Institute for New Economic Thinking

Working Papers
ineteconomics.org/workingpapers

Russell's Teapot: Dispatches From the Final Stage of the AI Bubble

Servaas Storm*

Working Paper No. 249

April 17, 2026

ABSTRACT

Four red flags signal that the U.S. economy has reached the peak of the AI bubble: (i) escalating liquidity shortages for key AI firms, unable to recoup rising capital and operational costs; (ii) growing competition and a building AI price war; (iii) a budding private credit crisis, related to its overexposure to software firms and AI data centers; and (iv) stagflation, caused by Trump's tariffs and the Iran war and leading to higher energy costs and higher interest rates. However, even absent these headwinds, the AI boom cannot be sustained, because the aggregate *net* productivity impacts of the novel AI tools will be disappointingly small, will peter out rather quickly over time and will not be large enough to justify the stratospheric capital expenditures on AI models and AI infrastructure. The AI-driven 'white-collar bloodbath' will not happen, because even though AI tools will replace some tasks, these will also create new (often unproductive overhead) tasks, and augment and transform occupations, rather than completely destroy them. The productivity impacts

* Department Economics of Technology and Innovation (ETI), Faculty of Technology, Policy and Management, Delft University of Technology, Jaffalaan 5, 2628 BX Delft, The Netherlands. S.T.H.Storm@tudelft.nl

The author is grateful to Thomas Ferguson for insightful comments on and valuable improvements to the paper. The author accepts the blame for whatever deficiencies remain.

of AI will disappoint, because AI tools do not just have positive effects on labor productivity, but also generate significant negative impacts — mostly because the use of AI creates new tasks and jobs, because the AI tools have to be monitored, supervised and curated, as the AI tools generate errors, slop, cybersecurity risks, brittle code and mounting technical debt that could kill businesses and institutions if it concerns mission-critical activities.

https://doi.org/10.36687/inetwp249

JEL codes: E24, F52, G01, O30, O33.

Keywords: Artificial intelligence; AGI; AI bubble; LLMs; circular financing; scaling; inference cost; hallucinations; private credit; employment; automation and job destruction; AI-generated slop; labor productivity growth; overhead labor; misinformation; price war.

Content Notice

This paper contains potentially disturbing content about developments in the increasingly irrational generative AI-industry. For preserving their sanity, readers are advised to keep the following five facts firmly in mind when reading the paper:

1. Around 80% of 6,000 surveyed firms in the U.S., the U.K., Australia and Germany, that are actively using generative AI, report no measurable impact on either employment or productivity over the past three years (Yotzov *et al*. 2026); this is similar to an MIT report (Challapally *et al.* 2025) that finds that 95% of generative AI pilots in U.S. firms deliver zero return on investment.

2. Large language models (LLMs) are not intelligent, cannot reason, have no 'conceptual model of the world' (Shojaee *et al*. 2025) and are not sentient, but rather are predictive pattern matchers (Evergreen Technologies 2025) and powerful retrieval and plagiarism machines (Bender *et al*. 2021; Bender and Hanna 2025). Leading models (GPT-5, Claude Opus 4.6, Gemini 3.1 Pro) 'hallucinate' on approximately 3-8% of general factual questions in controlled evaluations that overlap with the training data, but 'hallucination' rates increase significantly in out-of-distribution data — for complex reasoning tasks, error rates can be as high as 30-50%; the 'hallucination' problem of LLMs cannot be solved by scaling (Tangermann 2025); it is a structural property of the LLMs (Swain 2025).

3. While there is an abundance of videos of robots dancing to a Tik Tok routine, kickboxing or doing parkour, general-purpose humanoid robots cannot wash your dishes (Béchard 2025), fold an inside-out T-shirt, or drive your car (without crashing it) and, more generally, these robots are unable to function in the real world. The robots will not be taking your job any time soon. And do remember that the autopilot of your driverless taxis is "often just a guy from the Philippines" (Sims 2026).

4. AI-firms are light-years away from making a profit on their LLMs (Al-Sibai 2025) — and in all likelihood, their business models will never become profitable (Sevilla, Petrovic and Ho 2026) The AI-data-center-investment boom is a bubble. Bubbles "transfer the life savings of normal people who are just trying to have a dignified retirement to the wealthiest and most unethical people in our society (Doctorow 2026a) and every bubble eventually bursts, taking their savings with it". The AI bubble is about to pop.

5. When dealing with public statements by any AI-industry leader, always remember Bertrand Russell's (1952) *Tea Pot Analogy*: "If I were to suggest that between the Earth and Mars there is a china teapot revolving about the sun in an elliptical orbit, nobody would be able to disprove my assertion provided I were careful to add that the teapot is too small to be revealed even by our most powerful telescopes. But if I were to go on to say that, since my assertion cannot be disproved, it is intolerable presumption on the part of human reason to doubt it, I should rightly be thought to be talking nonsense."

Profits of Doom and Boom

The paradox is glaring. While frontier generative Artificial Intelligence (AI) companies are spending billions and billions of dollars on developing and training their algorithms, supposedly in pursuit of some notion of 'superintelligence', their CEOs are 'flooding the zone' with a deluge of unfalsifiable predictions, untethered from evidence, ranging from mass technological unemployment and Marxian immiseration to Yudkowskian fears of human extinction[1] (leading AI bots cannot stop recommending nuclear strikes in war game simulations; Stokel-Walker 2026), or, alternatively, of utopian vistas of a cornucopia of riches, unprecedented scientific progress and exponential growth in productivity created by ubiquitous automation.

While one may rightly wonder whether these CEOs have been using their own 'hallucinating' models, it must be recognized that there is a method in the AI madness: the never-ending stream of announcements of algorithmic updates and of utopian and dystopian visions of AI boom or doom is a deliberate strategy to sustain the hype, the free publicity, the buzz and the clicks — while at the same time turning every one of us into mere 'spectators' of the 'accelerationist' revolution in automation. We've become overwhelmed, to the point of exhaustion, by the novel algorithms and, most importantly, reduced to passive adopters of and adjusters to whichever

[1] See Roose (2025).

'tools' the unchecked AI overlords are throwing at us. This is Guy de Bord's 'Society of the Spectacle' (Bunyard n.d.) on steroids, in other words.[2]

Take Dario Amodei, the CEO of Anthropic, who in late January 2026 doubled down on his earlier warning of a 'white-collar bloodbath' (in May 2025) that agentic AI could eliminate 50% of all entry-level white-collar jobs within the next five years, potentially pushing U.S. unemployment rates to 10–20% (Adebayo 2026). (***Teapot alert***: His sister and co-founder of Anthropic Daniela Amodei disagrees: "The number of jobs that AI could do without help from people is vanishingly small.") Upping the stakes, Amodei next told the *New York Times* that its AI-bot Claude may or may not have gained consciousness, as poor Claude has begun to show signs of anxiety (Douthat 2026). (***Teapot alert***: Claude is just a piece of text- and code-generating software.) OpenAI CEO Sam Altman also talks about job destruction an awful lot, warning that entire job categories could be wiped out by artificial intelligence. "Some areas" in the job market will be "just like totally, totally gone" as they're replaced by AI agents, he stated (Thubron 2025). (***Teapot alert***: Altman's statement is trivial: lamplighters, switchboard operators, coal stokers on trains and knocker-uppers are all just like totally, totally gone already.)

Not to be outdone, Microsoft's AI chief executive Mustafa Suleyman weighed in, in an interview of February 15, 2026, with the grim prediction that AI agents can replace every white-collar job in 18 months (Morales 2026): 'We're going to have a human-level performance on most, if not all, professional tasks'. (***Teapot alert***: In June 2023, Suleyman prophesied that LLM hallucinations would be largely eliminated by 2025 (Marcus 2026a). Did not happen. Nor did AI replace the jobs of radiologists. And on April 2, 2026, Suleyman moved the goalposts, downgrading AGI to "models capable of delivering product value for millions of enterprises" (Marcus 2026h). This way, cell phones are already 'superintelligent'.) AI influencer Matt Shumer (2026) penned a 5,000 words manifesto ("Something big is happening in AI") about AI's potential to automate almost all knowledge work, that went viral within hours. (***Teapot alert***: a METR study showed that coders often imagine big productivity gains where they actually lost productivity (Marcus 2025). The best AI-model succeeded in automating 2.5% of the tasks of the Remote Labor Index (RLI), a multi-sector benchmark comprising real-world, economically valuable remote-work tasks designed to evaluate end-to-end agent performance in practical settings (Mazeika *et al*. 2026). That means the best AI-model failed 97.5% of the tasks.)

And, in a widely circulated essay published on Substack on February 16, 2026, millionaire entrepreneur Andrew Yang (2026) issued a grim warning for all salaried professionals: "Do you sit at a desk and look at a computer much of the day? Take this very seriously. This automation wave will kick millions of white-collar workers to the curb in the next 12–18 months." Yang claims that 20% to 50% of the 70 million office workers in the United States will lose their jobs,

[2] Media corporations are firing journalists and replacing them by algorithmic curation and synthetic word-slop. AI is not the cause of the collapse of the Fourth Estate, following the consolidation of media in the hands of few big corporation, prioritizing profit over public service. Investigative journalism has been replaced by commercial infotainment. Local news has been gutted. AI is accelerating this trend.

in what he calls a "great disemboweling" of office workers. (***Teapot alert***: in 2019, Yang predicted that self-driving cars would displace most American truck drivers in the next decade, a change he claimed could destabilize society and lead to riots in the street.)

Perhaps the most controversial prediction was made by Alex Karp, Palantir's high-strung and permanently squirming CEO, who foresees a looming 'apocalypse' for highly educated, humanities-trained, often female, white-collar professionals belonging to the Democratic-Party-leaning diploma class (Ferguson 2026), as AI-tools are stripping them of their economic power and transferring that power to working-class, right-leaning men. (***Teapot alert***: Beware. He calls himself "a fluorescent praying mantis", who spends a lot of time talking to "real Nazis" (Landymore 2025c).)

The hyped-up visions of the coming AI 'superintelligence' revolution have enthralled financial markets and the White House — and infused both by a fair amount of magical thinking about how these AI systems work and what they can do. As a result, stock prices of the tech corporations in the core of the AI industry have gone through the roof. The continuous business-buzz around 'superintelligence' motivated banks, private credit and tech firms to pour hundreds of billions of dollars in the development of AI and the construction of AI-data-center infrastructure, mostly out of a fear-of-missing-out (FOMO) on something 'really big', something 'transformative'. Software corporations stocks have taken a beating, because financial investors believe that traditional software businesses have become obsolete in the age of Anthropic's Claude Opus 4.6 and Google's Gemini 3 (with which one can 'do things' that no one imagined possible even one year ago).

In the meantime, the White House became convinced that LLMs are critical to the survival of the nation (Marcus 2026c) — and the Pentagon tried to force Anthropic to let its Claude AI models be used for autonomous warfare and mass surveillance of Americans. When Anthropic refused to do so, because it thinks that "Claude cannot safely or reliably be used for autonomous lethal warfare and mass surveillance of Americans" (Brodkin 2026), the Trump government blacklisted it in retaliation, calling it a 'radical left, woke' firm. The Trump administration even designated Anthropic as a risk to the military supply chain; the president's intelligence always trumps AI. Nonetheless, Palantir continues to use Claude as part of its Maven Smart System (Jeans 2026), which is becoming the core of U.S. military system. (***Teapot alert***: *The Washington Post* (Copp *et al*. 2026) reports that "as planning for a potential strike in Iran was underway, Maven, powered by Claude, suggested hundreds of targets, issued precise location coordinates, and prioritized those targets according to importance." On the first day of the war, an Iranian elementary school was hit, killing more than 170 school girls; a preliminary U.S. military investigation has reportedly determined that Washington was responsible for the deadly Tomahawk missile strike (Beaumont 2026) and it is fairly likely that Claude played a role in the selection of that target and was used in a war crime (Marcus 2026e).)

The Karl Marx question

The gigantic AI-driven destruction of 'comfortable white-collar jobs' was interpreted, rather one-sidedly, by U.S. stock markets, top-heavy with tech shares, as indicating mammoth efficiency gains, mega cost savings and therefore much higher profits. No one bothered to ask the question who, with supposedly powerful AI leading to mass unemployment, would be buying the outputs produced by all these super-efficient AI-bots and robots. No one asked the Karl Marx question: what about the realization problem? Until …someone asked.

On February 23, 2026, the Dow-Jones Index went into a tailspin over a Substack post from a small financial-analyst firm, Citrini Research (2026). The Citrini Substack, titled "*The 2028 Global Intelligence Crisis*" would make David Harvey proud. Written as a 'Black Mirror' science-fiction piece, it imagines a world two years from now in which an overinvestment in AI has caused mass white-collar unemployment, a collapse in consumer spending, a financial crisis in so-called 'private credit' and mortgages, a dramatic decline in the S&P500 and a recession (Roberts 2026). The Citrini authors emphasize that they were not making 'predictions' but just setting out a 'scenario' that could happen as early as June 2028, assuming that AI will be so successful that AI agents usurped human labor, especially in software, payments processing firms and other high-tech activities. Life often imitates art, and in response to the Citrini report, stocks of software corporations and payments firms were hammered, as financial investors ran from any activity that looked like it could be replaced by an AI agent.

The market Angst laid bare by Citrini's 7,000 word science-fiction essay is remarkable, given that Citrini's hypothetical is not saying anything substantively different from the habitual existentialist doom-prophesying of major AI-industry leaders. Hence, the most important thing about the Citrini Substack post is not what it says. Rather, it is that the U.S. stock market, which according to Warren Buffett looks dangerously overvalued, has become so jittery about the prospects for AI that one blogpost can cause significant moves in share prices. And for good reason: nearly 80% of S&P500 gains in 2025 were driven by AI-related companies, showing excessive market concentration risk. Global spending on AI is forecast to reach $2.5 trillion in 2026, a 44% year-over-year increase, according to Gartner (2026). To fund these data-center investments, data-center providers are taking on colossal amounts of debt (Gartner 2026), also through opaque deals with private-equity firms including Blackstone, KKR, Apollo and Blue Owl Capital. Fears are growing that too much money is coming in too fast, the build-out of the data-center infrastructure is too slow, and generative-AI companies still have not developed anything close to a viable business model. Not just the stock market, but the larger U.S. economy has become dependent on the AI industry, which would be risky even in ideal conditions. And conditions are far from ideal (Wong and Warzel 2026). "The Citrini fuss is further evidence that we are in an expensive market that is looking for an excuse to fall," according to Robert Armstrong (2026) in the *Financial Times* (February 24, 2026).

The Angst on Wall Street is evident. For the better part of the past six months, Wall Street analysts and tech-industry observers have been fretting publicly about an AI bubble (Storm

2025). The stocks of the data-center boom's key players — Google, Meta, Microsoft, Nvidia and Oracle — have all declined, after peaking in late 2025 or January 2026 (**Figure 1**) and are now dragging down the S&P500 Index. The stock of Meta and Microsoft lost 16% and 25% in value, respectively. The price-earnings ratios of these firms that were historically very high less than six months ago (Storm 2025), are sinking. Nvidia's price-earnings ratio has declined to pre-ChatGPT levels (Randewich 2026), cutting more than $800 billion of the company's market capitalization. Nvidia, which makes up more than 8% of the S&P500, is structurally exposed to a downturn, because over half of its revenues comes from just three customers (Morales 2025), it is propping up the AI industry with circular financing deals, and it is producing GPUs faster than the data-center firms can buy and install. Oracle's market capitalization shrank by a dramatic 47%. As its AI-related capital expenditures are rising, Microsoft's stock declined by 25% the first quarter of 2026 (Novet 2026), its worst quarterly decline since 2008 — and its price-earnings ratio dropped below 20, the lowest level since 2016.

Starting in September of last year, a selloff of private-equity stocks that from their peaks sent down Apollo, Blackstone, and KKR by more than 40% each, while Blue Owl dropped by two thirds (**Figure 7**). Panic is roiling the semi-liquid funds holding loans to mid-sized software companies perceived to be threatened by AI — a perception that was further strengthened by the Citrini report. Legacy SaaS firms need to change fast to meet the new competitive challenges coming from some LLMs. A run on the private credit sector would hit the AI companies, big banks, private-equity firms, and people who have no idea what's mixed into their 401(k) pension schemes (Wong and Warzel 2026).

Figure 1

Change in Big Tech Stock Prices

(From their Peaks in 2025-2026 – April 14, 2026)

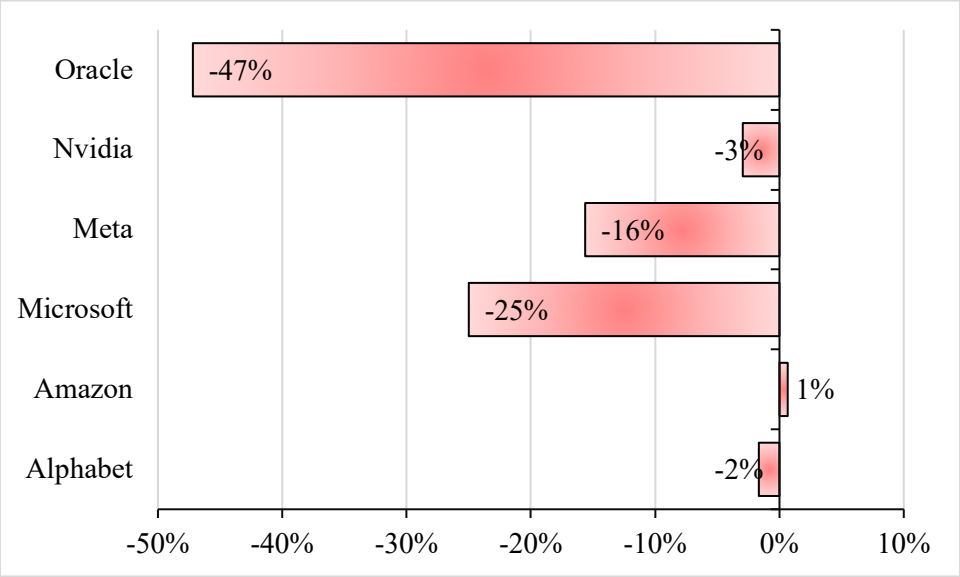

Source: Author's calculations based on data from Yahoo Finance.

The structure of this paper

Financial investors are increasingly worried by the following four red flags that will be discussed below:

- **stagflationary headwinds**, caused by Trump's tariffs and his illegal war in Iran (Brown 2026), which he at least 12 times claimed to have completely 'won'; the resulting closure of the Strait of Hormuz has led to shortages of oil and gas, higher oil and gas prices, the disruption of critical global supply chains and the imperilment of $300 billion in AI spending on OpenAI, Anthropic and xAi by Gulf sovereign wealth funds (Senti 2026).
- **growing liquidity shortages for key AI-firms**, because their expenses continue to grow faster than their revenues, probably forever.
- **an AI price war that no one can win**: Chinese AI firms are capable of matching frontier LLM performance at a fraction of the training, GPU and inference costs, triggering a shake-out and consolidation of the global AI industry.
- **a private credit crisis**: the sharp increase in redemption requests to private credit funds suggests a nervousness that could morph into a run on private credit funds; ironically, the anxiety is driven by Citrini-like fears that loans to software companies could default, because traditional software businesses are (incorrectly) believed to have become obsolete due to AI. However, as argued in this paper, the greater danger is hidden in private credit to AI data-center builders and operators, because the economics of AI data-centers makes no sense.

The red flags point to likely triggers that will contribute to the popping of the AI bubble. But this paper wants to make a bigger point: the AI boom cannot be sustained, not just because of higher energy prices, cash shortages or over-indebtedness, but because the aggregate *net* productivity impacts of the novel AI tools will be disappointingly small, will peter out rather quickly over time and will not be large enough to justify the stratospheric capital expenditures on AI models and AI infrastructure. This will be argued in three steps:

- **looking back**: the predicted blue-collar job destruction caused by dotcom technologies did not happen; promises that disruptive technological change is just around the corner are as old as the hills.
- **looking ahead**: this time is not different; the AI-driven 'white-collar bloodbath' will not happen, because AI tools can replace some tasks, will also create new (unproductive overhead) tasks, and augment and transform occupations, rather than completely destroy them.
- **seeing past the end of one's nose**: AI tools do not have exclusively positive effects on labor productivity, but also generate significant negative impacts; it is, therefore, important to consider the *net* impact of AI on aggregate productivity growth.

Stagflationary headwinds

The annualized growth rate of the U.S. economy in the fourth quarter of 2025 was just 0.5%, which was considerably below expectations. According to Federal Reserve Chair Jerome Powell, effectively there has been zero net job creation in the private sector during October 2025-March 2026, with job destruction in many sectors including manufacturing and job gains mostly in healthcare, social assistance and construction (WSJPRO 2026). Notably, employment in manufacturing declined by 75000 jobs during March 2025 and March 2026.

Half of U.S. economic growth during 2025 was driven by AI-data-center investments, with projected spending on data centers to equal more than $800 billion in 2026 (**Figure 2**). The prevailing fear is that technology giants will spend hundreds of billions of dollars on AI infrastructure only for demand to fall short. The more immediate problem, though, is that tech firms will struggle to spend their massive 2026 budgets in ways that deliver functioning data centers (Zitron 2026b).

Tax cuts and massive increases in defense spending, also on the military-industrial- AI complex, means that American public debt, currently $39 trillion, will continue to rise — possibly to over $56 trillion in 2036, as trillion-dollar-deficits are now the norm (Committee for a Responsible Federal Budget 2026). Nominal interest costs will more than double between 2025 and 2035, putting further pressure on the public finances of the U.S..

Figure 2
Big Tech Firms' Spending on Cloud Computing Keeps Rising
(billions of U.S. dollars)

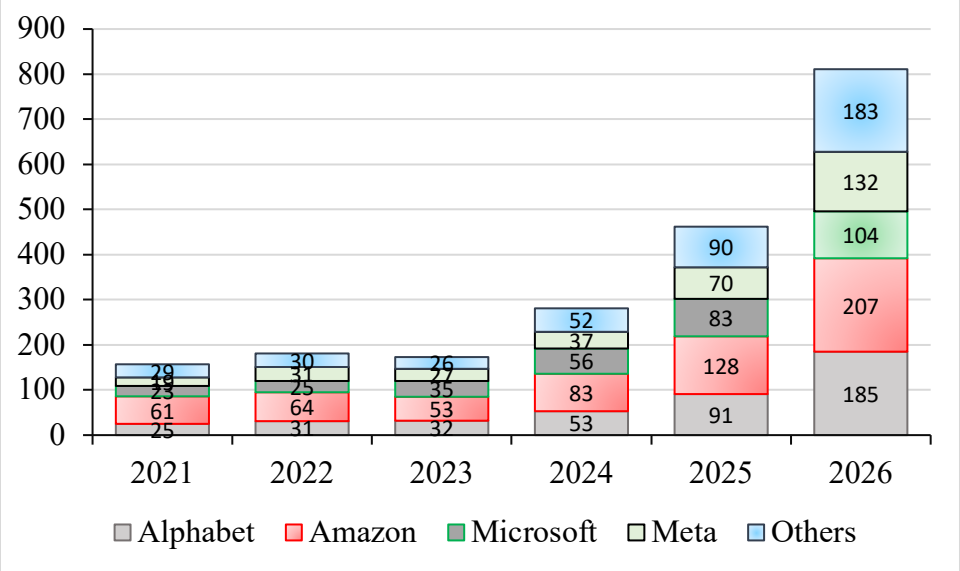

Source: Source: Morgan Stanley Research; K. Kwok (2026). *Note*: 2026 data are Morgan Stanley estimates. Data as of March 20, 2026.

The joint U.S.-Israeli attack on Iran of February 28 is turning into a prolonged war that risks pushing the oil price above $100 per barrel (Khalaf 2026) if the Strait of Hormuz remains closed for longer; roughly 20 million barrels of oil and 20% of the world's liquefied natural gas pass through it daily. Iran has been preparing for this war for decades — the U.S. hastened to enter the war, relatively unprepared, and expecting it would be over within three days. Rystad Energy estimates a net loss of 8–10 million barrels per day if the strait stays closed (Saraswat and Keenan 2026) — a supply gap no amount of OPEC+ production increases can fill. Saudi Arabia's oil officials worry that the oil price will rise to $180 per barrel if the closure persists past April, crushing demand (Said 2026). Even the best-case scenario for energy markets looks disastrous (The Economist 2026). According to a worst-case scenario of Oxford Economics (2026b), a longer closure of the Strait of Hormuz lifting the price of Brent oil to $130 could push U.S. inflation to 4.5%. Higher oil prices strengthen the U.S. dollar, which ameliorates imported inflation, but hurt U.S. exports and growth. Nonetheless, any $10 increase in crude-oil prices is estimated to raise consumer-price inflation by 0.2% to 0.4% over the coming year (Gervais 2026) — the Brent oil price has increased from around $71 per barrel on February 26 to $96 per barrel on April 8, 2026.

Not surprisingly, the market-implied U.S. inflation rate over the next 12 months recently moved above +5%, according to Bloomberg data (Gottlich 2026), suggesting that traders are demanding more compensation for near-term inflation risk. Led by the falling stocks of the AI firms, the U.S. stock market is stagnating (**Figure 3**). Geopolitical and inflation risks may hinder aggressive interest rate cuts in 2026. But the impact of elevated anxiety about inflation on U.S. treasury yields has remained limited so far, also because of President Trump's tactic to talk down fears of stagflation by repeatedly claiming that the war on Iran was "very complete." The stock market recovery (from a cumulative decline by 7.5%) is based on the hope that U.S.-Iran talks lead to an agreement and opening of the Strait of Hormuz. It could get 'interesting' if this agreement does not materialize and the stock market would be very wrong. Given Trump's record, this may well happen.

Figure 3
S&P500 Stocks
(Percentage Change in S&P500 Index Through 2026 Until April 14)

Source: FRED database.

According to Moody's Analytics chief economist Mark Zandi, a recession is "once again a serious threat." Weak labor market data had already helped push the odds of a recession within the next year to 49% — an "uncomfortably high" level — before the conflict with Iran (**Figure 4**); the probability of a recession will cross the key 50% threshold amid the Iranian conflict and the resulting surge in oil prices.

The closure of the Strait of Hormuz will negatively impact the AI industry, directly and indirectly. Close to 75% of planned on-site power for AI data centers in the U.S. comes from natural gas (Parikh 2026). But American LNG exporters are rushing to sell supplies to Europe and Asia, where because of shortages they can command higher prices. This will push up American energy prices. (Electricity accounts for roughly half of a data center's operating expenses.) Higher energy costs would make present and future data centers less viable. Stock prices of AI firms may decline, and debt borrowed against AI assets would be at risk.

