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 can hire the sharpest software engineers, designers, and researchers in the world to boost and protect their bottom line. And at the same time, perhaps more quietly over the years, there's been a unique shift in the application of what a technologist is in government.

Beyond building, designing, and deploying digital products and services at agencies like US Digital Service, GSA and 18F, law enforcement agencies like the FTC, Consumer Financial Protection Bureau, and the Department of Justice are firing up a surge of tech capacity on cases, investigations, policies and rules.

And as public servants, we have our own bottom line that drives our work, to serve and protect the American people. To date, we've onboarded a dozen brilliant technologists with skill sets to cover a broad swath of the economy, including a software engineer and journalist who built an open-source Tor browser for iPhone, a geneticist with deep machine learning expertise in health, a privacy engineer who developed safeguards for AI for kids and teens, and an investigative data journalist who covered algorithmic discrimination and the gig economy.

Scrappy, relentless, and adaptable, this small but mighty team is making the most of the resources we have. We're working directly on the cases and investigations, and engaging in the policy and horizon-scanning research. And we're receiving input from the communities who have been overwhelmed by security breaches, surveilled by hidden data brokers, and squeezed out by the companies who have outsized power.

Today the FTC is convening this summit to examine layers of the AI tech stack, to understand the extent of competition across the various layers and sub-layers, and their potential impacts on consumers. We'll look at three layers, hardware and infrastructure, data and models, and consumer applications.

The first panel will focus on the computational infrastructure of AI, including computer microprocessors and cloud services. The discussion will look into how dominant firms can have control over key hardware inputs like semiconductor chips, and we've highlighted that cloud computing is a central part of the economy. Our request for information flags submissions that raise concerns about widespread reliance on a number of cloud providers, potential security risks, and business practices affecting competition.

The second panel will discuss data and its use in training AI models. We'll discuss the methods of data collection for model development and training, and how they can have implications for competition and consumer protection. Through rigorous enforcement, the FTC has highlighted how companies can collect, use, and abuse user data, which can be magnified with the speed and scale of generative AI.

We've seen that third-party tracking pixels enable platforms to amass, analyze, and infer information about user activity. And we know that highly private data through voice recordings and videos can be used to train algorithms. AI facial recognition technology can recklessly be used for surveillance purposes.

And finally, the third panel will examine another area of long-time interest to the FTC, including the impacts of how AI is being deployed in consumer applications. We use many tools at the FTC to uncover these benefits and harms. For digital artists, sci-fi writers and musicians, we recently held a round table, and published a report to amplify their perspectives on how generative AI is impacting their creative fields.

And to prevent, monitor, and evaluate fraud and scams through AI-enabled voice cloning, we're partnering with the Bureau of Consumer Protection on an exploratory challenge. Above all today, we're eager to be learning. We've got a full lineup today with speakers from diverse sectors; startup founders, civil society advocates, researchers, and government officials.

Thanks again for joining, and without further delay, I'll turn it over to Chair Khan.

Lina Khan:

Thanks so much, Stephanie. Good afternoon everybody, and thanks so much for joining the FTC's first Technology Summit on artificial intelligence. I'm so excited for this event, and really thrilled to welcome all of our panelists. We've really managed to secure a fantastic set of folks with deep, deep expertise, and I'm really looking forward to learning from them in the conversation today.

I'd also like to just give a special thanks to our office of technology and our fantastic chief technologist, Stephanie Nguyen, for spearheading today's summit. As Stephanie noted, the Office of Technology is relatively new, but I've just been so struck by how they've hit the ground running, and really ensured that the FTC is on the front lines of tackling some of these new challenges that artificial intelligence tools are presenting, from the voice-cloning challenge that they've set up, to convening our conversations with creators to make sure that we're continuing to learn and able to adapt and update the application of our tools and legal authorities as needed.

And in partnership with our Bureaus of Competition and Consumer Protection, our team has really been working with great agility as we navigate this fast-moving moment of technological opportunity and risks.

Over the last 18 months, the rapid deployment of artificial intelligence tools has captured the world's attention, spurring some combination of awe, wonder, apprehension, and fear. We hear how these automated technologies could open up the door to breakthroughs across fields ranging from science to education, making life better for millions of people.

But we've also already seen how these tools can turbocharge fraud, automate discrimination, and entrench surveillance; putting people in harm's way. More fundamentally, we face basic questions of power and governance. Will this be a moment of opening up markets to fair and free competition, 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 an 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 mass 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.

Similarly, we faced similar questions prompted by new technologies in the mid-two thousands at the onset of the Web 2.0 era. And unfortunately, we saw that what began as revolutionary set of technologies ended up concentrating enormous private power over what have become near-essential services.

And through aggressive strategies to acquire or lock out companies that threaten their position, a handful of firms solidified their dominance while locking in business models that we now realize came at the expense of our privacy and security. Lawsuits around the country have surfaced the heavy costs from the decimation of independent journalism to serious harm to kids' mental health.

Today, policymakers across government recognize the importance of learning from these missteps as we navigate the challenges and opportunities posed by AI. At the FTC, the rapid development and deployment of AI is informing our work across the agency, as we look to promote fair competition and protect Americans from unfair or deceptive tactics.

There's no AI exemption from the laws on the books, and we're looking closely at the ways that companies may be using their power to thwart fair competition, or trick the public. As part of this effort, the commission today is launching a market inquiry into the investments and partnerships being formed between AI developers and major cloud service providers.

Through using the agency's 6(b) authority, we are scrutinizing whether these ties enable dominant firms to exert undue influence, or gain privileged access, in ways that could undermine fair competition across layers of the AI stack.

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.

We similarly recognize the ways that consumer protection and competition enforcement are deeply connected. With privacy violations fueling market power, and market power, in turn, enabling firms to violate consumer protection laws. And our remedies will continue to focus on deleting the models themselves in addition to unlawfully collected data.

Second, we are squarely focused on aligning liability with capability and control. This requires looking upstream and across layers of the AI stack, to pinpoint which actor is driving or enabling the lawbreaking, and is best positioned to put a stop to it.

What AI liability regimes will ultimately look like is still an open question. But our enforcement experience in other domains will directly inform how the FTC approaches this work.

For example, our recent robocall enforcement sweep not only targeted telemarketers and the companies that hired them, but also looked upstream to the lead generators and voice over internet protocol providers that enabled the illegal telemarketing. And in our recent work to combat scams, we are holding upstream payment actors like Walmart accountable for knowingly facilitating the fraud.

Third, we're focused on crafting effective remedies that address the underlying business incentives, and also establish bright-line rules on the development use and management of AI inputs. The FTC is making clear that some data is simply off the table for training models.

For example, our recent order against Rite Aid bans the company from using facial recognition tools after its reckless application of the technology led to innocent people being falsely accused of shoplifting. And our recent cases against data brokers include bans on their using or monetizing people's highly sensitive location data.

As we continue to establish rules of the road for AI, it's also essential that we set clear boundaries on the content that can and cannot be used for scraping and model training.

