The Real Reason Elon Made a Deal with Anthropic

29 MIN · YOUTUBE · HTTPS://YOUTUBE.COM/WATCH?V=-Q9SVNVFOHM
https://youtube.com/watch?v=-q9SvNvfOHM

This is the most ridiculous [music] turn of events. No one saw this coming. Two of the biggest rivals in AI, >> [music] >> SpaceX and Anthropic, are teaming up to take down Sam >> [music] >> Altman at OpenAI. And the reason why this is so ridiculous is literally a month ago Elon was posting tweets like this. Anthropic [music] hates Western civilization. What else have we got? We've got your AI hates white Asians, heterosexuals, and men. This is misanthropic and evil all the way to today or yesterday where he announced that I've reviewed Dario, I've spoken to the senior team at Anthropic, and I have determined that they are in fact not evil. And Claude, I quote, will probably be good. This is a a huge turn of events. They've signed a big deal where Anthropic is going to buy $5 billion worth of SpaceX's GPUs, also known as Colossus 1, their first major data center. And in turn, they're going to double everyone's Claude restriction limits, and Claude is finally going to catch up with OpenAI. I am shocked that this partnership is happening, but it's great to see.

You got to love a CEO that changes his mind. The guy who hated Anthropic, he spent a week with them and he was like, "You know what? These guys aren't that bad." What did he say here? He says, "So long as they engage in critical self-examination, Claude will probably be good." He said, "Everyone I met was highly competent and cared a great deal about doing the right thing. No one set off my evil detector." So, he went, he did some background checks, and he said, "Yeah, we'll give you the entire Colossus 1 data center. I'm now your landlord." It's a really funny turn of events. I think it's important to highlight the context of why this is so important for anyone who has been following along. Anthropic has had a really difficult time keeping up with their compute demands, mostly because they have grown so quickly. In fact, yesterday at Claude's coding event, we had a clip from Dario who was kind of doing just a general state of the union, and he shared with the public that Anthropic was off by a factor of eight in projecting their growth. So, over the course of the year, they projected for about a 10x multiple on the amount of compute they would need. The reality is is that this year alone, they are now guiding for 80 times what they were expecting. A full 8x multiple higher than what they were projecting. So, therefore, they had a big problem. They couldn't serve the compute. The quality of Opus was getting much worse. They needed compute. Where is the only large coherent cluster that's sitting dormant?

Well, it's owned by XAI. It's owned by the SpaceX team. And they signed a deal to give them the entire thing. So, the whole Colossus 1 cluster that XAI was using to train their early models is now run by Anthropic. And as a result, like you mentioned you guys, they doubled the rate limits. And now Colossus 2, which is the new largest cluster, that's where SpaceX AI is going to be training their new data centers. Now, you'll notice I said SpaceX AI and not XAI because as of yesterday, XAI has been depreciated. It no longer exists. SpaceX is kind of rolling everything up into one. And now they are working together. Right down the street. Anthropic is serving AI from one data center and SpaceX AI is serving data from the other. And this is a pretty big change of heart in the AI race. This is a big deal. There's so many stories wrapped into this one story. Both positive and perplexing to me that like I don't know. I'm just going to mind dump on you for a bit Josh and like let me know where where you think I end up, okay? So, number one, SpaceX has been or XAI has been trailing behind the frontier labs, right? They didn't have enough compute. We haven't had a good pioneering Groq release for like ages at this point. Over well over a year. And so, the idea was, okay, if you amass enough GPUs, you can create a better model, but that hasn't been the case and an article broke last week that said that XAI hasn't even been utilizing all their GPUs. In fact, 11% of the million GPUs that they have hasn't been used for Groq. And the idea here would be that Groq just isn't as demand as in demand as some of these other models like Claude and ChatGPT. So, now we see Elon pivoting to a neo cloud to become an AI compute infrastructure provider.

