



Dotcom on Steroids: Part II

| OpenAI: Still hallucinating?

21 Nov 2025

GQG Research

Key Takeaways

- > OpenAI's business appears commoditized and the model is not sustainable in the long run, in our view. It is a hyper capital-intensive company facing growing compute costs and significant competition from other players. We believe it will struggle to build a sustainable business over time
- > Open-source AI labs are producing models with efficacy comparable to that of closed-source offerings, such as those from OpenAI. This development is likely to lower switching costs for users and erode the stickiness and competitive advantage that closed-source groups like OpenAI have relied on
- > OpenAI is at the center of the AI infrastructure buildout—if capital markets begin to doubt that AGI is attainable within a reasonable timeframe, the company could face higher borrowing costs, or worse, a reduction in funding altogether. In our view, the spillover effects of such a scenario could be massive

In our previous GQG Research, [*Dotcom on Steroids*](#), we explored the striking parallels between the current AI landscape and the euphoric rise—and subsequent fall—of the late 1990s technology, media, and telecommunications bubble. In this follow-up piece, we delve into OpenAI, and take a hard look at the numbers, narratives, and risks separating the technological marvel from the economic reality.

| SCALING THE HEIGHTS, IGNORING THE CRACKS

On the surface, OpenAI may sound like a Mag 7 company in its early days—unprofitable, burning through cash, and a magnet to investors who think they have found the next pot of gold. Peel one layer or dare to bring out your calculator for a closer look, and OpenAI goes from being the world's "most valuable start-up" to what we believe is one of the most overvalued and overhyped companies in history.

In our view, the company's financials are broken and unrealistic even if you lower the standards and look at them through the lens of a start-up with potentially revolutionary technology. OpenAI is more capital intensive than any other start-up we have seen, lacking the stickiness, the moat, and the network effect that have paved the way for other tech success stories. In addition, the disconnect between OpenAI's revenue and valuation is alarming by any historical standard. As a comparison, consider Amazon or Google in their earlier days: when these companies had a \$500B valuation, their revenues were roughly 10x and 4x as large, respectively, in comparison to OpenAI's current run-rate revenue of ~\$20B.¹

As a long-only large cap manager investing in listed equities on behalf of its clients, most of the time we have the luxury of ignoring the less transparent world of private companies and some of their egregious valuations. So why is that different today? We are facing a unique situation where most of the large public players in tech are intertwined in a web of circular financing where money and deals flow through OpenAI. These relationships have direct implications for not just the large tech companies in our investable universe, but also many other companies in the S&P 500 in our view.


OpenAI is also playing a key role in driving the "hopes and dreams" factor, which we believe is helping to fuel this AI bubble. They believe Large Language Models (LLMs) will continue seeing meaningful step-function improvements toward this goal of Artificial General Intelligence (AGI), a notion a number of AI experts have started to pour cold water on.^{2,3} In fact, just recently we attended an AI conference in Miami, in which OpenAI's CFO Sarah Friar conveyed to a room filled with stakeholders, including investors and lenders, that the company's "true north" is "AGI for the benefit of humanity." In short, we see OpenAI as a capital-intensive business model that offers a product that faces meaningful risks of becoming a commodity. Yet, it continues to magically prop up valuations of its "partners" after each deal announcement.

EXPLOSIVE GROWTH, UNPRECEDENTED SPENDING

To give credit where it is due, OpenAI has firmly established itself as a leading innovator with its LLMs. Its market leadership is evidenced by rapid user adoption of their ChatGPT product which has swelled to over 800M weekly active participants and many large enterprise customers, in addition to revenue of billions of dollars, in the span of a few years.^{4,5}

OpenAI Revenue


Projections beyond 2024



Source: GQG Partners LLC (chart). TheInformation.com (data). The illustrations for years 2025 through 2030 include projections. Actual results may differ from any projections illustrated above.

Cash Crunch

OpenAI is now projecting much higher cash burn due to cloud computing and data center related expenses



Source: GQG Partners LLC (chart). TheInformation.com (data). 2024 are actual cash spend. The illustrations for years 2025 through 2030 include projections. Actual results may differ from any projections illustrated above. Free Cash Flow (FCF): a company's cash left over after operating expenses and capital expenditures.

However, success is rarely free.

This explosive growth has been financed by record-breaking capital infusions and expenditures to build the infrastructure to support it.⁶ This unprecedented level of spending is underwritten by market forecasts of what we view as extraordinary economic gains that may ultimately prove to be unsustainable. While these investments signal widespread enthusiasm, they obscure the structural challenges that threaten OpenAI's path to durable profitability. Despite the hype, the translation of this breakthrough technology into a resilient business is far from certain.

While we will not go into the details here, refer to the links at the end of this piece to read about what has been widely reported on both the circularity and domino effect OpenAI is having on other companies.^{7,8,9} Instead, we will delve into OpenAI's dynamics from an economic and technology perspective to illustrate why we believe it fails the proverbial "smell test" given that most analysts seem to be ignoring the facts.