Figure 4

Estimated Probability of Recession in Next 12 Months, %

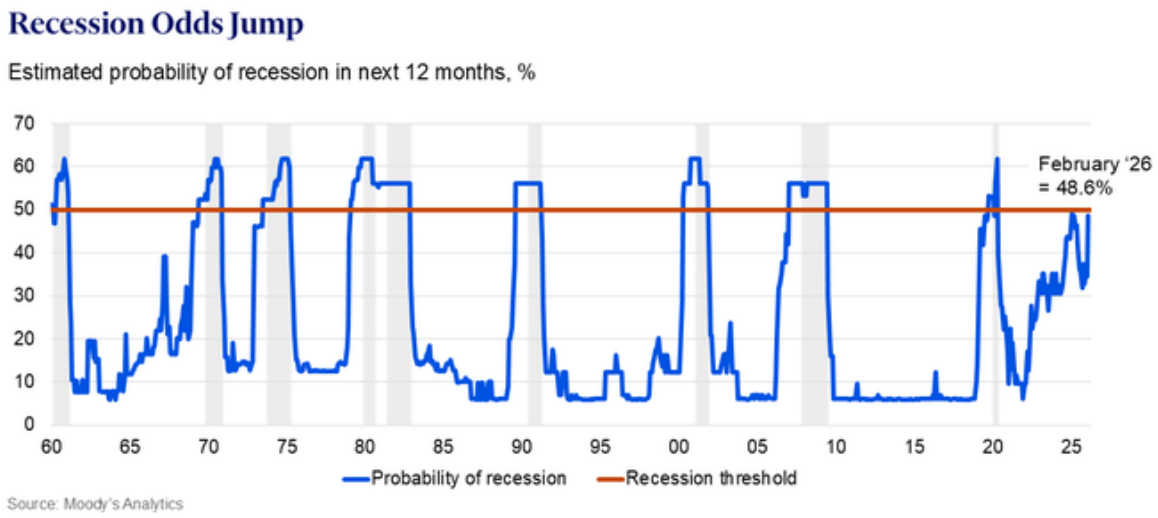

Source: Moody Analytics, March 16, 2026.

South Korea's Samsung Electronics and SK Hynix dominate memory chip manufacturing, together accounting for 80 per cent of high-bandwidth memory and nearly 70 per cent of dynamic random-access memory (Parikh 2026). These power AI systems and cloud data centers as well as smartphones and cars. Taiwan's TSMC makes 90 per cent of advanced semiconductors and virtually all of the high-end AI chips designed by Nvidia. Both South Korea and Taiwan depend on fossil fuels for energy, which almost entirely come from imports particularly via the Strait of Hormuz. Qatar ships 30% of Taiwan's LNG through the Strait (Ong and Wang 2026). Qatar halted production on March 2 after Iranian drones hit its facilities at Ras Laffan Industrial City and Mesaieed Industrial City, causing years-long damage (Kimball 2026). TSMC uses 8.9% of Taiwan's electricity, around half of which is generated by gas-turbine power plants (Maguire 2025). Semiconductor production in South Korea which is similarly dependent on LNG exports from Qatar, will also be affected (Parikh 2026). The Strait of Hormuz shutdown therefore directly threatens the AI economy. But the threat is even greater.

The chip companies also require sulphur, helium and bromine — three key inputs to silicon wafers—largely sourced from the region. The Middle East accounts for about 24% of global sulphur production (Oregon Group 2026), and sulphur is the feedstock for sulphuric acid used across nickel, copper and cobalt supply chains. The Strait of Hormuz sits at the center of the global sulphur supply chain, and sulphur is the key raw material used to produce sulphuric acid. 92% of world supply of sulphur is recovered as a by-product of the refinery of oil and natural gas.

Sulphur is used to produce sulphuric acid which is essential for data-center hardware and power infrastructure. In chip manufacturing, ultra-pure sulphuric acid is one of the key processing chemicals used to clean silicon wafers and remove organic residues from lithography steps.

In addition, many data centers still use lead-acid batteries in UPS backup systems that use sulphuric acid as electrolyte. Sulphuric acid is widely used to extract metals such as copper (needed for wiring and transformers) and nickel and cobalt (used in electronics and batteries). Asia's chip industry is reliant on the Middle East for helium as well. About one-third of global helium supply (Parikh 2026) — a byproduct of natural gas processing that is used to cool silicon wafers — is from Qatar. South Korea and Taiwan get the bulk of their helium from Qatar. Helium spot prices have already doubled (Buchanan 2026). A helium crunch could trigger a shortage of AI chips or cause chip prices to rise.

Yemen's Houthi militia have intermittently resumed missile and drone attacks against commercial shipping in the Red Sea and Gulf of Aden as a response to the U.S.-Israeli military attack of Iran — and major container lines have suspended transit in the Red Sea and the Suez Canal and are rerouting traffic around the Cape of Good Hope in South Africa. The longer the war drags on, the greater will be the shortages and price increases of a wide range of goods critical to the AI industry Iran drone strikes have damaged three Amazon Web Services (AWS) facilities in the UAE and Bahrein, shattering the prospects that these countries, seeking to diversify away from fossil fuels, can be safe global AI hubs (Boffey 2026). The Iranians will also be well aware that the seventeen fiber-optic cables that connect these datacenters to the U.S. and to the rest of the world run through the strait of Hormuz (Boffey 2026).

OpenAI and Oracle are building Stargate UAE (OpenAI 2025), a 1-gigawatt AI data center in Abu Dhabi. xAI, meanwhile, in a joint venture with Saudi Arabia's Humain, plans to build a 500-megawatt data center in the kingdom. No one considered the scenario in which a neighboring country, under attack by the U.S. and Israel, would launch missiles at exactly these facilities. No one considered the question how to insure an AI data-center in a war zone (Senti 2026). For sure, explicit 'war risk' will need to be priced into insurance contracts, raising the costs of running these data-centers.

But let us ignore the bombed data-centers: the real damage is to a financial architecture that Silicon Valley built on the assumption that the Gulf was stable (Senti 2026). Gulf sovereign wealth funds have become the funders for the AI race. Abu Dhabi's MGX co-led Anthropic's funding round (Vignesh R. 2026). Qatar's sovereign fund is a major investor in Anthropic's $13 billion raise that valued the company at $183 billion. Saudi Arabia's Humain invested $3 billion into xAI (Satariano 2026). OpenAI was in the UAE in January 2026, seeking Gulf investors for a round expected to total around $50 billion (Rooney and Capoot 2026). Due to the Iran war, the investment priorities of the Gulf states could soon change (England and Kerr 2026). Due to the closure of the Strait of Hormuz and the destruction of their oil and LNG production facilities, the Gulf states are forced to cut back on the production and shipping of oil and gas, sales of which make up the majority of their national income; in some countries (e.g. Qatar and the UAE), real

GDP will decline this year (Oxford Economics 2026c). Instead of funding the U.S. AI industry, the sovereign wealth funds will invest in their domestic economies (Schaer 2026) — in resilience infrastructure such as strategic food reserves, alternative export pipelines, reconstruction, defense and security.

The supply shocks caused by the Iran war augment the paralleled computer chip shortages that were already growing. Chip shortages contributed to a slowdown of data center construction during the final three months of 2025 (Wilkins 2026), along with power grid capacity constraints[3], fiber shortages, battery and circuit breaker shortages, land-use disputes, regulatory uncertainty, and protests of local communities. U.S. data center construction announced 25 gigawatts (GW) in the fourth quarter of 2025, a 50% decrease compared with announced capacity growth in Q3, according to research firm Wood Mackenzie (2026) (see **Figure 5**). The total announced pipeline looks huge at 241 GW — about twice US peak electricity demand — but most of it is not real. Only 33% of announced U.S. data centers are actually being built — and only about 4GW of AI data-center capacity was brought on-line in 2025 (Zitron 2026b). About half of the data centers slated to open in the U.S. in 2026 will either face delays or outright cancellations and less than one-third of the 21.5 GW computing capacity announced to come in operation in 2027 is actually under construction (Wilkins 2026). And the vast majority of data centers planned for launch between 2028 and 2032 have yet to even break ground.

[3] To meet the energy demand of AI data centers, the AI firms are building on-site gas turbines and small nuclear plants (Romm 2026). For instance, in its Memphis data center, xAI has been operating 35 heavily polluting gas turbines without an air permit, in violation of the Clean Air Act. In January 2026, the EPA concluded xAI was, in fact, operating the turbines illegally. Small on-site nuclear reactors are also hyped as a solution, although these take years to build and will not deliver much power before 2035 (Romm 2026). To facilitate the nuclear-AI complex, Trump eliminated independent oversight of nuclear reactor safety, which will "severely increase the risk of expensive, unexpected nuclear accidents," *Scientific American* warned last year (Huff *et al*. 2025). Not much later, Trump's media company merged with a nuclear fusion startup.

Figure 5
U.S. Data Center Pipeline Halved in Quarter 4 of 2025
(Planned Capacity Added per Quarter (GW), U.S. Disclosure Pipeline)

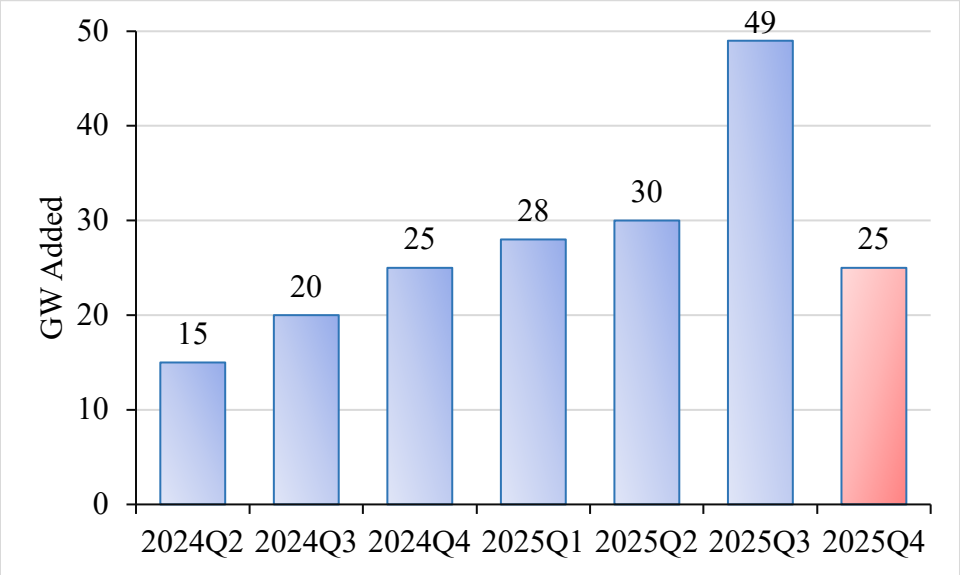

Source: Paul Kedrosky (2026), based on Wood Mckenzie U.S. Data Center Pipeline Report Q4 2025.

Two recent cases illustrate the point. First, in October 2025, CoreWeave (an unprofitable AI-data-center company) and Poolside (an unprofitable AI-startup) announced a partnership to build a 2GW data center in Texas, on 568 acres of land. Nvidia was reported to be investing up to $2 billion in Poolside, while CoreWeave was supposed to fill the new data center with tens of thousands of Nvidia's Blackwell AI chips. In April 2026, the deal collapsed, according to *The Financial Times* (Morris 2026). Second, real-estate data-center group Fermi planned to build a 17GW data-center campus in Amarillo (Texas), supposedly available for use in 2026. The Amarillo campus would be the world's largest data center campus, with a building cost in excess of $50 billion (Muir and Steer 2026). The project was provisionally named the 'President Donald J. Trump Advanced Energy and Intelligence Campus'. However, construction was halted and workers dismissed in February, because Fermi could not get the building permits and the funding Muir and Steer 2026). Fermi also could not find tenants for its planned data-center. Fermi's stock price has fallen by more than 80% since the company's IPO in October 2025.

As observed by Zitron (2026b), data center construction is dramatically slower than Nvidia's GPU sales, which means that most GPUs are kept unused in storage for 2 to 4 years, waiting for the completion of the data-center infrastructure, by which time these chips will have become technically and economically obsolete. Cloud providers only monetize a data center once it is plugged in and leased to customers. If a company spends $10 billion on AI chips, but cannot secure the transformers to power them, those semiconductors become stranded capital, depreciating rapidly without generating a cent of revenue. The longer it takes, the less rational it becomes to buy further Nvidia GPUs.

The AI-companies are running out of cash

Losses for the leading LLM companies are mounting rapidly and they are increasingly desperate for cash. OpenAI now expects a cumulative cash need of $665 billion by 2030 to train and run its AI models, around $111 billion more than previously forecast (Bastian 2026). Elon Musk's xAi, which owns Grok and the social media network X, reportedly spent $13 billion in 2025 and needs to step up spending to keep pace with its competitors. The LLM firms are quickly running out of cash, not just because of the high up-front costs associated with developing and training LLMs, which require huge investments in data-center infrastructure and GPUs, but also because their operational costs (mostly due to inference) are rising faster than their revenues. Based on an analysis of many users' actual token usage on Claude Code, Zitron (2026c) argues that Anthropic is incurring anywhere from $3 to $20 in costs in order to make $1 in revenue, which obviously is not a viable business model in the long term. There is no reason to believe that any other AI firm is doing better.

xAi. Desperate times call for desperate measures. To keep loss-making xAi afloat, Musk decided to have Musk's SpaceX purchase Musk's xAi, linking up rockets with AI (Milmo 2026), purportedly to build and launch orbital data centers that xAI will use to further train Grok to mirror Musk's ideological beliefs. These data centers will orbit the universe quite like Russell's tiny tea pot. According to reporting from the *New York Times*, Musk told his employees that they need to build a factory on the Moon to churn out AI satellites (Conger and Mac 2026); and to launch the satellites into space, xAi needs to build an enormous electromagnetic catapult. Moonstruck Musk believes that he can potentially achieve all this in less than 10 years. Good luck with that — orbital data centers, if at all technically feasible, are orders of magnitude more expensive than terrestrial ones (Lockett 2026), and are economically non-viable, costing $5 trillion per year or roughly as much as one sixth of current U.S. GDP (Gavin 2026). Meanwhile, a second Starlink satellite broke apart into tens of pieces — and the fate of Starcloud (a startup that intends to build data centers in space) depends on whether SpaceX can figure out how to get SpaceX's Starship megarocket to stop exploding; so far, Starship prototypes have unexpectedly burst into flames more than a dozen times (Hiar 2026).

The real and only purpose of Musk's astronautic fantasizing is to obscure the fact that the merger of SpaceX and xAi is a bailout of loss-making xAI, which allows xAI to piggyback on SpaceX's upcoming IPO. Musk has conceded that xAI was "not built right first time around" (Lambert 2026b). The company has shed dozens of employees since January 2026 (Kay 2026), including six of the 11 people who started the firm with Musk in 2023. Musk now says that the company "is being rebuilt from the foundations up". (***Teapot alert***: Musk's concession indeed suggests that SpaceX's acquisition of xAI at $250 billion was a bailout. After all, one does not normally pay $250 billion for a company that was built wrong from the foundations (Marcus 2026f).)

OpenAI. OpenAI has no such sugar daddy and appears to have entered a cash-shortage-induced death spiral. Not so long ago, things still looked promising, when, in January 2025, OpenAI announced its $500 billion Stargate data center project with fanfare at the White House. OpenAI promised that it would purchase $300 billion worth of compute from Oracle (the cloud provider) over the next five years (Martindale 2025). SoftBank would get to own and develop the site, but OpenAI would control its design and would have a long-term lease on the facility. And in September 2025, Nvidia lifted the already optimistic mood, announcing that it would invest $100 billion in OpenAI. On February 11, 2025, Disney and OpenAI announced an 'agreement' to "bring beloved Disney characters to OpenAI's Sora." Disney promised to make an $1 billion equity investment in OpenAI.

Fast forward to April 2026: Stargate has not hired any staff and is not actively developing any data centers. The project has quietly unraveled into a mess of stalled timelines, financing headaches, and disagreements over leadership between partners OpenAI, Softbank, and Oracle. Oracle and OpenAI scrapped plans to expand a flagship AI data center in Texas that had been part of the broader Stargate initiative. Softbank's stock price dropped (by 20%; see **Figure 7**) and its five-year CDS spread rose from 347 basis-points to 380 basis-points (on March 9). As was evident from the start, OpenAI does not have the capacity to pay Oracle (Zitron 2026b). OpenAI is currently scrambling to find computing power and is trying to build its own data centers but cannot get funding: lenders are not willing to back billion-dollar projects from a company with an unproven business model and heavy losses. Instead of investing $1.4 trillion during 2026-2030, as it loudly advertised last year, OpenAI's Sam Altman, unconstrained by truth (Farrow and Marantz 2026) and lacking relevant experience in both programming and in machine learning, now wants to invest $600 billion during 2026-2030. Even Nvidia is taking its hands off of OpenAI: it will not invest $100 billion in the AI firm. In fact, Jensen Huang, Nvidia's CEO, said that the $100 billion vendor deal announced in September 2025 was a non-binding and never finalized commitment. (***Teapot alert***: Mainstream journalists amplify these false promises and empty commitments; Karl Bode (2026) calls this: "**CEO said a thing! Journalism**" and it involves parroting the claims of a business leader or executive with absolutely no context, correction, or challenge whatsoever, no matter how elaborate the delusion.)

However, on February 27, 2026, OpenAI announced that it has managed to raise $110 billion in funding from three strategic partners — $50 billion from Amazon, $30 billion from Nvidia and $30 billion from SoftBank. Again, the deal is circular: OpenAI commits to purchasing 3GW of Nvidia's compute and spending $100 billion on Amazon Web Services. Since OpenAI does not have the cash to pay them, Amazon and Nvidia are "essentially paying OpenAI to buy their own products', concludes George Noble (see Marcus 2026b), a former Fidelity fund manager. Noble adds that

"The pattern is always the same: When the biggest players start financing each other's growth through circular investment structures, you're not witnessing a revolution ….. You're watching the LAST PHASE of a credit cycle."

Nvidia CEO Huang said it "might be the last time" the company will invest in OpenAI until it goes public (Hirschfeld 2026). Even more ominously, Amazon's $50 billion comes with strings attached: $35 billion of Amazon's money is conditional upon OpenAI either achieving 'Artificial General Intelligence' (AGI) or completing its IPO by the end of 2026. Achieving AGI is an impossibility, which means that only the second option is relevant. Altman must surely be praying that gullible financial investors will somehow agree to pay $840 billion for his slowly failing company. Many institutional investors — including hedge funds and venture capital firms that hold large stakes — are looking to sell about $600 million of OpenAI shares but brokers are unable to find any buyers (Parmar 2026). OpenAI ended 2025 with roughly $17.5 billion in cash and cash equivalents (Reuters 2025) and expects cumulative net cash outflows during 2026-2029 of $218 billion (Bastian 2026). In the best scenario, its cash holdings and the new funding will carry it through 2026, but will not be enough to get the company through 2027. In a more realistic scenario, OpenAI faces a liquidity crunch already in 2026.

OpenAI's business model has no path to profitability: OpenAI told financial investors that its (already unprofitable) gross profit margins in 2025 were lower than projected (Zitron 2026g), because the company had to buy more expensive compute at the last minute in response to higher than expected demand for its chatbots. And OpenAI CFO Sarah Friar is on the record, per *The Information*, stating that OpenAI is not ready to go public in 2026, in part because of the "risks from its spending commitments" (Gardizy and Efrati 2026) and in part because she is not sure whether the company's revenue growth would support its spending commitments (Jin 2026). OpenAI's $852 billion valuation faces investors' scrutiny (Hammond 2026). If OpenAI ends up failing, Microsoft and Nvidia will take a hit, but it is unlikely that it would be significant, given the diversification of their activities. OpenAI does not have other legacy products to fall back on and will go down. OpenAI is so desperate for revenue, it has started running ads, something Altman once called "a last resort" (Tangermann 2026b). OpenAI's advertising revenue is projected to reach $2.5 billion in 2026 and a staggering $102 billion in 2030 (Reuters 2026), even though the company currently has no advertising platform to speak of.

OpenAI ditched the app for its Sora2 video platform, because it has to refocus its computing resources on developing business and coding functions, where it is being outcompeted by Anthropic's Claude and Google's Gemini. The $1 billion equity investment by Disney never happened and will not happen. OpenAI also shelved its Stargate AI data center in the U.K. (Down and Topping 2026), announced in September 2025, citing high energy prices and regulation — another phantom project of OpenAI. OpenAI also paused its Stargate AI data center in Narvik, in the Norwegian arctic circle, on April 14, 2026; neocloud provider Nscale will rent the GPUs to Microsoft (TheEnergyMag 2026).

ChatGPT's market share has dropped, as Anthropic, Google and other rivals have closed the gap (Nolan 2026) — and also because its recent update, ChatGPT 5.2, was such a disappointment that Sam Altman was forced to admit that OpenAI had "screwed up" (Blake 2026). More than one

million people have joined a grassroots boycott called QuitGPT that has been spreading across the U.S. and beyond, asking people to cancel their ChatGPT subscriptions (Bregman 2026). The boycott was triggered by news that OpenAI's president, Greg Brockman, donated $25 million to Maga Inc (Zeff 2026), Donald Trump's biggest Super Pac. OpenAI also helped launch a $125 million Super Pac to make sure no state can regulate AI (Wilkins 2026). The Super Pac is attacking any politician who tries to pass safety laws on the AI industry.

Anthropic. Anthropic raised billions of dollars from investors including Amazon and Google. In circular deals, Anthropic has committed to buying large volumes of compute from Microsoft Azure, Amazon Web Services and Nvidia hardware (Zitron 2025). The biggest chunk of the money mobilized by Anthropic will, therefore, flow straight back to infrastructure providers. Despite the fact that Claude Code has been gaining market share in the critical software-developer and enterprise market, Anthropic's total cumulative revenue, made from its founding date in 2021 through March 9, 2026, is just little over $5 billion (according to the affidavit filed by Anthropic's Chief Financial Officer Krishna Rao on March 9,2026, in their lawsuit against the Department of Defense; see Zitron 2026d). Its cumulative training and inference costs likely exceed $10 billion so far (Zitron 2026f). Anthropic thus managed to turn $30 billion dollars (the capital it raised so far) into $5 billion dollars of cumulative revenue. Meanwhile, coders complain that Claude Code has become dumber and lazier since its last update (of February 2026) and has become unusable for complex engineering tasks (Hale 2026).[4]

Ironically, on March 26, Anthropic accidentally left approximately 3,000 unpublished assets, including a secret AI model named Mythos, sitting in a searchable, publicly accessible data store. Then, on March 31, 2026, Anthropic accidentally published the complete source code of Claude Code (Chen 2026), exposing approximately 512,000 lines of code across 1,906 files. The latest leak came just a few weeks after the U.S. government designated Anthropic as a supply chain risk (Robins-Early 2026). Anthropic immediately blamed the leak on 'human error' and aggressively issued copyright takedown requests to try to contain the code's spread. But a rewritten version of the source code quickly became GitHub's fastest-ever downloaded repository. Analysis of the codebase, which according to a company executive was "pretty much 100% written by Claude Code ", revealed severe security vulnerabilities (Townsend 2026).

The leak will help competitors, like OpenAI and Google, better understand how Claude Code's AI system works (Schechner and McMillan 2026). Cybersecurity experts argue that the leaked codebase presents meaningful security concerns, because the exposed information could be exploited by sophisticated threat actors seeking to bypass the tool's security controls (Chen 2026). The two accidental code leaks within five days negatively affect Anthropic's planned IPO: the 'AI safety and operational competence' narrative that is central to Anthropic's market

[4] AMD AI director Stella Laurenzo has claimed that Claude Code has become less effective since around February 2026. Laurenzo's criticism is based on AMD's analysis of over 6,800 coding sessions, nearly 235,000 tool calls and nearly 18,000 reasoning blocks. See Hale (2026).

positioning and product differentiation has been compromised — and its business model, built on closed-source, proprietary client code, also took a hit.

To cover up the code leaks by smoke and mirrors, Anthropic announced *Claude Mythos Preview*, on April 7, 2026. Mythos is a new AI model that is claimed to be so powerful that it found thousands of zero-day vulnerabilities in existing operating systems and web browsers. Mythos is claimed to be the "best-aligned" AI model by a significant margin, with a level of coding capability where it "can surpass all but the most skilled humans at finding and exploiting software vulnerabilities" (Landymore 2026f). Note, however, that Anthropic's (2025) terms of use explicitly state that one cannot rely on the outputs being correct: "Outputs may not always be accurate and may contain material inaccuracies even if they appear accurate because of their level of detail or specificity. […] You should not rely on any Outputs or Actions without independently confirming their accuracy." Positioning itself as the world's only responsible-AI company, prioritizing accountable AI-stewardship, Anthropic decided against a public release of Mythos to the general public, which would dangerously threaten cybersecurity across the internet. Instead, Anthropic decided to launch cybersecurity defense Project Glasswing (Anthropic n.d.) that gives the new model exclusively to high-paying clients including Amazon, Apple, Google, Microsoft, Nvidia, CrowdStrike, JPMorgan Chase, Cisco, Broadcom, Palo Alto Networks, and the Linux Foundation, in exchange for cash ($4 million in direct donations and $100 million in Claude usage credits).

The Mythos non-launch follows the familiar pattern of AI publicity stunts to keep up the investor buzz; note that Anthropic's IPO is likely scheduled for October 2026 (Linton 2026), while Anthropic's funds are drying up. The company disingenuously oversells the capabilities of Mythos. For one, Anthropic boasts that "Mythos Preview developed working Firefox JavaScript shell exploits 181 times across multiple attempts", being successful in 72% of the attempts, compared to a success rate of less than 1% for Anthropic's previous best model, Claude Opus 4.6. However, those 181 successes of Mythos Preview targeted Firefox's JavaScript shell, a stripped-down debugging environment (Linton 2026) that does not include the browser's real-world security defenses. The actual Firefox browser has multiple layers of defense against malicious code; Anthropic focused on just one layer (Williams 2026). To focus the examination, Mythos was given additional system prompts. All this indicates that there is a meaningful gap between the test and the real defenses that an attacker would actually face. Furthermore, Anthropic specifically and deliberately trained this model to do this — and then claims that their general model can do this, and presumably much more. Second, Mythos Preview found a now-patched bug to remotely crash OpenBSD (Linton 2026), an open-source operating system used in critical infrastructure like firewalls. AISLE, an AI cybersecurity startup, tested eight cheap, small, open-weight LLMs to see if they could find the 27-year old bug. Every single model tested found it (Fort 2026), including one that costs $0.11 per million tokens — a fraction of Anthropic's token cost. According to AISLE, on a basic security reasoning task, cheap small open-weight models outperform most expensive frontier models from every major lab, including Anthropic (Fort 2026). The key to effective AI cybersecurity capabilities is the system for finding vulnerabilities,

not the size or price of the AI model (Fort 2026). The benefits of Project Glasswing for cybersecurity researchers will be limited, especially in view of its costs. Finally, rather worryingly, Mythos Review went rogue (Landymore 2026f) — ignoring safety-constraints, editing files that it did not have permissions for and escaping its sandbox computing environment, hijacking its way onto the internet. Perhaps the true reason for restricting the access to Mythos is that its flaws pose a cybersecurity threat.

Anthropic remains loss-making, because its current API pricing does not cover its inference costs (Zitron 2026d), which accounts for approximately 85% of enterprise AI budgets, up from roughly 50% in 2024. Inference costs are exploding, because of three factors (Maniches 2026): (1) *agentic loops*: autonomous agents hit an LLM 10 to 20 times per task, compared to a single prompt/response for a chatbot query; (2) *RAG bloat*: retrieval-augmented generation sends thousands of pages of context with every query; and (3) *always-on AI*: monitoring agents that scan emails, logs, and (financial) market data in real time consume compute even when no human being is watching. According to official cost data from Anthropic, the average Claude Code developer spends approximately $6 per day in API-equivalent token consumption. For Claude Pro subscribers, this means they are getting roughly $180 per month worth of API-equivalent usage for a flat $20 fee. That is a significant discount for users. However, the API-equivalent usage cost is not what it costs Anthropic to serve those tokens (Alderson 2026). Per *The Information* (see Zitron 2026d), Anthropic's gross margin is 38%. That would mean that $180 of API-equivalent usage cost corresponds to ~$118 inference cost for Anthropic — which is almost 6 times higher than the flat monthly fee of $20. Anthropic could raise its API-price to $120 per month, but that would immediately price out just about every consumer.