And this isn't the first time that he's done it. Literally 2 weeks ago, he signed a deal with Corsa to do the exact same thing. And that, by the way, was for Colossus 2. So, although we're saying he's using Colossus 2 to train Grok models, he's also sectioning off a portion of that to allow Corsa to build their own coding model. Now, that's a bit of a different deal because the agreement that they've signed is if Corsa builds a really good coding model, hey, could you like lend that to X AI for Grok as well so we can pass it off as a Claude code competitor? So, that's a little more symbiotic between the two companies. But for Anthropic, it's pretty clear. Anthropic pays $5 billion.

Elon gets to register that as revenue so that, you know, X AI makes a bunch of money. And for Anthropic, they get access to 300 MW of power this month. That is a crazy amount of power to have online, and it's pretty clear from Darius' words in that clip that we just showed, um which is that Anthropic has been lagging behind on compute. We finally have the answer as to why we don't have Claude me those. In my opinion, it's not because it was too dangerous. It's because they didn't have enough compute to serve it in the first place. And the living proof is literally their biggest rival, OpenAI, which delivered 5.5 to everyone, and it is Claude me those level. So, the story writes itself. Yeah, now Anthropic, they're basically serving compute from every single hyperscaler. They have a deal with AWS, with Google, with Microsoft, with Nvidia, with Fluidstack, and now SpaceX. And SpaceX is the most exciting, the most noteworthy of these because this seems like a pretty large pivot in their business model. Now, they have Colossus 2, which is currently 500,000 coherent Blackwell chips in a single cluster. That's the most powerful in the world. So, hopefully we start to see some amazing stuff coming out of this SpaceX AI team. But more importantly, when you look at the SpaceX business model that's performed so well, it's very much a services industry. They're taking satellites and they blast them into outer space, and they have their own business built on top of it with Starlink, but they are that service provider and owning the rail has proved to be very lucrative as a business model. Now, what SpaceX AI is doing currently is kind of the same thing, but early stages where they have this infrastructure, they're incredible at building data centers faster than everyone else, and now they're leasing it out to other people who need it. This applies further out with space as well.

Now, in the announcement that they spoke about, they very clearly mentioned space AI data centers. Anthropic is very excited about this, and it seems as if they're going to be working together as well for the space part of the deal. So, once that Starship 3 rocket comes online this year, once they start getting payloads into space and they start launching satellites, Anthropic is going to have data centers on those rockets into outer space to help with their compute needs. So, not only is SpaceX AI currently sitting on the largest coherent cluster of GPUs with

Class 2, but they now own all of the rails and the only way of getting compute into outer space for themselves, but also for every other company in the world who's going to need more data centers that are going to need to go out into outer space. It's a pretty amazing business they have. I I I think this is a great move for Elon and XAI. I'm I'm in fact more bullish on XAI after they've made this announcement because Elon is just doing what like he does best. He's He like does a zero-to-one innovation and then he scales the crap out of that. He did that with electric cars, he's doing that with humanoid robots, and now he's going to do that hopefully with compute. And no one has built a data center as large and as powerful as he has in the time that he has. Do you know what I mean? So, he he just has a very impressive knack for scaling things. Um, I want to share this tweet which I think summarizes this entire story with a nice bow tie, which is XAI is a middling frontier lab and an incredible world-class infra provider.

Anthropic is an incredible frontier lab and a middling infra provider. They had to join forces in order to defeat OpenAI, which is both a good frontier lab with good infra. And that's very much how I see this playing out. Anthropic needed compute. X AI has a lot of compute needed to sell it, so it works vice versa. Um, it's important to note that having Anthropic as one of your major customers for compute is probably one of the biggest compliments you can get. Um, I just want to point out that Amazon, Google, and I believe Microsoft's 50% of their compute backlog is purchased by Anthropic themselves.

So, Anthropic is trying to get their hands on all the compute. That makes sense. And this is just another partnership that enables that. So, cool to see. Yeah, they're trying to play catch-up and it seems as if they have this It's funny, this was literally the only way they could have gotten This is the only compute cluster of this size available on Earth for them. And they somehow managed to secure it. >> Well, actually, Josh, there is another company whose stock has been underperforming for the last month who also has the second biggest cluster or data center and might also do the same thing because they their AI products aren't in You guess?