THE MYTH OF THE METRIC

While impressive on the surface, high usage metrics (of which OpenAI has many to showcase) can often mask fundamental weaknesses in the quality of demand for a product and its ability to retain users. Specifically, a significant portion of OpenAI's current user engagement is largely driven by free-tier users using it for non-productive reasons.^{10,11} This type of usage inflates activity metrics used for raising copious amounts of capital, but we believe it holds little signaling power for future revenue, making it a poor proxy for sustainable customer value. Not to mention that 90% of ChatGPT users are outside of the United States, according to Friar. Furthermore, we believe the translation of enterprise use into durable and sticky revenue is undermined by several key factors:

Persistent Reliability Gaps: LLMs continue to struggle with hallucinations.¹² Their lack of consistency erodes trust and makes it difficult for the models to be embedded in mission-critical, multistep applications or operations that generate recurring, high-value revenue. The primary workaround, involving re-runs (increased token¹³ usage) and human-in-the-loop verification, directly counteracts the core promise of LLM efficiency and cost reduction, ultimately slowing adoption.¹⁴

Fragile Adoption and Low Switching Costs: The current user base is heavily weighted toward free-tier participants who have no financial commitment to the platform.¹¹ Even for paying enterprise customers, switching costs remain low for now. Currently, 28% of OpenAI's API usage flows through low-code platforms like Zapier, Bubble, and Retool.⁵ These companies are model agnostic and can easily redirect workflows to better or more cost-effective models as they become available, including open-source models.

Questionable Performance Benchmarks: Extreme growth expectations are often justified by referencing the improving model performance on AI benchmarks.¹⁵ However, this too gives us pause. Some studies are beginning to show how these models have “gamed” benchmark results by “memorizing” the questions rather than understanding the subject.¹⁶ Beyond this technicality, the benchmarks create what we believe is a conceptual mismatch between what they measure and what enterprises value. They reward a model's breadth of capability, making it an impressive generalist, yet corporations tend to be structured around the reliability of specialist employees performing a certain set of tasks with consistency. Because standard benchmarks do not adequately measure this consistency, a model's top-tier score in a benchmark test can be a misleading indicator of its real-world utility.

THE CHEAPER IT GETS, THE MORE YOU PAY

Even if user engagements were a perfect proxy for value, the basic economics of generative AI present a challenge. While the cost per token has plummeted due to hardware and software innovations, the cost per query has not seen the same dramatic decrease.¹⁷

With the mainstream adoption of reasoning models, the economics of LLM usage have shifted dramatically.¹⁸ Paying consumers are charged for each token the model takes as input and each token that the model outputs. It is important to note that users are charged more for tokens generated by the model than the tokens that are input. While the general cost per token has plummeted due to hardware enhancements, the number of tokens consumed per task has not.

An analysis by an AI gateway provider revealed an interesting trend: simple, single-turn queries, which constituted 80% of enterprise usage in early 2024, dropped to just 20% by year end. They were replaced by multi-step, complex workflows that drive up token consumption for each request, with most of them being the costlier output tokens.¹⁹


This shift is stark when comparing different model generations. Early models were often tasked with simple Q&A, where a query might consume a few hundred tokens. Today's reasoning models are used for complex, multi-turn tasks like analyzing a lengthy document, searching the web, or creating large pieces of working software, which increases the token count by a couple orders of magnitude.

This “token inflation” is further exacerbated by the models' tendency to be overly verbose and the operational necessity of re-running queries to suppress hallucinations and ensure output consistency. Ultimately, these factors combine to drive up the cost per answer, leaving users to question true efficiency gains.

Will generative AI get sufficiently better from here?

It is not unusual for companies to invoke the promise of AGI when pressed on whether today's use cases justify the massive amount of AI infrastructure investment companies are making. What should be unsettling is that the notion of near-term AGI is increasingly challenged by researchers and operators (including OpenAI founders themselves), with evidence of continued scaling and deployment realities pointing to a longer, harder road for advancement.²

Free vs Paid Consumers AI Use



Source: GQG Partners LLC (chart). Menlo Ventures- 2025: The State of Consumer AI (data). Data as of 26 June 2025.

But let us lower the standards a bit and put AGI to the side—can AI, as we know it today, meaningfully improve? Was GPT-5 materially better than its predecessor, and was its rate of improvement comparable to the step up from GPT-3 to GPT-4? A cursory glance at release cadence, benchmark deltas, sharply rising compute, and power requirements suggests step-function gains are becoming rarer and more expensive—an uncomfortable backdrop for valuations and hype that implicitly assume (need?) continued leaps.

Rising Costs, Shrinking Gains

Version	Release Date	Time Since Previous	Estimated Training Cost	Benchmark Improvement vs Prior
GPT-2	Feb 2019	~7 months	~\$0.05 M	~+100% (vs GPT-1)
GPT-3	June 2020	~16 months	~\$5 M	~+60%
GPT-4	Mar 2023	~33 months	~\$100 M	~+40%
GPT-5	Aug 2025	~29 months	~\$750 M	~+20%

Source: GQG Partners LLC (chart). Paul Kedrosky (data). Data as of 17 November 2025.

