


The Daily: Problems with AI, part 2—Job security, regulation, and the likelihood it will cause human extinction

Audio



On today's episode, we discuss whether most of what AI writes is useless, if AI is coming for your job, the likelihood of human extinction, and AI rules that could help us avoid that scenario. "In Other News," we talk about what Nvidia is and why it just joined the hyper-exclusive trillion-dollar club and what happens now that Neuralink can test its brain implants on humans. Tune in to the discussion with our analysts Jacob Bourne and Gadjó Sevilla.

Subscribe to the "Behind the Numbers" podcast on [Apple Podcasts](#), [Spotify](#), [Pandora](#), [Stitcher](#), Podbean, or wherever you listen to podcasts. [Follow us on Instagram](#).



Episode Transcript:

Marcus Johnson:

Insider Intelligence subscribers can now take advantage of the exclusive opportunity to meet directly with our team of thought leaders in the financial services space. Through the new

Analyst Access Program, expect to leave equipped with what you need to drive profitable sales, legitimize strategic ideas, and develop models of addressable markets and opportunities. Visit insiderintelligence.com/analystaccess to learn more.

Jacob Bourne:

By definition, a super intelligence wouldn't be something that we could predict what it's going to do next or even control what it's going to do next. And by the same notion, we also can't say for sure that AGI will result in human extinction.

Marcus Johnson:

Hey, gang, it's Tuesday, June 6. Gadjo, Jacob, and the listeners, welcome to the Behind the Numbers Daily. It's eMarketer podcast. I'm Marcus. Today, I'm joined by same two folks I was joined by yesterday. Let's meet them again. We start with one of our analysts on the connectivity and tech briefing based out of California. It's Jacob Bourne.

Jacob Bourne:

Hello, Marcus. Hey, Gadjo. Thanks for having me.

Marcus Johnson:

Of course. Of course. We're also joined by someone else on the connectivity and tech briefing, senior analyst based out of New York, the other coast, it's Gadjo Sevilla.

Gadjo Sevilla:

Hey, Marcus. Hey, Jacob. Happy to be here.

Marcus Johnson:

Hello, hello. Gents, today's fact, I'll throw it as a question. What is the name of the highest waterfall in the world? It's not Niagara, hint.

Jacob Bourne:

It's in South America I know that much.

Marcus Johnson:

Yes, very nice. It is. Angel Falls in Venezuela. So at 3000 feet high, the waterfall drops over the edge of the Auyán-tepui mountain in the Canaima National Park. So for folks who have been to Niagara... Jacob, Gadjo, have you been to Niagara Falls?

Gadjo Sevilla:

I have, yes.

Jacob Bourne:

Yes.

Marcus Johnson:

Okay. Oh man, I haven't, unbelievable. Thanks for the invitation. So when you guys went to Niagara Falls, it's pretty astonishing, right? I feel like people say it is quite remarkable to see.

Gadjo Sevilla:

It's stunning. Yeah.

Marcus Johnson:

Okay.

Gadjo Sevilla:

Especially from the Canadian side.

Marcus Johnson:

That's what everyone says.

Gadjo Sevilla:

Yeah. Yep.

Marcus Johnson:

The world's just better from the Canadian side. Jokes about America. Kidding America, I love you. I live here. So if you've been to Niagara Falls and think that's amazing, Angel Falls is 17 times higher, which is just impossible to comprehend. I feel like I want to work remotely from there, although I speak chaos loud. Hi, welcome to the show. Victoria's going to hate that, but

I'm moving anyway. Today's real topic, Problems with AI, part two, Taking your Job, Regulation, and Human Extinction.

In today's episode, first in the lead, we'll cover some of the more profound concerns around AI. Then for In Other News, we'll discuss who Nvidia are and why they joined the trillion-dollar club and what happens now Neuralink can test its chip implants on humans. So the lead, well, yesterday we covered some of the less sinister, but still pressing concerns around AI, copyright, putting labels on AI, suing AI companies. Today we're talking about some of the apprehensions around AI that have a stronger gravitational pull. And so misinformation and bias are two such concerns. Oxford's Institute for Ethics in AI, senior research associate Elizabeth Renieris said she worried more about risks closer to the present like AI magnifying the scale of automated decision making that is biased, discriminatory, exclusionary, or otherwise unfair. And AI driving the exponential increase in the volume and spread of misinformation, thereby fracturing reality and eroding the public trust.

