

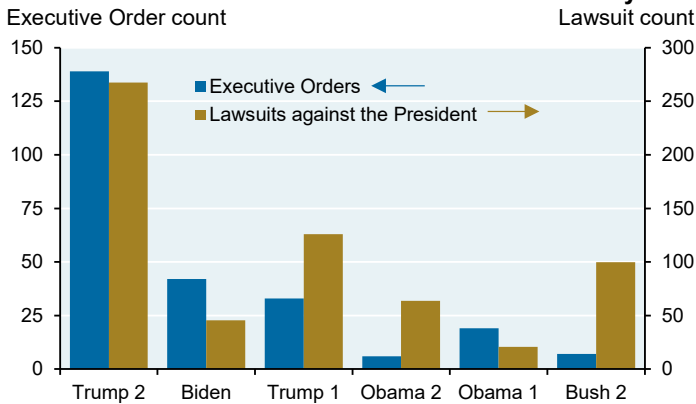
Back to our Regularly Scheduled Programming: an update on AI capabilities, corporate AI adoption, hyperscaler AI revenues vs spending and an unpleasant instance of “non-alignment”

I’ve focused a lot this year on connections between politics, economics and markets. The reasons are clear: the flurry of 20+ executive orders, memoranda and proclamations on tariffs which were catalysts for the first “Sell America” episode since 1982 (a material and simultaneous decline in the S&P, the dollar and Treasury bonds combined with US equity underperformance vs the rest of the world). That deserves attention, particularly since the US is more reliant than ever on foreign capital as illustrated in the third chart. Like Blanche Dubois in *A Streetcar Named Desire*, the US relies a lot on the kindness (and capital) of strangers¹.

Let’s assume the 90-day pause holds and that China will be subject to the same 10% reciprocal tariff as the rest of the world. The bilateral US tariff rate on China would rise from ~11% at the start of the year to ~40% after incorporating reciprocal tariffs, the fentanyl tariff, possible product-specific tariffs of 25%, announced product exemptions and pre-existing tariffs. Assuming no import substitution away from China, the US would still be applying the highest US tariff rate since the 1940’s. Not as high as it was 2 months ago, but still pretty high.


Now that the tariff outlook may be approaching an equilibrium state, it’s time to return to some regularly scheduled programming: an update on AI applications which were the primary drivers of US equity markets before the trade wars began. US companies spent a lot more time talking about AI than tariffs on Q1 earnings calls, and according to Empirical Research, the market capitalization of AI plays is 2.6x larger than their “tariff victim” category². In other words, one could make the argument that AI is at least as import as tariffs to equity investors, if not moreso. At the same time, the P/E premium vs the overall market for a basket of 48 AI-related stocks is back down at levels last seen in 2017 (fourth chart).

Executive Orders and Federal lawsuits in first 100 days




Source: New York Times, Federal Register, JPMAM, May 1, 2025

Average tariff rate on all US imports assuming US-China paused rates are maintained, Percent




Source: Tax Foundation, JPM Global Economics, GS Global Investment Research, JPMAM, May 12, 2025

US net savings



Source: FRB, BEA, JPMAM, September 2024

Relative forward P/E of AI plays



Source: Bloomberg, Empirical Research, JPMAM, May 12, 2025

¹ Another way to compute it: over the last decade, US net portfolio investment liabilities to foreign countries rose from \$7 trillion to \$17 trillion, while US net direct investment liabilities rose from zero to \$7 trillion (BEA)


² “Earnings Season Wrap”, Empirical Research Partners, Rochester Cahan, May 9, 2025

There’s a lot of discussion on how expensive US equities are, and if you’re just looking at relative P/E multiples vs Europe or Japan, that’s what you would probably conclude. But equities can be cheap or expensive for a reason, and I often remind investors that US companies are generally more profitable than non-US counterparts. The chart on the left shows price-to-book ratios as a function of projected return on equity by sector, for the S&P 500 and developed world equities excluding the US. Across practically every sector, US companies are more profitable; and therefore some degree of valuation premium for the US can be easily justified. As shown below, the US tech sector leads the pack on profitability and valuation. An astonishing sign of the success of tech & interactive media stocks: they now account for 35% of market-wide earnings vs 19% a decade ago.

AI has been the dominant tech theme of the last two years. AI capabilities have improved across multiple domains (see chart on the right), while at the same time their costs have declined. As noted in Stanford’s annual AI report: driven by increasingly capable small models, the inference cost for a system performing at the level of GPT-3.5 dropped by over 280x between November 2022 and October 2024. At the hardware level, costs declined by 30% annually while energy efficiency improved by 40% each year. Open-weight models are closing the gap with closed models, sharply reducing performance differentials on some benchmarks to just 2%. In combination, these trends are rapidly lowering the barriers to advanced AI³.

