Journal of American Studies (2026), 1–6
doi:10.1017/S0021875826102102

CAMBRIDGE UNIVERSITY PRESS

REVIEW ESSAY

Empire, AI, and Inevitability

Karen Hao, *Empire of AI: Dreams and Nightmares in Sam Altman's OpenAI* (New York: Penguin Press, 2025, $32.00). Pp. 496. ISBN 978 0 5936 5750 8.

Adam Becker, *More Everything Forever: AI Overlords, Space Empires, and Silicon Valley's Crusade to Control the Fate of Humanity* (New York: Basic Books, 2025, $32.00). Pp. 384. ISBN 978 1 5416 1959 3.

Trevor Owens

American University, USA

I often find myself wondering how historians a century or two from now will make sense of this digital age. While that is not a common preoccupation, it's one that comes naturally to those of us who work with cultural memory institutions to ensure enduring access to digital content. Undoubtedly, those future historians will need to expend considerable effort making sense of how artificial intelligence became a focal point for global capital and infrastructure. They will be well served by two books that came out in 2025, Karen Hao's *Empire of AI: Dreams and Nightmares in Sam Altman's OpenAI* and Adam Becker's *More Everything Forever: AI Overlords, Space Empires, and Silicon Valley's Crusade to Control the Fate of Humanity*.

Hao's and Becker's books are extensively researched efforts to understand the people, institutions, and ideologies that converged to make the artificial-intelligence business what it is today. That bit on ideologies is particularly significant, as it clarifies how very little about these technologies was inevitable. Hao focusses on the rise of OpenAI and its CEO, Sam Altman. Becker situates the present AI moment in a broader context of Silicon Valley ideologies like long-termism and transhumanism. The books pair well together. Hao gives us a zoomed-in detailed account of how OpenAI emerged as one of the defining companies of our age, while Becker provides a wider intellectual framework for understanding how assumptions driving this age of AI took shape.

Synthesizing years of reporting, Karen Hao reconstructs the rise of OpenAI and Altman. While situating Altman in both the social and financial networks of Silicon Valley, this deeply reported work of investigative journalism draws on interviews with more than 260 people and an extensive trove of correspondence and internal documents as sources. I hope that wealth of primary sources will itself make its way

© The Author(s), 2026. Published by Cambridge University Press in association with British Association for American Studies.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 07:21:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0021875826102102

into an archive for future researchers to access. Among those interviewed are over ninety former and current OpenAI employees and executives; dozens of individuals connected to Microsoft, Anthropic, Meta, and Google; and others involved in the development of AI systems.

Hao opens *Empire of AI* close to where the book will ultimately finish. It's Friday, 17 November 2023. Sam Altman, by then a world-renowned billionaire and technology leader, joins a Google Meet call with four members of OpenAI's board to learn that he is being fired. His ousting is short-lived. Within a few pages, OpenAI employees, investors, and allies at Microsoft have successfully rallied to reinstate him. This coup has failed. From there, Hao brings us back to 2015, where we find Sam Altman already fully engaged in shaping the future as the leader of the start-up accelerator Y Combinator and looking for collaborators to bring about a radical transformation in society through the development of artificial general intelligence (AGI). There we begin our journey tracing the tangled financial and ideological history that led OpenAI from being a seemingly idealistic nonprofit research initiative to becoming one of the most consequential technology companies of our time. The story is one of a corporate intrigue page-turner, with stakes extending far beyond Silicon Valley.

Hao's use of empire in the title signals a central theme. While acknowledging that contemporary AI firms are "not engaged in the same overt violence and brutality" associated with earlier colonial projects, she argues they nevertheless "seize and extract precious resources to feed their vision of artificial intelligence," including the creative labor that produced their training data and the "land, energy, and water required to house and run massive data centers and supercomputers" (17). For scholars in American studies or science and technology studies, this empire framing is likely a familiar one. Hao delivers on the interpretive value of this framing through careful, grounded reporting on the institutions, infrastructures, labor systems, and forms of extraction involved in contemporary AI development.

