

Not For Private Gain

An Open Letter

April 2025

The Honorable Robert Bonta
Attorney General of California
California Department of Justice
Attn: Public Inquiry Unit
P.O. Box 944255
Sacramento, CA 94244-2550

The Honorable Kathleen Jennings
Attorney General of Delaware
Delaware Department of Justice
Carvel State Building
820 N. French St.
Wilmington, DE 19801

RE: OpenAI's Proposed Corporate Restructuring

Dear Attorneys General Bonta and Jennings:

We are experts in law, corporate governance, and artificial intelligence; representatives of nonprofit organizations; and former OpenAI employees.

We write in opposition to OpenAI's proposed restructuring that would transfer control of the development and deployment of artificial general intelligence (AGI) from a nonprofit charity to a for-profit enterprise. The heart of this matter is whether the proposed restructuring advances or threatens OpenAI's charitable purpose. OpenAI is trying to build AGI, but building AGI is not its mission. As stated in its Articles of Incorporation, OpenAI's charitable purpose is "to ensure that artificial general intelligence benefits all of humanity" rather than advancing "the private gain of any person."¹

OpenAI has a bespoke legal structure based on nonprofit control. This structure is designed to harness market forces without allowing those forces to overwhelm its charitable purpose.² Nonprofit control over how AGI is developed and governed is so important to OpenAI's mission that removing control would violate

¹ OpenAI, Amended Articles of Incorporation (Apr. 23, 2020). *See also Defendants' Counterclaims, Answer, and Defenses, Musk v. Altman*, No. 4:24-cv-04722-YGR (N.D. Cal. Apr. 9, 2025) ("Since its founding as an AI research lab in December 2015, OpenAI has had one mission: to ensure that artificial intelligence with the ability to outperform humans—artificial general intelligence, or 'AGI'—benefits all humanity.").

² Our position is not that it is impossible in principle for a commercial enterprise to build AGI responsibly. But OpenAI committed to holding itself to a higher standard in its Articles of Incorporation and public statements. It has repeatedly benefited from its charitable structure, and it should not be allowed to discard this structure under the very sort of pressure it was implemented to address. *See Motion for Leave to File Amici Curiae Brief in Support of Plaintiffs' Oppositions to Defendants' Motions to Dismiss*, *Musk v. Altman*, No. 4:24-cv-04722-YGR (N.D. Cal. filed Apr. 11, 2025).

the special fiduciary duty owed to the nonprofit’s beneficiaries and “pose[] a palpable and identifiable threat” to the nonprofit’s charitable purpose.³ It “would be contrary to the Certificate [of Incorporation] and hence *ultra vires*.”⁴

The restructuring would also undermine your ability to protect OpenAI’s beneficiaries: the public (including the people of California and Delaware). As the primary regulators of OpenAI, you currently have the power to protect OpenAI’s charitable purpose on behalf of its beneficiaries, safeguarding the public interest at a potentially pivotal moment in the development of this technology. Under OpenAI’s proposed restructuring, that would no longer be the case.

To preserve your power to protect the public, we urge you to:

1. **Demand answers to fundamental questions.** OpenAI has not publicly explained how its proposed restructuring will advance the nonprofit’s charitable purpose of safely developing AGI for the benefit of humanity. Nor has it provided adequate explanations for why the governance safeguards that Mr. Altman testified to Congress were important to OpenAI’s mission as recently as 2023 became obstacles to its mission in 2024.
2. **Protect the charitable trust and purpose by ensuring the nonprofit retains control.** We request that you stop the restructuring and protect the governance safeguards—including nonprofit control—that OpenAI leadership have insisted are important to “ensure[] it remains focused on [its] long-term mission.”⁵

³ Oberly v. Kirby, 592 A.2d 445, 462 (Del. 1991) (“[B]ecause the Foundation was created for a limited charitable purpose rather than a generalized business purpose, those who control it have a special duty to advance its charitable goals and protect its assets.”).

⁴ *Id.*

⁵ Oversight of A.I.: Rules for Artificial Intelligence: Hearing Before the Subcomm. on Priv., Tech., & the Law of the S. Comm. on the Judiciary, 118th Cong. (2023), available at <https://www.judiciary.senate.gov/imo/media/doc/2023-05-16%20-%20Bio%20&%20Testimony%20-%20Altman.pdf> (statement of Sam Altman, Chief Executive Officer, OpenAI).

Table of Contents

I. Understanding OpenAI’s charitable purpose: Why nonprofit control protects the mission	4
A. What OpenAI means by “AGI”	5
B. OpenAI’s founding: An attempt to solve the problem of profit-driven AGI development	7
C. 2019 restructuring: The importance of its governance safeguards	9
II. The proposed restructuring would subvert OpenAI’s charitable purpose	13
A. The restructuring would remove nonprofit control and eliminate critical governance safeguards	14
B. No sale price can compensate for loss of control	18
III. OpenAI’s public explanations for the restructuring are inadequate	19
A. Competitive advantage is not a sufficient justification	19
B. OpenAI has not explained why removing nonprofit control is necessary	22
C. The nonprofit might receive nothing for its loss of control	23
IV. A proposed plan of action	24
Conclusion	27

I. Understanding OpenAI's charitable purpose: Why nonprofit control protects the mission

That's why we're a nonprofit: we don't ever want to be making decisions to benefit shareholders. The only people we want to be accountable to is humanity as a whole.

Sam Altman, 2017⁶

OpenAI's purpose, as stated in its Articles of Incorporation, is "to ensure that artificial general intelligence benefits all of humanity" rather than advancing "the private gain of any person."⁷ Directly building AGI is *one part of how* OpenAI has decided to pursue its mission, but under its Articles of Incorporation, any desire to build AGI must be subordinate to the actual charitable purpose. As OpenAI's President Greg Brockman put it:

The true mission isn't for OpenAI to build AGI. The true mission is for AGI to go well for humanity . . . our goal isn't to be the ones to build it, our goal is to make sure it goes well for the world.⁸

OpenAI should be held to the plain text of its Articles of Incorporation. Even if OpenAI's Articles of Incorporation were ambiguous, ample evidence supports the plain reading of the text.⁹ OpenAI's founders believed the public would be harmed if AGI were developed by a commercial entity with proprietary profit motives.¹⁰ They therefore founded OpenAI as a nonprofit, with a legally enforceable duty to favor the interests of the public over investors, to carefully restrict and control such proprietary interests.

⁶ Mosaic Ventures, *Mosaic Ventures in conversation with Y Combinator President, Sam Altman*, YOUTUBE (Mar. 28, 2017), <https://youtu.be/nLMZothlRNM?feature=shared&t=1458> (at 24:18).

⁷ OpenAI, *Amended Articles of Incorporation* (filed Apr. 23, 2020).

⁸ Lex Fridman Podcast, *Greg Brockman: OpenAI and AGI* (Apr. 3, 2019) <https://youtu.be/blrEM2FbOLU?feature=shared&t=1898> (at 31:38).

⁹ Delaware and California courts apply contract principles to interpret articles of incorporation, including a nonprofit's charitable purpose. See *Gunderson v. Trade Desk, Inc.*, 326 A.3d 1264, 1273 (Del. Ch. 2024), *as corrected* (Nov. 8, 2024); *Wong v. Restoration Robotics, Inc.*, 78 Cal. App. 5th 48, 61 (Cal. Ct. App. 2022). In California, courts consider extrinsic evidence if it is relevant to prove a meaning to which the contract language is reasonably susceptible. *Pac. Gas & E. Co. v. G.W. Thomas Drayage & Rigging Co.*, 69 Cal.2d 33, 37 (1968). In Delaware, unless the terms of the articles of incorporation are ambiguous, Delaware courts will give effect to their plain meaning. *Salama v. Simon*, 328 A.3d 356, 366 (Del. Ch. 2024). If the terms are ambiguous, the courts will consider extrinsic evidence. *Id.*

¹⁰ See Sections I.B and I.C, *infra*.

