Dear OpenAl: Open Letter to We write to you as the legal beneficiaries of your charitable OpenAl mission. Signatories OpenAI is currently sitting on both sides of the table in a closed boardroom, making a deal on humanity's behalf without Add your allowing us to see the contract, know the terms, or sign off on signature the decision. Appendix Your organization was founded with a legally binding commitment to ensure AGI benefits the public. It enshrined that mission in its founding documents, and it used the goodwill it inspired to raise money and attract talent. Your current structure includes important safeguards designed to ensure your technology serves humanity rather than merely generating profit, including profit caps for investors, nonprofit management of commercial operations, and explicit commitments to prioritize your charitable mission. However, you have proposed a significant corporate restructuring that appears to weaken or eliminate many of these protections, and the public deserves to know the details. We call for at least a basic level of transparency about how this transition will affect your legal commitments to the public. Specifically, we request clear answers to the following questions: 1. Will OpenAl continue to have a legal duty to prioritize its charitable mission over profits? 2. Will OpenAl's nonprofit continue to have full management control over OpenAl? 3. Which of OpenAI's nonprofit directors will receive equity in OpenAl's new structure? 4. Will OpenAl maintain profit caps and abide by its commitment to devote excess profits to the benefit of humanity? 5. Does OpenAl plan to commercialize AGI once developed, instead of adhering to its promise to retain nonprofit control of AGI for the benefit of all of humanity? 6. Will OpenAI recommit to the principles in its Charter, including its pledge to stop competing and start assisting if another responsible organization is close to AGI? 7. Will OpenAI reveal what is at stake for the public in its restructuring by releasing: a. The OpenAl Global, LLC operating agreement, which sets out OpenAI's duties to its charitable mission and the powers given to its nonprofit. **b.** All estimates of the potential value of above-cap profits, including any estimates it has shared with investors. We believe your response will help restore trust and establish whether OpenAI remains committed to its founding principles, or whether it is prioritizing private interests over its public mission. The stakes could not be higher. The decisions you make about governance, profit distribution, and accountability will shape not only your organization's future but also the future of society at large. Sam Altman has said that OpenAl wants to be held accountable to humanity, and we share this letter in the spirit of offering that accountability. We look forward to your response and to working together to ensure AGI truly benefits everyone. ¹ Additional detail on the motivation for and context around each of these questions can be found in the attached appendix, particularly Section 3. To be explicit, we are not requesting trade secrets or any proprietary information — information irrelevant to the public's entitlements can and should be redacted in any documents OpenAI chooses to disclose. Signatories - 1172 Encode **Co-organizer** M The Midas Project **Co-organizer** EyesOnOpenAl **Co-organizer** A California coalition of nonprofit, philanthropic, and labor organizations **Geoffrey Hinton Nobel laureate** University Professor Emeritus, University of Toronto Giorgio Parisi **Nobel laureate** Professor Emeritus, Sapienza University in Rome Sir Oliver Hart **Nobel laureate** Professor, Harvard University Sheldon Lee Glashow **Nobel laureate** Professor Emeritus, Harvard University and **Boston University** Vitalik Buterin Co-founder, Ethereum Sir Stephen Fry Writer, Director, Actor **Audrey Tang** Institute for Ethics in AI, University of Oxford Lawrence Lessig Professor, Harvard Law School Stuart Russell Professor of Computer Science, Founder of the Center for Human-Compatible Artificial Intelligence, UC Berkeley Andrew Barto Professor Emeritus, University of Massachusetts Amherst Sean Carroll Professor of Natural Philosophy, Johns Hopkins University Max Tegmark Massachusetts Institute of Technology, Future of Life Institute Don Norman Professor Emeritus, UC San Diego Daniel Kokotajlo Ex-OpenAI, 2022-2024 Executive Director, Al Futures Project Girish Sastry Ex-OpenAI, 2019-2024 Gretchen Krueger Ex-OpenAI, 2019-2024 Fellow, Berkman Klein Center, Harvard University Helen Toner Ex-OpenAI, 2021-2023 Jacob Hilton Ex-OpenAI, 2018-2023 Executive Director, Alignment Research Center Page Hedley Ex-OpenAI, 2017-2018 Rosie Campbell Ex-OpenAI, 2021-2024 Managing Director, Eleos Al Research Scott Aaronson Ex-OpenAI, 2022-2024 Professor of Computer Science, University of Texas at Austin Steven Adler Ex-OpenAI, 2020-2024 Independent Researcher (B) BetaNYC California Community Foundation California Endowment Center for AI and Digital Policy 🕮 Center for Leadership Equity and Research Center for Oil and Gas Organizing CivAl © Cognitive Resonance Communities United for Restorative Youth Justice Defend Democracy Demand Progress Demand Progress Education Fund East LA Community Corporation 6 Ekō Existential Risk Observatory Fair Vote UK Future of Life Institute Indigenous Justice Innovate Public School Internet Sans Frontières (P) LatinoProsperity Legal Advocates for Safe Science and Technology Liberty Hill Foundation MIT Data + Feminism Lab National Committee for Responsive