The uncomfortable truth is that this technology is running into hard constraints: the reservoir of high-quality human data is finite, returns to scale are slowing, and a recursive reliance on synthetic data risks degrading the signal. As a result, products like ChatGPT look a lot less like software with zero-marginal-cost scale and more like metered compute with rising variable inputs and operational choke points.


The illusion of scale

The unit economics of LLM providers do not seem to align with the high-margin, scalable models of successful Software-as-a-Service (SaaS) and marketplace companies. Unlike a SaaS model, the cost of adding a new LLM customer is not zero because the cost of compute for LLMs scales with the users, leading to an organic barrier to economies of scale that this business model can achieve. And unlike Uber, as an example, an LLM does not benefit from network effects in that the value of its service does not inherently improve as more users engage with it.

However, this flawed comparison appears to be driving a classic Venture Capital-subsidized “blitz-scaling” strategy that is now playing out globally. As shown in the graph, AI adoption has been unequal between geographies with varying GDP per capita. The lowest adoption rates can be seen in low GDP per capita countries, which tend to be more price sensitive, so providers seem to be using deeply subsidized pricing to capture users in these markets.

OpenAI, for instance, has introduced local plans in India, its biggest market after the US, at a fraction of the global price. Also in India, Perplexity, an AI-powered search engine, is bundling \$20 per month subscriptions for free in mobile plans where telecom providers generate as little as \$3 per month in revenue per user. It is difficult to rationalize how acquiring these highly price-sensitive users contributes to the future profitability that current valuations demand.

If growth within the consumer market relies on unsustainable subsidies, the hope for profitability must lie with enterprise adoption. This has been a comparatively bright spot, with enterprise spending on foundation models more than doubling in the first half of 2025. However, we believe this optimism is tempered by significant headwinds.



Source: GQG Partners LLC (chart). Empirical Research (data). “Internet population” uses 2023 estimates from the World Bank.

A few recent studies found that either companies remain in a perennial pilot phase for AI projects, or that most pilot programs for incorporating AI end up failing.^{20,21} Furthermore, even successful adoption may yield only marginal gains. A recent study estimated AI's net impact on enterprise profit margins by 2030 at a mere 50-70 basis points.²² While not immaterial, this margin uplift is expected to come primarily from job automation (which we think is becoming an increasingly questionable assumption²³) and offset by the significant depreciation expenses of the required AI infrastructure (which is a known and unavoidable cost).

The hard financials make this picture even bleaker. In the first half of 2025, OpenAI generated \$4.3B in revenue while posting a net loss of \$13.5B and burning through \$2.5B in cash. Despite this staggering burn rate, the company projects it will be profitable by 2030, with revenues soaring to \$200B and gross margins over 60%.²⁴ This type of long-range, hockey-stick forecast for a business with such a murky path to profitability is eerily reminiscent of the dotcom era, where many companies extrapolated short-term hype into extreme future earnings that never materialized.²⁵


For context, a \$200B revenue target would exceed the 2024 sales of Nvidia, a company that operates as a near-monopoly with immense barriers to entry and is years ahead of its closest competitors, something we believe OpenAI can hardly claim considering how easily interchangeable frontier language models are.²⁶ To suggest that OpenAI—a company burning through cash by the billions amid fierce competition with seemingly no durable moat—will achieve a similar financial profile in just a few years seems less like a forecast and more like a work of speculative fiction in our view.

This opacity is a private company's privilege, but it also raises a simple question: if the underlying economics were as strong as the hype, why not disclose them?

The reluctance to share a complete financial picture suggests the story of hyper growth may be masking a reality different from the hype. And when leadership can float a federal 'backstop' one day and swear off bailouts the next, it does not quite read like conviction. In our opinion it reads like pre-wiring a safety net while pretending not to need one.^{27,28} We found Sam Altman's defensive reaction when pressed by Brad Gerstner of Altimeter Capital on the math behind the planned \$1.4 trillion spend to be similarly telling.²⁹ And prior court depositions and board accounts corroborate the concern about there being a pattern here, not to mention OpenAI's previous board of directors fired Altman in November 2023 for not being "consistently candid in his communications" with the board!^{30,31}

Large-Capitalization Stocks

Estimated Impact of AI on Net Profit Margins 2025E through 2030E



Source: GQG Partners LLC (chart). Empirical Research Partners Analysis (data). Actual results may differ from any projections illustrated above. Data as of 30 September 2025.

We believe there is another underappreciated risk: the intense competition for talent among the AI labs. OpenAI is on track to spend close to \$6B on stock-based compensation in 2025 alone.³² Since these are all based on the company's ~\$500B valuation, even a modest slowdown could heighten the challenge of retaining key employees. However, despite the possible headwinds, OpenAI continues to raise its revenue expectations.

This all leads back to one unavoidable consequence: the math seemingly does not work, and the capital-intensive nature of the business creates relentless liquidity pressure. Despite record-breaking revenue, OpenAI has raised over \$40B this year, up from \$6.6B in 2024, all in addition to revolving credit lines and infrastructure partnerships. To keep up their growth, arguably by subsidizing the cost of compute for users and maintaining their high-in-demand workforce, we think they will need to keep raising money at record-breaking rates.

