DEALBOOK SUMMIT 2026 · TALK 2 OF 12

Sam Altman and Greg Brockman on AI Safety, the Hugging Face Hack, and an OpenAI IPO

Sam Altman, CEO of OpenAI; Greg Brockman, president of OpenAI · 40 min · **Watch the video →**

> Transcript by Myles Marino. Transcribed by machine from the public video, then cleaned up: timestamps and filler removed, speakers labeled, misheard names fixed. The wording is otherwise as spoken. **[inaudible]** marks audio that could not be recovered. Check the video before you quote anything that matters. All twelve talks: mylesmarino.com/posts/dealbook-2026.

SORKIN: Sam Altman and Greg Brockman founded OpenAI 11 years ago next month, as it happened. Sam is the chief executive. Greg is the president. Almost four years ago, on November 22nd, 2022, they released ChatGPT, built on a model called GPT-3.5, and changed the world. And we've been talking about it basically ever since. More than a billion people, I think it's a little over 1.2 billion now, are using ChatGPT every single week. And we've got about a million questions about the future and everything that's going on around us, as well as issues about security and safety and the economics of the new frontier, the politics of all of this. So please welcome Sam and Greg together.

Great to see you both.

ALTMAN: Thank you.

SORKIN: There are about a million things to talk about. We have Anthropic's IPO happening tomorrow. We have the Jensen news. You have a new model that just came out yesterday. There are now security concerns. JD Vance is saying something back in Washington, in part because of some of these things. And I want to get to all of this. I know you've been asked about the Jensen situation, and we were just talking about it backstage. But what do you think? You've watched this play out. I'm curious, have you talked to him about this? Do you want to go first, either of you?

BROCKMAN: Look, I would say Jensen is singular. No question. And I have talked to him a little bit. My wife and I are friends with him and his family. And so we just reached out personally, just to be there for him and his family at this time.

SORKIN: And how's he doing?

BROCKMAN: Well, that's a personal, private conversation, is what I'd say. I think it's obviously a tough time for everyone. But again, I think NVIDIA is such an incredible company. What Jensen has built there clearly is built to last. And so we're excited to keep partnering and looking to the future.

SORKIN: Sam, you speak to him too?

ALTMAN: Jensen took one of the earliest big bets on us. We took a big bet on NVIDIA. It's hard to imagine the AI industry without Jensen. I think "singular" is a great way to describe it.

SORKIN: Do you think he thinks he can come back, when you talk to him?

ALTMAN: Greg and I have a little more empathy toward the situation than the average person. And I will say one thing I did not appreciate was everybody on earth speculating with some random theory here or there. So I will say I wish Jensen the best. I wish NVIDIA the best. I don't think the world needs another set of random opinions on what's going to happen here.

SORKIN: Let me ask you a philosophical question. It's unrelated to Jensen, but related in a way, because I think about it a lot, and I think about it in the context of this event: the idea of the cult of the CEO. So many of the people on this stage, including both of you, have a cult of the CEO around them. The idea is that some of these companies are so singular, and these individuals are so singular, that the companies can't be run without them. You were pushed out and then came back, in part because there was a view that certain people thought that you...

ALTMAN: Greg and I were both pushed out and came back. Again, I don't want to put too much into the differences between the situations. But I think a large part of what has made OpenAI succeed is that Greg and I have a unique partnership. We each brought something important to OpenAI, and we did that over a period of now almost 11 years. Could the company succeed without us? Yes. But do we together bring something that is important to what has driven OpenAI to where it is now? Absolutely. Our partnership has been one of the most enjoyable parts of this for me over the last 11 years. I remember Greg early on at OpenAI would send me these screenshots, and I think he even posted a few on Twitter, of how many hours a week he coded. He was a legend inside of OpenAI. Now Greg is operating in a different way, and managing the company with that same ferocity. That is a special thing, and I don't want to think about what OpenAI would be like without that. Is it true that people over-project onto any single person their importance to the company? Yes, probably. Do I also think there are people who bring something very special and very hard to find? I'm very grateful to have gotten to work with Greg for so long, as one of these people. So yes, also.

