



FTC Tech Summit | January 25, 2024

Stephanie Nguyen:

Hello everyone, and welcome to the Federal Trade Commission's Tech Summit on AI. I'm Stephanie Nguyen, Chief Technologist of the Office of Technology. A year ago, Chair Khan and the commissioners voted to establish the Office of Technology to ensure that the FTC remains nimble and keeps pace with evolving markets.

Thanks to staff and the leadership across the agency, we built on solid foundations and hit the ground running to execute our mandate against a quickly-evolving tech landscape. And since then, we've hired some of the most talented technologists in the country, committed to amplifying the benefits of tech, curbing its harms, and enforcing the law.

This moment comes at a critical time. We're stacked against deep-pocketed corporate incumbents who

unleashing the full potential of emerging technologies? Or will a handful of dominant firms concentrate control over these key tools, locking us into a future of their choosing?

Which of these potential trajectories AI will take is not an inevitability. The outcome will be a direct result of policy choices that we make now.

Zooming out for a moment, virtually every large firm going back to US Steel in the early 1900s, to Alcoa in the 1930s, to IBM and AT&T in the 1970s, to Boeing in the 1990s, and to dominant technology platforms today, have argued that their market power is good for America, and that government officials have offered and agreed.

In the 1990s, officials even reportedly threatened the Europeans with sanctions if they would not allow the merger of Boeing and McDonnell Douglas. As one White House adviser put it, "Aerospace was the only sector where we have a de facto national champion, and you can be out and out advocate for it."

And yet when you concentrate production, as Boeing for example has done, you also concentrate risk. And so today we see firsthand some of the real implications of that.

To quote United Airlines CEO Scott Kirby, "That 1997 mega-deal is what led directly to the transformation of Boeing from a highly profitable, world-class engineering enterprise, to an ossified money-losing corporation with dangerous quality issues that we're now seeing firsthand."

Boeing's journey unfortunately is no different from that of many large corporations that policymakers have historically insulated from competitive challenges, and whose market power masks the decline of internal capacity. The difference between Boeing and many of these companies is that there's simply no masking airplanes falling apart in the sky.

And so I think this recent experience just really underscores the stakes of the decision we face as we choose between a future where we're continuing to consolidate control and power, versus really enabling and empowering open markets in fair competition.

As we continue this work, a few key principles are top-of-mind. First, we are squarely focused on how business models drive incentives. Just as we've seen behavioral advertising fuel the endless collection of user data, model training is emerging as another feature that could further incentivize surveillance.

The FTC's work has made clear that these business incentives cannot justify violations of the law. The drive to refine your algorithm cannot come at the expense of people's privacy or security, and privileged access to customer's data cannot be used to undermine competition.

-2 (c)-8h 2 (at (We pih(a1) -t (gs) z(e (w) 2 (7) 2) -0 (io) 7 (t)-5 (is) 7-2 (m) 4 (e)-2 mare h plof(ng)-3 (-14.) 4 aue2mare h plof(2 rt9,2 (bu)

AS

And with that, I'll turn it over to our Deputy Chief Technologist, Alex Gaynor, who will be leading our first panel.

Alex Gaynor:

Thank you, Chair Khan. My name is Alex Gaynor, and I'm a deputy CTO in the FTC's Office of Technology. One of the hallmarks of AI, for the history of computer science, has been that AI is about what is at the forefront of computing. Whether that's research into Automata Studies, which, believe it or not, was the original name for the research that became artificial intelligence, machines that can play chess, or large language models.

One of the features of the era of AI we are living in today is that they're uniquely demanding of cutting-edge computing resources. That's why I'm thrilled to introduce a fantastic set of panelists, representing

Thanks so much, Alex.

So this is Amazon, this is Microsoft, this is Meta. Even Tesla is making their own chips. And what this does is, it further makes it difficult for new entrants to come into the market. And on top of that, when you look at, as a startup founder, you look at the access to capital, you have to go into the room and explain to investors, "Okay, so we're going to compete with the largest companies in the world, and we're not sure if they're going to ever buy our chips because they're already making their own chips."

And it's extremely challenging even more so because, when you look at the way that a hyperscaler can operate, when they're making their own chips, it gives them unparalleled access to surveil. And a sort of form of innovation surveillance where they can see what their customers are doing. They can look into the memory inside of the chip itself, and they can see what is actually running on it. And what this means is, they can figure out what needs to be made before it needs to be made.

Now, in the semiconductor business, it is the apex of human science and technology collaboration and coordination. It survives on robust and strong partnerships. Nvidia has been a partner with TSMC since 1998, and it has been an enormously productive partnership for both companies. And the challenge here is when the partnership ceases to be a partnership, and it becomes a competition between a company and its own customer.

And as it relates to the semiconductor business, this means that you're kind of at a non-starter when you get out of the gate. And beyond it, it's really a function of, okay, well how do we get more venture capitalists to invest in more chip startups?

