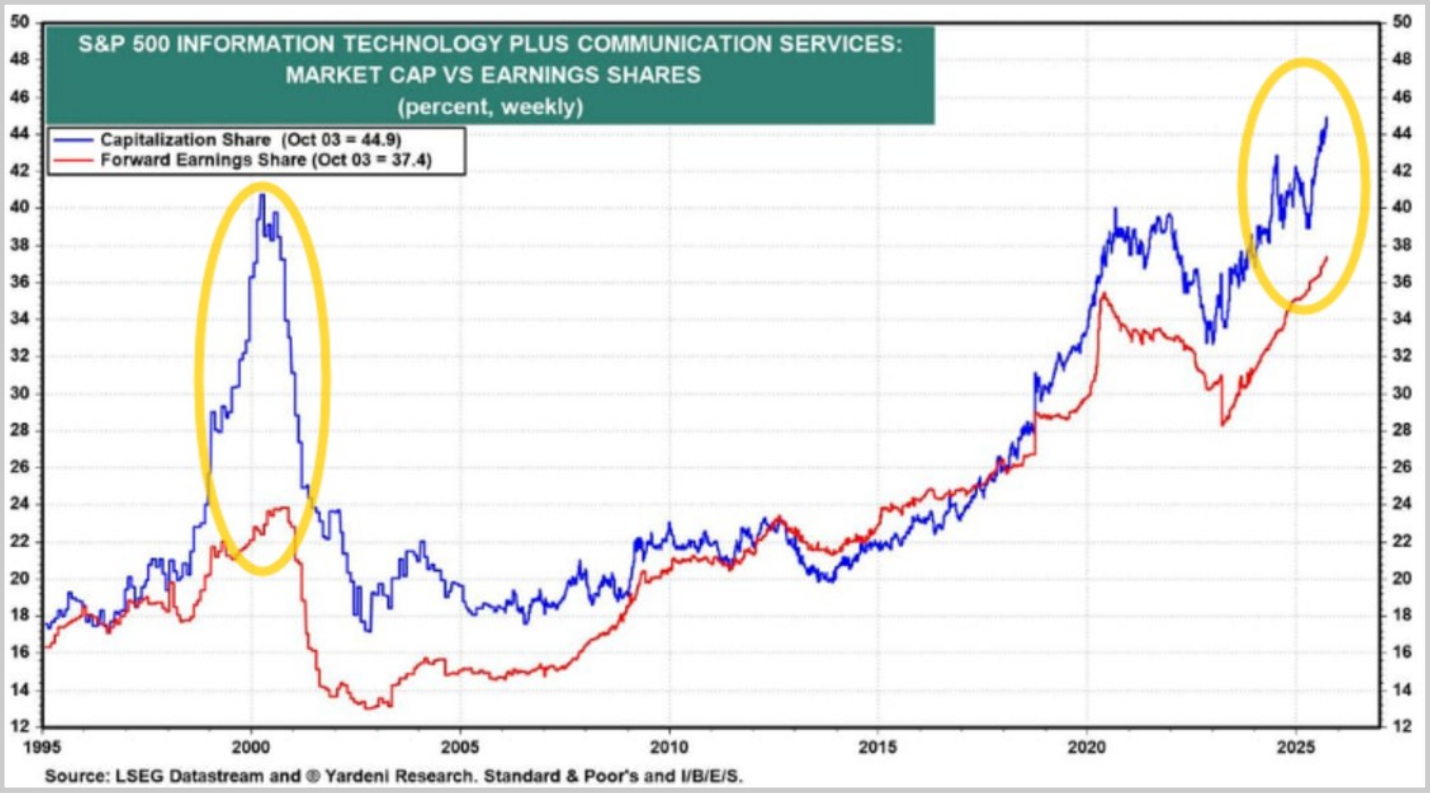

Ed Yardeni

Furthermore, there are clear pockets of excess. Palantir carries a $432 billion market cap on just $3.4 billion in revenue over the past 12 months, not income but revenue. Speculative energy names are appreciating at exponential rates. Oklo, a fission-reactor startup, has risen 536% to a $19.9 billion valuation, despite having zero revenue. Fermi, a data center developer still in "phase zero," went public at a $13 billion valuation; its main asset is 5,236 acres under a 99-year ground lease and a promise to deliver 11 gigawatts of future capacity. The good news: these kinds of sky-high valuations represent only a small fraction of the market. But they exist, and they evidence the market's current appetite for risk.