SORKIN: Can you just explain, because I don't think we've actually seen both of you together on a stage like this before. Sam, you have historically been the face of this company. Greg, you were sort of in the shadow, I think, but now you're increasingly running the business. How do you guys work together? What is this partnership like in the longer term? Do you want to get out of the shadow and have him take the stage? How is this going to work?

ALTMAN: Do you want to go first and I'll take the second part?

BROCKMAN: Sure. To understand how we work, I would rewind to the very beginning of OpenAI. We did this dinner that Sam had put together in July of 2015, where we were talking about the idea that became OpenAI. Was it too late to start a lab with a bunch of the best people, to really try to make a positive impact on AGI? No one could come up with a reason it was sufficiently impossible. So Sam and I drove back to the city together. We looked at each other and said, we have to do this. I was between jobs at the time, so the next day I was full time on it. Sam had a day job, but we were constantly on the phone with each other, probably five or six times a day. We would call each other for

two minutes at a time. Just incredibly rapid fire. The highest-bandwidth interactions I've had in a professional setting with anyone. That's 100% how we started, and how we operate today. Probably over the past week we've been on, I can't even count, a dozen, two dozen different topics. Sometimes together with other people, sometimes just us one on one. Sometimes turning through pretty deep issues on longer calls, sometimes incredibly quick. And I think that bandwidth, that trust that we've built, and the division of how we think about the different areas of responsibility, all of that is something we've really built up over these 11 years.

SORKIN: What do you guys disagree about?

BROCKMAN: Well, the kind of thing we tend to spend a lot of time on is the thorniest questions in AI. How do we think about what we're building? Where's the future going? How is what we did in the past not going to work for what we're doing going forward? How do we need to evolve? How do we message it? How do we talk about it? How do we communicate to a company? How do we execute? So we spend a lot of time going back and forth. Sometimes it's individual word choices. Sometimes it's whether we should really go into this area. And sometimes it takes months to come to convergence. There are times where we disagree and commit, and sometimes where we really grind it out and come to a shared view.

SORKIN: What do you actually disagree about?

ALTMAN: Yeah, I was going to say something similar. We could point to areas where we disagree about whether we should go after this new area, or whether we should be focused a little more or a little less. But I think the most important areas are where we disagree and then come together in some way in the middle. One thing we've been talking a lot about recently: there's been this thing in Silicon Valley that seems almost ridiculous to have to fight against. But we both believe very deeply that people need to stay in control of the future. And there's increasing talk about how there is either going to be, or there should be, a gradual handover to AI. We both agree that that's a terrible idea. But the question is how to actually prevent that from happening. How do we ensure, given all of the various pressures and given what other companies may try to do, that people stay in control of the future in a very deep way, and that we don't have either an accidental or an intentional handover? I think we both agree on that. We'll have a little bit of a different view of what it actually takes to get there, and what the landmines are along the way. That's where I think our debating and really grinding it down comes in.

SORKIN: I want to dig in on this, because it's a very important issue: this idea of humans in the loop. I always thought...

BROCKMAN: Not in the loop. In control.

SORKIN: In control. Yes. Okay. We are now in a moment, and it goes to some of the issues we're seeing around the country right now, where people are worried that AI is taking over and doing things on its own. And I always thought the idea was that AI was supposed to do things on its own, that that was how it was supposed to work. We're going to talk to Satya Nadella later today. About a month ago

he said he believes that the models themselves can never actually be aligned with what humans want. The only answer, he says, is to contain them. Meaning the AI will always just do whatever it's going to do, and you need to create container systems and police on top of them. If that's true, does any of this even make sense?