And we want to mention these because they're very significant concerns when it comes to AI, but we've covered misinformation to a certain extent on this show and how bad this could be if you mix generative AI's ability to produce things quickly with misinformation. But some other major concerns that we're going to be talking about today, one of them includes the question of words and what ChatGPT does to words. Is most of what ChatGPT writes Useless? So a founder of TeachingStartup.com and Chief product Officer at an on-demand vehicle care company, Joe Procopio just wrote in Inc that, "The real immediate problem with ChatGPT in generative AI is that words mean things unless you are using ChatGPT. Then words become useless." Mr. Procopio thinks the broadest issue with generative AI is not that humans won't be able to recognize it as machine written and accept it as the real thing. The problem is that humans will be able to recognize it as machine written and still accept it as the real thing. Gents, what did you think of this article?

Gadjo Sevilla:

Yeah, I would not say what generative AI generates is useless since it really depends on the quality and scope of the prompts it's given. So really the best way to optimize something like ChatGPT is to refine the prompts, ask a variety of questions, add qualifiers, and ask the AI to ask you questions to make results more accurate. Now, if you're just simply giving it one prompt to, say, write a script for a TV show, it's going to do that, but it's not going to be very

good. And that's where the problem lies, when people just sort of automate for automation's sake and expect AI to produce something that's cohesive and coherent.

Marcus Johnson:

Right.

Jacob Bourne:

Well, I was going to say... I mean, Gado's a hundred percent right on that. It comes down to the prompt engineering question. It's so important to getting the best output, and I think the best way to think about chatbots like ChatGPT is as assistants or brainstorming partners. They're not professionals that can do your job for you. And I think we are seeing a lot of inappropriate use cases that do make them seem useless, like the lawyer that just used ChatGPT in a court filing and ChatGPT ended up citing six fake cases. So that's an example of an inappropriate use case. I think the problem comes down to people thinking that machines are infallible, and so ChatGPT is infallible too.

The difference here is that ChatGPT is not like a calculator where you put in two times two and always get four. It's a large language model that's more like the human brain and how it functions. It's a neural network that like the human brain can make mistakes and ChatGPT has very broad knowledge, but it doesn't reach expert level knowledge on any area. And so if you want expert level output, you as a human user, you're the one who needs to push the quality forward. So it's really about using it appropriately.

Marcus Johnson:

And using it as an associate. You guys said you're supposed to be using this thing to help you with things, not to do the job entirely, but some folks have been concerned that AI, generative AI, but AI more broadly will just take a lot of people's jobs. But how will AI come for your job? Charlie Warzel of the Atlantic thinks instead of being replaced by robots, office workers in particular will soon be pressured to act more like robots. He thinks we've seen this movie before. Some technology promises to increase productivity by chipping away at inefficiencies in our lives to give us more time, but that time is then reinvested into more labor.

Mr. Warzel explains that Frederick Winslow Taylor ruthlessly optimized the factory floor at Bethlehem Steel by surveilling workers, cutting breaks, and streamlining their motions. The principles of Taylorism changed business and management forever, but its gains weren't to

the benefit of the worker who was simply driven to produce more each shift. Can you see that happening with AI, with generative AI, that it doesn't give us more time back at all, we just end up having to work more because of the efficiencies it's able to create?

Gadjo Sevilla:

Yeah, the problem there is the time that you save using generative AI could not be offset by the time you need to sort of fix what generative AI produces. If you want a good result-

Marcus Johnson:

Good point.

Gadjo Sevilla:

You need to have oversight and sometimes it doesn't balance out. You're better off doing the job as an expert, as a person who's been doing the job, whatever that job is, and rather than correcting AI's missteps or falsehoods.

Marcus Johnson:

Same as giving a job to somebody else. A lot of the time we think to ourselves, "I'll just do it myself."

Gadjo Sevilla:

Right.

Marcus Johnson:

Basically my parents, anything they asked me to do was "I'll just do it myself because I can't be bothered to correct all of your mistakes and end up taking more time to do the job." Yeah, I mean, Mr. Warzel explains that other workplace tools haven't made us more efficient or freed up our time either. Email, he says, didn't dismantle the culture of interoffice memos or workplace correspondence. Slack, supposed to be the corporate email killer, hasn't unclogged our inboxes. Instead, it's just added another channel for workers to check. And now when you say, "Where did you send that thing? Is it Slack, is it email, is it Zoom chat?" So these things don't automatically free up all of this extra time.