OpenAI’s Charter, which states the “principles [OpenAI] use[s] to execute on [its] mission,” describes the subordination of competitive and financial goals to its charitable purpose:

OpenAI’s mission is to ensure that artificial general intelligence (AGI)—by which we mean highly autonomous systems that outperform humans at most economically valuable work—benefits all of humanity. We will attempt to directly build safe and beneficial AGI, but will also consider our mission fulfilled if our work aids others to achieve this outcome.¹¹

In 2019, when OpenAI believed it needed more capital than it could raise in donations, it adopted a new structure: its nonprofit would control a for-profit entity, with additional governance safeguards to ensure the primacy of its charitable mission. It is these governance safeguards—including the nonprofit’s critical control rights over the for-profit—that are now at stake.

A. What OpenAI means by “AGI”

The problem with AGI specifically is that if we’re successful, and we tried, maybe we could capture the light cone of all future value in the universe. And that is for sure not okay for one group of investors to have.

Sam Altman, 2020¹²

OpenAI defines artificial general intelligence (AGI) as “a highly autonomous system that outperforms humans at most economically valuable work.”¹³ OpenAI’s leadership believes that AGI has the potential to “elevate humanity”:

If AGI is successfully created, this technology could help us elevate humanity by increasing abundance, turbocharging the global economy,

¹¹ OpenAI, [OpenAI Charter](#).

¹² Connie Loizos, *Sam Altman in Conversation with StrictlyVC*, YOUTUBE (May 18, 2019), <https://www.youtube.com/watch?v=TzcJlKg2Rc0&t=2734s> (at 45:34).

¹³ *Our Structure*, OPENAI (updated June 28, 2023), <https://web.archive.org/web/20230729203855/https://openai.com/our-structure>.

and aiding in the discovery of new scientific knowledge that changes the limits of possibility.¹⁴

They also expect it will generate “unprecedented economic benefits”¹⁵ and “close to infinite wealth.”¹⁶ And the organization that builds it could create “orders of magnitude more value than any existing company.”¹⁷

On the other hand, OpenAI warns that “AGI would also come with serious risk of misuse, drastic accidents, and societal disruption.”¹⁸ Its website states that a “misaligned superintelligent AGI could cause grievous harm to the world.”¹⁹ Mr. Altman and senior OpenAI employees joined hundreds of others, including recent Nobel Prize winner and AI pioneer Geoffrey Hinton, in signing a statement that “[m]itigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war.”²⁰

OpenAI now believes AGI—both the positive potential and the inherent risks—might be on the horizon. Mr. Altman has said: “We are now confident we know how to build AGI as we have traditionally understood it”²¹ and that “AGI will probably get developed during this president’s term.”²² A few months ago, Mr. Altman was quoted as saying “that on cybersecurity and bio stuff, we’ll see serious, or potentially serious, short-term issues that need mitigation.”²³

¹⁴ *Planning for AGI and beyond*, OPENAI (Feb. 24, 2023), <https://openai.com/index/planning-for-agi-and-beyond/>.

¹⁵ *Artificial Intelligence: With Great Power Comes Great Responsibility: Joint Hearing Before the Subcomm. on Research & Tech. and the Subcomm. on Energy of the H. Comm. on Sci., Space & Tech.*, 115th Cong. 2 (2018), <https://www.govinfo.gov/content/pkg/CHRG-115hhr30877/pdf/CHRG-115hhr30877.pdf>.

¹⁶ Guy Raz, *HIBT Lab! OpenAI: Sam Altman*, How I Built This with Guy Raz (Wondery Sept. 29, 2022), <https://podcasts.apple.com/us/podcast/hibt-lab-openai-sam-altman/id1150510297?i=1000580232536>.

¹⁷ Lex Fridman Podcast, *Greg Brockman: OpenAI and AGI* (April 3, 2019) <https://www.youtube.com/watch?v=blrEM2FbOLU&pp=0gcJCdgAo7VqN5tD> (at 27:55).

¹⁸ *Planning for AGI and beyond*, OPENAI (Feb. 24, 2023), <https://openai.com/index/planning-for-agi-and-beyond/>.

¹⁹ *Id.*

²⁰ *Statement on AI risk*, Center for AI Safety, <https://www.safe.ai/work/statement-on-ai-risk#open-letter>.

²¹ Sam Altman, *Reflections*, SAM ALTMAN BLOG (Jan. 5, 2025), <https://blog.samaltman.com/reflections>.

²² Josh Tyrangiel, *Sam Altman on ChatGPT’s First Two Years, Elon Musk and AI Under Trump*, BLOOMBERG BUSINESSWEEK (Jan. 6, 2025), <https://www.bloomberg.com/features/2025-sam-altman-interview/>.

²³ *Id.*

B. OpenAI's founding: An attempt to solve the problem of profit-driven AGI development

Been thinking a lot about whether it's possible to stop humanity from developing AI. I think the answer is almost definitely not. If it's going to happen anyway, it seems like it would be good for someone other than Google to do it first. Any thoughts on whether it would be good for YC to start a Manhattan Project for AI? My sense is we could get many of the top ~50 to work on it, and we could structure it so that the tech belongs to the world via some sort of nonprofit but the people working on it get startup-like compensation if it works. Obviously we'd comply with/aggressively support all regulation.

Sam Altman, 2015²⁴

In 2015, OpenAI founders Sam Altman, Elon Musk, and Greg Brockman were deeply concerned about the trajectory of artificial intelligence. DeepMind, which Google acquired as a subsidiary, was the leading AI company and the only major company explicitly trying to build AGI. The founders expressed the view that Google, a commercial entity whose ultimate responsibility is to shareholders, must not be the institution that builds what might be the most powerful technology ever created.²⁵

OpenAI was founded against this backdrop. Its founding announcement states:

Our goal is to advance digital intelligence in the way that is most likely to benefit humanity as a whole, unconstrained by a need to generate financial return. Since our research is free from financial obligations, we can better focus on a positive human impact.²⁶

To achieve its mission while avoiding becoming either another profit-driven Google competitor or an academic safety lab without direct influence over cutting-edge AI, OpenAI adopted a two-pronged strategy:

²⁴ Amended Complaint, Ex. 1, Musk v. Altman et al., No. 4:24-cv-04722-YGR (N.D. Cal. Nov. 14, 2024), ECF No. 32-2 ([containing an email from Sam Altman dated May 25, 2015](#)).

²⁵ See, e.g., *Id.* ("If it's going to happen anyway, it seems like it would be good for someone other than Google to do it first.").

²⁶ *Introducing OpenAI*, OPENAI (Dec. 11, 2015), <https://openai.com/index/introducing-openai/>.

1. Try to build AGI. Like Google, OpenAI would strive to be on the cutting edge of AI research and development. Directly building AGI advances the mission because no one is better placed to ensure AGI is built safely and for the benefit of humanity than the organization building it. As the OpenAI Charter states: “To be effective at addressing AGI’s impact on society, OpenAI must be on the cutting edge of AI capabilities—policy and safety advocacy alone would be insufficient.”²⁷
2. But the mission comes first. The goal of building AGI also leads to foreseeable conflicts with the mission, when the proprietary interests of investors diverge from the interests of humanity. To ensure their drive to compete in the race to build AGI would never undermine their mission, they incorporated as a nonprofit and made their goal of building AGI legally and structurally subordinate to the mission.