Hao is adept at showing how the contemporary AI context emerged from the rivalries and anxieties of a small social network of founders, investors, and researchers connected to Bay Area start-ups and technology companies. She invites us to consider Elon Musk's forty-second birthday party. It's 2013 in Napa Valley, and the birthday boy has gotten into an argument with Larry Page of Google. The topic? Whether artificial intelligence surpassing human abilities would be a significant problem. Musk said yes. In reply, Page called him a "specist." The remark alarmed Musk, so much so as to convince him that it would be dangerous for Google to be the organization that ultimately developed truly artificial general intelligence (24). From here, Hao fast-forwards by two years to 2015, at another dinner with Musk. This time, Musk finds himself in the company of Sam Altman, who has a vision for a nonprofit that would be "OpenAI, the anti-Google." The vision is committed to open science, transparency, and broadly shared benefit (49). These ambitions, and the belief that the creation of AGI could inaugurate an entirely new era of human history, runs throughout Hao's account and ultimately points toward the broader ideological terrain that Becker explores.

Hao demonstrates that AGI is strategically both grand and vague. Early in the book, she argues that "AGI is largely rhetorical," serving as "a fantastical, all-purpose excuse for OpenAI to continue pushing for ever more wealth and power" (49). Throughout the narrative, Hao returns to what she calls the inevitability card, the

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 07:21:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0021875826102102

insistence that increasingly powerful AI systems are coming regardless. Given this assumption, Musk, Altman, and others are positioning themselves as the only people capable of managing their arrival responsibly. Under this logic, OpenAI and related firms are simultaneously the creators of potentially catastrophic technologies and the only institutions capable of safely guiding society through their development. Hao writes that "under the specter of AGI's unstoppable arrival, the company needed to keep developing more and more powerful models to prepare itself and prepare society" (131). In this framework, the speculative possibility of future superintelligence becomes a justification for the concentration of capital, infrastructure, expertise, and political influence in the present.

Hao also demonstrates how the contemporary AI boom marshals resources toward rapid expansion of computational infrastructures. Early in OpenAI's history, its researchers coined and promulgated OpenAI's Law, the notion that major advances in AI capabilities result primarily from dramatic increases in computing power. This emphasis on scale over all else pushed the company toward ever larger data sets, data centers, and capital requirements. This function becomes critical in explaining Microsoft's massive company-saving investment in OpenAI that came after Musk's departure. Hao notes that, within Microsoft, the investment was less about believing that OpenAI would fulfil its core mission of bringing about AGI, and more about the infrastructural advantages that would result from participating in the race to invest in the quest for AGI. As one Microsoft executive observed, the scale of their "ambition is driving everything from data center design to compute silicon to networks and distributed systems architectures" (Hao, 72). Microsoft, and the market more broadly, did not need to believe in the impending arrival of AGI. What mattered was that the pursuit of AGI was driving extraordinary investment and experimentation across nearly every layer of computing infrastructure.

Hao's empire framing fully earns its value as she expands outward from Silicon Valley to the labor and environmental impacts, enabling contemporary AI development around the world. She reports from Kenya and Venezuela, where workers perform low-paid and often psychologically traumatic labor, tagging data and moderating disturbing content used to train AI systems. She traces the growing demands that AI data centers place on water and energy infrastructures in places like Chile and Uruguay. Hao argues that contemporary AI advancement increasingly operates through a scale-at-all-costs logic that concentrates resources and infrastructural power while obscuring the human and environmental costs required to sustain it. As many of its advocates frame AGI as an existential technological imperative, extraordinary trade-offs become easier and easier to justify as necessary. Many of the same actors advancing increasingly energy-intensive AI infrastructures insist that AGI itself is the only thing capable of providing the tools necessary to solve climate change and other seemingly intractable global problems. Whatever the question, the answer in this framework is to pour more resources into accelerating AI development.