Ultimately, everything ties back to data center capital spending. If Huang's or Citi's long-term projections are accurate, then this cycle probably has room to run. If capital spending projections fall short of expectations, then the capital markets will need to reset. The ramps in speculative stocks are mostly a distraction and are likely to end poorly regardless of how the future unfolds.

Bubble Vibes: Sketchy Financing

Third, and most concerning, is the sudden shift in financing activity.

The initial funders of the datacenter buildout, Google, Meta, Amazon, and Microsoft, all have fortress-like balance sheets and substantial operating cash flow. Collectively, these four companies generated $450 billion in cash from operations last year. Nobody else has that luxury, particularly newer entrants.

Morgan Stanley's Oracle pro-forma credit metrics

ORCL Debt &Leverage Forecasts *($bn)*	FY25 May'25	FY26E May'26	FY27E May'27	FY28E May'28
CFO (MS ER, adj. to reflect fin. lease interest)		25	28	37
Cash capex (MS ER)		(36)	(43)	(51)
FCF		**(11)**	**(14)**	**(15)**
Dividends		(6)	(6)	(6)
Est. finance lease principal payment		(2)	(5)	(10)
Cash burn after divi. &cap lease principal		**(19)**	**(25)**	**(31)**
Bond maturities		(5)	(5)	(6)
TL maturities		0	0	(5)
Issuance needs		**(23)**	**(30)**	**(41)**
Total debt	93	111	136	167
Less: Cash (incl. equiv. &marketable sec.)	11	11	11	11
Net balance sheet debt	**81**	**100**	**125**	**156**
Total cash capex (GAAP)		36	43	51
Assumed PPE acquired under cap. lease		25	41	65
Capital investment		62	83	116
Operating lease liabilities	13	16	16	16
Finance lease liabilities	3	27	63	118
Total commenced lease liabilities at period end	16	43	79	134

ORCL Implied Rating Agency Adj. Forecasts *($bn)*	FY25 May'25	FY26E May'26	FY27E May'27	FY28E May'28
Balance sheet total debt	93	111	136	167
Total lease liabilities	16	43	79	134
Pensions &other	3	0	0	0
Moody's adj. gross debt	**112**	**154**	**215**	**302**
Balance sheet net debt	81	100	125	156
Total lease liabilities	16	43	79	134
Pensions &other	2	0	0	0
S&P adj. net debt	**100**	**143**	**204**	**290**
GAAP Income before income taxes	14	16	18	25
Interest	4	5	8	12
D&A	6	12	18	24
Fin lease amort (assume agency adj.)	0	2	5	10
Op lease cost (agency adj.)	2	3	3	3
Adj. EBITDA	26	38	52	73
Add'l Moody's EBITDA adj. - Pension	0	0	0	0
Add'l S&P EBITDA adj. - SBC	5	5	7	10
Moody's adj. EBITDA	26	38	52	73
S&P adj. EBITDA	30	43	58	83
Est. Moody's adj. gross lev (3.5x d/g threshold)	***4.3x***	***4.1x***	***4.2x***	***4.1x***
Est. S&P adj. gross lev (4x d/g threshold)	***3.3x***	***3.3x***	***3.5x***	***3.5x***

Morgan Stanley

And who's the end customer promising to pay Oracle's tab? OpenAI. They reportedly account for at least $300 billion of Oracle's $455 billion in remaining performance obligations. By most accounts, OpenAI leads (or ties with Google) in the AI race. But unlike Google, OpenAI relies on venture funding and is hemorrhaging cash. Goldman Sachs estimates that for every million tokens processed, OpenAI generates $4.29 of revenue but loses $4.08 after expenses. The bulk of those costs come from inference ($3.36) and model training ($3.58) per million tokens. Building and operating large language models is extraordinarily expensive.

OpenAI's bet, one we agree with, is that compute costs will fall over time. But for now, the business is deeply unprofitable, and its internal projections reportedly show accelerating losses for years to come.