Jacob Bourne:

While that's true, we have seen some instances where companies are already citing AI as a reason for layoffs. So while it might not be a good idea, unfortunately we're seeing some examples of them doing just what we're calling putting too much faith in the AI and using it inappropriately, I think is unfortunately a thing already.

Marcus Johnson:

So getting to the more serious concerns here, we've talked about misinformation quickly and bias as well. We've talked about words meaning less. We've talked about how AI might come from your job, but as Kevin Roose of the New York Times outlines, "Leaders from AI companies, OpenAI who made ChatGPT, Google's Deep Minds, Anthropic, and other AI labs just said that the AI tech they're building might one day pose an existential threat to humanity and should be considered a societal risk on par with pandemics and nuclear wars." Great. Mr. Roose notes that, "Some believe AI could become powerful enough that it could create societal scale disruptions within a few years if nothing is done to slow it down." Tremendous. Jacob, what's the likelihood that AI leads to human extinction? There's a lot of pressure writing on this question and people are listening to it very intentionally-

Jacob Bourne:

It's a deep question.

Marcus Johnson:

So don't scare everyone.

Jacob Bourne:

It's an intense question. So before I answer, let me give a few data points. So an NYU survey at AI Experts showed that 36% think that AI could cause catastrophic outcomes like an all out nuclear war. Then there's a former OpenAI researcher, Paul Christiano, who thinks that there's a 10 to 20% chance that an AI takeover will eliminate many or most humans, and that there's a 50% chance that there will be a catastrophe shortly after we achieve AGI. Then there's another researcher, Eliezer Yudkowsky, who thinks that human extinction is the obvious thing that will happen once we reach AGI.

So I think the essential thing to take away from all of this is that by definition, a super intelligence wouldn't be something that we could predict what it's going to do next or even control what it's going to do next. And by that same notion, we also can't say for sure that

AGI will result in human extinction because if we could say that, then it wouldn't be a super intelligence. And so I think the big takeaway here is that building AGI is very risky and it's probably a bigger risk than a lot of people want to take, and we might not even need it. Instead of building these advanced general models, instead we could build small focus models that are also very powerful, but don't have that kind of general broad intelligence that would be needed to, say, take over the world.

Marcus Johnson:

On kind of mass scale.

Jacob Bourne:

We could also put more investment in quantum computing that promises some of the same benefits of AGI without the same kind of catastrophic risk.

Marcus Johnson:

So when Jacob's talking about AGI, it's artificial general intelligence, and as we explained yesterday, that's when artificial intelligence basically can reason like a person. So generative AI, ChatGPT, we said sounds like a person, AGI reasons like a person. And then next you'd have sentient AI which thinks like a person. Yeah, Geoffrey Hinton and Yoshua Bengio, so the two of the three researchers who won the Turing award for their pioneering work on neural networks, they're often considered as the godfathers of the modern AI movement, they signed the open letter of concern that we're talking about. And then in terms of how soon, you mentioned a few folks saying that the likelihood... Mr. Hinton, one of the three godfathers of AI said AI programs are on track to outperform their creators sooner than anyone anticipated, saying, "I thought for a long time that we were like 30 to 50 years away from that. Now, I think we may be much closer, maybe only five years away from that." Great.

Let's move on to how we might avoid that catastrophe. Sam Altman, the CEO of OpenAI, just called for the US to regulate artificial intelligence at a recent Senate committee hearing. Why? Because of everything we've been talking about today pretty much, and in yesterday's episode as well, along with the dangers of election misinformation, impersonation of public and private figures, job disruption and economic displacement. So Gadjo, when will AI get regulated and what will it look like?

Gadjo Sevilla:

So I think more than when, I think we should think about where it'll be regulated. Different regions have their take on how it's going to be. In China, it's already regulated to some extent. Beijing has control over algorithms. They can sort of dictate the scale of influence that the technology has. But in other regions like the European Parliament, they have a draft of changes that would place restrictions, but mostly on the foundational models, which right now are all American companies, the ones you mentioned, OpenAI, Anthropic, Google. I think there will be a race to regulate. However, it's not going to happen overnight. It'll take quite a lot of back and forth and I think it needs to be a collective effort with companies that have called out the need to regulate as well as government regulators. So they say, "Yeah, we need to regulate AI," but they don't say how. I think that is the next step.