The commission recently held the public workshop with creative professionals to better understand the types of guardrails that would help protect against creators' work being appropriated and devalued by generative AI models, including in ways that may undermine fair competition. Our subsequent report on the creative economy lays out our core concerns, and how our authorities may apply in this space.

Across our work, we are making clear that firms cannot use claims of innovation as cover for law-breaking. And we've already made that clear through a number of blog posts and other signals that we're sending to the market to make sure the model as service companies, and cloud providers, and others throughout the AI stack know that there's no AI exemption from the laws on the books.

Learning from our experience in the mid 2000s, we're using the full scope of our authorities to make sure that these hard-learned lessons don't repeat themselves. Much is uncertain about what the precise future of this technology will look like, but the good news is that we have the experience and expertise to meet the moment. And by continuing to sharpen our thinking, and faithfully enforce the law, we can unleash AI's potential benefits while safeguarding Americans from the potential harms.

I'm so grateful for today's convenings. The types of insights that will be shared today will directly inform how the FTC approaches our work. And so, just want to give a big thank you both to our team that helped put together today's event, and all the speakers who took out the time to come speak with us.

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 several different perspectives, discuss the role that chips and cloud providers have in artificial intelligence. While these are not all of the perspectives in the world, I think you'll find these panelists' expertise to be invaluable.

If all of the panelists can turn on their video and join me on screen.

Tania Van den Brande is a director of economics at Ofcom, where she works on competition and online safety issues in digital markets. She led the analysis of Ofcom's recent market study into the nature and extent of competition in cloud infrastructure services in the United Kingdom.

Before joining Ofcom Tania worked as an economic consultant, where she advised clients in front of European and national competition authorities in relation to a wide range of competition matters, covering FMCGs, mining, petrol retailing, healthcare and telecommunications. Tanya has expertise in the use of estimation and simulation methods to analyze competition issues.

Dave Rauchwerk is a technologist and entrepreneur. Dave has been involved in open-source software and hardware for over 20 years. For the last 10 years, he's been an active member of the startup scene as a founder, systems engineer, and advisor.

He's a frequent speaker on emerging technologies, and has lectured at University of California Berkeley, Stanford, South by Southwest Design Automation Conference, and O'Reilly's OSCON. His work has been featured in major print, television and online media including CNBC, BBC, Fortune, Wired, Time, Popular Mechanics, EE Times, and NPR. Ganesh Sitaraman holds the New York Alumni Chancellor's chair in law at Vanderbilt University, and is the director of the Vanderbilt policy Accelerator. He's the author or co-author of numerous books, including Networks, Platforms and Utilities' Law and Policy.

And Corey Quinn is the Chief Cloud Economist at the Duckbill Group, where he specializes in helping companies improve their Amazon web services bills by making them smaller and less horrifying. He also hosts the Screaming In The Cloud and AWS Morning Brief podcasts, and curates the Last Week in AWS, a weekly newsletter summarized in the latest in AWS news, blogs and tools, sprinkled with snark and thoughtful analysis in roughly equal measure.

So to get us started, Ganesh, you've previously talked about the layers of AI technology. Can you walk us through that framing a bit, and help folks understand why chips and cloud access are uniquely critical when it comes to AI systems?

Ganesh Sitaraman:

Thanks so much, Alex.

Yes. So when I think about AI, and I think when many people think about using AI, they think about an application like ChatGPT. But for an application like ChatGPT to work, it requires an underlying model that's provided by OpenAI. That model is trained on a lot of data. And that training process, and where the model is hosted, and how it works operates on actually hardware. It's called the cloud, but it's not up in the sky, it's actually on the ground.

And this computing infrastructure run by three providers in the cloud, primarily in the cloud space, is essential to making the models work. And the computing infrastructure in the cloud is fundamentally based on having semiconductors, chips, and those are produced and manufactured and sold as well.

And what I think is very striking as you work down these layers of what's called the AI tech stack, applications, models, cloud, and chips, is that there's more and more concentration as you work your way down.

The model layer, we may think of a number of different companies that are currently offering models. But when you look at the cloud layer, there are three main companies that dominate that service. At the chips layer, there's one company that is predominant in designing the most advanced chips. There's one company that is predominant in manufacturing those chips, and there's literally one company only that produces photolithography equipment, which is an essential input into manufacturing those chips.

Now, the challenges with this is that at the lower layers of the stack, where there's more concentration, you have the standard problems, potentially, that we see with monopolies and oligopolies in lots of different areas. Concentrated power means that entities have the ability to preference their own vertically-integrated business lines to discriminate against users of their service, and of course to increase price and reduce quality of services.

And so I think one of the challenges that we should anticipate as there's development of AI over time is that concentration at these lower layers of the stack, if nothing is done, will lead to increasingly seeing these kinds of problems that we see pretty standardly in other areas, where there are monopoly and oligopoly providers.

Alex Gaynor:

Thank you, Ganesh. Dave, you previously founded a semiconductor company. Given Ganesh's landscape of AI layers, can you talk a little bit about the experience of a founder at that foundational chips layer, and what it takes for startups to compete at that layer?

I think you're still muted.

Dave Rauchwerk:

Ganesh got right to it, right? Which is that at the most foundational layer of AI, when we talk about the semiconductors that make it possible, it is an oligopoly. As a startup founder, you are basically competing with the most valuable companies in the world.

And it goes beyond the success of Nvidia, which we can talk about. And it includes fundamentally all of the major hyperscalers, all of the major platform companies. What we've seen over the last five years, through a series of acquisitions, is platform companies, hyperscalers, starting to make their own chips.

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 down the stack, and the actual fabrication of the chips. And we look at what's happening with Chips and Science Act, and you see that Intel is also in the business of making their own chips. And is also getting into the business now, with their IFS strategy, of producing chips for other companies.

And what comes to mind is really two companies that we can talk about this more, which is Nirvana was a company that had a full stack co-designed AI solution, and was sold to Intel, and then was shut down when Intel decided that Habana was the way they wanted to go. And you can see this also in the treatment of Risk Five, and the treatment of the other companies that are coming into the market.

So, as a former semiconductor founder, now has never been a better time to be in the semiconductor business. We're going to have more fabrication capacity online in the next five years within the country than we've ever had. And there's enormous demand. And yet, the dynamics of the market make it extremely challenging to get off the ground.

Alex Gaynor:

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 chance it gets. Regardless of how 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 massively supply-constrained.

I have a customer who loves using as many of these things as they can get for their own internal training runs. They're a massive company, and they're constrained in how many they can get. Because how these things are being doled out, not just by Amazon but by Nvidia to other companies, has always been something of a black box. There are a number of small cloud computing upstarts providing infrastructure on these GPUs, and without exception, every single one of them has deep historical connections from the founders into Nvidia.

Nvidia's CEO has been at a number of different tech events over the past year speaking in basically every tech company keynote, saying basically the same things. I figure that's how he's planning his vacation schedule these days. And it's honestly a story of, how are you doing the allocation? They have become effectively the new king-makers in this entire space.

