GITTERMAN
ASSET MANAGEMENT
A PERIGON PARTNER

379 Thornall Street | 2nd Floor | Edison, NJ 08837
www.gittermanasset.com | Fax: 888-476-1314

Can Hyperscalers Generate Enough Revenue to Justify Their CAPEX Investments?

By Adam Bernstein, Investment Analyst

As concerns surrounding an AI market bubble, S&P 500 concentration, and the sustainability of unprecedented capital spending continue to grow, we believe it is worth examining how the AI infrastructure buildout is unfolding. This is the first in a two-part series exploring whether the hyperscalers[1] making today's unprecedented AI infrastructure investments can ultimately monetize those investments in a way that generates attractive long-term returns for shareholders. In this first piece, we examine whether AI revenues can grow rapidly enough to justify today's infrastructure spending. In part 2, we turn to the costs, margins, and industry dynamics that will determine who captures those returns.

The context

Over the past several years, the hyperscalers' alongside AI cloud infrastructure providers Coreweave and Nebius have committed unprecedented levels of capital to AI infrastructure, which includes servers, data centers, and power generation. The hyperscalers spent much of the last decade returning billions of dollars to shareholders through share buybacks, and they are increasingly redirecting free cash flow toward AI investments. Additionally, they are accessing debt and, in some cases, equity markets, to help finance the spend.

Hyperscaler and AI infrastructure CAPEX

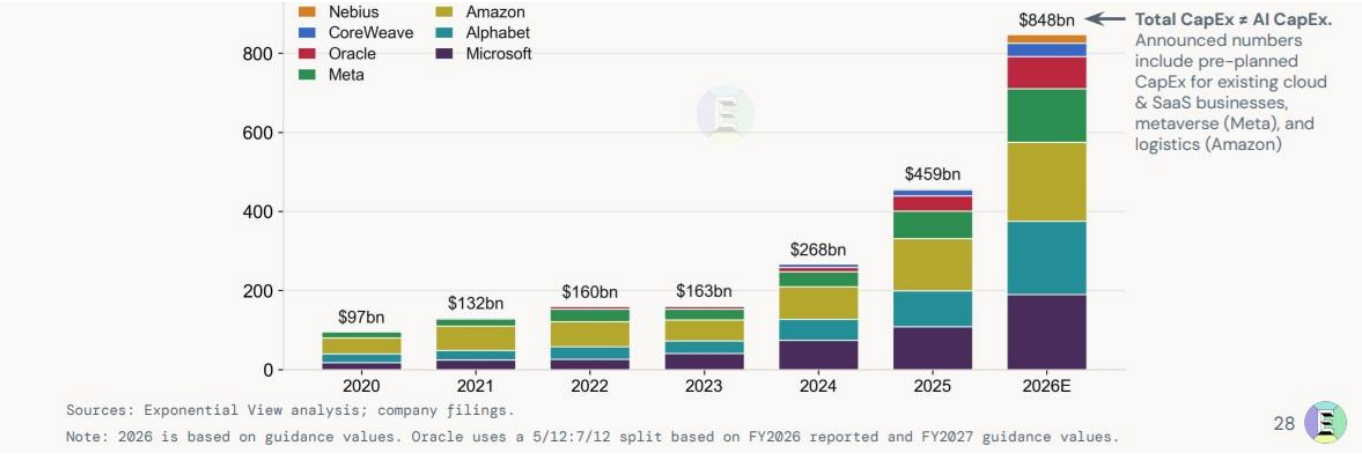

Source: Exponential View

JPMorgan recently estimated that, in the medium term, the cost of the AI buildout could exceed $5 trillion by 2030. As such, generating "a 10% return on current AI investments could require $650

379 Thornall Street | 2nd Floor | Edison, NJ 08837
www.gittermanasset.com | Fax: 888-476-1314

billion in AI-related annual revenue."[2] Based upon a comprehensive bottom-up assessment of AI monetization by Exponential View, a leading AI technology research firm, the global generative AI ecosystem is currently generating ~$175 billion in annualized revenue over the 12-month period to the end of June 2026.[3] This figure represents U.S. based AI companies, excluding revenue from Chinese firms like Alibaba, Tencent, and ByteDance. While these estimates are not perfectly comparable, each uses a different definition of the AI ecosystem, and they illustrate the scale of monetization still required relative to the capital being deployed.