If we all want to continue to lead in AI, which we do, we need to have more companies than just the few, and just the hyperscalers, producing these chips. So in order for there to be a market for these chips, there has to be investment. And the problem now is, I looked at the data recently, and there's about 5,000 venture capitalists who have made investments or actively investing in AI startups, but there's only 300 that are investing in semiconductor startups.

And so when we look at this, it's this situation where the dynamics of the market basically make it a non-starter not just for entrepreneurs, but also for the investors themselves. Now, this is not to discount the success and the emergence of some really great AI-focused semiconductor companies, but they have primarily had to compete with hyperscale companies that have enormous volumes, and themselves can finance the development of their own chips.

So Ganesh also talked about further d8gn ches-6 (her)8 (elo)5 (pmsye)-3 (s)5 (c)-9(t)-6 (hes)5 (e)-2 (c)-8 (hi)5 (ps)7 (. S)s

Thank you. Beyond just startups, Corey, what are you seeing in terms of the needs of companies building their own AI systems to offer higher-level products and services, with needs like training workloads? That is, not the clouds themselves, the people who are trying to build AI products on top of them.

Corey Quinn:

People tend to miss across the board just how much work it is to train one of these large models. Last week there was a talk given by James Hamilton, SVP at Amazon, talking about how the servers that they're doing their training jobs on cost roughly a third of a million dollars a piece. And a recent unnamed training run they did internally cost \$65 million.

This is not the sort of thing that almost any startup is going to be able to raise money from, since the Vision Fund isn't doing what the vision Fund Used to do. So this does fall to larger companies.

And as we see it, all roads lead through one company, and that is Nvidia. Now, there are people who would say otherwise. That no, there are other GPU manufacturers who are serious players in this. AWS themselves, where I focus most of my professional energy, has been building their own chips for a while.

They make their own arm-compatible general-purpose computing chip called Graviton. And those are decent chips, don't get me wrong. And we never can forget that they make them because Amazon does not, and cannot stop, running its corporate gums about them every day (s63p27R33fad)4 (i28ha bys W)4 (e)JTJ E germane to the conversation it happens to be.

They also make two other types of chips that are in the GPU space, Tranium and Inferentia. Because if there's one thing that Amazon is remarkably consistent on, it's naming things terribly. But they don't talk about that in the context of their own generative AI services. So we know they're not using them. We know that they're using Nvidia chips to do this.

The presentation last week mentioned all the economics around Nvidia's, H100s. Each GPU from that costs \$30,000, give or take. These things are extraordinarily expensive, and hard to come by, and they're azon d73 (it)-k E

money to wind up getting one of those things instead, just because the back orders are stretching into years.

Alex Gaynor:

Thank you, Corey. Tanya, Ofcom recently conducted a market study on competition in the cloud computing space. Can you explain why Ofcom pursued that study, and what the findings were? What's the impact of cloud competition on downstream users like AI?

Tania Van den Brande:

Thanks, Alex. So the purpose of the study we did at Ofcom was really to understand and get to the bottom to some of the potential competition issues in cloud. And that was really a priority for us because as we were saying earlier, cloud is becoming quite essential to how the UK economy operates, right? It's reshaping the way businesses work, and we see it as an enabler of technological innovation that then creates new business opportunities.

And when we first launched into our study, we didn't really have a preconceived view on whether there were any competition issues, but we particularly wanted to explore whether some of the outcomes we saw in UK markets were signals that competition wasn't working well.

And the first one of those was the observation that cloud in the UK is very concentrated towards AWS and Microsoft. And secondly, we were starting to see some evidence emerging that customers were struggling to switch.

Now, we published our findings last fall. And in a nutshell, we identified what we thought were a number of factors

Tania Van den Brande:

... sectors that suggest competition isn't working well in cloud. And it's really for that reason that we've referred cloud to our competition authority in the UK for an in-depth investigation.

Now, let me unpack a little bit what drove our concerns in relation to competition. And in a nutshell, that was really some of the barriers to switching that we saw for customers, but also the difficulties that they have to use more than one cloud provider. And first of all, one of those barriers we saw was egress fees. These are the cost that customers pay to move their data out of a cloud. And according to the work we did, it can really get quite expensive for a customer when it's running a multi-cloud architecture, particularly if it's moving data between different clouds in that process. And we also saw scenarios where those egress fees can make it really expensive, the switch, particularly if that switching needs to happen gradually and customers using a multi-cloud architecture during that switching process.

Now, a second barrier we looked at were the costs and the effort that customers need to put in when they need to re-engineer an app and move it from one cloud to another, and that makes switching hard.

centers things went down a lot more, but the failures weren't correlated. It wasn't effectively every business having a problem.

Even today, if you decide that you want to build an e-commerce store and I'm going to build it on Azure so I don't have to deal with AWS in any way, well, if you're using Stripe to handle your checkouts, they're a full in AWS company. So if AWS has a bad day, no one's buying anything on your shop. Those dependencies wind up happening across the board.

