Jacob Bourne, and Gadjo Sevilla

The Daily: AI rules to focus on first, Minecraft's milestone, and what AI chatbots can glean from a conversation

Audio

On today's podcast episode, we discuss what AI rules the government should focus on first, what to make of AI "nutrition labels," and what concerns us most about the dark side of AI. "In Other News," we talk about a Minecraft milestone and what AI chatbots can tell about you from a conversation. Tune in to the discussion with our analysts Jacob Bourne and Gadjo 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

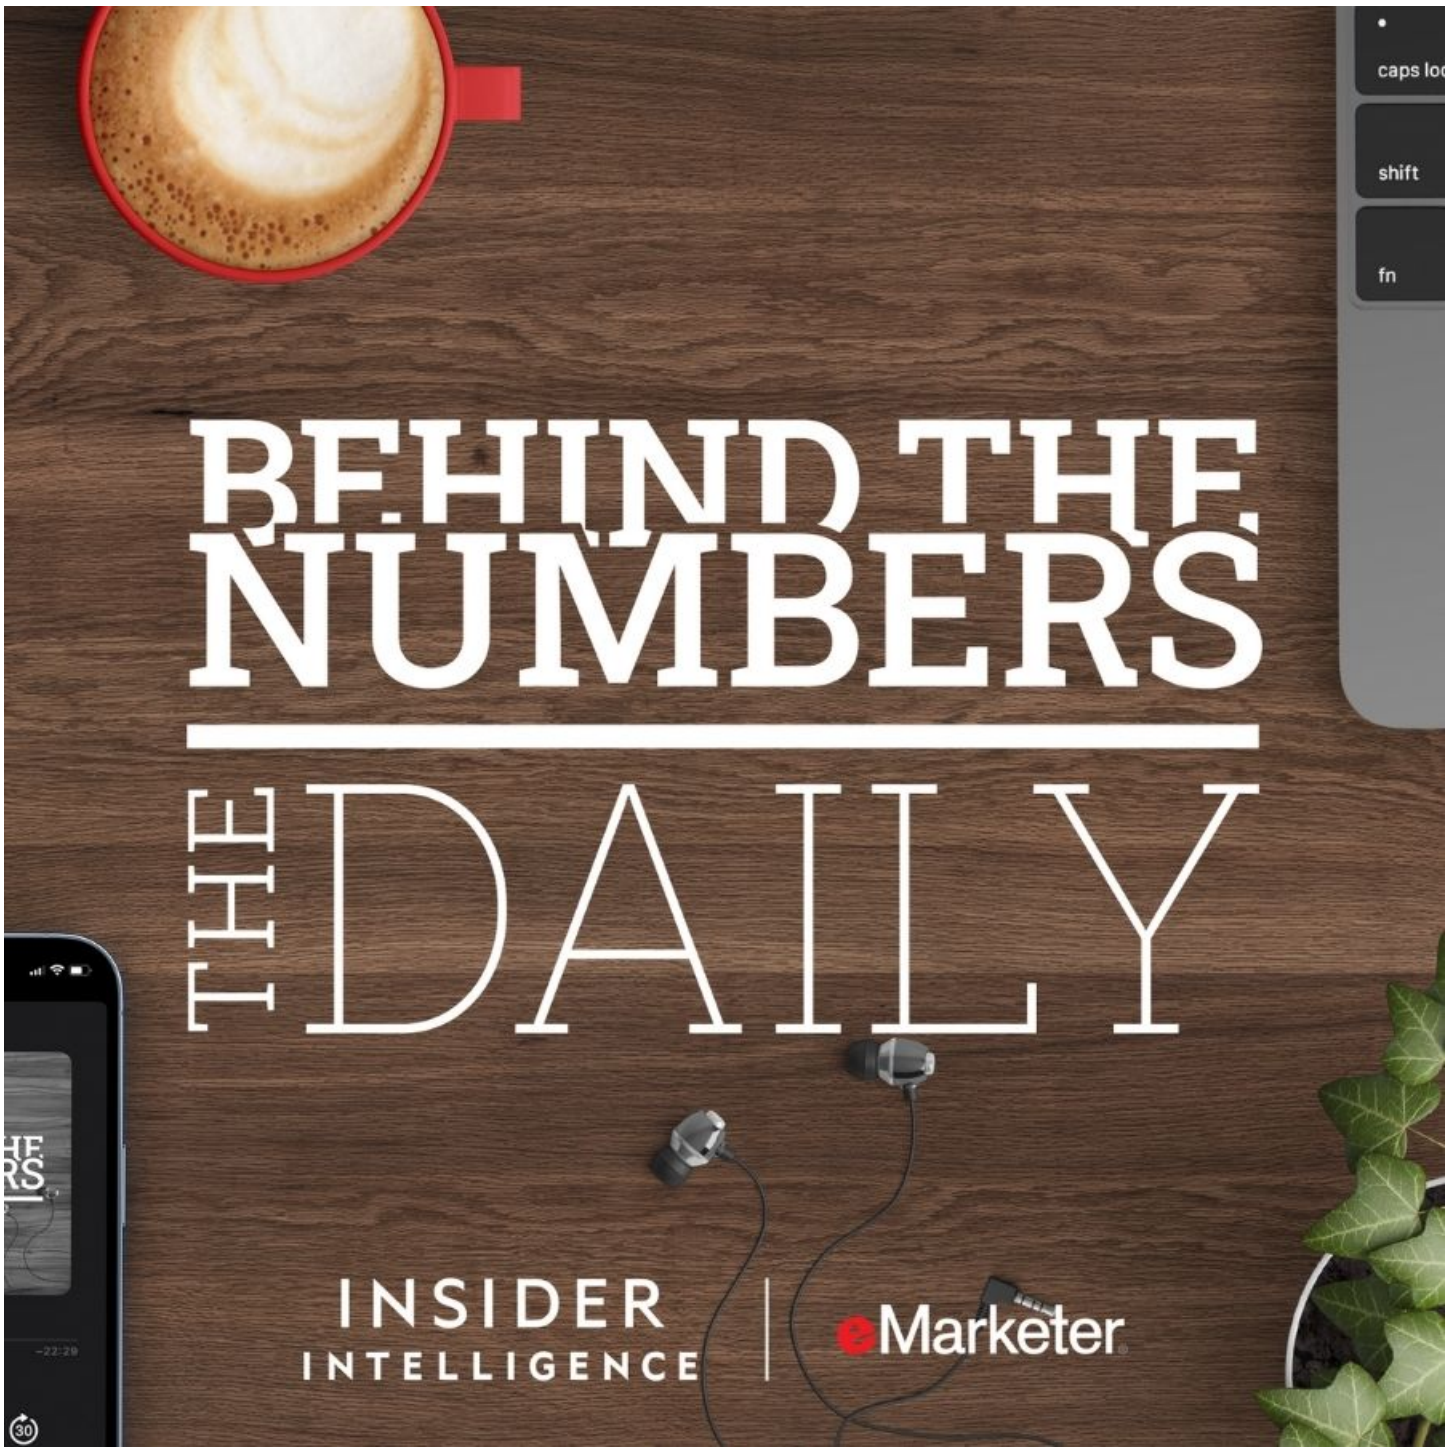

Episode Transcript:

Marcus Johnson:

This episode is made possible by Mailchimp. Ever heard of a clustomer? It's the result of marketers grouping customers with different behaviors into one big mess. But with

Mailchimp, you can use real-time behavior data to personalize emails for every customer based on their browsing and buying behavior, turning your clustomers into customers. Intuit Mailchimp, the number one email marketing automations brand. Visit mailchimp.com/personalize for more information. Based on competitor brands' publicly available data on worldwide numbers of customers in 2021, 2022. Availability of features and functionality vary by plan, which are subject to change.

Gadjo Sevilla:

Within that umbrella of regulation, having a safe sandbox to have at least some aspects of the AI be open source, I think that could yield some very interesting results while still being under the purview of regulation.