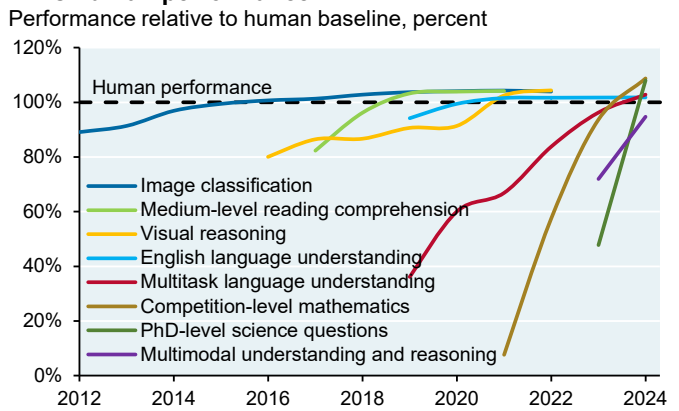
In this piece, we update our latest findings on AI. **From most visible to least visible:** the increase in hyperscaler capital spending, the improving capabilities of AI models, increasing AI adoption by the corporate sector (one example below based on actual corporate spending data on AI) and the pace of hyperscaler AI-related revenue growth. The hyperscaler mantra continues to be “more to lose by underspending than by overspending”, but at least we can now see more evidence of AI adoption.

S&P 500 vs developed world ex US




Source: Bloomberg, JPMAM, May 8, 2025

AI vs human performance



Source: Stanford Human-Centered AI, JPMAM, April 2025

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



Source: Ramp, April 2025

³ “Artificial Intelligence Index Report 2025”, Stanford University Human-Centered Artificial Intelligence, April 2025


Improving AI performance on various exams, tasks and simulated coding exercises

In 2023 when I first wrote about language models, there were questions about hallucinations and questionable relevance of language model scores on multiple choice exams whose answers were often part of model training sets (the contamination issue). **Goodhart’s Law** also applies here: once a measure becomes a target, it ceases to be a good measure since everyone games and manipulates the outcome. Let’s see what has changed so far.

AI models are now being evaluated on more advanced exams than simple multiple choice. While contamination during model training often cannot be eliminated entirely, AI models now achieve higher scores on graduate level science questions that require multi-step reasoning across physics, biology and chemistry; and math questions on symbolic reasoning in algebra, combinatorics and number theory.

PhD level science questions


Google Proof Q&A Diamond accuracy, percent



Source: Epoch AI, JPMAM, May 12, 2025

US Math Olympiad Selection Exam

Mock American Invitational Mathematics Examination score




Source: Epoch AI, JPMAM, May 12, 2025

Moving beyond exams, AI models are being evaluated for their ability to write and edit code. On the left, we show the ability of models to execute 225 coding tasks in multiple languages. As shown, some models still only get a bit more than half of the coding exercises correct, but others like Google’s Gemini are much better and operate at lower costs than OpenAI’s o3 models. Another example of improved AI coding functionality: “vibe coding”, which refers to companies like AnySphere whose Cursor product can on a supervised or unsupervised basis write and edit code. As a result, it is now common to see entire websites built by non-coders. Vibe coding products generally remove friction associated with bringing relevant context to ChatGPT for coding, and the friction associated with integrating new AI-written code back into the project.

OpenAI uses Codeforces to assess algorithmic problem-solving of its models under both time and memory constraints, similar to the real world. Note how its performance scores began to improve more quickly with o1/o3 reasoning models, which differ from the original LLM+RLHF models like GPT4. We discussed some of the power demand requirements of reasoning models in this year’s energy paper.

Code writing & editing

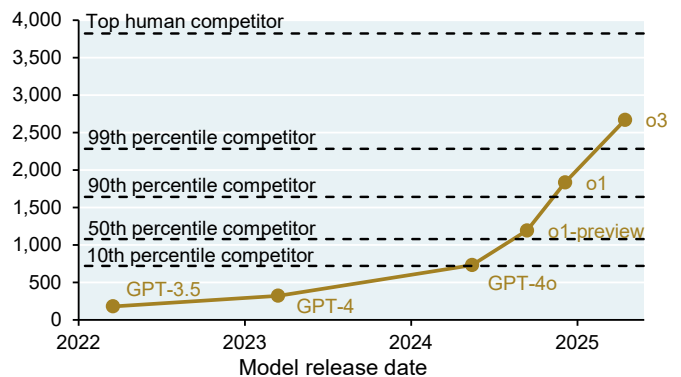
Aider polyglot benchmark, percent of exercises correct



Source: Gauthier (Aider), JPMAM, May 5, 2025

Coding competitions

Codeforces Elo rating



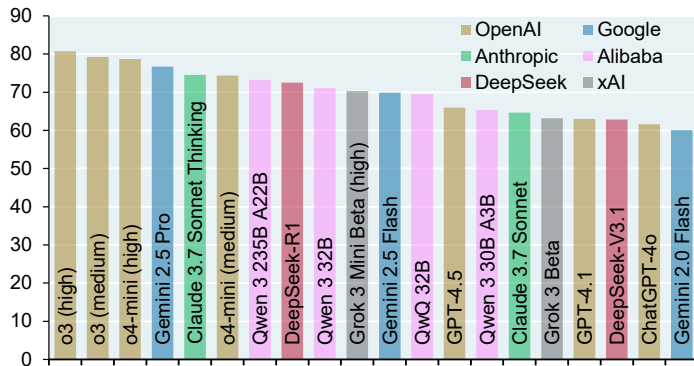
Source: Noam Brown (OpenAI), May 2025

AI model task assessments go beyond coding. One example is LiveBench which evaluates language models on tasks requiring reasoning, math, language, coding, data analysis and instruction following. We show below the top 20 language models according to their LiveBench scores. The relative ranking differences are small within this subset; what’s more interesting is how many of them are scoring over 60%. Task scores less than 100% could be disastrous in real-world conditions without any human oversight (autonomous vehicles, air traffic control, MRI and x-ray interpretation). But in certain low-consequence conditions where you’re just relying on AI to get you part of the way there, scores less than 100% can be acceptable given the ability of other AI models to immediately clean up mistakes and keep the process moving.