Now, speaking to egress fees as well, there's a misunderstanding in many parts of the industry around them where it means you can't switch from one provider to another. They're high, but they're not that high. Storing the data inside of a cloud versus sending it somewhere else, the transmission of it out of that cloud costs the same is storing the data for roughly four months. So it's high, but it's not egregiously

Alex Gaynor:

Tania... Yes, please.

Tania Van den Brande:

Yeah, thanks Alex. Just to build on what we've discussing already, I think what creates incentives for innovation, right? It's that ability to attract new customers and to make money for the innovations that you've made. And really for me, what gets to the heart of all of this is that you need to have a possibility for customers to move around easily and to benefit from those innovations that sort of fits what they need.

And so building on some of what Ganesh said, I think that means it's got to be easier for customers to move around and choose a cloud provider that has a set of AI solutions that they're most interested in, and not necessarily just the ones from the cloud provider themselves, but from third parties that might build those apps on top of those clouds.

So in the UK context, currently, our competition authority is looking at these issues and to the extent that it finds concerns, we'll be able to make interventions, and you can imagine that that might focus particularly in part on some of that ability and that incentive for customers to switch. And I think what's

you can talk about your 5 and 10 year roadmaps. You don't want to be talking about those things with someone who is in the position to do it before you.

So beyond that, there is the possibility for a future where there are more clouds, where there are more chip companies, where the open source community has more options than just Nvidia, where an AMD can say, you know what? It's actually worth prioritizing even more investing in the ROCm and making sure that there is portability. Those things will happen, but we have to take action, we have to do something about it.

Corey Quinn:

From my (m) ye 6 (et)-TJETQ9o4tk o4that t biininyuhtuaine t knoher (o)4 eopting3 (pe(ns)7 (t)-6 (ha)4 (t)39 (ir)-23 (f)10

And then the third point, which we haven't talked as much about but I will use as use my time here to say something about, is that I think there are important ways for enforcement and regulation to address some of these problems. And I just want to specify a few because we didn't have a chance to talk about them earlier, but I think they're important to get out on the table. These are all quite traditional different kinds of legal rules that have been applied across many different sectors.

But the first is a structural separation. And this is basically an idea that you restrict the lines of business operations within an entity so that they can only operate one or a set number of lines of business as opposed to being vertically integrated across many different lines of business or a conglomerate that applies across many different lines of business. That is a very clean and administrable way to prevent things like self preferencing and other kinds of harms that may emerge from integration.

Second are non-discrimination rules. These are rules that say first that an entity that is operating a kind of service has to treat everyone on similar terms and really create a level playing field that again, would apply to self preferencing, but it also applied to other kinds of preferences that could be both related to price, it could also be related to terms and conditions or different other kinds of types of orders or applications. So I think having a sense of non-discrimination rules is another important way to ensure that there's confidence for users of a service that they are not going to be effectively price gouged or appropriated out of all the potential benefits that they might have for their innovative idea to make profits and money by the entity that has a bottleneck over an essential service.

A third and related point to non-discrimination rules is transparency of some of these terms and conditions so that we really know, and I think this is something that Corey has referenced a couple of times, that in some cases we may not even know what is going on in some of these areas.

And then the last point that I'll . raise is one about what kinds of rules there are around interoperability. Again, something we haven't talked as much about, but the ability to actually be able to switch between different providers. So I think those are some of the things that are kind of standard legal tools both in the remedy context in antitrust cases or in the regulatory context that we've seen imposed in many different sectors where there are similar concerns about monopoly or oligopolies being dominant and the harms that can come from them. And so I just want to make sure that to viewers and listeners, that they understand there are real solutions that have been workable across many, many different contexts and those could work in this context too.

Alex Gaynor:

Well thank you to all of our panelists for this fantastic discussion. I'd encourage the folks at home to give them a virtual round of applause.

To recap a few of the points that we just heard. The panelists discussed ways that dominant firms may have control over key infrastructure inputs such as cloud computing and access to hardware such as GPUs, and this may be exacerbated by obstacles to migrating between offerings. This may in turn allowBT7T1 1 Tf11 0 0 1

Thank you for tuning in and see you shortly.

Amritha Jayanti:

Hi everyone. Welcome back. My name is Amritha Jayanti. I'm the deputy chief technologist in the FTC's Office of Technology and emcee for our event today. We are now back from our short break and I'd like to turn it now over to Commissioner Slaughter for some remarks. Commissioner Slaughter, over to you.

Rebecca Slaughter:

Thank you Amritha and thank you especially to all of today's panelists, attendees, and to Stephanie and the whole team at the Office of Technology for putting on this important summit.

I'm really grateful for this opportunity to speak before the incredibly distinguished group of panelists coming on the AI and data models panel. As Stephanie's team has said, we're at a pivotal moment in the emergence and rapid deployment of AI technology and other advanced algorithms. There are stories nearly every day about their potential to transform industries, upend markets, and even change the

I'd like to hear about the lessons we can learn from that first digital revolution so that in this era we can really get it right. Thanks again to the team at the FTC and to our guests for being here. And I'll now pass it over to Krisha Cerilli in our Bureau of Competition who will be leading the second panel. Over to you Krisha.