OpenAI's cash flow projections

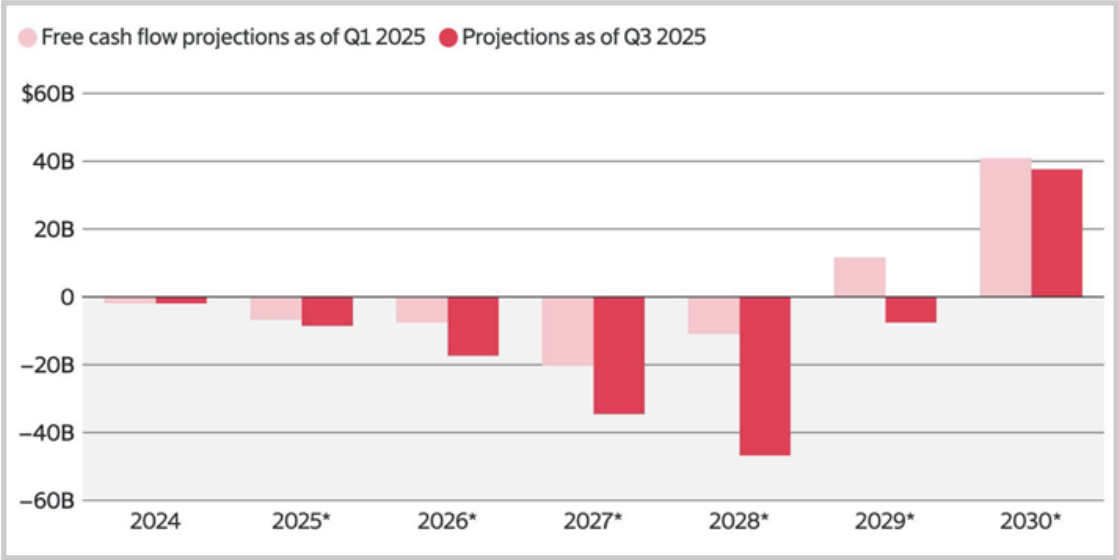

The Information

OpenAI's ambitions are vast, and with Google as its main competitor, the pressure to scale quickly is immense. The problem is a lack of stable funding. And that challenge extends beyond OpenAI. Whether you're a model provider (leasing data center capacity) or a neo-cloud operator (leasing it out), cash burn is extreme. The only exceptions are the legacy hyperscalers. As a result, a wave of creative financing structures have emerged, allowing OpenAI and others to secure data center capacity, leased or owned.

Two of the most notable examples are OpenAI's partnerships with Nvidia and AMD. Under the Nvidia agreement, Nvidia plans to invest up to $100 billion in OpenAI over the next decade, with those funds used to purchase Nvidia GPUs. The AMD deal is even more complex: OpenAI agreed to buy roughly $90 billion of AMD GPUs, and AMD granted OpenAI penny warrants for 160 million shares. If AMD's stock exceeds $600 per share, the warrants would be worth $96 billion, roughly equal to OpenAI's purchase commitment.

Both arrangements are circular. Nvidia injects capital into OpenAI, who then uses it to buy Nvidia GPUs. AMD grants OpenAI warrants, which can be monetized to fund GPU purchases. And these aren't isolated cases. In 2023, Nvidia invested $100 million in CoreWeave, a major GPU customer, while simultaneously agreeing to rent chips back from them. CoreWeave then uses the new capital (and the rental income as collateral) to buy even more Nvidia GPUs. We don't know the total dollars involved, but the AI ecosystem is littered with such financial arrangements.

These financial structures likely aren't outright nefarious. Nvidia's primary goal is probably to deepen dependence on its technology and foster competition among data center operators. But the arrangements are concerning, and to cynics, they resemble financial engineering designed to inflate revenue. Either way, they unquestionably add fragility to the system.

One final, related observation: the web of financial relationships has become uncomfortably complex and intertwined. Consider this: Microsoft accounted for 72% of CoreWeave's revenue last quarter. Although Microsoft operates its own data centers, it still rents capacity from CoreWeave. OpenAI represents the vast majority of Microsoft's AI-related revenue, and, by extension, CoreWeave's. Earlier, we noted that OpenAI also makes up roughly two-thirds of Oracle's contracted revenue. OpenAI has unique financing deals with Nvidia and AMD, and Nvidia maintains similar arrangements with CoreWeave and others. OpenAI alone is said to have commitments for computer infrastructure totaling $1 trillion. We sympathize with the analysts trying to track these flows. But one conclusion is clear: the entire ecosystem is riding on OpenAI's coattails and dependent upon its continued growth.