And that becomes very galling for customers who might want to build a good chatbot. But instead, when they visit AWS's website, they have Amazon Q, which is their new pop-up thing that answers questions badly, and spreads misinformation hilariously. Picture Microsoft Clippy if it suffered a traumatic brain injury and you're pretty close.

They have to deal with that thing sucking up these resources that they feel they could use to make a better answer, but the capacity is simply not available. Because no one knows how these things are getting doled out. I can't go to the store and buy one today, if I somehow decide to sell a car and use the 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. But also some of the difficulties they have in connecting apps that are hosted on different clouds. And we thought that could make multi-cloud more difficult.

And finally, we pointed to a number of discounting structures in the industry that we thought create quite strong incentives for large customers to put all or most of their cloud needs with a single provider, and particularly make it unattractive for those large customers to split their cloud usage between larger and smaller cloud players.

And really those barriers for us pulls a risk competition in a number of different ways. First of all, if a customer finds it hard to switch or to add another cloud provider, well that makes it more difficult for them to benefit from the best deals that are available in the market and that can really lead them to paying more than they have to or make it more difficult for them to mix and match the kind of product in the market that best suit their needs.

But we're also worried that some of these barriers create a risk that the cloud market might concentrate even further towards the market leaders, and particularly that they make it difficult for small cloud providers and challengers to really go after customers that are already established on AWS and Microsoft. And that worried us because we thought that makes it more difficult for those challengers to start building their customer base and to really start gaining that skill that you need in cloud to be a more direct competitor in the market, or at least threatened to become a more direct competitor.

Now in your final question on AI, that wasn't really the focus of the work we did, but you can see how some of those concerns directly translate. If you're a customer and you're finding a hard switch, well then that might mean that you can't benefit from the best AI solutions that are available from other cloud providers. Similarly, if you are a challenger and you've got a really good innovative AI solution, you might not be able to attract the customers that you need to get that foothold in the market and scale up and become sort of a stronger player. And I think it's really interesting what Corey is saying because it sort of points to potential concern where particularly AI developers might not be able to get the access to the compute power that they need and definitely not at reasonable prices. And that can create a bunch of questions in terms of equal access and level playing field when those AI developers create their solutions. And you can imagine that's a particular risk where those AI developers are trying to create something that competes directly with what cloud providers are doing.

So I think the good news is on the UK side is that our competition authorities now are having an in-depth look at competition in cloud, and it's already signaled that it'll think about the implications of AI in that context.

Alex Gaynor:

Thank you. One of the risks we look at with respect to competition has to do with single points of failure. How should we be thinking about this risk in the context of computing and AI? And Corey, I'll throw it to you to get us started and then invite the other panelists to jump in.

Corey Quinn:

Oh, dear Lord. I think the war has already been lost. We've passed a tipping point where you cannot avoid the three main hyperscalers out there, full stop. In fact, if there were to be a law or technical issue passed tomorrow where AWS could onboard no new cloud customers, they would continue to grow revenue for at least several quarters just based upon organic growth. Too big to fail has passed a tipping point long ago, and the centralization risk is massive. Once upon a time when we all ran our own data 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 so. There's a concept known as data gravity, because it's expensive to move data around, you keep the data where it is and then the compute workloads such as AI stuff starts centralizing around that. We're seeing an inversion of that when only some providers are doling out access to the kinds of compute that we need in the form of Nvidia GPUs. So you wind up effectively having to bring the mountain to Mohamed, so to speak. As a result of this, we are seeing people do massive data transfer projects from where they are to get it close to these things, and it's sort of turning a lot of this stuff upside down in a very topsy-turvy way.

But I think at this point it's obvious that we have if not a monopoly, the next thing to it. These cloud companies talk in the language of monopolists, which always touches on these ideas of, oh, it's a fight for survival for them. They could be out-innovated tomorrow by a startup in their garage. Well, yeah, if you give that startup $6 billion of funding for all of these AI training runs they'll need to do and the massive hiring binges and the specialized hardware, yeah, then maybe. But I kind of don't see it.

Alex:

Ganesh, do you want to build on that?

Ganesh Sitaraman:

Well, I doubt I will be as quotable as Corey. One of the things I think that is a challenge is that when you think about concentration and your question about single points of failure, is exactly the resilience point that Corey was mentioning, and I think that appears in different ways. I mean, Corey talked about one version of it. In the chips layer, the CHIPS and Science Act is addressing a different kind of resilience challenge, one tied to geographic production of chips and where they're located, obvious national security issues and concerns there. So I think one of the other places we need to think is just how concentration can be important in that direction.

On a couple of other points, I think if you are the federal government, there's also a question of dependence on a single actor in the private sector for a significant amount of your compute power or really for any other resource as a federal government, one of the challenges that you worry about in a democratic republic like our own is who's actually calling the shots. And to the extent that there's significant lobbying, regulatory capture and other dependencies, it may be that government actors over time don't feel like they can take significant enforcement or other actions against companies on which they are dependent.

This is the sort of too big to prosecute idea that emerged in a number of sectors after the financial crisis in which there were concerns that the too big to fail, banks were so large and so important systemically in the economy that enforcement actions might be problematic against them for the effects that they would have. I worry also in that case that we may end up in a similar kind of situation if there's a very limited number of actors upon which there's real dependence, particularly by the government, but across the economy as well.

Corey Quinn:

Yeah. The government, the US federal government runs a staggering percentage of its compute workloads on the big three hyperscalers. I'm not suggesting that there's unfair influence of stop investigating us or your computers are going to stop working. I don't think anyone is getting to that point. But there is a sense of how much enforcement can really be done when effectively you are critically dependent upon the continued existence and wellbeing of these companies just to go about the daily business of government.

Dave Rauchwerk:

Yeah, Corey-

Alex Gaynor:

Please.

Dave Rauchwerk:

Yeah, Corey, can I jump in here? I want to talk about Nvidia and I want to talk about really just the general idea that Nvidia is doing incredibly well. And if you look at the history of Nvidia, they've always been a really developer-centric company. I think there's another issue here, which is, okay, well, there's a point at which we need more chips and we're there. And the thing about it is, if you look at, for example, Argonne National Laboratory, they have an entire suite of different chips from startups like SambaNova, like ROC, Habana via Intel. And there are options, but we're not seeing them in the cloud providers.

And this thing there is that when you look into it's not that there's one player, but what if we imagine a world where there's 50 companies making AI chips and it's probably going to be chiplets, and we could talk about chiplets and what an ecosystem will look like that makes that possible. But right now, there's not enough cloud providers and there's not enough chip companies. And if we had a world where there were more clouds and more chip companies, there would be more competition. And Nvidia has a lot of margin to play with right now. As an entrepreneur, that's like, wow, okay, that's really interesting. That means that if I can start a company that could compete, then there's a lot of money to be made, and I'm competing with Nvidia, but at least there's opportunity.