Krisha Cerilli:

Thank you, Commissioner Slaughter. Good afternoon. My name is Krisha Cerilli. I'm the deputy assistant director in the FTC's Technology Enforcement Division. We are the division that investigates and litigates potential antitrust violations by technology companies. It is my privilege to host our next panel which is dedicated to the role of data and AI technologies and models. I'll ask my fellow panelists to please join me on screen at this time.

While they're joining, let me set the stage for our discussion. We just heard about the importance of cloud computing and specialized chips to the deployment of AI. Fair to say that data is also an important input to AI development. For context, public reports indicate that certain AI foundation models involve hundreds of billions of distinct parameters that have been traced using many terabytes and trillions of tokens of data. The use of data at this scale raises a host of legal and policy issues including related to competition and consumer protection.

On the consumer protection front, for instance, what are the privacy implications of companies using consumer data to train and produce content in a generative AI model? On the competition front, for instance, is there an even an open playing field with respect to accessing the data needed to compete in AI? Or is there a meaningful risk of market concentration and market power?

Thankfully, I have a distinguished panel here with me to help unpack and discuss those issues. The topic of how data is used in AI is obviously fairly broad, and of course we can't cover every nuance or every voice and perspective in just an hour, but I hope this discussion will help advance the conversation and surface valuable insights even though we can't cover everything.

So let me now briefly introduce our panelists and we'll jump into the discussion. First, we have Cory Doctorow, who is a science fiction author, activist, and journalist. He is the author of many books and has been inducted into the Canadian Science Fiction and Fantasy Hall of Fame. Cory is a paid special advisor to the Electronic Frontier Foundation, an organization that campaigns on issues related to digital privacy, free speech and innovation.

Next, we're also joined by Jonathan Frankle, who is a chief scientist of neural networks at Databricks, which offers a data intelligence platform powered by AI. At Databricks, Jonathan leads a research team toward the goal of developing more efficient strategies for training neural networks. He recently completed his PhD at MIT during which he empirically studied deep learning. He is also actively involved in policymaking issues related to AI challenges.

Next, we're joined by Amba Kak, who is a trained lawyer and technology policy expert with over a decade of experience in roles across government, academia, and the nonprofit sector. Amba is currently the executive director at the AI Now Institute, a research organization that focuses on policy responses for artificial intelligence. She also currently serves on the board of directors of the Signal Foundation and on the AI committee for the board of directors for the Mozilla Foundation.

And lastly, we're joined by Stephanie Palazzolo. She covers artificial intelligence at The Information where she also writes the publication's daily newsletter in AI. Stephanie has broken news about OpenAI,

And even though there's obviously a lot of activity, I do think that things have calmed down since the kind of AI funding frenzy of early 2023. And I think even more now, I kind of see this tale of two cities narrative emerging. So on one hand you have a lot of companies that are getting tons of venture capital at really insane prices, but on the other hand you do have a number of startups that are really struggling to find any funding at all. And so I think there are a couple of things that are determining which one of those buckets you fall into as a up and coming AI startup. So I think the first thing is a lot of investors because things are so early, they're kind of grasping onto these characteristics of these different startups to understand if they think it's a good startup or not.

And one of the big ones is talent. So I've noticed a lot of investors going after companies that are founded by ex-OpenAI researchers or maybe scientists that were at Google or from very... some of the top colleges in the U.S. And that kind of makes it harder for founders that maybe come from other types of backgrounds to get funding and to get capital from these investors. And I think OpenAI, even though it is in many ways still just a startup, it has kind of fallen into this role as a market leader and it's indirectly kind of choosing which startups win or lose in this AI wave. So I've noticed a lot of VCs are really hesitant to back things that either directly compete with open AI or even are in areas that open AI could maybe go into at some point. I think a lot of investors were also burned by some early generative AI bets that they made that haven't quite worked out.

And so I think they have a little bit less appetite now to back startups that are maybe more research oriented, will take a bit more time to get their product to market or are just less proven out. But sadly within this group kind of falls the startups that I mentioned earlier. So the ones that are building new types of models, which again are very important to challenge market leaders like OpenAI and Google in terms of developing these foundation models.

And so all of this is very important because I think capital is crucial in AI even more than in other tech spaces because that determines whether you can strike deals with third parties for valuable data, like Cory was just talking about. Whether you can pay for chips to train or run these models. And I think just the last thing I wanted to mention is this interesting thing that I'm keeping an eye on moving forward, which is how the incentives of investors are going to play a bigger role, especially as these companies get older.