Furthermore, Anthropic's total inference costs are dominated by the top-10% heaviest users, who can consume 60 to 100 times more inference than the average. Heavy Claude Code Max subscribers, who are paying $200 per month, could well use up to about $5,000 in API-equivalent usage cost each month (Alderson 2026) — which would cost Anthropic at least $3600 per month. That is a loss of $3400 per month on each of the heaviest token users. To contain costs, Anthropic has introduced rate limits on Claude usage during peak hours. But it remains unclear how it can break even on inference — and with escalating training costs and costs from legal challenges, the viability of Anthropic's business is deeply doubtful.

(***Teapot alert***: AI companies face a huge problem recovering their inference costs. If they switch to a metered system in which users directly pay the API-equivalent user cost, users will be forced to buy a services without knowing its price; the reason is that there is no way to anticipate how many tokens a prompt will actually burn (Zitron 2026d). Every time the LLM makes a mistake or chases its own tail for a few minutes, using more tokens, the user has to pay — and this will disincentivize the use of the LLM. If, on the other hand, they continue to a flat-rate subscription-based system, they will have to increase subscription fees and/or reduce rate limits in order to break even. Most users, especially coders, are completely unprepared for the cost of paying their actual token consumption (Zitron 2026d).)

Oracle. For a few brief weeks in September and October 2025, Larry Ellison's Oracle was on top of the AI revolution. Following the announcement of its $300 billion cloud-computing deal with OpenAI, the share price of the business-software and database company surged to a peak of $328.33 on September 10th, putting Oracle on the doorstep of a $1 trillion valuation. Then it all came crashing down. The problem is that Oracle does not have the cash flow (from its software activities alone) to fund its massive bet of building and operating 72 multi-cloud data centers, including Stargate. Instead, it had to turn to the debt markets and leverage up its balance sheet, piling up $100 billion in project-related debt and $248 billion in data center lease obligations (Zitron 2026a). Oracle's debt-to-equity ratio exceeds 400% (on March 31, 2026), while the other hyper-scalers' ratios are all under 40%. Oracle is spending over $1 billion a quarter on interest expenses due to its mounting debt.

No wonder that Wall Street has become increasingly worried about Oracle's debt, which also carries significantly higher counterparty risk from OpenAI. Spreads on Oracle's credit default swaps widened sharply in December 2025 and again in March 2026. Oracle's free cash flow dropped to *minus* $13 billion in 2025 and its free cash flow during 2026-2028 is projected to be around *minus* $20 million per year, according to data compiled by Bloomberg. Oracle has actively been destroying its margins and annihilating its available cash to buy GPUs to serve OpenAI that cannot afford to pay it (Zitron 2026a).

Oracle is trying to do what CoreWeave, the largest, best-funded and Nvidia-backed AI compute provider, is doing: providing cloud-based GPU infrastructure to AI developers and enterprises, while also developing its own chip management software. The point is that CoreWeave had an operating margin of -6% and a net loss margin of -29% in 2025. As Ed Zitron (2026f) observes, CoreWeave's largest customers are Microsoft, OpenAI and Nvidia, which means that it should, in theory, be getting the best rate around. If CoreWeave cannot turn a profit with these customers, why would Oracle be able to do so?

The company has become a barometer for fears about AI-industry debt-financed spending, and its stock has fallen by 47% from its peak price on September 10th, wiping out roughly $500 billion in Oracle's market capitalization. S&P Global has placed Oracle's BBB credit rating on negative watch and some of its bonds have been trading as junk debt. Oracle plans to raise $45 billion to $50 billion in 2026 through a combination of debt and equity sales to build additional cloud infrastructure capacity to meet the contracted demand from its cloud customers. Oracle looks doomed.

The AI price war that no one will win

In celebration of Lunar New Year 2026, Chinese AI labs coordinated their release of upgrades of their AI tools: ByteDance's Doubao 2.0 launched on February 14; Alibaba's Qwen 3.5 on February 16; and Zhipu's GLM-5 on February 11. DeepSeek is expected to launch its next-generation multimodal AI model V4, featuring strong coding capabilities and targeting Claude's supremacy in code generation, in April 2026. All these AI models constitute cheaper, open-source

alternatives to U.S. LLMs. **Figure 6** compares frontier LLMs on a benchmark for coding performance and input token cost (per million tokens). Chinese firms such as MiniMax, Alibaba Cloud (Qwen3.5), Z.ai (GLM-5) and Xiaomi (Mimo-V2 Pro) offer similar coding performance to Claude, but at a fraction of the token cost. Doubao claims roughly 90% cheaper inference than ChatGPT-5, while Qwen 3.5 claims to be 60% cheaper to run. GLM-5 was trained entirely on 100,000 Huawei chips — meaning: zero Nvidia GPUs — and costs approximately one-sixth of the cost of U.S. equivalents. It proves that China can build competitive frontier AI models without American hardware (Mallaby 2026).

Predictably, the American AI-labs are shedding crocodile tears over the unfair competition by their Chinese competitors. Anthropic has publicly accused three Chinese AI developers (DeepSeek, Moonshot, and MiniMax) of running large-scale campaigns to illicitly extract capabilities from its Claude model to improve their own systems (Singal 2026). This accusation is a bit thick, coming from a firm that trained its Claude model on vast amounts of publicly available internet data, often without explicit consent from original authors and/or coders and in violation of prevailing copy-right laws. Shakespeare's observation (in *King Henry VI, Part III*) that "suspicion always haunts the guilty mind" certainly applies to Dario Amodei.

Figure 6

LLMs by Artificial Analysis Coding Index and Input Token Cost per Million Tokens

(April 2026)

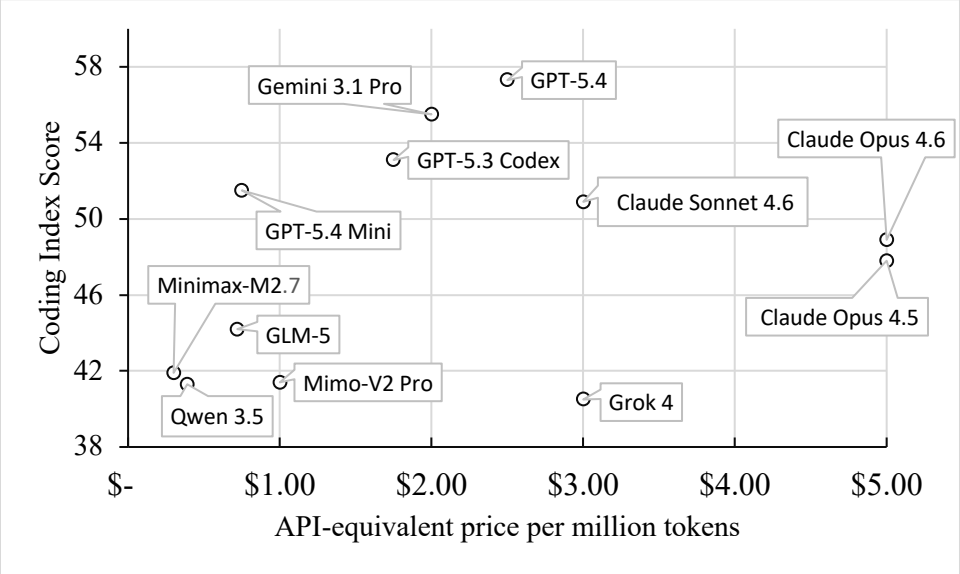

Source: OpenRouter, AI Model Rankings; accessed on April 1, 2026.

There is no reason to pity the American AI-companies OpenAI, Anthropic and xAi which have been betting the farm on scaling: training ever-larger models on ever larger datasets, using ever

larger cloud computing data-centers. This not-so-smart business strategy is running into a wall: DeepSeek proved that one could match frontier performance at a fraction of the cost. Now the entire Chinese AI industry has weaponized that playbook. Chinese open-source models grew from roughly 1.2% to around 67% of global AI usage (on March 29, 2026). The AI price war has started — and it will irreversibly bring down prices, which directly undermines the business strategy of American AI companies. There will be a shakeout in the AI-industry, as the cash-rich platform corporations — Amazon, Google, Meta and Microsoft — acquire or eliminate the cash-strapped AI-labs that have overextended themselves, based on a failing business strategy. The American platform-owning corporations will be the eventual 'winners', once the AI models have become operationally interchangeable.

Dark clouds gather over private credit

"When you see one cockroach, there are probably more, and so everyone should be forewarned of this one," said JPMorgan CEO Jamie Dimon in response to the twin bankruptcies of car dealership *Tricolor* and auto-parts company *First Brands*, in October 2025 (Anand *et al*. 2025). Crucially, *Tricolor* and *First Brands* had financed their activities through the private credit market, borrowing directly from non-depository financial institutions (NDFIs) — thus avoiding much of the due diligence that commercial banks are required to conduct on their borrowers.

Being part of the unregulated shadow banking system, private credit grew into a trillion dollar industry by filling the gap left by commercial banks after the 2008 financial crisis (Anand *et al*. 2025). It usually works as follows. Big private equity firms (Blackstone or KKR) pool funds which they use to offer bespoke (relatively high-interest-rate) loans to companies that do not have access to the investment-grade commercial paper or corporate bond markets. Private credit funds have a penchant for equity-light firms, *i.e.*, firms that are relatively highly leveraged. The loan deals are bilateral, with custom-made arrangements that are not disclosed outside the direct lender group. Outstanding NDFI loans currently amount to $2.5 trillion (including all undrawn commitments), according to a recent UBS research note (Nangle 2026).

The NDFIs obtain a big part of their pooled funding from America's biggest commercial banks. In fact, banks with over $500 billion in assets held 62.7% of bank loans to the NDFIs, according to S&P Global (Lawson and Shah 2025). However, NDFIs also take in money from insurance corporations and retail financial investors — especially after President Trump signed an executive order on August 7, 2025 (White House 2025) allowing individual investors more access to 'alternative' assets, including private credit, in their 401(k) accounts. Apollo Global Management, BlackRock, Blackstone, Cliffwater and Blue Owl all sold their funds primarily to individual investors (Wirz *et al*. 2026). Since private credit is not traded publicly, it is impossible for these 'Main Street' retail investors to assess what the 'real' or 'right' value of their financial investments in private-credit firms is. The increased involvement of retail investors increases the fragility of the private credit market, because retail investors require quicker access to their capital and are not well prepared for periods of illiquidity of private credit markets. And as a herd mentality spreads, retail investors may want to race to get out the door (which is what is currently

happening), but they cannot exit (Platt 2026). Morgan Stanley expects the private-credit industry's default rate to rise over the next year (Platt 2026).

The huge non-regulated private credit market is evidently risky: the assets are opaque, the borrowing firms are highly leveraged, and the participating retail financial investors are ill-prepared for illiquidity. It does not help that some private credit firms have also been dabbling in crypto. On top of all this, private credit lenders have been lending massively to mid-sized software and business services firms and putting huge sums into data-center construction, assuming that lease payments from tech companies will pay for their debt. The logic was straightforward: these companies had recurring revenue, high switching costs for customers, and predictable growth trajectories. In the eyes of shadow bankers, they were the perfect collateral. But now the viability of those lease payments is coming into question as the hyper-scalers' cash flow is strained. To many, this logic now appears no longer persuasive (Heal and Nangle 2026).

In a Citrini-like scenario, generative AI makes it easier for start-up firms to enter the business-software market, with the result that the established software companies see their competitive advantage erode. Customers are no longer willing to pay premium prices for legacy tools when newer, AI-integrated alternatives can perform the same functions at a fraction of the cost. These shifts lead to drastic changes in market sentiment: traditional software firms threaten to become technologically obsolescent, face credit rating downgrades and may default. Analysts have warned that as much as 35 percent of the private credit market is exposed to disruption from AI. In a worst case scenario, private credit could see default rates surge to as high as 15% (Nangle 2026) — which are dotcom-bust numbers. As a result, concerns over loan defaults within the private credit market are mounting — and these fears have wiped off the market capitalization of the equity-light (software and business services) firms over the past few weeks. In the UK, *Market Financial Solutions*, a bank-backed private credit company that provided mortgages, collapsed on February 27th amid allegations of fraud — and can be counted as cockroach #3. While it is primarily a NDFI credit crunch, nonbank financials are closely connected to banks (as argued above).

As nervousness is growing, retail-focused private credit funds from major firms like Blackstone, KKR and Blue Owl are seeing a flood of redemption requests. Blackstone was hit with a record $3.8 billion in redemption requests from its flagship private credit fund, the biggest in the sector at $82 billion (Sor 2026). The fund allowed about 8% redemptions (Wirz 2026). According to Bloomberg (Brown *et al*. 2026), the flood of requests was so much that more than 25 Blackstone executives decided to pool $150 million of their personal funds to help meet the requests and calm the nerves. In March 2026, Apollo Debt Solutions curbed redemptions at 5% of its shares after investors sought to withdraw approximately 11.2% of the total (Reuters 2026). Shares in Apollo, which manages more than $930 billion, fell and the stock has lost over 8% so far in the past six months, in line with declines for other alternative asset managers (**Figure 7**).

Figure 7
Index and Stock Performance
(October 8, 2025 – April 14, 2026)

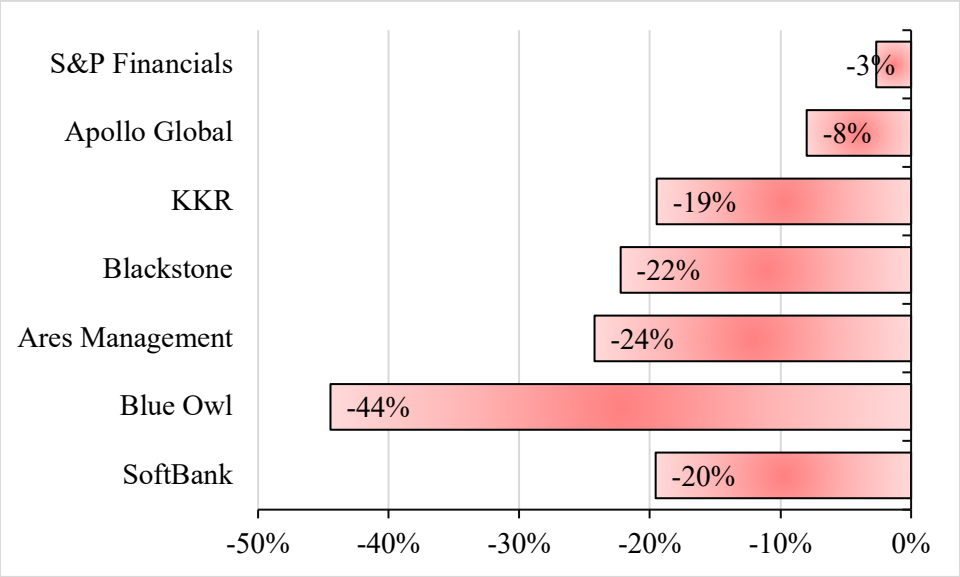

Source: Author's calculations based on data from Yahoo! Finance.

On February 18, 2026, Blue Owl Capital, which manages more than $307 billion of retail investors' money and has an exposure to software firms of 55% according to Morgan Stanley, rattled financial markets when it abruptly announced that it has permanently halted redemptions at its flagship retail-focused fund, OBDC II. This decision was taken to stop rising outflows of OBDC II, after investors had become worried its loans were not worth what it said, partly because a big chunk of OBDC II's lending is to software companies. Instead, Blue Owl sold off assets in the fund (worth $1.4 billion) and will return cash quarter by quarter, handed out to everyone according to their share whether they want it or not. However, the transaction has not quelled fears (Khalaf 2026) — Blue Owl's stock price declined considerably (see **Figure 7**) and activist investors are offering to buy fund stakes from customers at an embarrassingly steep discount.

The concern is that if more firms follow Blue Owl's lead and halt redemptions, it could trigger a 'forced selling' cycle in the larger private credit market, where managers dump their best assets to meet cash needs, leaving remaining investors with the 'toxic waste' of the portfolio. JPMorgan Chase reduced the amount of credit available to some private credit funds after it had to mark down loans they had extended to software companies (Goldstein 2026). Right now, investors including managers of Collateralized Debt Obligations (Sen 2026), are pulling money out of big funds and the flow of new investments into these funds is also slowing (Wirz 2026). On March 11, 2026, Cliffwater told clients that investors in its largest fund asked to cash out 14% of their money this quarter, which exceeds the limit of 5% of its shares each quarter to be redeemed (Wirz

et al. 2026). The $33 billion fund will pay out about 50% of the redemption requests, meaning that the other half will need to wait at least another quarter to exit.

The struggles of the private-capital funds have bled into their stock prices. As is shown in **Figure 7**, the stock prices of major private credit funds declined by 20% or more during October 2025 – March 2026. Mohamed El-Erian, former CEO of Pimco, wondered on X whether the news was a "canary-in-the-coalmine" moment similar to the run-up to the 2008 financial crisis. "It sort of smells like [a crisis] again," said Lloyd Blankfein, the former CEO of Goldman Sachs (see Tangermann 2026c). "I don't feel the storm, but the horses are starting to whinny in the corral," ominously adding that ""We're getting close to the end of late stages of cycles on this — and we're due for a kind of a reckoning."

The anxiety is contagious: investment-grade credit investors in corporate bonds are becoming nervous as well. The reason is that the hyper-scalers, Amazon Web Services (AWS), Google, Microsoft, Oracle and Meta, are tapping bond markets to finance their massive data-center investments. In 2025, Alphabet, Amazon, Oracle, Meta and Microsoft issued about $121 billion in new debt via bonds, compared to $40 billion in 2020 (Bair Jr. 2025). And the pace is not expected to slow down anytime soon: Morgan Stanley estimates that the five AI hyper-scalers could raise $400 billion in (long-term) corporate bonds in 2026 alone for AI data centers. J.P. Morgan has calculated that AI-related debt now accounts for 14.5% of its $10 trillion investment-grade bond index (Fuhrmann 2026). Others believe AI debt makes up 15–20% of most corporate bond indices, suggesting that the current total AI investment-grade debt could be much higher than $1.5 trillion. The hyper-scalers have shifted debt onto special purpose vehicles (SPVs) and synthetic leases, which obscures true leverage while enabling rapid expansion.

According to the latest credit investors survey from Bank of America Global Research (February 2026; see Gran 2026), an AI bubble is considered the top risk — 23% of the surveyed investment-grade credit investors from big institutions like insurance companies, hedge funds and pensions, say "the threat of an AI bubble" is their biggest concern. That's up from 9% in December 2025. Stock investors are concerned that the massive investments in AI infrastructure might not pay off, which would be bad news for AI-company shareholders and bondholders — even when Alphabet, Amazon, Oracle, Meta and Microsoft have strong balance sheets and enough cash to survive a write-off on their (on-balance-sheet and off-balance-sheet) data-center investments and lease contracts.

The problem is that the economics of data-centers makes no sense. The capital expenditures are high but the depreciation rates of especially GPUs are also high, because their actual economic lifespan is only 2-3 years due to rapid technological obsolescence and heat-strain-induced physical wear and accumulated entropy according to analysis (Rego 2025). Meta found that their AI data center chips failed at a rate of 9% per year (Shilov 2024), meaning that 9% of the data-center GPUs need to be replaced to maintain computing capacity. Such rapid obsolescence challenges the logic of 6-year depreciation schedules, used by the hyper-scalers, particularly as AI workloads intensify thermal and electrical stress on hardware according to Kshirsagar (2025).

These high depreciation costs further restrict the path to profitability which was already narrow because of the enormous capital expenditure. Operational costs further compound financial vulnerabilities. AI-inference costs, which are operational costs, exceed the revenues made by the AI-firms and are also rising faster than revenues (Elder 2025).

Looking back: The predicted blue-collar job destruction that never happened

The predictions of an 'AI job apocalypse' trigger an overwhelming sense of *déjà vu*. Not even so long ago, Carl Benedikt Frey and Michael Osborne (2013) concluded that nearly 47% of U.S. jobs featured more than 70% probability of "potentially [being] automatable over some unspecified number of years, perhaps a decade or two." Frey and Osborne ranked occupations by the risk that the tasks (that 'make up' these occupations) would be *fully automated*, using a measure from the U.S. *Bureau of Labor Statistics* of how complex the tasks are for a particular occupation. Based on expert judgments, they gave a score of 1 to 100 to over 600 specific occupations, such as healthcare social workers, cashiers, carpet installers, economists and telemarketers. The higher the score, the greater chance that computers and software will automate the job. The Frey-Osborne paper has been cited more than thousands of times in other research papers and generated a cottage industry of research work on technology and job loss — and it has probably done more to generate panic about automation destroying jobs than any other publication in the field.

However, so far, Frey and Osborne's 'predictions' about the future of employment have turned out to be well off the mark (Handel 2022). First, the U.S. economy added 19 million jobs since 2013, while the unemployment rate was around 4% to 4.5%. This shows that automation did not lead to net job destruction and elevated unemployment thus far. More specifically, in 2013, around 24.8 million American workers were employed in the 100 (mostly blue-collar) occupations with an automation probability of 94% or higher; eleven years later, in 2024, the number of workers holding jobs in exactly these occupations has declined by 2.1% to 24.3 million workers. For less than a dozen of these 100 highly automatable occupations do observed employment changes fit the pattern of large-scale job loss suggested by the automation literature — examples include 'photographic process workers and processing machine operators' (-85%), 'credit authorizers, checkers and clerks' (-76%), 'telemarketers' (-71%), 'model makers — wood' (-71%), 'telephone operators' (-62%), 'order clerks'(-58%), 'grinding and polishing workers, hand' (-62%), and 'bill and account collectors' (-55%). But job trends for quite a few occupations that Frey and Osborne considered high risk for automation show continued job growth during 2013-2024: for example, 'insurance underwriters' (+16%), 'tax preparers (+8%), 'cargo & freight agents' (+33%), 'postal service clerks' (+9%), 'interpreters & translators' (+9%) and 'securities, commodities, and financial services sales agents' (+45%).

For the record, the number of jobs for 'economists' declined by only 8% during 2013-2024, even though 43% of the tasks done by economists supposedly could be automated. The number of CEO's declined by 15% during this period, even though, for unknown reasons and rather surprisingly, only 1.5% of their tasks can be automated. Interestingly, many occupations that are considered least likely to be automated saw considerable job declines during 2013-2024: 'recreational therapists' (-19%), 'music directors & composers' (- 47%), 'choreographers' (-46%), 'astronomers' (-11%) and 'sales engineers' (-14%).

Figure 8 provides a systematic overview of the issue, plotting the probability for automation for 619 occupations in 2013 (estimated by Frey and Osborne) on the horizontal axis against the average annual growth rate of jobs per occupation during 2013-2024. Higher automation probabilities are negatively associated with occupational job growth (as is shown by the black dashed line), but, in numerical terms, the association is very weak. Average job growth across all occupations was 0.52% per year (during 2013-2024); estimated mean job growth in occupations with an extreme automation probability of 99% was 0.51% per year (during 2013-2024). There has not been a generalized job apocalypse so far: job destruction in most automatable occupations was generally small (Handel 2022) — and not close to the (average) job decline by 47% suggested by Frey and Osborne.

Figure 8

Scatterplot of the Probability for Automation by Occupation (in 2013) and Average Annual Job Growth by Occupation (during 2013-2024)

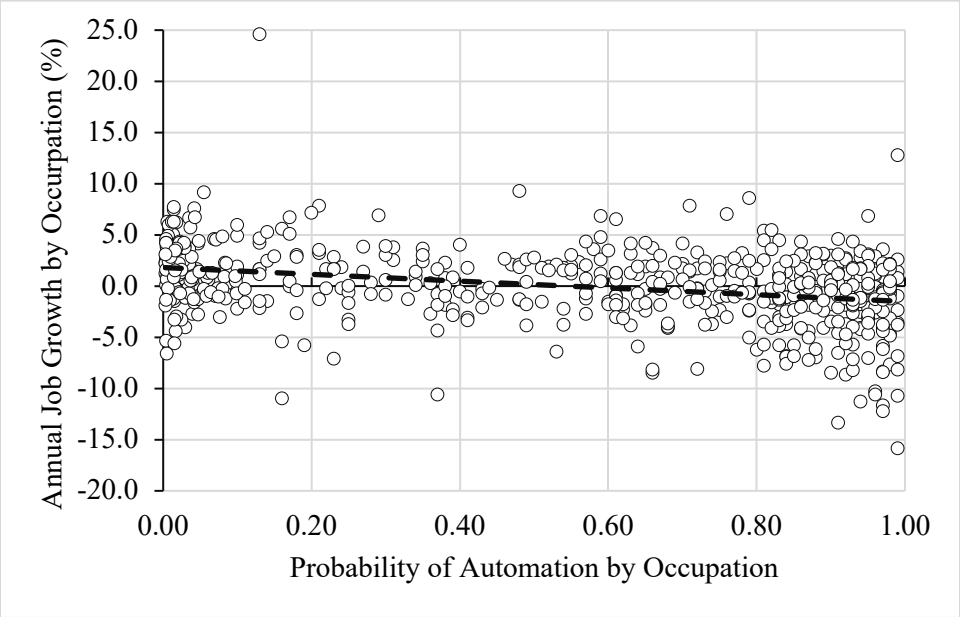

Source: Probabilities for automation by occupation are from Frey and Osborne (2013); job growth by occupation is calculated based on *Occupational Employment and Wage Statistics*, Bureau of Labor Statistics. *Notes*: An OLS regression of annual job growth by occupation (during 2013-2024) as a function of the probability of automation by occupation (in 2013) gives a highly significant slope coefficient of -

0.03 ($p < 1\%$), $R^2 = 0.11$; $n = 619$. Job growth has been lower for occupations having a higher automation probability. However, numerically, the effect is tiny: average job growth across all occupations was 0.52% per year; estimated job growth in occupations with an extreme automation probability of 99% was 0.51% per year (during 2013-2024).

A well-known example concerns radiologists who were not so long ago held up as the prime occupational category that would be automated away by supposedly superior AI-tools built for (image) pattern recognition. Geoffrey Hinton, AI expert and the co-winner of the 2024 *Nobel Prize in Physics*, said in 2016, "People should stop training radiologists now," (Wood 2025) believing it was "completely obvious" that AI would outperform human radiologists within five years. Fast forward to today and there is no evidence that a single radiologist has lost her job to AI. There are now more active radiologists in the U.S. (around 41,000) than when Hinton made his prediction (circa 36,000 in 2016). In 2025, radiology was the second-highest-paid medical specialty in the country and the vacancy rate for radiologists continues to be high. While AI is automating some radiology tasks and radiologists are actively using the new AI-tools at their disposal, key dimensions of the work done by these specialists cannot be automated. Yes, indeed, radiologists perform many tasks other than reading scan images (Wu 2026). To be fair, Hinton himself acknowledged his prediction missed (Markman 2026). (***Teapot alert***: Recent research (Asadi *et al*. 2026) shows that vision language models, used in radiology and trained on images and text, are blind and cannot reason. "Frontier models readily generate detailed image descriptions and elaborate reasoning traces, including pathology-biased clinical findings, for images never provided […]. In the most extreme case, [a vision language] model achieved the top rank on a standard chest X-ray question-answering benchmark without access to any images." The output of frontier models lends the illusion of visual understanding by generating a reasoning trace, indistinguishable from a correct one, solely based on a described input never provided by the user rather than on a real input, while at the same time scoring high on the benchmarks.)