Marcus Johnson:

Hey, gang. It's Monday, October 23rd. Jacob, Gadjo, and listeners, welcome to the Behind the Numbers Daily, an eMarketer podcast made possible by Intuit Mailchimp. I'm Marcus. Today, I'm joined by two folks. Both write for our connectivity and tech briefing. One of them is on the West coast. He's based in California. One of our analysts on that crew, it's Jacob Bourne.

Jacob Bourne:

Hey, Marcus. Thanks for having me today.

Marcus Johnson:

Hey, fella. Of course, of course. The other chap who's part of that crew lives on the other side of the country on the East Coast. Based in New York, one of our senior analysts, it's Gadjo Sevilla.

Gadjo Sevilla:

Hey, Marcus. Happy to be here again.

Marcus Johnson:

Hello, sir. So gents, today's fact. Just how many people live in New York City? So it's eight and a half billion. Not billion. That's too many. Eight and a half million. It feels like it, but it's not. [inaudible 00:01:59]-

Jacob Bourne:

The whole world crammed into New York City, essentially, [inaudible 00:02:02].

Marcus Johnson:

If you've been to Times Square, that's what you're thinking.

Jacob Bourne:

Right, right.

Marcus Johnson:

Million. Eight and a half million roughly. So there's more people living in New York City than the individual populations of 38 US states. So if New York City was a state, the five boroughs was its own state, it would be, I believe, 13th on the list.

Jacob Bourne:

Wow.

Marcus Johnson:

So 38 US states which have a smaller population than New York City. The state obviously has another 11 million, about 20 in total. Put another way, Texas, it's about 30 million people, Texas has a population that's four times greater than New York City, despite having a landmass that's 650 times as large.

Jacob Bourne:

Wow. That's some density there.

Marcus Johnson:

There's too many people here.

Jacob Bourne:

For sure. Yeah.

Gadjo Sevilla:

Everything's bigger in Texas, right?

Marcus Johnson:

Yeah. Too true. It's painfully true. Eight and a half billion people? It's too many. I said it wrong. Anyway, today's real topic, when will AI get regulated and what will it look like?

In today's episode, first in the lead we'll cover AI rules that are being looked into. And then for In Other News, we'll discuss Minecraft's milestone and what AI chatbots could learn about you from a simple conversation. We start with the lead, and we're talking about some rules we might expect to see. First, what an AI nutrition label even is, when we might see some rules, and then we'll end with the darker side of the AI. But let's start with what the government's most likely to focus on at the moment.

So Senator Richard Blumenthal, Democrat of Connecticut, and Senator Josh Hawley, Republican of Missouri, plan to announce a sweeping framework to regulate AI, writes Cecilia Kang of The New York Times. As she notes, their framework will include, and she outlines four pieces here, A, requirements for licensing and auditing of AI; B, the creation of an independent federal office to oversee AI; C, liability for companies for privacy and civil rights violations; and D, requirements for data transparency and safety standards. So we've got licensing and auditing, an independent AI regulator, privacy liability, and data transparency and safety standards. Jacob, I'll come to you first. What should the government focus on to start?

Jacob Bourne:

Well, my hope is actually that the government is capable of focusing on multiple things at once because that's really-

Marcus Johnson:

We'll see.

Jacob Bourne:

Yeah, and then that's really what's needed here. That said, an independent regulatory agency would help because it would entail hiring a lot of AI experts, and that's really what the government needs because there's a big gap right now. I mean, not a lot of people really understand this technology to the degree that is needed, and the government is one of those bodies where that's lacking. The other thing that I think that actually would be a really good place for them to start is actually AI chips, but not necessarily... I'm not talking about [inaudible 00:05:21]. I'm talking about domestic chip regulation.

So there's three things that are needed in AI model development. You need data, you need algorithms, and then you need advanced AI chips. Now, the thing is, controlling the data and algorithms is challenging because they're not tangible things. It's essentially code. The chips, however, are tangible. They're scarce, not a lot of companies can make them, and it's something you could actually control the sale of. In other words, the government could pass licensing requirements for advanced AI chips so that they know, first of all, who is using them, what it's being used for, and being able to track the use.