And so a lot of the earlier stage investors care most about having huge growth or having a really talented team or a great idea, but later stage investors and especially those on the public market, they care a lot more about whether a company can generate cash and their margins. And so we've actually written a number of stories recently that dive into the margins and the cash generation capabilities of these companies. And because it takes so much capital for them to buy chips, to buy data, to hire people, it's actually much harder for them to generate cash and they tend to have lower margins than traditional software businesses that we've seen in the past. So this could obviously change as we move forward and as chips get more efficient, for instance. But part of me does wonder if VCs these might come to kind of regret their actions of funding a number of these startups at very insane prices.

Krishna Cerilli:

Thank you, Stephanie, that's definitely helpful and useful insights. I have a specific follow-up question. You mentioned potential investors being deterred from going up against current market leaders. Do you have a sense of why that is or what might be deterring that investment?

Stephanie Palazzolo:

Yeah, I mean, for instance, I know this was mentioned in a number of stories that we've written and other papers have written too, but OpenAI had a developer day last year where they announced a lot of new products. And I remember talking to startups kind of in the wake of that day and many of them... So for instance, one new feature that OpenAI announced is this GPT store where you can make customized chatbots for different use cases. Like oh, a chatbot to help a middle school teacher write a lesson plan.

And I remember talking to one startup in the weeks following that where that was their entire concept leading up to that event. And now they're kind of just like, "What are we supposed to do? We raise money on this idea, but now OpenAI, which has \$13 billion of funding is going after the same thing. How are we supposed to compete against that?" So I think a lot of investors are just wary of backing startups that could even be in an adjacent area to what OpenAI is doing now because they're worrying that if OpenAI or Google or some other big tech giant goes up and tries that as well, that their investment just won't be able to survive.

Krishna Cerilli:

Thank you. Let me now turn to Jonathan. Jonathan, you have experience studying AI models and also work for an AI technology company. What do you currently observe related to the competitive dynamics, related to development of AI models, including access to data, and are there any sort of consumer protection or consumer harm issues that you 4.9am[(w)10 (o)4 (r)8

to touch on the question of the risks of providing open source models that anyone can customize, for goals you may or not agree with. I think that's a much more complicated discussion. But it's something that I think I imagine other folks are going to have more to say than I do, given that's not my area of expertise. But it's certainly when you give someone control over an artifact, you give them control over an artifact. The last thing I'll mention briefly is we do tend to think of open sourcing as a binary, or at least making access to a model freely available as a binary.

And it's something I personally think a lot about just in my day-to-day work. In my one hand, I'm at a business building AI. In my other life, I'm an academic, working with PhD students, and writing papers. And I'm always trying to think about, "Are there middle grounds in open sourcing? Are there ways to say get more transparency by having limited access to a model, where someone like my PhD students could get access by applying for it, but we don't just put it out there freely?" And I do wonder whether there's a design space there. I don't have the answer. But it's a question I think it's worth asking.

Krishna Cerilli:

Thank you. Jonathan. Jonathan offered an important, I guess, qualification that even the concept of open and closed may not be, or is not completely binary. It might be a matter of degree. But Stephanie, picking up on that, even to the extent that there's models that are open to a degree, what have you seen or observed related to their impact on competition in particular? Have you heard of any limits that they have in relation to somebody's ability to compete using an open source model?

Stephanie Palazzolo:

Yeah. Well first I was going to say, I'm going to have to ask Jonathan after this what he thinks we should call open source models instead of open source. Maybe we can start a new trend. And call them open access models, or something. No. But I thought those were really, really great points. And I think to your question about competition, I think something that Jonathan mentioned was really interesting to me, which is, "How are these businesses sustainable in the long run?" And the question of business model, and how startups that are making open source models make money is very important if you're one of their customers that's depending on them to build a product on top of their open source model. And I guess one question I have is, this is a question that's come up a lot with other types of open source software and tech as well, which is, "How are they going to sell this?"

"How are they going to make a product on it that people are willing to pay money for?" And I think, although I agree with Jonathan that in a lot of practical use cases, you don't always need the latest and greatest GPT-4 model from OpenAI, and you can personalize an open source model for very cheap, for a specific use case, if, for instance, if an open source model developer decides that they want to take on a business model where they have a larger, more advanced version of their open source model that you have to pay for, just given the amount of funding that companies like OpenAI, and Google, and Anthropic have, it just seems very difficult for me to imagine that these open source model developers will be able to compete on the most bleeding edge, advanced models out there today.

And I think it's de4 (XT1 l)38don a

they have 350. So that's literally such a huge scale of difference. I think there's just so much room for professors and researchers at these universities to be part of this conversation. And they don't have the same incentives that profit driven companies do. And I think we really need to be encouraging that a lot more.

Krishna Cerilli:

Thank you, Stephanie. So our last couple of questions have done a little bit more of a deeper dive on competition issues. I want to switch back to some of the consumer protection and privacy concerns that were raised earlier, and turn to you Cory, and ask, "From your perspective, do changes need to be made to better protect consumers related to privacy? And in particular, what do you think the private sector can do? And how might government enforcers and policymakers approach those issues?"