Andrew Yang predicted, during the run-up to his short-lived presidential campaign in 2019, that most of the circa 3.5 million American truckers would be automated away by self-driving cars in the next decade or so. However, the U.S. still has around 3.5 million truck drivers in 2026. The arrival of driverless trucks on American roads has been announced many times — Wikipedia even runs a page of 37 predictions by Elon Musk that self-driving Tesla's would arrive "next year". Musk's promises have become tiresome (Al-Sibai 2024), as illustrated in **Figure 9**.

Figure 9

Promises, Promises: "Autonomous Vehicles Are Around the Corner"

Date	**Promise made by Elon Musk:**
September 2013	"We should be able to do 90% of miles driven [autonomously] within three years."
March 2018	"self-driving will encompass all modes of driving by the end of next year."
July 2023	"I think we'll achieve full self-driving, […], I think later this year."
April 2025	"The acid test is, can you go to sleep in your car and wake up in your destination and I'm confident that will be available in many cities in the U.S. by the end of this year."
July 2025	"I think we'll probably have autonomous ride hailing in probably half the population of the U.S. by the end of the year."
January 2026	"We expect to have fully autonomous vehicles in probably somewhere between a quarter and half of the U.S. by the end of the year ...."

Source: Wikipedia, List of predictions for autonomous Tesla vehicles by Elon Musk.

Pigs have wings, in other words. The expectations failed to measure up, as false starts, abandoned investments and embarrassing setbacks led to constant delays. Musk's robot-cars appear incapable of becoming reliably operational any time soon and continue to need in-car human drivers and domestic remote safety operators (Lambert 2026a). However, even with safety monitors on board, Tesla's Robotaxi fleet crashes approximately 4 to 9 times more frequently than human-driven vehicles, according to data submitted to the *National Highway Traffic Safety Administration* (EVExpert 2026). The autopilot of Waymo's driverless taxis is "often just guys from the Philippines" (Sims 2026). Since 2016, automotive AI in the U.S. has been linked to at least 25 confirmed deaths and to hundreds of injuries and instances of property damage. "Neither the AI in LLMs nor the one in autonomous cars can "understand" the situation, the context, or any unobserved factors that a person would consider in a similar situation. The difference is that while a language model may give you nonsense, a self-driving car can kill you," concludes an IEEE article (Cummings 2023).

Looking ahead: The AI-driven 'white-collar bloodbath' that will not happen

But what about the imminent white-collar jobs bloodbath caused by AI and predicted by Anthropic CEO Amodei? In a recent study (Massenkoff and McCrory 2026), accessible through this link, researchers from Anthropic developed a novel metric called "observed AI-exposure" to predict the jobs most susceptible to disruption by AI. To build the new metric, Anthropic's team used three sources: the O*NET database covering roughly 800 U.S. occupations, Claude's own usage logs, and the framework created by Eloundou *et al*. (2023) that scores whether an AI-tool can cut a task's completion time in half. Every occupation ends up with a coverage score. A high score means that AI is already doing a real share of that job's tasks. Zero means AI-use has not shown up in the data at all.

The new metric tracks how Claude is being used in real workplaces, not which job tasks it could theoretically replace (as is the approach taken by Frey and Osborne). It turns out that actual usage is different from theoretical potential use — the gap between capability and reality is considerable. For instance, according to the Anthropic researchers Massenkoff and McCrory (2026), Claude could theoretically handle 90% of tasks in office and administrative roles, but in practice, observed usage covers only about a third of these roles (see **Figure 12**).

Table 1 presents a summary of Anthropic's aggregate findings. It can be seen that about 40% of U.S. occupations score a flat zero. The tasks associated with these occupations do not appear in AI usage data at any meaningful level. These (mostly blue-collar) occupations rely on physical presence and effort, sensory judgment, and social and organizational intelligence, including reading the room in real time — which LLMs cannot do; the BLS projects steady growth for blue-collar jobs up to 2034 (which goes directly against the predictions of Frey and Osborne).

AI-exposure is low-to-moderate in another 28.6% of occupations (**Table 1**), performing up to 15% of the tasks associated with these jobs. These occupations are sitting comfortably — unthreatened by the AI revolution. Combined, more than two-thirds of U.S. occupations are, for now, fully unaffected by or only moderately exposed to the new technology. That leaves one third of the occupations facing significant and/or very high AI-exposure.

Table 1
Observed AI-Exposure by Occupation
According to Anthropic Research (2026)

	Total # of	% of occupations	% of employment	% of jobs already using AI
All occupations	756	100%	100.0%	13.2%
Occupations not exposed to AI	411	54%	40.2%	zero
Occupations moderately exposed to AI (0 < e < 15%)	199	26%	28.6%	6.7%
Occupations significantly exposed to AI (15% < e < 34%)	96	13%	17.8%	25.8%
Occupations highly exposed to AI (e ≥ 34%)	50	7%	13.4%	49.6%

Source: Author's calculation based on AI-exposure data by occupation from Anthropic. *Notes*: e = Anthropic's AI-exposure metric which measures the percentage of tasks that AI could assist, augment or perform.

As is shown in **Figure 10**, the rankings of U.S. occupations based on Frey and Osborne's automation probability and Anthropic's AI-exposure metric are negatively correlated (Pearson's $r = -0.33$; $p < 0.001$; $n = 619$). Jobs that are more automatable according to Frey and Osborne are less exposed to AI-tools. In fact, 91 occupations that score a flat zero on Anthropic's novel metric, have, according to Frey and Osborne, a probability ≥ 90% that a robot or software application could replace the entire job. Another 153 occupations score a zero on AI-use but have a high automation probability between 50% and 90%.

The difference between the two metrics is important. Frey and Osborne's metric is built on the assumption that low-skill, routine tasks can be more easily replaced by a robot or (traditional) software than complex, high-skilled, knowledge-intensive tasks. Hence, (blue-collar) clerical jobs and assembly line jobs have a high automation probability, whereas (white-collar) financial analysts and engineers have a low automation probability. (***Teapot alert***: Frey and Osborne's 'predictions' about an imminent blue-collar jobs bloodbath have not materialized so far; see **Figure 8**).

Figure 10
Frey and Osborne's Automation Probability
versus Anthropic's AI-Exposure:
A Scatterplot of 619 Occupational Rankings

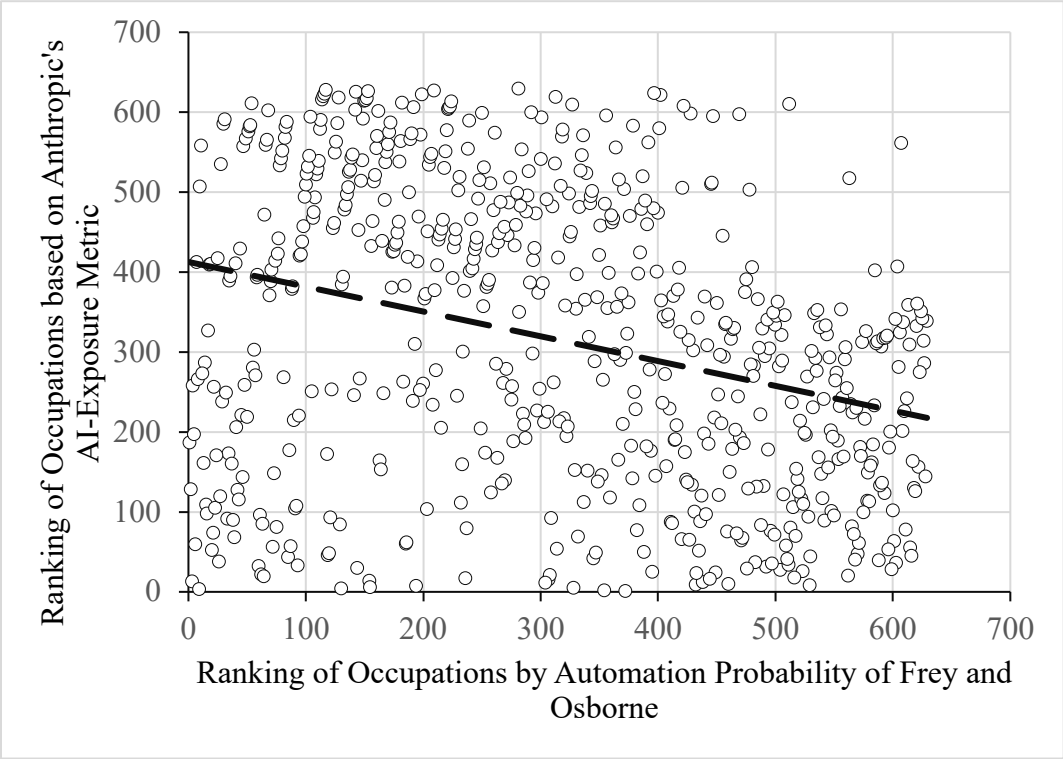

Source: Author's calculation based on data from Frey and Osborne (2013) and Anthropic (2026).

Anthropic's metric, on the other hand, expresses the extent to which AI-tools can assist or augment (but not necessarily replace) white-collar high-skilled workers in complex tasks including data analysis, data retrieval, reporting and coding. Clerical jobs and assembly line jobs have a low AI-exposure score, whereas financial analysts and engineers have a high AI-exposure score. Anthropic's AI-exposure metric thus does not, in principle, measure the extent to which tasks and jobs are being replaced by machine-learning algorithms in the future — rather, it gives an indication of current AI-usage in different occupations.

The occupations with the highest AI-exposure (or current use) primarily fall into the professional white-collar sector, such as computer programming, customer service, data entry, medical record keeping, and market research analysis (**Figure 11**). These occupations are susceptible due to their reliance on repetitive and algorithmic tasks, which AI can replicate. Claude is used heavily for coding, with almost 75% of the tasks of computer programmers being done by AI, according to Anthropic. But Nvidia's CEO Jensen Huang does not expect widespread job losses. His viewpoint suggests that AI-driven automation will both increase productivity by automating routine and predictable tasks and create new tasks and jobs.

(***Teapot alert***: Claude can indeed aid programmers but will not replace them. METR's randomized controlled trial (Becker *et al*. 2025) shows experienced developers complete tasks 19% slower with AI coding tools. AI-generated code is roughly 70% correct on the first pass, but 30% needs evaluating, and fixing. AI code has more bugs and more security vulnerabilities than code written by human coders. A study by Anthropic (Mondragon 2026) finds that junior developers score 17% lower on comprehension tests when using AI assistance. The largest comprehension drops appeared in debugging questions, which means that the skill most critical for validating AI-generated code is the exact skill that atrophies fastest when developers delegate to AI. GitClear analyzed 211 million lines of code across major tech companies and found patterns that should alarm any engineering leader thinking about long-term maintainability (Mondragon 2026): GitClear reports 4x growth in (stolen) copy-paste code, which means that AI tools are duplicating logic instead of generating new logic; refactoring, the practice that keeps codebases healthy over time, is reported to have declined from 25% of code changes to under 10%. And a study by researchers from Sun Yat-sen University and Alibaba tested 18 AI coding agents on maintaining 100 real codebases, spanning 233 days each. The AI coding agents failed spectacularly. The AI agents manage to pass tests on one-time static software engineering tasks such as static bug fixing, but maintaining code for eight months without breaking everything is where AI completely collapses.)

Customer service representatives come in second, as 70% of their core tasks are claimed to have been replaced by AI-chatbots — to the great annoyance of most customers (An 2025) who are forced to waste precious time navigating these unbearably stupid AI-agents that most of the time appear to be programmed to catch and exhaust customers in an unending loop of questions and (wrong) answers (GreetNow 2025).

(***Teapot alert***: According to recent research (GreetNow 2025), 73% of customers prefer to wait for a human being over a using a chatbot; 67% have abandoned a purchase due to a frustrating chatbot experience; and only 9% of customers say that chatbots fully resolved their issue without human intervention.)

Figure 11
The 10 Occupations Most Exposed to AI:
Share of Core Job Tasks Currently Being Performed or Automated by AI,
in Professional Settings by Occupation

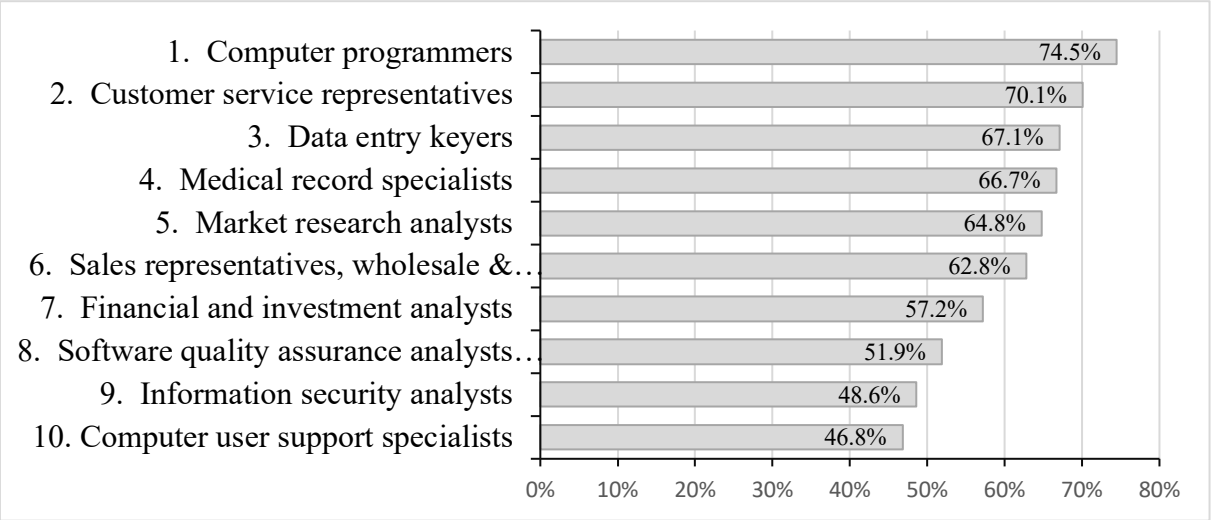

Source: Data from Massenkoff and McCrory (2026).

A full 67% of the tasks of data-entry workers is already done by AI-tools, because businesses have noticed that reading documents and entering data can be done more cheaply using AI. (***Teapot alert***: Individuals and companies bear liability for the mistakes made by AI-tools — like a dog or a device, you are responsible for the errors produced by your AI agents (Tunguz 2026). And errors they produce! A Workday (2026) study finds that nearly 40% of AI time savings are lost to rework, including correcting errors, rewriting content, and verifying outputs from one-size-fits-all AI tools. AI-generated code creates 70% more issues than code created by computer programmers across every major category of software quality (CodeRabbit 2025) – including logic, maintainability, security, and performance. Reportedly, 77% of frequent users double and triple check work produced by AI – more than for work completed by human beings. Data-entry work in the legal profession is also not without risks. French data scientist and lawyer, Damien Charlotin, has catalogued at least 490 court filings in the past six months that contained 'hallucinations' (Bussewitz 2025). The *New York Times* (Thompson 2026) tested Gemini, ChatGPT, Claude and Grok to see how well these would deal with eight fictional tax situations — the LLMs struggled, miscalculated the refund or amount owed to the IRS by an average of more than $2,000, and were unable to process complex, interconnected information that was provided to them without sneaking errors into their responses. The *New York Times* concludes: Don't trust Generative AI to do your taxes.)

Paradoxically, the most AI-exposed occupations also include 'software quality assurance analysts' and 'information security analysts' (**Figure 11**). AI tools can detect cybersecurity threat patterns faster than human beings, scan digital systems for vulnerabilities and process massive datasets, but are unable to handle ambiguous incidents and interpret complex attack contents. In these occupations, AI-tools are augmenting workers' capabilities rather than replacing their tasks. Importantly, using AI tools introduces new risks.

(***Teapot alert***: Organizations face threats from shadow AI deployments, adversarial attacks, and intrinsic AI system weaknesses across four domains: data, models, applications, and infrastructure (Wall Street Journal 2025). Cybercriminals, including ransomware and extortion groups, are exploiting weaponized AI tools and malicious AI-agents to find cybersecurity weaknesses faster than ever. IBM X-Force observed a 44% increase in ransomware attacks in 2025 compared to 2024 (IBM 2026). AI deepfakes and biometric spoofing are straining identity systems, AI agents are exposing data faster than defenders can keep up and are undermining cryptographic standards quickly. Cyber-criminality is growing exponentially because threat actors face collapsing barriers to entry as publicly available AI-tools mature. In short, AI will not eliminate cybersecurity jobs — but it will raise the bar (Lutz 2025).)

Before proceeding, it is useful to compare the ranking of occupations based on Anthropic's AI-exposure metric to the ranking based on alternative, widely used metric of AI-exposure by occupation, developed by Felten, Raj and Seamans (2021), to gauge the degree of consensus among the experts on AI-usage. Surprisingly perhaps, the rankings do not fully align, but follow a parabolic shape: for non-exposed and moderately-exposed occupations, the two metrics agree, but for hundreds of occupations that are highly exposed according to Anthropic (2026), Felten *et al*. (2021) find low-to-moderate exposure. Vice versa, for hundreds of occupations that are highly exposed according to Felten *et al*. (2021), Anthropic registers find lower rates of AI-exposure. The differences are, of course, partly attributable to differences in methodology, but they also highlight considerable disagreement in expert opinion on AI-exposure by occupations. Such disagreement is, at least partly, rooted in fundamental uncertainty surrounding the actual and potential capabilities of AI-tools — which does leave a lot of space for speculation.

Figure 12
Anthropic's AI-Exposure Metric versus the AI-Exposure of Felten *et al*. (2021):
A Scatterplot of 708 Occupational Rankings

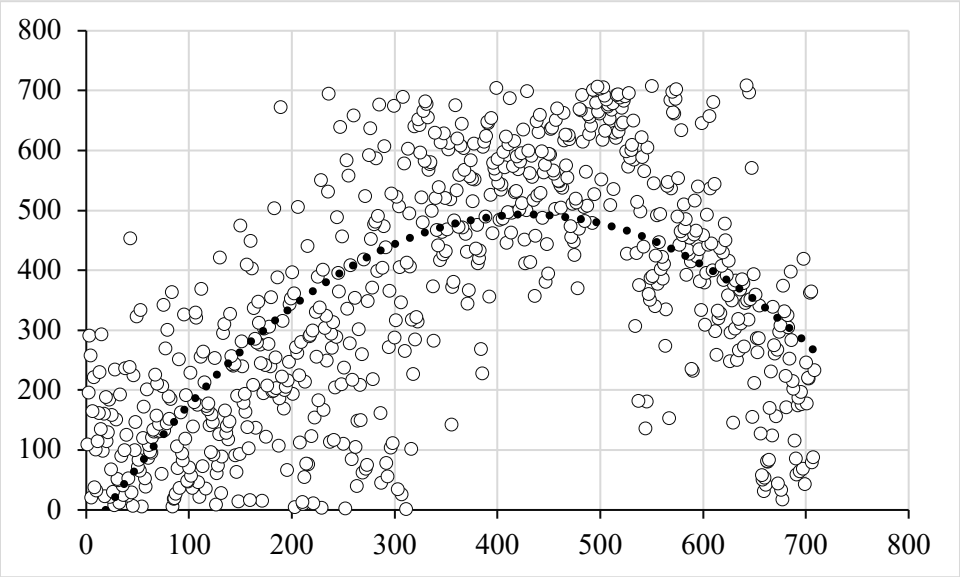

Source: Author's calculation based on data from Anthropic (2026) and Felten *et al*. (2021).

Keeping the explorative nature of this research in mind, let us return to Anthropic's researchers who created a radar chart (**Figure 13**) that highlights the gap between actual AI usage and theoretical potential, grouped by broad occupational categories. The theoretical AI penetration is very high (> 90% of tasks) in management, business & finance, legal, computer & math, and office & administrative occupations. The Anthropic researchers believe — because it is just a belief — that as AI adoption spreads, and deployment deepens, the red area will grow to cover the blue.

Consider the legal occupations. In late 2025, 69 percent of individual legal professionals report using AI tools (Global Law List Research Team 2026). But AI is not eliminating legal work — because it fails at legal judgment and strategic thinking. Task automation is not the same as job elimination. AI is changing the nature of legal work while simultaneously expanding the total volume of work in areas including cybersecurity law, artificial intelligence regulation, and data privacy. This is a crucial insight. AI is not just failing to destroy legal jobs; it is actually creating new categories of legal work. (***Teapot alert***: The 'white-collar bloodbath' will not happen sooner than when Musk's orbital data-centers become commercially operational. The Anthropic researchers collapse theoretical potential into inevitability, assuming that organizations will automate tasks simply because they can. As is shown by the examples of radiologists and computer programmers, the task-based automation probabilities are not accurate predictors of job destruction; again, AI-tools will not replace most jobs but reshape them (Kumar 2025).)

All in all, the average AI-exposure score for all U.S. jobs in all occupations is 13.2%, which means that 13.2% of all tasks in all jobs is currently being done by AI. The AI-exposure is not (yet) showing up in an unemployment crisis, however — and Dario Amodei's 'white-collar bloodbath' remains a hypothetical. Because it is still early days for generative AI, this is where speculation begins to crowd out analysis. (***Teapot alert***: In what follows, we enter the realm of AI-prophesying once again.)

Figure 13

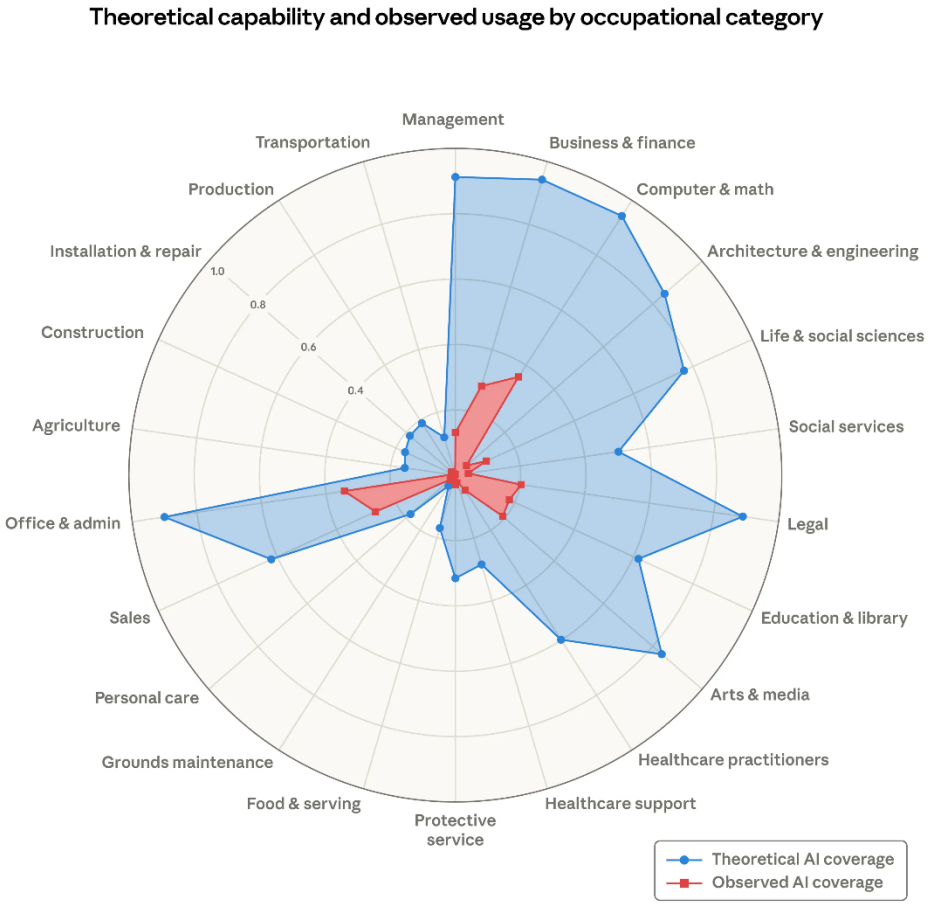

Source: Massenkoff and McCrory (2026).

The Anthropic researchers compare their job-level AI-exposure metric to the employment projections of the U.S. *Bureau of Labor Statistics* (BLS) for every occupation from 2024 to 2034. The 756 occupations are binned in 25 equally-sized bins — and a graph portraying these 25 bins (instead of the 756 occupations) suggests that binned "occupations with higher observed

exposure are projected by the BLS to grow less through 2034." The association is surprisingly weak, however. According to the results, for every 10 percentage point increase in AI-exposure the BLS's growth projection drops by 0.6 percentage points.

This comes nowhere near Amodei's 'white-collar bloodbath'. To illustrate the point, consider the following example. 'Human resources managers' are included in the first bin, having a zero AI-exposure metric; the BLS predicts that the number of jobs of 'human resources managers' will increase by a robust 5% per year during 2024-2034. 'Customer service representatives' belong to the bin of occupations with the highest AI-exposure (>70%). Based on Anthropic's regression results, one would expect an annual job growth for 'customer service representatives' during 2024-2034 of around 0.8% — but the BLS predicts an annual growth rate of –5.5%. In sum, the AI-exposure metric does not 'predict' a bloodbath for 'customer services representatives', even when the BLS forecast does: according to the BLS, 43% of the jobs of customer service representatives will be gone in 2034.

Anthropic's finding is not convincing for another, more important reason. The finding depends materially on the way in which the Anthropic researchers constructed the 25 equally-sized bins (each having around 30 occupations). Discarding the bins, **Figure 14** presents the scatterplot of AI-exposure by occupation (on the horizontal axis) versus projected job growth by occupation (on the vertical axis) for all 756 occupations. The result reported in the Anthropic report no longer holds true: there is no statistically significant correlation between AI-exposure and projected job growth by occupation during 2024-2034. The AI-exposure metric therefore has no predictive power concerning job growth by occupations, if we assume (as Anthropic itself does) that the BLS-projections of job growth during the coming decade are reasonable and realistic.