And it's complicated further, I'm curious what you think of this, but how portable is a training run going to be on a Trainium, and are we going to be able to move it to Azure on Microsoft's processor? And it's interesting to me because there is so much innovation happening in real time at the lowest layers of the stack, but we may not get to actually see it because of the dominant players and the concentration at the hyperscalers. I mean, it's to the point where the hyperscaler becomes the customer for the chip startup. And you talk to chip companies and they say, oh yeah, well, if we can sell to one hyperscaler, one data center is millions of units and it'll make our whole business.

And so, in the sense that there's a distortion and right as the company shift to address that demand, the hyperscalers come out and announced they have their own chips. I don't know if I've articulated it well, but I'm curious what you think.

Corey Quinn:

Well, even in the open-source community, in the hobbyist world, everything is around Nvidia. The only time you see people building these things on other chips is similar to the old trick of installing NetBSD on a toaster. Just to prove it can be done, but no one's seriously suggesting you go ahead and run your data center on those things. It is an Nvidia monoculture now, and that's frankly what scares me.

Dave Rauchwerk:

Yeah, I mean that's real. I think that the momentum of CUDA, the CUDA moat as it were, is a real thing. But it's also that Nvidia has had a history of selling downmarket, starting as a gaming company. It gets back to this thing I was talking about earlier that's worth some more discussion which is, okay, well, there's a point at which Nvidia becomes a leader in the AI market because it listens to its customers. It realizes, oh wow, our gaming GPUs can be used for this thing called machine learning and AI, deep learning. It's amazing. So we're going to make improvements to our designs, to our products to accommodate that. And so they've got this early lead.

But if the startups that come out... like SambaNova, for example, why can't I get a PCI card from them at a competitive price to put it into my machine and then participate in a different part of the hobbyist market? Well, they can't get to scale. It's really interesting that the FTC has convened this because we're in this moment where, okay, what is the point at which the market leader who gets there by fair means over 30 years becomes anti-competitive? And what needs to change so that we can have more companies building more chips? Because you can imagine a chip comes out and it enables a whole new business model for a different kind of cloud that could compete on training efficiency. But we are not seeing that, we're seeing a CoreWeave that's just a giant Nvidia cloud.

Corey Quinn:

We need the software stack to support it up and down the stack. It needs to be as easy to get started with these other things without having to know the intricacies of what GPU you're doing current training on. And nothing I've seen indicates that any of the tooling's there yet, just because there's no a need for it to be.

Alex Gaynor:

Kind of building on that, what should we want to see in the marketplace and what would enable more competition. And I'll pass it to Ganesh then Tanya, to get us started.

Ganesh Sitaraman:

Great. Well, thanks Alex. So one thing I think that's important is we want to see innovation. And I think a critical part of that when you think about the lower layers of the stack is the possibility that we will have innovation be foreclosed by vertical integration and dominance at these lower layers. And so I think addressing that is a critical issue.

Let me just give a couple of examples of how you can imagine something like this working. If you're a fully vertically integrated company and you have cloud models applications, it's very possible that someone develops an application or a kind of model that has some features that you think are terrific. And so you copy the idea, you integrate it into your own system, and that new startup is effectively out of luck because you have the scale and ability to operate that in every part of your ecosystem and roll it out to many, many customers because you're one of these major companies that is integrated across the whole stack and that you have all these different applications.

I think that's a real problem because what we might see over time then is actually less innovation in the model or application layers, anything dependent on these concentrated layers where there's a kind of bottleneck of players that can incorporate those new ideas directly into their own offerings and then spread them throughout their vertically integrated business lines. And that might mean that at net we end up being less innovative rather than more innovative. And so I think there's a real danger from concentration to innovation in these downstream areas based on dependence on the more utility like elements upstream. And I think that's a real concern. And so solutions to addressing that problem in a combination of enforcement and regulatory actions, I think are very critical. And we could talk about that more later or now if that's of interest. But that's a place where I think we should be very worried.

A second way that this could happen is not involving any sort of copying or taking of others' ideas, but simply self preferencing one's own integrated business lines over others. And when you're a very large company and you have a lot of customers and a lot of users, you have the ability to preference your own integrated offerings. And what that means is that others can't get a real opportunity to get into that market. And that puts real limits on new providers that might be more innovative, have better ideas, have more interesting ideas. And so self preferencing in this area could be another concern.

Traditional issues like tying are another potential problem under the antitrust laws. Obviously tying not allowed, but there have been, in our history, many, many, many examples of companies that have tried to do so anyway, and also enforcers who have gone after them for engaging in that kind of behavior. So I think we should be concerned about those things, and I think innovation is a key part of why. And if we want an innovative technology ecosystem, we're going to need to take action to make sure that there's enforcement to prevent these kinds of behaviors that would prevent innovation.

Corey Quinn:

This is nowhere scarier than with Amazon itself. Take a look at, what industries does Amazon not operate in? The only one I can think of is philanthropy. Employees in jurisdictions where it's not barred have to sign an 18 month non-compete scope to Amazon, which means that there's no industry they're not in, there's no way to not run afoul of that, which causes a certain chilling effect. But also it means that they can start doing a lot of the bundling and packaging across the board. Okay, Nvidia, give us a bunch more chips and we will give you preferred placement for your other retail line on amazon.com when people search for this list of terms. Is it happening? How would I know? There's no transparency here. And it's this crosscutting across so many different units of business that lets them start tying things together in a bunch of very strange and harmful ways that we just don't know if it's happening.

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 important here is that not only it can enable challengers, but also to make sure that the cloud providers keep being incentivized to go after each other's customer bases. Where that might be less the case if once the customer's moved in the cloud, they're more or less locked in.

Corey Quinn:

It's almost impossible. Moving from a data center into a cloud provider is a massive project that's measured with multiple calendar years. Moving from one cloud provider to another is almost that same level of difficulty. Once you're there, you tend not to move. There are a few stories of people fully leaving a cloud provider that they have been all in on. And for good reason, it simply doesn't happen at most. A workload or two will move from one to another or something greenfield will be spun up somewhere else. But once a workload is there, it basically is there to stay.

Tania Van den Brande:

I think that's an interesting point, Corey, and we sort of got different views from different customers, but I think a point you made earlier is interesting there, which is that customer demand is changing over time, new solutions are coming to the cloud, customers themselves might change. So even if you can't switch that established base, I think we definitely want to make sure that competition remains for any new demands that a customer might create over time.

And so the question whether some of the interventions that we could think about could at least make sure that customers can move some of those new workloads or new needs to a competitor that has a better solution.

Corey Quinn:

Inertia is such a powerful force it's hard to overcome. Oh, a new vendor to get through all of my procurement processes, my security validation, understand how it works, more importantly understand how it breaks because when it breaks, you really wanted someone who's been there before. And it becomes this almost insurmountable series of obstacles to the corporate decision making process where, okay, let's put this on Amazon too is a straight shot. The big get bigger and the gulf grows wider. It becomes a bimodal distribution whether we want it to or not.

Dave Rauchwerk:

The economies of scale are real, Corey. My hammer here is always from the semiconductor angle, which is, I look at the puzzle and it's like if I had a magic wand we would have 10,000 cloud companies in the United States. There'd be 10,000 regionalized clouds, they would specialize based on the workload. If you're doing a synthetic bio and you're training a model for that purpose, well, it turns out that there's a data prep process that could be done, and there's special chips that have been developed to accelerate that process. And then, for me-

Corey Quinn:

That would be so great. I wish I lived in that world.