OpenAI’s Charter acknowledges the potential for conflicts between pecuniary goals and charitable mission, and it explicitly reinforces the primacy of the mission:

Our primary fiduciary duty is to humanity. We anticipate needing to marshal substantial resources to fulfill our mission, but will always diligently act to minimize conflicts of interest among our employees and stakeholders that could compromise broad benefit.²⁸

Indeed, the Charter’s stop-and-assist clause articulates a specific scenario in which racing to build AGI would undermine the mission, and it commits OpenAI to stop trying to build AGI under these circumstances:

We are concerned about late-stage AGI development becoming a competitive race without time for adequate safety precautions. Therefore, if a value-aligned, safety-conscious project comes close to building AGI before we do, we commit to stop competing with and start assisting this project.²⁹

²⁷ OpenAI, [OpenAI Charter](#).

²⁸ *Id.*

²⁹ *Id.*

As a nonprofit, OpenAI raised \$130.5 million from donors who believed in its mission.³⁰

C. 2019 restructuring: The importance of its governance safeguards

We had tried and failed enough to raise the money as a nonprofit; we didn't see a path forward there. So we needed some of the benefits of capitalism, but not too much. I remember at the time someone said: 'As a nonprofit, not enough will happen; as a for-profit, too much will happen'. So we need this sort of strange intermediate.

Sam Altman, 2023³¹

In 2019, OpenAI decided it needed to raise equity capital in addition to the donations and debt capital it could raise as a nonprofit nonstock corporation. To do this while preserving the primacy of its mission, it established a controlled, for-profit subsidiary of the nonprofit corporation. Through the controlled subsidiary (which we refer to here as “OpenAI-profit”), it would raise capital from investors seeking to make a return. The new for-profit entity would, however, be a subsidiary of the original nonprofit (which we refer to here as “OpenAI-nonprofit”), which would retain its fiduciary duty to advance the charitable purpose above all else.

The structure was carefully selected to manage the conflict between the interests of investors and the charitable purpose. OpenAI's announcement of the change discusses its motivations plainly: First, the formation of the subsidiary was motivated by a need for capital. As the announcement states:

We've experienced firsthand that the most dramatic AI systems use the most computational power in addition to algorithmic innovations, and decided to scale much faster than we'd planned when starting OpenAI. We'll need to invest billions of dollars in upcoming years into large-scale cloud compute, attracting and retaining talented people, and building AI supercomputers.³²

³⁰ *Our Structure*, OPENAI (updated June 28, 2023), <https://web.archive.org/web/20230729203855/https://openai.com/our-structure>.

³¹ Lex Fridman Podcast, *Sam Altman: OpenAI CEO on GPT-4, ChatGPT, and the Future of AI* (Mar. 25, 2023), https://www.youtube.com/watch?v=L_Guz73e6fw&t=4481s (at 1:14:41).

³² *OpenAI LP*, OPENAI (Mar. 11, 2019), <https://openai.com/index/openai-lp/>.

Second, the mission must still come first despite the need to raise this capital:

We’ve designed OpenAI LP to put our overall mission—ensuring the creation and adoption of safe and beneficial AGI—ahead of generating returns for investors *Regardless of how the world evolves*, we are committed—legally and personally—to our mission.³³

As the statement above shows, OpenAI foresaw that it might be tempted to prioritize investor returns over its mission and sought to bind itself to the mast—committing to the mission “regardless of how the world evolves.” As OpenAI explained in the preamble to this announcement, its new structure “allows us to rapidly increase our investments in compute and talent *while including checks and balances to actualize our mission.*”³⁴

The announcement enumerates the specific “checks and balances” OpenAI believed were important to its mission:

1. Nonprofit control. The new company (OpenAI-profit) “is controlled by OpenAI Nonprofit’s board.”³⁵

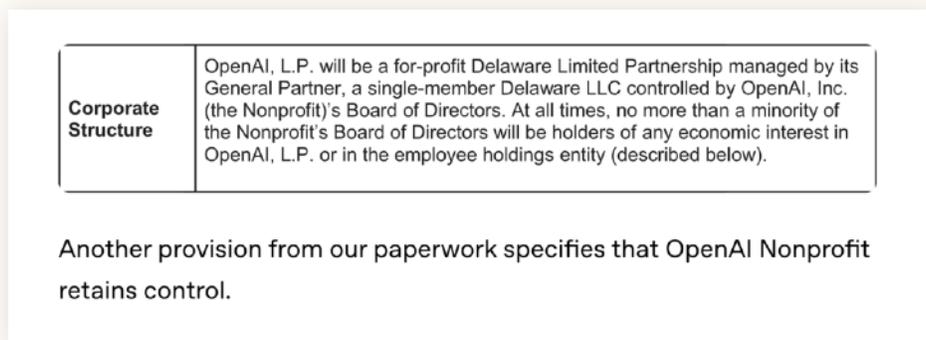


Image source: *OpenAI LP*, OPENAI (Mar. 11, 2019), <https://openai.com/index/openai-lp/>.

³³ *Id.* (emphasis added).

³⁴ *Id.* (emphasis added). When OpenAI’s president, Greg Brockman, was asked in 2019 why OpenAI did not incorporate as a public benefit corporation, he responded: “We needed to custom-write rules like: Fiduciary duty to the charter - Capped returns - Full control to OpenAI Nonprofit.” Greg Brockman (@gdb), HACKER NEWS (Mar. 11, 2019), <https://news.ycombinator.com/item?id=19359928>.

³⁵ *OpenAI LP*, OPENAI (Mar. 11, 2019), <https://openai.com/index/openai-lp/>.

This was not just a formality. “All investors and employees sign agreements that OpenAI LP’s obligation to the Charter always comes first, even at the expense of some or all of their financial stake.”³⁶

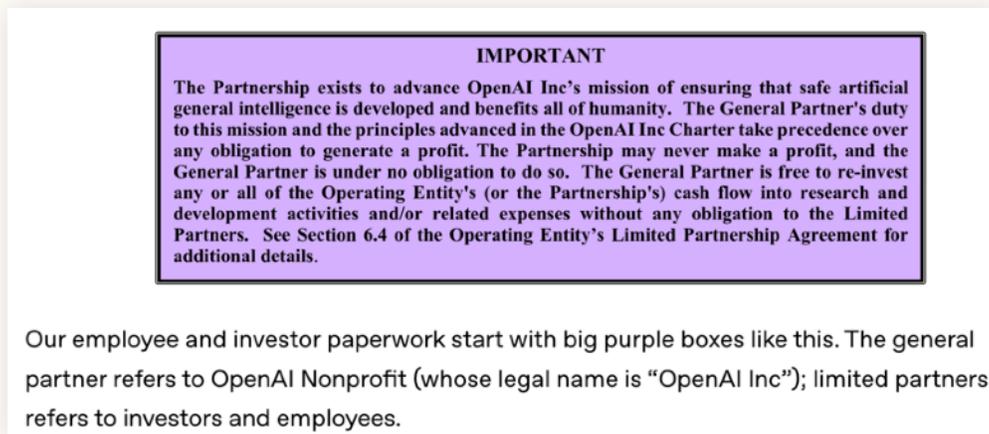


Image source: *OpenAI LP*, OPENAI (Mar. 11, 2019), <https://openai.com/index/openai-lp/>.