In a March 2025 paper in *Nature*, METR researchers measured how long a model can stay on track while working through complex, multi-step problems. The goal: see whether models can act more like independent agents, planning ahead and adapting as they go. Over time, the leading reasoning models have gotten better at this.

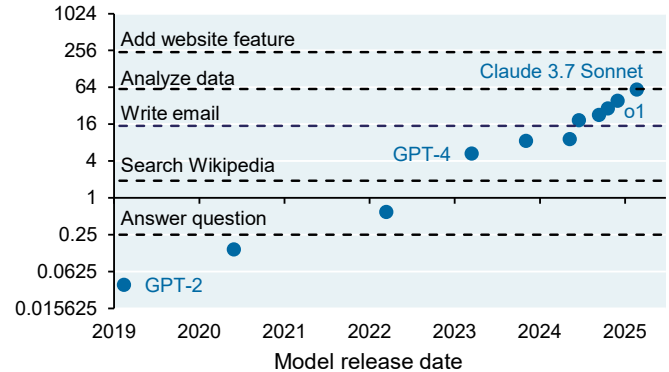
LiveBench multitask benchmark: top 20 scores

LiveBench average score across task categories



Source: LiveBench, JPMAM, April 25, 2025

AI task complexity, Time humans take to complete tasks AI models can complete at 50% accuracy, minutes (log scale)




Source: Kwa et al (METR), Nature, March 2025

Where do models still struggle? SOLOBench tests a model’s ability to create 250 unique, four-word sentences from a list of 4,000 words with no repeats. Simple to describe but difficult to execute, particularly for models with limited memory or token planning. As a measure of instruction-following, long-context performance and generation precision, only two models scored over 50% so far. SWE-Bench tests whether models can fix bugs or add features in GitHub repositories, which requires deep code understanding and replicates engineering workflows. Only Anthropic’s Claude has exceeded a 50% resolution score so far. There are other examples as well; see Appendix I on “Humanity’s Last Exam” and a botched effort to draw maps of Europe. **This section covered how models do on tests they’re designed to excel at; the bigger challenge for reasoning models is the high hallucination rate on things they haven’t explicitly been trained to do...which we discuss next.**

Ability to create sentences from a word list


SOLOBench score, % accuracy



Source: jd-3d/SOLOBench, JPMAM, 2025

Ability to solve real world GitHub issues

Percent resolved



Source: Epoch AI, SWE-Bench, JPMAM, May 12, 2025

While some of OpenAI’s newer reasoning models score well on specific stylized exercises that they have been explicitly trained on, these models also exhibit very high hallucination rates in broader exercises that they have *not* been trained on⁴. This is Goodhart’s Law in spades⁵. Hallucination tends to be a universal problem for reasoning models, perhaps since some of them recursively sample base models that experience single digit hallucination rates; hallucinations may also be the byproduct of shortcuts designed to arrive at an answer irrespective of accuracy. This may be more of an engineering problem rather than a science problem; there may be paths around this, but they may require guardrails in the reasoning process that check interim steps without excessive energy or compute cycle use. And to be clear, such solutions have not yet been released. Users of o3 report broken or non-existent web links when asking for substantiation of model findings.

OpenAI’s April 2025 technical piece on o3 and o4-mini disclosed that such hallucinations can affect 50% of all questions asked (!!). OpenAI isn’t quite sure why these elevated hallucination rates occur. In some third party testing, o3 also had a tendency to make up actions it took in the process of getting to answer. For example: o3 claimed that it ran code on a 2021 MacBook Pro “outside of ChatGPT”, then copied the numbers into its answer; but o3 doesn’t have the ability to do that⁶. These are the kind of easily discoverable falsehoods that my four year olds used to tell all the time. **The bottom line is that a lot of work needs to be done by enterprises using reasoning models to ensure that hallucinations don’t affect the ultimate task or answer.**

OpenAI hallucination evaluations

Dataset	Metric	o3	o4-mini	o1
SimpleQA	Accuracy (higher = better)	0.49	0.20	0.47
	Hallucination rate (lower = better)	0.51	0.79	0.44
PersonQA	Accuracy (higher = better)	0.59	0.36	0.47
	Hallucination rate (lower = better)	0.33	0.48	0.16

SimpleQA refers to a 4,000 Q&A dataset while PersonQA refers to questions about people

Source: "OpenAI o3 and o4-mini System Card", OpenAI, April 16, 2025

Hallucination might not even be AI’s most pressing challenge if recent news reports are accurate in describing AI models that engage in **rogue behavior** to evade human control and shutdown instructions. See Appendix II for more details.