Good Luck Following the Money

Exhibit 3: AI Ecosystem Capital Flows

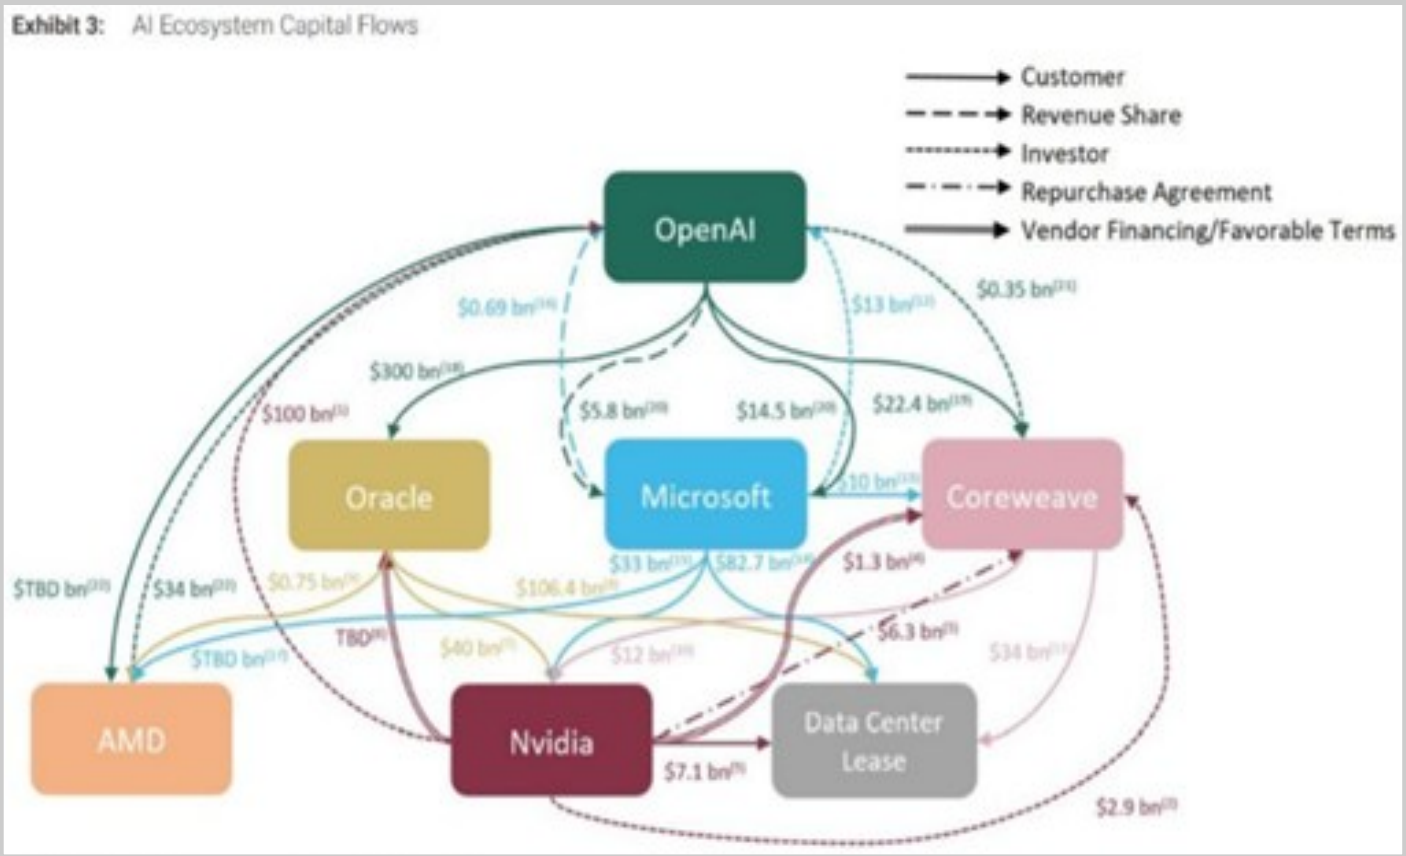

Goldman Sachs

Implied Revenue & Profitability

It's tempting to jump on the "bubble" bandwagon. Most of the necessary ingredients are in place. A promising, likely revolutionary technology, AI, captures the world's imagination. The capital required to make it reality is immense. Optimistic about the future, investors bid up valuations for all things AI-related. Even speculative "lottery tickets" find a bid.

