

LBRIS

We know
books

KAREN HAO

Empire of AI

*Inside the Reckless Race
for Total Domination*



PENGUIN BOOKS

CONTENTS

AUTHOR'S NOTE	xi
----------------------	----

PROLOGUE A Run for the Throne	1
--------------------------------------	----------

I

1 Divine Right	23
2 A Civilizing Mission	46
3 Nerve Center	73
4 Dreams of Modernity	88
5 Scale of Ambition	117

II

6 Ascension	141
7 Science in Captivity	158
8 Dawn of Commerce	175
9 Disaster Capitalism	189

III

10	Gods and Demons	227
11	Apex	256
12	Plundered Earth	271
13	The Two Prophets	301
14	Deliverance	326

IV

15	The Gambit	343
16	Cloak-and-Dagger	361
17	Reckoning	377
18	A Formula for Empire	399

EPILOGUE	How the Empire Falls	409
-----------------	----------------------	-----

ACKNOWLEDGMENTS	423
------------------------	-----

NOTES	427
--------------	-----

INDEX	469
--------------	-----

LBRIS

We know
books

Empire of AI

Prologue

A Run for the Throne

On Friday, November 17, 2023, around noon Pacific time, Sam Altman, CEO of OpenAI, Silicon Valley's golden boy, avatar of the generative AI revolution, logged on to a Google Meet to see four of his five board members staring at him.

From his video square, board member Ilya Sutskever, OpenAI's chief scientist, was brief: Altman was being fired. The announcement would go out momentarily.

Altman was in his room at a luxury hotel in Las Vegas to attend the city's first Formula One race in a generation, a star-studded affair with guests from Rihanna to David Beckham. The trip was a short reprieve in the middle of the punishing travel schedule he had maintained ever since the company released ChatGPT about a year earlier. For a moment, he was too stunned to speak. He looked away as he sought to regain his composure. As the conversation continued, he tried in his characteristic way to smooth things over.

"How can I help?" he asked.

The board told him to support the interim chief executive they had selected, Mira Murati, who had been serving as his chief technology officer. Altman, still confused and wondering whether this was a bad dream, acquiesced.

Minutes later, Sutskever sent another Google Meet link to Greg Brockman, OpenAI's president and a close ally to Altman who had been the only board member missing from the previous meeting. Sutskever told Brockman he would no longer be on the board but would retain his role at the company.

The public announcement went up soon thereafter. "Mr. Altman's departure follows a deliberative review process by the board, which concluded that he was not consistently candid in his communications with the board, hindering its ability to exercise its responsibilities. The board no longer has confidence in his ability to continue leading OpenAI."

On the face of it, OpenAI had been at the height of its power. Ever since the launch of ChatGPT in November 2022, it had become Silicon Valley's most spectacular success story. ChatGPT was the fastest-growing consumer app in history. The startup's valuation was on the kind of meteoric ascent that made investors salivate and top talent clamor to join the rocket-ship company. Just weeks before, it had been valued at up to \$90 billion as part of a tender offer it was in the middle of finalizing that would allow employees to sell their shares to said eager investors. A few days before, it had held a highly anticipated and highly celebrated event to launch its most aggressive slate of products.

Altman was, as far as the public was concerned, the man who had made it all happen. He had spent the spring and summer touring the world, reaching a level of celebrity that was leading the media to compare him to Taylor Swift. He had wowed just about everyone with his unassuming small frame, bold declarations, and apparent sincerity.

Before Vegas, he had once again been globe-trotting, sitting on a panel at the APEC CEO Summit, delivering lines with his usual dazzling effect.

"Why are you devoting your life to this work?" Laurene Powell Jobs, founder and president of the Emerson Collective and Steve Jobs's widow, had asked him.

"I think this will be the most transformative and beneficial technology humanity has yet invented," he said. "Four times now in the history

LBPIS

We know
books

of OpenAI—the most recent time was just in the last couple of weeks—I have gotten to be in the room, when we sort of push the veil of ignorance back and the frontier of discovery forward, and getting to do that is, like, the professional honor of a lifetime.”