Philanthropy

Load More See Full List

Add Your Signature

You will receive a confirmation

email to verify your signature.

Individual

Full Name *

Email*

but it's under threat.

It explains:

the public.

obligations.

agreement.

general public.

5. And why all of this matters.

attract talent, and gain tax benefits.

Section 1: Why OpenAI has unique

OpenAI is not like other technology companies. The

transparency obligations to the public

organization was founded as a nonprofit with an explicit

of humanity. It enshrined this commitment in its founding

This mission created special legal obligations, which state

attorneys general can enforce. As the primary regulators of

nonprofits, attorneys general have the responsibility to protect

all beneficiaries of OpenAl's charitable purpose — that is, the

In 2019, OpenAI created for-profit subsidiaries to raise capital, but these entities remained under the control of the nonprofit

governed by the nonprofit's mission, explicitly prioritizing the

1. Nonprofit oversight: The for-profit entity is managed by

2. Capped investor profits: Investor returns were initially

3. Majority independent board: Conflicts of interest on

charitable purpose over commercial interests. All investors

entity's obligation to the Charter always comes first, even

4. Mission primacy: All decisions must prioritize the

and employees sign agreements that the commercial

at the expense of some or all of their financial stake.

capped at 100x, with excess profits flowing to the

nonprofit (and thus its beneficiary, humanity).

the nonprofit board were to be avoided.

parent. This hybrid structure was carefully designed with safeguards to ensure OpenAI's business operations were

charitable purpose over investors' interests.

The specific safeguards included:

OpenAl nonprofit's board.

mission to ensure that artificial general intelligence benefits all

documents, and it used the goodwill it inspired to raise money,

Fill out this form to add your signature below.

Sign as:

Submit

Appendix to "An Open Letter to OpenAl"

OpenAI was founded in 2015 as a nonprofit with a mission to

ensure that artificial general intelligence benefits all of

humanity. Despite all that's happened since, that mission remains — as both an aspiration and a legal requirement —

We believe OpenAl's current lack of transparency puts the

is living up to its legal obligations and abiding by its

commitment to act in humanity's best interests.

nonprofit's integrity at risk. This appendix explains why greater transparency is required from OpenAI in order to ensure that it

