LETTER FROM SILICON VALLEY

SAM ALTMAN WON IN COURT AGAINST ELON MUSK. BUT, REALLY, WE ALL LOST

The cases of Elizabeth Holmes and Sam Bankman-Fried at least offered a pleasant sense of comeuppance. But in Musk v. Altman, to root against Tweedledum was effectively to root for Tweedledee.

By Gideon Lewis-Kraus

May 20, 2026



Illustration by Joan Wong; Source photographs from Getty

Save this story

Listen • 23 minutes

A famous logic puzzle takes place on a mythical island divided between the knights, who never lie, and the knaves, who always do. A foreign traveller encounters a fork in the road: one way guarantees safe passage, the other certain death. A member of each tribe is present, though it isn't clear which is which, and the traveller is granted only one question. The solution is well known: ask either of them what the other would advise, and then to choose the opposite path. (An accurate account of a lie and an inaccurate account of the truth amount to the same wrong answer.) But this works only if *someone* is honest. What if nobody can be trusted? The Cretan philosopher Epimenides inspired an alternative scenario set on his own island, when he supposedly said that "all Cretans are liars." Logicians call unstable statements like these "self-referential paradoxes," or utterances that undermine their own claims. Epimenides would presumably have felt at home at the trial of Musk v. Altman, which over the past few weeks turned an Oakland courtroom into an island of lying cretins.

In theory, the trial was about the good-faith control of artificial intelligence. In 2015, Elon Musk and Sam Altman founded OpenAI together as a nonprofit. Its mission—"to ensure that artificial general intelligence benefits all of humanity"—was explicitly intended to counter Google's potential dominance of the technology, which seemed almost foreordained at the time. Musk pledged up to a billion dollars to prevent that outcome. It didn't take long for the two men to disagree over the chain of command. Each thought he alone deserved to run the show. About two years and thirty-eight million dollars later, Musk took his remaining nine hundred and sixty-odd million dollars and went home. In his valedictory e-mail, he wrote, "My probability assessment of OpenAI being relevant to DeepMind/Google without a dramatic change in execution and resources is 0%. Not 1%. I wish it were otherwise." OpenAI needed another source of largesse. With investors in mind, it opened a for-profit subsidiary and secured billions of dollars from Microsoft and others. This past October, the company completed a lengthy process of restructuring and recapitalization. Today, the subsidiary is worth something close to a trillion dollars.

Musk's lawsuit alleged that Altman, along with other OpenAI executives and in collusion with Microsoft, "stole a charity." He believes that they solicited his generosity on false pretenses, exploiting the cover of a humanitarian cause to build one of the world's most valuable companies, and, in the process, enriching themselves beyond measure. The remedies he sought include the unwinding of OpenAI's transformation into a for-profit company, the disgorgement of a hundred and fifty billion dollars in damages to be paid to the original nonprofit, and the final exile of Altman from the organization. It would effectively destroy the venture as such. The suit was an act of vengeance, and its primary function seemed to be to make everyone involved look heinous.

At the very least, it promised to be entertaining. During jury selection, one prospective juror assessed Musk to be "a greedy, racist, homophobic piece of garbage," while a more restrained prospect deemed him only "a world-class jerk." Musk's lawyers argued that such sentiments were blatantly prejudicial. Yvonne Gonzalez Rogers, the district-court judge and one of the few trial participants who managed to acquit herself honorably, told them to suck it up. "The reality is that people don't like him," she said. "Many people don't like him—but that doesn't mean that Americans nevertheless can't have integrity for the judicial process." In the first week of testimony, Musk took the stand and couldn't help but get tetchy. He seemed to have the impression that he alone understood the finer points of trial law. The *Verge* reporter Elizabeth Lopatto, who seemed to be live-blogging her own disgorgement (not of money but of bile), wrote, "I have never been more sympathetic to Sam Altman in my life."

Musk's attorneys hoped that the second week—which rehashed an imbroglio from around Thanksgiving 2023, when the OpenAI board fired Altman only for Altman to return and fire almost the entire board—would mitigate Altman's contrasting appeal. These allegations, for anyone unfortunate enough to have paid close attention, had long been rehearsed to death, and the iterative incantation of

now-canonical lines from e-mails and text had the quality of operatic leitmotif: Altman had been removed, we heard again and again, for having been "not consistently candid in his communications." The few novel revelations made him look less like a mastermind than kind of a loser. The existence of the lawsuit was almost redeemed by the release of a text thread between Altman, who was conferring with the Microsoft C.E.O., Satya Nadella, and Mira Murati, who briefly replaced Altman as chief executive. Altman plays the role of the clueless boyfriend who can't accept that his partner is leaving him:

ALTMAN: can you indicate directionally good or bad? satya and others anxious

MURATI: Directionally very bad

"Directionally" is Silicon Valley jargon for "generally." But Altman still doesn't get it:

ALTMAN: can you wrap up soon? lots of pressure from msft for an update

MURATI: Sam this is very bad

ALTMAN: can i come in?