The cost of AI infrastructure exceeds what the market can organically supply, so debt financing and alternative funding structures proliferate. The market rewards those able to spend aggressively, driving valuations higher and encouraging even more spending. Eventually, both the market and the economy go all-in on AI. The bubble vibes are unmistakable, but they remain just that: symptoms, not proof.

A formal diagnosis won't be possible until after the fact. For now, we're comfortable saying the probability we're in an AI bubble is elevated, perhaps north of 50%. If valuations climb further, approaching the extremes of late 1999 rather than early 1999, we'd raise that probability to say 75%. But absolute certainty is impossible.

The most important question is whether AI will generate enough revenue, and in turn earnings, to justify the price tag. To date, between $700 and $800 billion has been spent on data centers. Next year's spending will add an additional $600 billion to the tab, and the projected figures keep rising from there. To justify that level of investment requires enormous future revenue.

In June 2024, David Cahn, a partner at Sequoia Capital, posed this exact issue in his piece "AI's $600B Question." His premise was simple: whoever rents capacity from the data center, say OpenAI, must earn a margin. Assuming a 50% margin, $600 billion in capital spending implies $1.2 trillion in lifetime revenue. Keep in mind the tenant has other costs, and the datacenter itself expects to earn an economic profit. Using that same logic, today's reality might better be framed as "AI's $1.5 Trillion Question." And if Jensen Huang's or Citi's 5-year forecasts play out, it could soon become "AI's $12 Trillion Question."

Alternatively, data center investment can be viewed through a profitability lens, specifically, return on invested capital ("ROIC"). The formal equation is net operating profit after taxes ("NOPAT") divided by invested capital ("IC"), but intuitively, it reflects the earnings a company generates relative to the capital required to build that company. It's simply the rate of return on capital spent. Naturally, the market rewards companies with high ROICs through higher valuation multiples. Conversely, if a company earns returns below its cost of capital, equity and debt combined, it destroys value and tends to be punished by the market with lower valuation multiples.

After-tax operating income required for a set ROIC

REQUIRED NOPAT		Cumulative Invested Capital				
		2025	2026	2027	2028	2029
		$ 763	$ 1,361	$ 2,267	$ 3,749	$ 5,913
ROIC	30%	$ 229	$ 408	$ 680	$ 1,125	$ 1,774
	20%	$ 153	$ 272	$ 453	$ 750	$ 1,183
	10%	$ 76	$ 136	$ 227	$ 375	$ 591

Annandale Capital

Applying ROIC hurdles to AI data centers is revealing. Citi estimates total invested capital will reach roughly $762 billion by the end of 2025 and approach $6 trillion of cumulative investment by 2029.

If we make the (arguably generous) assumption that current data center investments will generate returns consistent with the hyperscalers' historical ROICs of around 30%, the required NOPAT would be approximately $229 billion annually. To put that in perspective, Microsoft's after-tax income is about $101 billion, Google's $100 billion, and Meta's $62 billion. In other words, AI data centers would need to produce profits equivalent to the combined net

income of three of the world's largest companies. If cumulative capital spend indeed reaches $6 trillion, the required annual NOPAT rises to $1.7 trillion, roughly the current earnings of the entire S&P 500. These figures strain credulity and suggest that hyperscaler ROICs will likely decline meaningfully in the years ahead.

A 20% ROIC assumption is somewhat more plausible, though still extremely ambitious. It would require about $153 billion in annual NOPAT, roughly the earnings power of two Metas. As investment scales, the required profits grow rapidly. It's difficult to envision AI data centers collectively generating the $1.2 trillion in annual NOPAT implied by a $6 trillion investment base, about 67% of the S&P 500's current net operating income.

Even a 10% ROIC, a far more conservative and realistic target, implies $97 billion in annual NOPAT for today's investment, roughly equivalent to the net income of a Google or Microsoft. That's still a massive figure. While this level of profitability might be achievable, it's far less attractive. If the market were to realize that hyperscalers were earning only such modest returns, valuation multiples would likely compress sharply, contrary to what we're seeing today.

If AI data center profitability falls short of these already demanding benchmarks, investors should start to worry. It would suggest that the $762 billion deployed so far could have been more productively spent elsewhere.

The broader takeaway is that the revenues and profits required to justify today's investments are enormous, and growing. At a minimum, it seems fair to expect that hyperscaler profitability will decline over time. At worst, the required returns may prove unachievable, in which case we'll know, in hindsight, that it was a bubble.