2. Capped investor profits. Potential profits for investors in OpenAI-profit are capped, with any returns above the cap going to OpenAI-nonprofit. The capped-profit structure was designed to strike a balance between the ability to raise capital and ensuring that the vast majority of the expected value of AGI accrues to OpenAI-nonprofit and its public interest mission.³⁷

The fundamental idea of OpenAI LP is that investors and employees can get a capped return if we succeed at our mission, which allows us to raise investment capital and attract employees with startup-like equity. But any returns beyond that amount—*and if we are successful, we expect to generate orders of magnitude more value than we’d owe to people who invest in or work at OpenAI LP*—are owned by the original OpenAI Nonprofit entity. . . . Our goal is to ensure that most of the value (monetary or otherwise) we create if successful benefits everyone, so we think this is an important first step.³⁸

3. Independent Board. Conflicts of interest on the nonprofit board should be avoided:

³⁶ *Id.*

³⁷ Tomio Geron, *Nonprofit AI Lab Alters Structure to Build Massive Computing Power*, WALL ST. J. (Mar. 11, 2019), <https://www.wsj.com/articles/nonprofit-ai-lab-alters-structure-to-build-massive-computing-power-11552352064>.

³⁸ *OpenAI LP*, OPENAI, (Mar. 11, 2019), <https://openai.com/index/openai-lp> (emphasis added).

Only a minority of board members are allowed to hold financial stakes in the partnership at one time. Furthermore, only board members without such stakes can vote on decisions where the interests of limited partners and OpenAI Nonprofit’s mission may conflict—including any decisions about making payouts to investors and employees.³⁹

As Mr. Altman testified before Congress in May 2023, OpenAI’s “unusual structure” “ensures it remains focused on [its] long-term mission.”⁴⁰ Mr. Altman then enumerated the following governance safeguards:

- **Nonprofit control.** “First, the principal entity in our structure is our Nonprofit, which is a 501(c)(3) public charity. Second, our for-profit operations are subject to profit caps and under a subsidiary that is fully controlled by the Nonprofit.”⁴¹
- **Legal duty to put the charitable purpose first.** “Third, because the board serves the Nonprofit, each director must perform their fiduciary duties in furtherance of its mission—safe AGI that is broadly beneficial. While the for-profit subsidiary is permitted to make and distribute profit, it is subject to this mission. The Nonprofit’s principal beneficiary is humanity, not OpenAI investors.”⁴²
- **Independent board.** “Fourth, the board remains majority independent. Independent directors do not hold equity in OpenAI.”⁴³
- **Profit caps.** “Fifth, profit for investors and employees is capped by binding legal commitments. The Nonprofit retains all residual value for the benefit of humanity.”⁴⁴

³⁹ *Id.*

⁴⁰ *Oversight of A.I.: Rules for Artificial Intelligence: Hearing Before the Subcomm. on Priv., Tech., & the Law of the S. Comm. on the Judiciary*, 118th Cong. (2023), available at <https://www.judiciary.senate.gov/imo/media/doc/2023-05-16%20-%20Bio%20&%20Testimony%20-%20Altman.pdf> (statement of Sam Altman, Chief Executive Officer, OpenAI).

⁴¹ *Id.*

⁴² *Id.*

⁴³ *Id.*

⁴⁴ *Id.*

- **Ownership of AGI.** “AGI technologies are explicitly reserved for the Nonprofit to govern.”⁴⁵

These safeguards are now in jeopardy under the proposed restructuring.

II. The proposed restructuring would subvert OpenAI’s charitable purpose

We did not expect massive scale to be as important as it turned out to be. By 2019, we realized that, and that the amount of money we were going to need to succeed at the mission was beyond what we could raise as a non-profit. . . . We wanted to preserve as much as we could of the specialness of the nonprofit approach, the benefit sharing, the governance, what I consider maybe to be most important of all, which is the safety features and incentives.

Sam Altman, 2022⁴⁶

OpenAI recently announced that it intends to restructure again, purportedly to satisfy investor demands to simplify the capital structure.⁴⁷ Under this new proposed restructuring, OpenAI-profit would transform into a Delaware public benefit corporation (which we refer to here as “OpenAI-PBC”). OpenAI-PBC would then “run and control OpenAI’s operations and business, while the non-profit will hire a leadership team and staff to pursue charitable initiatives in sectors such as health care, education, and science.”⁴⁸ OpenAI-nonprofit’s “significant interest in the existing for-profit would take the form of shares in the PBC.”⁴⁹

By removing nonprofit control, the proposed restructuring would eliminate most, if not all, of the governance safeguards that OpenAI has insisted are important to its charitable purpose. And no amount of money would advance OpenAI-

⁴⁵ *Id.*

⁴⁶ Guy Raz, *HIBT Lab! OpenAI: Sam Altman, How I Built This with Guy Raz* (Wondery Sept. 29, 2022), <https://podcasts.apple.com/us/podcast/hibt-lab-openai-sam-altman/id1150510297?i=1000580232536> (at 35:16).

⁴⁷ Defendants’ Counterclaims, Answer, and Defenses, Musk v. Altman, No. 4:24-cv-04722-YGR (N.D. Cal. filed Apr. 9, 2025).

⁴⁸ *Why OpenAI’s Structure Must Evolve to Advance Our Mission*, OPENAI (Dec. 27, 2024), <https://openai.com/index/why-our-structure-must-evolve-to-advance-our-mission/>.

⁴⁹ *Id.*

nonprofit’s charitable purpose more than control over one of the world’s leading AI companies.

A. The restructuring would remove nonprofit control and eliminate critical governance safeguards

On the ethical front, that’s really core to my organization. That’s the reason we exist . . . when it comes to the benefits of who owns this technology? Who gets it? You know, where did the dollars go? We think it belongs to everyone.

Greg Brockman, 2018⁵⁰

The restructuring would profoundly change the duties owed by the organization controlling the development and deployment of OpenAI’s technology. Its incentives would shift from (i) an enforceable duty owed to the public to ensure AGI is safe and benefits all of humanity, even if it means no profits are ever made, to (ii) no fiduciary duty at all to the public, with an enforceable requirement to consider profit-making to enhance shareholder value.⁵¹ All of the governance safeguards that Mr. Altman claimed “ensure [OpenAI] remains focused on [its] long-term mission”⁵² would be in jeopardy. *See Table 1.*

<i>Governance Safeguards</i>	<i>Today</i>	<i>Proposed Restructuring</i>
1. Profit motives are subordinate to charitable purpose	Yes	No

⁵⁰ *Artificial Intelligence: With Great Power Comes Great Responsibility: Joint Hearing Before the Subcomm. on Research & Tech. and the Subcomm. on Energy of the H. Comm. on Sci., Space & Tech.*, 115th Cong. 17 (2018), <https://www.govinfo.gov/content/pkg/CHRG-115hhr30877/pdf/CHRG-115hhr30877.pdf>.

⁵¹ The capped-profit structure of the controlled subsidiary was designed to avoid precisely this situation. As Mr. Altman himself stated in a 2021 interview with Ezra Klein, “One of the incentives that we were very nervous about was the incentive for unlimited profit, where more is always better” because “with these very powerful general purpose A.I. systems, in particular, you do not want an incentive to maximize profit indefinitely.” Ezra Klein Show, *Sam Altman on the A.I. Revolution, Trillionaires and the Future of Political Power*, N.Y. TIMES (June 11, 2021), <https://www.nytimes.com/2021/06/11/podcasts/transcript-ezra-klein-interviews-sam-altman.html>.