Shocked employees learned about Altman’s firing just as everyone else did, the link to the public announcement zipping from one phone to the next across the company. It was the chasm between the news and Altman’s glowing reputation that startled them the most. The company was by now pushing eight hundred people. These days, employees had fewer opportunities to meet and interact with their CEO in person. But his charming demeanor on global stages was not unlike how he behaved during all-hands meetings, at company functions, and, when he wasn’t traveling, around the office.

As the rumor mill kicked into a frenzy and employees doomscrolled X, formerly Twitter, for any shreds of information, someone in the office latched on to what they saw as the most logical explanation and shouted, “Altman’s running for president!” It created a momentary release of tension, before people realized this was not the case, and speculation started anew with fresh intensity and dread. Had Altman done something illegal? Maybe it was related to his sister, employees wondered. She had alleged in tweets that had gone viral a month before that her brother had abused her. Maybe it wasn’t something illegal but ethically untoward, they speculated, perhaps related to Altman’s other investments or his fundraising with Saudi investors for a new AI chip venture.

Sutskever posted an announcement in OpenAI’s Slack. In two hours, he would hold a virtual all-hands meeting to answer employee questions. “That was the longest two hours ever,” an employee remembers.

Sutskever, Murati, and OpenAI’s remaining executives came onto the screen side by side, stiff and unrehearsed, as the all-hands streamed to employees in the office and working from home.

Sutskever looked solemn. He was known among employees as a deep thinker and a mystic, regularly speaking in spiritual terms with a force

IBDIS

We know
books

of sincerity that could be endearing to some and off-putting to others. He was also a goofball and gentlehearted. He wore shirts with animals on them to the office and loved to paint them as well—a cuddly cat, cuddly alpacas, a cuddly fire-breathing dragon—alongside abstract faces and everyday objects. Some of his amateur paintings hung around the office, including a trio of flowers blossoming in the shape of OpenAI's logo, a symbol of what he always urged employees to build: "A plurality of humanity-loving AGIs."

Now, he attempted to project a sense of certainty to anxious employees submitting rapid-fire questions via an online document. But Sutskever was an imperfect messenger; he was not one that excelled at landing messages with his audience.

"Was there a specific incident that led to this?" Murati read aloud first from the list of employee questions.

"Many of the questions in the document will be about the details," Sutskever responded. "What, when, how, who, exactly. I wish I could go into the details. But I can't." Anyone curious should read the press release, he added. "It actually says a lot of stuff. Read it maybe a few times."

The response baffled employees. They had just received cataclysmic news. Surely, as the people most directly affected by the situation, they deserved more specifics than the general public.

Murati read off a few more questions. How did this affect the relationship with Microsoft? Microsoft, OpenAI's biggest backer and exclusive licensee of its technologies, was the sole supplier of its computing infrastructure. Without it, all the startup's work—performing research, training AI models, launching products—would grind to a halt. Murati responded that she didn't expect it to be affected. They had just had a call with Microsoft's chief executive Satya Nadella and chief technology officer Kevin Scott. "They're all very committed to our work," she said.

What about OpenAI's tender offer? Employees with a certain tenure had been given the option to sell what could amount to millions of dollars' worth of their equity. The tender was so soon that many had made plans to buy property, or already had. "The tender—we're, um, we're going to see," Brad Lightcap, the chief operating officer, waffled. "I am in

touch with investors leading the tender and some of our largest investors already on the cap table. All have committed their steadfast support to the company.”

After several more questions were met with vague responses, another employee tried again to ascertain what Sam had done. Was this related to his role at the company? Or did it involve his personal life? Sutskever once again directed people to the press release. “The answer is actually there,” he said.

Murati read on from the document. “Will questions about details be answered at some point or never?”

Sutskever responded: “Keep your expectations low.”

As the all-hands continued and Sutskever’s answers seemed to grow more and more out of touch, employee unease quickly turned into anger.