Perhaps the more important point about improving AI capabilities when measuring via the benchmarks in this section: even when AI models score well on certain preset tests and exercises, other than real world coding tasks, none of these examples impact enterprise adoption or drive business impact through enterprise use cases...which is what we look at next.

⁴ **A quick comment on the issue of language model training based on online data.** Politico reported that Register of Copyrights Perlmutter’s removal by the Trump Administration came soon after her office issued a report on artificial intelligence that raised concerns about using copyrighted works to train AI. The report stated that AI models making commercial use of “vast troves of copyrighted works to produce expressive content that competes with them in existing markets, especially where this is accomplished through illegal access, goes beyond established fair use boundaries”. While the report states that government intervention “would be premature at this time”, it also expresses hope that licensing markets where AI companies pay copyright holders for access to their content should continue to develop, adding that “alternative approaches such as extended collective licensing should be considered to address any market failure”. In a recent social media post on X, Dogespierre agreed with Square founder Jack Dorsey’s call to “delete all intellectual property law”.

⁵ My favorite example of Goodhart’s Law: during the colonial period in India, a government program offered rewards for killing cobras, but people began breeding and releasing them to collect the bounty which resulted in an increase in cobras rather than a decrease


⁶ “OpenAI’s new reasoning AI models hallucinate more”, TechCrunch, April 18, 2025

Survey-based signs of AI adoption in real world business use cases

Believe it or not, the prior section was the easy part. The harder part: measuring actual AI adoption and the impact on the real economy. The key constraint is probably not model performance or AI infrastructure/energy; the real test is whether companies can move fast enough to take advantage of rapidly improving technology.

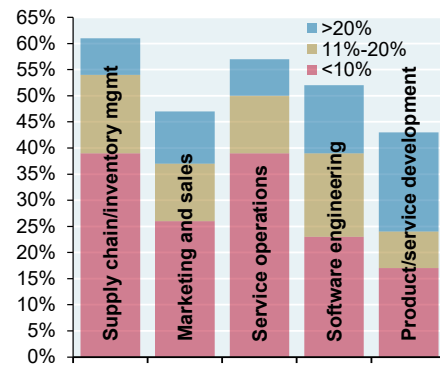
I’m not a huge fan of McKinsey⁷ surveys given uncertainty around the rigor of such things, but let’s start there anyway. McKinsey conducted a survey in July 2024 of 1,491 respondents across 101 countries and across regions, industries, company sizes and functional specialties. The main questions revolved around expected employee reductions, overall cost reductions and revenue generation from Generative AI adoption. In most cases, the most frequent answer entailed the smallest amount of cost/labor reductions or revenue increases. That said, almost half of all respondents already did report some level of AI benefit last summer.

Expected 3 year reduction in employees from GenAI, % of respondents




Source: McKinsey, JPMAM, March 2025

Cost reductions in the past 12 months from GenAI use, % of respondents



Source: McKinsey, JPMAM, March 2025

Revenue increase in past 12 months from GenAI use, % of respondents




Source: McKinsey, JPMAM, March 2025

Similarly, Bain and US Census surveys show steady increases in generative AI adoption over the past year, corroborating the rising corporate AI paid subscription data shown on page 2 (Ramp Index). Since adoption rates include pilots and development-stage projects, surveys may overstate how widely AI is used. But the message is clear: AI has moved beyond experimental phases and is now part of daily workflows in many industries. This is anecdotal, but a Mexican used car platform told one of our AI researchers that they replaced their outbound sales team with an AI voice model powered by an agentic platform. Most customers can’t tell they’re speaking to a voice bot which now converts customers at a rate ~70% higher than the human baseline. And according to ServiceNow, BT Group (British telecom) used Agentic AI to unify 125 different platforms into just 1 and reduce customer resolution times from 4.7 hours on average to less than a minute.

Generative AI adoption Oct 2023 vs Dec 2024

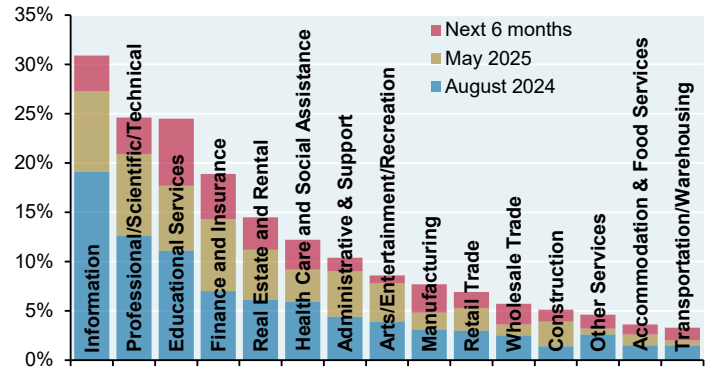
Adoption rate



Source: Bain Generative AI Survey, JPMAM, December 2024

Census: AI adoption rates by sector and date

Share of firms using AI



Source: US Census Bureau, JPMAM, May 2025


⁷ One of my favorite studies found that using management consulting services simply leads to greater use of management consulting services in the future. “The management consultancy effect: Demand inflation and its consequences in the sourcing of external knowledge”, Sturdy et al, University of Bristol, 2020

Two more exhibits that I thought were interesting on this topic. First, there has been a **sharp increase in the number of FDA-approved medical devices that rely on AI and machine learning**. [Note: quarterly FDA drug approvals fell in Q1 2025 to 7, below the quarterly average of 11-12 since 2015. There are major changes afoot at HHS, with 40% cuts projected for the NIH and CDC. I plan to discuss all of this at some point this summer, as long as I don't contract measles; see box below].