>> Oh. Can I take a guess? Um, yeah, speaking of products that are not in demand, we have Meta who has also spent a tremendous amount of money on AI and compute and has nothing to show for it or not much at least. So, is that who comes next? Is this the new type of business model? Did like SpaceX just start the new thing? Like, hey, if you suck at AI, perhaps you can just like use our compute until we figure it out and you pay us a lot of money for that. I I can do you one better. Um, Meta actually announced on their previous earnings uh Zuck in a note to to the public said, "We're making this massive bet because we can't afford to lose the compute race, but in the event that our AI products aren't in as demand as we project, we can just resell that compute to all the other AI labs that need it." So, he's preempted what Elon and Anthropic announced yesterday. So, I think Meta's next and it might actually be bullish for their stock, to be honest. And this gets back to our bubble episode from yesterday, which if you haven't watched, I highly advise. But basically, so long as the demand is there for tokens, and people are willing to pay for that, everything works. So, as long as we can continue to build valuable tools, people will figure it out, the tokens will get allocated, and things will be good. Now, what I found interesting, most interesting about this Elon post, is the fact that he said he spent last week with the Anthropic team. Because he he had a very busy week, clearly, because he was doing other things as well, most notably the trial that he's currently involved in.

So, he has been going to trial with Sam Altman, OpenAI against each other, debating whether or not it is okay for OpenAI to exist as a for-profit entity when it started off as a nonprofit. Now, thankfully, as we always do with these court cases, we get some good banter, we get some good clips, and we get some good intel. In this case, we have texts from Mira Murati, which, for those of you who are not familiar, she is one of the co-founders of OpenAI, who has since left to start their own business called Thinking Machines. But, we have these texts that happened in real time as Sam Altman was getting fired from OpenAI. If you all remember, this was what, 2 years ago now? Uh Sam was actually fired as CEO of OpenAI and completely kicked off the company.

Shortly after they brought him back, but these texts I read were so funny because it gives you an inside look at exactly what was happening as the news was going down. So, Sam Altman texts Mira, he says, "Can you indicate directionally, good or bad? Satya and others are anxious." As in Satya Nadella, CEO of Microsoft, who is the largest stakeholder of OpenAI. Mira replies, "Directionally, very bad." Sam goes, "Okay, can you wrap up soon? Lots of pressure from Microsoft for an update."

And then Mira just goes, "Sam, this is very bad." And this is a series of a lot. Sam's like, "Can I come in?" And then Mira says, "They don't want you to." And then Sam says, "What do you want to make it better? I'm still willing to just walk away if that helps. If they are ramped up for crazy lawsuits against me, then I'm not sure what." It's just so funny to see into his mind as this is going on, and you kind of just see the the chain of thought of a CEO. It's really funny. He asks like, "Can I come in and talk about [clears throat] a path forward with them?" They're saying, "No, they need more time."

And it's it's really funny. And then the one line that has turned into a meme is Mira mentioning the new CEO who they put in place. And that new CEO was um she says, "This new guy is a rando Twitch guy." And it turns out that random Twitch guy was Emmett Shear, who is the guy who turned Twitch into a billion-dollar exit and very successful. And he turned it into his ex-header profile background. And it was just it's a very funny story seeing the inside look at what communications looks like during a time as chaotic as this one. That That just all human. That's That's the best part about these lawsuits is we just the public gets to gossip on all of their texts, which just get revealed. Um now, the the other personal text that was revealed uh especially over the this week was Greg Brockman's diary. Now, we referenced this in a previous episode that we did literally last week covering the first couple of days of the trial. A lot has happened since. Um a few things a few highlights that I want to point out. Um number one, this uh this lady that you're looking at on the screen right now, uh her name is Shivon Zilis, and she is a a partner to Elon. They've had like four kids together, but she's also a genius. She works at Neuralink, I believe. And she was brought on as a witness by the uh Musk team to basically testify that, you know, in the early days cuz she was involved in the whole OpenAI founding, and she was like did a bunch of work for Elon at that point. Um she wanted to point out that like, you know, they they being Sam and Greg Brockman had issues or intentions to make it a for-profit eventually and to testify against Sam's nature. She was brought to the witness stand as well as a host of other Musk team witness candidates. Now, the other revealing thing was uh Greg Brockman's diary was exposed, and he explicitly says in that diary or in his list of exchanges with Sam that we should kind of be strategic about when we turn OpenAI into a for-profit. I don't want to do it whilst Elon is involved or is a stakeholder, otherwise it would seem like we are ousting him. So, literally the case that Elon Musk's team is making is being proven in the words that Greg is saying.