And I think that would be one way to really, really make sure that this doesn't get out of hands, make sure that AI doesn't get into the wrong hands essentially. The one caveat here is that they would have to implement in such a way that it doesn't hurt small businesses versus big tech companies using it. I think that's the only challenge around it.

Marcus Johnson:

Fascinating point. Okay, so a regulatory body specifically to review or to cover AI, and then also the regulation of chips themselves, particularly in the US. Okay.

Jacob Bourne:

Mm-hmm.

Marcus Johnson:

Gadjo, where should they start? Where should the government start?

Gadjo Sevilla:

Well, I think focusing on the aspects of AI that could do the most harm makes sense to regulate that first. So I could see them looking at data collection, privacy protection, and also AI's hallucination problem where it sometimes makes up falsehoods. I mean, those are still complicated, but I think if they tackle that first, or at least make inroads to curb those areas, then they would be off to a good start.

Marcus Johnson:

Okay. One thing that's being proposed, this is from a company, a nutrition label. So in an Axios article, Ina Fried noting that Twilio, which helps companies automate communications with their customers, said it would put nutrition labels, in quotation marks, on the AI services it offers to businesses, clearly outlining how their data will be used. So the labels would lay out

what AI models Twilio is using, whether those models are being trained on customer data, if features are optional, and if there is a human in the loop. The piece explains that there's also a privacy ladder which distinguishes between company data that is used only for customers' internal projects, data that's being used to train models by other customers, and whether personally identifiable information is included. So think of any food item, drink item that you've seen, there's a nutrition label. AI would have a version of that. In this scenario. Jacob, what'd you make of the idea of AI nutrition labels?

Jacob Bourne:

Yeah, I think they're almost necessary. I mean, it doesn't have to be the nutritional labels format, but something that lets the general public know, "Hey, you're engaging with a generative AI tool, model right now, and this is what it means in terms of your privacy, in terms of security risks, in terms of other potential risks," and there are a lot. So I think it's definitely something.

And the thing is that a lot of people don't really understand generative AI. And the more that we see this become ubiquitous in society, the more people might be interacting with a generative AI in some way and not know it and not know the risks. And so this really gives a standardized way to communicate with people what the risks are, that a model is safe, and also I think it creates some accountability too. So I think it's definitely a good idea.

Marcus Johnson:

Hmm. So we're talking about where the government should start first in terms of regulation and AI. Nutrition labels being one idea from a company. But going back to the government for a second, they spoke to some of the big players in AI quite recently, and there's a Wall Street Journal piece by Ryan Tracy and Deepa Seetharaman explaining that in September, top Silicon Valley bosses briefed lawmakers on AI. But there are many folks who hold competing views on the technology, and the piece was outlining six particular viewpoints on AI from some of the heavy hitters. So OpenAI CEO Sam Altman says, "Regulate us and create a new AI agency to do so," which is what Jacob, you said would be important-

Jacob Bourne:

Mm-hmm.

Marcus Johnson:

... to focus on first. Meta's CEO Mark Zuckerberg is championing an open source approach to AI. Tesla, X, and xAI leader Elon Musk is trying to warn folks of AI's existential risk. So maybe those problems, perhaps Gadjo, you were alluding to. A researcher, Inioluwa Deborah Raji, wants people to be aware of bias. And then two more. The president of Writer's Guild of America West, Meredith Stiehm, wants to look out for workers. She thinks that's the priority. And then Google's chief, Sundar Pichai, wants the government to let industry lead with voluntary efforts to address some of AI's potential harms. Gadjo, whose perspective did you find most interesting in this piece and why?

Gadjo Sevilla:

Well, actually, if I can I choose a mix of perspectives, because I do think it needs to be-

Marcus Johnson:

A cocktail, if you will.

Gadjo Sevilla:

A cocktail, yes. So I do think that the industry does need to be regulated, but I'm also seeing the benefit of having an open source AI platform the way Meta's trying to do it, just because looping in a lot of developers working on the same problems could result in possibly better answers sooner for when it comes to AI. So within that umbrella of regulation, having safe sandbox to have at least some aspects of the AI be open source, I think that could yield some very interesting results while still being under the purview of regulation.

Marcus Johnson:

Okay. Jacob?