Marcus Johnson:

I mean, you said it's not going to happen overnight, it's going to take a long time. I mean if you look across the pond, so the European Union, so the EU, they're working on what it calls the AI Act, which you alluded to.

Gadjo Sevilla:

Right.

Marcus Johnson:

So Martin Coulter and Supantha Mukherjee of Reuters were explaining that lawmakers have proposed for the EU's AI Act, classifying different AI tools according to their perceived level of risk from low to unacceptable, so it'd be the highest level of risk. High risk AI would include critical infrastructure, law enforcement, education, things like that. And that level would face rigorous risk assessment. They'd have to log their activities, make data available to authorities to scrutinize. They said that those using AI systems which interact with humans are used for surveillance purposes or can be used to generate deep, fake content, would face strong transparency obligations as well.

And this bill would work in tandem with other laws like GDPR. If these companies are caught doing something wrong, Sam Schechner of the Wall Street Journal noting the current draft of the bill would impose fines of up to 6% of a company's global revenue or over \$30 million, whichever's higher in the case of noncompliance. For some companies, this could mean billions of dollars, though. Whether that would be a deterrent, who knows? But to your point about how quickly this could happen, after the terms are finalized for the EU, which is way

ahead of the US, there would be a grace period of around two years to allow affected parties to comply with regulations. So we're still a ways off. Jacob, what's the state of AI regulation in the AU?

Jacob Bourne:

Well, the US is far behind China and the EU on this. Congress recently had a hearing with Sam Altman and a representative of IBM. I mean they talked, they floated ideas, but so far we're not really seeing much beyond talk. And so we can assume that it's going to take quite a while for the US government to really take meaningful action on this. What we might see eventually is the creation of a federal agency to govern AI. And within that, what we might see happen is that they require that companies building certain types of advanced risky models get a license in order to do so.

Marcus Johnson:

Yeah.

Jacob Bourne:

Now passing the law is actually going to be the easy part. And as we see, it's not all that easy. The hard part is actually once they pass it, how do you enforce it? I mean, this is new territory for the government. And one of the things that has AI experts really concerned is this proliferation of open source models. And so this pretty much started earlier this year. Meta has had an open source model called LLaMA. That actually got leaked online. And so now you have a situation where people are just building upon and deploying these very powerful models in their basement, in their garages. It can run off of a regular consumer device. And so the question is how is the government going to actually effectively enforce the laws there?

Marcus Johnson:

Yeah. This agency, Mr. Altman of OpenAI was one of the folks who suggested that to the government. He said, "Hey, a new agency to license AI companies could be good." There's also the idea of a worldwide agency. So Ryan Heath of Axios is saying that the founders of OpenAI think the International Atomic Energy Agency, folks might have heard of the IAEA, which exists to ensure nuclear technology is used for peaceful purposes, is a good model for limiting AI that shows super intelligence. So we've got to think not just what's going on within countries, but at the global level as well.

The US, to your point, is way behind and still in kind of a think about what legislation looks like phase. However, the White House's website does say that it's taken some steps on regulating AI, not passing anything but some suggestions and recommendations, some bills. They include a blueprint for an AI bill of rights. That's a set of protections Americans should have from AI systems, including being told when AI is being used and being able to opt out. Again, these are recommendations, not laws. There's also the AI risk management framework is working on and a roadmap for standing up the national AI research resource as well. We can't even get Congress to agree on national privacy legislation. So good luck with an AI bill. Anyway, that's what we've got time for. The lead time for the halftime report. Gadjo, let's start with you. What's worth repeating from the first half?

Gadjo Sevilla:

What's worth repeating from the first half is the idea that AI could come and take your job. Just understanding that it's not a one-to-one thing, it's a very nuanced reality. Like AI is only as good as what you put into it. And so I think the faulty thinking there would be to just automate for automation's sake and allow AI to just become a person or developer, a programmer, a writer. It'll take a lot of work to ensure the results from those AI jobs are actually even close to what people are used to.

Marcus Johnson:

Jacob?