“When a group of people grow through a difficult experience, they often end up being more united and closer to each other,” Sutskever said. “This difficult experience will make us even closer as a team and therefore more productive.”

“How do you reconcile the desire to grow together through crisis with a frustrating lack of transparency?” an employee wrote in. “Typically truth is a necessary condition for reconciliation.”

“I mean, fair enough,” Sutskever replied. “The situation isn’t perfect.”

Murati tried to quell the rising tension. “The mission is so much bigger than any of us,” she said.

Lightcap echoed her message: OpenAI’s partners, customers, and investors had all stressed that they continued to resonate with the mission. “If anything, we have a greater duty now, I think, to push hard on that mission.”

Sutskever again attempted to be reassuring. “We have all the ingredients, all of them: The computer, the research, the breakthroughs are astounding,” he said. “When you feel uncertain, when you feel scared, remember those things. Visualize the size of the cluster in your mind’s eye. Just imagine all those GPUs working together.”

An employee submitted a new question. “Are we worried about the hostile takeover via coercive influence of the existing board members?” Sutskever read.

“Hostile takeover?” Sutskever repeated, a new edge in his voice. “The OpenAI nonprofit board has acted entirely in accordance to its objective. It is not a hostile takeover. Not at all. I disagree with this question.”

That night, several employees gathered at a colleague’s house for a party that had been planned before Altman’s firing. There were guests from other AI companies as well, including Google DeepMind and Anthropic.

Right before the event, an alert went out to all attendees. “We are adding a second themed room for tonight: ‘The no-OpenAI talk room.’ See you all!” In the end, few people stayed long in the room. Most people wanted to talk about OpenAI.

Brockman had announced that afternoon that he was quitting in protest. Microsoft’s Nadella, who had been furious about being told about Altman’s firing only minutes before it happened, had put out a carefully crafted tweet: “We have a long-term agreement with OpenAI with full access to everything we need to deliver on our innovation agenda and an exciting product roadmap; and remain committed to our partnership, and to Mira and the team.”

As rumors continued to proliferate, word arrived that three more senior researchers had quit the company: Jakub Pachocki and Szymon Sidor, early employees who had among the longest tenures at OpenAI, and Aleksander Mądry, an MIT professor on leave who had joined recently. Their departures further alarmed some OpenAI employees, a signal of a bleeding out of leadership and talent that could spook investors and halt the tender offer or, worse, ruin the company. At the party, employees grew more and more despondent and agitated. A dissolution of the tender offer would snatch away a significant financial upside to all their hard labor, to say nothing of a dissolution of the company, which would squander so much promise and hard work.

Also that night, the board and the remaining leadership at the company were holding a series of increasingly hostile meetings. After the all-

LIBRIS

We know
books

hands, the false projection of unity between Sutskever and the other leaders had collapsed. Many of the executives who had sat next to Sutskever during the livestream had been nearly as blindsided as the rest of the staff, having learned of Altman's dismissal moments before it was announced. Riled up by Sutskever's poor performance, they had demanded to meet with the rest of the board. Roughly a dozen executives, including Murati and Lightcap, had gathered in a conference room at the office.

Sutskever was dialed in virtually along with the three independent directors: Adam D'Angelo, the cofounder and CEO of the question-and-answer site Quora; Tasha McCauley, an entrepreneur and adjunct senior management scientist at the policy think tank RAND; and Helen Toner, an Australian-born researcher at another think tank, Georgetown University's CSET, or Center for Security and Emerging Technology.

Under an onslaught of questions, the four board members repeatedly evaded making further disclosures, citing their legal responsibilities to protect confidentiality. Several leaders grew visibly enraged. "You're saying that Sam is untrustworthy," Anna Makanju, the vice president of global affairs, who had often accompanied Altman on his global charm offensive, said furiously. "That's just not our experience with him at all."