Cory Doctorow:

Well, I think Americans often underestimate just how primitive the state of American privacy law is. The last time we got a really big muscular improvement to our national federal privacy regime was in 1988, when Congress got worried about video store clerks leaking their video store rental history, and passed a law prohibiting that activity. The Internet's come a long way since then. And the threats to our privacy have come a long way since then. And at the Electronic Frontier Foundation, we've been talking about something called Privacy First, which identifies a potential political coalition for federal privacy law, that

Thank you, Amba. As we approach our last 15 minutes, I'm going to have one last question, which is to invite everybody to offer the audience one thing they'd like to take away from this discussion. You can make it anything you'd like. But I might offer one specific thought to prompt some ideas, which would be, "From your perspective, what does success look like for consumers and other stakeholders as AI is increasingly deployed across the economy?" Cory, do you want to kick us off?

Cory Doctorow:

Sure. I'm happy to. So as I've said a few times through my interventions in this panel, which again, I want to express my gratitude for getting a chance to speak to you about this, we need to think about the problems of data beyond a property rights regime, beyond the idea that if you make data, it's your property. And someone else has to pay you, and get your permission before you use it. Because what we want to make sure of is not that everything in the models is paid for, but that the public and other stakeholders aren't harmed.

And if it's possible to pay for the data, and still enact the same harms, still displace creative workers with the work that they've done for you, still possible to produce grotesque privacy invasions in the form of non-consensual pornography, still possible to harm people by mining their data to make inferences about them that are adverse to their interests, then we have managed to fail to solve the problem, while still creating a bunch of law, and wasting a bunch of time, and incidentally, also creating a regime in which people who have the mon4.(ngm454.35 47hu0 612862 r.n11 72.025 431.63 Tm[(in)5 (w)9 (hi)5 (c)-8 (h p)3 (eop

Stephanie Palazzolo:

Yeah. No. I also just wanted to echo what Cory said, just say thank you so much for letting me speak on this. It was great to chat with everyone. I think I'll cheat a little bit. And I'll say two quick things. I think the first thing that I love from this panel is this idea that everything is not black and white. It's not open versus closed source. It's not more laws versus letting people innovate. It's not the academic side versus industry. There's a lot of gray area here, as we all talked about today. And I think it's the responsibility of, especially the media, and people like me, even though it's much more easier to write stories, and just say, "Oh. It's X versus Y," I think it's up to all of us to make sure that we're discussing this, keeping those gray areas and nuances in mind.

I think the second quick takeaway, and leading into what I think success will look like is that venture capital, and the tech giants have a very large role in picking what startups are going to win and lose, coming out of this AI boom. As I mentioned earlier, for instance, just the fact that OpenAI exists is stopping investors from backing certain types of companies. And I think for me, success means leveling that playing field both within industry, and within academia, that more players have a chance to compete with the big industry labs.

Krishna Cerilli:

Thank you, Stephanie.

Krishna Cerilli:

Okay. Jonathan, what are your closing thoughts for the audience?

Jonathan Frankle:

I think I want to be a little forward-looking, I guess. I'm an engineer by training and I try my best to, where there's work to be done I try to have interesting answers, and try to be creative about this. There are two things that have really been on my mind in that respect effect.

I had mentioned before that one of the consequences of competition is that people really do feel a lot of

The other big thing, again, I'll mention, which I'd mentioned before, is thinking about open sourcing in a more flexible way. It really is treated as a binary right now, in terms of either you open source your model and throw it out there to the world, or you keep it secret and never tell anyone about it, or how you built it, or share it with anyone. I have to believe there are good middle grounds. It's certainly something I personally plan to experiment with at Databricks, just to see if we can get more things into the hands of the academics sooner. Putting on my academic hat, I like the idea that academics have the ability to interrogate and ask hard questions about the kinds of technology we're putting out there. I think we should all be a little more open-minded to ways that we can try to raise the bar, try to create new possibilities, and allow this technology to get out there and do it in a thoughtful way.

Krishna Cerilli:

Thanks very much, Jonathan. Amba, would you like to leave us with some closing thoughts?

Amba Kak:

No pressure. Yeah. Again, thank you for convening this conversation. Maybe I'll go back to where we started with Commissioner Slaughter's remarks, which is that there is nothing that is inevitable about the current trajectory of AI. That is really important to keep remembering, and reminding everybody. Because I think this current AI race is based on certain assumptions about both scale and speed as a proxy for progress. And it's a view that's based on narrow benchmarks, it's one that never really properly contends with the longer term environmental, or labor impacts, or the impacts on our

needed to compete. Certainly, that raises questions about whether the AI models will be developed and deployed in a way that fosters competition and introduces new competitive pressures on incumbents, or whether those challenges associated with access to data and other resources might steer AI

LM 7A
11/11/11

9

Tel: 0161 275 1111
Fax: 0161 275 1111
The Trafford Centre

Part 1 of 1

Part 1 of 1
Bibliography

9

is being used to make decisions about them, people knowing why those decisions were made, people knowing what they can do if those decisions misfire, and people generally appreciating what is happening around them and why. Success looks like people being in control of technology, not the other way around.