Now, his response to that in court was, "I was just referring to it being when it turns into a billion dollars, not my $29 billion Sorry, not my $30 billion stake that I currently own at OpenAI right now." And Musk's team is basically saying, "You own $30 billion worth of OpenAI. How is this not a for-profit? Like, how are you not somewhat incentivized to cash this out and make more money or turn it into a for-profit?" And his response was, "I was just talking about making a billion dollars, not not the extra, you know, $29 billion." To which Elon's team responded saying, "Why don't you donate the $29 billion then? Like, and we can settle this lawsuit up." And so, there's a lot of like back and forth going on between them. Now, it's been very Elon-centric in the news for this court case in in particular because Elon's been the first to step up and produce his witnesses. But, next week is when the script officially flips and Sam takes the witness stand as well as OpenAI's witnesses, and we might see a similar case where they end up shooting themselves in the foot similarly to how Elon has happened for the first week of the trial.

And what I'm really enjoying is how now they're available via audio, so you can actually tune in and listen. You can't watch the presentations, but you can listen to the testimonies. People have been doing it and kind of sharing the hot takes and the information from it. It's very invasive in a way, but in the sense that you just mentioned earlier where people are just human. Like, obviously Greg Brockman wants to make a billion dollars, and sure you want the nonprofit, but he's a human. Why would you not? And it's it's interesting to see these dynamics play out, particularly with Shivon. Shivon is such a sweetheart. She For people who don't know, she actually joined OpenAI in 2016, which I think was when it first got started, as a general advisor, later moved onto the board of directors, then went over to Neuralink. She was at Tesla for a small stint of time. She is the mother to a few of Elon's children.

There's this whole lore and background. She actually left the OpenAI board of directors when Elon began xAI because of conflict of interest. So, she has become this kind of key feature in this early part and it's been interesting to see it play out. We're going to continue to follow along. Next week we'll hopefully have some updates from the OpenAI side and then the following week, I believe it's we're getting close to a verdict. So, we're going to start to see things move quickly over the next few weeks as it relates to this trial. Shifting gears slightly. Um Anthropic, aside from the partnership that they signed, uh has had a pretty promising week of developments, product development specifically, uh as is normal for Anthropic style. Um finance, they released a big finance update where basically they are creating a bunch of ready-to-run agent templates for the most time-consuming work in financial services. Now, if any of you uh any of you listeners have worked or are working in the financial industry, you'll know that there's a bunch of admin work. There's a bunch of memos you need to review, a bunch of financial modeling that you need to do, a bunch of paperwork and administration that just takes up a bunch of time and in an ideal sense could be automated but still requires a lot of human input and judgment. Well, Claude, specifically Claude co-work and Claude code, has now become smart enough to be able to assess and do this work for you. And they've released a bunch of templates, specifically AI agent workflows, that can do this for you. Now, it sounds kind of unsexy, right? Claude code and Claude specifically has been able to run Microsoft Excel, PowerPoint, and Word pretty well now. In fact, they've integrated directly into Microsoft suite for all of those software programs, but there hasn't been a combined package.