MURATI: They don't want you to

Most of the regular courtroom observers eventually gave themselves over to listlessness. *Wired* ran a story about the butt pillows used by OpenAI's phalanx of lawyers and executives, including the president and co-defendant, Greg Brockman,

to insulate themselves from the discomfort of the court's benches. On a Tuesday morning of the third and final week of trial arguments, in the vacant hours of predawn Oakland, I arrived at the Ronald V. Dellums Federal Building to find a half-dozen journalists and interested civilians in the disorderly semblance of a self-organized queue. (One woman, who had severe bangs and a medieval-looking corona of braids, reminded Lopatto, the quick-witted *Verge* reporter, of a "stern German nanny"; she declined to provide her name or purpose and refused to recognize the authority of the line.) The general topic of desultory conversation was not the dispiriting trial of the present but the livelier intrigue of courtroom tech-dramas past—of Elizabeth Holmes, which inspired particular nostalgia, or Sam Bankman-Fried. Neither of those performances featured anybody to pull for, exactly, but at least they held out the pleasurable promise of comeuppance.

Had Musk v. Altman merely been a petulant matter of injured vanity, it might have played as a diverting farce. It was instead a travesty. The underlying issues—of how A.I. ought to be governed, and by whom, and how—are of great consequence. But in this trial, to root against Tweedledum was effectively to root for Tweedledee. It was a no-win situation.

The butt pillow might have begun as a symbol of the trial's frivolity, but it was clear soon enough that it was also a powerful metaphor for the collective failures that got us here. It was difficult, sitting in the unyielding pews, not to feel personally implicated. These were the leaders our society had somehow been assigned. Mike Isaac, a veteran tech reporter for the *Times*, wasn't ashamed to admit that Brockman had inspired him to secure his own butt pillow. Isaac, a magnanimous man who looks like the actor Wilford Brimley styled as a member of the hardcore band Minor Threat, offered to share the cushion, but it struck me as somehow more appropriate to sit in the docks as a penitent. The courtroom filled up quickly in anticipation of Sam Altman, who was set to appear on the stand that day under oath. The OpenAI C.E.O. has long been known for his

boyishness, but the past few years have coarsened his features and frosted his spiky hair with gray tips. He looked like a lesser vocalist for 'N Sync on a reunion tour. His presence in the courtroom had the mournful air of someone who no longer qualified as precocious.

The basic question of the case, which is also the basic question of Altman's career, is whether the transmogrification of OpenAI from a safety-minded nonprofit into a ravenous corporate behemoth was cynical in intention or merely in outcome. Recently, my *New Yorker* colleague Andrew Marantz appeared on a podcast to discuss the alternative ways to model his behavior: the "always-a-master-plan, 3-D-chess view" and the "improvisatory-checkers-all-along view." There was a clever bit of trollery in Altman's decision to hire as lead counsel the lawyer William Savitt, who had earlier forced Musk to follow through on his impulse to buy Twitter. Over hours of direct questioning, Savitt elicited from his brow-furrowed client a defense narrative that combined the most flattering elements of each version of the story. The part of the scheme that involved the creation of what he praised as "one of the largest charities in the world"—the nonprofit parent, by virtue of its equity stake in the for-profit subsidiary, has assets valued at more than two hundred billion dollars—was the result of what Altman repeatedly called "hard work" or "incredible work." But the part of the scheme that involved the creation of one of the largest and most powerful for-profit companies in the world was extemporaneous—the by-product of having been "open to creative structures." Altman said, "So this sounds a little silly to say now, but at the time, we almost didn't start this effort because we thought Google was so far ahead that it might be hopeless to compete."

The decision to stand up a profit-making entity was a matter of facts on the ground: the future of humanity required that OpenAI prevail in an existential battle against Google; this battle could not be fought without access to enormous pools of capital; it was impossible to court investors without the promise of returns. On these three points, everyone involved was in agreement: a dinky

donor-funded charity would be taking an abacus to a data-center fight. It was acknowledged only in passing that the introduction of a fiduciary motive might create perverse incentives, and even then the worry was primarily about optics. As one of Musk's consiglieri wrote in an e-mail, "I'm a super fan of capitalism and making tons of money doing great things, but not sure if this correlates with the 'noble cause for humanity, not doing it to make money' narrative." What divided Musk and his lieutenants, on one side, from Altman, Brockman, and the OpenAI chief scientist Ilya Sutskever, on the other, was the unresolved issue of which special man got to wear the pants. In September, 2017, Musk e-mailed Sutskever and Brockman to describe a scenario in which he "would unequivocally have initial control of the company." He insisted that he had no interest in retaining unilateral power over the destiny of the species. At some unspecified time, he continued, the authority vested in him, by him, would devolve upon an expanded board: "The rough target would be to get to a 12 person board (probably more like 16 if this board really ends up deciding the fate of the world)."

