

Business | Rift valley

The Sam Altman drama points to a deeper split in the tech world

Doomers and boomers are fighting for AI dominance



Nov 19th 2023 (Updated Nov 20th 2023)

Even by the pace of the tech world, the events over the weekend of November 17th were unprecedented. On Friday Sam Altman, the co-founder and boss of Openai, the firm at the forefront of an artificial-intelligence (ai) revolution, was suddenly sacked by the company's board. The reasons why they lost confidence in Mr Altman are unclear. Rumours point to disquiet about his side-projects, and fears that he was moving too quickly to expand Openai's commercial offerings without considering the safety implications, in a firm that has also pledged to develop the tech for the "maximal benefit of humanity". Over the next two days the company's investors and some of its employees sought to bring Mr Altman back.

But the board has stuck to its guns. Late on November 19th it appointed Emmett Shear, former head of Twitch, a video-streaming service, as interim chief executive. Even more extraordinarily, the next day Satya Nadella, the boss of Microsoft, one of Openai's largest investors, posted on X (formerly Twitter), that Mr Altman and a group of employees from Openai would be joining the software giant to lead a "new advanced ai research team".

The events at Openai are the most dramatic manifestation yet of a wider divide in Silicon Valley. On one side are the “doomers”, who believe that, left unchecked, ai poses an existential risk to humanity and hence advocate stricter regulations. Opposing them are “boomers”, who play down fears of an ai apocalypse and stress its potential to turbocharge progress. The camp that proves more influential could either encourage or stymie tighter regulations, which could in turn determine who will profit most from ai in the future.

Openai’s corporate structure straddles the divide. Founded as a non-profit in 2015, the firm carved out a for-profit subsidiary three years later to finance its need for expensive computing capacity and brainpower in order to propel the technology forward. Satisfying the competing aims of doomers and boomers was always going to be difficult.

The split in part reflects philosophical differences. Many in the doomer camp are influenced by “[effective altruism](#)”, a movement that is concerned by the possibility of ai wiping out all of humanity. The worriers include Dario Amodei, who left Openai to start up Anthropic, another model-maker. Other big tech firms, including Microsoft, are also among those worried about ai safety, but not as doomery.

Boomers espouse a worldview called “effective accelerationism” which counters that not only should the development of ai be allowed to proceed unhindered, it should be speeded up. Leading the charge is Marc Andreessen, co-founder of Andreessen Horowitz, a venture-capital firm. Other ai boffins appear to sympathise with the cause. Meta’s Yann LeCun and Andrew Ng and a slew of startups including Hugging Face and Mistral ai have argued for less restrictive regulation.

Mr Altman seemed to have sympathy with both groups, publicly calling for “guardrails” to make ai safe while simultaneously pushing Openai to develop more powerful models and launching new tools, such as an app store for users to build their own chatbots. Its largest investor, Microsoft, which has pumped over \$10bn into Openai for a 49% stake without receiving any board seats in the parent company, is said to be unhappy, having found out about the sacking only minutes before Mr Altman did. That may be why the company offered Mr Altman and his colleagues a home.

Yet there appears to be more going on than abstract philosophy. As it happens, the two groups are also split along more commercial lines. Doomers are early movers in the ai race, have deeper pockets and espouse proprietary models. Boomers, on the other hand, are more likely to be firms that are catching up, are smaller and prefer open-source software.

Start with the early winners. Openai’s Chatgpt added 100m users in just two months after its launch, closely trailed by Anthropic, founded by defectors from Openai and now valued at \$25bn. Researchers at Google wrote the original paper on large language models, software that is trained on vast quantities of data, and which underpin chatbots including Chatgpt. The firm has been churning out bigger and smarter models, as well as a chatbot called Bard.