Is this gap a long-term problem?

There are significant risks given the size of the gap between spend and revenue combined with broader AI ecosystem uncertainties. However, we are experiencing the rapid build out of a new industry. History has plenty of examples, such as the railroads and cloud computing, but there is no parallel for the pace and the dollar numbers associated with AI. Unsurprisingly, this makes some people nervous.

Regardless, the nature of the buildout is arguably necessary. Scale and compute are key drivers of success in this space. Pre-ChatGPT, Richard Sutton, a research scientist at Google DeepMind and professor of computing science at the University of Alberta, wrote in *"The Bitter Lesson"* (2019) that major advances in artificial intelligence consistently come from applying more compute, more data, and more scale, rather than focusing on increasingly sophisticated human-designed algorithms.[4] As such, competitive advantage should be attained by companies capable of deploying unprecedented amounts of capital into AI infrastructure.

What are the risks?

But even if the AI buildout ultimately generates the roughly $650 billion of annual recurring revenue that JPMorgan estimates is needed to produce a 10% return on the industry's cumulative capital investment, investors must still ask whether that is an attractive outcome. Cloud computing ultimately became one of the most profitable businesses in history, but it took more than a decade for cloud infrastructure companies to build the scale, customer adoption, and operating leverage necessary to consistently generate operating margins of 20%–30% or more. AI infrastructure spending is occurring on a much more compressed timeline and on a significantly larger scale. While the investment may ultimately prove justified, timing matters. It remains uncertain whether AI services can achieve cloud-like economics quickly enough to generate returns that are commensurate with the risk being taken, especially in the short term.

379 Thornall Street | 2nd Floor | Edison, NJ 08837
www.gittermanasset.com | Fax: 888-476-1314

In addition, there are concerns about commodification and low margins. In 2025, Jeff Currie argued that tech firms are "now producing a commodity called 'AI compute' priced in $/hour, much like electricity priced in $/MWh or oil in $/bbl," comparing the "'growth at all costs' capex cycle" of the shale boom with tech spending. He noted that, "Some argue that it makes sense for Big Tech AI to spend more than 100% of cash flow if a single generative AI model can achieve monopolistic profits; however, thus far AI looks more like a basic commodity industry." In conclusion, he states that, "[i]t's perfectly fine for AI to be a commodity, of course: just don't be surprised when industry multiples get priced like one."[5]

David Cahn made a similar point in 2024, stating, "GPU computing is increasingly turning into a commodity, metered per hour." He challenged the idea that railroads are an accurate analogy: "If you own the tracks between San Francisco and Los Angeles, you likely have some kind of monopolistic pricing power, because there can only be so many tracks laid between place A and place B. In the case of GPU data centers, there is much less pricing power."[6]

How are AI revenues currently looking?

Jeff Currie and David Cahn could be right, but if usage volumes are sufficiently high from enough demand, tech firms could generate the revenues to close the gap and justify valuations over the longer-term.

And, while the gap between investment and monetization is high, AI revenues are scaling approximately three times faster[7] than those earned under previous major technology waves (Internet, mobile apps, cloud). However, much of the revenues generated today are from companies responsible for infrastructure buildout and hardware alongside rewards for early enterprise adoption. We are not yet seeing the mature, recurring software economics that should ultimately make AI profitable for the companies incurring most of the spend. Still, this implies that historical analogies such as cloud computing or the internet may not be sufficiently accurate in terms of the opportunity as well as the pace of infrastructure investment.[8] That said, investments could result in an overbuild, whereby oversupply creates downward pricing pressure, at least in the shorter term.