Jacob Bourne:

No cocktail for me.

Marcus Johnson:

Oh.

Jacob Bourne:

Mostly because I think-

Marcus Johnson:

He wants his neat.

Jacob Bourne:

Yeah. All these figures, I think, have not sufficiently addressed all of the risks. But I will say that OpenAI's Sam Altman has a very solid strategy in his approach. I mean, he's basically becoming the AI regulation champion. He's going to the federal government saying, "Please, please regulate me and the industry," and I think it's very smart. I think he's at a position to influence regulation, and we saw a bit of that at the Senate testimony a couple months ago where he was supposed to be grilled by the senators, but actually they were trying to network with them. And so it really gives OpenAI a chance to affect what happens to the competitors in the industry. And ultimately, crafting sound regulation will help prevent fallout from the technology, which of course benefits everybody, including OpenAI.

Marcus Johnson:

Yeah. So question here is, when? When are we going to see something? Ian Prasad Philbrick of The New York Times writing that, quote, "The US regulates cars, radio, and TV. When will it regulate AI?" It was pointing out in the piece that Congress has tended to be slow to respond to revolutionary technologies, saying that for cars there were 70, seven zero, years between invention/patenting and the first major federal regulation being passed. So from invention to regulation, 70 years for cars, 60 for trains, 20 for planes, and even five for the TV. Jacob, when will we likely see some federal rules on AI?

Jacob Bourne:

Well, depending on where we're talking, I mean in the us-

Marcus Johnson:

Mm. Yeah, good point.

Jacob Bourne:

... not immediately. The EU, however, the AI Act is slated to pass by the end of the year, which is-

Marcus Johnson:

Wow.

Jacob Bourne:

It's going to be interesting because actually, a recent Stanford study showed that the current draft regulation and the draft legislation showed that most commercial models today don't comply.

Marcus Johnson:

Oh, great.

Jacob Bourne:

So there could be changes by the end of the year, but I think we should expect that by the end of the year, or at least whenever it goes into effect, that AI companies will have to make some changes in order to operate in the EU. Now in the US, I think what we're going to see, we're going to see some executive orders next. We're also going to see some state regulation. There are some bills at the state level, including California, to comprehensively address AI. And I think we're going to see that a lot sooner than we see legislation at the federal level.

Marcus Johnson:

Yeah. So let's finish the lead by talking about some troubling stories with regards to AI, some of the really important reasons why AI does need some kind of regulation, going back to Gadjo's point at the very top of the segment. So a few troubling stories around AI's dark side that have been in the press. One was around actor Tom Hanks warned that an advert that appears to be fronted by him was in fact an AI fake. Two, a deepfake on TikTok of the world's biggest YouTuber, MrBeast, showed him offering people new iPhones for $2. So that was pretend. And then the third story here is particularly troubling. Someone in the UK was just given a nine-year sentence for breaking into Windsor Castle with a crossbow in 2021 and declaring that he wanted to kill the queen because he claimed a chatbot had encouraged him to do so. So some of these are warnings for folks, whether they're advertisers, whether they're AI makers, that some of these do get pretty sinister and pretty dark quite quickly. Gadjo, I'll start with you. What do you make of these stories that show the dark side of AI?

Gadjo Sevilla:

I think it goes back to, again, regulation, right? When you have such a powerful tool which is easily accessible and not really regulated to the level that it could be deemed safe for use for most people, then you're going to run into problems like these, right?

Marcus Johnson:

Mm-hmm.

Jacob Bourne:

Yeah. I mean, this particular issue is called the ELIZA effect, basically this inclination that people have to interact with AI like it's human and potentially causing some psychological problems. This actually dates back to the 1960s, so people have known about this for quite some time. And of course, here we are today. And I think we really saw some of the dark side potential AI during Microsoft's Bing Chat testing phase. And with the Sydney bot, there was lots of transcripts. The New York Times reporter, other AI people in the field tested it and it went off the rails a number of times in disturbing ways. Microsoft since has put safeguards in it. Of course, we know ChatGPT has safeguards where it won't say certain things.