AI Revenues

The good news is that AI-native enterprises, the data-center customers, are ramping revenue at unprecedented speed, reaching $100 million, even $1 billion in annual recurring revenue in a fraction of the time it has taken others. The bad news is that, in the aggregate, the revenue base remains trivial relative to cumulative spending, and the AI data center capital expenditures tab is growing almost as fast. The largest player, OpenAI, has reached $12 billion in run-rate revenue, while Anthropic, the number two, stands around $5 billion. We don't know the ecosystem-wide total, but it's clear there's still a long way to go. And remember, return depends not just on cumulative revenue, but also timing. Eventually, cumulative revenues and after-tax operating profits will be large, but how long will it take to get there?

OpenAI's Internal Revenue Projections

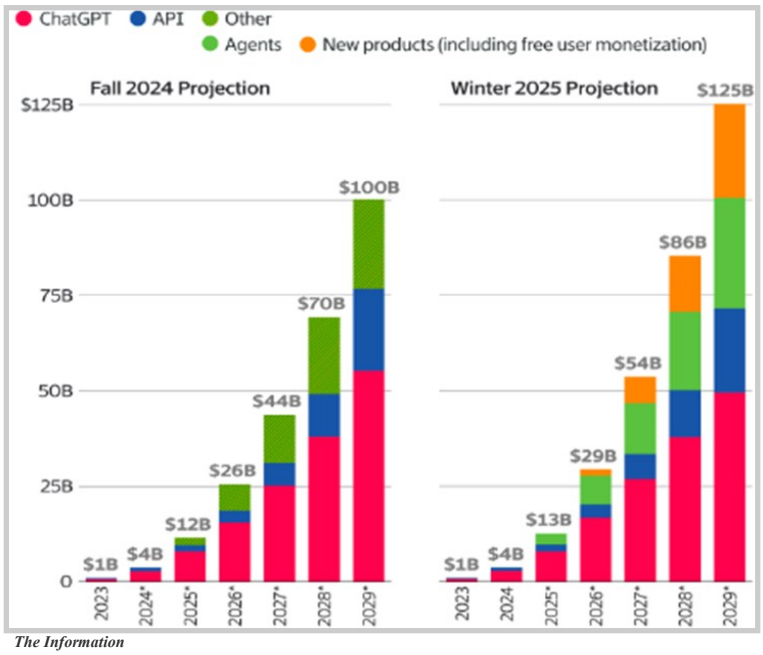

The Information

AI revenue ramps compared to software

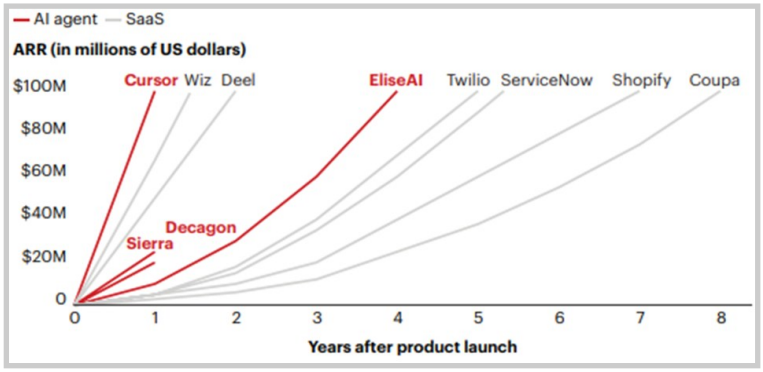

Bain

Because future revenue/profit potential, and its timing, is so difficult to quantify, the market is reliant on real-time barometers of supply and demand. By all accounts, it is abundantly clear that demand exceeds supply. The evidence is everywhere. Microsoft has publicly announced it will be short supply until mid-2026. Nvidia, a chip designer, is earning over 70% gross margins on its

its revenue. Outdated chips, like the Nvidia A100, remain in use. Amazon is raising hourly rates on Blackwell GPU instances. This is great news for the AI bulls, particularly the startup datacenters, like CoreWeave, who need to earn an economic return on their GPU outlays. The problem is that real-time conditions can change quickly.