379 Thornall Street | 2nd Floor | Edison, NJ 08837
www.gittermanasset.com | Fax: 888-476-1314

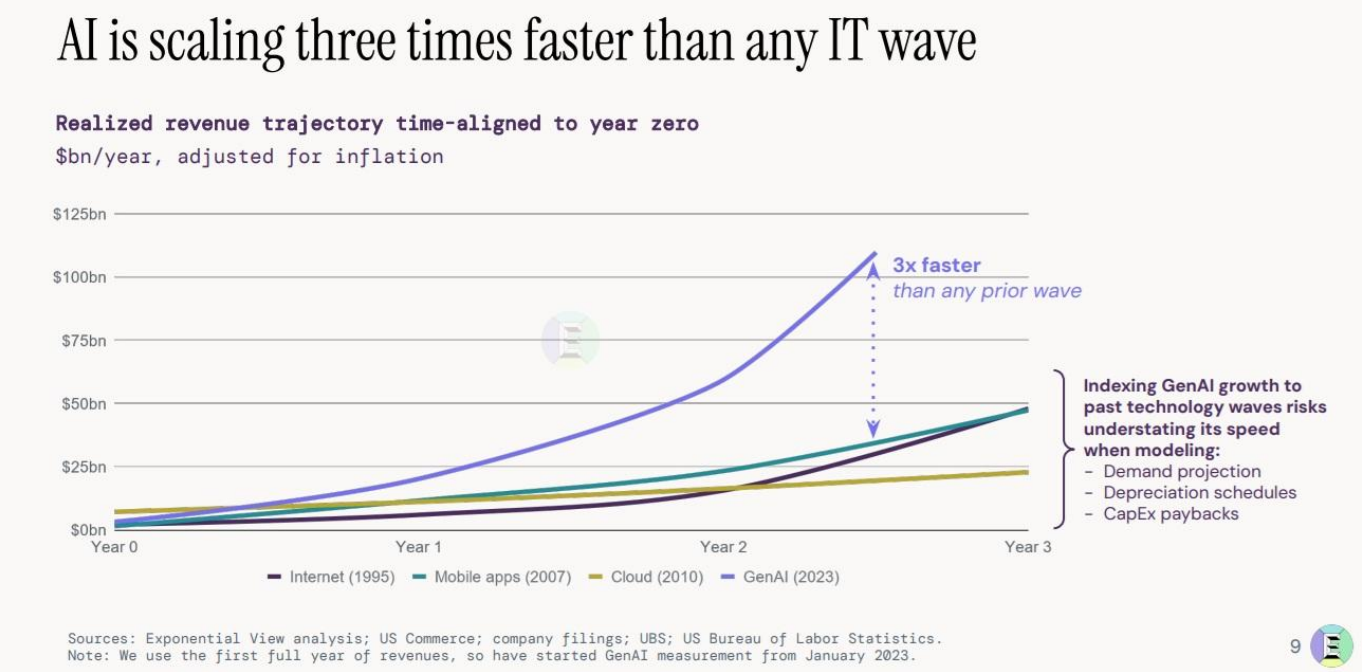

Source: Exponential View

At present, leading AI companies are continuing to produce exceptional earnings underpinning their valuations, with the potential for AI monetization across various products and services. Among private foundation model developers, OpenAI's annualized revenues as of February 2026 were $25 billion, up from $21.4 billion annualized at the end of 2025. Anthropic's annualized revenues were $9 billion at the end of 2025.[9] In May 2026, Anthropic reported that its revenue run-rate surpassed $47 billion, noting the increase in enterprise level adoption.[10]

Token revenue is a highly visible measure of direct AI monetization, but revenue increases are also available where AI is integrated into existing business lines. Among the hyperscalers, for example:

- Meta has used AI to better tailor the ads presented to users increasing advertising revenues by "a few billion dollars."[11]
- Alphabet, while not separately disclosing AI's contributions, is monetizing AI across its various products and services. Per the Q4 2025 earnings call, overall full-year revenues "exceeded $400 billion for the first time," with search-related revenues increasing by 17% for the quarter, adding $9 billion, and quarterly YouTube advertising revenues reaching $11.4 billion, an increase of 9%.[12]

That said, as shown below, hyperscaler valuations are down from recent highs and markets are skeptical about upcoming IPOs, especially in light of the recent SpaceX public offering.