Using a quick back-of-the-envelope calculation, taking the 136M US private workforce at the going subscription rate of \$20 per user per month yields less than \$3B in revenues—and of course that is before discounts, usage caps, or reseller splits. To put this into perspective, Google raised just \$26M (~\$60M in today's terms) before becoming profitable, while Meta raised about \$480M (~\$800M today).^{33,34} By comparison, OpenAI is projected to burn through a staggering \$115B before it even hopes to reach free cash flow profitability in 2030.³⁵ The scale of this capital consumption is without precedent.


Not so open to other AIs

This constant need for capital exposes a deeper vulnerability: the absence of a durable competitive moat. In a head-to-head matchup against other well-funded labs like Google and Anthropic, any performance advantage is fleeting. Minute differences between top models and low switching costs prevent the formation of a durable price premium or sticky user base, as evidenced by Anthropic's recent capture of a major share of the software coding market.³⁶

This vulnerability is magnified by the disruptive and often understated force of the open-source community. While closed-source models have historically maintained a performance lead, that gap is closing. Notably, China's open-source Kimi K2 Thinking model has now demonstrated performance results that match or exceed the leading closed-sourced reasoning systems from the likes of OpenAI and Anthropic, underscoring just how quickly the frontier is shifting.³⁷


As the graph shows, this is far from anomalous: open-source models consistently catch up to (and now seemingly surpass) the state-of-the-art performance of their closed-source counterparts.

Rising Revenue



Source: GQG Partners LLC (chart). TheInformation.com (data). Data as of 5 September 2025. 2023 and 2024 are actual revenue. The illustrations for years 2025 through 2030 include projections. Actual results may differ from any projections illustrated above.

Closed-Source vs Open-Source Models



Source: GQG Partners LLC (chart). LMArena (data). Data as of 7 November 2025.

Furthermore, we have seen that open-source models optimize resource utilization that ultimately improves efficiency and lowers the cost of inference resulting in even more adoption and innovation.

This dynamic mirrors the rise of Linux, which, despite being perpetually behind the polished proprietary systems of its day, steadily grew a loyal following. That following eventually built the foundation for Android, an open-source project that now dominates the mobile operating system market. **The lesson is clear: pricing power is not always tied to having the absolute state-of-the-art model. A “good enough” alternative can become the market standard just because it is accessible to more people.**

Ultimately, we believe this will democratize access and distribute the costs of development to more companies and builders, making it increasingly difficult for closed-source players to obtain and keep market share at sufficiently attractive economics. We believe this will drive the industry to become more like a marketplace, aiding in more innovation as more builders and consumers come together with aligned incentives. This organic growth tailwind will be difficult to replicate with the centralized, capital-intensive nature of the closed-source AI labs. Moreover, closed-source models face challenges with the way established SaaS incumbents like Microsoft and Salesforce are positioned to extract value for their users. Their value proposition is not the raw capability of the AI model, but the productivity increase delivered through an integrated workflow that they already serve to their clients. They can treat the increasingly commoditized AI models as an interchangeable component, while capturing the lion’s share of the value creation.

This leaves the foundational model providers like OpenAI bearing the immense and perpetual cost of R&D while the incumbents capture the lion’s share of the profits, relegating the innovators to the position of a low-margin utility supplier. This is not dissimilar to what we saw (and indeed recently wrote about in [Dotcom on Steroids](#)) during the TMT bubble of the late 1990s where the big spenders generally did not end up capturing the lion’s share of the economics which that infrastructure enabled.³⁸

IMMORTAL TECH, MORTAL COMPANY

As Howard Marks once said, “There are no bad assets, only bad prices.” As the saying goes, LLMs represent a genuine technological innovation with the potential to create immense value for humanity. However, as an investment proposition, we think the sticker price to participate in them is untenable given the uncertain path to translating the innovation into a resilient, capital-efficient business. The current market, fueled by a venture capital culture that has long prioritized growth over fiscal discipline, appears to be conflating headline usage, benchmark victories, and run-rate revenues with durable economics.

The current exuberance is particularly jarring because participants seem to be completely discounting signs that the exponential improvements driven by scaling laws may be slowing. Research from major tech labs has highlighted the inherent limitations of LLMs like the ‘illusion of thinking’ paper from Apple.³⁹ Other studies have also cast doubt on whether simply scaling current architectures can lead to true AGI, the very goal that was supposed to justify this level of investment.⁴⁰ This is a critical point: **while LLMs are impressive, in our view, it is still just a technology with its own inherent constraints**, not a magic pill to solve the world’s problems.

This exuberance also seems to overlook the fundamental fragility of the current business model. OpenAI’s explosive revenue growth, on track to have delivered ~\$13B by the end of 2025, has been matched by an equally explosive burn rate, requiring massive and repeated infusions of capital just to sustain operations.⁴¹

Meanwhile its \$500B valuation rests on the assumption that this growth can be sustained and eventually converted into profit.⁴² This premise appears to be increasingly questionable given the structural headwinds and fierce competition from both closed- and open-source AI labs alike.

