π PLAINTIFF π
United States District Court
Northern District of California
Case No. 4:24-CV-04722-YGR
Case Title Musk et al. v. Altman et al.
Exhibit No. 359
Date Entered
Mark Busby, Clerk
By: , Deputy Clerk

From: Greg Brockman [gdb@openai.com]
Sent: 12/30/2023 5:07:13 AM
To: Jason Kwon [jason@openai.com]; Sam Altman [sama@openai.com]
Subject: Fwd: Benefit corp

[privileged]

---------- Forwarded message ---------
From: **Greg Brockman** <gdb@openai.com>
Date: Thu, Dec 28, 2023 at 10:03 AM
Subject: Re: Benefit corp
To: Larry Summers <lhs@openai.com>
CC: Sam Altman <sama@openai.com>, Bret Taylor <btaylor@openai.com>, Adam D'Angelo <ad@openai.com>

Great to hear, I am also intrigued by benefit corp. Annual report within the current structure could be interesting!

As part of that, would definitely like to align on how we think about the mission. Sam and I have been working on Charter 2.0 as mentioned at the board meeting. Here's our current draft, would love any early feedback on the direction:

Charter 2.0

We started OpenAI with the goal "to advance digital intelligence in the way that is most likely to benefit humanity as a whole"; this remains our fundamental guiding principle.

Over the years, we've increasingly learned how to further this goal in practice, such as by making breakthroughs and deploying them as products like ChatGPT. We've also realized we were wrong at times. For example, we originally set the company up as a non-profit "unconstrained by a need to generate financial return"; over the past eight years we've grown to regard capitalism not as a constraint but instead as a positive force that, managed well, can align different stakeholders to pool their resources to create far more value than they keep.

While we know we'll make changes as the technology continues to unfold, the following is our articulation of how we pursue our fundamental goal today:

OpenAI's mission is to provide AGI that benefits all of humanity.

Empower humanity

We provide AI tools that allow others to shape the future. We have deep conviction that much more good than bad will come of this, since we believe:

(1) AI can benefit every aspect of our lives — from helping with daily tasks to one day helping humanity solve our largest challenges.
(2) The collective effort of humanity will be required to figure out how to best apply AI, and

(3) AI's downsides can be mitigated by collectively working on topics such as technical alignment, trading off misuse mitigations with liberty, maximizing race-to-the-top and minimizing race-to-the-bottom dynamics on safety, predicting AI capabilities, responding to societal impacts, etc..

Mission success requires that humanity as a whole, our customers, and our shareholders all win. That is, that everyone's quality of life is radically better (imagine a world where everyone can figure out how to be happier and more fulfilled than anyone is today), our customers are accomplishing things they couldn't have dreamed of otherwise, and our employees and investors are making a return commensurate with making all of this possible.

We focus on producing the most capable models even if they require too much compute to serve at humanity scale today. We are working in parallel to scale up the world's supply of compute, which we expect to become a critical limited resource in the future.

Deploy iteratively [[better name needed here: meaning is "don't let an overhang of capability build up" or "keep changes in the world continuous to match the continuity of technological progress".]]

We choose projects in order to maximize our rate of learning about and preparation for AGI. This generally means our top priority is our research, but we also deploy useful services, build large-scale inference infrastructure, work with governments and the public on understanding and responding to the upsides and downsides of this technology, explore new methods for getting democratic input on what an AI should do, etc..

We believe the mission is more likely to succeed when society can co-evolve with AI advances, so we work to iteratively deploy our technology broadly and are excited to see an emerging constellation of AI efforts with a broad range of conflicting perspectives. Since next year's AI will be much better than today's, we see risks increasing the longer a given system exists only in a lab or is known only to a select few.

We regard building an AGI-level system as our first milestone. We plan for our activities to continue uninterrupted — though with increasing global coordination — into a world with a spectrum of AGIs and as humanity contemplates approaching superintelligence.

- gdb

On Thu, Dec 28, 2023 at 7:24 AM Larry Summers <lhs@openai.com> wrote:

> I note that musk s thing is like anthropic is a benefit corp.
>
> I think we are agreed that looking into this or things related to it should be an early priority for the new board.
>
> Reading a bit about benefit corps it appears that one salient aspect involves an annual report on meeting public responsibilities and the like. It occurs to me that within our current structure we could prepare and issue such a report. It might preserve flexibility to do such a thing and it would not be much of a lift.
>
> Sent from my iPad

--

- gdb