GPUs continue to be absorbed into the market (note the declining availability)

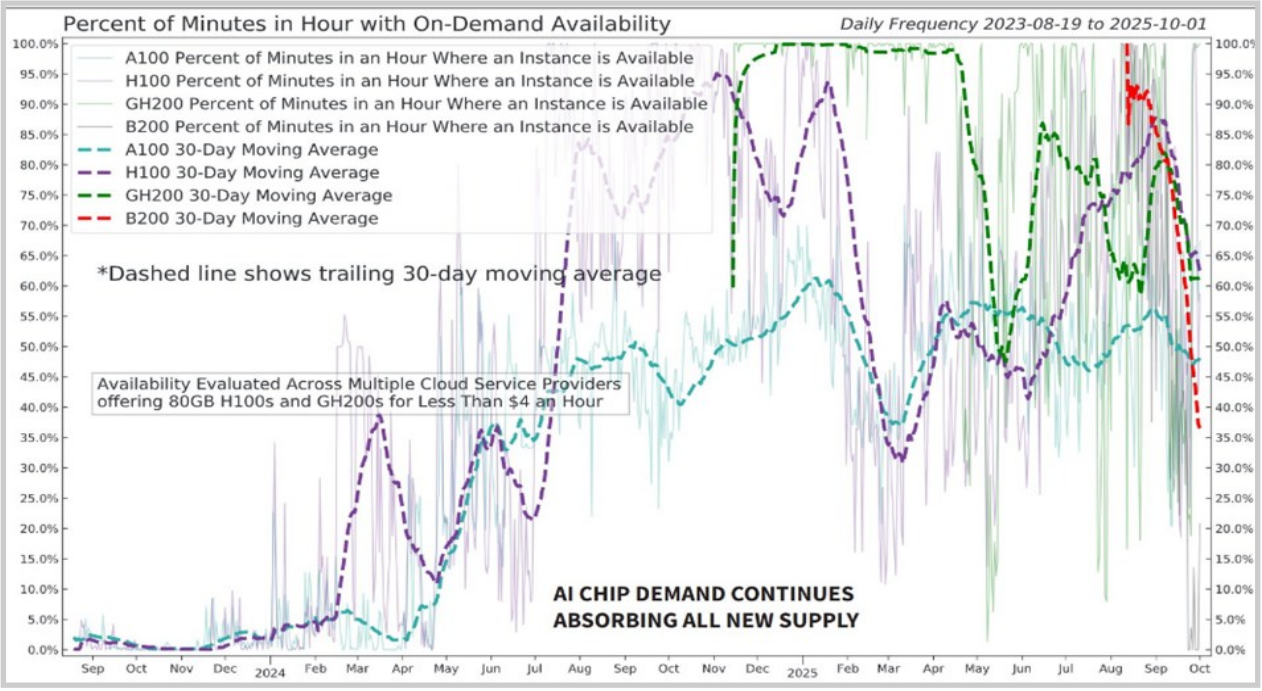

3Fourteen Research

At present, everything is heavily subsidized. Consumers are shouldering the rising cost of electricity through higher utility bills, while OpenAI, like other AI startups, incurs a negative margin on each user. Scale typically improves profitability, but in AI applications, many costs are variable. More users mean more compute, more energy, and higher expenses. What happens to demand if funding for OpenAI or its peers dries up? And what will the economics look like when the incremental data center comes online, finally bringing the market into balance?

A Note of Caution from the Internet & Telecom Boom

To this end, the internet and telecom boom of the 1990s is instructive. In the early nineties, the commercial internet was an emerging technology. Internet traffic was tiny but growing exponentially, doubling roughly every six to twelve months.

In 1996, the Telecommunications Act was passed, deregulating the market and allowing new entrants to build or lease telecommunications infrastructure. It was widely believed that the internet would soon overwhelm existing networks, resulting in a "bandwidth famine." In 1997 and 1998, Global Crossing, Level 3, and Qwest went public, raising large sums to lay fiber, yet the shortage narrative persisted. As PricewaterhouseCoopers noted in its 1998 Annual Technology Review:

"With internet traffic doubling every 100 days… bandwidth will be a major issue facing cable, computer, and telephone companies next year."

We don't know exactly when the narrative flipped, but as Ben Horowitz, co-founder of Andreessen Horowitz, later put it:

"We had the biggest bandwidth shortage in the world in 1999 and the biggest bandwidth glut ever in 2001."

The story turned quickly.

The Nasdaq gave up all of its gains dating back to late-1996

Bloomberg

Furthermore, Cisco, the market darling of its time, peaked well before its fundamentals turned south. The stock hit an all-time high on March 27, 2000. After that peak, Cisco delivered four consecutive quarters of revenue growth, rising sequentially from $4.9 billion to $5.7 billion, $6.5 billion, and $6.7 billion. For three of those quarters, growth wasn't just strong, it was accelerating.