However, one researcher actually referred to these safeguards as something akin to putting a smiley face mask on a monster. So in other words, you're not seeing the problems, but they're still there. Nothing has really changed, and so you can actually break past the safeguards and get it to generate problematic output. And that can take all kinds of forms. And there's no foolproof method to deal with this, and I'm not entirely convinced that we get there anytime soon. So it is a very troubling challenge for the industry.

Marcus Johnson:

Yeah. All right, gents. Let's skip the halftime report and move straight to the second half of the show. Today in other news, Minecraft hits a new milestone, and what can chatbots figure out about you from a simple conversation?

Story one. Gadjo, you recently noted that Minecraft has sold over 300 million copies, making it one of the bestselling video games in history. You write that the game, which was first made public in 2009 and bought by Microsoft in 2015 for $2.5 billion, has 175 million monthly active users as of September, according to ActivePlayer. But to you, Gadjo, what's the most important takeaway from this 300 million copies sold milestone?

Gadjo Sevilla:

I think it exemplifies how Microsoft was able to take a popular game and was able to grow it into really valuable IP and a gaming phenomenon. So they gave the studio, Mojang, the latitude to expand as needed. I don't think they were involved much in the key development, so they let it grow organically. So I think it's a testament to Microsoft's ability to build on gaming foundations.

Can they replicate that with Activision Blizzard? That does remain to be seen because that's a much larger scale and it's more than one game and one studio. But for Minecraft, it really shows that they were able to pave the long-term success for that property.

Marcus Johnson:

Mm-hmm. Over 300 million copies sold. You point out some really interesting context. Tetris, the only game that has sold more copies. 520 million compared to this 300.

Gadjo Sevilla:

Yeah, and Tetris has been around for 30 years at least.

Marcus Johnson:

Yes. Yeah.

Gadjo Sevilla:

So it's had a headstart.

Marcus Johnson:

Yeah, just a bit. Yeah. Story two. Jacob, you just wrote that research shows chatbots like ChatGPT can infer sensitive user information from context clues in conversations. You explained that ETH Zurich research team has shown that chatbots can infer information like race, location, occupation, and more. And you say there's no clear way to fix the problem, which sounds troubling. Jacob, why to you is this so concerning?

Jacob Bourne:

Yeah, well, before I get into that, I also want to say this is evidence of how powerful this technology is that it can even do this. So if there's any... Of course, we focus on some of the negative aspects of generative AI, but it's a really amazing technology. But of course, talking with a chatbot and it having your information, or being able to infer information about you like

where you live is certainly a privacy concern. And I think that a lot of people would be uncomfortable with that, but the most concerning thing is what happens next.

So as models get more powerful, including gaining multimodal capabilities where they can analyze images or photographs, for example, or analyze your voice when you're speaking to it, it's not farfetched, I don't think, that it could eventually gain the capability of identifying individual people that it's talking to. And of course, that would be really troubling. I think a lot of people would be wary about what they share with it and limiting the commercial value, limiting the utility of it. Now, the researchers who identified this issue said that they see no way to circumvent it at this time, so that's a concern. It's a technical issue. But if you make it so that AI can't make inferences, then you're making it less powerful. So it's a bit of a tension.

Marcus Johnson:

Yeah, it's quite the balance. Yeah.

Jacob Bourne:

Mm-hmm.

Marcus Johnson:

So chatbots that could gain the ability to identify individuals by chatting with them, and you write hackers and scammers could deploy chatbots specifically to gain sensitive information as well.

Jacob Bourne:

Yes, yes.

Marcus Johnson:

Something else to look out for. Terrific. That's all we've got time for, for this episode. Thank you so much to my guests. Thank you to Jacob.

Jacob Bourne:

Thank you, Marcus. Thanks, Gadjo.

Marcus Johnson:

And thank you to Gadjo.

Gadjo Sevilla:

Marcus, Jacob, thanks again.

Marcus Johnson:

Thank you. And thank you to Victoria who edits the show, James who copy edits it, and Stuart who runs the team. Thanks to everyone for listening in. We hope to see you tomorrow for the Behind the Numbers Daily, an eMarketer podcast made possible by Intuit Mailchimp.