The second answer I would give here is more from the competition lens, which I know is rightfully the focus of most of our panels today. I think that from a competition side, success looks like us using a technology, not because the biggest or most powerful company put it out, and put it right in front of us using the platforms we already use, but rather us using a technology because it has proven itself to be the best technology in the marketplace on its merits. Success looks like people like Jonathan, or other

fee, and you can see...
side of the...
the open...
one model, you can download...
if you compute...
on your computer...
can't...
efficiently...
everyone in the...
And so...
specialized...
they're...
his...
really looked into...
phones...
the cost...
his...
They're competing...
pricing, but...
not...
of the...
as...
art...
ing on cost...

Andy Has...
And one last...
quick question...
le bit...
of heart...
before...

Conrad Kamer...
Yeah, I'm really excited about...
because...
I'm really excited...
repetitive...
going to...
That's...
in part...
it...
to get...
easier...
it's...
one of the...
forms...
of hope...
next...
potential...
entirely...

Andy Has...
Thank you, Conrad. Kaen, I'm going to...
experiencing on the...
creating...
going to...
them, and we...
what...
and support...

Kaen Hao...
Absolutely...
helpful...
Difficult...
it's...
also...
Thank you for...
excited about...
their...
generation...
an achievement...
seen...
the...
concept...
des...
building...
ignoring...

And they might not even know that there's a model there. The commissioner's point, this has also been a problem. Not just with large language models, this has been a problem with predictive AI models, where someone might've gotten falsely arrested by a facial recognition algorithm and they didn't know that it was a facial recognition algorithm at work.

So I think that both these problems have created a layering effect in terms of the lack of transparency that consumers are dealing with. And from my perspective, as an investigative journalist, it is really, really complicated and it takes me a lot of time to and unspool all of the different ingredients, all of the different vendors, all of the different challenges of these tools. So as a consumer, it feels really impossible.

Andy Hasty:

Thanks very much, Karen. Ben, turning to you, you recently co-authored a report on generative AI harms. What risks do you see, and all of us, how might we address them?

Ben Winters:

Hi, yeah, thanks for having me. Yeah, I think it's really difficult to try to grapple all of the risks that the generative AI explosion in cultures are promoting and exacerbating. When we wrote our report, we tried to boil this down as well as we can, but there's a few different ways of looking at it from the outset.

So we looked at thinking about social harms, which are often a little bit less tangible, like the impact on the environment, the fracturing and stressing of the information ecosystem that the availability of these text generation tools has. The impact on elections, the impact on that competition ecosystem. And then there are also individual harms. So that is the data theft, the data security, the victims of non-consensual intimate imagery that people can create with these available tools. And so that's one binary way of thinking about it.

And then the other places, the other ways we try to approach it is thinking about where that harm is felt, and then where in the process is the cause of that harm coming into the model? So in terms of the functional harms that we identified, we had nine ones that I'll run really quickly through.

Up top, I mentioned the information manipulation part. So that's how mis- and disinformation is being supercharged by the availability of these tools. Whether it's just on ChatGPT or whatever, or one of the countless other checks generator tools that every company is racing to do, whether they're making it themselves or purchasing it from another vendor and just pretending it's their own. And that goes to Karen's point about how it's so impossible, even if you are expert, to realize how many different vendors you are interacting with in a given moment.

The second big problem is the harassment, impersonation and extortion. This is a really big problem when you think about non-consensual intimate imagery, otherwise known as revenge porn. But that also comes up with potentially intimate images of children. It also just comes up when you have someone using the voice of Barack Obama, and then connecting that to a robocaller and telling voters that they shouldn't vote, or they should vote on the wrong day or something to that effect. Those are all fracturing the information ecosystem, and there's a lack of transparency and accountability between how those outputs can be made and how they can be disseminated. So there's a lot of different actors throughout the creation and dissemination of it that can and should be responsible.

they do not constitute legal inter

That was a lot, so I'll stop for the moment.

Andy Hasty:

Thanks, Atur. I wanted to circle back to an issue I think Karen raised, and I'm going to direct this question to you, Karen, and you, Conrad. Karen, you mentioned marketing, it seems like companies are using different terms in their marketing to signal to users that they care about things like privacy and safety, labels like AI safety, or privacy enhancing. It sounds like you've noticed that, so I can skip that part of the question. What have you seen, and what do companies mean when they use these terms? How do

being weaponized for disinformation in this sort. But at the same time, it's a relatively singular, and clear and defined concern." And what this point showcases is in the context of AI safety as defined by AI developers, there's this constant comparison that's being made between the potential, hypothetical existential risks of AGI, and the present day risks of the models that we have, like disinformation. Where the hypothetical risks are always going to be worse. If all of humanity ends, that is always going to be worse than any other thing.

