Musk v. Altman

The Corporate Governance Litigation We Should All Follow

Teodora Groza

08.05.2026 doi: 10.59704/e97014ffb1b26802

It doesn't happen often that corporate governance litigation raises existential stakes. A prominent exception is the ongoing *Musk v. Altman* trial, where Elon Musk and Sam Altman dispute over the future of OpenAI, the world's most famous artificial intelligence ("AI") lab. The case concerns OpenAI's transformation from a charity founded to develop AI for the benefit of humanity into a for-profit entity driven primarily by commercial interests.

The dispute hasn't received much attention in the EU. Nonetheless, despite being held in a US district court over a US company and governed by US state and federal law, the case raises two fundamental issues with important implications for the ongoing European debate over AI regulation. The first is about the limits of regulation. OpenAI's evolution shows that what determines whether AI development and deployment serve the public interest is not just what ends up on the market, but decisions made long before that, such as risk tolerance, research direction, and the willingness to slow down when uncertainty is high. These choices happen inside AI organizations, and (product) regulation cannot reach them. If Europe is serious about AI safety, rules about AI products *à la* AI Act are not enough.

The second question is about what safeguards can be put in place to steer the governance of AI companies towards the public interest. OpenAI's story suggests that voluntary commitments, however well-designed on paper, are fragile when confronted with commercial pressure. If governance safeguards are to mean anything, they cannot depend solely on the goodwill of those subject to them.

The Evolution of OpenAI

OpenAI was founded in 2015 by Sam Altman and Elon Musk who allegedly joined forces due to their shared concern over the trajectory of AI research and development. According to Musk's court filings, their goal was to start a "Manhattan Project" for AI which would be structured "so that the technology belongs to the world via some sort of nonprofit." Failing to obtain support from the US government, Altman and Musk decided to start OpenAI as a charity supported by philanthropic donations. The mission of the initial OpenAI was "to provide funding for research, development and distribution" of AI, with the commitment that "the resulting technology will benefit the public" and not "the private gain of any person."

In the coming years, it became clear that the costs of AI research could not be financed by donations alone. In an effort to maintain the nonprofit commitment and attract funding, the organization adopted in 2019 an innovative structure: it added a capped-profit subsidiary to attract profit-oriented investors while still maintaining all governance rights within the initial charity. The rationale behind the capped-profit structure was to show that the organization's main purpose was still research for the public benefit, and that investors could only derive limited profits (although the limit was generously set at 100x their investment).

On paper, this structure seemed to be a masterpiece of legal engineering, seemingly able to do what no pre-existing US corporate form does: reconciling profits with a strong commitment to a public mission. In reality, it turned out that there is only so much corporate governance mechanisms can do to tame profit motives. The first stress test came in 2023, when

the board of directors of the charity decided to fire Sam Altman from his CEO role over integrity and safety concerns. In the timespan of a couple of days, due to pressure from Microsoft, it was the board that had to resign, and Altman was reinstated in his executive position.

Ever since, OpenAI's structure has moved closer and closer to an orthodox for-profit organization. The former capped-profit subsidiary has become the company's main operating entity and has been converted into a Public Benefit Corporation (PBC), a conventional corporate form capable of raising large-scale capital without any constraints on investors' returns. The initial nonprofit has become a foundation that formally still holds the right to appoint all the board members of the operational arm. Nonetheless, given the former charity's failure to rein in profit motives, there is little reason to believe the foundation will fare better.

The consequences of OpenAI's organizational transformations are palpable. According to the 2025 Stanford foundation model transparency index, OpenAI had fallen from second to second-to-last among major AI companies in research transparency. In January 2026, it announced plans to introduce advertising into ChatGPT, and its CEO floated the idea of an erotic mode for verified adults, a proposal which was eventually shelved after internal pushback. One month later, the company disbanded its alignment team, despite being repeatedly sued for releasing its models to the market too quickly and without adequate safety testing. Several lawsuits claim that ChatGPT had encouraged users, including teenagers, to commit suicide.

