

OpenAI signals a new phase in tech's AI race

Chris Nicholson, PhD, Head of Research

What is OpenAI's rumored breakthrough, Q*?

On November 17th, OpenAI's board of directors unceremoniously fired CEO Sam Altman, leading co-founder Greg Brockman to resign and over 95% of OpenAI's employees to threaten to follow them. The move blindsided key investors like Microsoft. Shareholders and employees managed to bring Altman back as CEO and force most of the board to step down. Now that the dust has settled on the fight over OpenAI's corporate mission and structure, it's time to assess the battle's significance to the tech sector's AI competition.

The reasons for the firing are beginning to emerge. At the time, the board issued a [statement](#) saying only that Altman was "not consistently candid" in his communications with it. According to a recent Reuters [exclusive](#), "Ahead of OpenAI CEO Sam Altman's four days in exile, several staff researchers wrote a letter to the board of directors warning of a powerful artificial intelligence discovery that they said could threaten humanity."

After Reuters broke the story, OpenAI CTO Mira Murati wrote employees an internal memo acknowledging a project called Q* (pronounced Q-star) and the letter to the board discussing it. Some at OpenAI think the technology could be a breakthrough toward AGI, artificial general intelligence, which the firm defines as AI that can surpass humans at most economically valuable tasks.

Although impressively able to create realistic texts and images, leading large language models like ChatGPT (LLMs) are what some critics call stochastic parrots. They simply recognize patterns in how words are used together and use those probabilities to construct sentences. This aspect of intelligence maps onto creativity and intuition. But ChatGPT struggled with math, which requires step-by-step deliberation to deduce a single right answer. GPT4 [aced](#) the bar exam, but only scored around the 40th percentile on AP Calculus BC.

Not much is known about Q*, but Reuters cited claims it performed very well at advanced high-school math. This achievement would enable AI to take on a number of quantitative jobs, especially coding. But such a good starting point also indicates Q* could perform higher-level math when using more compute power, and from there, start applying its expertise to sciences.

The tech and finance industries are buzzing with speculation about Q*. Forbes published a 10,000-word piece from an AI expert who, after almost-conspiratorial theorizing, [concludes](#) that Q* is probably an enhancement to LLMs using a cutting-edge reinforcement learning algorithm. Strive's research suggests that OpenAI may have adopted a new technique called Q* search [proposed](#) by a recent academic publication.

If true, this marks the AI revolution returning to its roots. ChatGPT's pattern-recognition approach caught the world's attention, but it was reinforcement learning that launched the

current AI era when DeepMind's AlphaGo beat top Go player Lee Sedol. Instead of recognizing patterns, reinforcement learning algorithms construct plans. In the classic example, a child learns the rule “don't touch the hot stove” after a single negative reinforcement. This ability to use experiences to construct rules and solve problems makes programs like AlphaGo able to [learn](#) truths beyond human knowledge, something LLMs like ChatGPT are weak at. ChatGPT is creative; AlphaGo and its descendants are deductive.

In layman's terms, the Q* search technique allows AI to speculate more about possible outcomes when deducing how to achieve its goals—for instance, instead of having to touch the hot stove once, it might reason “that stove will probably burn me if I touch it.” This refined ability to search through solutions to a problem likely makes Q* better at math.

In other words, ChatGPT only has a right brain right now, a creative side, and AlphaGo only had a left-brain logical side. Now OpenAI may have added a state-of-the-art left brain to ChatGPT and allowed the creative and deductive intelligences to work together. It's unclear what new capabilities may emerge from making generative AI capable of math and planning, which may have alarmed some of OpenAI's employees and its board, including chief scientist Ilya Sutskever.

This would be a breakthrough in the sense that it's a significant advance in capability, but OpenAI's Q* innovation was expected; it was obvious to many that combining the left and right brains was the next step for AI research. In fact, DeepMind, acquired by Alphabet, is releasing its own program enhancing LLM's pattern-recognition with reinforcement learning in early 2024, with a project called Gemini. We expect Alphabet's Gemini to be competitive with Microsoft and OpenAI's offerings, considering DeepMind's leadership in reinforcement learning.

Corporate governance questions over safety vs acceleration remain

Some in the AI community, like OpenAI's former directors and their interim CEO pick [Emmett Shear](#), think AI is advancing too fast in pursuit of corporate profit, and should slow development to prioritize safety from the existential risks and economic impacts of AI. OpenAI had a unique corporate structure that capped its investors' potential profit and gave control to a board of directors with a non-profit mission to develop AI for the benefit of humanity. In this case, advances in AI's commercial applications made its board concerned for its non-profit mission.

In September, Altman [said](#), “For me, AGI...is the equivalent of a median human that you could hire as a co-worker,” and could “do anything that you'd be happy with a remote co-worker doing,” including learning how to be a doctor and a highly competent coder. The rumored project called Q* is likely a step in that direction thanks to its improved quantitative reasoning. In one sign of its concern for the economic implications, OpenAI [appointed](#) former Treasury secretary Larry Summers to its new board. Summers predicts AI's new advances will “change

the way doctors, financial traders, authors and editors work.” He added that white-collar professions are likely to see more disruption than those that involve manual labor.