As investors, our role is to allocate capital efficiently, which requires separating the technological marvel from its economic impact. The current approach of funneling capital into a few leaders in a winner-take-all race overlooks the risk of creating entities so large that their potential failure becomes a systemic threat with second-degree effects on the wider economy, in our view.

We believe the real test will come when the current wave of enthusiasm and subsidies fades. As capital tightens and competition accelerates, the market will be forced to distinguish between speculative and sustainable enterprises. LLMs are not a bad asset, but the price being paid for a business with such immense capital costs, no obviously durable moat, and questionable unit economics may prove to be fatal. We believe the story to watch is not whether the technology is immortal (it most likely is), but whether the companies building it are.

Recommended Reading

> [How OpenAI Uses Complex and Circular Deals to Fuel Its Multibillion-Dollar Rise – The New York Times](#)

Sam Altman, the chief executive of OpenAI, says that technological revolutions are driven by more than just technology. They are also driven, he argues, by new ways of paying for them.

> [Why Sam Altman Won't Be On The Hook For OpenAI's Massive Spending Spree](#)

OpenAI has gone wild with compute deals this year, committing to spend far more than its balance sheet can currently sustain. So who takes the fall if it can't pay? It won't be Altman.

> [How high are OpenAI's compute costs? Possibly a lot higher than we thought](#)

Tech blogger Ed Zitron has an interesting post about OpenAI's cash burn. The gist is that OpenAI's running costs may be a lot more than previously thought, and that its main backer Microsoft is doing very nicely out of their revenue share agreement.

> [Oracle is already underwater on its 'astonishing' \\$300bn OpenAI deal](#)

It's too soon to be talking about the Curse of OpenAI, but we're going to anyway. Since September 10, when Oracle announced a \$300bn deal with the chatbot maker, its stock has shed \$315bn* in market value. OK, yes, it's a gross simplification to just look at market cap. But equivalents to Oracle shares are little changed over the same period (Nasdaq Composite, Microsoft, Dow Jones US Software Index), so the \$60bn loss figure is not entirely wrong. Oracle's "astonishing quarter" really has cost it nearly as much as one General Motors, or two Kraft Heinz.

> [When Will the AI Bubble Burst?](#)

Professor Gary Marcus (Professor Emeritus of Psychology and Neural Science, New York University) and Murad Hemmadi (Journalist, The Logic) at Attention: Govern Or Be Governed, a two-day international gathering, dedicated to understanding how the world's democracies can chart a new path forward in the digital age.

> [OpenAI, Nvidia Fuel \\$1 Trillion AI Market With Web of Circular Deals](#)

A wave of deals and partnerships are escalating concerns that the trillion-dollar AI boom is being propped up by interconnected business transactions.

> [At a major AI conference, one startup got voted most likely to flop](#)

In a city obsessed with betting on the next big thing, the Cerebral Valley AI Conference turned that instinct inward. Before panels featuring AI heavyweights like Anthropic and xAI wrapped up on Wednesday, founders and investors in the audience were asked a question that Silicon Valley rarely says out loud: Which billion-dollar AI startup would you bet against?

> [Who is OpenAI's auditor? \(Update: it's Deloitte\)](#)

OpenAI is committed to spending about \$1.4 trillion on data centres over the next decade or so. It accounts for about two-thirds of unfulfilled contracts at Oracle, which is valued at about \$630bn, and two-fifths of unfulfilled contracts at CoreWeave, which is valued around \$36bn. It has around \$375bn of unfulfilled contracts with Microsoft, the world's third-most-valuable company.

> [Andrej Karpathy — AGI is still a decade away](#)

Andrej explains why reinforcement learning is terrible (but everything else is much worse), why model collapse prevents LLMs from learning the way humans do, why AGI will just blend into the previous ~2.5 centuries of 2% GDP growth, why self driving took so long to crack, and what he sees as the future of education.