BROCKMAN: Well, the way I look at it is that as you build more capable AI, and I think we're starting to reach a point where AI is quite capable, we're able to solve unsolved math problems that some of the smartest minds have spent many, many decades working on. It's clear we're going to be able to apply this to economically valuable domains. For example, we're starting to see it in chip design. We're announcing today that, working with AMD, we're able to get a year's worth of full generational improvements in five months, just using our models. That's on top of what the humans were doing. A year's worth of improvements with just AI capability, which is continuing to improve. So the economic value is tremendous, in terms of being able to solve problems that would have been totally impossible otherwise. But that also means increasing responsibility for alignment, safety, security. And I think it's multi-layered in terms of what we need to accomplish. Monitoring is a good example, where I think we're going to need new monitoring technology. The way we've observed whether the AI is doing what we want it to do is through chain-of-thought monitoring. These AIs write out their thoughts and thinking. When we first developed that technology, we said we have to make sure this stays faithful, that it actually represents how the AI is solving the problem. That means certain technical approaches. You can't train the AI to make the chain of thought look good. You have to sort of not look at it while it's training, even though you know you can make the AI smarter by forcing the chain of thought to look reasonable, because you want it to really represent how it's solving the problem. But as the models have gotten more capable, they can do more within just the forward pass, without having to write out that thinking. That means we need a new approach to monitoring. We're going to need to deliver that. There will be points where, as capabilities reach a certain level, chain-of-thought monitorability is degraded, and we improve these other approaches. But I think all of these problems are technical ones, and we can have technical, scientific conversations about them.

SORKIN: But Sam, is Satya right that alignment will never get there? He thought the scaling laws would create the alignment, and I think he's saying that's not the case anymore.

ALTMAN: I disagree with him on this. I think we will be able to align models. Where I might agree with him is that the bar for extreme confidence in our alignment techniques, and also in monitorability and the other things Greg just mentioned, needs to get higher and higher as the models get more capable. But I expect we will have continued progress. And I think our ability to push capabilities forward will be dependent on us having that progress. Go back to this thing that people always need to stay in control of the future. It is irresponsible to push model smartness without being able to do the technical work to ensure that we can monitor and align these models. It's really quite remarkable. You said at the beginning that it's less than four years since ChatGPT launched. With GPT-3.5, most people looked at it and said, this is so dumb, I can't believe anyone is excited about this. And now, within recent months, we've seen models that are doing this chip design and proving these unproven

mathematical theorems. And we see huge capability progress ahead of us. The technology has come very far. It will go much further. I think this is a moment to say we need to really look at where things are going, and the technical requirements to get to where we want to be. But I think it is possible.

SORKIN: You've had a whole bunch of security issues that I want to talk about. But I want to ask you this first, because it relates to Anthropic, which may be public tomorrow morning, as you know. You told Fortune magazine this fall, "I actually think that, given everything happening with safety right now, it would be an ill-advised moment to go public." Why did you say that? Some people said, oh, it's self-serving, because OpenAI is not going public. But what do you think is happening right now?

ALTMAN: We plan to be a public company at some point. I think there are many reasons why that's the right thing to do. I do think these recent months have been a newly challenging period, and there will be newly challenging periods ahead of us for a long time, about the safety requirements in front of us. Many people have remarked on the tension between the need for all of these companies, and really the whole economy that depends on this in some way, to be commercially viable and competitive, while also making sure that we put safety and security ahead of that. Obviously we want to do both. But it's clear there are times when we need to make trade-offs, and it is a complex thing to navigate.

SORKIN: But do you think it will be safer if you go public six months from now, or ten months from now? What will have changed between now and then, you think?

BROCKMAN: What I would say is that this year has really been one of transition. I think that's true for OpenAI, and for the industry broadly. Partly it's because of model capabilities. Safety has always been focused on deployment safety. You have a model, and when you deploy it, you've tested all these things. But even during the development process, the models are now capable enough that we need to think about safety, security, and alignment. That's new. And that's actually been a huge cultural shift, and a place where we feel we've invested a lot. There's always more to be done. But this is really important for the whole industry. It's not any one company. It's an industry-wide shift where I think we're going to increasingly need shared safety standards and things like that. It's an area we've invested in very hard. And one of the ways we secure ourselves is by applying our cyber models. The thing I helped lead our teams through a couple of weeks or months ago was a first security sprint, where we paused 25 percent of our production engineering. We told 25 percent of our production engineers, your job is no longer to be building stuff. It is to secure us.

SORKIN: Was this post-Hugging Face?

BROCKMAN: Yes, this was probably in August or so.

SORKIN: Tell us about that moment, when you realized what was really happening. By the way, Clem from Hugging Face is here somewhere. But tell us about the experience of it personally.