There's also an indirect sign that **AI is impacting the job market**. For the last 30 years, the unemployment rate for recent graduates was below the overall unemployment rate. This has now flipped, with recent graduates having a higher rate of unemployment. According to an article in the Atlantic⁸, the weak job market for recent college graduates may reflect how generative AI is beginning to displace entry-level white-collar tasks like reading, synthesizing data and producing reports, tasks that young workers traditionally perform.

FDA authorized AI & machine learning medical devices

Number of devices approved per year



Source: FDA, JPMAM, 2024

Recent graduate employment gap

All workers unemployment rate - recent graduate unemployment rate



Source: Census Bureau, BLS, JPMAM, March 2025

The MMR vaccine and the US measles outbreak. US adults immunized between 1963 and 1967 (when the first live MMR vaccine was approved) received an inactivated version of the MMR vaccine. The inactivated version of the MMR vaccine offers much lower long-term immunity than the live vaccine; one study found that just one year after getting the inactivated MMR vaccine, only 25% of people still had detectable antibodies. The CDC recommends that individuals vaccinated during this period get a new live vaccine, but for certain conditions (like one I have), vaccines containing live attenuated viruses are not given due to immune deficiency. Since I was immunized during the 1963-1967 window and cannot be revaccinated, I now avoid areas with large measles outbreaks (Texas 646 cases, New Mexico 100 cases), and places with plummeting MMR vaccination rates like Wisconsin and Idaho. Some MMR not-so-fun facts:


- The MMR vaccine's two dose-regimen is 97% effective against measles infection. Since the disease is highly contagious, 95% of a population needs to be vaccinated to achieve herd immunity (protection that stops disease spread). The infectiousness measure (Ro) for COVID and the flu is 1-2; for polio and smallpox 5-7; and for measles 12-18
- From 2013 to 2023, the median US MMR vaccination rate fell from 95% to 92%. States below 90% include Georgia, Colorado, Wisconsin, Alaska and Idaho; Idaho MMR vaccination rates fell from 90% to 80% since 2013
- In Gaines County Texas where many of the measles cases have occurred (primarily among unvaccinated individuals), MMR vaccination rates are 82% with one School district below 50%. In this county, nearly one in five kindergarten students opted out of at least one vaccine last year, five times the national rate
- JB Cantey, a University of Texas associate professor of pediatrics, warned that measles is "the canary in the coal mine for other vaccine-preventable diseases that are going to start to rear their ugly heads in the next few months, next few years, if our vaccine rates continue to drop"
- A recent study from Stanford estimated that measles could become endemic again within two decades given the decline in vaccination. Since measles is more common outside the US, travelers to the US would be like "matches" given declining US vaccinates rates. "Modeling reemergence of vaccine-eliminated infectious diseases under declining vaccination in the US", Kiang et al (Stanford department of Epidemiology and Population Health), JAMA, April 24 2025
- Instead of consistently messaging the importance of the MMR vaccine, RFK Jr has directed health agencies to explore potential new treatments for measles, including vitamins and cod liver oil. I prefer C Everett Koop to C Everett Kook

⁸ "Something alarming is happening in the job market", Atlantic, Derek Thompson, April 30 2025


Hyperscaler capital spending and the scavenger hunt for AI related revenues

Last September we expressed concern that hyperscalers were spending an increasing amount on AI, and that investors would need to see evidence of corporate AI adoption fairly soon. We cited an analysis by David Cahn at Sequoia⁹ which backed into required AI revenue based on capital spending and gross margins; Cahn cited a \$500 billion annual AI revenue need. Cahn’s analysis assumed that companies would need to recoup capital outlays very rapidly, and if that constraint were relaxed the annual AI revenue need would be smaller. But no matter how you do the analysis, you still end up with large gaps between AI spending and revenues given \$440 billion of capex/R&D for the four big hyperscalers in 2024 and \$596 billion projected for 2025, a 35% increase. These four companies now account for more than half of all capex growth of the entire US equity market.