Usually when you get access to these things, you can be like, "Okay, this is cool, but I don't know what to prompt it. What do I say to this? How do I connect it to my database in my company? How do I make it do the thing that I want it to do?" These templates basically solve that and it is, once again, Anthropic's step towards taking over many different sectors that are adjacent to what they do. First, they started off with code. Now, they're focusing on enterprise specifically financial services. Tomorrow could be law. It's just a continuing trend of Anthropic doing this particular pattern taking over sectors.

If you're in college right now and you're in finance and your your goal is to make it to Wall Street, what are the kind of jobs and the skills that you're optimizing for for the entry-level job, right? You're you're building pitches, you're preparing for meetings, you're reviewing earnings, you're building out models, you're doing market research. You're kind of doing all of these basic tasks. Every single one of these has been turned into one of the financial templates. And I'm actually reading from the list of things that Claude has automated because it sounds like a job description. And I think this is one of the really powerful things that when people are going to see they they're going to say, "Oh my god, it's taking all the entry-level jobs." But the reality is is like now that child who is in school, who is in college trying to be working Wall Street, be in finance, has these tools at his disposal to go and accelerate progress to do the next level of things. What are the things that come after you do these market research reviews or the valuation reviews or whatever it is? That's where I think people should spend a lot of their time. And instead of viewing this as this detriment to the entry-level class of Wall Street, it's like think about how more powerful you can be. You could just use all of these. You have your own employees now. So so think bigger, build better. And I think this Anthropic update empowers people to do that. And it's it's cool. I I want to test it out. I want to apply this to our episode that we talked about yesterday of the AI stack and have it kind of build out models and valuations on what I should look for when I'm investing in a >> model. We should get it to do Yeah, for the stock. Dude, it's like we we don't need to be on Wall Street or pay hundreds of thousands of dollars to hire this entry-level job. We just have the templates now on Claude. And you just ask it questions, you use their plugins, and it will generate like pretty proficient research. So I think it's it's very empowering.

>> I wonder if we could hook that up to like a a like a temporary bank account that we can spin up that can then just like invest in a bunch of stocks for us. That that might be an episode in itself. That would be a fun I mean again, your imagination is the only the only fixed variable here. It's like the only limitation. It could do basically anything you want. I like that. Well, in order to pull all of this off, right?

Let's assume financial services goes crazy and and AI is applied to pretty much everything. There's a trend. Tokens are being consumed, Josh. We need more compute. And the one company that has all of the demand, but not enough compute, is, like we mentioned earlier, Anthropic. But it's had an inverse side effect, which is Google is now the most valuable company in the world, more valuable than Nvidia, all right? So we have a new king of the of the throne, I guess. Um Anthropic signed a $200 billion deal with Google to get access to that compute via Google Cloud and through their TPUs. Now, it's important to point out that this is spread over 5 years, but it is a locked-in deal. In fact, it makes up over 50 to actually 60% of Google's revenue backlog for the next couple of years. Now, what's important to point out here is Anthropic has already been a big fan and customer of Google. And it works both ways.

Number one, Anthropic uses a bunch of Google's TPUs, which is their own version of GPUs that competes directly with Nvidia, to train and serve Claude. And it's actually helped Claude improve massively. Mythos was trained on Google's TPUs. Now, the other side is Google is one of the biggest stakeholders and investors in Anthropic themselves. In fact, I think they own roughly right now, after this deal and after a few other deals that happened over the past month, around 15 to 17% of Anthropic. Now, it's Google and Amazon combined stake, I think it's like 43% in total of Anthropic. So it's just insane.

Anthropic is selling a lot of their stake to these competitors or infrastructure providers, but there's now a symbiotic relationship where Google is more inclined to serve up compute to Anthropic. The one thing I couldn't get my head around, Josh, I guess is similar to the SpaceX AI and Anthropic partnership is if you're working at Google on the Gemini team, if you're Demis Hassabis, right? And you know you need compute to make Gemini a better model, you must be thinking, "Could you please just give us that compute so we can train the number one frontier model instead of Anthropic and Claude?" It seems like Google's just kind of like giving them over the reins at this point. Well, what's the problem?