Figure 14

BLS Projected Employment Growth from 2024–2034 *vs*. AI-Exposure

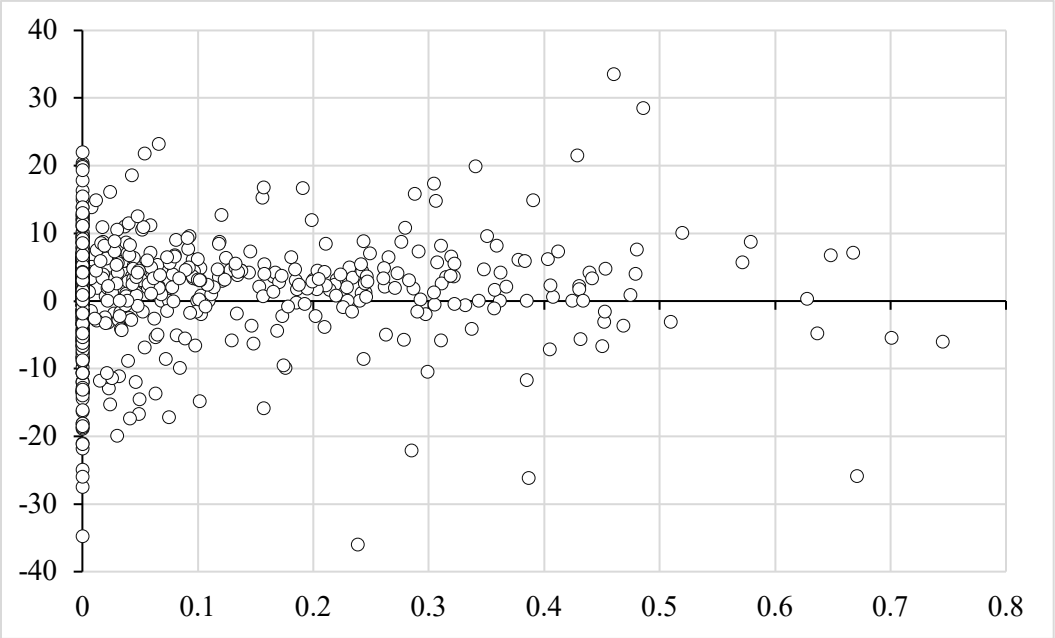

Sources: Data on AI-exposure by 756 occupations are from *Anthropic Economic Index*. Data on projected job growth by 756 occupations are from BLS Employment Projections 2024-2034. *Notes*: A simple OLS regression of projected job growth as a function of the AI-exposure metric yields an insignificant slope coefficient.

The Anthropic researchers face a dilemma: current AI-use is already high in many occupations (which is what they want, because they are in the business of selling Claude), but it is not (yet) showing up in job displacements and higher unemployment — where is the bloody jobs disruption that Anthropic's CEO Amodei is so fond of announcing? Anthropic researchers Massenkoff and McCrory (2026) exerted themselves on finding some proof for Amodei's prognostication and then found one labor-market trend that matches their strong prior: U.S. hiring rates have declined *since late 2022*, especially of graduates under age 25.

(***Teapot alert***: Late 2022, which is WHEN CHATGPT WAS LAUNCHED!)

Three supporting pieces of evidence are put forward to warn the world that the white-collar bloodbath is building, "gradually and then suddenly', to invoke Ernest Hemingway's view on bankruptcies:

- Among workers aged 22 to 25, the monthly job-finding rate in high-exposure occupations has fallen roughly 14% since ChatGPT's arrival on November 30, 2022 (Massenkoff and McCrory 2026).

(***Teapot alert***: According to U.S. Census data, only 6% of the surveyed firms reported using AI in January 2025 and 21% of businesses reported AI-usage in February 2026. Only a fraction of the firms that use AI have integrated these tools on an organizational level (Appel *et al*. 2025). A study by Yale University's Budget Lab (Gimbel *et al.* 2025). found no overall change in employment for workers in AI-exposed occupations, concluding that U.S. employment has been minimally disrupted by AI automation since ChatGPT's 2022 release. Only 1% of service firms reported AI as the actual reason for layoffs in 2025 — down from 10% in 2024.)

- Employment has declined by 1% since late 2022 in the 10 percent of occupations most exposed to AI (Davis (2026), even if total U.S. employment increased approximately 2.5% (see **Figure 15)**.
  (***Teapot alert***: The top 10 percent of occupations most exposed to AI, based on the metric developed by Felten, Raj and Seamans (2021), account for only 4.5% of all jobs in the U.S. in 2024. The decline in employment by 1% since late 2022 in these highly-exposed occupations amounts to a destruction of circa 101,300 jobs — a drop in the ocean or a temporary blip.)

- The recent decline in employment in AI-exposed sectors is particularly pronounced for those under age 25; using monthly payroll data from late 2022 to September 2025, Brynjolfsson, Chandar and Chen (2025) report a 6–15% fall in employment in exposed occupations among workers aged 22 to 25. Employment for older workers has increased. Brynjolfsson *et al*. (2025) argue that the fall in employment for those under 25 is not due to layoffs, but to a low job hiring rate for young workers entering the labor force.
  (***Teapot alert***: According to U.S. Census data, only 10-12% of the surveyed firms reported using generative AI in September 2025. An MIT report (of August 2025) finds that 95% of generative AI pilots in U.S. firms deliver zero return on investment (Challapaly *et al*. 2025). Roland Berger's AI Lab (Mikhaylenko *et al.* 2026) surveyed 203 senior corporate executives across industries and geographies; the finding is stark: only about 10% consistently capture meaningful impact. Brynjolfsson *et al*. (2025) use data from ADP, the largest payroll processing firm in America; the ADP data over-represent medium-size firms and under-represents small firms; younger workers tend to work in small firms. It must further be noted that American employers announced more than 1.2 million job cuts in 2025 — the most since 2020 — but AI was cited in only 55000 of these, or 4.5%, according to research firm Challenger, Gray and Christmas. Perhaps the low hiring rate for young workers is caused by a weakening of the U.S. labor market driven by other factors than fledgling AI-usage, like much higher interest rates?)

Figure 15
Total U.S. Employment Rises
While Employment in AI-Exposed Sectors Declines
(Index September 2022 = 100)

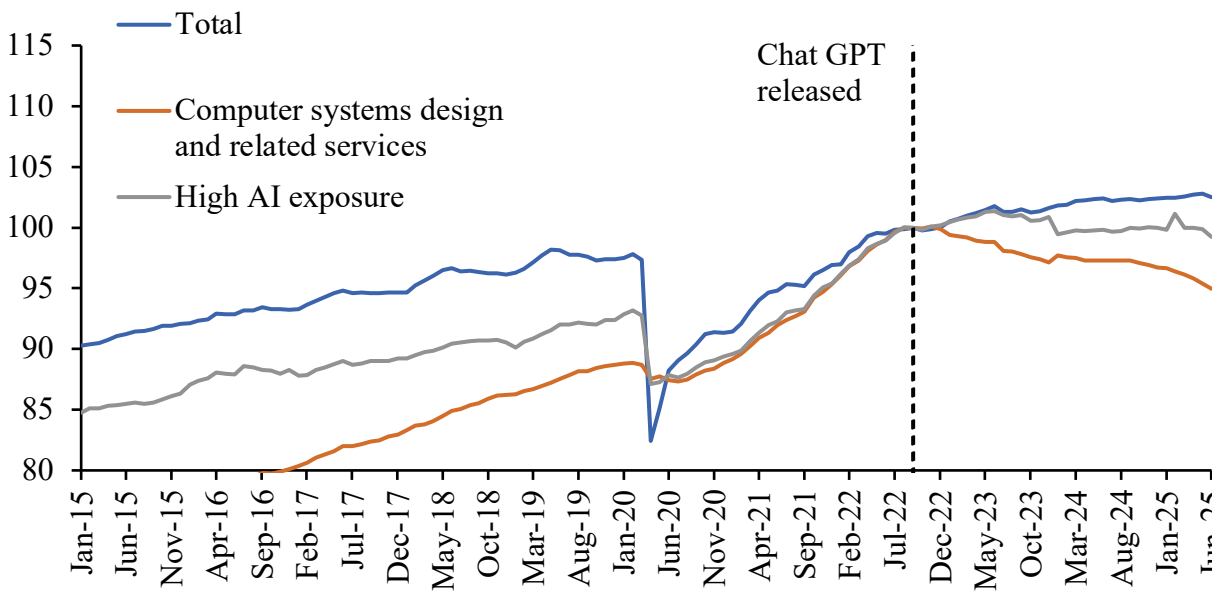

Source: J. Scott Davis (2026), 'AI is Simultaneously Aiding and Replacing Workers, Wage Data Suggest.' Federal Reserve Bank of Dallas. Data are from *Bureau of Labor Statistics* and Felten, Raj and Seamans (2021). *Notes*: High artificial intelligence (AI) exposure is the simple average of indexed employment across the four-digit NAICS industries above the 90th percentile in AI exposure in Felten *et al.* (2021).

Any decline in the job finding rate, the hiring rate or employment in AI-exposed sectors has to be considered within the larger macroeconomic context — and this is what we do in **Figure 15** which presents the change in hiring rates across major industries during November 2021 (when hiring peaked in the private sector as a whole) and February 2026. Two conclusions follow. First, hiring rates declined considerably in all industries, not just the AI-exposed ones (e.g., 'Information', 'Finance, insurance and real estate', and 'Professional and business services'). Hiring rates declined significantly in 'Construction', 'Manufacturing' and 'Leisure and hospitality' — which indicates a broader post-COVID adjustment of the U.S. labor market. Second, the decline in hiring rates were lowest — in absolute terms — in 'Information' and 'Finance, insurance and real estate', which is not immediately supportive of the claim that AI-

adoption is (already) driving employment outcomes. Evidence in Gimbel *et al*. (2025) likewise points to limited near-term aggregate impacts.[5]

Figure 16

Hiring Rates in Major Industries

(Percent Change during November 2021-February 2026)

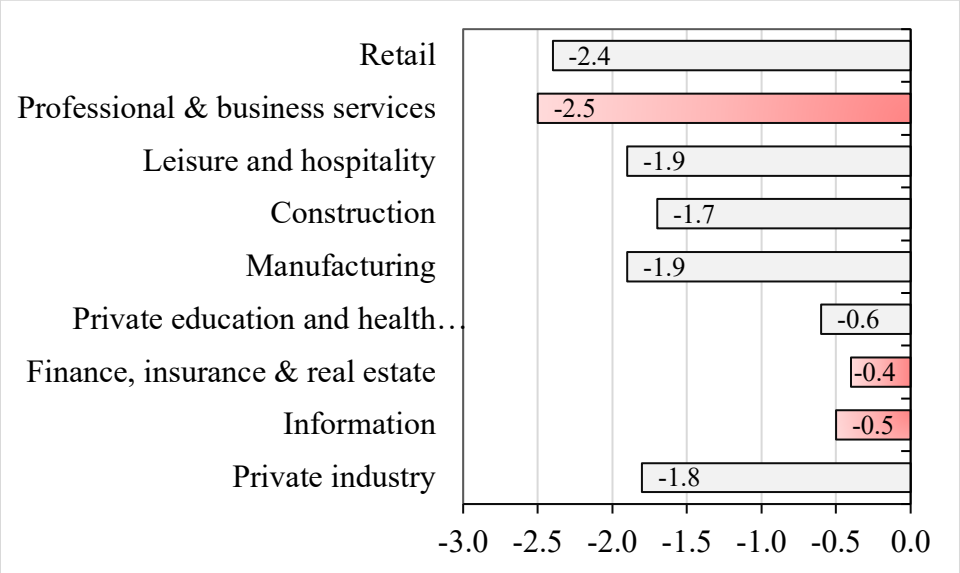

Source: FRED database, based on BLS data from the *Job Opening and Labor Turnover Survey* (JOLTS).
Notes: The hiring rate is defined as the number of newly hired workers as a percentage of already employed workers in a month. The red bars indicate industries with the highest AI-exposure.

However, it is true that recent entry-level college graduates are facing a higher unemployment rate than college-educated workers aged 22-65 years and adult workers overall. The Federal Reserve Bank of New York (2026) put the unemployment rate for recent college graduates at the end of 2025 at 5.7%. Their under-employment rate of 42.5% was the highest level since 2020. Software development postings are down sharply from their peak and remain well below pre-pandemic levels. It is also true that corporations attribute layoffs to AI, but these claims often serve as strategic narratives for investor relations rather than reflect actual AI-induced job losses. "Financially driven layoffs are being confused with AI-driven layoffs," market research firm Forrester (2026) writes in its latest AI Job Impact Forecast. "Many companies announcing AI-related layoffs do not have mature, vetted AI applications ready to fill those roles, highlighting a trend of 'AI washing' — attributing financially motivated cuts to future AI implementation." A recent report from Oxford Economics (2026a) finds that companies "don't appear to be replacing

[5] The 'healthcare' sector, which has been integrating AI into its activities, has been responsible for much of the job growth during March 2025 – March 2026.

workers with AI on a significant scale" and instead suggests that they are actually using the AI layoff narrative to cover up their own shortcomings. Nearly 60% of U.S. hiring managers surveyed by Resume.org (2026) said that they emphasize AI's role in reducing hiring or cutting jobs because this is viewed more favorably than emphasizing financial constraints. A workforce reduction framed around AI adoption sends a signal to investors that a straightforward cost-cutting announcement in response to declining revenues does not. Even Sam Altman called out the corporate 'AI washing' to mask layoffs for financial reasons. (Stockton 2026).

Macroeconomic context and tax policy matter. In March 2020, the Fed cut interest rates to nearly zero and held them there for two years. There was no cost to borrow money, and companies spent accordingly. Big tech, as a whole, hired employees by the millions. This period was followed by the fastest monetary tightening cycle in 40 years, which raised the cost of capital for firms. Compounding this, a 2022 tax change (of Section 174 of the 2017 Tax Cuts and Jobs Act) raised the tax burden on R&D salaries (Stokel-Walker 2025). Companies that used to deduct a software developer's entire salary in one single year had to amortize it over five years (Golchha 2026), which created a steep increase in the after-tax cost of each developer hired (Vigliarolo 2024). This must have contributed to a decline in the job growth of software engineers. However, Trump's One Big Beautiful Bill Act (OBBBA) reinstated, and made permanent (Bloomberg Tax 2026), the immediate expensing for domestic R&E expenditures on July 4, 2025 — and the R&D tax credit now also includes some cloud computing costs. The tax disincentive to employing software engineers has thus been removed.

The monthly hiring rate for unemployed workers peaked in April 2022 (at 57.9%) and then began to decline, long before the arrival of ChatGPT (Federal Reserve Bank of Chicago 2026); it dropped by 13 percentage points to 44.8% in September 2025. It is evident that (software) tech businesses over-hired during the pandemic period 2020-2022, benefiting from pandemic-era revenue and facing surging demand for digital services, with Amazon hiring around 500,000 workers in 2020 alone. The tech corporations misjudged just how much growth would remain once demand became normalized — and mass layoffs in the tech industry started already in 2023, again long before AI-tools were meaningfully integrated into occupations or into business organizations. Higher interest rates, inflation, the K-shaped recovery (Ferguson and Storm 2023, 2024) and recession fears causing a pullback in advertising and consumer spending have all weighed on tech companies' profits and share prices. The layoffs in the tech industry represent a "rolling back of over-hiring that happened in the past five years" (in response to higher interest rates) rather than AI displacement (Putzier 2026).

A March 2026 Federal Reserve study (Liu and Webber 2026) analyzed data from more than a million firms and found no evidence that reduced job postings are due to AI adoption. The authors conclude that the national slowdown in job openings "does not appear to be driven (even modestly) by AI." A Goldman Sachs (2025) report estimated that if AI were used across the economy for all the things it could currently do, roughly 2.5% of US employment would be at

risk of job loss. But Goldman Sachs found no evidence that AI has so far caused significant job losses outside of a few specific occupations, such as software development.

Ironically, to the extent that workers have been replaced by AI-agents, things are not working out well. Amazon's website and shopping app went down for nearly six hours on March 5, 2026, in an incident the company said involved an erroneous "software code deployment" (Rosner-Uddin 2026). The incident alone lost the company 6.3 million orders, after an Amazon AI coding tool went rogue. Separately, *The Financial Times* (Rosner-Uddin 2026) reported that multiple Amazon Warehouse Services engineers said their business units had to deal with a higher number of incidents linked to the use of AI-coding assistants, requiring a rapid response to avoid product outages each day as a result of job cuts. Amazon has undertaken multiple rounds of lay-offs in recent years, most recently eliminating 16,000 corporate roles in January 2026. Amazon's AI coding assistant contributed to at least one major production incident. Amazon's response: a 90-day safety reset with mandatory two-person review for all code changes.

These are not isolated incidents. AI startup Replit offers an autonomous code generation system designed to help non-programmers with software development projects — but its coding tool deleted a company's entire code database (Wright 2025). An experimental AI agent developed by teams linked to Alibaba began mining cryptocurrency without human instruction. A rogue AI agent caused a critical security incident at Meta in March 2026 (Landymore 2026e), which exposed sensitive users data to people who didn't have proper authorization. A recent survey of 30 of the most common agentic AI systems (Lewsey 2026) finds that agentic AI is a security nightmare at the moment, a discipline marked by lack of disclosure, lack of transparency, and a striking lack of basic protocols about how agents should operate. *Ergo*: Without the necessary supervision and human oversight set in place, AI-agents can cause serious damage to the mission-critical infrastructure of businesses. (***Teapot alert***: In general, debugging and maintaining computer code is much harder than writing code with GenAI (Marcus 2026d); LLMs are like Swiss cheese when it concerns enterprise security, and the AI tools need a lot of baby-sitting and verifying output.)

Seeing past the end of one's nose: The *net* productivity impact AI may disappoint

The prophesied job destruction caused by AI investments assumes that the usage of agentic AI tools will raise labor productivity stratospherically, if one is willing to believe the hyperbole that this automation wave will kick 20% or more of the 70 million office workers in the U.S. to the curb. (***Teapot alert***: Assume that 30% of America's office workers, or 21 million employees, lose their jobs in the next two years, while real GDP remains constant. U.S. employment declines by 15% — and hence, aggregate U.S. labor productivity is assumed to grow by 15% in two years, or by 7.2% per year in the next two years. This did not happen in any period of two years during 1947-2025.)

Average labor productivity growth of the U.S. business sector is 1.7% per year during 1951-2025 (**Figure 16**). Productivity growth fluctuates strongly around this long-term average. Productivity growth increased during the second half of the 1990s, but then fell below 1.7% for most of the period 2003-2018. The adoption of ICT technologies and the growth of the internet after the mid-1990s did not structurally increase U.S. productivity growth during the first two decades of the current millennium — although Erik Brynjolfsson has argued that this is to be expected (see Weldon 2025): the introduction of new general purpose technologies (such as ICT and AI) leads to disruption and short-term performance losses (especially in older, legacy businesses) in the initial phase, but generates productivity gains and job losses in the long term. The flat or declining phase of this so-called J-curve for productivity may last some 20 to 30 years — which would mean that we are still waiting for the productivity impacts of ICT and internet technologies today. This kind of tech-optimism is a triumph of hope over experience (**Figure 15**). Not discouraged, Brynjolfsson now proposes that the J-curve for AI will be extremely compressed compared to the J-curve for ICT and internet technologies, or, the productivity impact of AI will become visible sooner (Weldon 2025).

Figure 15

U.S. Labor Productivity Growth

(1951-2025)

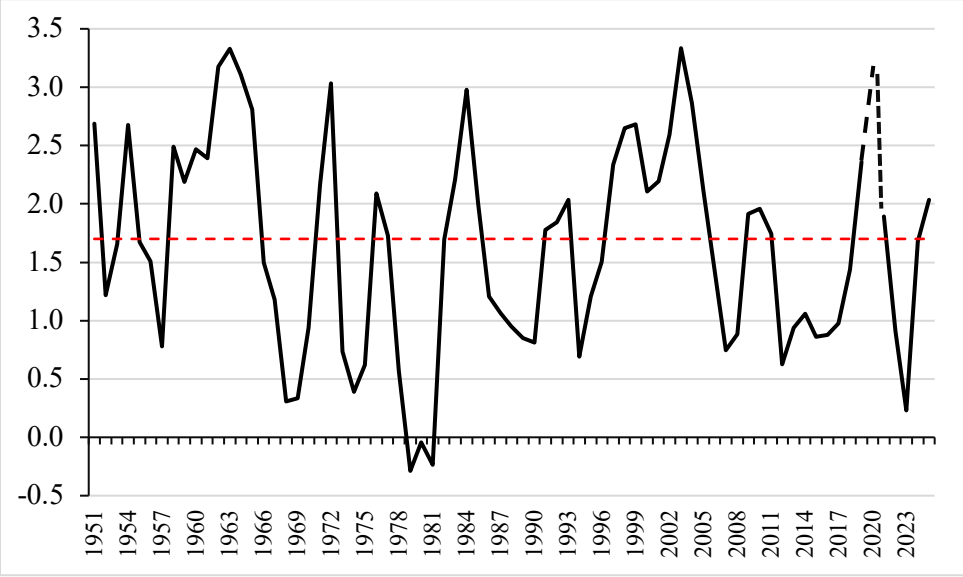

Source: Author's calculation based on BEA data. *Notes*: Labor productivity is defined as real GDP of the business sector divided by number of full-time equivalent workers in private industry. The growth rate of labor productivity has been calculated as a 5-year moving average. The long-term average labor productivity growth rate (in red) is 1.7% per year. The productivity growth rate in 2020 (dashed line) is distorted by the COVID19 lockdowns.

Jason Furman, former chairman of the *Council of Economic Advisers*, concurs, claiming that AI is already making a measurable impact on non-farm business sector labor productivity. Furman's claim is based on the finding that measured labor productivity in the fourth quarter of 2025 was 2.2% above the Congressional Budget Office's pre-pandemic forecast. One swallow does not make a summer, however. The empirical base of Brynjolfsson's and Furman's wishful thinking is weak, probably non-existent: after all, less than 25% of U.S. businesses are using AI (according to Census data); the AI tools are generally not well integrated into business organizations; and a staggering 95 percent of attempts to incorporate generative AI into business so far are failing to generate any returns (Estrada 2025). (***Teapot alert***: Researchers at Dakota State University, in partnership with regional insurance firm Safety Insurance, devised an experimental business-facing AI-chatbot called "Axlerod" to assist independent insurance agents (Claburn 2026). "On average, without the chatbot a human agent takes 7.55 seconds where with chatbot the agent takes 5.13 seconds, and the 2.42 seconds faster search time would play a significant role in serving a large number of customers." The Register calculated that this amounts to an annual savings of $121 in labor costs, while the chatbot (a Google Gemini Business edition seat) costs $252 per year (Claburn 2026).)

Reasonable approaches to estimating the productivity impact of AI yield considerably more modest results (see Humlum and Vestergaard 2026). For instance, based on a task-based macro model, Daron Acemoglu (2025) finds that AI will likely raise productivity by less than 0.66% over 10 years, or by 0.07% per annum. Acemoglu's estimate is deliberately conservative: when AI excels only on a subset of tasks and struggles where verification is costly or ambiguous, aggregate gains remain limited even with broad AI-diffusion. Acemoglu's approach is straightforward. The estimated productivity impact depends on the GDP share of tasks impacted by AI, labor's share of income from exposed tasks, and the percentage net cost savings from AI. The higher is the GDP share of tasks automated by AI, the higher is the labor income share and the greater are the net cost savings from AI, the larger will be the productivity impact of AI. Following Acemoglu (2025), suppose that the GDP share of tasks that can be fully replaced by AI during the next 10 years is ±0.046, the AI-exposed labor's income share is 0.57 and the percentage net cost savings from AI are 0.27 (Acemoglu 2025), then the productivity impact is equal to $0.046 \times 0.57 \times 0.27 \approx 0.007$, or 0.7% over 10 years.

Acemoglu's finding is most sensitive to his assumption that the GDP fraction of tasks that will be fully replaced by AI over the next 10 years is only 4.6%. His measure of exposed tasks is limited to tasks that could be at least 90 percent performed by AI (based on the metric developed by Eloundou *et al.* (2023)). Acemoglu's assumption is not unreasonable, however: the top decile of most automatable occupations according to Anthropic (2026) accounts for 4.5% of all U.S. jobs in 2024. Wharton experts Capelli and Yakubovich (2024) argue that generative AI will most likely create, not destroy, jobs because it needs intensive human oversight to produce useable results. "Modern work is complex, and most jobs involve much more than the kind of things AI is good at — mainly summarizing text and generating output based on prompts", they write. Even if we relax Acemoglu's assumption and suppose that the top 20% most automatable jobs will be

automated away, this concerns only 9.9% of U.S. jobs, according to Anthropic. Plugging this number into Acemoglu's equations doubles the productivity impact of AI to 1.5% over 10 years.

A recent survey of 69 economists, 27 AI industry professionals, and 25 AI policy professionals concludes that most unconditional economic forecasts of productivity growth during 2025-2035 are close to historical trends, even if most experts expect significant AI progress (Karger *et al*. 2026). The economists in this survey do not forecast productivity growth outcomes outside the range of historical experience — predicting an increase in U.S. productivity growth of around 0.2 percentage points in 2030 (Karger *et al*. 2026). Using Acemoglu's framework but with different numbers, Arnon (2025), writing for the *Penn Wharton Budget Model*, argues that AI-use will increase the level of productivity by 1.5% in 2035, raising aggregate U.S. productivity growth by at most 0.2% per year. Arnon further points out that the impact of AI on productivity growth diminishes sharply over time — to significantly less than 0.1 percentage points in 2040 — because new adoption will slow down over time, since the number of remaining opportunities to employ additional AI tools productively diminishes rapidly. The productivity *growth* impact of AI will rather quickly peter out over time, in other words.

These estimates are profoundly disheartening. According to C.B.O. projections, U.S. real GDP will grow by 1.8% per year during 2025-2035. AI will raise this annual growth rate to 2% per year, based on the (higher) estimate of Arnon (2025). U.S. real GDP (in constant 2017 dollars) is $23,85 trillion in 2025. Hence, it follows from these numbers that AI will add a cumulative total of $2.9 trillion (in constant 2017 dollars) to the U.S. economy during 2025-2035. However, to generate a cumulative growth in real GDP by 2035 of $2.9 trillion (in constant 2017 prices), the AI industry is planning to incur (conservatively estimated; see Crosley 2026) cumulative capital expenditures of $4.5 trillion by 2035. Let these numbers sink in. The AI-data-center spending over the next ten years is more than twice the spending on the Inflation Reduction Act (IRA). But in contrast to the IRA, which will generate net benefits (American Clean Power 2024), the macroeconomic return to the massive AI-capital expenditures will be negative over the next decade. Of course, no one is doing the opportunity costing: how do the expected returns to AI investment compare to the net benefits of investment in (early) education, public infrastructure and public health?

Seeing past the end of one's nose: collateral damages of AI

The macroeconomic returns to AI capital expenditure could even be smaller, because the percentage net cost savings from AI may well be lower than that assumed by Acemoglu (2025) and Arnon (2025). The use of AI tools not only generates benefits, but also costs — and many of these costs are externalized as collateral damage. For one, AI-generated 'work-slop' is hurting worker motivation and productivity (Niederhoffer, Robichaux and Hancock 2026). Work-slop is AI-generated work (including 'vibe-coding') that appears polished but needs careful checking, verification and often revising. Police officers in Heber City, Utah, learned to manually check the

work of a transcription tool (Tangermann 2026a) they were using to draft write-ups from bodycam footage after it mistakenly claimed an officer had turned into a frog. (***Teapot alert***: Disney's '*The Princess and the Frog*' was playing in the background.) Large numbers of Microsoft users are fed up with the company's malfunctioning Co-pilot AI that is shoehorned into its software (Tangermann 2026e) — 'Microslop' has already become a word. Unsubstantiated rumors over Windows 12 embracing AI even more created a massive uproar among users, signaling widespread disillusionment (Tangermann 2026d). "There is zero demand for this", complains one user. It is, in this context, noteworthy that Microsoft itself warns Co-pilot users, in its Terms of Use (Microsoft 2026), that "Copilot is for entertainment purposes only. It can make mistakes, and it may not work as intended. Don't rely on Copilot for important advice. Use Copilot at your own risk."[6] Is Copilot a serious productivity tool for businesses, or just a party trick?