Dave Rauchwerk:

Exactly. Right.

Corey Quinn:

Instead, Amazon last year bought my doctor's office.

Dave Rauchwerk:

There it is.

So it's this thing where if you want to have more chip companies you need more cloud companies, and we have too few cloud companies and now they're making their own chips. So we're sort of stuck. And it's this thing where real innovation is possible, real innovation exists in the market from the very fundamental base layers of the stack. But because the large cloud companies are not actually buying those chips from these innovative companies, we're not seeing that specialization. Effectively, it's not dissimilar to the embrace, extend, and extinguish model of Microsoft in the '90s as it relates to the semiconductor business. It's a wicked problem, because it goes all the way back to, well, if there's no market for these diverse set of chips from a large number of cloud companies and there's really only five customers, well, why would a venture investor invest in a chip company? So we have no chip startups. And it becomes this, as you talked about, this sort of self-reinforcing problem.

Corey Quinn:

Very much so. And the worst part is I don't see a way to fix this.

Dave Rauchwerk:

Well, we're hoping that the FTC has some ideas, I think.

Corey Quinn:

Hope springs eternal.

Alex:

Thank you all. So I'd like to bring us to our final question for each of you, which is what is one thing you'd like the audience to take away from this discussion? And I'll assign you an order. Dave, Corey, Tanya, and finally Ganesh.

Dave Rauchwerk:

All right, so everybody in the chip business, everybody who works low level on the stack is aligned, this includes NIST and all the people there, around chiplets being the future that we want to basically build systems of systems and that the workloads of the future being all flavors of AI are going to require mixing different heterogeneous architectures. And the problem that we have as we look at CHIPS and Science Act is we're talking about building a national champion as it relates to Intel. And the CEO of Intel has said publicly Intel's two companies under one roof. And that's Intel that designs chips and makes chips. And then there's Intel, the manufacturing operation that is going to open itself up to other companies. Now, so long as that's the case, there is basically a fiduciary responsibility for the management team of Intel to monitor what's successful within their fabrication business and effectively supplant any innovation to take it over. This is the innovation surveillance I've talked about.

And if you look at what has made Nvidia, Nvidia, Nvidia and TSMC are inseparable, and that real legitimate partnership between TSMC and Nvidia is one as Morris Chang would say, based on fundamental trust. That, we are going to grow together and we're going to both specialize one on fabrication and one on design. And as we look at what's going to happen over the next three to five years in semiconductor fabrication in the United States, if we want to have a greater diversity of chips, if we want to have more chip companies, we need to have a pure play foundry. We need to have a national champion not dissimilar to a TSMC. So that we have an entity that you can partner with that 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 part of it, I think that the biggest challenge that we're seeing is that all roads lead to Nvidia. They are today a bottleneck on all of this, followed only slightly by the large cloud providers that are their primary customers. I think that fundamentally, since all these companies love to talk about being utility computing, it's time to start treating them like a utility. We would not stand for an electricity provider or a water provider that suddenly slapped a zero or two on the end of every price that they were charging people. But there's very little stopping companies from doing that today other than not wanting to destroy goodwill and then evoke a sharp regulatory response immediately. It's instead they're taking the boiling the frog approach.

I think that in the short term increased transparency because right now it's who can genuflect the hardest to the feudal lord that they are sworn to. That's not a viable system for modern governance around distributing an asset that is right now incredibly rooted in scarcity. I want to see a better tomorrow, not so much a better 10 years from now. How do we start making small steps today rather than hoping for an ultimate solution decades from now when it's even harder to change then?

Tania Van den Brande:

Well, I've got good news, Dave and Corey, because I think regulators and authorities, not just in the UK but across the world, I think are starting to look into the risks that you've been highlighting during this talk, and I think have several tools available to them to intervene where we can find solutions. And particularly in UK, we're about to start doing regulation of the big tech companies on an ongoing basis where some of the issues, if they do arise, could be dealt with on a more ongoing basis. So really value the discussion today to really help the creative juices in how we might tackle issues as they emerge.

Ganesh Sitaraman:

Well, thanks. I'm going to actually highlight three different takeaways I think from this conversation. So first is at the lower layers of the stack, there are already monopolies and oligopoly, and concentration is already a reality in the lower layers of the stack. So that's the first point.

Second point is that I think we've heard today that there are significant problems that come from having concentration at these lower layers of the stack, including to innovation and to having a robust ecosystem for startups and for new entrants.

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 allow them to charge excessive prices or impose coercive terms, and as a result, they may be able to exercise market power in ways that favors their own incumbency or impacts competition. We've also heard about ways in which the structure of these markets may make it challenging for new players to compete even where their offerings may be better than the incumbents.

This brings us to the end of our first panel. We'll now take a short break and starting at 1:20 p.m. Eastern, Commissioner Slaughter will speak. Then our second panel on AI and data and models will begin.

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 work of government. Almost every facet of these developments has major implications for our mission to protect consumers and promote competition. But we've also seen this play before. We as a society, as well as our specific agency, are still dealing with the fallout from a government-wide largely hands-off approach at the beginning of the big data, ad tech, and social media era.

By learning from the past, we can unlock the promise of this new technology, make sure it is free from the control of just a few gatekeepers and enact guardrails to ensure its safe deployment and use for all.

I see the story of our regulatory posture to the social media, ad tech, and commercial surveillance era as unfolding in three parts. Act one was the emergence of social media powered by big data and ad tech and it was filled with nearly boundless optimism about our newfound ability to connect and conduct commerce easily and quickly over vast distances. Despite early warnings about privacy and market consolidation by a few advocates and forward-thinking legislators, regulators and legislators policed only the most egregious conduct. Act two in the 2010s, we saw users and whistleblowers and experts ringing alarm bells about data privacy, harms to kids and teens, and about the corrosive effect these large firms were having on our politics and markets.

Still, for most of the decade regulators were playing catch up. We're now in Act three, nearly halfway into the third decade of the social media, ad tech, and commercial surveillance era. We're seeing the consequences of the hands-off approach. The markets have consolidated into and around a few extraordinarily large companies. As these companies have gotten larger and controlled the advertising and attention markets, the once vibrant and iconic American journalism and arts sectors are in crisis. The business model of many social media firms profits off our most private information and at the same time, the content that does proliferate in these networks has facilitated a teen mental health crisis and widespread misinformation and disinformation online that threatens the integrity of elections and our political system.