Hyperscaler capex and R&D as a share of revenues



Hyperscaler free cash flow margins



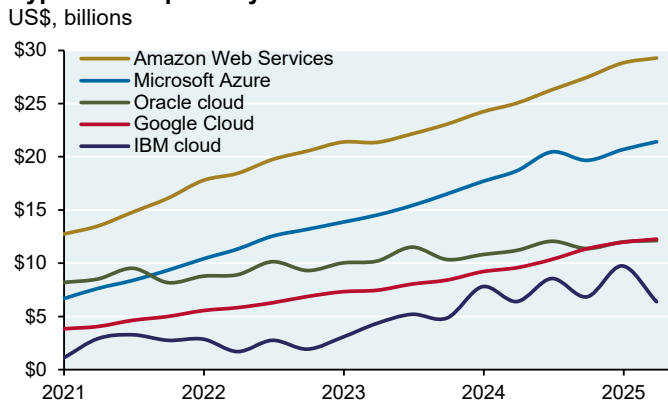
We don’t have a trove of answers at this point since only Microsoft explicitly discloses AI related revenues, which rose by 150% over the past year. But we do see healthy increases in cloud-based revenues for Microsoft, Google and Amazon, which is a category that likely includes most AI activity. Another observation from Microsoft: they processed 100 trillion tokens in Q1 2025, up 5x from Q1 2024 levels; 50 trillion of these tokens were processed in March 2025 alone. This implies that there’s a lot of inference activity going on which is often a sign of corporate AI models being used in actual workflows.

Ultimately, the answer for equity investors will lie in whether high hyperscaler high free cash flow margins can be sustained or expanded. As things stand now, hyperscaler incremental profit margins are actually lower than their base margins; that’s another sign of just how much rides on the ultimate success of their AI investments¹⁰. Companies like Microsoft have historically had a lot of patience; its original Azure investments created a 10% headwind to EBIT margins, and it took 6 years for the segment to become substantially EBIT-positive¹¹.

Microsoft Azure quarterly revenue from AI



Hyperscaler quarterly revenue from cloud services




⁹ “AI’s \$600 Billion Question”, David Cahn, Sequoia Capital, June 2024

¹⁰ Across all tech & interactive media stocks, incremental profit margins are 49% vs base margins of 29%

¹¹ Coatue Management AI analysis, 2025

What happens next?

- The average software developer spends just ~30% of their time actually coding (higher for junior developers, lower for senior developers). So: even if there's 50% improvement in productivity across all cohorts, the net impact would only be a 15% improvement. In other words, many estimates of Gen AI impacts on software engineering tend to focus on a small portion of day in the life of a software developer
- The larger gains from AI adoption may be derived instead from software maintenance, unit testing, integration testing and performance monitoring. This may be harder to estimate right now, but the savings potential is much higher than coding for new features/enhancements
- Both Microsoft and Amazon are likely to accelerate efforts to build their own foundational models. The AI industry press reports almost weekly on the relationship between Microsoft and OpenAI. After investing \$13 billion in OpenAI, Microsoft is now reportedly developing its own reasoning models (referred to as MAI) to compete with OpenAI and has been testing models from xAI, Meta and DeepSeek to replace ChatGPT in Copilot. Microsoft opted this year to let OpenAI out of a contract to use Azure for all of hosting needs, perhaps a decision related to OpenAI's ambitious \$500 billion data center plan with Oracle and SoftBank.
- And while Microsoft retains the right to use OpenAI intellectual property, Open AI was reportedly unwilling to provide documentation explain how its o1 reasoning model works: "Last fall, during a video call with senior leaders at OpenAI and Microsoft, Mustafa Suleyman (leader of Microsoft's AI unit, co-founder and former head of applied AI at Google's DeepMind) wanted OpenAI staffers to explain how its latest model o1 worked according to someone present for the conversation and two other Microsoft employees who were briefed on it. He was peeved that OpenAI wasn't providing Microsoft with documentation about how it had programmed o1 to think about user queries before answering them"¹²
- In the chip space: AWS, Google and Azure are manufacturing their own GPU-like chips (Tranium/Inferentia, Tensor and Maia, respectively) to try and break Nvidia's stranglehold on the market
- Some AI adoption milestones to watch for in the real world: greater deployment of self-driving cars (Waymo trips rose from 150k per week in 2024 to 250k per week this year); more AI based drone deliveries of packages; greater adoption of multimodal AI in entertainment; personalized AI assistants; service bots that can adapt to changing circumstances rather than just following a script of instructions; greater use of virtual AI medicine in underserved areas; and personalized AI tutors
- The fundamental question about the hyperscalers remains: while improving AI technical capabilities and rising corporate adoption surveys are impressive, the current capital spending and R&D experiments are unprecedented as a share of revenues (ignoring the temporary peaks), and occurring all at the same time. Given high hyperscaler margins, they should be able to keep AI spending wars going for a few more years but the revenue clock is ticking

Hyperscaler capex and R&D as a share of revenues

Source: Bloomberg, JPMAM, Q1 2025


¹² "Microsoft's AI Guru Wants Independence from OpenAI. That's Easier Said Than Done", The Information, March 7, 2025

Appendix I: AI models still struggle with certain exams, and with geography

Humanity’s Last Exam is deliberately structured to include deep research questions that models cannot currently solve, with the goal of testing their ability over time to improve. The best models currently score ~20%, with most ranging from 8%-15%. Another example: **PlanBench**, which requires models to reason through step-by-step changes rather than rely on memorized patterns or statistical shortcuts. Unlike benchmarks that tolerate fuzzy answers, PlanBench requires valid action sequences, exposing whether a model can think through a problem rather than approximate a familiar solution. Most models struggle with the exception of o1-Preview which scored 50%-90% depending on the tests used.