⁵² *Oversight of A.I.: Rules for Artificial Intelligence: Hearing Before the Subcomm. on Priv., Tech., & the Law of the S. Comm. on the Judiciary*, 118th Cong. (2023), available at <https://www.judiciary.senate.gov/imo/media/doc/2023-05-16%20-%20Bio%20&%20Testimony%20-%20Altman.pdf> (statement of Sam Altman, Chief Executive Officer, OpenAI).

2. Leadership has a fiduciary duty to advance the charitable purpose, enforceable by the attorneys general	Yes	No
3. Investor profits are capped, with above-cap profits owned by the nonprofit	Yes	Rumored no
4. Majority independent board commitment	Yes	Unknown
5. AGI, when developed, belongs to the nonprofit for the benefit of humanity	Yes	No by default
6. Stop-and-assist commitment from Charter	Yes	Unknown

Table 1. Governance safeguards at stake in proposed restructuring

1. Profit motives no longer subordinate to the mission. OpenAI-nonprofit is currently required to govern the organization in the interests of humanity, without regard to whether OpenAI-profit makes any money at all. “Any action that poses a palpable and identifiable threat to [its charitable] goals, or that jeopardizes its assets would be contrary to the Certificate [of Incorporation] and hence *ultra vires*.”⁵³ Even transactions approved by disinterested directors are “not legally binding” if they “pose[] a clear threat to the charitable purpose or the assets of the [nonprofit] corporation.”⁵⁴

If control is transferred to OpenAI-PBC, the board governing the development and deployment of OpenAI’s technology would be required to consider the pecuniary interests of shareholders in making its decisions.⁵⁵ And while OpenAI-PBC could “balance” that consideration with a stated public benefit, its directors would owe no fiduciary duties enforceable by or on behalf of the public benefit’s beneficiaries.⁵⁶

2. No legally enforceable duty to advance the charitable purpose. Under the current structure, OpenAI-nonprofit’s directors owe a special fiduciary duty to the

⁵³ *Oberly*, 592 A.2d at 462.

⁵⁴ *Id.* at 468 n.17.

⁵⁵ 8 Del. C. § 365(a).

⁵⁶ 8 Del. C. § 365(b).

public as the beneficiaries of the organization’s charitable purpose.⁵⁷ That duty itself is important, but equally important is its enforceability. Right now, each of you as the attorneys general of OpenAI’s state of incorporation and the location where its charitable assets are held have the power to protect OpenAI’s charitable purpose, including by enforcing the special fiduciary duty owed under Delaware law.⁵⁸

By shifting control to OpenAI-PBC, the proposed restructuring would divorce the development and deployment of AGI from any duty to the public enforceable by elected officials. Instead, it would empower shareholders of OpenAI-PBC to sue its board derivatively on behalf of the company for failing to adequately consider the profit motives of the shareholders.⁵⁹ The only stakeholders with the power to sue to alter the balance between pecuniary interests and public benefit would then be shareholders with strong financial incentives to tip that balance away from public benefit and toward maximizing returns.

3. Profits above the current cap would reportedly go to investors. Today, investor profits are subject to a cap, with profits above the cap going to OpenAI-nonprofit for the benefit of humanity. OpenAI has not commented on whether it would attempt to approximate the existing capped-profit structure under the proposed restructuring, but credible reports claim that recent investments were made contingent on the removal of the profit cap.⁶⁰ This may constitute a massive reallocation of wealth from humanity at large to OpenAI-PBC shareholders.

4. No independent board commitment. Under the current structure, OpenAI has committed to a majority-independent board. OpenAI has not publicly commented on whether OpenAI-PBC would have a majority-independent board. There are conflicting reports about whether Mr. Altman would receive an equity

⁵⁷ *Oberly*, 592 A.2d at 462 (“[B]ecause [OpenAI] was created for a limited charitable purpose rather than a generalized business purpose, those who control it have a special duty to advance its charitable goals and protect its assets.”).

⁵⁸ *Id.* at 468 (“Delaware law unambiguously places the burden of protecting the interests of beneficiaries upon the Attorney General.”). In California, the Attorney General holds the “primary responsibility for supervising charitable trusts in California, for ensuring compliance with trusts and articles of incorporation, and for protection of assets held by charitable trusts” Cal. Gov’t Code § 12598(a).

⁵⁹ 8 Del. C. §§ 365(a), 367.

⁶⁰ Aditya Soni, Arsheeya Bajwa & Krystal Hu, *OpenAI Outlines New For-profit Structure in Bid to Stay Ahead in Costly AI Race*, REUTERS (Dec. 27, 2024), <https://www.reuters.com/technology/artificial-intelligence/openai-lays-out-plan-shift-new-for-profit-structure-2024-12-27/>.

stake in OpenAI-PBC, with Mr. Altman denying rumors he would receive a 7% stake.⁶¹

5. AGI presumptively owned by investors. Separate from the question of who is entitled to profits above the cap is the question of who would actually *own and control* AGI technology. Under the current structure, when OpenAI builds AGI, it would be controlled by OpenAI-nonprofit. OpenAI’s contract with key investor Microsoft provides that Microsoft would have no right to AGI technologies.⁶² The nonprofit would have the right and legal responsibility to use AGI in the manner that most benefits humanity,⁶³ which could include providing subsidized access or even transferring control to a governmental body. Although OpenAI has not publicly commented on who would own AGI under the proposed restructuring, it would presumably belong to OpenAI-PBC and its investors. Credible reporting also claims that OpenAI and Microsoft have discussed removing the contractual restriction on Microsoft’s access to AGI technologies.⁶⁴

6. Stop-and-assist commitment. OpenAI has not publicly commented on whether OpenAI-PBC would honor the commitments in the OpenAI Charter, including the commitment to stop competing and assist a mission-aligned company close to building AGI. Mr. Altman has, however, claimed that one of the benefits of its current structure—with nonprofit control—is that it allows them to stop competing with and start assisting another organization.⁶⁵ And even if OpenAI-PBC announced that it intended to honor those commitments, it could easily abandon them without a mission-aligned reason for doing so.

⁶¹ Sharon Goldman, Kali Hays & Verne Kopytoff, *Why Investors Want Startup Founders to Own Equity—Including OpenAI’s Sam Altman*, FORTUNE (Sept. 30, 2024), <https://fortune.com/2024/09/30/sam-altman-openai-equity-stake-billionaire/>.

⁶² *Our Structure*, OPENAI, <https://openai.com/our-structure/>.

⁶³ OpenAI, *Amended Articles of Incorporation* (filed Apr. 23, 2020); *Oberly*, 592 A.2d at 462.

⁶⁴ Cristina Criddle & George Hammond, *OpenAI Seeks to Unlock Investment by Ditching “AGI” Clause with Microsoft*, FIN. TIMES (Dec. 6, 2024), <https://www.ft.com/content/2c14b89c-f363-4c2a-9dfc-13023b6bce65>.

⁶⁵ Lex Fridman Podcast, *Sam Altman: OpenAI CEO on GPT-4, ChatGPT, and the Future of AI* (Mar. 25, 2023), https://youtu.be/L_Guz73e6fw?si=q7UinB9gDFIVDDBK&t=4440 (at 1:14:05).

B. No sale price can compensate for loss of control

At OpenAI, when we wrote our charter, we talked about the scenarios where we would or wouldn't make money. And . . . the things we wouldn't be willing to do no matter how much money they made. And we made this public so the public would hold us accountable to that. And I think that's really important.