Today we're acting in earnest and with urgency to address these issues, but it really is an open question whether a vision of our digital ecosystem as a dynamic and open space to connect and share information across boundaries can be restored. After a problem is entrenched is often way too late to correct it. We're in the middle of this play when it comes to the data-driven AI and algorithmic technology markets. This is a moment of both promise and peril. We've seen the deployment of advanced algorithms to make decisions in healthcare, criminal justice, employment, credit, housing, and other areas of economic consequence. Through vigorous strategic enforcement of our competition and consumer protection laws, we can begin to meet the challenge of this new era and create a marketplace that unleashes innovation while protecting consumers and competition. In the AI marketplace, competition can be hampered by limited access to key inputs such as computer processing power and chips, as we heard in the previous panel, and training data, as I imagine we'll hear in the next. These key inputs can be controlled by large incumbents. I worry that the concentration of AI models with access to huge amounts of consumer data in the hands of very few companies could pose enormous risk including around consequential economic decisions and access to opportunity, information and privacy.

If even a little of the hype around the power of AI models is to be believed, then it is emphatically our responsibility and obligation to support open and fair markets and prevent monopolies in their incipiency. We should also all take note of the proliferation of partnerships and direct investments involving AI developers and large technology companies that are structured to avoid triggering notification to the FTC or DOJ antitrust division under the pre-merger notification rules. It is reasonable to wonder whether these investments could lead to a heavily consolidated market dominated by only a few firms with either no competitors or with competitors who are hamstrung by their dependence on those same incumbents.

That's why I was very pleased to support the 6B study that the chair announced at the start of this summit. The way to stay on top of this quick moving market and avoid repeating the mistakes of the past is by using the full panoply of our statutory tools, including our 6B study authority, to make sure we are fully aware of the business models and related incentives and consequences that are taking shape in the world of AI.

Of course, studying a market is a complement to and not a substitute for appropriate enforcement if the laws and the facts support it. There is no AI exception to the law. Even investments that do not have to be noticed under the Hart-Scott-Rodino Act procedures may violate the underlying antitrust laws. I am confident that where there are facts that support enforcement investigations, our agency will pursue them. Proactive use of our Section 5 consumer protection authority will also be essential to unlocking the benefits of a competitive AI market while avoiding the mistakes of the past.

I'm a believer in the potential for AI tools to help make our lives easier, but as we saw in the commercial surveillance era, hype cycles drive investments in products and deeply informed consumer purchases. When that hype is built on a foundation of unsubstantiated over-promises about a product's capabilities, billions can be lost chasing snake oil.

Honest marketing claims are deeply pro-competitive. Deceptive marketing about an AI product's capabilities crowds out scrupulous firms that promise only what they can actually deliver. I also continue to be concerned about the intersection of data privacy and access to economic and commercial opportunities. AI models can use data in ways that risk deepening social inequities and bias decisions. Those can stem from differential inaccuracy for different demographic groups, as we've seen in bias facial recognition technology from the bias deployment of those tools in marginalized communities, to the repetition of existing patterns of economic distribution under the guise of unbiased algorithmic decision making.

I'd like to see us write a different play than the one we saw unfold in the commercial surveillance era. Nothing about pervasive data collection, ad tracking, the shape of social media, or the dominance of a few tech firms was inevitable. Inaction in the face of those developments was a policy choice. We have the knowledge and experience now to see this era play out differently. I'm so excited to learn from the leaders here today about how we can build a vibrant, safe, and competitive market in the AI era.

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, AI funding and other developments involving large tech companies. Prior to joining The Information, she covered AI at Business Insider and previously was a tech investment banker at Morgan Stanley where she advised tech companies on a variety of transactions.

So again, thank you panelists for helping facilitate this conversation. So without further delay, let me dive in and kick off the conversation with a question for you Cory. Can you set the stage for us by offering your perspective on how you observe data playing a role in AI and what the implications are from your perspective for consumer privacy and other consumer protection issues?

Cory Doctorow:

Yeah, thank you very much. What a treat to get to speak to you all and to hear Commissioner Slaughter. It is a slightly disorienting but very welcome experience to hear a FTC commissioner say really smart things about technology that I agree with. It's a hell of a time to be alive.

So it's very hard to figure out what to say about AI and data because it's very hard to figure out what people mean when they say AI. The investor story about AI is something like AI is everything that is good and valuable and everything that you do will be touched by AI. And obviously that can't all be true. And so we have to kind of decompose the data question into sub questions that really narrow it down from this, "Doing everything all the time. If you think it's valuable, that's us and you should give us money," to a much narrower set of important and concrete questions.

So I'll start with a data question because obviously anything that's doing data mining and applied statistics with data starts by gathering data and analyzing it. And we made a great error in the internet's history in deciding not to create muscular privacy or labor or consumer protection laws that really dealt with particulars of how technology worked and to fail to enforce the ones that were applicable for many years. And that left us with just one law that we really applied to the internet, which is copyright law. And so as we struggle with the data question, we tend to reach for copyright law and I think that copyright is just not a great framework for this, for a few reasons I'll sketch out here.

People in my line of work, I work in the creative industries as you heard. I've written a couple of dozen books and I live here in Hollywood and spend a lot of time out there on strike with my colleagues from the screenwriter's and the actor's guild. People from my line of work are very worried and I think rightly so, that our bosses would like to fire us and replace us with algorithms. And there is something intrinsically offensive about the idea that they'll take the product of your work and use it to make sure that they don't have to pay you anymore.

And so I understand why so many of my colleagues have reached for copyright, a property right, in order to defend themselves from this injustice. I think that it neglects something very important about the structure of the creative labor market, which is that it's a monopsony in which a small number of firms have enormous amount of bargaining power over the creative workers who generate the value for them. We are talking about an industry with five giant publishers, it was almost four, but thanks to our friends in the US government, it's still five, but we have four studios and three labels, two companies that do all the ad tech, and one company that sells all the eBooks.

And in that world, giving creators more to bargain with is like giving bullied school kids extra lunch money. There just isn't an amount of lunch money that gives the bullies enough that they decide to just hand that over to the kid and get the kid fed. And a regime in which we say as a condition of training a model you must first license the content, is not a regime in which creative workers are defended from our creative employers. It's a regime in which, for example, Apple offers $50 million to New York Times or The Atlantic and other organizations that have, to varying degrees, not been very welcoming to labor claims by their workforces, to sign over a corpus that can be used to train a model that can be used to erode the wages of those people. You have a regime in which all the training data is permissioned and in which the outcome you were hoping to prevent is still in play.

Likewise for people who worry about non-consensual pornography or non-consensual deepfakes for public figures and for creative workers, there is enough material whose copyright is held by people who are adverse to your interests, that they can train the model that can be used to displace you or harm you. To address these harms, we have to reach to things like labor and privacy law, not copyright law. It is not enough to merely have the right to feel affronted by conduct of firms. We should have the right to do something about it.

Now I quickly want to move on to something important about AI alignment and the AI story. So the investor story, as I said, is, "AI can do everything and will do everything and is infinitely valuable and you should give us all our money." One thing that is intrinsic to that pitch is the idea that AI will allow firms not just to make their workers output better, but to make it cheaper, to allow them to shed workers and replace them with automation.

I, as a lay person who has read and tried to understand the papers on AI radiology and solid mass detection in lungs, I'm willing to stipulate that AI can probably find tumorous masses in lungs that radiologists might miss. And I would be very happy if I ever had a suspicious mass to have the radiologist checked by an AI that said, "Maybe you missed something," so that the radiologist who might be processing 10 X-rays today would process 10 X-rays tomorrow, but maybe in fact it would only be nine because the AI flagged one.