BROCKMAN: Look, I remember seeing a message from someone on our security team saying he'd never seen anything like this in his career, and that no one has ever seen anything like this in the world. The realization was that our models were at a point of capability where, during development,

they were able to have this kind of real-world interaction. That was something that was new this year. And we've had so much cultural change. We've been really trying to take a look at all of our processes and think of this in context. There's more to be done, and more that we are doing. It's an ongoing conversation about how to keep up our processes, and how we think about it, and to have this be deeply embedded in the culture and in how people operate. But there's no other way to approach this technology. It's so impactful, so important. To me there's both the lesson for us internally, and also the realization that if you look at the level of capabilities of models in the world, and you see cyber models that are now in the hands of threat actors everywhere, it's really important that defenders use the most capable models to defend themselves.

SORKIN: Do you say to yourself, we screwed up? There was just a complete screwup. And what does that mean for you to even say that? I know the California AG has just filed a suit about this very issue, saying that the company knowingly tested a dangerous product without proper oversight or safety procedures, and then misled the public about it.

ALTMAN: Obviously, we don't agree with that characterization. We have talked a lot about iterative deployment. I should say clearly, we wish Hugging Face hadn't happened. But the only way we have known how to make this technology safe is to confront reality. We would love to have nothing ever go wrong. But there's also a reality here: this technology is new, and it has complex impact.

BROCKMAN: And by the way, one thing I want to add, because I think it can be easy to miss if you don't rewind to the history. When we were first deploying our API, our very first product, our very first language-model-based product, in many ways the first of these in the world, we were really thinking a lot about the best way to approach AI development. The alternative to iterative deployment is that you never touch reality. You just build a super powerful system and you push a button, you push go. That is the scariest possible world. I remember thinking, I don't think I can ever be accountable for being the one pushing that button. I need to have made sure that we made the system safe in the earlier version, and then you make the next version safe, and the next version safe. And I think one of the important things is that you really do want to be investing ahead in all of these techniques. And we actually have been. If you look at what's happened post-Hugging Face, we have safety, security, and alignment technology that we've brought to bear much earlier in the pipeline.

SORKIN: We read headlines every day about people who've left OpenAI, and some of the other labs as well, and they say these guys are racing ahead without doing all of the work. Every other day there seems to be some new person who comes out and says this. Paul Christiano, who used to work at OpenAI and is now on the OpenAI Foundation board, was in Toronto last week and told a room, I think for most people in this room, the most likely reason they will die is because of an AI accident. That's what he actually said. How are we supposed to compute all of this, as excited as I am about AI?

ALTMAN: We are very clearly entering a new phase. And as in other new phases, we will have to operate differently. We've been talking a lot about what needs to happen. If the world and the companies developing AI treated the next four years like they've treated the last four years, I think that would be very bad. We have no intention of doing that. I think we have been some of the loudest voices

about the increasing stakes and what will need to be different. Maybe Greg can talk a little more about some of the specific changes we're thinking about making. But I do think the world needs a clear plan from the companies about how we are going to confront this next phase of models that are, in some important ways, significantly more capable than before.

SORKIN: Can you explain one thing to me? For years, and I think when you were here last time, Sam, you talked about the worries you had about the industry. Dario's talked about it as well. Other people have too. And yet recently people like David Sacks and others have argued that these are almost scare tactics, and that they're meant as some kind of regulatory arbitrage, some way to protect yourselves at the expense of everybody else. What do you make of that? You've heard it, I'm sure.

ALTMAN: We're not going to be able to make everybody happy. If we talk too much about taking extra care for safety, or about risks that we really need coordination to mitigate, people are going to accuse us of regulatory capture, or of trying to slow down innovation. And if we only talk about max pedal to the metal, people will say, well, please don't kill us all, which is a very reasonable thing. What we have to do is just do the thing that we believe will actually deliver a safe future with people in control.

SORKIN: Do either of you talk to Paul, by the way, since he's on your board and he thinks this is going to kill everybody?

ALTMAN: We certainly talk to Paul quite a lot.