379 Thornall Street | 2nd Floor | Edison, NJ 08837
www.gittermanasset.com | Fax: 888-476-1314

Company	Symbol	52-Week High	Price as of June 30, 2026	% Change
Alphabet	GOOGL	$402.62	$357.37	(-11.24%)
Amazon	AMZN	$274.99	$254.96	(-7.28%)
Meta	META	$738.31	$681.31	(-7.72%)
Microsoft	MSFT	$483.62	$373.02	(-22.87%)
Oracle	ORCL	$248.15	$146.55	(-40.94%)

Source: Bloomberg

Not all AI beneficiaries have participated equally, however, as reflected in Oracle's larger share price decline despite its continued AI infrastructure positioning.

Is long-term AI demand underpinned by broader economic ROI?

A current risk to shorter-term demand is that value creation for end users is far from guaranteed. A Bain & Company survey of corporate AI users found "that while 37% targeted cost reductions of 11% to 20%, nearly 40% of those who measured outcomes landed in the 0% to 10% bucket instead."[13] A lot of technology budgets are being gobbled up without generating the expected, and required, outcomes.

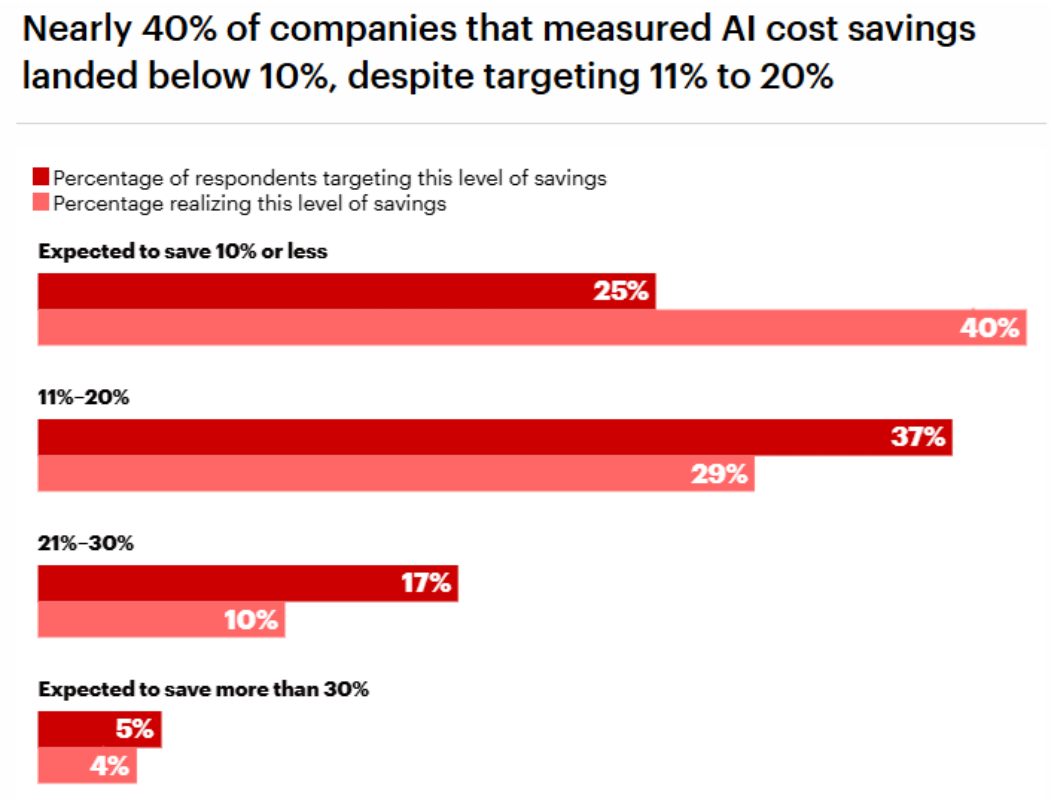

Source: Bain Automation and AI Pathfinder Survey 2026 (n=951)

379 Thornall Street | 2nd Floor | Edison, NJ 08837
www.gittermanasset.com | Fax: 888-476-1314

However, beyond the headline findings, the more interesting story is arguably about the organizations as opposed to AI's failings. The companies "capturing transformational savings have leaders who recognize that the human operating model is as important to redesign as the process itself. Typically, they have funded and sponsored that change explicitly at the CEO level, before deployment rather than after."[14] Therefore, there is no single story about whether AI works or not. There are instead stories about companies that have created the conditions for success and those that have not (yet).