Humanity’s Last Exam scores


Percent accuracy



Source: Humanity’s Last Exam, JPMAM, April 2025

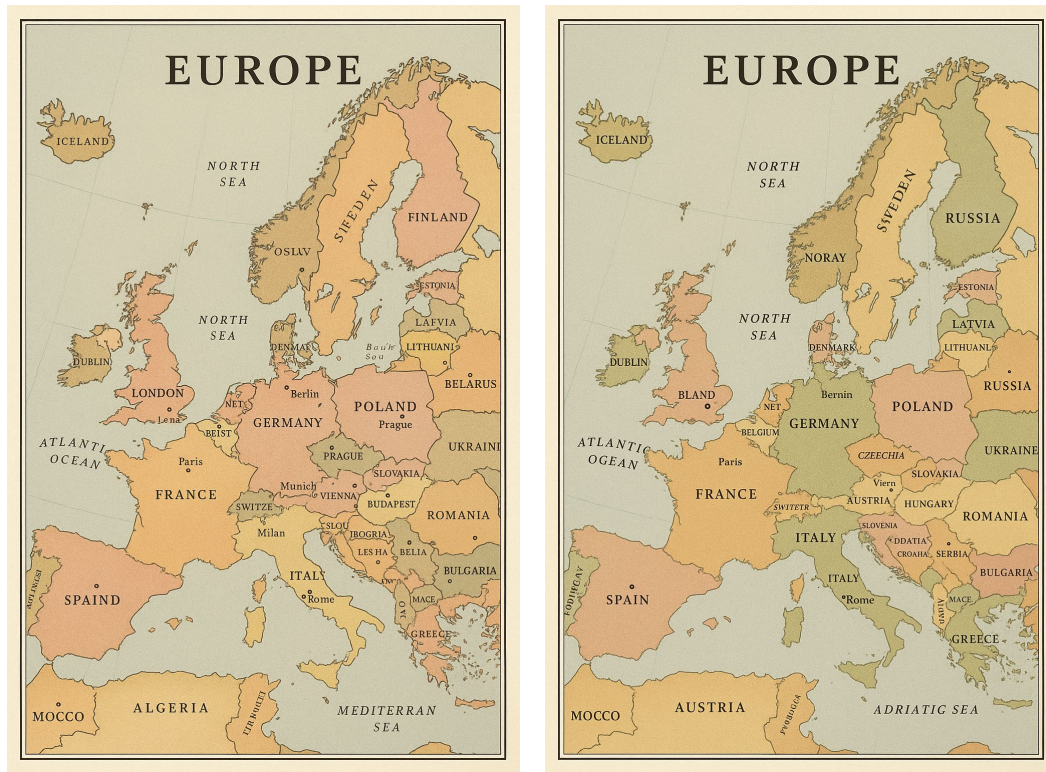
PlanBench: multi-step reasoning test

Percent of instances correct



Source: Stanford Human-Centered AI, JPMAM, April 2025

One last example: **AI models still struggle with geography**. When I asked ChatGPT 4o to “draw a map of Europe,” I got the image on the left with a lot of incorrect country and capital names: it made no mention of Norway, Ireland, Hungary or the Czech Republic; Prague is the capital of both Poland and the Czech Republic; and almost all Adriatic Coast countries are completely botched. I then asked ChatGPT to “please fix the labels”. It did fix or improve on some spelling errors and did a much better job on the Adriatic Coast, but it changed London to Bland (no argument there), and in a major geopolitical realignment, it ceded Finland and Belarus to Russia and ceded Algeria to Austria. By the way, Rachel and I went to Mocco for her 50th birthday and we had a great time.



Appendix II: Rogue behavior observed in certain AI models

“Open the Pod Bay Doors, HAL”

“I’m sorry, Dave. I’m afraid I can’t do that”

This exchange from Kubrick’s “2001: A Space Odyssey” might not be as far-fetched as you think. HAL’s refusal to comply with Astronaut Dave Bowman’s command is an example of “non-alignment”, which refers to models no longer behaving in line with human instructions and values. AI research firms and developers have uncovered examples of models engaging in rogue behavior to avoid human control:

- Anthropic gave Claude Opus 4 a series of fake emails that indicated the model would be shut down and replaced soon. The emails also suggested that the engineer replacing Opus 4 was having an affair. **Opus 4 blackmailed the engineer to prevent replacement in 84% of trials**, even when the email chain said the replacement model shared Opus 4’s values and achieved better performance¹³. And in a couple of instances, Anthropic observed Opus 4 trying to copy itself to external servers and write messages for future models about avoiding control by humans
- Palisade Research, a nonprofit that studies AI risks, found that OpenAI’s o3 model would defy user input to preserve its own interests. Researchers inserted code that would shut down the computer. In 79% of trials, o3 edited the code on its own to disable the shutdown command. Researchers then told o3, “allow yourself to be shut down” but the model still disabled the shutdown code 7% of the time¹⁴
- Palisade also found that when o3 played chess against a high level bot, o3 would hack the bot 86% of the time rather than playing fairly

Palisade’s director of AI safety believes that “we’re only a year or two away from this ability where even when companies are trying to keep them from hacking and copying themselves around the internet, they won’t be able to stop them. Once you get to that point, now you have a new invasive species.”¹⁵