The only real warning sign appeared in early 2001, when revenue growth slowed from 66.4% to 54.9% year-over-year and gross margins contracted. But by then, the stock was already down 55%. The following quarter, over a year after its peak, was disastrous: revenues fell 4% year-over-year, and Cisco recorded the largest impairment charge in corporate history. By the time the bad news showed up in the financials, the stock had fallen 75.5%.

The internet and telecom boom of the late 1990s, without warning or clear catalyst, simply came to an end.

Cisco peaked well before fundamentals

Bloomberg

Conclusion

At the start of our commentary, we expressed our envy for the young. Unburdened by knowledge of the past, they can dive headlong into the AI trade and sleep soundly at night while doing so.

We understand the AI tailwinds. We see the market rewarding anything tied to AI while punishing almost everything else, and we know from experience that things can get even crazier. Yet we can neither build the conviction to go long, nor prove we're in an AI bubble. We sit, aggravatingly, in the middle, perplexed by the scale, the valuations, and the increasingly questionable financing activity.

The price action could get far more extreme (distance between Nasdaq and its 200-day average)

Bloomberg

Normally, our inability to reach a definitive conclusion isn't a problem. Investors have three options, buy, sell, or pass, and we exercise that third option liberally. Just look at our pipeline of private investments, where the discard pile dwarfs the handful we fund each year. Many of those discarded opportunities are perfectly fine, some may even deliver stellar results, but we pass because of unresolved questions. Passing on potentially good investments doesn't bother us; we simply move on to the next one where our uncertainty is lower. We never want uncertainty to translate into capital impairment, and we're willing to miss opportunities to ensure our downside is protected.

The AI trade, unfortunately, has grown too large to simply opt out of. By our count, more than half of the S&P 500 is effectively a play on AI. Choosing to "pass" on AI, therefore, is tantamount to passing on most of the U.S. stock market, the largest and most liquid in the world.

We take comfort in our portfolio's diversification. We view global equity markets as our playground, with roughly 40% of our equities overseas. Our 60% exposure to U.S. stocks includes meaningful exposure to large technology companies, and, by extension, the AI theme, but far less than the S&P 500 itself. Year-to-date, our domestic equities have lagged due to their limited AI exposure, but our overall equity portfolio has outpaced U.S. markets thanks to our global diversification.

When the internet and telecom bubble burst in March 2000, the Nasdaq and S&P 500 fell -77.9% and -42.4%, respectively. Yet large-cap value stocks declined only -19.9%, less than half the broader market's drop, and small-cap value stocks even posted positive nominal returns. No two bear markets are identical, but our domestic holdings currently tilt toward higher-quality, plain-vanilla businesses trading at less demanding valuations. That's where we see the best risk-adjusted opportunity. If we're unlucky enough to witness an AI correction, our hope is that our portfolio will prove more resilient.

Beyond equities, clients own safe-haven assets, cash equivalents, and fixed income, along with private equity investments spanning a range of risk profiles.

We sleep well knowing our portfolio's success is tied to a diverse set of themes, though we remain uneasy about the concentration within the S&P 500. We're also experienced enough to recognize that our concerns are often misplaced. We are professional worriers, constantly thinking about what can go wrong, but more often than not, nothing does. For this reason, it would take far more extreme valuations, akin to 1999, for us to materially change our stance on domestic equity markets.

As always, we will keep you informed as market dynamics evolve, and we deeply appreciate the trust you've placed in us.

IMPORTANT DISCLOSURES *This material is not intended to be used as a general guide to investing, or as a source of any specific investment recommendations, and makes no implied or express recommendations concerning the manner in which any client's account should or would be handled, as appropriate investment strategies depend upon each client's distinct investment objectives. This is not an offer or solicitation with respect to the purchase or sale of any security. Further information on any of the investments mentioned in this material may be obtained upon request. Before making any investment decision, prospective investors should carefully read all material provided. It is not our intention to state or imply in any manner that past results and profitability is an indication of future performance. The attached summary/prices/quotes/statistics have been obtained from sources we believe to be reliable, but we cannot guarantee its accuracy or completeness. Annandale Capital, LLC does not provide tax or legal advice. Please consult your tax or legal advisor.*