Hey ChatGPT, Which One of These Is the Real Sam Altman?

Two journalists explore the artificial intelligence company OpenAI and present complementary portraits of its notorious co-founder.

May 19, 2025

The OpenAI executive Sam Altman testifies at a Senate hearing in May 2025 about the global artificial intelligence race.Chip Somodevilla/Getty Images

By Tim Wu

Tim Wu is a law professor at Columbia University and the author of the forthcoming "The Age of Extraction: How Tech Platforms Conquered the Economy and Threaten Our Future Prosperity."

EMPIRE OF AI: Dreams and Nightmares in Sam Altman's OpenAI, by Karen Hao

THE OPTIMIST: Sam Altman, OpenAI, and the Race to Invent the Future, by Keach Hagey

The "paper clip problem" is a well-known ethics thought experiment in the world of artificial intelligence. It imagines a superintelligent A.I. charged with the seemingly harmless goal of making as many paper clips as possible. Trouble is, as the philosopher Nick Bostrom put it in 2003, without common-sense limits it might transform "first all of earth and then increasing portions of space into paper clip manufacturing facilities." The tale has long served as a warning about objectives pursued too literally.

Two new books that orbit the entrepreneur Sam Altman and the firm he co-founded, OpenAI, suggest we may already be living with a version of the problem. In "Empire of AI," the journalist Karen Hao, who has worked for The Wall Street Journal and contributes to The Atlantic, argues that the pursuit of an artificial superintelligence has become its own figurative paper clip factory, devouring too much energy, minerals and human labor. Meanwhile, "The Optimist," by the Wall Street Journal reporter Keach Hagey, leaves readers suspecting that the earnest and seemingly innocuous paper clip maker who ends up running the world for his own ends could be Altman himself.

“Empire of AI” is the broader and more critical of the two. Hao profiled OpenAI in 2020, two years before its most famous product, the intelligent chatbot called ChatGPT, debuted publicly. She portrays OpenAI and other companies that make up the fast-growing A.I. sector as a “modern-day colonial world order.” Much like the European powers of the 18th and 19th centuries, they “seize and extract precious resources to feed their vision of artificial intelligence.” In a corrective to tech journalism that rarely leaves Silicon Valley, Hao ranges well beyond the Bay Area with extensive fieldwork in Kenya, Colombia and Chile.

“The Optimist” is a more conventional biography, concentrated on Altman’s life and times. Born in Chicago to progressive parents named Connie and Jerry — in the 1980s, Jerry innovated a way to stir investment in affordable housing — Altman was heavily influenced by their do-gooder spirit. (“You can’t out-nice Jerry,” his friends would say.) Altman’s relentlessly upbeat manner and genuine technical skill made him a perfect fit for Silicon Valley. Charming and smart, he tells people what they want to hear and has a knack for talking big in exactly the way 2010s Bay Area investors liked.

The arc of Altman’s life also follows a classic script. He drops out of Stanford to launch a start-up that fizzles, but the effort brings him to the attention of Paul Graham, the co-founder of Y Combinator, an influential tech incubator that launched companies like Airbnb and Dropbox. By age 28, Altman has risen to succeed Graham as the organization’s president, setting the stage for his leadership in the A.I. revolution.

As Hagey makes clear, success in this context is all about the way you use the people you know. The author supplies a meticulous account of 21st-century networking culture in Silicon Valley, where Altman’s technical talents end up being less important than some of the qualities usually

associated with religious leaders — dare-to-dream boldness, self-effacing geniality and, she writes, "a skill for convincing people that he can see into the future." Not unlike ChatGPT, Altman molds himself into whatever people want him to be. As Graham once quipped, "You could parachute him into an island full of cannibals and come back in five years and he'd be the king."

During the 2010s Altman joined a group of Silicon Valley investors determined to recover the grand ambitions of earlier tech eras. Tired of start-ups based on incremental tweaks to social-media platforms or gig-work apps, they sought to return to outer space, unlock nuclear fusion, achieve human-level A.I. and even defeat death itself.