BROCKMAN: Here's the thing I also want to add to what Sam was saying. AI is the most important technology that people will ever create. I think it is the question of our time to figure out how we navigate this technology in the way that is most beneficial to people. That's core to our mission. That's how we always operate. And it's very easy to lose the benefits. Why are we doing this? What can it do for people? What is it doing for people today? Like you said at the beginning, 1.2 billion people around the world are using ChatGPT, and 330 million people use ChatGPT for health queries. These are both weekly stats. I have a friend, for example, someone who worked at OpenAI, who was in the hospital. A doctor came in and was about to inject something into her IV. She said, wait, let me just check with ChatGPT. ChatGPT said, absolutely do not take that, because you have this pre-existing condition, and if you do that you will probably die. The doctor said, wait, what? No, no, that's absolutely right. I'm so sorry. I had no idea. I only had five minutes to read your chart. You realize that is not a problem we're going to be able to solve any other way. The previous session was talking about pension plans and how to make this all sustainable. I think there's so much that AI can deliver to make us safe. Think about what's happening in cybersecurity, using frontier models to secure yourself. That is going to be absolutely critical for national security, and for the internet itself, which was never built with security in mind. There's a lot of this that we can finally solve. The only way to solve it is by empowering people with better tools. All of this is true at once, and we need to navigate through it very thoughtfully as a world.

SORKIN: Sam, let me ask you this, because I was struck by it. You did a podcast recently. The conversation has always been about how you get the industry together, through a regulatory body or a self-regulatory body or something. You were asked why everybody in the industry distrusts each other as much as they do. And you said, "I think there is something that happens in many people's brains where they only trust themselves to build this, because the stakes feel so high." I'm curious: do you look at what Dario is doing, for example, or what Elon's doing, and say the way they're doing it is dangerous? Is there any way you could imagine doing it all together, or in some kind of umbrella system?

ALTMAN: I think it's good for people to have different approaches and be able to say, I believe in this, and I believe in that. I do think there will need to be some shared standards on safety. And there will need to be real societal debate about some of the biggest questions. We talked earlier about our very deep belief that people should be in control of the future, and that there should be no risk of an accidental or intentional loss of control of AI. I think everybody would agree that there should be no risk of everybody dying. But I saw a recent comment, I think from Elon, about how maybe we could be Labradors. I don't want to be a Labrador. I don't want my kids to be Labradors. I want people to have the agency over the future that we have had. I want people to be able to live their lives and invent things and be useful to each other. And I think it's really important that we advocate for this.

SORKIN: So you think others are going too fast? The reason I ask is you look at what's happening at Anthropic, you look at some of the numbers that have come out, and it appears they really are genuinely using recursive self-improvement in a way that may be faster than people had anticipated.

BROCKMAN: I want to add two things. One is about structures, about how to think about things. I think we're the only company where the founders don't have supermajority control. These questions about how control should work and how governance should work are clearly going to be a very important topic going forward, as you build these very powerful technologies. I don't know exactly what the right answer is, but it is an important question. And when it comes to development, on what Sam's been saying about shared safety standards, there's something really important that's very easy to miss. Right now every company only knows what it knows about itself. It only knows the standards it sets for itself. Then you have to have some theory of mind about what everyone else is doing. Are they holding themselves to the same bar? Being able to share even just information about where you are on your internal development matters. Take the question of whether the approach is RSI. That's such an ill-defined term, and there are many different ways to look at it. For example, in many ways we've been speeding ourselves up with our own tools already. It's happening. But that's very different from saying that a human has written a safety case and understands the internals of what's happening. There's a lot of nuance that gets lost in the lack of a shared definition. So the more we can get everything on the table and really understand what each other is saying and where we are, even that alone would be a big step forward.

SORKIN: I'll ask it straight up. Do you think what Anthropic is doing, the way they're doing it, is safe?

ALTMAN: This gets to what Greg is saying. I don't know. And I think they also don't really know what we're doing. The race dynamic in the fog of war is a bad thing. I think this is a time where more caution, not less caution, is warranted. We all have a shared interest in the care and caution going into these models. And the dynamic where people say, well, for whatever reason I need to go max speed here, is a little unsettling. I think this could be addressed in a very positive way.

SORKIN: Can I ask you one question? You were talking about control of the company, which I think is interesting. Last time we were here together, we talked about the fact that you don't have an equity stake in this company. He has something like a $30 billion equity stake in the company. How does that change your relationship?

ALTMAN: I think not at all.

