

# **EXHIBIT 15**

**From:** Elon Musk <erm@spacex.com>

**To:** Ilya Sutskever <ilyasu@openai.com>

**Cc:** Sam Altman <shga@ycombinator.com>, Greg Brockman <gdb@openai.com>, "Sam Teller" <steller@spacex.com>, Shivon Zilis <Shivon@spacex.com>

**Subject:** Re: Honest Thoughts

**Date:** Wed, 20 Sep 2017 14:17:02 -0700

**Importance:** Normal

---

Guys, I've had enough. This is the final straw.

Either go do something on your own or continue with OpenAI as a nonprofit. I will no longer fund OpenAI until you have made a firm commitment to stay or I'm just being a fool who is essentially providing free funding for you to create a startup.

Discussions are over.

On Sep 20, 2017, at 2:08 PM, Ilya Sutskever <ilyasu@openai.com> wrote:

Elon, Sam,

This process has been the highest stakes conversation that Greg and I have ever participated in, and if the project succeeds, it'll turn out to have been the highest stakes conversation the world has seen. It's also been a deeply personal conversation for all of us.

Yesterday while we were considering making our final commitment given the non-solicit agreement, we realized we'd made a mistake. We have several important concerns that we haven't raised with either of you. We didn't raise them because we were afraid to: we were afraid of harming the relationship, having you think less of us, or losing you as partners.

There is some chance that our concerns will prove to be unresolvable. We really hope it's not the case, but we know we will fail for sure if we don't all discuss them now. And we have hope that we can work through them and all continue working together.

**Elon:**

We *really* want to work with you. We believe that if we join forces, our chance of success in the mission is the greatest. Our upside is the highest. There is no doubt about that. Our desire to work with you is so great that we are happy to give up on the equity, personal control, make ourselves easily firable — whatever it takes to work with you.

But we realized that we were careless in our thinking about the implications of control for the world. Because it seemed so hubristic, we have not been seriously considering the implications of success.

- The current structure provides you with a path where you end up with unilateral absolute control over the AGI. You stated that you don't want to control the final AGI, but during this negotiation, you've shown to us that absolute control is extremely important to you.
  - As an example, you said that you needed to be CEO of the new company so that everyone will know that you are the one who is in charge, even though you also stated that you hate being CEO and would much rather not be CEO.

- Thus, we are concerned that as the company makes genuine progress towards AGI, you will choose to retain your absolute control of the company despite current intent to the contrary. We disagree with your statement that our ability to leave is our greatest power, because once the company is actually on track to AGI, the company will be much more important than any individual.
- The goal of OpenAI is to make the future good and to avoid an AGI dictatorship. You are concerned that Demis could create an AGI dictatorship. So do we. So it is a bad idea to create a structure where you could become a dictator if you chose to, especially given that we can create some other structure that avoids this possibility.

We have a few smaller concerns, but we think it's useful to mention it here:

- In the event we decide to buy Cerebras, my strong sense is that it'll be done through Tesla. But why do it this way if we could also do it from within OpenAI? Specifically, the concern is that Tesla has a duty to shareholders to maximize shareholder return, which is not aligned with OpenAI's mission. So the overall result may not end up being optimal for OpenAI.
- We believe that OpenAI the non-profit was successful because both you and Sam were in it. Sam acted as a genuine counterbalance to you, which has been extremely fruitful. Greg and I, at least so far, are much worse at being a counterbalance to you. We feel this is evidenced even by this negotiation, where we were ready to sweep the long-term AGI control questions under the rug while Sam stood his ground.

### **Sam:**

When Greg and I are stuck, you've always had an answer that turned out to be deep and correct. You've been thinking about the ways forward on this problem extremely deeply and thoroughly. Greg and I understand technical execution, but we don't know how structure decisions will play out over the next month, year, or five years.

But we haven't been able to fully trust your judgements throughout this process, because we don't understand your cost function.

- We don't understand why the CEO title is so important to you. Your stated reasons have changed, and it's hard to really understand what's driving it.
- Is AGI *truly* your primary motivation? How does it connect to your political goals? How has your thought process changed over time?

### **Greg and Ilya:**

We had a fair share of our own failings during this negotiation, and we'll list some of them here (Elon and Sam, I'm sure you'll have plenty to add...):

- During this negotiation, we realized that we have allowed the idea of financial return 2-3 years down the line to drive our decisions. This is why we didn't push on the control — we thought that our equity is good enough, so why worry? But this attitude is wrong, just like the attitude of AI experts who don't think that AI safety is an issue because they don't really believe that they'll build AGI.
- We did not speak our full truth during the negotiation. We have our excuses, but it was damaging to the process, and we may lose both Sam and Elon as a result.

There's enough baggage here that we think it's very important for us to meet and talk it out. Our collaboration will not succeed if we don't. Can all four of us meet today? If all of us say the truth, and resolve the issues, the company that we'll create will be much more likely to withstand the very strong forces it'll experience.

- Greg & Ilya