Alongside the varying results in corporate AI investments, and despite the ongoing influence of Moore's Law,[15] the rising cost of AI deployment is leading to some demand destruction. Some companies have established boundaries on AI spend. For example, Uber has set an upper limit of $1,500 per employee per month.[16] Another unnamed company inadvertently spent $500 million in a single month because spending limits had not been established.[17]

But there is a conundrum here. Some companies are noticing that high spending might be driving value. As reported by TechCrunch, "Jellyfish, an engineering management platform... found engineers who used the most tokens were about twice as productive as those who used AI less, but they spent 10x the number of tokens to get there."[18] Similar to the Bain & Co findings, big spending carries risks but also the potential for outsize performance improvements if clear goals and internal capabilities exist.

What about earnings quality?

While hyperscalers are still posting strong earnings, current earnings quality and their repeatability are questionable.

Other income

Gains related to private investments are a meaningful contributor to the reported profitability of tech firms in recent quarters. For example, "Alphabet and Amazon generated "other income" totaling [sic] $53 billion in Q1 2026, which accounted for nearly 60% of those two companies' incomes in Q1 and 34% of the total $155 billion in income this quarter across the five largest hyperscalers." The source of this is the rising valuations of private investments including OpenAI and Anthropic.[19]

Income from unrealized gains is fundamentally different from recurring operating earnings generated by customers. In addition, the valuations used to book the gain are driven off

379 Thornall Street | 2nd Floor | Edison, NJ 08837
www.gittermanasset.com | Fax: 888-476-1314

investments made by Alphabet and Amazon themselves: “In plain English, the more they invest in Anthropic, the more profit they can report—without Anthropic ever having to pay them a dollar.”[20]

Risks may rise as OpenAI and Anthropic go to IPO, given their currently profitability or lack thereof. If the firms book losses after they go public, this could depress their valuations, impacting future financial reporting, should the hyperscaler investors continue to hold the shares.

Revenue circularity

Portions of the AI ecosystem are now increasingly circular. As noted above, companies building out AI such as OpenAI and Anthropic are receiving investment capital from hyperscalers while also being their customers. Nvidia also invests capital into AI companies that then use its GPUs, which are recorded as Nvidia revenue.

“Circularity can be a winner for all involved if things go well… When demand is rising and capital is readily available, the combination of investments and purchase commitments can act like a flywheel.”[21] However, from a risk perspective, such transactions can “skew the balance of incentives to encourage bad decision making. A company that has one of its suppliers as a major shareholder may be more likely to keep buying its stuff whether or not it makes commercial sense. That increases the risk of money being spent to secure business that fails to materialize. The circularity can be most risky when a handful of buyers are responsible for a large share of the market — as is the case with AI.”[22]

How is GAM thinking about this space?

In summary, we believe that, over time, the broader AI ecosystem can generate sufficient revenues to justify today's CAPEX investments, and we remain long-term bullish on the AI landscape. However, the investment case differs across participants. AI infrastructure providers, such as CoreWeave and Nebius, alongside picks-and-shovel providers, including NVIDIA, TSMC, AMD, SanDisk, and Broadcom, have already captured substantial economic value from mismatched supply and demand dynamics within the AI buildout.

The key question now lies with the hyperscalers, which continue to fund the majority of AI infrastructure investment. While demand for AI continues to grow, it remains to be seen whether they can translate record capital spending into sufficient recurring AI revenues to generate attractive long-term shareholder returns. However, in the same way that the spend has no true historical analog, the potential rewards could be unlike those of any previous technology cycle.

379 Thornall Street | 2nd Floor | Edison, NJ 08837
www.gittermanasset.com | Fax: 888-476-1314

Our base case is that AI revenues will close the gap over the next 5-8 years between infrastructure investment and earnings. While earnings quality and profitability may remain volatile during this transition, we believe that the flywheel effect is more likely to be realized than the downside risks from circular revenues. AI-related revenues have grown far faster than those generated under previous technology cycles. and monetization is increasingly spreading beyond infrastructure into enterprise software and AI applications.