END NOTES

- ¹Capoot, Ashley. "Sam Altman says OpenAI will top \$20 billion in annualized revenue this year, hundreds of billions by 2030". CNBC. 6 November 2025.
- ²Roytburg, Eva. "Did an OpenAI cofounder just pop the AI bubble? The models are not there". Fortune. 21 October 2025.
- ³Patel, Dwarkesh. "Richard Sutton – Father of RL thinks LLMs are a dead end". Dwarkesh Podcast. 26 September 2025.
- ⁴Bellan, Rebecca. "Sam Altman says ChatGPT has hit 800M weekly active users". TechCrunch. 6 October 2025.
- ⁵Elad, Barry. "OpenAI Statistics 2025: Adoption, Integration & Innovation". SQ Magazine. 7 October 2025.
- ⁶Tap Twice Digital Team. "8 OpenAI Statistics (2025): Revenue, Valuation, Profit, Funding". Tap Twice Digital. 18 May 2025.
- ⁷Weil, Jonathan. "Microsoft Needs to Open Up More About Its OpenAI Dealings". The Wall Street Journal. 27 October 2025.
- ⁸"Nvidia's \$100bn bet on OpenAI raises plenty of questions". The Economist. 25 September 2025.
- ⁹Waters, Richard. "How OpenAI put itself at the centre of a \$1tn network of deals". The Financial Times. 10 October 2025.
- ¹⁰"OpenAI's Shocking Revelation: 70 % of ChatGPT Use Isn't for Work". Open Tools.ai. 22 September 2025.
- ¹¹Carolan, Shawn, et al. "2025: The state of Consumer AI". Menlo Ventures. 26 June 2025
- ¹²Dilmegani, Cem and Daldal, Aleya. "AI Hallucination: Comparison of the Popular LLMs". AIMultiple. 20 October 2025.
- ¹³Microsoft.com. "Understanding Tokens". A token is a unit of text, like a word, part of a word, or character, that an AI language model processes to understand and generate language.
- ¹⁴Todorov, Ventsi. "Human-in-the-Loop For LLM Accuracy". NineTwoThree. 4 June 2024.
- ¹⁵"AI Benchmarking". epoch.ai.
- ¹⁶"The Sequence Opinion #485: What's Wrong With AI Benchmarks". TheSequence. 6 February 2025
- ¹⁷Spark, Zoe. "The LLM Cost Paradox: How 'Cheaper' AI Models Are breaking Budgets". IKANGAI. 21 August 2025.
- ¹⁸IBM.com. "What is a reasoning model?". Reasoning models are a new type of LLMs designed to solve complex problems by performing structured, multi-step thinking before providing an answer. They typically use more tokens while processing their answers.
- ¹⁹Sambharia, Siddharth. "LLMs in Prod 2025: Insights from 2 Trillion+ Tokens". Portkey. 21 January 2025.
- ²⁰Singla, Alex, et al. "The state of AI in 2025: Agents, innovation, and transformation". QuantumBlack AI by McKinsey. 5 November 2025.
- ²¹Nanda, Mit, et al. "The GenAI Divide. The State of AI in Business in 2025". mlq.ai. July 2025.
- ²²Empirical Research Partners. "AI and Margins: A New, New Economy?". September 2025.
- ²³Kinder, Molly, et al. "New data show no AI jobs apocalypse—for now". Brookings. 1 October 2025.
- ²⁴Muppidi, Sri. OpenAI Forecasts Revenue Topping \$125 Billion in 2029 as Agents, New Products Gain". The Information. 23 April 2025.
- ²⁵Smriti. "Pets.com failure: Learn from one of the biggest innovation failures". InspireIP. 4 October 2023.
- ²⁶King, Ian. "Why Is Nvidia the King of AI Chips, and Can It Last?". Bloomberg. 5 November 2025.
- ²⁷Ghaffary, Shirin. "Altman Says OpenAI Doesn't Want a Government Bailout for AI". Bloomberg. 6 November 2025.
- ²⁸Capoot, Ashley. "OpenAI CFO Sarah Friar says company isn't seeking government backstop, clarifying prior comment". CNBC. 6 November 2025.
- ²⁹Tangermann, Vicotr. "Sam Altman Loses His Cool When Asked About OpenAI's Minuscule Revenue". Futurism. 3 November 2025.
- ³⁰Lanz, Jose Antonio. "Inside the Deposition That Showed How OpenAI Nearly Destroyed Itself". Decrypt. 4 November 2025.
- ³¹Milmo, Dan and Montgomery, Blake. "OpenAI staff threaten to quit en masse unless Sam Altman is reinstated". The Guardian. 20 November 2023.
- ³²Palazzolo, Stephanie, et al. "OpenAI's First Half Results: \$4.3 Billion in Sales, \$2.5 Billion Cash Burn". The Information. 20 September 2025.
- ³³"How Much Did Google Raise? Funding & Key Investors". Clay. 9 May 2025.
- ³⁴"How Much Did Meta Raise? Funding & Key Investors". Clay. 24 March 2025.
- ³⁵"OpenAI expects business to burn \$115 billion through 2029, The Information reports". CNBC. 6 September 2025.
- ³⁶Govil, Nidhi. "Anthropic Surpasses OpenAI in Enterprise AI Market, Dominates Coding Sector". TheOutpost.ai. 21 July 2025.
- ³⁷Franzen, Carl. "Moonshot's Kimi K2 Thinking emerges as leading open source AI, outperforming GPT-5, Claude Sonnet 4.5 on key benchmarks". VentureBeat. 6 November 2025.
- ³⁸"Dotcom on Steroids". ggg.com. 11 September 2025.
- ³⁹Shojaee, Parshin, et al. "The Illusion of Thinking: Understanding Strengths and Limitations of Reasoning Models via the Lens of problem Complexity". Apple Machine Learning Research. June 2025.
- ⁴⁰Hannas, William, et al. "Chinese critiques of Large Language Models". CSET. January 2025.
- ⁴¹Palazzolo, Stephanie. "OpenAI Generated \$4.3 Billion in Revenue, Burned \$2.5 Billion in First Half". The Information. 13 November 2025.
- ⁴²Sigalos, MacKenzie and Capoot, Ashley. "OpenAI boosts size of secondary share sale to \$10.3 billion". CNBC. 3 September 2025.