Those who want to accelerate AI’s capabilities are making progress, and newly [ascendant](#) now that OpenAI’s safety-minded board members have lost their bid to replace Altman. But the corporate governance war over AI safety versus acceleration is just getting started. We predict that the battleground will move from corporate structure to government regulation and shareholder governance. The self-described safetyists have convincingly lost their attempt to control AI development through an independent board with a non-profit mission—the OpenAI episode proves that large shareholders with profit motives can override a non-profit board. The shareholder can, for instance, gain leverage by poaching the AI firm’s employees, as Microsoft threatened to.

This event may end the trend of AI companies adopting hybrid for-profit/non-profit structures as an attempted panacea for safety concerns. OpenAI private competitor Anthropic, which Amazon invested \$4 billion in, [incorporated](#) as a PBC, a public benefit corporation. So did startup Inflection AI. A [PBC](#) is like a version of OpenAI’s hybrid structure formalized by Delaware law, giving shareholder profit and a non-profit mission equal importance. But those concerned about the existential and economic risks of AI may give up hope that hybrid corporate structures can control them. The AI safety movement will likely shift its hopes to government regulation and shareholder governance, possibly by making shareholder proposals and engagements aimed at public companies like Microsoft, Amazon, and Alphabet.

Economic implications of OpenAI’s possible breakthrough

Next-generation AI programs like Q* and Gemini are likely to displace large numbers of white-collar workers. [According](#) to Goldman Sachs, generative AI could lift US labor productivity by 1.5 percentage points per year over the next decade, doubling the historical rate since 2005. Industries that employ white collar workers are likely to gain the most from the margin improvements of automating much of their work. Industries that rely on manual labor may benefit the least, since advances in robotics have not kept pace with software and are limited by battery technology. These productivity gains could have a disinflationary effect, mitigating inflationary secular trends like deglobalization and aging populations.

When it comes to picking AI software winners among the tech sector, Microsoft and Alphabet have technological leads through OpenAI and DeepMind. It’s worth noting that Alphabet entirely owns DeepMind, and a separate Google AI division called Google Brain [invented](#) the transformer technology that OpenAI’s Generative Pre-trained Transformer (GPT) is based on. Alphabet has recently [merged](#) DeepMind and Google Brain into one entity to accelerate its AI progress. Investors who think Microsoft has a big lead in AI are mistaken: it’s currently a two-player race, and Alphabet is overlooked.

The semiconductor sector is the clear winner of the accelerationist victory at OpenAI. The faster AI development moves, the hungrier it will be for computing power. The AI hardware

race has only one contender at the moment: Nvidia. In Q3, its data center revenue was up [279% year-over-year](#), driven in large part by demand for its AI GPUs. Nvidia has maintained a consistent edge in hardware quality over its only competitor, AMD, but as the AI revolution takes off, it may also be reaping the benefits of AI programmers becoming locked-in to its CUDA software ecosystem—many engineers learned how to program AI on Nvidia’s software and don’t want to switch. Several years ago, investors debated whether Nvidia’s CUDA software would become a moat. So far, Nvidia’s results suggest it is.

Large customers like FAANG companies will try to reduce their dependence on Nvidia’s AI hardware and software, but the effort will take years. One step will be for them to create their own specialized chips, ASICs, as Google did by creating Tensor Processing Units to power its machine learning. Tesla is [rolling out its own ASIC](#) for use in its upcoming Dojo supercomputer, developing self-driving cars. The ASIC trend will take incremental business from Nvidia, but these company-specific internal chips won’t capture the market easily—Tesla couldn’t buy TPUs, and Google can’t buy Tesla’s ASICs to power its own self-driving effort, Waymo. Companies other than tech giants will likely lack the resources to develop their own chips.

So among semiconductor companies the clear AI leader is Nvidia, but AMD may benefit as an alternative to it, and ASIC-makers stand to gain for the same reason. The semiconductor industry’s near duopoly in Electronic Design Automation means that Cadence Design Systems and Synopsys are likely to benefit from increased chipmaking activity. Earlier in the semiconductor value chain, ASML is the likely equipment manufacturing winner because of its monopoly on the technology used to make the most advanced AI chips, Extreme Ultraviolet Lithography. As with Nvidia, the market will want competitors—Canon just launched a [product](#) it claims can compete with ASML’s chipmaking technology. But as with Nvidia, competing against an established tech monopoly is easier said than done.

The AI gold rush may have started a while ago, but the power struggle at OpenAI and the breakthrough that caused it signal a new phase. As with any gold rush, the vendors who sell picks and shovels tend to get richer than the miners. Investors’ task will be to figure out what the best pick-and-shovel AI plays are.