Jacob Bourne:

Yeah, I mean, I think it's stunning that we're starting to hear this discourse around this technology causing something like human extinction. At the same time, we see that the regulation to potentially prevent some type of catastrophic outcome is so far behind pace at which this is going and the pace at which this technology is advancing. And so that's the real elephant in the room, is getting that caught up in time to really fend off some of the worst outcomes.

Marcus Johnson:

Yeah, the irony of sitting in front of Congress and saying, "This thing I'm working on at the office could potentially kill everybody. Just I just wanted to let you guys know that. Thank you

for your time. I have to really go now because I've got to get back to the office to work on that thing I was talking about that could potentially kill everyone." That's brazen.

Gadjo Sevilla:

You can't blame them if they said, "I told you so," right?

Marcus Johnson:

I know. Exactly.

Jacob Bourne:

But there'll be no one around to blame then, right?

Marcus Johnson:

Talk about a disclaimer. Anyway, time for the second half of the show. Today In Other News, Nvidia joins the trillion-dollar Club and Neuralink wins FDA approval for the human study of brain implants. Story one, "Nvidia joins the trillion-dollar club," writes Dorothy Neufeld of Visual Capitalist. Five companies are currently in that club. Apple, Microsoft, Aramco, Alphabet, and Amazon have all reached market caps of over a trillion dollars and are currently in that club. Two others, Meta and Tesla, used to be in that club back in 2021, but have since been kicked out because their market caps have fallen below that trillion-dollar watermark. "Chipmaker Nvidia recently became the sixth member of the club, thanks to strong earnings and hype around AI," notes Miss Neufeld. The company is now worth nearly as much as Amazon, but Gadjo, the most interesting thing about Nvidia is what to you and why?

Gadjo Sevilla:

Okay, so first of all, Nvidia is usually associated with gaming. They make GPUs. To some extent, they've done chips for robotics and they're really known for graphics chips. But more surprisingly, they've been able to use those chips to sort of be the backbone of the AI revolution that we're seeing. What's surprising to me is they were able to add a \$184 billion to their market value in one day and that just clearly shows how that shift is happening. And we talk about AI as a cloud service mostly, but you need hard performance for that. And these GPUs and CPUs that Nvidia has really been refining for the past decade, they come at a perfect time.

Marcus Johnson:

Yeah, that's stunning, adding \$184 billion to its market value in a single day.

Gadjo Sevilla:

Yeah.

Marcus Johnson:

Only two other companies, Amazon and Apple have added more in one day. Story two, "Elon Musk's Neuralink wins FDA approval for the human study of brain implants," writes Rachel Levy, Marisa Taylor, and Akriti Sharma of Reuters. Mr. Musk said he envisions brain implants curing a range of conditions including obesity, autism, depression, and schizophrenia, as well as enabling web browsing and telepathy. Neuralink, however, is facing federal scrutiny following Reuters' reports about the company's handling of animal experiments. But Jacob, the most interesting sentence in this article about Neuralink and being able to now test their brain implants on humans is what and why.

Jacob Bourne:

Yeah. The most fascinating thing here is that the FDA acknowledged approving the clinical trials, but then declined to provide more details about it. And then for Neuralink, they declined to respond to Reuters' requests for comment. And in light of the fact that Neuralink is under two federal investigations, the one for the animal cruelty and then another one for allegedly unsafe transport of infectious disease contaminated equipment, it's not a good look given that we're in this post-pandemic era, of course, and so they're trying to sell something that's hoping to put an implant into people's brain. And the thing is they're not the only company doing this. And so if other companies, rivals want to say, "Okay, we're going to be more transparent," well, then that's not going to bode well for Neuralink.

Marcus Johnson:

That's all what we've got time for today's episode, folks. Thank you so much to my guests. Thank you to Jacob.

Jacob Bourne:

Thanks, Marcus. Thanks, Gadjo.

Marcus Johnson:

Thank you to Gadjjo.

Jacob Bourne:

Marcos, Jacob, always a pleasure.

Marcus Johnson:

Yes indeed. Thank you for putting up with me for back-to-back days. Thank you to Victoria who's put up with me for three back-to-back years. Thank you for that. She edits the show of course. Thank you to James who copy edits. It's Stuart who runs the team and thanks to everyone listening to the Behind the Numbers Daily, an eMarketer podcast. You can tune in tomorrow for the Re-Imagining Retail show with host Sarah Libo as she chats with senior analysts, Carina Perkins and Zach Stambor all about Ikea.