This organizational structure might have struck a reasonable person as a trifle undemocratic, all things considered, but what was readily clear from the trial was that Musk and Altman agreed that A.I. governance was much too serious to be left in the hands of non-player characters such as the nine assembled jurors. Altman, at times, spoke to them like children: Microsoft built them a "big computer," but they needed "more capital to keep building larger computers." (The ongoing effect was like the scene in "Airplane!" where Julie Hagerty's stewardess character, upon hearing that a passenger needs to go to the hospital, asks, "A hospital? What is it?" and Leslie Nielsen's character treats her like a ditz: "It's a big building with patients, but that's not important right now.") In his defense, it seemed as though the main lesson he'd gleaned from his dealings with Musk is that many grownups are best treated as toddlers. Altman testified that Shivon Zilis, a Musk confidant, onetime OpenAI board member, and the mother of some of Musk's many children, had "counselled me over the years when dealing with

Elon to remind him of things that happened in the past, because he was often upset." The chief prerequisite for Musk's employment seemed to be a talent for tantrum avoidance.

But Musk deserved such condescension, and the jurors did not. With the exception of Microsoft's C.T.O., Kevin Scott, a Silicon Valley engineer of the classic "Whole Earth Catalog" variety, not a single witness seemed to regard the jurors as the sorts of people with brains. David Schizer, the former dean of Columbia Law School, provided expert testimony at a rate of fifteen hundred dollars an hour—for a total he ballparked as somewhere north of three hundred grand—to describe the relationship between the OpenAI nonprofit and its subsidiary as that of a museum to its gift shop. The implication (in a trial of freely mixed metaphors) was that the profit-seeking tail of the shop had come to wag the patrimony-preserving dog of the museum. In response, the defense produced Daniel Hemel, a law professor at N.Y.U., who was paid seventeen hundred and fifty dollars an hour to argue that the gift-shop analogy was all wrong. It would be more accurate, he said, to compare the OpenAI corporation to the Newman's Own brand, which directed its profits to support a philanthropic network of summer camps. The dog of outdoor adventures for seriously ill children was not, in other words, being wagged by the tail of the popular salad-dressing company.

The testimony consistently deployed a cavalier attitude about money. Bret Taylor, the chair of the OpenAI board, responded to an inquiry about his director's compensation by saying, "I'm embarrassed to say I don't know, but on the order of two hundred thousand dollars." OpenAI's chief futurist, a researcher named Joshua Achiam, wasn't quite sure if he'd sold some of his equity last fall for more than twenty million dollars or less. These witnesses seemed to be hoping that the treatment of their financial windfalls as an afterthought would make them appear committed to elevated principles, but the net effect was to remind the jurors and the public that these sums represented little more than couch-cushion change to them. The winner of the expert-witness sweepstakes was the U.C. Berkeley computer-science professor Stuart Russell, the co-author of the canonical A.I.

textbook and an outspoken proponent of A.I. safety. Musk's attorneys brought him on to argue that OpenAI had played fast and loose with the technology, but Judge Gonzalez Rogers ruled that talk of existential risk was off-limits. (Musk's lead attorney, Steven Molo, argued, "We all could die! We all could die as a result of artificial intelligence!" The judge observed that Musk might've endorsed this argument more sincerely if he hadn't funded xAI as a competitor.) Russell struggled to find something relevant to say. To prep for his unused time, he was paid for a full forty-hour workweek at a rate of five thousand dollars an hour.

The acknowledgment of A.I.'s risk—existential threat or otherwise—was present only outside the courthouse, courtesy of a small cohort of the kinds of genteel retirees one might see a half-dozen miles north, in the Berkeley precincts of the worker-owned Cheese Board or Chez Panisse. One protester, with his face obscured by a ghoulish, oversized cartoon Musk mask, wore a Cybertruck suit made of corrugated cardboard with "Swastitruck" scrawled in runic letters. He brandished a quart-size ziplock bag labelled "Ketamine." Three women danced around him, singing a capella about Musk and Altman's shared depravity to the tune of "The Battle Hymn of the Republic."