It's like it is is Gemini having an issue? Are they hitting a wall where they haven't quite figured out where to spend that compute on? Because it's clear Gemini doesn't have nearly as much compute constraints in terms of serving up inference to users that Anthropic does. So, I wonder if this is a hedge against DeepMind of the Google business. Like in terms of if you were the CEO of Google and you have all your chips on DeepMind, are you going to continue to bet 100% of your TPUs that DeepMind is going to be the one to build this incredible AI or are you going to start hedging against that by offering your compute to others? Now, there's a few reasons why they would offer the compute to others as that hedge, but also just because they're having the xAI problem where they have good models, but they haven't quite matched the users with that compute and they they don't really need all those TPUs. So, we could speculate for many reasons, but I think the idea is that capitalism rocks, man. All of the arbitrage is going down across all of these companies. If you have any dark GPUs, they get snapped up immediately and people are willing to pay a very large market price in order to do so.

So, the market is efficient. TPUs are being bought and sold all over the place, same with GPUs, and we're just going to continue generating these tokens at whatever the cost takes. Oh my god, you just reminded me of something. Did you see that crazy stock announcement this week, Josh, where they're putting Nvidia GPUs and attaching it to your home, like to the outside wall, so that your home becomes its own silo data center, which then gets interconnected with a bunch of other homes in your neighborhood, effectively like turning a home neighborhood into a data center. Did you see this?

Uh no, but that sounds a little A freaky. It was the craziest suggestion ever and like Nvidia's investing into this startup just to point that out. Um but someone pointed out which was like, if I see this outside of your home, I'm probably going to rob you, right? It's the biggest advertising to just to say like, "Hey, I have all this money. I can afford these GPUs. I come" Anyway, like to to to talk about your capitalist comment, I was like, "Well, we might be going too far at this point."

>> I'll do you one better quickly, which is uh this is a plan loosely for the Tesla network of autonomous vehicles.

As they roll out AI 5 AI 6 chips in these cars, they're going to essentially have supercomputers that are connected to the internet and what better use of idle driving time than to turn them into a mega cluster. So, the idea is out there. I'm sure someone is going to capitalize it and I'm very excited for whatever the first person who does that and what it looks like cuz man, that'd be pretty cool. I mean, you run some GPUs in your house, you you get paid for it, you maybe get free electricity, whatever it is. I don't know. And then another great business model, which seems like it's not going to work, is ElevenLabs who somehow manages to continue to surprise and delight everyone in the market.

ElevenLabs, for those who aren't familiar, you could think of them as the audio AI company. They are the flagship AI company when it comes to anything that enters your ears, whether it be voices, whether it be music, anything that you hear, ElevenLabs has specialized in it. And now, they're not the only ones. Google has an audio model that makes music really well. There's a few other companies that are trying, but for some reason ElevenLabs is doing it the best. And that's thanks to their CEO, Mati, who had an announcement of $500 million ARR this year. This is a huge announcement from the team. Yeah, super impressive. Like one thing that ElevenLabs has done well is they've managed to maintain and build a really cool attractive image on the outside.

So, many celebrities are partnered or invested in ElevenLabs, the main one being Matthew McConaughey. We've got Jamie Fox joining us, Ronda Eva Longoria. So, that helps with your distribution mode of making ElevenLabs cool and appealing to the people that they're eventually going to be replicating or replacing, right? Like there's a whole fear in Hollywood that it's like, you know, "Uh AI's going to replace us. It's not authentic." But, if you can kind of license your IP, your voice, your character in a way that you can control, such as with Eleven Labs, you can end up, you know, making a bunch of money.