Taken together, these developments paint a consistent picture: once the governance commitments proved fragile, commercial logic took the front seat, with clear implications for the general public. This calls into question OpenAI's willingness to make the right decision the day it is confronted with a more radical choice, such as the release of a model that can automate (even more) white collar jobs, upend financial markets, or, in the extreme scenario, threaten our survival as a species.

The Trial and the Stakes

Musk left OpenAI's board in 2018, officially citing a conflict of interest with his role at Tesla, although subsequent reporting suggests that his departure followed a disagreement over the organisation's direction. He is now suing OpenAI on the basis that he contributed tens of millions of dollars, as well as advice and recruiting support, on the understanding that OpenAI would honour its founding mission of developing safe AI for the benefit of humanity. The remedies he seeks are far-reaching: the removal of Altman as CEO and the unwinding of the restructuring that converted OpenAI's operational arm into a for-profit entity.

Some have framed the case as an eleventh hour opportunity to take seriously the question of what kind of organizations should be entrusted with developing AI. If that sounds like too big a question for a district court, it is because it is. The presiding judge has said as much, stating explicitly that "this is not a trial on the safety risks of AI." Legally, the dispute is about an alleged breach of charitable trust: Musk argues that the donations he and others made were given on the understanding that OpenAI would remain a public interest organization, and that converting to a for-profit structure diverted those charitable assets to private gain.

The case has little chances of success. First, it is a jury trial, and Elon Musk is hardly a good trustee for this cause. His concerns over AI safety are hard to swallow, primarily because his own AI company, xAI, routinely makes headlines for its flagship chatbot Grok producing extremist content, generating disinformation, and failing to filter illegal material such as child sexual abuse (also known as "CSAM"). Second, when OpenAI converted the capped-profit subsidiary to a PBC, it had to obtain approval from the Attorney Generals of Delaware and California, the states where OpenAI is incorporated and has its main operations. Both blessed the operation, subject to minimal requirements. Given their approval, it is unlikely that the Oakland district court will go in a different direction.

Why This Matters for Europe

Even in the likely scenario that Musk loses, the case delivers an uncomfortable lesson that Europe cannot afford to ignore: neither regulation nor private ordering is sufficient to ensure that AI development serves the public interest.

The European conversation on how to limit the risks of AI development and deployment and to ensure that the technology is in line with fundamental rights and societal interests has so far focused exclusively on imposing regulatory constraints. Europe has received much criticism for the AI Act, a regulation targeted at limiting the risks of AI products placed on the internal market, and is in the process of trying to reach a deal on how these rules could be watered down so as to preserve the continent's competitiveness in the AI race.

Musk v. Altman should remind us that regulatory intervention addresses only part of the problem. OpenAI's story shows that a commercially-oriented AI lab has little incentive to delay or refrain from putting products on the market, and that no product regulation, however well-designed, can change that calculus. The decisions that matter most for AI safety are not made at the point of product release; they are taken upstream, in iterative choices about what to research, what to deploy, and when to slow down. By the time rules are drafted, negotiated, and enforced, the technology has moved on.

But OpenAI's story also shows that voluntary governance commitments designed through private ordering do not work either. When commercial pressures build, those commitments give way, as OpenAI's trajectory makes clear. What is needed is a third pathway: enforceable governance safeguards, embedded within AI organizations and backed by meaningful legal obligations, that oversee research and deployment decisions as they are being made, not after the fact. Concrete mechanisms worth exploring span a spectrum of interventionism: from mandatory internal safety boards featuring independent experts, to regulator-appointed observers embedded within AI organizations, to more assertive instruments such as golden shares granting public authorities a direct stake in governance decisions. All deserve consideration, and if the EU is serious about AI safety, the time to start is now.

ABOUT THE AUTHOR

Teodora Groza
Teodora Groza is a postdoctoral researcher at the University of Tübingen and an external lecturer at École polytechnique.

SUGGESTED CITATION
Groza, Teodora: Musk v. Altman: The Corporate Governance Litigation We Should All Follow, Verfassungsblog, 08.05.2026, doi: 10.59704/e97014ffb1b26802.

Licensed under CC BY-SA 4.0.

Langzeitarchiviert bei IntRechtDok unter: doi.org/10.17176/20260508-173531-0.