Altman didn't exactly come off well, though he did get in some dryly amusing lines about having to endure Musk's desire to show everyone memes on his phone. It didn't really matter insofar as the evidence against Musk was challenging to refute. As document after document showed, Musk was well aware of the various plans under way to outfit their "summer-camp charity" with a substantial arm dedicated to the wanton sale of salad dressing. Musk had no problem going forward with a for-profit structure; his preference was to fold it into Tesla. (Altman testified that he'd looked askance at this plan: "Tesla is a car company.") Even if Altman and others had in fact stolen his charitable valor, it was plain that his complaint did not fall within the three-year statute of limitations. But neither Musk nor his lawyers seemed to have embarked upon this case with the goal of

juridical victory. It instead provided them with an opportunity to cast Altman, in advance of OpenAI's expected I.P.O., as an untrustworthy guy. Molo, Musk's attorney, is a stooped, pallid, funereal man with enormous hands, and he carried out his cross-examination with the solemnity of an undertaker. The first few minutes were supposed to be, I'm sure, a tour de force:

MOLO: Are you completely trustworthy?

ALTMAN: I believe so.

MOLO: But you don't know whether you're completely trustworthy?

ALTMAN: I'll just amend my answer to yes.

MOLO: Should the jury believe your testimony?

ALTMAN: I think that's up to them, but I believe so.

MOLO: You believe so, or they should?

ALTMAN: Sir, I'm not gonna tell the jury what to think.

MOLO: Do you always tell the truth?

ALTMAN: I believe I'm a truthful person.

MOLO: That wasn't my question, sir. Do you always tell the truth?

ALTMAN: I'm sure there is some time in my life when I have not.

Molo went on to recite previous testimony from various former employees or board members, who offered that Altman had dissembled about OpenAI's safety protocols and created a "toxic culture of lying." Altman suggested that he hadn't heard their comments. Did Altman know that Dario Amodei, who worked alongside Altman at OpenAI before he left to found its rival Anthropic, had accused him of "misrepresenting the terms" of the initial Microsoft investment? Altman, who unlike Musk has a recognizable sense of humor, remarked, "Dario has accused me of many things." At one point, perhaps ten minutes into this droning line of questioning, Molo asked Altman if he'd actually read the lawsuit. "Many versions, sir," Altman said. The dubiousness of Altman's character is of paramount importance to the world in general, but it seems the case—at least directionally—that it has, by this time, been priced into his reputation. To relitigate the matter on behalf of Musk only served to underline the conviction that all of these jerks deserved one another.

Besides, the cumulative effect of Molo's sustained effort to demand Altman's self-incrimination seemed to remind the courtroom that most of us, irrespective of our own faults, tend to think of ourselves as credible people doing our best. This, of course, is the real problem. Musk attempted to color Altman as a uniquely unsuitable supervisor of this technology, but this invariably invited further scrutiny into his own abject unfitness for the role. The only proper response to any of this is to point out that something as flimsy and corruptible as individual character was always going to be an insufficient basis for A.I. governance.

To claim that OpenAI's mission of cultivating beneficial A.I. was compromised by Sam Altman is to let the entire industry off the hook. Yes, Altman seems to have a

rather casual relationship with the truth. But it is far more interesting, complicated, and useful to take his self-defense at face value—to interpret the many sins of OpenAI, and its competitors, as the result of a good-faith exercise in futility. What if we imagined that he did in fact set out to do good? And discovered—or, perhaps better, failed to discover—along the way that this was structurally impossible? There are signs that some part of him anticipated his enterprise's destiny from the beginning. In the earliest correspondence about what would become OpenAI, Altman seemed vaguely aware that the collective-action problem of A.I. governance had nothing to do with individual heroism. In 2015, he e-mailed Musk to say that the development of advanced A.I. was virtually an inevitability. He wrote that "if it's going to happen anyway, it seems like it would be good for someone other than Google to do it first." What he had in mind, he continued, was not exactly a rival outfit. It was, instead, something like a "Manhattan Project for A.I." Perhaps the most salient thing about the Manhattan Project, however, was that it was not entrusted to the private sector. To launch such a transformational and dangerous program as a commercial initiative was reckless and hubristic; a "Manhattan Project run by a wealthy cabal" was as much a self-referential paradox as "all Cretans are liars."

On Monday, the jury took only two hours to reach its verdict. Musk's complaint, the panel found, had indeed exceeded the statute of limitations. If Musk had really thought that his beloved charity had been stolen, it would have occurred to him to raise the issue long before OpenAI had become what it is today. If he hadn't cared before the launch of ChatGPT, he had no right to pretend to care now.