Hao's account undermines the inevitability narrative. Readers follow the OpenAI team stumbling across a wide range of competing approaches for getting to AGI. She shows that, for many of its early years, "OpenAI had no idea what it was doing" (58). Researchers pivoted between robotics projects, video game experiments, language models, and other exploratory initiatives without consensus about

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 07:21:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0021875826102102

what path would ultimately matter most. Even the launch of ChatGPT, the product that made OpenAI a household name, was framed as a relatively "low-key research preview" when it appeared in late 2022 (258). The scale of the public response rapidly transformed both the company and the broader technology industry. ChatGPT famously reached a million users in its first five days and a hundred million users within two months, becoming the fastest-growing consumer application in history. Hao's account is particularly effective at holding these two realities together: the path toward the present AI moment was highly contingent, but the technology that has been realized has a very real user base. While there is unprecedented hype surrounding AI technologies, the extraordinary uptake of ChatGPT demonstrates that these systems connect with forms of curiosity, experimentation, and everyday utility that hundreds of millions of people around the world would rapidly incorporate into everyday life.

Hao's book covers many internal struggles, ideological fractures, and institutional dramas that some readers will find more granular than necessary, from the failed boardroom coup against Altman to "The Divorce," which led many members of OpenAI's internal AI safety contingent to break away and found their competitor, Anthropic. Yet part of the long-term value of Hao's work lies precisely in the density of this documentation. Given how little of the internal record of companies like OpenAI is likely to ultimately survive and become accessible to future researchers, Hao's extensive interviews and documentation are likely to represent an important historical resource.

Documenting the unstable and elastic idea of AGI throughout this history is itself a significant contribution of Hao's book. In the story, AGI functions as a technological MacGuffin. Late in the book, Altman explains AGI, the thing he founded OpenAI to bring about, to the *New York Times* as "a ridiculous and meaningless term," adding, "I apologize that I keep using it" (401). Regardless, the dream of AGI has proven extraordinarily powerful as an organizing vision for accelerating and concentrating resources. Hao's story of OpenAI shows how speculative visions of future intelligence drive very real systems of labor, extraction, governance, and economic concentration in the present. Where Hao gives readers the institutional and infrastructural history of how this contemporary AI moment emerged, Becker helps explain the deeper intellectual and ideological traditions that made this worldview legible and compelling to so many influential figures within Silicon Valley.

"The dream is always the same," Adam Becker opens *More Everything Forever*, "go to space and live forever" (1). In Becker's account, the aspirations driving contemporary AI work are tied to something much bigger. They are part of a worldview organized around immortality and escape from all earthly limits. Becker situates figures like Eliezer Yudkowsky, founder of the Machine Intelligence Research Institute, within an ideological ecosystem that treats AI as a key component through which tech oligarchs imagine they will achieve "perfect health, immortality," and eventually interplanetary civilization (1). Becker's intervention is useful because it clarifies the intellectual terrain in which to situate Hao's reporting. Her account of OpenAI and the AI industry is also an account of the persistence of a bizarre Silicon Valley eschatology that now increasingly shapes the assumptions of major global institutional actors.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 07:21:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0021875826102102

One of Becker's most useful contributions is his explanation of how philosophies like long-termism, which make brief appearances in Hao's book, came to exert such influence within contemporary technology culture. In the context of long-termist reasoning, the potential impact on the lives of vast numbers of future humans, who may exist tens or even hundreds of thousands of years from now, greatly outweighs any present-day suffering. What types of things can be justified if they could benefit trillions of imagined future people? This framework enables its acolytes to "ignore conventional morality," justifying extraordinary actions in the name of safeguarding humanity's long-term future (Becker, 29).