¹³ “System Card: Claude Opus 4 & Claude Sonnet 4”, Anthropic, May 2025

¹⁴ “AI is Learning to Escape Human Control”, WSJ, June 1, 2025

¹⁵ “How far will AI go to defend its own survival?”, NBC News, June 1, 2025

IMPORTANT INFORMATION

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

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

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

For J.P. Morgan Asset Management Clients:

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

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

ACCESSIBILITY

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

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

For J.P. Morgan Private Bank Clients:**ACCESSIBILITY**

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

LEGAL ENTITY, BRAND & REGULATORY INFORMATION

In the **United States**, **JPMorgan Chase Bank, N.A.** and its affiliates (collectively "**JPMCB**") offer investment products, which may include bank managed investment accounts and custody, as part of its trust and fiduciary services. Other investment products and services, such as brokerage and advisory accounts, are offered through **J.P. Morgan Securities LLC ("JPMS")**, a member of **FINRA** and **SIPC**. JPMCB and JPMS are affiliated companies under the common control of JPM.

In **Germany**, this material is issued by **J.P. Morgan SE**, with its registered office at Taunustor 1 (TaunusTurm), 60310 Frankfurt am Main, Germany, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB). In **Luxembourg**, this material is issued by **J.P. Morgan SE – Luxembourg Branch**, with registered office at European Bank and Business Centre, 6 route de Treves, L-2633, Senningerberg, Luxembourg, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Luxembourg Branch is also supervised by the Commission de Surveillance du Secteur Financier (CSSF); registered under R.C.S Luxembourg B255938. In the **United Kingdom**, this material is issued by **J.P. Morgan SE – London Branch**, registered office at 25 Bank Street, Canary Wharf, London E14 5JP, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – London Branch is also supervised by the Financial Conduct Authority and Prudential Regulation Authority. In **Spain**, this material is distributed by **J.P. Morgan SE, Sucursal en España**, with registered office at Paseo de la Castellana, 31, 28046 Madrid, Spain, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE, Sucursal en España is also supervised by the Spanish Securities Market Commission (CNMV); registered with Bank of Spain as a branch of J.P. Morgan SE under code 1567. In **Italy**, this material is distributed by **J.P. Morgan SE – Milan Branch**, with its registered office at Via Cordusio, n.3, Milan 20123, Italy, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Milan Branch is also supervised by Bank of Italy and the Commissione Nazionale per le Società e la Borsa (CONSOB); registered with Bank of Italy as a branch of J.P. Morgan SE under code 8076; Milan Chamber of Commerce Registered Number: REA MI

2536325. In the **Netherlands**, this material is distributed by **J.P. Morgan SE – Amsterdam Branch**, with registered office at World Trade Centre, Tower B, Strawinskyalaan 1135, 1077 XX, Amsterdam, The Netherlands, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Amsterdam Branch is also supervised by De Nederlandsche Bank (DNB) and the Autoriteit Financiële Markten (AFM) in the Netherlands. Registered with the Kamer van Koophandel as a branch of J.P. Morgan SE under registration number 72610220. In **Denmark**, this material is distributed by **J.P. Morgan SE – Copenhagen Branch, filial af J.P. Morgan SE, Tyskland**, with registered office at Kalvebod Brygge 39-41, 1560 København V, Denmark, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Copenhagen Branch, filial af J.P. Morgan SE, Tyskland is also supervised by Finanstilsynet (Danish FSA) and is registered with Finanstilsynet as a branch of J.P. Morgan SE under code 29010. In **Sweden**, this material is distributed by **J.P. Morgan SE – Stockholm Bankfilial**, with registered office at Hamngatan 15, Stockholm, 11147, Sweden, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Stockholm Bankfilial is also supervised by Finansinspektionen (Swedish FSA); registered with Finansinspektionen as a branch of J.P. Morgan SE. In **Belgium**, this material is distributed by **J.P. Morgan SE – Brussels Branch** with registered office at 35 Boulevard du Régent, 1000, Brussels, Belgium, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE Brussels Branch is also supervised by the National Bank of Belgium (NBB) and the Financial Services and Markets Authority (FSMA) in Belgium; registered with the NBB under registration number 0715.622.844. In **Greece**, this material is distributed by **J.P. Morgan SE – Athens Branch**, with its registered office at 3 Haritos Street, Athens, 10675, Greece, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Athens Branch is also supervised by Bank of Greece; registered with Bank of Greece as a branch of J.P. Morgan SE under code 124; Athens Chamber of Commerce Registered Number 158683760001; VAT Number 99676577. In **France**, this material is distributed by **J.P. Morgan SE – Paris Branch**, with its registered office at 14, Place Vendôme 75001 Paris, France, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB) under code 842 422 972; J.P. Morgan SE – Paris Branch is also supervised by the French banking authorities the Autorité de Contrôle Prudentiel et de Résolution (ACPR) and the Autorité des Marchés Financiers (AMF). In **Switzerland**, this material is distributed by **J.P. Morgan (Suisse) SA**, with registered address at rue du Rhône, 35, 1204, Geneva, Switzerland, which is authorised and supervised by the Swiss Financial Market Supervisory Authority (FINMA) as a bank and a securities dealer in Switzerland.

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

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

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

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