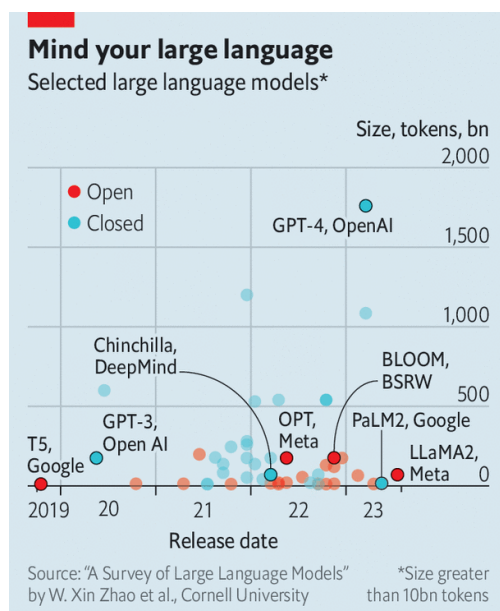
Microsoft’s lead, meanwhile, is largely built on its big bet on Openai. Amazon plans to invest up to \$4bn in Anthropic. But in tech, moving first doesn’t always guarantee

success. In a market where both technology and demand are advancing rapidly, new entrants have ample opportunities to disrupt incumbents.

This may give added force to the doomers' push for stricter rules. In testimony to America's Congress in May Mr Altman expressed fears that the industry could "cause significant harm to the world" and urged policymakers to enact specific regulations for ai. In the same month a group of 350 ai scientists and tech executives, including from Openai, Anthropic and Google signed a one-line statement warning of a "risk of extinction" posed by ai on a par with nuclear war and pandemics. Despite the terrifying prospects, none of the companies that backed the statement paused their own work on building more potent ai models.

Politicians are scrambling to show that they take the risks seriously. In July President Joe Biden's administration nudged seven leading model-makers, including Microsoft, Openai, Meta and Google, to make "voluntary commitments", to have their ai products inspected by experts before releasing them to the public. On November 1st the British

government got a similar group to sign another non-binding agreement that allowed regulators to test their ai products for trustworthiness and harmful capabilities, such as endangering national security. Days beforehand Mr Biden issued an executive order with far more bite. It compels any ai company that is building models above a certain size—defined by the computing power needed by the software—to notify the government and share its safety-testing results.



Another fault line between the two groups is the future of open-source ai. llms have been either proprietary, like the ones from Openai, Anthropic and Google, or open-source. The release in

February of llama, a model created by Meta, spurred activity in open-source ai (see chart). Supporters argue that open-source models are safer because they are open to scrutiny. Detractors worry that making these powerful ai models public will allow bad actors to use them for malicious purposes.

But the row over open source may also reflect commercial motives. Venture capitalists, for instance, are big fans of it, perhaps because they spy a way for the startups they back to catch up to the frontier, or gain free access to models. Incumbents may fear the competitive threat. A memo written by insiders at Google that was leaked in May admits that open-source models are achieving results on some tasks comparable to their proprietary cousins and cost far less to build. The memo concludes that neither Google nor Openai has any defensive "moat" against open-source competitors.

So far regulators seem to have been receptive to the doomers' argument. Mr Biden's executive order could put the brakes on open-source ai. The order's broad definition of "dual-use" models, which can have both military or civilian purposes, imposes complex

reporting requirements on the makers of such models, which may in time capture open-source models too. The extent to which these rules can be enforced today is unclear. But they could gain teeth over time, say if new laws are passed.

Not every big tech firm falls neatly on either side of the divide. The decision by Meta to open-source its ai models has made it an unexpected champion of startups by giving them access to a powerful model on which to build innovative products. Meta is betting that the surge in innovation prompted by open-source tools will eventually help it by generating newer forms of content that keep its users hooked and its advertisers happy. Apple is another outlier. The world's largest tech firm is notably silent about ai. At the launch of a new iPhone in September the company paraded numerous ai-driven features without mentioning the term. When prodded, its executives lean towards extolling "machine learning", another term for ai.

That looks smart. The meltdown at Openai shows just how damaging the culture wars over ai can be. But it is these wars that will shape how the technology progresses, how it is regulated—and who comes away with the spoils. .