Hyperscalers continue to expand capital investment, having the best visibility into customer demand. So, while risks remain, these developments suggest the ecosystem is evolving from one driven primarily by infrastructure spending toward one supported by a broader base of recurring end-user demand. There is a meaningful possibility that, in the race to establish AI leadership, the ecosystem builds more infrastructure than near-term demand requires. If so, periods of flat earnings, margin compression, or even losses should not be surprising. This would be consistent with previous infrastructure investment cycles but would not invalidate a long-term investment thesis. We believe that improvements in AI monetization will continue with demand continuing to expand as users derive more value. As such, AI revenues can grow to earn enough to exceed the illustrative 10% return hurdle discussed earlier, while creating the potential for substantially higher long-term returns as AI adoption matures. We believe any excess capacity that can create short-term price volatility will ultimately be absorbed over time.

We do not treat AI as a standalone thematic allocation and do not believe that investors should enter overweight positions at this time. However, neither should investors be underweight. As a manager of models with broad market exposure, we have a market weight allocation to U.S. large-cap equities. Given the concentration of AI leadership therein, our models have meaningful exposure to AI leaders.

Some of our core themes: Energy Addition and Climate Adaptation, have been reinforced by AI. Exploding electricity demand accelerates investment in power generation, grid technologies, and electrical equipment. Data centers increase demand for cooling technologies, water infrastructure, and resource efficiency. We believe it is important to allocate to a diversity of potential return drivers that are less dependent on the performance of a handful of companies.

In part two of this series, we will focus on the costs, competition, and supply chain considerations before stating our final conclusion.

GITTERMAN
ASSET MANAGEMENT
A PERIGON PARTNER

379 Thornall Street | 2nd Floor | Edison, NJ 08837
www.gittermanasset.com | Fax: 888-476-1314

Disclosures

This document does not constitute advice or a recommendation or offer to sell or a solicitation to deal in any security or financial product. It is provided for information purposes only and the understanding that the recipient has sufficient knowledge and experience to be able to understand and make their own evaluation of the proposals and services described herein, any risks associated therewith and any related legal, tax, accounting or other material considerations. To the extent that the reader has any questions regarding the applicability of any specific issue discussed above to their specific portfolio or situation, prospective investors are encouraged to contact Gitterman Asset Management or consult with the professional advisor of their choosing.

Certain information contained herein has been obtained from third party sources and such information has not been independently verified by Gitterman Asset Management. No representation, warranty, or undertaking, expressed or implied, is given to the accuracy or completeness of such information by Gitterman Asset Management or any other person. While such sources are believed to be reliable, Gitterman Asset Management does not assume any responsibility for the accuracy or completeness of such information. Gitterman Asset Management does not undertake any obligation to update the information contained herein as of any future date.

This presentation is confidential and is intended only for the person to whom it has been directly provided.

Certain information contained herein constitutes "forward-looking statements," which can be identified by the use of forward-looking terminology such as "may," "will," "should," "expect," "anticipate," "project," "estimate," "intend," "continue," or "believe," or the negatives thereof or other variations thereon or comparable terminology. Due to various risks and uncertainties, actual events, results, or actual performance may differ materially from those reflected or contemplated in such forward-looking statements. Nothing contained herein may be relied upon as a guarantee, promise, assurance or a representation as to the future.

Gitterman Asset Management is a dba of Perigon Wealth Management, LLC ("Perigon"), a registered investment advisor registered with the U.S. Securities and Exchange Commission.

---

[1] Defined primarily as Microsoft, Amazon, Alphabet, Meta. Also includes Oracle, Coreweave and Nebius where specifically stated.
[2] "How is AI being monetized?", Aliaga, Stephanie, January 2, 2026, J.P. Morgan Asset Management, https://am.jpmorgan.com/us/en/asset-management/adv/insights/market-insights/market-updates/on-the-minds-of-investors/how-is-ai-being-monetized/
[3] "The State of the AI Economy", June 25, 2026, Exponential View, https://intelligence.exponentialview.co/
[4] "The Bitter Lesson", Sutton, Rich, March 13, 2019, Available at: https://www.cs.utexas.edu/~eunsol/courses/data/bitter_lesson.pdf
[5] "How Much Does AI Rhyme with shAIe?", Currie, Jeff, November 25, 2025, Carlyle, https://www.carlyle.com/carlyle-compass/ai-compute-market-price-analysis
[6] "AI's $600B Question", Cahn, David, June 20, 2024, Sequoia Capital, https://sequoiacap.com/article/ais-600b-question/
[7] "The State of the AI Economy", June 25, 2026, Exponential View, https://intelligence.exponentialview.co/
[8] "The State of the AI Economy", June 25, 2026, Exponential View, https://intelligence.exponentialview.co/