The investor Peter Thiel was a major influence, but Altman's most important collaborator in the field of A.I. was Elon Musk. The early-2010s Musk who appears in both books is almost unrecognizable to observers who now associate him with black MAGA hats and chain-saw antics. This Musk, the builder of Tesla and SpaceX, believes that creating superintelligent computer systems is "summoning the demon." He becomes obsessed with the idea that Google will soon develop a true artificial intelligence and allow it to become a force for evil.

Altman, dining regularly with Musk, mirrors his anxieties and persuades him to bankroll a more idealistic rival. "If it's going to happen," Altman emailed Musk in 2015, "it seems like it would be good for someone other than Google to do it first." He pitched a "Manhattan Project for A.I.," a nonprofit to develop a good A.I. in order to save humanity from its evil twin, just as the actual Manhattan Project sought to outrace the Nazis to the atomic bomb. Musk guaranteed $1 billion and even supplied the name OpenAI.

Hagey's book, written with Altman's cooperation, is less critical, but no hagiography. "The Optimist" lets the reader see how thoroughly Altman

outfoxed his patron, leveraging Musk's paranoia into enormous sums of money while slowly making OpenAI his own. It's striking that, despite providing much of the initial capital and credibility, Musk ends up with almost nothing to show for his investment.

Hao's 2020 profile of OpenAI, published in the M.I.T. Technology Review, was unflattering and the company declined to cooperate with her for her book. She believes that OpenAI was "begun as a sincere stroke of idealism," but she wants to make its negative spillover effects evident. Hao does an admirable job of pulling the camera back, telling the stories of workers in Nairobi who earn "starvation wages to filter out violence and hate speech" from ChatGPT, and of visits to communities in Chile where data centers siphon prodigious amounts of water and electricity to run complex hardware.

Both books climax with the weekend in November 2023 when Altman was abruptly fired by his company's board, only to be reinstated days later after staff members and investors revolted. From the outside, many critics saw the coup as a last-ditch effort to stop OpenAI from becoming the very Eye of Sauron it was founded to restrain.

Hagey renders this moment as a conventional board mutiny: Directors had tired of Altman's "duplicity and calamitous aversion to conflict." One of them, the OpenAI co-founder Ilya Sutskever, recalls that Altman "would tell him one thing, then say another, and act as if the difference was an accident." Sutskever said he regretted his vote to oust Altman, but after Altman returned to OpenAI, Sutskever left the company.

Hao's version is darker. Relying on a lot of the same sources as Hagey, she presents Altman as a peddler of "many little lies and some big ones," who helped create "a directionless, chaotic and back-stabbing environment." In her book, Sutskever comes to see Altman as engaging in what some of his

colleagues call "psychological abuse."

Together, these two excellent and deeply reported books form a diptych. On one panel stands Altman as the secular prophet preaching human progress and boundless optimism. Hagey calls Altman a "brilliant deal maker with a need for speed and a love of risk, who believes in technological progress with an almost religious conviction." On the other panel is Altman the opportunist. He uses idealism as a tool, harnessing the concept of human progress to build an empire the way Europeans once used Christianity to justify conquest.

Altman recently told the statistician Nate Silver that if we achieve human-level A.I., "poverty really does just end." But motives matter. History suggests that some technologies aimed at growth have taken a bad situation and made it worse. The efficiencies of the cotton gin, for instance, saved on labor but made slavery even more lucrative. If the aim is not, in the first place, to help the world, but instead to get bigger — better chips, more data, smarter code — then our problems might just get bigger too.

(The New York Times has sued OpenAI and its partner, Microsoft, accusing them of copyright infringement regarding news content related to A.I. systems. OpenAI and Microsoft have denied those claims.)

EMPIRE OF AI: **Dreams and Nightmares in Sam Altman's OpenAI** | By Karen Hao | Penguin Press | 482 pp. | $32

THE OPTIMIST: **Sam Altman, OpenAI, and the Race to Invent the Future** | By Keach Hagey | Norton | 367 pp. | $31.99