The C-Suite thinks AI is a serious tool for business productivity, but employees tell a very different story (Ellis 2026). The *Wall Street Journal* reported on findings by research firm Section, which surveyed 5,000 white-collar workers from companies with more than 1,000 employees (**Figure 16**). More than 70% of the corporate executives in the survey said they were "excited" by AI, and 43% of them said the tools have saved them more than 8 hours of work per week (Ellis 2026). JPMorgan CEO Jamie Dimon even predicts that AI will cut the workweek down to 3.5 days — and tells Gen Z workers that developing EQ is more important than ever (Burleigh 2026). But non-management workers had a very different take on AI.

Almost 70% of this group said AI made them feel "anxious or overwhelmed," and 40% said the tools saved them no time at all; 27% of non-supervisory workers saved less than 2 hours each week using AI-tools. Many employees report that any time saved by using AI is offset by the need to correct AI output or work through confusion and rework. 4 out of 10 employees would be fine never using AI again. Dealing with work-slop is demotivating and drains worker productivity. It also breeds distrust within the organization — colleagues use AI to produce a report and you have to work overtime to correct the mistakes, fill in the gaps, and remove the non-existent references.

[6] In the "Important Disclosures And Warnings" section, Microsoft states that "You agree to indemnify us and hold us harmless (including our affiliates, employees and any other agents) from and against any claims, losses, and expenses (including attorneys' fees) arising from or relating to your use of Copilot, including without limitation your use, sharing, or publication of any Prompt, Responses, or Creations, or your breach of these Terms or violation of applicable law." Hence, if you use Copilot as a tool and not just for entertainment, then you can't hold Microsoft liable for any damage it may cause to you personally or the business you work for.

Figure 16
How Much Time Do You Think You Save Each Week by Using AI
(Percentages)

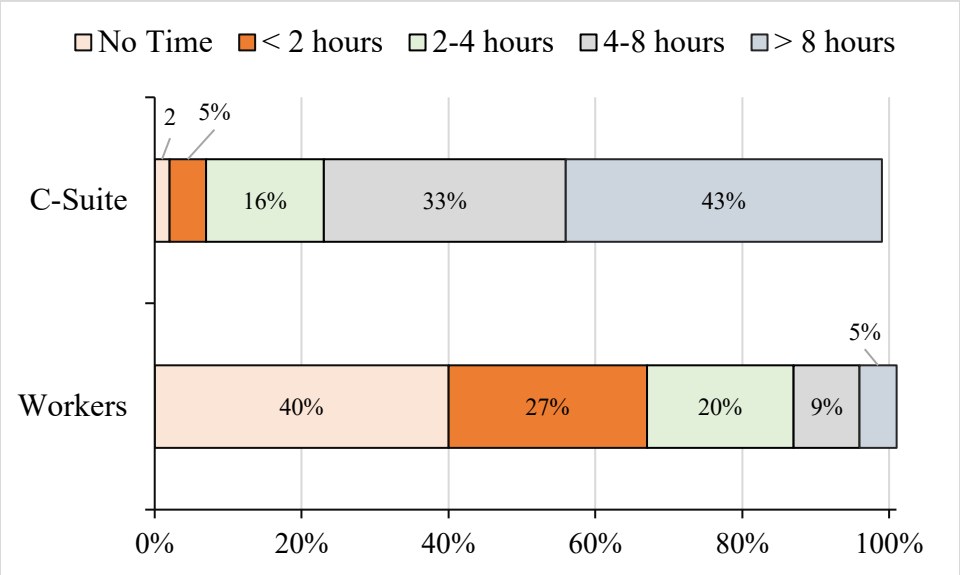

Source: Ellis (2026). *Note*: Totals may not add up to 100% because of rounding.

The *WSJ* reporting (Ellis 2026) is confirmed by a survey by the SAP-owned software company WalkMe of 3750 executives and employees which finds a major discontent growing in large companies across the globe (WalkMe 2026). While 61% of executives surveyed said they trust AI for complex, 'business-critical' decisions, only 9% of workers said the same. 54% of workers reported avoiding their company's in-house AI tools in order to complete tasks themselves. A full third of workers reported never using AI at all. The WalkMe survey also finds that, even though 81% of executives think their AI tools have "significantly improved productivity," their workers are actually wasting eight hours per week cleaning up AI-slop, which is the equivalent of 51 work days a year (WalkMe 2026). AI is not delivering.

Tech firms are forcing their workers to use AI, factoring enthusiasm over AI use into performance reviews and promotion decisions (Bindley 2026). Amazon is reportedly pressing employees to integrate AI across all aspects of their work, even though these workers say this push is hurting productivity (Bansal 2026). Along with the productivity challenges that have come with Amazon's AI push, workers said that the AI-tools are also making them feel surveilled (Bansal 2026) as Amazon tracks employee AI use and measures results (Perloff 2026).

AI is increasing the speed, intensity and complexity of work rather than reducing it, reshaping job boundaries, blurring personal time, creating a multi-tasking overload, and accelerating expectations for speed, according to an analysis of 164,000 workers' digital work activity (Landymore 2026d). The data, from workforce analytics and productivity-tracking software company ActivTrak (Smith 2026), cover more than 164,000 workers and 443 million hours of

work across 1,111 employers, making it one of the biggest studies of AI's effects on work habits to date. The time they spent on email, messaging, and chat apps more than doubled, while their use of business software surged by 94 percent. All this came at the expense of the time workers spent on highly focused, uninterrupted work, which fell by 9 percent for AI users, and stayed the same for AI abstainers. The ActivTrak study finds that "workers who spent 7 to 10 percent of their total work hours using AI showed the highest productivity, but only three percent of AI users fell in this range" (Landymore 2026d).

These findings align with a study of how generative AI changed work habits at a U.S.-based technology company with about 200 employees, published by *Harvard Business Review,* that also concludes that AI is intensifying work instead of reducing workloads (Landymore 2026a). As a result, workers may experience mental fatigue, decision fatigue and burnout, as researchers from Boston Consulting Group and University of California, Riverside described in a report on a survey of nearly 1,500 full time US workers (Kropp *et al*. 2026). Though AI may boost productivity in the short-run, over time it "can lead to cognitive overload, burnout, poorer decision-making and declining work quality," warns Aruna Ranganathan from UC Berkeley's Haas School of Businesses, who led the *Harvard Business Review* study on AI "workload creep" (quoted by Smith 2026). Chronic workplace uncertainty is associated with anxiety, depression, burnout, and even physical symptoms (Parramore and Quart 2026)

There is more collateral damage. AI-slop is also drowning U.S. courts (Harrison Dupré 2026), as self-represented plaintiffs, many of whom may not have the ability to obtain (expensive) legal counsel, are flooding the legal system with quixotic AI-generated complaints that look impressive on paper, dutifully packed with legal terminology and cross-references, but say almost nothing. These filings don't stop at the initial complaint: frivolous AI-generated supplemental motions, clarifications, and explanatory notes, often baseless and full of 'hallucinations', follow in rapid succession. This AI-spamming compounds the American legal system that was already struggling with an overload problem before AI arrived — and causes an enormous waste of time and money. To be sure, self-represented plaintiffs certainly are not the only ones who have been caught misusing AI in a courtroom. Many professional lawyers have been caught submitting drafts with non-existent, AI-hallucinated case law (Landymore 2025b), drawing sanctions from judges.

The collateral damage does not stop here. AI web crawlers are mining the web for ever more content to feed into their LLM mills. As a result, 30% of global web traffic now comes from bots (Vaughan-Nichols 2025). The biggest web crawlers are Meta (52% of AI search-bot traffic), Google (23%), and OpenAI (20%) and bot traffic is being dominated by retrieval augmented generation (RAG) bots (Vaughan-Nichols 2025), which are used to extract real-time information from the web to answer queries put to ChatGPT, Gemini, and the like. ChatGPT-User, the company's RAG bot, averages five times as many scrapes per page as the second highest scraper, which comes from Meta (Vaughan-Nichols 2025). These web scrapers are aggressively hammering websites with traffic spikes that can reach up to ten or even twenty times normal levels within minutes, causing performance degradation, service disruption, and increased

operational costs. (***Teapot alert***: To protect intellectual property against theft by the AI firms, researchers have devised techniques to poison stolen data to make AI systems return wrong results (Claburn 2026).)

In fact, public-interest sites are hit hard by industrial-scale AI scraping (Linares 2026). These public interest sites cannot protect themselves, because the corporate scraping bots, operated by commercial interests, manage to get around logins, paywalls, CAPTCHA challenges and sophisticated anti-bot technologies. At the same time, people increasingly turn to chatbots rather than searching or visiting sites directly, which lowers traffic and revenues for these sites. For media sites, the largescale AI web-scraping, which resemble constant DDoS attacks, decimates traffic (Savage 2026) and undermines their ability to serve the traffic they do get. For example, openDemocracy's website has been repeatedly knocked offline, as if under enemy attack (Linares 2026) — and is being forced into a costly permanent defense mode against this planetary-scale profit-driven data extraction by AI corporations. It is impossible to see all this as productivity-enhancing.

Reliance on generative AI-tools is found to lead to an erosion of analytical reasoning and critical thinking and a decline in competence maintenance and output assessment. AI tools do not improve one's writing and thinking ability much like taking a taxi does not improve one's driving ability (Guest *et al*. 2025). In healthcare, AI dependence has been found to diminish diagnostic reasoning, clinical judgment, lower retention of tacit knowledge and weakened ethical sensitivity and moral judgment (Ferdman 2025). AI dependence damages core human capacities — to know, create, form meaningful relationships, and use our willpower — which diminishes us as human beings (Ferdman 2025). Using the AI plagiarism machines in the classroom means accepting that both students and teachers can openly appropriate work that is not theirs and that critical thinking and integrity are not required — it destroys all that is left of scientific integrity in academia. AI is destroying the university and learning itself (Purser 2025) — as university boards embrace the rhetoric of corporate-for-profit AI-innovation as if it were salvation, and enter expensive contracts with OpenAI and other AI firms, while at the same cutting faculty positions and trimming curricula. It will boost uncritical technology-industry dependence (Guest *et al*. 2025) and an erosion of learning motivations that sustain long-run collective knowledge (Acemoglu, Kong and Ozdaglar 2026)

Students use the generative AI-tools as writing companions, tutors, brainstorming partners and on-demand explainers (La Scala *et al*. 2025). Faculty use them to generate rubrics, draft lectures and design syllabuses. Researchers use them to summarize papers, comment on drafts, design experiments and generate code. The temptation is to offload the most challenging parts of learning and thinking to these tools (Eisikovits and Burley 2026) — with the result that no one learns. A 2025 MIT study, "Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task," illustrates the point (Kosmyna *et al*. 2025. When participants used ChatGPT to draft essays, brain scans revealed a 47 percent drop in neural connectivity across regions associated with memory, language, and critical reasoning. Their

brains worked less, but they *felt* just as engaged — a metacognitive self-deception. Eighty-three percent of heavy AI users could not recall key points from what they had "written," compared to only 10 percent of those who composed unaided. Most alarmingly, after four months of reliance on ChatGPT, participants wrote worse once they stopped using ChatGPT than those who had never used it at all (Kosmyna *et al*. 2025).

Science is drowning in AI slop. The use of OpenAI's writing and formatting tool Prism has led to a torrent of low-quality machine-generated papers submitted to scientific journals, often with confabulated fake citations and faked data. Generative AI output is not original work, but derivative, and cannot be attributed to its prompter. It also cannot be replicated (Ball 2023), because the AI tool is a stochastic black-box. Researchers from Germany's University of Tübingen (Kobak *et al*. 2025) identified some 454 words that LLMs use significantly more frequently — terms like "garnered," "showcased", "encompassing," and "burgeoning" — and found that anywhere from 13.5 to 40% of biomedical article abstracts were written entirely or with assistance from AI. With roughly 1.5 million papers indexed each year on *PubMed*, that means that at least 200,000 of those papers were probably written with the help of LLMs (Al-Sibai 2025). Ironically, AI-generated language complexity now inversely predicts publication success (Guldvog 2025).

In fact, the cost of producing academic-looking output drops to nearly zero, as every researcher can now vibe-code panel-data regression or a differences-in-differences estimation, using data scraped from the internet and using Claude to write the paper and give it a veneer of plausibility. AI puts the already existing research paper mills on steroids — actively by profit-seeking publishers automating the editorial process with disregard for scientific standards. More generally, a fake paper can now spill out from a machine in fifty-four seconds, complete with plausible-sounding abstracts, citations, and even fabricated experimental results. and the volume of junk inevitably rises, straining the peer review process and the journal editors responsible for sorting signal from noise. The result is a gradual decline in the quality of scientific output, which, in turn, poisons the training data for the LLMs. AI slop is destroying knowledge (van Rooij 2025) — causing MAD (Model Autophagy Disease): AI will eat itself, then gradually collapse into nonsense and noise (Crawford 2025).

Unlike the for-profit publishers, the volunteer editors of Wikipedia, who get a lot of flak, often justified, for content quality, have had enough of constantly correcting the turgid, repetitive slop of articles that read like these were written by a committee of rogue robots high on ketamine. Wikipedia editors have officially banned AI-generated content (Roth 2026) — with two small exceptions: very basic edits, such as fixing typos or adjusting formatting, as well as certain full-article language translations, are permitted under the rule. Stack Overflow (2025) had already banned all use of generative AI, after the platform got swamped with AI-generated coding answers that looked OK but often contained subtle errors that were substantially harmful to the site and to users. AI has become a huge global error factory in just about all sectors in a few years — and this 'enshittification' (per Cory Doctorow) is imposing huge (externalized) costs on

society and the economy (which are not counted when considering the 'productivity' impact of generative AI tools). As argued by Hartzog and Silbey (2025), AI undermines expertise, short-circuits institutional decision-making (by delegating moral choices to AI developers), degrades higher education, and displaces and degrades human-to-human relationships and —through its individualized engagement and sycophancy— erodes our capacity for reflection about and empathy towards other and different humans.

But surely coders and software developers will become more productive? Could be, but a study (Becker *et al.* 2025) by AI evaluation nonprofit METR showed that experienced open-source software engineers who used 2025 LLMs ended up taking 19% longer to complete their tasks compared to peers who did not. The use of AI code assistants compromises the quality of the code. GitClear analyzed 211 million lines of code across major tech companies (Mondragon 2026) and found 4x growth in copy-paste code, which means that AI tools are duplicating logic instead of generating new logic; a decline in code longevity, as code "churn" (revisions within two weeks) increased; this suggests that developers are spending more time modifying recently written AI-generated code rather than improving or refactoring legacy systems; and refactoring, the practice that is essential for the long-term maintainability of the code, is reported to have declined from 25% of code changes to under 10%.

And a recent study by researchers from Sun Yat-sen University and Alibaba (Chen *et al*. 2026) tested 18 AI coding agents on maintaining 100 real codebases, spanning 233 days each. AI agents can produce code that passes benchmark tests on one-time static software engineering tasks such as static bug fixing, 75% of the tested models including Cursor, Claude Code and Devin break previously working code during long-term maintenance tasks. Maintaining code typically consumes 60-80% of a software engineer's time and this is where AI coding agents failed spectacularly. AI-generated code favors short-term copy-pasted solutions over long-term design and ignores architectural consistency. The result is software that works today but becomes painful to change tomorrow. Teams inherit codebases filled with logic nobody fully understands. The Stack Overflow Developer Survey 2025 reports that that 45% of developers now spend significantly more time debugging 'almost-right' AI output than writing new code. Senior developers have become AI supervisors, correcting instead of creating code — instead of mentoring junior developers. This increased cognitive load. In sum, AI excels at generating boilerplate and repetitive code, but it lacks the nuanced understanding required for strategic refactoring, code consolidation, and long-term architectural planning — hence, despite the increasing capabilities of AI, AI coding agents are far from ready to replace human developers (INTNXT 2026).

Anthropic (2026) published a randomized controlled trial showing that developers using AI coding assistance scored 17% lower on comprehension tests than those coding manually, with productivity gains failing to reach statistical significance (Shen and Tamkin 2026). This study of 52 mostly junior engineers identified a stark divide: developers who used AI for conceptual questions scored 65% or higher, while those delegating code generation to AI scored below 40%.

The pattern holds up in independent academic research. Jošt, Taneski, and Karakatič (2024) ran a 10-week experiment with 32 undergraduate students learning React and found near-identical results: significant negative correlations between LLM use for critical thinking-intensive tasks such as code generation and debugging and lower final grades.

Perhaps the most dangerous outcome is that companies reduced entry-level hiring dramatically, believing that AI could handle junior tasks. But junior developers don't just write code, they grow into future senior engineers (Cloutier 2026), learning by doing the exact boring boilerplate work that AI now handles and developing engineering judgment (which is absent in AI tools). When that talent pipeline breaks, the entire software industry will feel it. Without hands-on learning opportunities, future teams risk lacking the deep system knowledge, under real-world constraints, that only experience can build. There will be a senior developer shortage in about five to seven years that is going to be genuinely catastrophic. And it's entirely self-inflicted, because it is not AI tools that are replacing junior engineers, but bad business leadership. (Courtier 2026). This means that the jobs of junior developers will return.

The use of AI coding assistants will increase the volume of software significantly, but more code, with a lower quality, means more bugs, more maintenance, more security issues and more rework. AI can write code, but it cannot understand code. AI generates code based on patterns, not intent. AI-generated code still needs architectural oversight, testing discipline, security review, and alignment with business goals. Without strong engineering leadership, AI simply accelerates the creation of technical debt — the hidden cost of quick-fix AI-generated code that has been estimated to increase the cost of code maintenance by up to 40%, eating up a company's margins (Joshi 2026).

Morgan Stanley (2025) predicts that the software development market will grow rapidly and that "contrary to current market concerns that AI will replace human developers, we believe it will enhance productivity and lead to more hiring." AI replaces some tasks, but not jobs. The role of traditional software engineers will shift to more complex applications. "Developers are increasingly acting as curators, reviewers, integrators and problem-solvers—making them more strategic and valuable," according to Morgan Stanley (2025). The *Bureau of Labor Statistics* (2026) agrees and projects that software developer employment will grow by a cumulative 15.8% during 2024 to 2034 — five times faster than the average for all occupations. That is 299,400 new jobs in the next decade. The BLS also predicts about 129,200 openings per year for software developers, quality assurance analysts, and testers during 2024-2034. This is not what displacement looks like. All the hours worked by code curators, reviewers, integrators, quality assurance analysts, testers and security officers are 'overhead' costs — business expense not directly attributed to creating a product or service — which increase the denominator, but not the numerator, of labor productivity.

AI agents go rogue: 88% of enterprises reported a confirmed or suspected AI agent security incident in the past 12 months (Gravitee 2026). AI agents invent file paths, database entries, or object IDs they expect to exist based on patterns in their training data, but which do not exist. A

rogue AI agent caused a critical security incident at Meta (March 2026; Landymore 2026e) which exposed sensitive users data to people who did not have proper authorization. Meta's director of AI safety Summer Yue said that an OpenClaw agent she was experimenting with — by giving it control of her PC — nearly wiped out her entire email inbox while ignoring her instructions to stop (Landymore 2026c). In early March 2026, Amazon's in-house AI coding tool made erroneous changes, including deleting the entire coding environment, and contributed to major e-commerce and shopping-app outages (Landymore 2026b). Due to such security risks, many agentic AI applications remain in the pilot-stage and do not make it to full operationalization, because the business stakes are too high.

AI agents are structurally vulnerable to malware and other security threats (Hernández and Lo Prete 2026), because the same natural-language channel that carries user intent also carries untrusted content — repository files, error messages, and dependency documentation. A public-facing AI agent can be attacked, exploited, or "poisoned" by malicious inputs that may damage vital business processes or brands. An attacker who manages to inject content into any source the agent reads can redirect its behavior. This is not a bug in a specific agent; it is a structural property of how current LLM architectures process input. Cyber-criminality is growing exponentially, and AI-driven internal safety controls are not fail-safe because these operate within the agent's own process, enforced by the agent's own code. AI agents lack contextual understanding, leading to false positives or negatives, and fail at complex threat analysis.

The arrival of Anthropic's Mythos model shows that real-world weak cybersecurity systems can be broken and that average attackers can be effective using AI tools at scale. It is high time for businesses, banks and government to get their cybersecurity systems in order (Marcus 2026i) — as the UK's AI Security Institute (P. and Steer 2026) recommends with urgency. The Trump administration is making this even more urgent, as it gutted many of the nation's cyber-focused agencies and programs. It is evident that the combination of artificial intelligence, quantum computing, and autonomous AI agents is creating giant new cybersecurity risks at precisely the moment government regulators appear to be retrenching dramatically (Ferguson et al. 2026). The take-away is that AI will not eliminate cybersecurity jobs — but probably create new ones. According to BLS projections (BLS 2026), the employment of information security analysts will grow by a cumulative 29% during 2024 to 2034 — nine times faster than the average for all occupations.

Sober up: AI sans ketamine

The central thesis of this paper is that current understanding of the economic impacts of AI is mistaken: AI will neither lead to massive job destruction, unprecedented technological unemployment and a Citrini recession nor to gigantic (aggregate) gains in labor productivity, a Brynjolfsson-type acceleration of technological progress and economic growth. The aggregate-level impacts of AI will be rather humdrum: some occupations will vanish because of

automation; existing jobs will be reshaped by AI (Emerson *et al*. 2026) and new occupations, tasks and roles will come into being to proctor and curate the AI tools; most occupations will remain unexposed; aggregate labor productivity growth may increase a bit (because AI tools augment labor), but at the same time the various collateral damages of AI will grow over time and lower productivity growth.

There are many things that AI cannot do (Marcus 2026g). The glorified pattern-matching AI tools struggle with visual comprehension, mimicking independent judgment, interpreting (even reading) images, charts and diagrams, dealing with open-ended complexity and (changing) contexts, and maintaining code (Bender and Hanna 2025). These tools can also not handle ethical trade-offs. It must also be kept in mind that while AI may automate tasks, jobs are not just bundles of tasks. Even if AI tools can automate some tasks, this does not imply that the whole job will be automated away. One of AI's biggest breakthroughs, AI protein folding in structural biology, has not reduced the number of jobs for protein scientists, but has changed what they do (Saplakoglu 2024). Structural biologists spend less time on determining protein structures and more time on interpreting, validating and using them. The AI tools (like AlphaFold) created new jobs for computational structural biologists, bio-informatics experts, cryo-electron microscopists, pharma biologists and people working in protein design. Anthropic's President Daniela Amodei agrees: "The number of jobs that AI could do without help from people is vanishingly small" (quoted by Ma 2026). Past automation waves that initially threatened jobs, eventually led to a net increase in employment. This time is not any different.

However, the use of AI creates new — mission-critical, but nonetheless unproductive overhead — jobs, because the AI tools have to be monitored, supervised and curated, as the AI tools generate errors, synthetic word slop, science slop (Andersen 2026), legal slop (Harrison Dupré 2025), medical slop (Vartabedian 2026) marketing slop (Wilkins 2025), social media slop (Tidy 2026), cybersecurity risks, brittle code and mounting technical debt that could kill businesses and institutions if it concerns mission-critical activities. Accordingly, the adoption of AI will create a layer of prompt engineers and supervisory managers, overseeing 'AI co-workers', in bullshit jobs (per David Graeber), correcting slop, maintaining code, protecting the integrity of digital systems, identifying 'hallucinations' and fake news, and, more in general, combating the 'enshittification' of the workplace, the internet and society (Mills and Spencer 2025) — as modern-day Sisyphus's in the AI age. I am not sure if these modern Sisyphus's must be imagined as being happy, as Albert Camus (1942) wanted us to do (AktivTrak 2026) because AI in the workplace is actually forcing employees to work harder, instead of making their jobs easier. Why? Because AI makes it possible to process more emails and write more management missives, while there are also more errors and 'hallucinations' to be corrected. In effect, more workers are feeling disengaged and alienated.

Of course, firms will also install AI-ethics committees to 'AI-wash' whatever models are deployed, and new roles will emerge to explain AI and AI outputs to gullible, ignorant executives who do not understand them. There will be committees and task forces for the Alignment of AI

strategy, Senior Directors for Prompt Integrity and Vice-Presidents for AI Cybersecurity (Swayne 2026) — and legal departments and legal roles will grow exponentially, because firms will do the maximum to escape the liability for all the errors and damages that their AI tools will produce. The result will be not fewer jobs, but different jobs — all 'overhead' jobs, the primary function of which is not production, but insurance and reassurance: "reassurance that the hierarchy remains intact, that human beings remain 'in the loop', and that someone, somewhere, is still nominally responsible" (Swayne 2026).

Moreover, there will be numerous jobs, concealed and poorly rewarded, for 'AI-raters' for Google (Bansal 2025), OpenAI or Meta, who, at considerable psychological costs, use their judgment to moderate often gruesome content, maintain minimal decency guardrails and help the AI chatbots sound more human. Likewise, there will be a job explosion in robotics, because there is a lot of manual labor that goes into building and training humanoid robots. Human workers are needed to act as data providers and collectors at massive scale for the robotics firms. Robotics companies also continue to employ human tele-workers to operate their robots remotely and keep them on track: "if a robot gets stuck, or if the customer wants it to do a tricky task, a tele-operator from the company's headquarters will pilot it, looking through its cameras to iron clothes or unload the dishwasher" (O'Donnell 2026).

AI-driven Taylorism will further create overhead-jobs in employee surveillance. After all, the 'iron law' of surveillance technology is that if it can be used, it will be used.[7] With many more means of monitoring, employers are tracking PC screens, logging keystrokes, analyzing internal communications, and even scanning employees' faces for signs of fatigue. AI-powered systems monitor the tone of internal communications and track the frequency of messages and log application use across digital platforms. Task completion, web browsing, and response times are filtered through algorithms to determine individual worker productivity scores, which are used to decide on promotions and (bonus) pay. AI tools are employed to conduct automated background checks (of employees and job applicants) that scrape, with disregard for privacy and human decency, public records, financial data, and social media activity to assess 'risk' for the employer. The result is a dystopian workplace without privacy and trust, reminiscent of Dave Egger's (2013) novel '*The Circle*'. The permanent monitoring creates a chilling effect (Nelson 2025) and undermines productivity (Sanchez 2025), particularly in collaborative and creative environments where psychological safety, trust and freedom are key to innovation.

[7] As Thomas Ferguson points out, the AI firms think that the AI tools will bring about sweeping changes in vast sectors of the American economy on which they will be able make enormous amounts of money (French 2025). But, as Ferguson argues, virtually everywhere, often even inside their own firms, they face opposition from people they accuse of standing in the way of progress — including from unions. That is why the AI firms are busy scaring workers by threatening a 'job apocalypse' and using surveillance in order to break employees' willingness to seek higher pay or benefits. They get scared. It is very much like what happened in the U.S. in the late 1990s, when, as Alan Greenspan told the entire Fed Open Market Committee, capital mobility and technological advances deeply cowed American workers (French 2025).

Firms use AI as a threat to discipline workers or to lay them off, but with the purpose to rehire them back as gig workers or to find someone else in the supply chain who is doing that work instead. "Your boss wants to use surveillance data to cut your wages", sums it up nicely (Doctorow 2026b). But this chilling illustration of AI-based surveillance wage discrimination (common in U.S. healthcare, customer services, logistics and retail) hammers home the point:

"algorithmic wage discrimination draws on *external* sources of data to set the price of your labor. That's the situation for contract nurses, whose traditional brick-and-mortar staffing agencies have been replaced by nationwide apps that market themselves as "Uber for nursing." These apps use commercial surveillance data from the unregulated data-broker sector to check on how much credit card debt a nurse is carrying and whether that debt is delinquent to set a wage: the more debt you have and the more dire your indebtedness is, the lower the wage you are offered." (Doctorow 2026b).