DEFINITIONS

AI benchmarks: standardized tests that evaluate and compare the performance of artificial intelligence systems, using a specific dataset or set of prompts and a scoring method to measure performance on a particular task.

Capital-intensive: a business or industry requires a large investment in physical assets like machinery, equipment, and buildings to produce goods or services.

GDP per capita: a country's gross domestic product (GDP) divided by its population.

Hallucinations: When a large language model perceives patterns or objects that are nonexistent, creating nonsensical or inaccurate outputs.

Cost of inference: the total expense of using a trained AI model to generate an output, including resources like computing power, energy, storage, and network usage.

Closed Source vs Open Source: Open-weight models provide access to the model's parameters (weights), allowing for customization, while closed-weight models are proprietary and accessible only via an API. Key differences include customization and control, security and privacy, and innovation and flexibility. Open-weight models offer greater flexibility and security for sensitive data, while closed-weight models provide convenience and potentially higher initial performance through a managed API.

IMPORTANT INFORMATION

This document may be distributed by GQG Partners LLC and its affiliates (collectively "GQG").

The information provided in this document does not constitute investment advice and no investment decision should be made based on it. Neither the information contained in this document or in any accompanying oral presentation is a recommendation to follow any strategy or allocation. In addition, neither is it a recommendation, offer or solicitation to (i) sell or buy any security, (ii) purchase shares in any investment fund that GQG may sponsor, offer or manage, (iii) establish any separately managed account, or (iv) implement any investment advice. It should not be assumed that any investments made or recommended by GQG in the future will be profitable or will equal the performance of any securities discussed herein. Before making any investment decision, you should seek expert, professional advice, including tax advice, and obtain information regarding the legal, fiscal, regulatory and foreign currency requirements for any investment according to the law of your home country, place of residence or current abode.

This document reflects the views of GQG as at a particular time. GQG's views may change without notice. Any forward-looking statements or forecasts are based on assumptions and actual results may vary.

GQG provides this information for informational purposes only. GQG has gathered the information in good faith from sources it believes to be reliable, including its own resources and third parties. However, GQG does not represent or warrant that any information, including, without limitation, any past performance results and any third-party information provided, is accurate, reliable, or complete, and it should not be relied upon as such. GQG has not independently verified any information used or presented that is derived from third parties, which is subject to change. Information on holdings, allocations, and other characteristics is for illustrative purposes only and may not be representative of current or future investments or allocations.

The information contained in this document is unaudited. It is published for the assistance of recipients, but is not to be relied upon as authoritative and is not to be substituted for the exercise of one's own judgment. GQG is not required to update the information contained in these materials, unless otherwise required by applicable law.

The contents of this document are confidential and intended solely for the recipient. No portion of this document and/or its attachments may be reproduced, quoted or distributed without the prior written consent of GQG.

GQG is registered as an investment adviser with the U.S. Securities and Exchange Commission. Please see GQG's Form ADV Part 2, which is available upon request, for more information.

Any account or fund advised by GQG involves significant risks and is appropriate only for those persons who can bear the economic risk of the complete loss of their investment. There is no assurance that any account or fund will achieve its investment objectives. Accounts and funds are subject to price volatility and the value of a portfolio will change as the prices of investments go up or down. Before investing in a strategy, you should consider the risks of the strategy as well as whether the strategy is appropriate based upon your investment objectives and risk tolerance.

There may be additional risks associated with international and emerging markets investing involving foreign, economic, political, monetary, and/or legal factors. International investing is not for everyone. You can lose money by investing in securities.

Past performance may not be indicative of future results. Performance may vary substantially from year to year or even from month to month. The value of investments can go down as well as up. Future performance may be lower or higher than the performance presented, and may include the possibility of loss of principal. It should not be assumed that recommendations made in the future will be profitable or will equal the performance of securities listed herein.

GQG Partners LLC is a wholly owned subsidiary of GQG Partners Inc., a Delaware corporation that is listed on the Australian Securities Exchange (ASX: GQG). GQG Partners LLC and its affiliates provide certain services to each other.

NOTICE TO RECIPIENTS IN AUSTRALIA & NEW ZEALAND

The information in this web page is issued and approved by GQG Partners LLC ("GQG"), a limited liability company and authorised representative of GQG Partners (Australia) Pty Ltd, ACN 626 132 572, AFSL number 515673. This information and our services may be provided to wholesale and retail clients (as defined in section 761G of the Corporations Act 2001 (Cth)) domiciled in Australia. This web page contains general information only, does not contain any personal advice and does not take into account any prospective investor's objectives, financial situation or needs. In New Zealand, any offer of a Fund is limited to 'wholesale investors' within the meaning of clause 3(2) of Schedule 1 of the Financial Markets Conduct Act 2013. This information is not intended to be distributed or passed on, directly or indirectly, to any other person.