And 500 million AR kind of like plays into that. So, a massively successful model, but there is another story that I saw today. Well, rather not a story, Josh, but a clip. Have you ever seen a robot, right, crack an egg before with such precision? I'll do you one better. Have you ever seen a robot carefully take a pipette and pipette the right amount of stuff into a petri dish? Or have you ever seen it solve a Rubik's Cube? Or have you ever seen it What is it doing here? Is this like blending and making a shake? Anyway, the point I'm saying is this new startup announcement from Genesis AI has created a new generalized robotics model that can be applied to any kind of robotic machinery, but instantly give it a very precise amount of movement and precision in the way that it interacts with different objects to a point that we have never seen before. I think robotics is finally reaching escape velocity.

We've seen a bunch of announcements from the likes of Figure and Unitree, which, you know, the robots doing cool things, but it's mainly at a manufacturing process level, like they work in factories. This is the first time where I'm like, "Oh, the robot looks pretty human, and I I might have it in my home." It's just pretty cool. Yeah, okay. So, we're going to speed around the rest of these cuz there's a few interesting news topics that are left. One is a new round that is led by Peter Thiel for a company named Pantalassa. I think I'm saying that right. Um a $140 million round led by Peter Thiel, and their entire basis of this company is to build planetary-scale energy in the middle of the ocean. So, basically, they have these You could think of a computer cluster, a data center that we have kind of like in Texas or that we have in Tennessee. What they're doing is they're taking that data center, they're making it float, and then they're driving it off into the middle of the ocean to function independently. So, you can unlock the ocean as a planetary-scale resource, and this is pretty mind-blowing when we think about how AI is going to scale. I mean, we we're very familiar now with the AI data centers in space conversation, but the AI in the ocean seems pretty noteworthy, pretty interesting. And finally, we have some, I guess, AI-adjacent news, but not really. We just we we were just blown away by this and we we had to call it out, guys. The next coronavirus, apparently, is around the corner and it's appropriately named hantavirus.

There's a luxury cruise ship right now where it has an outbreak, very of coronavirus. I'm getting PTSD. And nearly 150 people are still on board with no symptoms, but there are some seriously ill passengers. Josh, what's going on here? Dude, I woke up this morning, this was my entire feed online. I'm like, what is going on? Is the world ending? And when we when we recorded our episode yesterday on the the threats to the AI bubble, this was not baked into the equation. So, we're just sharing this just to bake it in. This is highly speculative. According to Polymarket, there's only a 9% chance this turns into a pandemic, but it is like it's funny and not funny, but like it is scary to see the kind of rerun of this traumatic event that happened, but in a different flavor, in this different context. And I just wanted to just highlight as another thing to, you know, keep your eyes on.

Everything's fine for now, but people in at least three US states are reportedly being monitored for possible exposure and like a series of people have died and things are happening. So, just, you know, want to make sure everyone's aware, everyone's okay. Things are going to be fine. According to Polymarket, the odds are low. They're still in our favor. The up only pump is going to continue, but this should be hedged into your models at least. Add this to your L line. I'm telling them, hey, create this as a parameter when you're evaluating risk going forward. Yes, yes. Well, there you have it, folks. You are officially caught up by the end of the week of everything that's been happening in AI, everything from the drama in the courtroom to the latest partnerships between the rivals that are trying to kill the other rival, all the way to breaking virus news. You can get everything here at Limitless. We read a bunch of episodes this week, which uh jumped into the depths of why all the semiconductor stocks are pumping massively. We covered OpenAI's new phone that's releasing as early as 2027. And we compared the mother load of all AI models, GPT 5.5 versus Claude Opus 4.7.

Which one's the best? We did a breakdown. We did a bunch of demos. We may or may not have recreated Mario. Definitely go check out all those episodes. Um if you are listening to us on whatever platform you are on, YouTube, Spotify, or Apple, please leave us a review. Give us a rating. It helps us out massively. Josh, any final parting thoughts? Thank you so much for watching. If you made it this far, you're a real one. I very much appreciate all the support of getting caught up. Go enjoy your weekend. Have a great time touching grass cuz you're all caught up. Thank you so much for watching and we will see you guys next week. See you guys.

>> [music]