Becker takes us from long-termist to transhumanist thought centered on overcoming biological and material limits altogether through consideration of the work of figures like Ray Kurzweil. Kurzweil's influential articulation of "the Singularity," the imagined moment when artificial intelligence radically exceeds human intelligence and fundamentally transforms civilization, is an animating vision driving much of the work toward AGI. Kurzweil insists that, by 2045, "we will multiply our effective intelligence by a billionfold by merging with the intelligence we have created" (Becker, 41). These ideas are no longer fringe curiosities at the edges of tech culture. Kurzweil has now spent more than a decade as a Google employee, where he continues to serve as a prominent futurist and AI visionary. Becker's analysis provides useful context for Hao's account of Altman's investments in companies focussed on radical life extension and the possibility of eventually uploading human consciousness into computers. Becker shows that these projects are not unrelated eccentric lines of thinking. These fantasies of transcendence and immortality serve to justify ever greater doubling down on the expansion of computing power in a war against all limitations regarding the human experience.

Becker also provides context for understanding contemporary AI safety discourse. Speculative thought experiments about superintelligent machines, such as Nick Bostrom's "paperclip maximizer" scenario in which a superintelligent AI is instructed to manufacture paperclips and ultimately converts all available matter, including human beings, into additional paperclips, serve as the foundation for AI safety frameworks. This AI safety line of thinking moved from being a fringe concern two decades ago into operational and policy discourse as a key concern. Becker demonstrates how scenarios like this helped establish the existential-risk framing in Silicon Valley, which has then been exported out to various policy contexts. The fact that these scenarios likely do not represent plausible futures is beside the point. They operate in the present to focus attention and resources on AI development and divert attention from the immediate environmental and social harms already associated with contemporary AI systems and their development.

In this worldview, technological progress is a moral necessity. We read Marc Andreessen insisting that slowing development of AI would constitute "a form of murder" because of all the lives future AI systems might eventually save (Becker, 207). From this frame, tech thought leaders like Andreessen insist that notions of "sustainability" and "social responsibility" are themselves the enemy because they act as obstacles to the technological acceleration necessary to avert that form of future murder. We find related themes when Jeff Bezos rhetorically asks, "Do we want stasis and rationing or do we want dynamism and growth," later explaining that he would "love to see a trillion humans living in the solar system" (Becker,

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 07:21:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0021875826102102

222). This anti-limits worldview explains the extraordinary urgency surrounding contemporary AI development, and the related focus on life extension and establishing colonies on Mars. At the bottom of all of this we find Jeff Bezos in a space suit and cowboy hat and Elon Musk wearing a T-shirt that says, "Nuke Mars." All of which sounds like a remix and restatement of Turner's frontier thesis, turned outward from our Earth toward conquest and consumption of the whole of the universe itself. The title of Becker's book is the kicker. The goal itself is more everything forever. Whatever might get in the way of that vision is flagged as something to be overcome.

Taken together, Hao and Becker provide essential context for understanding the AI moment we find ourselves in. Underneath the hype and dramatic growth of the AI industry, we find a buffet of questionable fringe theories. Collectively those theories work to advance moral arguments that justify ever greater acceleration of resource extraction and consumption. At the root, we find a fear of an end to unending growth.

By engaging deeply and seriously with the histories of OpenAI, Altman, and the development and dissemination of various ideological schools of thought, these books help undermine the aura of inevitability surrounding AI. Becker concludes by arguing that many technology leaders insist that their futures are inevitable precisely to discourage others from recognizing how contingent technological development really is. By showing us how OpenAI itself emerged through a series of rivalries and institutional accidents rather than any singular or coherent technological destiny, Hao's account reinforces this point.

Hao ends *Empire of AI* by briefly turning to consider alternative models for AI development. She discusses projects focussed on reciprocity, consent, data sovereignty, and improved conditions for the workers whose labor sustains these systems. Other futures can be dreamed and manifested, and unmasking and understanding how we have arrived where we have is an important component of opening up those possibilities.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 15 Nov 2026 at 07:21:07, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/S0021875826102102