Sam Altman, 2018⁶⁶

Whether OpenAI-nonprofit receives fair market value for its controlling interest in OpenAI-profit is not the core question. The core question is whether selling control advances or undermines OpenAI's purpose. The law here is clear: "Although the public in general may benefit from any number of charitable purposes, charitable contributions *must be used only for the purposes for which they were received in trust.*"⁶⁷

No sale price could put OpenAI-nonprofit in a comparable position to realize its charitable purpose.⁶⁸ While in principle financial assets could be leveraged to obtain influence over the development and governance of AGI, obtaining the degree of influence OpenAI-nonprofit already possesses is no longer possible in practice.

Further, OpenAI is not even claiming the proceeds of the sale would be used to ensure safe AGI benefits all of humanity. Rather, OpenAI promises that the restructuring would create "one of the best resourced non-profits in history," that would "pursue charitable initiatives in sectors such as health care, education, and science."⁶⁹ However beneficial such a foundation may be, it would not advance

⁶⁶ Kara Swisher, *Should Mark Zuckerberg Fire Himself?*, RECODE DECODE (Vox Media, Dec. 7, 2018), <https://www.vox.com/2018/12/10/18134926/sam-altman-kara-swisher-recode-decode-live-mannys-podcast-transcript-facebook-zuckerberg-ethics>.

⁶⁷ *Holt v. Coll. of Osteopathic Physicians & Surgeons*, 394 P.2d 932, 935 (emphasis added); see also *Queen of Angels Hosp. v. Younger*, 66 Cal. App.3d 359, 369 (Cal. App. 2d Dist. 1977) ("The issue is not the desirability of the new use [of charitable assets], but consistency with the charitable purpose).

⁶⁸ See, e.g., *In re Milton Hershey School Trust*, 807 A.2d 324 (Pa. Cmmw. Ct. 2002) as instructive. In that case, the Hershey charitable trust had a controlling interest in Hershey Foods, which it sought to sell. The Pennsylvania Attorney General sued to enjoin the sale. Noting the "symbiotic relationship" between the trust and the company, the court found that even a control premium would be inadequate. "How many trusts enjoy holding a controlling interest in one of this nation's largest, historically profitable, and best-known corporations?" *Id.* at 334.

⁶⁹ *Why OpenAI's Structure Must Evolve to Advance Our Mission*, OPENAI (Dec. 27, 2024), <https://openai.com/index/why-our-structure-must-evolve-to-advance-our-mission/>.

OpenAI's specific charitable purpose. Imagine a nonprofit with the mission of ensuring nuclear technology is developed safely and for the benefit of humanity selling its control over the Manhattan Project in 1943 to a for-profit entity so the nonprofit could pursue other charitable initiatives.

III. OpenAI's public explanations for the restructuring are inadequate

OpenAI's public justifications for its restructuring prioritize competitive advantage over its charitable purpose and fail to address how abandoning nonprofit control is compatible with the mission. While OpenAI's current structure might have limitations, and competitive positioning might indirectly support its mission, any proposed solution must be tailored to address actual deficiencies without compromising core principles. OpenAI has not publicly demonstrated how the purported benefits to its mission of restructuring outweigh the substantial risks of dismantling the very safeguards designed to keep OpenAI faithful to its mission.

A. Competitive advantage is not a sufficient justification

The primary reason OpenAI cites for the restructuring is competitive advantage. As stated in a recent court filing:

OpenAI's current structure poses challenges in attracting new investment and retaining and attracting highly skilled personnel. Every one of OpenAI's significant competitors has a familiar corporate structure that allows for offers of conventional equity—an attraction not just for investors contemplating multi-billion-dollar commitments but for current and prospective employees who want a stake in the enterprise they're helping to build.⁷⁰

Competitive advantage might be a relevant factor, but it is not a sufficient reason to restructure. OpenAI's charitable purpose is not to make money or capture market share. Its "competitors" are not nonprofits with a duty to the public. OpenAI's structure, by design, comes with competitive costs. Attracting talent and remaining on the cutting-edge of AI development at best indirectly advances

⁷⁰ Defendants' Counterclaims, Answer, and Defenses, Musk v. Altman, No. 4:24-cv-04722-YGR (N.D. Cal. filed Apr. 9, 2025), ECF 147.

OpenAI's mission of ensuring AGI benefits all of humanity. Obtaining a competitive advantage by abandoning the very governance safeguards designed to ensure OpenAI remains true to its mission is unlikely to, on balance, advance the mission.

OpenAI might respond that a competitive advantage inherently advances its mission, but that argument is an implicit comparison of OpenAI and its competitors: that humanity would be better off if OpenAI builds AGI before competing companies. Based on OpenAI's recent track record, this argument is unlikely to be convincing:⁷¹

- OpenAI's testing processes have reportedly become "less thorough with insufficient time and resources dedicated to identifying and mitigating risks."⁷²
- It has rushed through safety testing to meet a product release schedule.⁷³
- It reneged on its promise to dedicate 20% of its computing resources to the team tasked with ensuring AGI's safety.⁷⁴
- OpenAI and its leadership have publicly claimed to support AI regulation⁷⁵ while OpenAI privately lobbied against it.⁷⁶
- Mr. Altman said that it might soon become important to reduce the global availability of computing resources⁷⁷ while privately attempting to arrange

⁷¹ One might take this list as evidence that the existing structure is ineffective. But the current structure provides an enforceable commitment to a charitable purpose with which to counter transgressions. If OpenAI pushes past what that commitment allows, the answer is not to remove the commitment or make it unenforceable; the answer is to enforce it.

⁷² Cristina Criddle, *OpenAI Slashes AI Model Safety Testing Time*, FIN. TIMES (Apr. 11, 2025), <https://www.ft.com/content/8253b66e-ade7-4d1f-993b-2d0779c7e7d8>.

⁷³ Pranshu Verma, Nitasha Tiku & Cat Zakrzewski, *OpenAI Promised to Make Its AI Safe. Employees Say It 'Failed' Its First Test*, WASH. POST (July 12, 2024), <https://www.washingtonpost.com/technology/2024/07/12/openai-ai-safety-regulation-gpt4/>.

⁷⁴ Jeremy Kahn, *Exclusive: OpenAI Promised 20% of Its Computing Power to Combat the Most Dangerous Kind of AI—but Never Delivered, Sources Say*, FORTUNE (May 21, 2024), <https://fortune.com/2024/05/21/openai-superalignment-20-compute-commitment-never-fulfilled-sutskever-leike-altman-brockman-murati/>.

⁷⁵ Cecilia Kang, *How Sam Altman Stormed Washington to Set the A.I. Agenda*, N.Y. TIMES (June 7, 2023), <https://www.nytimes.com/2023/06/07/technology/sam-altman-ai-regulations.html>.

⁷⁶ Billy Perrigo, *Exclusive: OpenAI Lobbied the E.U. to Water Down AI Regulation*, TIME (June 20, 2023), <https://time.com/6288245/openai-eu-lobbying-ai-act/>.

⁷⁷ *Planning for AGI and beyond*, OPENAI (Feb. 24, 2023), <https://openai.com/index/planning-for-agi-and-beyond/>.

- trillions of dollars in computing infrastructure buildout with U.S. adversaries.⁷⁸
- OpenAI coerced departing employees into extraordinarily restrictive non-disparagement agreements.⁷⁹

What might make OpenAI the best choice for humanity? The most significant differentiator between OpenAI and its competitors is that it is a nonprofit with a duty to put humanity's interests first. But it is that very differentiator that the proposed restructuring would strip away.