And so what he was arguing is, you need to deploy the models so that you can learn from it. Deploy now, and then we can learn how to handle the potentially existential future that we're facing. But then this is all wrapped under the term safety. And so consumers think, oh, when a company is deploying now and iterating later, they are shoring up all of these other things that consumers associate with safety, they are shoring up privacy, they're shoring up security. So essentially, it provides this cover for companies to continue behaving in these dangerous ways, but consumers think that it's all for their own good.

And one last aspect I'll mention, and this is a bit of a perhaps tangent, but when I think about enforcement, one thing that we don't talk a lot about is whistleblowers. And I think, we at the bureau, we continue to encourage industry whistleblowers, those that, in the consumer financial products and services realms, who see potential misconduct, to report those concerns to us.

I'll stop there and turn it back to you, Andy.

Andy Hasty:

Thanks, Atur.

Okay. So we are approaching the end of our panel here, so I think we need to start going into wrap up mode. I'm going to take a cue from the earlier panels and open it up for each of you. What is one thing... And before I ask the question, I'll give you the order, I think. Conrad, if you could start, and then Karen... I'm sorry, Conrad and then Ben, Karen and Atur. But the question is, what is one thing you want the audience to take away from this discussion? What would be a win for us as consumers, as workers, as small businesses, as entrepreneurs, et cetera?

So, over to you, Conrad.

Conrad Kramer:

Yeah. I think the highest level takeaway is that, as an entrepreneur and as a startup, you can build a product that complies with all the laws. You can build a product that respects users' privacy. And you can build a product that is safe for them to use.

I think, the biggest takeaway is that, me building a product today, I'm not concerned with AI safety in terms of existential risk. I'm concerned with the risks of consumers using the product today, being confused and how it interacts with the world. And so, I think that a lot of other startups share this responsibility, and I think further that doing the minimum or doing what is required is one approach, but I think that startup to take this seriously and actually innovate ways to give privacy to users and to build safer models, I think, will ultimately succeed. Yeah.

Andy Hasty:

Thanks, Conrad.

Ben, I'm sorry. Over to you.

Ben Winters:

Good. Yeah.

Andy Hasty:

Closing thoughts?

Ben

Thanks, Karen.

Okay, Atur. So, Conrad, Karen and Ben were pretty concise here. We have, I don't know, four minutes or so. I think you mentioned that you had some thoughts on some earlier questions, so now would be a time to address those if you want before getting to a closing thought, or you could just go to your closing thought. But I'll leave it up to you.

Atur:

Sure. Let me try to run through my earlier thinking.

It really surrounds the fact that AI itself is just a marketing term. It's kind of an amorphous term that oftentimes describes things that are very simple models and also ones that are really complex. So we have this tendency or there is this tendency to add another layer of marketing, like safe AI and things of that nature, but it really just feels like a murky mishmash of words. So whenever I hear these things, what I think is ultimately important to keep in mind is that, there are laws and there's laws that have existed for a really long time that we have a lot of experience with, that have been used to prosecute companies when they make, for example, deceptive marketing claims. And the CFPB has actually been active in this space, and one public(dec)-8 (eptT8Tm.)5 548. 0 65 (s)l73.sive tpaive tpaime marketing, like safsddec w2onfCple

context and applications, and this requires deeply thoughtful resource allocation, risk assessment and triage, which leads to a third theme here.

I think we heard a clear call for more upfront risk mitigation. These powerful new AI computing technologies, decision-making technologies, content generating technologies are increasingly being

comprehensive rule to prohibit and deter the practice. We've made clear that firms can't retain kids' data forever, even and especially to train models. And we've established that firms must either take steps to ensure their AI tools don't harm consumers, by discriminating against them, or cease to use these tools altogether.

None of this is to suggest the more resources and authority are not needed. But what should be clear is that we are using every tool, enforcement, rulemaking, education, market studies, and more, to protect the public from emerging harms. Our most important asset is our people, our multidisciplinary teams of world-class attorneys, economists, investigators, consumer education specialists and technologists. And all of us benefit enormously from events like these that engage top experts from both inside and outside the government to better understand how AI is reshaping the marketplace. A generation from now, when my future bureau director, or their AI avatar, discusses the history of this era, I am confident they will recount an FTC that was active and engaged in harnessing the AI promised can be harnessed for the benefit of people, rather than a handful of tech giants.

Thank you again for joining us at this important summit, and I'll now turn things back to Amritha.

Amritha Jayanti:

Thanks so much, Sam, for those excellent closing statements. And thank you again to Henry for your participation. And with that, we come to the end of our FTC Tech Summit on AI. Before we fully close out, I want to communicate a few more thank yous. Firstly, thank you to the staff across FTCA's Office of Technology Policy, Bureau of Consumer Protection, Bureau of Competition, Office of International Affairs, the Commissioner's Offices, the FTC events team and other staff for their work in planning this event. Thank you to Chair Khan, Commissioner Slaughter, Commissioner Bedoya, and CTO Stephen Coyne.