NOTICE TO RECIPIENTS IN CANADA (Alberta, British Columbia, Manitoba, New Brunswick, Nova Scotia, Ontario, Québec, Saskatchewan (the "Canadian Jurisdictions"))

GQG Partners LLC relies on the (i) international adviser exemption pursuant to section 8.26 of NI 31-103 in each of the Canadian Jurisdictions, and (ii) non-resident investment fund manager exemption pursuant to section 4 of MI 32-102 in Ontario and Québec and is not registered as an adviser or investment fund manager in the Canadian Jurisdictions.

This web page has been prepared solely for information purposes and is not an offering memorandum or any other kind of an offer to buy or sell or a solicitation of an offer to buy or sell any security, instrument or investment product or to participate in any particular trading strategy. It is not intended and should not be taken as any form of advertising, recommendation, investment advice or invitation to trade. This information is confidential and for the use of the intended recipient only. The distribution of this web page in Canada is restricted to recipients who are qualified "permitted clients" for purposes of National Instrument 31-103 Registration Requirements, Exemptions and Ongoing Registrant Obligations. This information may not be reproduced, redistributed or copied in whole or in part for any purpose without the prior written consent of GQG. Upon receipt of this information, each Canadian recipient hereby confirms having expressly requested that all information evidencing or relating in any way to the information described herein be drawn up in the English language only. Par la réception de ce document, le détenteur au Canada de celui-ci confirme par les présentes avoir expressément exigé que tous les documents faisant foi ou se rapportant de quelque manière que ce soit aux informations présentées dans ce document soient rédigés en anglais seulement.

NOTICE TO RECIPIENTS IN SOUTH AFRICA

GQG PARTNERS LLC is an authorised financial services provider in the Republic of South Africa and regulated by the Financial Sector Conduct Authority (FSCA) with FSP number: 48881. Investors should take cognisance of the fact that there are always risks involved when buying or selling any financial product. Past performance of a financial product is not necessarily indicative of future performance. The value of financial products can increase as well as decrease over time, depending on the value of the underlying securities and prevailing market conditions. The investment value of a financial product is not guaranteed, and any illustrations, forecasts or hypothetical data are provided for illustrative purposes only. This web page does not constitute financial advice, a solicitation, invitation or investment recommendation. Prior to selecting a financial product or investment, it is recommended that South African based investors seek and obtain specialised financial, legal and tax advice.

NOTICE TO RECIPIENTS IN THE UNITED KINGDOM

GQG Partners LLC is not an authorised person for the purposes of the Financial Services and Markets Act 2000 of the United Kingdom ("FSMA") and the distribution of this web page in the United Kingdom is restricted by law. Accordingly, this web page is provided only for and is directed only at persons in the United Kingdom reasonably believed to be of a kind to whom such promotions may be communicated by a person who is not an authorised person under FSMA pursuant to the FSMA (Financial Promotion) Order 2005 (the "FPO"). Such persons include: (a) persons having professional experience in matters relating to investments; and (b) high net worth bodies corporate, partnerships, unincorporated associations, trusts, etc. falling within Article 49 of the FPO. The services provided by GQG Partners LLC and the investment opportunities described in this web page are available only to such persons, and persons of any other description may not rely on the information in it. All, or most, of the rules made under the FSMA for the protection of retail clients will not apply, and compensation under the United Kingdom Financial Services Compensation Scheme will not be available.

GQG Partners (UK) Ltd. is a company registered in England and Wales, registered number 1175684. GQG Partners (UK) Ltd. is an appointed representative of Sapia Partners LLP, which is authorised and regulated by the Financial Conduct Authority ("FCA") (550103).

NOTICE TO RECIPIENTS IN ADGM

GQG Partners Ltd, a company limited by shares, registered in Abu Dhabi Global Markets ("ADGM"), having its address at Unit No. 1 and 2, Floor 14, Al Maryah Tower, Abu Dhabi Global Market Square, Abu Dhabi, Al Maryah Island, United Arab Emirates. GQG Partners Ltd is licensed by the ADGM's Financial Services Regulatory Authority (FSRA) (license number 240015). GQG Partners Limited is licensed by the ADGM's Financial Services Regulatory Authority (FSRA) to conduct the regulated activities of Managing a Collective Investment Fund, Advising on Investments or Credit, Arranging Deals in Investments, Managing Assets, Shari'a-compliant Regulated Activities. This web page is intended for distribution only to persons of a type specified in the FSRA's Rules (i.e., "Professional Clients") and must not be delivered to or relied on by any other type of person. It is for the exclusive use of the persons to whom it is addressed and in connection with the subject matter contained therein. The FSRA, or any other regulatory authority, has no responsibility for reviewing or verifying this web page or any other information in connection with it. Accordingly, the FSRA, or any other regulatory authority, neither approved this information or any other associated information nor taken any steps to verify the information set out in this web page and has no responsibility for it.

© 2025 GQG Partners LLC. All rights reserved. This document reflects the views of GQG as of November 2025.

TLOPNAI 1125