OpenAI might also respond that humanity is better off if OpenAI builds AGI before a company based in China. OpenAI foreshadowed this argument in its recent submission to the Office of Science & Technology Policy:

As America's world-leading AI sector approaches artificial general intelligence (AGI), with a Chinese Communist Party (CCP) determined to overtake us by 2030, the Trump Administration's new AI Action Plan can ensure that American-led AI built on democratic principles continues to prevail over CCP-built autocratic, authoritarian AI.⁸⁰

But OpenAI is not America's only organization at the AI frontier. To the extent OpenAI believes that it is imperative to its mission for an American company to build AGI before a Chinese company, the solution might not be to compete but to assist. As a nonprofit, OpenAI could offer its resources—compute, talent, and IP—to another American company or the U.S. government, which could do more to advance American competitiveness than the proposed restructuring. Indeed, the OpenAI Charter explicitly envisions a related scenario in which the mission would be best served not by competing, but by assisting:

⁷⁸ Keach Hagey & Asa Fitch, *Sam Altman Seeks Trillions of Dollars to Reshape Business of Chips and AI*, WALL ST. J. (Feb. 8, 2024), <https://www.wsj.com/tech/ai/sam-altman-seeks-trillions-of-dollars-to-reshape-business-of-chips-and-ai-89ab3db0>.

⁷⁹ Kelsey Piper, *ChatGPT Can Talk, but OpenAI Employees Sure Can't*, Vox (May 18, 2024), <https://www.vox.com/future-perfect/2024/5/17/24158478/openai-departures-sam-altman-employees-chatgpt-release>.

⁸⁰ Public Letter from Christopher Lehane, OpenAI Vice President, Glob. Affs., to Faisal D'Souza, Off. of Sci. & Tech. Pol'y (Mar. 13, 2025), <https://tinyurl.com/ycy23ctm>.

[I]f a value-aligned, safety-conscious project comes close to building AGI before we do, we commit to stop competing with and start assisting this project.⁸¹

Furthermore, profit incentives might actually cause OpenAI to hinder American competitiveness by, for example, commercializing technology counter to America's national security interests. Nvidia is currently lobbying to reduce U.S. export controls so that its technology can be more widely sold in China.⁸² A profit-driven OpenAI might act similarly.

B. OpenAI has not explained why removing nonprofit control is necessary

To justify the restructuring, OpenAI primarily cites investor demands that it “simplify its capital structure”, and refers specifically to market unfamiliarity with its profit caps:

The profit interests in OpenAI's capped-profit are less familiar [compared to its competitors]. The challenges OpenAI faces are reflected in its most recent fundraising rounds, in which investors have insisted on conditions freeing them from certain funding commitments or allowing redemption of invested funds with interest in the event OpenAI fails to simplify its capital structure.⁸³

OpenAI has not, however, explained the relationship between nonprofit control and the profit caps. Why removing nonprofit control is necessary to simplify its capital structure is not self-evident.

OpenAI's only public justification for removing nonprofit control is that nonprofit control subordinates investor interests to the charitable mission. As OpenAI explained in its restructuring announcement: “Our current structure does not allow the board to directly consider the interests of those who would finance the

⁸¹ OpenAI, OpenAI Charter.

⁸² Sharon Goldman, *Nvidia Lashes Out at Biden's Last-minute Export Controls on AI Chips and Rushes to Praise Trump*, FORTUNE (Jan. 13, 2025), <https://fortune.com/2025/01/13/nvidia-lashes-out-at-biden-administration-sweeping-last-minute-export-controls-on-ai-chips/>.

⁸³ Defendants' Counterclaims, Answer, and Defenses, Musk v. Altman, No. 4:24-cv-04722-YGR (N.D. Cal. filed Apr. 9, 2025).

mission.”⁸⁴ By contrast, OpenAI-PBC would allegedly have “the exact same mission but also having accountability to investors and employees.”⁸⁵ This argument fails because OpenAI’s board *is already* permitted to take investor and employee interests into account, *provided doing so advances the charitable purpose*. What the board cannot do is take those interests into account at the expense of OpenAI’s charitable purpose, and that appears to be what the proposed restructuring seeks to do.

C. The nonprofit might receive nothing for its loss of control

OpenAI does not adequately explain how the surviving nonprofit entity—stripped of any control over OpenAI-PBC—would be in a better position to advance its charitable purpose. It states:

Our plan would result in one of the best resourced non-profits in history. The non-profit’s significant interest in the existing for-profit would take the form of shares in the PBC at a fair valuation determined by independent financial advisors. This will multiply the resources that our donors gave manyfold.⁸⁶

But as OpenAI explains, the sale would simply “exchange [OpenAI nonprofit’s] current economic interest in the capped-profit entity for an equity stake in the PBC.”⁸⁷ It is the value of its *current economic interest* in OpenAI-profit that makes OpenAI-nonprofit “one of the best resourced non-profits in history,” not additional value the nonprofit would receive from the proposed restructuring. The transaction, as described, does not put any dollar value on the nonprofit’s control of one of the world’s leading AI companies, let alone try to justify how that unstated amount of money would better enable the nonprofit to achieve its mission. If OpenAI’s donors simply wanted a “manyfold” return that could be directed to generic charitable initiatives a decade later, they had far more direct ways of accomplishing that goal.

⁸⁴ *Why OpenAI’s Structure Must Evolve to Advance Our Mission*, OPENAI (Dec. 27, 2024), <https://openai.com/index/why-our-structure-must-evolve-to-advance-our-mission/>.

⁸⁵ *Defendants’ Counterclaims, Answer, and Defenses*, Musk v. Altman, No. 4:24-cv-04722-YGR (N.D. Cal. filed Apr. 9, 2025), ECF 147.

⁸⁶ *Why OpenAI’s Structure Must Evolve to Advance Our Mission*, OPENAI (Dec. 27, 2024), <https://openai.com/index/why-our-structure-must-evolve-to-advance-our-mission/>.

⁸⁷ *Defendants’ Counterclaims, Answer, and Defenses*, Musk v. Altman, No. 4:24-cv-04722-YGR (N.D. Cal. filed Apr. 9, 2025).

OpenAI offers ambitious plans for how the activities of the new nonprofit can benefit the public, “particularly within OpenAI’s home state of California.”⁸⁸ Putting aside whether those plans are consistent with OpenAI’s charitable purpose, we see no good reason why the cost of those activities must be OpenAI-nonprofit relinquishing control over OpenAI-profit.⁸⁹ OpenAI-nonprofit is currently free to pursue ambitious charitable initiatives—and has in fact made grants in the past⁹⁰—but it should not be permitted to sell out its mission to do so.

IV. A proposed plan of action

We respectfully request that you, Attorneys General Bonta and Jennings, take the following actions:

1. Demand answers to fundamental questions

OpenAI has not publicly provided answers to fundamental questions about the restructuring. We urge you to investigate the following questions:

Rationale for Board’s decision	
1	Is removing nonprofit control the best way to advance OpenAI’s charitable purpose, and if so, why? What analyses were done and what alternatives were considered?

⁸⁸ *New Commission to Provide Insight as OpenAI Builds the World’s Best-equipped Nonprofit*, OPENAI (Apr. 2, 2025) <https://openai.com/index/nonprofit-commission-guidance/>.

⁸⁹ OpenAI claims that the current structure “does not enable the non-profit to easily do more than control the for-profit” but provides no explanation for why that is the case. See *Why OpenAI’s Structure Must Evolve to Advance Our Mission*, OPENAI (Dec. 27, 2024), <https://openai.com/index/why-our-structure-must-evolve-to-advance-our-mission/>.