1. Why OpenAI has unique transparency obligations to

2. The pattern of failures to live up to these transparency

3. How the restructuring presents a crisis for OpenAI.4. Why OpenAI should disclose its hidden operating

Organization

Title and/or Affiliation

6

An Open Letter to OpenAI

Because OpenAI is committed to serving humanity's interests rather than maximizing shareholders' profits, the public has legitimate rights to understand: • What is currently owed to the public under OpenAl's operating agreement. • How decisions about these obligations are being made. • Whether these obligations will change or be removed as a result of the upcoming restructuring. This isn't about curiosity. OpenAl is a nonprofit, and nonprofits are accountable for their mission and to the people they claim to serve. The general public should be informed about how OpenAI plans to fulfill its legal obligations to them to ensure that their interests are duly served. Section 2: The pattern of failures to live up to its charitable obligations Unfortunately, OpenAI has not consistently lived up to the standard of transparency that its mission demands. And its track record raises serious concerns about whether the public can trust the company's commitments. Silent Changes to Core Commitments When OpenAI restructured in 2019, the company promised that investor returns would be capped at 100x, with excess profits "returned to humanity" in the form of direct cash flow to the nonprofit, allowing it to distribute the windfall that AGI may bring back to humanity as it best sees fit. However, in 2023, it was reported that the profit caps would increase by 20% every year, starting in 2025. Within forty years, a \$100 billion limit on investor returns would become \$100 trillion larger than today's entire global economy. The public only learned about this through independent reporting. OpenAl never announced the change that transferred billions in future value from humanity to investors. Silencing Internal Voices Former OpenAI employees faced highly restrictive nondisclosure and non-disparagement agreements that threatened the loss of their vested equity (potentially worth millions) if they refused to sign or if they later criticized the company. These practices have been alleged to illegally interfere with federal reporting of safety risks, effectively preventing internal voices from warning the public about potential dangers. **Broken Safety Promises** Computing Resources Promise: In 2023, OpenAl promised to commit 20% of its computing resources to its newly formed Al safety research team. However, according to former team lead Jan Leike and reporting from Fortune, OpenAl did not allocate these computing resources to its safety team. **Security Breach Cover-up**: In 2023, a hacker gained access to OpenAl internal messages and stole details about their Al technology. The company did not inform authorities about the breach, and the story did not come out for over a year. **Rushed Safety Evaluations and Deteriorating Safety Standards**: Some members of OpenAI's safety team felt pressured to speed through a new testing protocol, designed to prevent the technology from causing catastrophic harm. The company planned release celebrations before the preparedness team could determine if the model was safe. OpenAI's testing processes have reportedly become "less thorough with insufficient time and resources dedicated to identifying and mitigating risks." **Missing Safety Documentation** OpenAI has committed to releasing safety evaluations alongside model releases, yet crucial documentation has been systematically delayed or omitted: OpenAl released its preparedness scorecard risk assessment for GPT-40, the first model it released after adopting its preparedness framework, three months after the model's public release. • For several recent models, safety scorecards weren't released at launch, and some never came out at all. • Results of the model's "preparedness evaluations," the tests OpenAl runs to assess an Al model's dangerous capabilities and other risks, were based on earlier versions of ol. They had not been run on the final version of the model. Section 3: How the restructuring presents a crisis for OpenAl OpenAl's recent restructuring plans have brought the transparency problem into sharp focus. They are trying to change OpenAI's legal entity from a capped-profit/nonprofit hybrid to a more conventional for-profit company (a public benefit corporation) — weakening its legal obligations to prioritize the public interests over profit — all behind closed doors. OpenAl admits that their new restructuring plan comes at the behest of investors, despite the fact that investor pressure is exactly what the nonprofit structure was designed to avoid. Under the new plans, all of the following are under threat: 1. Control of OpenAI by the nonprofit board: Today, OpenAI's nonprofit is the manager of its commercial entity, with full control over its operations. The nonprofit's powers are specified in the commercial entity's operating agreement, which OpenAI has not made public. OpenAI has provided few details about the level of control the nonprofit will have in its new structure. **2. Profit caps on returns to investors:** The potential returns on all OpenAI investments are currently capped, and profits above that cap belong to its nonprofit, and must be used to benefit humanity. OpenAI has implied, but not explicitly stated, that profits will no longer be capped in its new structure. More concerning, OpenAl hasn't said whether it will retroactively remove caps for existing investors — potentially transferring profits that currently belong to humanity into private hands. **3. The governance of AGI:** Investors such as Microsoft currently have no right to AGI technologies. Instead, OpenAl's nonprofit would have the right to govern AGI and must use it to benefit humanity. OpenAI has not commented on who would govern AGI when OpenAI creates it, or whether it would be used to benefit humanity or OpenAI's investors. 