379 Thornall Street | 2nd Floor | Edison, NJ 08837
www.gittermanasset.com | Fax: 888-476-1314

[9] "OpenAI tops $25 billion in annualized revenue, The Information reports", March 4, 2026, Reuters, https://www.reuters.com/technology/openai-tops-25-billion-annualized-revenue-last-month-information-reports-2026-03-05/
[10] "Anthropic raises $65B in Series H funding at $965B post-money valuation", May 28, 2026, Anthropic, https://www.anthropic.com/news/series-h
[11] "Power Overwhelming: dissecting the $1.5T AI revenue shortfall", Zhang, Kevin, December 4, 2025, LessWrong, https://www.lesswrong.com/posts/XZL9TqptvKsDtWHbi/power-overwhelming-dissecting-the-usd1-5t-ai-revenue
[12] "2025 Q4 Earnings Call", February 4, 2026, Alphabet, https://abc.xyz/investor/events/event-details/2026/2025-Q4-Earnings-Call-2026-Dr_C033hS6/default.aspx
[13] "Your AI Budget Is Growing. Your Returns Aren't. Here's Why", Heric, Michael et al, June 1, 2026, Bain & Company, https://www.bain.com/insights/your-ai-budget-is-growing-your-returns-arent-heres-why/
[14] "Your AI Budget Is Growing. Your Returns Aren't. Here's Why", Heric, Michael et al, June 1, 2026, Bain & Company, https://www.bain.com/insights/your-ai-budget-is-growing-your-returns-arent-heres-why/
[15] This is the concept that the magnitude of transistors on a semiconductor chip doubles annually, as predicted by Gordon Moore in 1965. See "Moore's law", (n.d.), Britannica, https://www.britannica.com/technology/Moores-law
[16] "Uber Caps Usage of AI Tools Like Claude Code to Manage Costs", Lung, Natalie, June 2, 2026, Bloomberg, https://www.bloomberg.com/news/articles/2026-06-02/uber-caps-usage-of-ai-tools-like-claude-code-to-cut-costs
[17] "One company spent half a billion dollars on Claude in a single month: Report comes as AI costs climb", Cramer, Jude, May 29, 2026, FastCompany, https://www.fastcompany.com/91550884/claude-ai-costs-climb-company-spent-half-a-billion-dollars-in-a-single-month-report
[18] "The token bill comes due: Inside the industry scramble to manage AI's runaway costs", Bellan, Rebecca, June 5, 2026, TechCrunch, https://techcrunch.com/2026/06/05/the-token-bill-comes-due-inside-the-industry-scramble-to-manage-ais-runaway-costs/
[19] "The mysterious $53bn 'other income' boost to AI hyperscaler earnings", Wigglesworth, Robin, May 7, 2026, Financial Times, https://www.ft.com/content/be97df0a-76b1-4cb0-9ba4-d1117d8d1450
[20] "Google and Amazon's biggest profit driver last quarter was their Anthropic stakes—which they haven't sold", Roytburg, Eva, April 30, 2026, Fortune, https://fortune.com/2026/04/30/google-amazon-ai-profits-anthropic-stake-bubble-earnings-2026/
[21] "A Guide to the Circular Deals Underpinning the AI Boom", Sam, Cedric et al, January 22, 2026 (Updated June 9, 2026), Bloomberg, https://www.bloomberg.com/graphics/2026-ai-circular-deals/
[22] "A Guide to the Circular Deals Underpinning the AI Boom", Sam, Cedric et al, January 22, 2026 (Updated June 9, 2026), Bloomberg, https://www.bloomberg.com/graphics/2026-ai-circular-deals/