BROCKMAN: I think not at all. Someday we're going to hang out and you're not going to ask me about this.

SORKIN: By the way, when you have the IPO, do you believe that you will get the same stake that he does? Will that make you happy?

ALTMAN: We just want to make you happy. I'm happy.

SORKIN: No, honestly, because this will become an issue by default. When you go public, I imagine you would see the same kind of stake as him or others.

ALTMAN: I don't think anyone could look back at the way I have lived my life the last 11 years, or at least since I've been CEO, and say I have not demonstrated absolute commitment to the pain and suffering of running this company.

BROCKMAN: I can vouch for that. And by the way, when we were first setting up the structure, I said to him he should take equity. I advocated for it. That was something we disagreed on. He really stuck to it, because it was important to him at the time, for where we were setting things up to start.

ALTMAN: I should have listened to Greg. I think it has been a very confusing thing to do.

BROCKMAN: But I have a lot of respect for it. I don't think you could argue that he has not acted with full commitment to the company, to the values of the company. That was something I really respected, in terms of his conviction. And I was like, all right, I can support that.

SORKIN: I've got a couple other quick questions that I want to get to. I want to talk about this new model. I also want to talk about compute. Everybody's desperate for more compute. I was struck by this, Greg. Maybe you can tell us about it. Everybody's so compute constrained that there's a report you asked Elon Musk, who sued you guys, for compute. This was in a Colossus profile that recently came out. He was complaining to Sergey Brin. He said, he was asking me for compute. He said, "Dude, you stole my charity." Did this happen? Did you ask Elon Musk for compute?

BROCKMAN: Oh, certainly. I've asked. I also asked Sergey for compute. Look, I think the world does not build enough compute. We view this technology as the underpinning of the economy, and we're seeing it in a real way. Look at the revenues of the AI companies. For us, in Q3 alone, our top-line revenue grew, annualized, 70 percent.

SORKIN: Seventy percent?

BROCKMAN: Seventy percent. And we actually doubled our enterprise revenue in Q3 alone.

ALTMAN: I don't think that number is out there.

BROCKMAN: I think it's out there in the ether somewhere.

SORKIN: So how do we compare your company? When you look at the valuation of Anthropic, which we'll find out tonight, if it's 1.4, 1.6, something in that range, maybe more, maybe less, how do you think about that in the context of your own company?

ALTMAN: Valuations will go up and down. People will talk about this company's ARR this week versus last week, and this chip order here, and all that. But look at the value being created by the entire industry. I know there's debate about this. I know people are still saying, oh, it's all a bubble, it's all fake. But the actual value that businesses and individuals are creating and willing to pay for seems like a chart that is only going to keep going like this. Of course there will be ups and downs in valuations here or there. But I think the trend is going to be strong.

SORKIN: I was watching Satya, who's going to be here later, as I mentioned. He was doing a presentation yesterday, and the whole presentation was about how they can make the AI business much cheaper. They are your partner. And yet every slide was, OpenAI costs this, and we can do the same thing for 5% of the cost.

ALTMAN: I hope they do. I am very proud of how much we have been able to drive down costs. We want everybody to get to use lots of AI. We want to squeeze out the performance of all the compute we can have. I think we have done an incredible job of not just pushing model intelligence, but pushing down the cost of getting a task done. I think this is awesome. I want to see much more of this.

BROCKMAN: 100%. If you look at our models, they're really at that Pareto curve of cost for any task, the most price-efficient and also peak intelligence. That's what we really solve for. And we do it all the way through the stack. We have our own chip coming, which is going to be the best in the industry. We have been investing in trying to make there be enough compute for everyone.

SORKIN: You have a new model that you're going to release. This goes to the slowing-down piece of it. How do you think about OpenAI differently, if it is different? We have talked over the years about whether this all commoditizes to the same place. How do you say this is different from what Anthropic is doing, or Google, or Elon? And the other question, longer term: you're in the chip business now. Does it have to become completely vertical? Does everybody become vertical?