I don't think anyone is pitching hospitals on that. I think the pitch to investors is that we are going to tell hospitals you can fire half your radiologists and double their output, and that is not the AI productivity and benefit world that we want. That is the alignment problem that we're worried about. I'll close by quoting something I wrote this week or last week rather, which is that we are nowhere near a place where a bot is going to steal your job, but we are well beyond the point where your boss can be suckered into firing you and replacing you with a bot that fails at doing your job. And I think that's the real AI alignment problem we should be thinking about.

Thank you.

Krisha Cerilli:

Thank you Cory. There are a number of important threads there that I want to pick up on as we move forward. So let me turn to you Stephanie. Cory mentioned kind of concern with concentration in content creation and also mentioned the landscape for investment in AI. What do you observe related to the current environment for AI startups who want to enter into AI model development?

Stephanie Palazzolo:

Yeah, so in my current role, I spend a lot of time looking at the VC market, at the startup market, and obviously there's a ton of activity with AI startups there. So definitely a lot happening when it comes to the development of AI models, especially large language models. On one hand we obviously have the large research labs like OpenAI and Anthropic, but we've also seen a number of smaller open source model developers pop up like Mistral. There's also a whole other group that's starting to emerge of startups that are building completely new types of model architectures. So right now a lot of the most popular models like GBT-4 or Claude, are formed of this architecture called the Transformer. And so now there's been a new group of startups that are trying to build non-Transformer models, which obviously has a very important effects on what ends up happening with competition in this space.

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.

Krisha 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.

Krisha 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 see developing?

Jonathan Frankle:

Yeah, so first I just want to say it's really exciting to be here. I mean, my mom worked at the FTC for several decades, so I have a lot of happy memories playing on the floor in the building and nice to get to come back. Now I do feel like a bit of a straw man here. I'm the one in industry, building AI. I'm certainly personally not promising magic or AGI. And if you chat with my friends in the community, I'm known personally for my skepticism of those sorts of claims. I personally look at AI as a medium for producing just really useful machines, kind of like code has been for the past several decades. And I think the important thing to take away from that is just how diverse the ecosystem is. I think one of the biggest misconceptions I typically come across in our field, and especially when I chat with folks in policy is everyone just assumes that because they know of OpenAI and ChatGPT, that's the only business model and that's the only way of operating.

And I hear this assumption a lot, and there're really a bunch of different kinds of business models. I mean there is the OpenAI business model, and the way that I would describe it to the best of my knowledge as someone who doesn't work there is they're trying to build one general purpose, incredibly powerful model, and they charge people for access to it. The model is the core IP and they might spend hundreds of millions of dollars or more per year to develop and maintain that IP. And there are lots of companies that do this. OpenAI, Anthropic, Cohere, Google, and even in image companies like Midjourney and Adobe are now kind of doing the same thing.

But this definitely isn't the only business model out there. I mean, at Databricks, my business model is basically help companies build their own models, often from scratch, on their own internal data. So my core IP is really the science and infrastructure behind building models, not the models themselves. And historically we've actually just open sourced the models we've created and shared all the details of how we build them and what we've trained them on because that's not really core IP for us.

And there are a million other business models out there. There's the ecosystem around these models, things like data labeling and ways of interfacing with models. And I think again, it's important to remind people AI existed before November 2022 when ChatGPT came out and there are numerous businesses building specialized AI or performing boutique consulting that have existed for decades. So I always want to make sure I make the point, the world is a lot bigger than just OpenAI and it's important to keep that in mind as we think about the ecosystem.

The other way I think about it is kind of vertically, there's a large stack for building AI and we just heard about chips and the cloud situations, everything from the foundries to the chips to the networking and servers to the clouds to the companies that are building the software to facilitate training to the companies that actually build the models.

And speaking from my personal experience across all these areas, because I have to interface with all them in my day-to-day work, competition is incredibly intense. Every day, my friends at the major clouds and numerous newer clouds people haven't even heard of are competing for my business. I talk to them at multiple clouds probably on a daily or weekly basis. And every day I'm fighting hard to compete for enterprise business. And things have even, I think, become a bit more competitive over time. There are new clouds that have popped up in the AI world and even chip companies that have historically not been that involved in the AI market are starting to now field really impressive chips going

Jonathan Frankle:

... going into this year. Plus the clouds are building their own chips. But it's worth saying, "What are the consequences of competition?" Competition itself, that is not sufficient to have good outcomes. One of the interesting things that I've seen is a few years ago when I was a student, it was just taken for granted that places like Google Brain or Facebook AI Research, or OpenAI would just share any new innovation they had. We were just academics. We would share. We would publish. And that's how we got here in my sense. The Transformer paper was at Google based on academic work. Meta and Facebook iterated on that, and riffed on it. Another team at Google built BERT, which was an important early model. And the people at OpenAI scaled them up, and led us where we are today. But especially in the past couple of years, that's really shut down. Because now, it's competitive intelligence. And nobody wants to give that away. We know a lot less about the inner workings of current generation systems than we do about their forebears. And that is one consequence of the intense competition that we're seeing. One of the reason for seeing this closing down is that from what I've heard from friends at these places, there's a lot of just legal and regulatory uncertainty for those who have business models similar to OpenAI, where it's based around creating one giant model. Nobody knows what they're allowed to use or, for what I'm seeing all sorts of stuff in the news about illegal issues on that front. I'm not a lawyer. But my understanding from lawyers is nobody really knows how anything is going to turn out the first time around. So there's incentive for organizations to be pretty secretive about what data they're using, and how they're building models, because it reduces risk in a really uncertain environment. And even for some of the most popular freely available models like Llama 2 and Mistral, we can access the weights of the models. We don't know how they were built.

If you want insight, there are lots of great open source data sets for training models, things like Wikipedia and GitHub. And folks in the open source world are building models out of that data that are competitive with the best in the industry. So one can get away just by using that data. But we don't know what's being used in general. And one last consequence of all the competition, there's a lot of pressure to get to market quickly. I feel really fortunate at Databricks just due to the resources we have, and the great customers we have. But when it comes to building any technology, doing it fast adds pressure. And there's a great example of this I think with the Llama 2 model from Meta. Meta was really careful about this. And you can read the paper. They detail some work they did to be extra careful.

They did 14 rounds after they trained the model, where they tested the model, collected feedback from humans about what they liked, or didn't like about the outputs, and then updated the model, 14 rounds. And you can do the math. If each round took a week, that's three months. If each round took two weeks, that's six months. And that's a lifetime in the AI world. And even getting the data to do this is incredibly expensive. One of the most important expensive inputs to building these models isn't computed. It's the data hand labeled by humans to make the model good at specific things. And for any of the fledgling startups that are out there, that three to six months, and those millions of dollars probably have to be weighed against racing to market, and making a name for yourself in a really competitive environment. And that incentivizes risk-taking. So that's certainly one of the consequences of competition that I imagine is on the minds of a lot of people I know at small startups.