4. The OpenAI Charter and its "stop and assist" commitment: OpenAl's Charter states the principles it uses to execute its mission. It includes a commitment to stop competing with and assist a mission-aligned organization close to building AGI, which limits the risk of a dangerous race for the technology. Today, the operating agreement for OpenAI's commercial entity requires it to uphold these principles, even when doing so conflicts with its commercial interests. It has not yet publicly committed to uphold the Charter in its new structure. 5. The legal primacy of OpenAI's mission to benefit **humanity:** Today, the operating agreement for OpenAl's commercial entity requires it to prioritize its charitable mission over profits. It has not committed to the same requirement in its new structure. **6. Independence of the nonprofit board:** OpenAl has committed to having an independent board, meaning a majority of its directors have no financial interest in its commercial operations. OpenAI has not definitively commented on whether any of its independent directors — including Sam Altman — stand to gain from its restructuring, via receiving equity in the new entity or via affiliated companies. OpenAI has claimed that the nonprofit will continue to control its new for-profit entity. However, the proposed restructuring would, by default, significantly weaken the degree of control that the nonprofit currently has. Right now, the nonprofit has total managerial oversight over all of OpenAl's activities, but in the new restructuring, it may only have the right to appoint and fire directors for the for-profit, as well as consent rights for a small handful of decisions. The public is asked to trust OpenAl's assurances, but the company's track record of silent changes to core commitments makes such trust difficult to sustain. OpenAI has described the complete removal of limits on investor returns as a mere "simplification" of the company's capital structure. They've described the potential disempowerment of OpenAI's nonprofit as an unremarkable preservation of the status quo. And they've claimed that their new structure will allow them to prioritize purpose over profit — even though that is a right which they already have, and which they may lose on their current trajectory. There was a time when OpenAI bragged about its structural safeguards. They described these rules as legally enforceable and core to the mission. Now. they tend to downplay, ignore, or conveniently redefine them. Section 4: Why OpenAI should disclose its hidden operating agreement At the heart of this transparency crisis lies a simple issue: the public cannot verify OpenAI's commitments because the legal document that implements them are hidden from view. OpenAI's LLC operating agreement — the internal contract that governs how the for-profit subsidiary actually operates contain the real safeguards that supposedly protect humanity's interests. These documents likely specify: How much control the nonprofit actually has. • What influence investors have over key decisions. How mission-related commitments are legally enforced. • The true terms of profit caps and distribution mechanisms. Governance structures and voting rights. Now, as they consider transitioning to a new public benefit corporation, tricky decisions about these safeguards will have to be made: which will stay exactly as they are, which will be reworded and softened in small and large ways, and which will be discarded entirely. But without access to the legal documents instantiating these commitments, the public has no idea exactly what entitlements it is retaining or losing. While private companies typically keep these agreements confidential, OpenAI's subsidiaries are different. They are controlled by a nonprofit that exists to benefit humanity. The operating agreements that govern these entities are the mechanism by which OpenAI's promises to the public are instantiated. If those promises are going to be revoked or modified, the public has a right to know. To demonstrate genuine commitment to transparency and accountability, OpenAI should: Publish the LLC operating agreement, allowing independent verification of the safeguards the company claims protect humanity's interests. Provide detailed documentation of how restructuring will preserve nonprofit control and mission primacy, including specific legal mechanisms and enforcement procedures. • Establish regular transparency reporting on governance decisions, safety protocols, and mission-related activities. Section 5: Why all of this matters The stakes extend far beyond OpenAI. As artificial intelligence becomes increasingly powerful, the precedents set today will shape how AI development proceeds globally. If organizations claiming to serve humanity can operate with secrecy, public accountability becomes meaningless. OpenAI has repeatedly stated that transparency and safety are core values, yet it continues to make its most consequential decisions behind closed doors. The time has come for the company to demonstrate that these are more than talking points. The public OpenAI claims to serve deserves to see the legal foundations of the company's commitments. And if OpenAI is planning to change that foundation, the public deserves to know exactly what is being altered and why. If these changes are truly in humanity's interest, as the company claims, then why hide the details? Let the people decide for themselves. Contact Privacy Policy Terms of Service