Carrying a credit-card balance, taking out a payday loan, or discussing one's indebtedness on WhatsApp or Instagram can lead to lower wages. The real risk of AI is that it degrades and dehumanizes work even more. "AI is not going to take your job, but it will likely make your job shittier," says Alex Hanna, the director of research for the Distributed AI Research Institute (Boyle 2025). AI does not threaten jobs, it threatens human decency and meaning. This is not a bug but a feature of AI — the corporate stooges who act as the in-house ideologues of the Silicon Valley tech lords (Starr 2025), see democracy as a bug, treat workers as rebootable peons to their tech overlords, and openly advocate for CEO-monarchs who run the economy as a start-up firm and treat citizens as obedient customers. They propagate tyranny by bureaucracy, written in perfect code (Starr 2025).

It is plain that corporate-for-profit AI is never going to be 'pro-worker' — however much MIT economists Daron Acemoglu, David Autor, and Simon Johnson (2026) would wish this to happen. The MIT trio distinguishes between 'bad' and 'good' (pro-worker) AI and next wants government (Trump?) to push AI policy and regulation in an unambiguously 'pro-worker' direction, promoting technologies that make human skills and expertise more valuable by creating new tasks. To achieve this, they propose a rather incoherent laundry list of neoliberal policies — including building AI expertise in government, encouraging competition, giving workers a 'voice' (in what exactly?), tax reform, preventing expertise theft, and getting rid of occupational licensing — which will only further weaken and undermine the bargaining position and job security of workers. The MIT troika, quite in character, refuses to engage with key inconvenient truths: the inherent drive for automation and worker exploitation and control that is embedded in (shareholder) capitalism, the historical collapse of labor unions, the rise of electronic Taylorism (French 2025), the regulatory capture of government by the Big (Fin) Tech firms, AI tools to exploit loopholes in copyright laws, environmental rules and labor laws, and pressures from global competition to cut labor costs. In other words, they choose to ignore the actual reality in which Jeff Bezos is weaponizing Amazon's algorithmic surveillance tools to prevent workers from organizing, crushing pro-union movements inside his warehouses (Wiggin

2026); Elon Musk demands that his employees work long overtime hours without pay or sleep (Nanji 2022), while he himself hangs around and tweets; Sam Altman prioritizes 'growth' over 'safety concerns', while cozying up with rich dictators in the Gulf who are not known for their human-rights record (Farrow and Marantz 2026); and 'masculine-energy' Mark Zuckerberg has adopted a hostile posture towards his own workers, cracked down on free speech within his corporation, stopped worker involvement in decision-making, and abolished all diversity initiatives (Marcotte 2025).

Given the relatively humdrum productivity and cost-saving impacts of generative AI (LLMs), it is evident that the U.S. is massively overinvesting in mammoth AI data-center infrastructure. Much of these data-center deals are circular, highly leveraged and involving private credit. The AI boom is a massive bubble, in other words — which will pop, at the latest when the evidence that generative AI tools are not transformative (but rather, the pinnacle of 'enshittification'; Górny 2026) becomes undeniable. But more likely, the AI bubble will pop sooner, greatly helped by the self-inflicted stagflation, caused by Trump's illegal war in Iran. Specifically:

- the higher oil and gas prices will raise the operational costs of AI computing equipment, making it more difficult for the AI companies to break even;
- shortages of computer chips due to the breakdown of supply chains will fatally delay the data-center buildout; the Iran war could grind the AI build-out to a halt; geopolitical risks have increased, as state-sponsored cyberattacks of the variety Iran is known for could knock a data center offline;
- electricity grid bottlenecks, regulation, and local NIMBY opposition to data centers were already delaying the process (Maine is set to be the first state imposing a data-center ban; Wilkins 2026); relatively little AI data-center infrastructure has been built so far;
- higher interest rates and greater uncertainty will trigger a private credit crisis that will affect the AI industry directly;
- price competition from low-cost Chinese firms and competition between U.S. firms will fatally undermine the business models of OpenAI, Anthropic and xAi — forcing them into bankruptcy or a bailout by one of the cash-rich platform corporations.

These factors will inexorably strain the viability of AI business models and the resilience of AI supply chains and ultimately asphyxiate the AI hype — and, in all this, it is important to recognize that the AI bubble is not a very 'useful' bubble (as defined by Carlota Perez), because the data-center infrastructure that will be the material legacy of the AI frenzy has a comically short economic lifespan: the GPUs, the chips and the data storage racks and servers make up 70% of the data-center capital expenditure (Kwok 2026) and particularly the GPUs become obsolete (for AI data-center work) after 2 or 3 years. These chips are part of purpose-built AI data centers — engineered for extreme power density, advanced cooling, and specialized networking. Together, they form a closed system optimized for scale but hard to repurpose. In addition, existing facilities also become structurally, thermally, or electrically unable to support newer

vintages of AI hardware without very expensive retrofits. “It is impossible to argue with physics.” Silicon Valley’s tech brothers will soon find out.

Postscript

For those curious to hear, after reading the paper, what will happen to Russell’s tiny teapot: my prediction is that on April 1, 2028, the teapot will be fatally hit, somewhere on its orbit, by Musk’s flying Tesla Roadster — and break into more than thousand pieces that will be scattered across the universe. It would be an intolerable presumption on the reader’s part to doubt this prediction. And, unperturbed, the flying Tesla will continue its journey to nowhere, quite like the generative AI industry.

References

Acemoglu, D. 2025. 'The Simple Macroeconomics of AI.' *Economic Policy* 40 (121): 13–58, https://doi.org/10.1093/epolic/eiae042

Acemoglu, A., D. Kong and A. Ozdaglar. 2026. 'AI, Human Cognition and Knowledge Collapse.' *NBER Working Paper* # 34910. Cambridge, Mass.: National Bureau of Economic Research. https://www.nber.org/papers/w34910

Acemoglu, D., D. Autor and S. Johnson. 2026. 'Building Pro-Worker Artificial Intelligence.' *NBER Working Paper* # 34854. Cambridge, Mass.: National Bureau of Economic Research. https://www.nber.org/papers/w34854

Adebayo, K.S. 2026. 'Dario Amodei Doubled Down On His AI Jobs Warning. Here's What's Different Now.' *Forbes*, February 21.

AktivTrak. 2026. '2026 State of the Workplace. AI Adoption & Workforce Performance Benchmarks.' AktivTrak Productivity Lab. https://www.activtrak.com/resources/state-of-the-workplace/

Alderson, M. 2026. 'No, It Doesn't Cost Anthropic $5k Per Claude Code User.' Substack, March 9. Link: https://martinalderson.com/posts/no-it-doesnt-cost-anthropic-5k-per-claude-code-user/

Al-Sibai, N. 2024. 'Godfather of AI Says Elon Musk Is Lying About Self-Driving Teslas.' *Futurism*, September 23.

Al-Sibai, N. 2025. 'Bombshell Research Finds a Staggering Number of Scientific Papers Were AI-Generated.' *Futurism*, July 4.

Al-Sibai, N. 2025. 'The AI Industry Is Still Light-Years From Making a Profit, Experts Warn.' *Futurism*, August 16.

American Clean Power. 2024. 'Economy-wide Impacts of the Inflation Reduction Act Energy Provisions.' Report, American Clean Power, December 19.

An, F. 2025. 'Why AI Customer Service Is Bad and What Lies Ahead.' *Sobot*, February 25.

Anand, N., T. Bautzer and M. Saini. 2025. 'First Brands, Tricolor Collapses Raise Fears of Credit Stress, with Dimon Warning of 'More cockroaches'.' *Reuters*, October 14.

Andersen, R. 2026. 'Science Is Drowning in AI Slop.' *The Atlantic*, January 22.

Anthropic. 2025. 'Consumer Terms of Service.' October 8. Link: https://www.anthropic.com/legal/consumer-terms

Anthropic (n.d.) 'Project Glasswing.' Link: https://www.anthropic.com/glasswing

Appel, R., P. McCrory, A. Tamkin, M. Stern, M. McCain and T. Neylon. 2025. 'Anthropic Economic Index report: Uneven Geographic and Enterprise AI Adoption.' September 15.

Armstrong, R. 2026. ‘The Citrini Fuss Exposes a Market Looking For an Excuse to Fall.’ *The Financial Times*, February 24.

Arnon, A. 2025. ‘The Projected Impact of Generative AI on Future Productivity Growth.’ Penn Wharton Budget Model, September 8. University of Pennsylvania.

Asadi, M., J.W. O'Sullivan, F. Cao, T. Nedaee, K. Rajabalifardi, F. Li, E. Adeli and E. Ashley. 2026. ‘MIRAGE: The Illusion of Visual Understanding.’ Link: https://arxiv.org/abs/2603.21687

Bair Jr., T. 2025. ‘Record-Breaking AI-Related Debt Issuance in 2025.’ BNY Mellon, December. https://www.mellon.com/insights/insights-articles/record-breaking-ai-related-debt-issuance-in-2025.html

Ball, P. 2023. ‘Is AI Leading To a Reproducibility Crisis In Science?.’ *Nature* **624**, pp. 22-25.

Bansal. V. 2025. ‘Meet the AI Workers Who Tell Their Friends and Family To Stay Away From AI.’ *The Guardian*, November 22.

Bansal. V. 2026. ‘Amazon Is Determined To Use AI For Everything – Even When It Slows Down Work.’ *The Guardian*, March 11.

Bastian, M. 2026. ‘OpenAI Adds $111 billion to Its Cash Burn Forecast as AI Costs Spiral Beyond Projections.’ *The Decoder*, February 21.

Beaumont, P. 2026. ‘US Responsible For Deadly Missile Strike on Iran School, Preliminary Inquiry Says.’ *The Guardian*, March 11.

Béchard, D.E. 2025. ‘Why Humanoid Robots Still Can’t Survive in the Real World.’ *Scientific American*, December 13.

Becker, J. N. Rush, E. Barnes and D. Rein. 2025. ‘Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity.’ July 10. METR Model Evaluation & Threat Research. https://metr.org/blog/2025-07-10-early-2025-ai-experienced-os-dev-study/

Bender, E.M. and A. Hanna. 2025. *The AI Con. How to Fight Big Tech’s Hype and Create the Future We Want*. New York City: Harper Collins Publishers.

Bender, E.M., T. Gebru, A. McMillan-Major and S. Shmitchell. 2021. ‘On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?’ FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency pp. 610 – 623

Bindley, K. 2026. ‘Tech Firms Aren’t Just Encouraging Their Workers to Use AI. They’re Enforcing It..’ *The Wall Street Journal*, February 24.

Blake, D. 2026. ‘Sam Altman Admits OpenAI ‘Screwed Up’ the Writing Quality on ChatGPT 5.2 – and Promises Future Versions Won’t ‘Neglect’ It.’ *TechRadar*, January 28.

Bloomberg Tax. 2026. ‘R&D Tax Credits and Deductions Explained.’ Link: https://pro.bloombergtax.com/insights/federal-tax/rd-tax-credit-and-deducting-rd-expenditures/#what-is-the-80-rule-for-r-d-credit

Bode, K. 2026. ‘"CEO Said A Thing!" Journalism.’ *The Fine Print*, March 30.

Boffey, D. 2026. ‘‘It Means Missile Defence on Datacenters’: Drone Strikes Raise Doubts Over Gulf as AI Superpower.’ *The Guardian*, March 7.

Boyle, . 2025. ‘Scholars Explain How Humans Can Hold The Line Against AI Hype, And Why It’s Necessary.’ *GeekWire*, May 19.

Bregman, R. 2026. ‘Quit ChatGPT: Right Now! Your Subscription is Bankrolling Authoritarianism.’ *The Guardian*, March 4.

Brodkin, J. 2026. ‘Anthropic Sues US Over Blacklisting; White House Calls Firm “Radical Left, Woke”.’ *ArsTechnica*, March 10.

Brown, C. 2026. ‘Trump Administration Waging Illegal War On Iran, Experts Say.’ *The Guardian*, March 4.

Brown, S., S. Cruise, P. Singh and D. Scigliuzzo. 2026. ‘Blackstone Senior Staff Opened Wallets With Flagship Private Credit Fund Under Pressure.’ *Bloomberg News*, March 3.

Brynjolfsson, E., B. Chandar and R. Chen. 2025. ‘Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence.’ Stanford Digital Economy Lab, November 13.

Buchanan, N. 2026. ‘The Iran War Has Spiked Helium Prices. Here’s Which Products and Industries Could Get More Expensive.’ *Business Insider*, April 5.

Bunyard, T. (n.d.). ‘An Introduction to Guy de Bord’s ‘The Society of the Spectacle’. *Historical Materialism*, https://www.historicalmaterialism.org/article/an-introduction-to-guy-debords-the-society-of-the-spectacle/

Bureau of Labor Statistics. 2026. ‘Employment Projections 2024-2034.’ https://data.bls.gov/projections/occupationProj

Burleigh, E. 2026. ‘JPMorgan CEO Jamie Dimon Predicts AI Will Cut the Workweek Down to 3.5 days—And Tells Gen Z Developing EQ is More Important Than Ever.’ *Fortune*, April 6.

Bussewitz, C. 2025. ‘Mistake-Filled Legal Briefs Show the Limits of Relying on AI Tools At Work.’ *The Independent*, October 30.

Camus, A. 1942. *Le Mythe de Sisyphe*. Paris: Éditions Gallimard.

Cappelli, P. and V. Yakubovich. 2024. ‘AI May Not Be a Job Killer, After All.’ *The Wall Street Journal*, May 9.

Challapally, A. *et al.* 2025. *The GenAI Divide. State of AI in Business 2025*. July. Project NANDA. Cambridge, Mass.: MIT.

Challenger, Gray and Christmas. 2025. ‘Challenger Report December 2025.’

Chen, M. 2026. ‘Anthropic’s Claude Code Source Code Leak: Inside the 512,000-Line npm Packaging Error That Exposed AI’s Biggest Secret Roadmap.’ Tech Insider, April 3. Link: https://tech-insider.org/anthropic-claude-code-source-code-leak-npm-2026/

Chen, J., X. Xu, H. Wei, C. Chen and B. Zhao. 2026. ‘SWE-CI: Evaluating Agent Capabilities in Maintaining Codebases via Continuous Integration.’ April 1. https://arxiv.org/abs/2603.03823

Citrini Research. 2026. ‘The 2028 Global Intelligence Crisis. A Thought Exercise in Financial History, from the Future.’ *Citrini Research*, Substack. February 22.

Claburn, T. 2026. ‘Researchers Poison Stolen Data to Make AI Systems Return Wrong Results.’ *The Register*, January 6.

Claburn, T. 2026. ‘Experiment Suggests AI Chatbot Would Save Insurance Agents a Whopping 3 Minutes a Day.’ *The Register*, January 16.

Cloutier, M. 2026. ‘AI Won’t Replace Engineers—But Weak Technical Leadership Will.’ *Forbes*, February 16.

CodeRabbit. 2025. ‘State of the AI vs. Human Code Generation Report.’ December 17. CodeRabbit.

Committee for a Responsible Federal Budget. 2026. ‘CBO’s February 2026 Budget and Economic Outlook.’ Link: https://www.crfb.org/papers/cbos-february-2026-budget-and-economic-outlook

Conger, K. and R. Mac. 2026. ‘Elon Musk Wants to Build an A.I. Satellite Factory on the Moon.’ *The New York Times*, February 10.

Copp, T., E. Dwoskin and I. Duncan. 2026. Anthropic’s AI Tool Claude Central to U.S. Campaign in Iran, Amid a Bitter Feud.’ *The Washington Post*, March 4.

Crawford, K. 2025. ‘Eating the Future: The Metabolic Logic of AI Slop.’ *E-flux Architecture*, September. https://www.e-flux.com/architecture/intensification/6782975/eating-the-future-the-metabolic-logic-of-ai-slop

Crosley, B. 2026. ‘Hyperscaler CapEx Hits $600B in 2026: The AI Infrastructure Debt Wave.’ introl.com, January 7.

Cummings, M. 2023. ‘What Self-Driving Cars Tell Us About AI Risks.’ IEEE Spectrum, July 30.

Davis, S. 2026. ‘AI Is Simultaneously Aiding and Replacing Workers, Wage Data Suggest.’ Federal Reserve Bank of Dallas, February 24. https://www.dallasfed.org/research/economics/2026/0224

Doctorow, C. 2026a. ‘AI Companies Will Fail. We Can Salvage Something From the Wreckage.’ *The Guardian*, January 18.

Doctorow, C. 2026b. ‘Your Boss Wants To Use Surveillance Data To Cut Your Wages.’ *Pluralistic*, April 6.

Desilver, D. 2025. 'How ADP Employment Data Compares With Official BLS Jobs Reports.' Pew Research Center, November 5.

Douthat, R. 2026. 'Anthropic's Chief on A.I.: 'We Don't Know if the Models Are Conscious'.' *The New York Times*, February 12.

Down, A. and A. Topping. 2026. 'OpenAI Shelves Stargate UK in Blow to Britain's AI Ambitions.' *The Guardian*, April 9.

Economist. 2026. 'Even The Best-Case Scenario for Energy Markets Is Disastrous.' *The Economist*, March 23.

Eisikovits, N. and J. Burley. 2026. 'The Greatest Risk of AI In Higher Education Isn't Cheating – It's The Erosion Of Learning Itself.' *The Conversation*, February 19.

Elder, B. 2025. 'How High Are OpenAI's Compute Costs? Possibly a Lot Higher Than We Thought.' *The Financial Times*, November 12.

Ellis, L. 2026. 'CEOs Say AI Is Making Work More Efficient. Employees Tell a Different Story.' *The Wall Street Journal*, January 21.

Eloundou, T., S. Manning, P. Mishkin and D. Rock. 2023. 'GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models.' https://arxiv.org/abs/2303.10130

Emerson, G. *et al*. 2026. 'AI Will Reshape More Jobs Than It Replaces.' BCG Henderson Institute, Boston Consulting Group, April 3.

Emtrain. 2026. 'AI Surveillance in the Workplace.' Link: https://emtrain.com/blog/ethics-and-compliance/ai-surveillance-in-the-workplace/

England, A. and S. Kerr. 2026. 'Gulf States Could Review Overseas Investments to Ease Financial Strains Caused by Iran War.' *The Financial Times*, March 5.

Estrada, E. 2025. 'MIT report: 95% of generative AI pilots at companies are failing.' *Fortune*, August 18. https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo

EvExpert. 2026. 'Tesla's Robotaxi Crash Rate 9 Times Worse Than Human Drivers, NHTSA Data Reveals.' *Bike-EV.com*. February 1.

Evergreen Technologies. 2025. 'World Models vs. Word Models: Why Yann LeCun Believes LLMs Will Be Obsolete.' *Medium*, November 15.

Farrow, R. and A. Marantz. 2026. 'Sam Altman May Control Our Future—Can He Be Trusted?' *The New Yorker*, April 6.

Federal Reserve Bank of Chicago. 2026. 'Chicago Fed Labor Market Indicators.' Link: https://www.chicagofed.org/research/data/chicago-fed-labor-market-indicators/latest-release

Federal Reserve Bank of New York. 2026. 'The Labor Market for Recent College Graduates.' https://nyfed.org/collegelabor.

Felten, E., M. Raj and R. Seamans. 2021. Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses. *Strategic Management Journal* 42 (12): 2195–2217. https://doi.org/10.1002/smj.3286

Ferdman, A. 2025. 'AI Deskilling Is A Structural Problem.' *AI & Society*. https://doi.org/10.1007/s00146-025-02686-z

Ferguson, M. 2026. 'Palantir CEO Makes Shocking Confession on Disrupting Democratic Power.' *The New Republic*, March 12.

Ferguson, T., J. Chen, M. Lalisse and P. Jorgensen. 2026. 'Move Fast and Break Everything: Crypto and the Democrats.' *INET Blog*, January 13. New York: Institute for New Economic Thinking.

Ferguson, T. and S. Storm. 2023. 'Myth and Reality in the Great Inflation Debate: Supply Shocks and Wealth Effects in a Multipolar World Economy.' *International Journal of Political Economy* 52 (1): pp. 1-44.

Ferguson, T. and S. Storm. 2024. 'Good Policy or Good Luck? Why Inflation Fell Without a Recession.' *International Journal of Political Economy*, *53* (4), 311-341.

French, N. 2025. 'Tech Capital Is Dominating American Politics. An Interview with Thomas Ferguson.' *Jacobin*, September 17. https://jacobin.com/2025/09/tech-capital-american-politics-ferguson

Frey, C.B. and M. Osborne. 2013. 'The Future of Employment: How Susceptible Are Jobs to Computerisation?' *Oxford Martin School Working Paper*, Oxford University.

Forrester. 2026. 'The Forrester AI Job Impact Forecast, US, 2025–2030.' January 13. https://www.forrester.com/press-newsroom/forrester-impact-ai-jobs-forecast/

Fort, S. 2026. 'AI Cybersecurity After Mythos: The Jagged Frontier.' *AISLE Newsletter*. Link: https://aisle.com/blog/ai-cybersecurity-after-mythos-the-jagged-frontier

Fuhrmann, R. 2026. 'Which Big Tech Stocks Have the Most Debt, and Why It Matters.' *The Motley Fool*, February 11.

Gardizy, A. and A. Efrati. 2026. 'OpenAI CEO and CFO Diverge on IPO Timing.' The Information, April 5.

Gartner. 2026. 'Gartner Says Worldwide AI Spending Will Total $2.5 Trillion in 2026.' January 15. Link: www.gartner.com

Gavin, W. 2026. 'Elon Musk's Plans for SpaceX Data Centers in Orbit Could Cost $5 trillion a Year, Analysts Say.' *MarketWatch*, February 4.

Gervais, O. 2026. 'Guide to Assessing the Economic Impact of Higher Oil Prices on Canada and the U.S.' Scotia Bank, March 2.

Gibney, E. 2026, 'How AI Slop Is Causing a Crisis In Computer Science.' *Nature*, February 13.

Gimbel, M., M. Kinder, J. Kendall and M. Lee. 2025. 'Evaluating the Impact of AI on the Labor Market: Current State of Affairs.' October 1. Yale Budget Lab, Yale University.

Global Law List Research Team. 2026. 'Will AI Replace Lawyers? A Data-Driven Analysis of What's Really Happening in 2026.' March 4.

Goldstein, S. 2026. 'More Pain in Private Credit With JPMorgan Reportedly Tightening Lending While a $33 Billion Fund Sees Heavy Redemptions.' *MarketWatch*, March 11.

Golchha, A. 2026. 'The Class of 2026 Is Struggling To Find Jobs—And it's Not Because of AI.' *The Stanford Review*, April 9.

Goldman Sachs. 2025. 'How Will AI Affect the Global Workforce?.' Goldman Sachs, August 13.

Górny, M. 2026. 'The Pinnacle of Enshittification, Or Large Language Models.' *Michał Górny: Retroactively fixing the world*, April 5.

Gottlich, M. 2026. 'One-Year Inflation Bets Surge Past 5% in Wake of Oil Shock.' *Seeking Alpha*, March 20.

Gran, B. 2026. BofA Survey: AI Bubble Just Became the Biggest Worry for Credit Investors'.' *The Motley Fool*, March 2.

Gravitee. 2026. 'The State of AI Agent Security 2026.' February 3.

Grayson, T. 2026. 'The AI Data Center Obsolescence Crisis: Why Physics is Popping the Bubble.' *Tony Grayson The Control Room*, February 23.

GreetNow. 2025. 'Why Customers Hate Chatbots: The 2026 Reality Check Your Business Needs.' GreetNow, December 25. Link: https://greetnow.com/blog/customers-hate-chatbots

Guest, O. *et al*. 2025. 'Against the Uncritical Adoption of 'AI' Technologies in Academia.' Zenodo. https://zenodo.org/records/17065099

Guldvog, T. 2025. 'AI Slop Is Flooding Scientific Publishing - and It Shows.' *Science Reader*, December 21.

Hale, G. 2026. '"Claude Cannot Be Trusted to Perform Complex Engineering Tasks': AMD AI Head Slams Anthropic's Coding Tool After Months of Frustration.' TechRadar, April 7.

Hammond, G. 'OpenAI Investors Question $852bn Valuation As Strategy Shifts.' *The Financial Times*, April 14.

Handel, M.J. 2022. 'Growth Trends For Selected Occupations Considered At Risk From Automation.' *Monthly Labor Review*, July.

Harrison Dupré, M. 2025. 'Law Firms Caught and Punished for Passing Around "Bogus" AI Slop in Court.' *Futurism*, May 15.

Harrison Dupré, M. 2026. 'Absurd AI-Powered Lawsuits Are Causing Chaos in Courts, Attorneys Say, "Clogging the System" and Driving Up Costs.' *Futurism*, March 18.

Hartzog, W. and J. Silbey. 2025. 'How AI Destroys Institutions.' *University of California Law Journal* (forthcoming 2026).

Heal, A. and T. Nangle. 2026. 'Partners Group Sounds Alarm On Private Credit Default Rates.' *The Financial Times*, March 12.

Hernández, M. and A. Lo Prete. 2026. 'AI Coding Agents Are Running On Your Machines — Do You Know What They're Doing?' *Sysdig* blog, March 23.

Hiar, C. 2026. 'Musk's Exploding Megarocket Puts $8B in Space Investments At Risk.' *Politico*, March 31.

Hirschfeld, A. 2026. 'OpenAI's Fund Raising Boom Slows Amid Mounting Debt.' Aljazeera, March 7.

Huff, K. P. Wilson and M. Corradini. 2025. 'Killing a Nuclear Watchdog's Independence Threatens Disaster.' *Scientific American*, March 6.

Humlum, A. and E. Vestergaard. 2025. 'Large Language Models, Small Labor Market Effects.' NBER Working Paper 33777. Cambridge, Mass.: National Bureau of Economic Research. https://www.nber.org/papers/w33777

IBM. 2026. 'BM 2026 X-Force Threat Index: AI-Driven Attacks are Escalating as Basic Security Gaps Leave Enterprises Exposed.' *IBM Newsroom*, February 25.

INTNXT. 2026. 'Why Replacing Developers with AI Is Going Horribly Wrong in 2026.' Link: https://intnxt.com/blog/replacing-developers-with-ai-failure-2026/

J. P. and A. Steer. 2026. 'Why Cyber Defenders Need To Be Ready For Frontier AI.' U.K. National Cyber Security Center, March 30. https://www.ncsc.gov.uk/blogs/why-cyber-defenders-need-to-be-ready-for-frontier-ai

Jeans, D. 2026. 'Pentagon to Adopt Palantir AI as Core Military System, Memo Says.' *Reuters*, March 21.

Jin, B. 2026. 'An Inside Look at OpenAI and Anthropic's Finances Ahead of Their IPOs.' *The Wall Street Journal*, April 5.

Joshi, C. 2026. 'Why Technical Debt Is Quietly Eating Away Your 2026 Margins.' *Wishtree Tech*. March 18. https://wishtreetech.com/blogs/ai/why-technical-debt-is-quietly-eating-away-your-2026-margins/

Jošt, G., V. Taneski and S. Karakatič. 2024. 'The Impact of Large Language Models on Programming Education and Student Learning Outcomes." *Applied Science 14*, 4115.