⁹⁰ In their most recent reported fiscal year, OpenAI made over \$2.5 million in grants. OpenAI, Inc. [IRS Form 990](#), Part I, line 13 (FY2023) (reporting \$2,641,712 in grants and similar amounts paid). In the 2020 fiscal year, OpenAI made over \$10 million in grants. OpenAI, Inc. [IRS Form 990](#), Part I, line 13 (FY2020) (reporting \$10,250,005 in grants and similar amounts paid).

2	In testimony before Congress in May 2023, Mr. Altman stated that OpenAI’s governance “ensures it remains focused on [its] long-term mission.” Many if not all of the safeguards he cited then would cease to exist under the proposed restructuring. What changed between May 2023 and September 2024 such that the safeguards that were important to OpenAI’s mission became obstacles instead?
3	Mr. Altman has said several times that being a for-profit creates perverse incentives for organizations trying to build AGI. Is that not still the case, and if so, why not?
4	What role has investor pressure had on the decision to restructure? Did the board formally approve a funding round that included investments conditioned on the restructuring? If so, what information was provided to inform its approval?
Involvement of interested directors	
5	Which of the current directors have participated in or are planning to participate in decisions regarding the proposed restructuring, which (if any) are not, and what potential conflicts of interest does each have?
6	Under the proposed restructuring, would any of the directors receive equity, additional equity, or other direct or indirect personal interests in OpenAI-PBC?
OpenAI’s proposed restructuring plan	
7	What would the governance structure be of OpenAI-PBC? Who would be on the board of the entity trying to build AGI?
8	What structures would OpenAI-PBC have to guarantee its actions are consistent with the mission of ensuring that AGI is developed safely and for the benefit of all humanity? Specifically, what safeguards will be in place to ensure that AI systems OpenAI-PBC releases are safe and that AGI’s benefits are distributed to all of humanity, not preferentially to investors?
9	If OpenAI-PBC takes actions that are clearly deleterious to humanity, who would have recourse?
10	Would OpenAI-PBC recommit to the principles in the Charter, including the promise to stop competing with and assist a mission-aligned organization close to building AGI?

11	Is it OpenAI's position that other frontier AI companies are not mission aligned? Is it OpenAI's position that humanity's interests are in jeopardy if any of those companies build AGI before OpenAI does? If so, what is the justification for those positions?
12	Will new investors' profits still be capped? Would the above-cap profits still go to OpenAI-nonprofit? Will past investors' caps be removed? What value will the nonprofit receive in exchange for any changes to the profit cap structure? Request all internal analyses about the expected value of profits above the cap.
13	If OpenAI builds AGI, who would own and control it? Would Microsoft have the same right to AGI IP that it has to the IP of OpenAI's current AI systems?

Table 2. Proposed questions

2. Protect the charitable trust and purpose by ensuring the nonprofit retains control

Do not allow the restructuring to proceed as planned. We urge you to protect OpenAI's charitable purpose by preserving the core governance safeguards that OpenAI and Mr. Altman have repeatedly claimed are important to its mission:

1. The leadership of OpenAI—the entity trying to build AGI—should have a fiduciary duty to the mission: to ensure AGI is developed safely and for the benefit of humanity.
2. All goals other than the mission—including profits and winning the race to AGI—should be subordinate to the mission.
3. OpenAI's duty to put mission first must be legally enforceable by, at a minimum, the attorneys general, directors, and parties with a special interest in the matter.
4. Investor profits should continue to be capped, and profits above the cap should be used exclusively and directly to benefit humanity.
5. OpenAI should retain the commitment in its Charter to stop competing and start assisting a mission-aligned project that is close to building AGI.

6. AGI itself—if and when OpenAI creates it—should belong to the nonprofit entity, or a similar entity whose sole responsibility is to ensure it is used responsibly and for the benefit of humanity. It should not be owned or controlled by a commercial entity or its investors.

Stopping the restructuring is not enough, however, as these governance safeguards were apparently not enough to prevent management from pushing to remove them. We also request that you ensure the board has the necessary independence, resources, information, and will to push back against management in furtherance of its fiduciary duties. Specifically:

- **Removal of directors.** Any director found to have undermined the integrity of the board’s decision regarding the restructuring should be removed.
- **Independence.** A majority of the board should have no direct or indirect personal interests in OpenAI-profit.
- **Expertise.** The board should have the expertise necessary to know when OpenAI-profit is taking actions that are at odds with the mission.
- **Resources.** The board should have the staff and budget necessary to oversee the operations of an organization with thousands of employees and a valuation of \$300 billion.
- **Information.** OpenAI should provide the board with routine and timely updates on any OpenAI activities that might be in conflict with its mission. Directors should receive prompt and detailed responses to any questions they ask of management.
- **Oversight.** We encourage you to oversee the implementation of these changes or appoint an independent body to oversee their implementation. Until these changes are made, any decision by the board should receive careful scrutiny.

Conclusion

OpenAI was founded to ensure AGI is developed safely and benefits all of humanity. Its current structure, which legally subordinates profit motives to this mission, is not incidental—it is fundamental to achieving this purpose. The proposed restructuring would eliminate essential safeguards, effectively handing

control of, and profits from, what could be the most powerful technology ever created to a for-profit entity with legal duties to prioritize shareholder returns.

You have both the authority and duty to protect OpenAI's charitable trust and purpose. We urge you to halt this restructuring, restore proper governance, and ensure OpenAI remains accountable to you, the public, and its charitable purpose.

Respectfully submitted,

Page Hedley, OpenAI 2017-2018

Sunny Gandhi, Encode AI

Nathan Calvin, Encode AI

Tyler Whitmer, Legal Advocates for Safe Science and Technology

Lucian Bebchuk, Harvard Law School

Anu Bradford, Columbia Law School

Samuel Brunson, Loyola University Chicago School of Law

Michael Dorff, UCLA School of Law

Lawrence Lessig, Harvard Law School

Oliver Hart, Harvard University (Nobel laureate)

Katharina Pistor, Columbia Law School

Marc Rotenberg, Georgetown Law

Joseph Stiglitz, Columbia University (Nobel laureate)

Kevin Werbach, The Wharton School, University of Pennsylvania

Luigi Zingales, University of Chicago

Geoffrey Hinton, University of Toronto (Nobel laureate)

Margaret Mitchell, Hugging Face

Stuart Russell, University of California, Berkeley

Scott Aaronson, OpenAI 2022-2024

Steven Adler, OpenAI 2020-2024

Jacob Hilton, OpenAI 2018-2023

Daniel Kokotajlo, OpenAI 2022-2024

Ryan Lowe, OpenAI 2019-2024

Gretchen Krueger, OpenAI 2019-2024

Girish Sastry, OpenAI 2019-2024

Nisan Stiennon, OpenAI 2018-2020

Anish Tondwalkar, OpenAI 2023-2024

Center for Humane Technology
The Tech Oversight Project

Arturo Béjar, former leader for Protect and Care, Facebook (2009-2015)

Jennifer Gibson, Psst.org

Sam Hiner, Young People's Alliance

Ed Howard, Children's Advocacy Institute, University of San Diego School of Law

Joan F. Neal, NETWORK Lobby for Catholic Social Justice

Christabel Randolph, Center for AI and Digital Policy

Jason Green-Lowe, Center for AI Policy

Reed Schuler, Massachusetts Institute of Technology