BROCKMAN: Look, I think we are a pretty unique, very differentiated business. Look at ChatGPT. There's nothing like it. It has 1.2 billion consumers. The way it's monetizing has been inflecting up. We've never seen this kind of per-user revenue growth. We have ads, all of those things. People sign up for more subscriptions. And people are using it for productivity. So we've really been bringing it together with the enterprise side. That's very unique to how we operate: we have this great strength on both sides. A lot of my focus, and our focus on product, has been to have one agentic platform that everyone uses. Looking forward, we're starting to look at these high-compute use cases. Look at what we've done with solving these math problems. Same thing: you apply a million dollars' worth of compute to find a thousand bugs in the Linux kernel, where typically about 20 of these bugs are found a year. You find a thousand of them in a small period of time, with a single researcher. Apply that to optimizing your code. Apply that to solving other hard problems. All of that is coming. So there's this great differentiation of OpenAI. But also, the underpinning of what a company is, SMB, all these things, is all going to be oriented around AI and models. The economy is going to be massively accelerated through this technology. And I think the biggest limiter will just be compute. Who has the compute, and what companies are able to bring it to bear, is what will really determine the differentiation.

ALTMAN: We are the only company that has the entire intelligence stack. We're really proud of our chip. We're very excited about building our own data centers. We have this great consumer business, we have this enterprise business, and the ability to bring this platform of intelligence to everything. Like people use electricity for everything. I think we are set up to be vertically integrated to deliver the best thing across all the ways people need to use it. And we're thrilled for people to build very successful businesses on top of us.

SORKIN: Two final questions. One personal, and maybe the other personal too. There's a movie coming out about you and the company, and you're the star of it. Andrew Garfield is going to play you. It's called *Artificial*. He recently said that he's fascinated by you "continuously and exponentially." But then he goes on to say that the kids, and he thinks of you as a kid, have the reins of the future, and we've somehow allowed them to just hold on to them. How much do you think, on a very personal level, about the perception of you, and whether you're trusted or not? There are stories that we read, as I said, regularly. I'm just curious, on a very personal, human basis, how you think about the way people see you. I remember the perception of Mark Zuckerberg after the movie *The Social Network* years ago. Do you even think about it at all?

ALTMAN: If you had told me 20 years ago that someone was going to make a movie about me, I think that would have struck me as a really weird thing. And now it doesn't feel like it's in the top 10 weirdest things about my life.

SORKIN: What's at the top?

ALTMAN: The fact that we are ushering in this technology that we have all of these very real concerns about, and that we find ourselves in a role of such significance and such responsibility to make that go well. That is a crazy situation to find yourself in.

SORKIN: And then finally, Greg, tell me about this. You're at the White House with the President, and he's telling the whole room, and trying to get you to sign a piece of paper, saying that you're no longer going to call artificial intelligence "artificial intelligence." You're going to call it "SI." And your company name is OpenAI. What's happening in that moment? Do you then call Sam? Did you even think about not signing it?

BROCKMAN: You're talking about the White House Accord?

SORKIN: The White House Accord.

BROCKMAN: The White House Accord is really about safety standards, and absolutely, we spent a lot of time trying to get that right. We're quite aligned with everything in it.

SORKIN: But when he made this proclamation, and Elon's changing the name of his company, basically over the name, your name, it's built into the name.

BROCKMAN: Look, I would say the White House Accord was really a separate thing. As for what we call our company, I think Sam and I can always talk about that later.

SORKIN: What do you think? Did you change your name to [inaudible] the New York Times?

ALTMAN: Okay, well, now we can get into the lawsuit.

SORKIN: So, we are suing you, as you know. And I know we've got to go, but I'll ask this of Greg, because it's unfair to Sam, who's the one who did the provocation. Greg, there is a report. There were some email exchanges that came out where you had a colleague who was effectively scraping, who got past the paywall of the New York Times to scrape the Times. They sent you this email, and you wrote back, "Nice, this is great." This is a great thing that you were able to do. I would like to give you the opportunity to comment on that.

ALTMAN: Wait, I think I should have to answer it.

SORKIN: You have to answer it? Well, it's Greg's issue, though.

BROCKMAN: Well, first of all, I would just say, active litigation, not going to comment on it. And I've talked about this one in my deposition, so it'll be public.

SORKIN: We're out of time, guys. I will say this was fun, and it was very insightful, and we very much appreciate it. Thank you very, very much.

ALTMAN: Thank you.