In a better world, Altman might emerge from this humiliating rite with some genuine humility. This is almost certainly too much to expect. Still, his courtiers seemed to be encouraging him to lean into sheepishness. On the last morning of testimony, Mike Isaac, of the *Times*, tweeted an image of his sad excuse for a courthouse lunch: a meat-based protein bar, Bumble Bee-brand Snack on the

Run! Tuna Salad & Crackers, and vessels of sugary caffeine. After the morning break, Isaac sat down next to me on his butt pillow and tweeted a follow-up: "a good samaritan who took pity on me (and whose identity i will protect) came and gave me a bagel and cream cheese." I'd been standing nearby when Altman made the offering of the paper-plated bagel. Looking like a chastened little boy, he said, softly, "My comms team told me to give you this." ♦

New Yorker Favorites

- What the midterms mean for Donald Trump.
- The doxing of Jane Doe and the fallout from the "Cornell Seven" rape case.
- "She Made Me Do It": a famed equestrian instructor shot his most annoying student.
- Parents disagreed on whether to vaccinate their children. Then the kids got sick.
- Play a new game of Laugh Lines every day of the week!

Want better search results? Make *The New Yorker* a preferred source in Google.

Sign up for our daily newsletter to receive the best stories from *The New Yorker*.



Gideon Lewis-Kraus is a staff writer at The New Yorker covering technology, academia, and books, among other topics. He is the author of the memoir "A Sense of Direction."

More: Artificial Intelligence (A.I.) Elon Musk Lawsuits Silicon Valley

The News & Politics Newsletter

Read the latest from Washington and beyond, covering current events, the economy, and more, from our columnists and correspondents.

gideonlk@gmail.com

SIGN UP

By signing up, you agree to our user agreement (including class action waiver and arbitration provisions), and acknowledge our privacy policy.

READ MORE

THE LEDE

What China's Vibe-Coding Capital Can Tell Us About the A.I. Boom

Tech workers have flocked to Liangzhu, a small town outside Hangzhou, looking to escape China's corporate culture and build their future with A.I.

By Afra Wang

THE LEDE

"An Unprecedented Case of Child Abuse"

In an L.A. criminal court, new details emerge about the mysterious couple who collected more than twenty babies through surrogacy.

By Ava Kofman

ANNALS OF LAW

Should I Stay or Should I Go?

The quandary faced by government lawyers who think that the Trump Administration is subverting justice.

By Eyal Press

THE POLITICAL SCENE

Who Won the Redistricting War?

Democrats and Republicans are in a race to the bottom—and the worst consequences will extend far beyond next month's midterms.

By Jon Allsop

THE CURRENT CINEMA

The Vertiginous Revelations of "You Can See Everything"

Nathan Fielder and Lance Oppenheim's interrogation of Elizabeth Holmes, the disgraced Theranos founder, is a compelling spectacle of manipulation, suspicion, and apparent sincerity.

By Richard Brody

THE LEDE

In "Musk," Alex Gibney Punctures Elon's Self-Mythology

"To make this film, I set out to solve a mystery," Gibney says in the documentary. "How could Musk assume so much power, and why had we given it to him?"

By Sheelah Kolhatkar

THE WEEKEND ESSAY

My Year of O.C.D.

A simple e-mail or kitchen appliance could send me into a pit of anxiety. The only way out, I learned, was in.

By Andrew Lipstein

THE NEW YORKER INTERVIEW

Jeremy Strong Is Not Quite Over It

One of the great actors of the era on finding sympathy for Mark Zuckerberg, Kendall Roy, and Roy Cohn—and why the *New Yorker* Profile of him still stings.

By David Remnick

ANNALS OF TECHNOLOGY

Computer Chips Are the Most Magnificent Things Humanity Has Ever Made

They power our phones, cars, dishwashers—and the A.I. systems that prop up the economy. They're also at the heart of fierce geopolitical rivalry.

By Gideon Lewis-Kraus

THE LEDE

The Humanism Driving the Anti-A.I. Movement

At the Pro-Human Assembly in Washington, D.C., attendees talked less about jobs and policy, and more about their philosophy of life.

By Becca Rothfeld

THE LEDE

There Is No A.I. "Race"

We're constantly told that Chinese advances would have catastrophic consequences. But the first to innovate is rarely the one who wins.

By Chang Che

INFINITE SCROLL

Caleb Hammer's Humiliation Economy

The host of "Financial Audit," a popular YouTube series, claims that his goal is to make his viewers more financially literate, with his down-and-out interviewees serving as extreme test cases for what not to do.

By Brady Brickner-Wood

Privacy Information