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. DX-0502
Date Entered
Mark Busby, Clerk
By: , Deputy Clerk
Δ DEFENDANT Δ

From: Elon Musk
To: Sam Altman
Subject: Re: AI lab
Date: Wednesday, June 24, 2015 11:05:44 PM

Agree on all

On Jun 24, 2015, at 10:24 AM, Sam Altman [redacted] wrote:

<!--[if !supportLists]-->1) <!--[endif]-->The mission would be to create the first general AI and use it for individual empowerment—ie, the distributed version of the future that seems the safest. More generally, safety should be a first-class requirement.

<!--[if !supportLists]-->2) <!--[endif]-->I think we'd ideally start with a group of 7-10 people, and plan to expand from there. We have a nice extra building in Mountain View they can have.

<!--[if !supportLists]-->3) <!--[endif]-->I think for a governance structure, we should start with 5 people and I'd propose you, Bill Gates, Pierre Omidyar, Dustin Moskovitz, and me. The technology would be owned by the foundation and used "for the good of the world", and in cases where it's not obvious how that should be applied the 5 of us would decide. The researchers would have significant financial upside but it would be uncorrelated to what they build, which should eliminate some of the conflict (we'll pay them a competitive salary and give them YC equity for the upside). We'd have an ongoing conversation about what work should be open-sourced and what shouldn't. At some point we'd get someone to run the team, but he/she probably shouldn't be on the governance board.

<!--[if !supportLists]-->4) <!--[endif]-->Will you be involved somehow in addition to just governance? I think that would be really helpful for getting work pointed in the right direction getting the best people to be part of it. Ideally you'd come by and talk to them about progress once a month or whatever. We generically call people involved in some limited way in YC "part-time partners" (we do that with Peter Thiel for example, though at this point he's very involved) but we could call it whatever you want. Even if you can't really spend time on it but can be publicly supportive, that would still probably be really helpful for recruiting.

<!--[if !supportLists]-->5) <!--[endif]-->I think the right plan with the regulation letter is to wait for this to get going and then I can just release it with a message like "now that we are doing this, I've been thinking a lot about what sort of constraints the world needs for safefy." I'm happy to leave you off as a signatory. I also suspect that after it's out more people will be willing to get behind it.

Sam

From: Sam Altman
To: Elon Musk
Subject: AI lab
Date: Wednesday, June 24, 2015 10:24:54 AM

--[if !supportLists]-->1) <!--[endif]-->The mission would be to create the first general AI and use it for individual empowerment—ie, the distributed version of the future that seems the safest. More generally, safety should be a first-class requirement.

--[if !supportLists]-->2) <!--[endif]-->I think we'd ideally start with a group of 7-10 people, and plan to expand from there. We have a nice extra building in Mountain View they can have.

--[if !supportLists]-->3) <!--[endif]-->I think for a governance structure, we should start with 5 people and I'd propose you, Bill Gates, Pierre Omidyar, Dustin Moskovitz, and me. The technology would be owned by the foundation and used "for the good of the world", and in cases where it's not obvious how that should be applied the 5 of us would decide. The researchers would have significant financial upside but it would be uncorrelated to what they build, which should eliminate some of the conflict (we'll pay them a competitive salary and give them YC equity for the upside). We'd have an ongoing conversation about what work should be open-sourced and what shouldn't. At some point we'd get someone to run the team, but he/she probably shouldn't be on the governance board.

--[if !supportLists]-->4) <!--[endif]-->Will you be involved somehow in addition to just governance? I think that would be really helpful for getting work pointed in the right direction getting the best people to be part of it. Ideally you'd come by and talk to them about progress once a month or whatever. We generically call people involved in some limited way in YC "part-time partners" (we do that with Peter Thiel for example, though at this point he's very involved) but we could call it whatever you want. Even if you can't really spend time on it but can be publicly supportive, that would still probably be really helpful for recruiting.

--[if !supportLists]-->5) <!--[endif]-->I think the right plan with the regulation letter is to wait for this to get going and then I can just release it with a message like "now that we are doing this, I've been thinking a lot about what sort of constraints the world needs for safefy." I'm happy to leave you off as a signatory. I also suspect that after it's out more people will be willing to get behind it.

ɜm