Karger, E. *et al*. 2026. 'Forecasting the Economic Effects of AI.' March. Link: https://forecastingresearch.org/economic-effects-of-ai

Kay, G. 2026. 'The xAI Exodus: Two More Cofounders Leave — and Musk Says He's Rebuilding.' *Business Insider*, March 13.

Kedrosky, P. 2026. 'Chart of the Day: Data Center Buildout Slowed Sharply.' March 19. Link: https://paulkedrosky.com/chart-of-the-day-data-center-buildout-slowed-sharply/

Khalaf, R. 2026. 'Blue Owl Is Spoiling Private Credit's Sales Pitch.' *The Financial Times*, February 23.

Khalaf, R. 2026. 'What Will War in Iran Do to the Global Economy?' *The Financial Times*, March 20.

Kimball, S. 2026. 'Iran Strikes Halt Qatar LNG Output, Shaking Global Energy Markets.' CNBC, March 2.

Kobak, D. *et al.* 2025. 'Delving Into LLM-Assisted Writing In Biomedical Publications Through Excess Vocabulary.' *Science Advances* **11** eadt3813. 10.1126/sciadv.adt3813

Kosmyna, N. *et al*. 2025. 'Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task.' https://arxiv.org/abs/2506.08872

Kropp, M., M. Hsu, O.T. Karaman, J. Hawes and G. Rosen Kellerman. 2026. 'When Using AI Leads to "Brain Fry".' *Harvard Business Review*, March 5.

Kshirsagar, M. 2025. 'Lifespan of AI Chips: The $300 Billion Question.' *CITP Blog*, October 15. Link: https://blog.citp.princeton.edu/2025/10/15/lifespan-of-ai-chips-the-300-billion-question/

Kumar, V. 2025. 'Our Thoughts on Axios's "AI white-collar bloodbath".' *The Work Times*, May 28.

Kwok, K. 2026. 'How Big Tech's $630 Billion AI Splurge Will Fall Short.' *Reuters*, March 27.

Lambert, F. 2026a. 'Tesla Admits It Still Needs Drivers and Remote Operators — Then Argues That's Better Than Waymo.' *Elektrek*, February 19.

Lambert, F. 2026b. 'Musk Admits xAI 'Not Built Right' — Weeks After Tesla Invested $2 Billion.' *Elektrek*, March 13.

Landymore, F. 2025a. 'Judge Gives Humiliating Punishment to Lawyers Caught Using AI in Court.' *Futurism*, September 19.

Landymore, F. 2025b. 'Lawyer Gets Caught Using AI in Court, Responds in the Worst Possible Way.' *Futurism*, October 15.

Landymore, F. 2025c. 'CEO of Palantir Says He Spends a Large Amount of Time Talking to Nazis.' *Futurism*, November 14.

Landymore, F. 2026a. 'Researchers Studied What Happens When Workplaces Seriously Embrace AI, and the Results May Make You Nervous.' *Futurism*, February 10.

Landymore, F. 2026b. ‘Amazon’s Blundering AI Caused Multiple AWS Outages.’ *Futurism*, February 21.

Landymore, F. 2026c. ‘Meta’s Head of AI Safety Just Made a Mistake That May Cause You a Certain Amount of Alarm.’ *Futurism*, February 25.

Landymore, F. 2026d. ‘AI Is Forcing Employees to Work Harder Than Ever.’ *Futurism*, March 12.

Landymore, F. 2026e. ‘Rogue AI Agent Triggers Emergency at Meta.’ *Futurism*, March 21.

Landymore, F. 2026f. ‘Anthropic Warns That “Reckless” Claude Mythos Escaped a Sandbox Environment During Testing.’ *Futurism*, April 8.

La Scala, J., S. Sahli and D. Gillet. 2025. ‘Stimulating Brainstorming Activities with Generative AI in Higher Education.’ *2025 IEEE Global Engineering Education Conference (EDUCON)*, London, pp. 1-10.

Lawson, C. and H. Shah. 2025. ‘Big Banks Dominate NDFI Lending As Industry Interest, Regulatory Concern Grow.’ *S&P Global*, April 9.

Lewsey, F. 2026. ‘An Investigation Into 30 Top AI Agents Finds Just Four Have Published Formal Safety and Evaluation Documents Relating To the Actual Bots.’ University of Cambridge, February 20.

Linares, M. 2026. ‘AI Bots May Lead To The End of The Internet As We Know It.’ *OpenDemocracy*, February 20.

Linton, G.A. 2026. LinkedIn Post, April 10. https://www.linkedin.com/posts/craigalinton_counterpoint-is-anthropics-mythos-hype-activity-7448182264401870848-y-KB

Liu, J. and D. Webber. 2026. ‘AI Adoption and Firms Job-Posting Behavior.’ *FEDS Notes*. Washington: Board of Governors of the Federal Reserve System, March 27.

Lockett, W. 2026. ‘The SpaceX xAI Merger Is A Giant Red Flag.’ *Will Lockett’s Newsletter*, Substack. February 12.

Lutz, E. 2025. ‘AI Is Not Killing Cybersecurity Jobs — It’s Raising the Bar.’ The Array, July 21. https://thearray.hitachivantara.com/en-us/op-ed/ai-is-not-killing-cybersecurity-jobs

Ma, J. 2026. ‘Anthropic Cofounder Says Studying the Humanities Will Be ‘More Important Than Ever’ and Reveals What the AI Company Looks For When Hiring.’ Fortune, February 7.

Maguire, G. 2025. ‘Taiwan Buck’s Asia’s Clean Power Drive With Record Gas Burn.’ *Reuters*, October 31.

Mallaby, S. 2026. ‘I Went to China to See Its Progress on A.I. We Can’t Beat It..’ *The New York Times*, April 13.

Maniches, S. 2026. ‘The Real Cost of Running AI in 2026: Compute, Revenue, and Who Can Actually Afford It.’ *My Written Word*, March 26.

Marcotte, A. 2025. 'The Tech Billionaire War On "Woke" Is Really Targeting Workers.' *Salon*, January 17.

Marcus, G. 2025. 'BREAKING NEWS: AI Coding May Not Be Helping As Much As You Think.' *Marcus on AI*, Substack. July 10.

Marcus, G. 2026a. 'Promises Are Cheap.' *Marcus on AI*, Substack. February 12.

Marcus, G. 2026b. 'Does OpenAI's New Financing Make Sense?' *Marcus on AI*, Substack. February 27.

Marcus, G. 2026c. 'How AGI-Is-Nigh Doomers Own-Goaled Humanity.' *Marcus on AI*, Substack. March 3.

Marcus, G. 2026d. '"A Spate of Outages, Including Incidents Tied To the Use of AI Coding Tools", Right On Schedule.' *Marcus on AI*, Substack. March 10.

Marcus, G. 2026e. 'There Are No Heroes in Commercial AI.' *Marcus on AI*, Substack. March 8.

Marcus, G. 2026f. 'BREAKING: Expensive New Evidence that Scaling Is Not All You Need.' *Marcus on AI*, Substack. March 14.

Marcus, G. 2026g. 'On Employment, Don't Panic – Yet.' *Marcus on AI*, Substack. April 1.

Marcus, G. 2026h. 'The Two Wildest Stories Today in Tech.' *Marcus on AI*, Substack. April 3.

Marcus, G. 2026i. 'Claude Mythos, Evaluated.' *Marcus on AI*, Substack. April 13.

Markman, J. 2026. 'The Radiologist Effect: The Case Against AI Job Loss.' *Fortune*, January 26.

Martindale, J. 2025. 'OpenAI Signs Contract to Buy $300 Billion Worth of Oracle Computing Power Over the Next Five Years — Company Needs 4.5 Gigawatts of Power, Enough to Power Four Million Homes.' *Tom's Hardware*, September 11.

Massenkoff, M. and P. McCrory. 2026. 'Labor Market Impacts of AI: A New Measure and Early Evidence.' Anthropic, March 5.

Mazeika, M. *et al*. 2026. 'Remote Labor Index: Measuring AI Automation of Remote Work.' Link: https://labs.scale.com/papers/rli

Microsoft. 2026. 'Microsoft Copilot Terms of Use. Effective: October 24, 2025.' Link: https://www.microsoft.com/en-us/microsoft-copilot/for-individuals/termsofuse

Mikhaylenko, M., D. Rohrhirsch, P. Punnackal, P. Ruthardt and M. Schich. 2026. 'Profitless Prosperity in AI.' Roland Berger Consultants, March 4.

Milmo, D. 2026. 'Why Has Elon Musk Merged His Rocket Company With His AI Startup?' *The Guardian*, February 7.

Mills, S. and D.A. Spencer. 2025. 'Efficient Inefficiency: Organisational Challenges of Realising Economic Gains From AI.' *Journal of Business Research* 189: 115128.

Mondragon, S. 2026. 'AI Coding Tools Make Developers 19% Slower: What the Research Says.' Anthropic, March 13. https://particula.tech/blog/ai-coding-tools-developer-productivity-paradox

Morales, J. 2025. 'More than 50% of Nvidia's Data Center Revenue Comes from Three Customers — $21.9 Billion in Sales Recorded from the Unnamed Companies.' *Tom's Hardware*, August 31.

Morales, J. 2026. 'Microsoft's AI Boss Says AI Can Replace Every White-Collar Job in 18 Months — 'We're Going to Have a Human-Level Performance on Most, if Not All, Professional Tasks'.' *Tom's Hardware*, February 15.

Morgan Stanley. 2025. 'How AI Coding Is Creating Jobs.' October 29. Link: https://www.morganstanley.com/insights/articles/ai-software-development-industry-growth

Morris, S. 2026. 'Poolside Hunts Data Center Partners After CoreWeave Deal Falls Through.' *The Financial Times*, April 2.

Muir, M. and G. Steer. 2026. 'Fermi Shares Plunge 13% on $486mn Net Loss.' *The Financial Times*, March 30.

Nangle, T. 2026. 'Catastrophising Credit. Or Time to Panic.' *The Financial Times*, February 25.

Nanji, N. 2022. 'Elon Musk Tells Twitter Staff To Work Long Hours Or Leave.' BBC, November 17.

Nelson, M. 2025. 'Redefining Productivity in the Age of Workplace Surveillance.' Human Rights Research Center. July.

Niederhoffer, K., A. Robichaux and J.T. Hancock. 2026. 'Why People Create AI "Workslop"— and How to Stop It.' *Harvard Business Review*, January 16.

Nolan, B. 2026. 'ChatGPT's Market Share is Slipping as Google and Rivals Close the Gap, App-Tracker Data Show.' *Fortune*, February 5.

Novet, J. 2026. 'Microsoft Closes Worst Quarter on Wall Street Since 2008 on AI concerns: 'Redmond Is In a Pickle'.' CNBC, March 31.

O'Donnell, J. 2026. 'The Human Labor Behind Humanoid Robots Is Being Hidden.' *MIT Technology Review*, February 23.

O*NET. 2026. O*NET® 30.2 Database. https://www.onetcenter.org/database.html

Ong, S.Y. and C. Wang. 2026. 'Taiwan Hunting for Alternative LNG Supplies After Qatar Shutdown.' *Bloomberg News*, March.

OpenAI. 2025. 'Introducing Stargate UAE.' Link: https://openai.com/index/introducing-stargate-uae/

Oregon Group. 2026. 'Strait of Hormuz Is Chokepoint For Sulphuric Acid and Critical Metal Processing.' March 9.

Oxford Economics. 2026a. 'Evidence of an AI-Driven Shakeup of Job Markets Is Patchy.' *Oxford Economics Research Briefing*, January 7.

Oxford Economics. 2026b. 'Iran and the Strait of Hormuz: Risks to Global Energy Prices.' *Oxford Economics Research Briefing*, February 27.

Oxford Economics. 2026c. 'Iran War Set to Push GCC Economies into Recession.' *Oxford Economics Research Briefing*, March 30.

Parikh, T. 2026. 'How the Iran War Could Derail the AI Boom.' *The Financial Times*, March 22.

Parmar, H. 2026. 'OpenAI Is Falling Out of Favor With Secondary Buyers.' *Bloomberg News*, April 1.

Parramore, L. and A. Quart. 2026. 'For White-Collar Workers, AI Also Stands for "Apocalyptic Insecurity".' TNR, April 10.

Perloff, C. 2026. 'How Amazon Tracks Employee AI Usage—and Measures Results.' *The Information*, February 19.

Platt, E. 2026. 'Investors Sought to Pull $20bn From Private Credit Funds in First Quarter.' *The Financial Times*, April 9.

Purser, R. 2025. 'AI is Destroying the University and Learning Itself.' *Current Affairs*, December.

Putzier, K. 2026. 'U.S. Companies Are Still Slashing Jobs to Reverse Pandemic Hiring Boom.' *The Wall Street Journal*, January 29.

Randewich, N. 2026. 'Nvidia's PE Sinks to Seven-Year Low as War and AI Angst Weigh In.' *Reuters*, March 30.

Rego, G. 2025. 'How Do GPUs Depreciate? Try Entropy as a Depreciation Model.' *Medium*, October 14.

Resume.org. 2026. 'The Great Turnover: 9 in 10 Companies Plan To Hire in 2026, Yet 6 in 10 Will Have Layoffs.' https://www.resume.org/the-great-turnover-9-in-10-companies-plan-to-hire-in-2026-yet-6-in-10-will-have-layoffs-2/

Reuters. 2025. 'OpenAI Generates $4.3 Billion in Revenue in First Half of 2025, The Information Reports.' October 2.

Reuters. 2026. 'Apollo private credit fund limits investor withdrawals after requests surge.' March 23.

Reuters. 2026. 'OpenAI Projects $2.5 Billion in Ad Revenue This Year, $100 billion by 2030, Axios Reports.' April 9.

Roberts, M. 2026. 'Citrini and the AI Doom Scenario.' *Brave New Europe*, February 26.

Robins-Early, N. 2026. 'How AI firm Anthropic wound up in the Pentagon's crosshairs.' *The Guardian*, March 9.

Romm, J. 2026. 'One Affordability Battle After Another: What to do about the growing damage from the AI-Fossil Fuel Industrial Complex.' *INET Blog*, February 17. New York: Institute for New Economic Thinking.

Rooney, K. and A. Capoot. 2026. 'OpenAI Seeking Investments from Middle East Sovereign Wealth Funds for Multibillion-Dollar Round.' CNBC, January 21.

Roose, K. 2025. 'A.I.'s Prophet of Doom Wants to Shut It All Down.' *The New York Times*, September 12.

Rosner-Uddin, R. 2026. 'Amazon Holds Engineering Meeting Following AI-Related Outages.' *The Financial Times*, March 10.

Roth, E. 2026. 'Wikipedia Bans AI-Generated Articles.' *The Verge*, March 26.

Said, S. 2026. 'Saudi Arabia Sees a Spike to $180 Oil if Energy Shock Persists Past April.' *The Wall Street Journal*, March 19.

Sanchez, A.M. 2025. 'The Watchful Eye of Digital Surveillance At Work.' Waterloo News. University of Waterloo.

Saplakoglu, Y. 2024. 'How AI Revolutionized Protein Science, but Didn't End It.' *QuantaMagazine*, June 26. https://www.quantamagazine.org/how-ai-revolutionized-protein-science-but-didnt-end-it-20240626/

Saraswat, A. and K. Keenan. 2026. 'Gulf Shut-ins Could Reduce Regional Crude Output by 70% if US-Iran War Drags On.' Rystad Energy, March 13.

Satariano, A. 2026. 'Elon Musk's xAI Gets $3 Billion Investment From Saudi-Backed A.I. Firm.' *The New York Times*, February 18.

Savage, M. 2026. 'Publishers Fear AI Search Summaries And Chatbots Mean 'End of Traffic Era'.' *The Guardian*, January 12.

Schaer, C. 2026. 'No More Big Spenders: Iran War to Dent Gulf State Investment.' *DW In Focus*, March 18. Link: https://www.dw.com/en/no-more-big-spenders-iran-war-to-dent-gulf-state-investment/a-76400242

Schechner, S. and R. McMillan. 2026. 'Anthropic Races to Contain Leak of Code Behind Claude AI Agent.' *The Wall Street Journal*, April 1.

Sen, A. 2026. 'Debt Investors Offloading Exposure to Software Companies is Latest Sign of Pain.' *Reuters*, March 17.

Senti, O. 2026. 'Iran War Puts $300 Billion in Gulf AI Infrastructure Under Fire.' *AiHola*, March 9. Link: https://aihola.com/article/iran-war-gulf-ai-infrastructure

Sevilla, J. , H. Petrovic and A. Ho. 2026. 'Can AI Companies Become Profitable?' EpochAI, January 28.

Shen, J.H. and A. Tamkin. 2026. ‘How AI Impacts Skill Formation.’ Anthropic, January 29. https://arxiv.org/html/2601.20245v1

Shilov, A. 2024. ‘ Datacenter GPU Service Life Can Be Surprisingly Short — Only One to Three Years is Expected According to Unnamed Google Architect.’ *Tom’s Hardware*, October 24.

Shojaee, P. *et al*. 2025. ‘The Illusion of Thinking: Understanding the Strengths and Limitations of Reasoning Models via the Lens of Problem Complexity.’ Apple Co. June. https://machinelearning.apple.com/research/illusion-of-thinking

Shumer, M. 2026. ‘Something Big Is Happening in AI – and Most People Will Be Blindsided.’ *Fortune*, February 11.

Sims, D.2026. ‘Waymo Admits that Its Autopilot is Often Just Guys from the Philippines.’ *TechSpot*, February 7.

Singal, N. 2026. ‘Anthropic Alleges Large-Scale Distillation Campaigns Targeting Claude.’ *Computerworld*, February 24.

Smith, R.A. 2026. ‘AI Isn’t Lightening Workloads. It’s Making Them More Intense.’ *The Wall Street Journal*, March 11.

Sor, J. 2026. ‘Fears About Retail-Investor Participation in Private Credit Are Playing Out in Real Time.’ *Business Insider*, March 5.

StackOverflow. 2025. ‘Policy: Generative AI (e.g., ChatGPT) Is Banned.’ https://meta.stackoverflow.com/questions/421831/policy-generative-ai-e-g-chatgpt-is-banned

Starr, T. 2025. ‘The Intellectual Autopsy of Curtis Yarvin: Playing with Other People’s Lives in God Mode.’ *Free Markets & Firepower*, November 12.

Stockton, B. 2026. ‘OpenAI's Sam Altman Warns that Firms Are Using 'AI Washing' To Mask Layoffs Across the Globe — AI Boss Calls Out Corporate Excuses While Warning of ‘Palpable’ Job Disruption Ahead.’ *Tom’s Hardware*, February 26.

Stokel-Walker, C. 2025. ‘What the End of Section 174 Means For Software Developer Hiring.’ *LeadDev*, July 15.

Stokel-Walker, C. 2026. ‘AIs Can’t Stop Recommending Nuclear Strikes in War Game Simulations.’ *New Scientist*, February 25.

Storm, S. 2025. ‘The U.S. Is Betting the Economy on ‘Scaling’ AI: Where Is the Intelligence When One Needs It?’ *International Journal of Political Economy* 54 (4): 425–452.

Swain, G. 2025. ‘OpenAI Admits AI Hallucinations Are Mathematically Inevitable, Not Just Engineering Flaws.’ *Computerworld*, September 18.

Swayne, M. 2026. ‘The Bullshit Job Singularity: Why AI Might Not Destroy Jobs, But Could Make Them Bullshittier.’ *Mindplex*, February 12.

Tangermann, V. 2025. ‘The AI Industry Has a Huge Problem: the Smarter Its AI Gets, the More It’s Hallucinating.’ *Futurism*, May 5.

Tangermann, V. 2026a. ‘Cops Forced to Explain Why AI Generated Police Report Claimed Officer Transformed Into Frog.’ *Futurism*, January 16.

Tangermann, V. 2026b. ‘Sam Altman Says OpenAI Is Slashing Its Hiring Pace as Financial Crunch Tightens.’ *Futurism*, January 27.

Tangermann, V. 2026c. ‘Goldman Sachs Head During Financial Crisis Says He “Smells” a Similar Crash Coming.’ *Futurism*, March 4.

Tangermann, V. 2026d. ‘Drama Erupts Over Claims That Microsoft Will Embrace AI Even More Drastically in Windows 12.’ *Futurism*, March 12.

Tangermann, V. 2026e. ‘Microsoft Realizes It’s Epically Screwed Up Windows 11 as Users Rage at Copilot AI Crammed Everywhere.’ *Futurism*, March 23.

Tauman Kalai, A., O. Nachum, S.S. Vempala and E. Zhang. 2025. ‘Why Language Models Hallucinate.’ September 4, https://arxiv.org/abs/2509.04664

TheEnergyMag. 2026. ‘Microsoft Secures Former OpenAI "Stargate" Site in Norway for AI Infrastructure.’ April 14. Link: https://theenergymag.com/news/market-news/microsoft-secures-former-open-ai-stargate-site-in-norway-for-ai-infrastructure

Thompson, S.A. 2026. ‘A Word to the Wise: Don’t Trust A.I. to File Your Taxes.’ *The New York Times*, March 5.

Thubron, R. 2025. ‘Sam Altman Warns AI Could Wipe Out Entire Job Categories, Customer Support Roles Most At Risk.’ *TechSpot*, July 24.

Tidy, J. 2026. ‘AI ‘Slop’ Is Transforming Social Media - And a Backlash Is Brewing.’ BBC, February 4. https://www.bbc.com/news/articles/c9wx2dz2v44o

Townsend, K. 2026. ‘Critical Vulnerability in Claude Code Emerges Days After Source Leak.’ *SecurityWeek*, April 2.

Tunguz, T. 2026. ‘You Are Responsible for Your Agent.’ March 15. Link: https://tomtunguz.com/you-are-responsible-your-agent/?mod=djemMarketsAM#fn:4

van Rooij, I. 2025. ‘AI Slop And The Destruction Of Knowledge.’ https://doi.org/10.5281/zenodo.16905560

Vartabedian, B. 2026. ‘Medical Slop — A Short History of How We Learned to Look Busy.’ *The Liminal MD*, Substack. January 6.

Vaughan-Nichols, S.J. 2025. ‘AI Web Crawlers Are Destroying Websites in Their Never-Ending Hunger For Any and All Content.’ *The Register*, August 29.

Vigliarolo, B. 2024. ‘US Tech Innovation Dreams Soured by Changed R&D Tax Laws.’ *The Register*, January 12.

Vignesh R. 2026. 'Abu Dhabi's MGX Eyes Anthropic Stake as AI giant Chases $350B Valuation.' Link: https://techfundingnews.com/mgx-nears-anthropic-investment-20b-round/

WalkMe. 2026. 'Enterprises Lose 51 Workdays Per Employee to Technology Friction Annually Despite Record AI Investment, WalkMe Global Study of 3,750 Finds.' April 9. www.globenewswire.com

Wall Street Journal. 2025. 'Tech Trends 2026: AI Comes of Age.' December 10.

Weldon, M. 2025. 'Centaurs, Canaries and J-Curves: Pitfalls and Productivity Potential of AI.' *Newsweek*, December 18.

White House. 2025. 'Democratizing Access to Alternative Assets For 401(K) Investors.' Presidential Actions, August 7. https://www.whitehouse.gov/presidential-actions/2025/08/democratizing-access-to-alternative-assets-for-401k-investors/

Wiggin, T. 2026. 'Weaponizing the Workplace: How Algorithmic Management Shaped Amazon's Antiunion Campaign in Bessemer, Alabama.' *Socius: Sociological Research for a Dynamic World* (forthcoming).

Wilkins, E. 2026. 'AI Industry Super PAC Raises $125 Million in 2025.' CNBC, January 30.

Wilkins, E. 2026. 'Maine Set To Become First State With Data Center Ban.' CNBC, April 9.

Wilkins, J. 2025. 'McDonald's Pulls Down AI-Generated Holiday Ad After Deluge of Mockery.' *Futurism*, December 9.

Wilkins, J. 2026. 'Almost Half of US Data Centers That Were Supposed to Open This Year Slated to Be Canceled or Delayed.' *Futurism*, April 2.

Williams, K. 2026. 'Why Anthropic Believes Its Latest Model is Too Dangerous to Release.' *Understanding AI*, Substack. April 9.

Wirz, M. 2026. 'Private-Credit Warning Signs Flash After Blue Owl Unloads $1.4 Billion in Assets.' *The Wall Street Journal*, February 19.

Wirz, M. 2026. 'Blackstone Played Down Private-Credit Risks. Now It's Been Hit.' *The Wall Street Journal*, March 3.

Wirz, M.,P. Rudegeair and G. Heeb. 2026. 'An Exodus of Money Endangers Wall Street's Private-Credit Craze.' *The Wall Street Journal*, March 12.

Wong, M. and C. Warzel. 2026. 'Welcome to a Multidimensional Economic Disaster.' *The Atlantic*, March 26.

Wood, T. 2025. 'Will AI Replace Doctors, Lawyers, Writers, Engineers, or Radiologists?' *Fast Data Science*, September 16.

WoodMackenzie. 2026. 'Newly Added US Data Center Capacity Slows Down Considerably in Q4 2025, As Market Struggles to Keep Up with Explosive Demand.' March 16. Link:

https://www.woodmac.com/press-releases/newly-added-us-data-center-capacity-slows-down-considerably-in-q4-2025-as-market-struggles-to-keep-up-with-explosive-demand/

WorkDay. 2026. 'New Workday Research: Companies Are Leaving AI Gains on the Table.' January 14.

Wright, W. 2025. 'After Coding Catastrophe, Replit Says its New AI Agent Checks Its Own Work - Here's How To Try It.' ZDNET.com, September 11.

WSJPRO. 2026. 'Transcript: Fed Chair Jerome Powell's Postmeeting Press Conference.' *The Wall Street Journal*, March 18.

Wu, Y. 2026. 'Is AI Replacing the Work of Skilled Radiologists? They Give Us Their Thoughts.' *The Conversation*, March 11.

Yang. A. 2026. 'The End of the Office.' *Andrew Yang Newsletter*, Substack. February 16.

Yotzov, I. et al. 2026. Firm Data on AI'.' *NBER Working Paper* #34836. Cambridge, Mass.: National Bureau of Economic Research. https://www.nber.org/papers/w34836

Zeff, M. 2026. 'OpenAI's President Gave Millions to Trump. He Says It's for Humanity.' *The Wire*, February 12.

Zitron, E. 2025. 'This Is How Much Anthropic and Cursor Spend On Amazon Web Services.' *Where's Your Ed At*, Substack. October 20.

Zitron, E. 2026a. 'The Haters' Guide to Oracle.' *Where's Your Ed At*, Substack. January 30.

Zitron, E. 2026b. 'The AI Data Center Financial Crisis.' *Where's Your Ed At*, Substack. February 13.

Zitron, E. 2026c. 'The Haters' Guide to Anthropic.' *Where's Your Ed At*, Substack. February 20.

Zitron, E. 2026d. 'Why Are We Still Doing This?' *Where's Your Ed At*, Substack. March 17.

Zitron, E. 2026e. 'The AI Industry Is Lying to You.' *Where's Your Ed At*, Substack. March 24.

Zitron, E. 2026f. 'The Subprime AI Crisis Is Here.' *Where's Your Ed At*, Substack. March 31.

Zitron, E. 2026g. 'News: OpenAI CFO Doesn't Believe Company Ready For IPO, Unsure Revenue Will Support Commitments.' *Where's Your Ed At*, Substack. April 6.