The gathered leadership pressed the board to resign and hand their seats to three employees, threatening to all quit if the board didn't comply immediately. Jason Kwon, the chief strategy officer, a lawyer who had previously served as OpenAI's general counsel, upped the ante. It was in fact *illegal* for the board not to resign, he said, because if the company fell apart, this would be a breach of the board members' fiduciary duties.

The board members disagreed. They maintained that they had carefully consulted lawyers in making the decision to fire Altman and had acted in accordance with their delineated responsibilities. OpenAI was not like a normal company, its board not like a normal board. It had a unique structure that Altman had designed himself, giving the board broad authority to act in the best interest not of OpenAI's shareholders but of its mission: to ensure that AGI, or artificial general intelligence, benefits humanity. Altman had long touted the board's ability to fire him as its most important governance mechanism. Toner underscored the

LBPIS

We know
books

point: "If this action destroys the company, it could in fact be consistent with the mission."

The leadership relayed her words back to employees in real time: Toner didn't care if she destroyed the company. Perhaps, many employees began to conclude, that was even her intention. At the thought of losing all of their equity, a person at the party began to cry.

The next day, Saturday, November 18, dozens of people, including OpenAI employees, gathered together at Altman's \$27 million mansion to await more news.

The three senior researchers who had quit, Pachocki, Sidor, and Mađry, had met with Altman and Brockman to talk about re-forming the company and continuing their work. To some, word of their discussions increased employee anxiety: A new OpenAI competitor could intensify the instability at the company. To others it offered hope: If Altman indeed founded a new venture, they would leave to go with him.

OpenAI's remaining leadership gave the board a deadline of 5 p.m. Pacific time that day: Reinstate Altman and resign, or risk a mass employee exodus from the company. The board members refused. Through the weekend, they frantically made calls, sometimes in the middle of the night, to anyone on their roster of connections who would pick up. In the face of mounting ire from employees and investors over Altman's firing, Murati was no longer willing to serve as interim CEO. They needed to replace her with someone who could help restore stability, or find new board members who could hold their own against Altman if he actually came back.

That night, after the deadline came and went, Jason Kwon sent a memo to employees. "We are still working towards a resolution and we remain optimistic," he wrote. "By resolution, we mean bringing back Sam, Greg, Jakub, Szymon, Aleksander."

Altman tweeted in his signature lowercase style. "i love the openai team so much."

Dozens of other employees began retweeting it with a heart emoji.

On Sunday, Altman and Brockman arrived back at the office to negotiate their return. Over the course of the day, more and more employees joined them to wait in suspense. By then, most employees, leadership, and the board had barely slept in more than thirty-six hours; everything was beginning to blur together. Altman tweeted a selfie, lips pursed, brows furrowed, displaying a guest badge in his hand. “first and last time i ever wear one of these,” he added as the caption. Leadership set another 5 p.m. deadline for the board to reinstate Altman and to resign.

The pressure was now piling on from all directions. Microsoft, OpenAI’s other investors, and heavyweights across Silicon Valley were publicly siding with Altman. A source relayed the playbook to the media: Not only would employees leave en masse if the decision were not reversed, but Microsoft would withhold access to its computing infrastructure, and investors would file lawsuits. The combination would make an OpenAI without Altman untenable.

Still, the board continued to resist. Nearing 9 p.m., once again well past the latest deadline, Sutskever posted a long message on Slack on behalf of the board. Altman was not returning; Emmett Shear, the former CEO of Twitch, was now the new interim head of OpenAI. He and Shear would arrive at the office in five minutes to give a speech about the company’s new vision.

“The board firmly stands by its decision as the only path to advance and defend the mission of OpenAI,” he wrote. “Put simply, Sam’s behavior and lack of transparency in his interactions with the board undermined the board’s ability to effectively supervise the company in the manner it was mandated to do.”

The Slack instantly lit up with dozens of angry replies from employees.

“You and what fucking army”

“you’re delusional”

“Emmett will be the CEO of nothing”

Roughly two hundred employees paraded out of the office to boycott the talk. Murati rushed the executives out the building. By the time