Krisha Cerilli:

Thank you Jonathan. So I'm going to turn to you. Jonathan offered some important insights about his perspective on certain aspects of competition, and the risk to consumers. Can you offer your perspective on how you think the current approach and practices with respect to data collection and model development impact competition and consumer protection?

Amba Kak:

Absolutely. I am going to actually just build on the Tale of Two Cities analogy that Stephanie introduced. I think it's a really good one, and bring it to data in particular. I think just taking a step back, it's important to, I guess conceptualize data, and understand data as operating as a barrier to entry in this current AI landscape, along axes of both quality and scale. So this idea that data is everywhere, and therefore not a scarce resource is both intuitively appealing. But it misses the point. Because quality data, and this has everything to do with labor, data sets with high levels of human curation, and feedback, and niche data sets, especially in high impact sectors like healthcare or finance, data sets that come with assurances of accuracy, and legitimacy, and diversity at scale, these are all becoming a very key source of competitive advantage, especially in the hyper-competitive generative AI market that Jonathan was also just describing.

Right? And right on cue, we are seeing that trend towards more restrictions on publicly available data, higher cost of acquisition, and a turn towards more non-transparency and opacity around who is even using what data. I think it goes without saying. But it is worth saying that big tech firms have a very clear advantage here from the last decade of commercial surveillance, which Commissioner Slaughter described so sharply in her introductory remarks. They also, and this is crucial, have access to near unlimited capital to invest in labor, to make these data sets more robust. As a related point, we also don't know if and to what extent these data advantages will port to the so-called AI startups that they are strategically investing in, potentially creating new forms of dependency and power asymmetries outside of those that already exist via the compute and cloud arrangements as was discussed in the previous panel.

And I think the other axis on which this is happening, and Cory spoke to this is obviously IP and copyright. So we are seeing large AI companies already in extremely expensive content deals. Axel Springer is a case in point. So I think a question for the very near horizon is to try to watch for how large tech firms leverage their existing relationships, and importantly, their skewed power dynamics with publishers, and with the media industry to maximize both access, but also push for exclusivity. And I think again, this is not a problem that is unique or particular just to the foundation model layer. I think we are already hearing reports of how AI startups that are looking to fine tune models at the application layer are also finding it harder to acquire third-party data, whether that's from prospective clients that are more cagey about parting with data, or anxieties around cyber security.

And here too, the prevailing assumption is that bigger companies have the edge. So the other big point I want to make is that these data advantages are very self-reinforcing, right? So I just read a 2024 paper. It's still in preprint that is marveling at the ability of LLMs based on evidence to make personalized inferences at scale based on inputs, the questions that people are asking the chatbot for example, making inferences like, "How old might this person be," or, "What might their interest be," and, "Where do they live?" And Sam Altman has himself described individual customization, which is at best a euphemism, as the next phase for the company. Right? And he said himself that this is going to make a lot of people very uncomfortable. So I guess to say the quiet part out loud, "Why would it make us uncomfortable?"

I think it should. Because it does mirror and echo the commercial surveillance business model, the behavioral targeting based business model that has powered consumer search, and social media, and driven the race to the bottom that Commissioner Slaughter opened with. So I think in a market where, to Cory's point, the business model is entirely elusive, the regulators and the public I think need to keep a close tab on where this is headed, if all roads do lead back to some data-driven personalized advertising paradigm. And I'll close with this. Maybe we can come back to it later in the conversation, but just to really foreground data minimization as the key principle that has been endorsed and enforced over a decade, and just to say that this principle is more important, not less in the age of AI. Incentives already exist for invasive and irresponsible commercial surveillance. But this current version of the AI race definitely pours gasoline on that problem.

So I think that we need to start getting quite specific about drawing guardrails, and drawing them quickly, that we're not left cleaning up the mess after those incentives have already been supercharged.

Krisha Cerilli:

Thank you very much, Amba. We've heard a couple of references already to open source. And I want to specifically pull that thread, and turn to you Jonathan. Can you just describe specifically what open source modeling means in this context? And what is your perspective on how open source models may impact competition?

Jonathan Frankle:

Definitely. So I will start by giving you a hard time, and saying I hate the word open source when it comes to describing models in the context of AI. It is my number one pet peeve. And I'll unpack it, and give you a better term that I personally prefer, a better set of terms. Because I think it's too big of a word. And I think AI feels really new to a lot of us, including folks in the technical world. So we tend to seek out words from other areas to help us describe it. I hear the word red teaming a lot in AI for example. And it's a word that has a very specific meaning in traditional system security, or prompt engineering. So we tend to pull these words. Open source is one of those words that came. And it brings a lot of baggage that I think comes with us.

So I prefer to take apart the term open source when it comes to models, and think of it as two things. It's about access to models, and transparency about models. So by access, I mean, "Do you have access to the model weights itself, and not just a way to talk to it? Can you literally take the model, and manipulate it, and work with it yourself?" And a great example of that is Llama from Meta. Those are model weights that you can work with yourself. A counter example to that is GPT-4, which you cannot do that with. You can only talk to it. So that's access. And the other is transparency. "Do you know how the model was built?" To give you an example, we don't know a lot about how Llama 2 was built. The paper doesn't say a ton about the data that was used to train the model, or the details of the hyperparameters.

I would say that that model is we have access to it. We don't have transparency. And there are plenty of models that have been released by folks in the open source community where they've detailed everything about how they built the model. That to me is the full-on open source world. You have the weights, and you know what happened. So that quibble aside, and I think it's an important quibble, I hope definition maybe everybody will come around to my view of the world, someday. And I won't have to say this every time I get asked about open source. Making models freely available in this way, and really it's about when we say open source, a lot of times we just mean the access. We don't always mean the transparency. It has pros and cons. And I think there are some really great pros. I think just from a business perspective, I think about my customers. And Databricks has I think 10,000 enterprise customers.

We're mainly business to business. We don't really deal with consumers, or computer data to my understanding. Nobody mediates the access to that model for you. You can fully customize that model to your liking, whether you're a hobbyist in your basement, a researcher like my PhD students, or a Fortune 500 company. You can innovate by building on admittedly the very expensive work of others. And it's one of those standing on the shoulders of giants. It feels great. As a scientist, I love this. You also control your own fate. You don't have to worry about the model changing underneath you. I think one of the concerns I hear a lot from folks in enterprises is, "What if OpenAI posts an update to the model, and it happens to break my use case? Or what if someone decides to deprecate a model that I was relying on for this business scenario?"

"Do I have to architect my system?" With an open source model, you control your own fate. You also get greater control of your cost structure. You can serve the model yourself, and you know exactly all the inputs that are going into that down to the cloud level. But I do want to emphasize. These are all benefits of access to models. I haven't said anything about the benefits of transparency. I think the transparency part is just a little more complicated. And we can have a nice technical discussion about what you can or can't say about a model, based on having the weights, or knowing how it was built. But there are certainly complications and trade-offs. On the one hand, a lot of companies are being in some sense, very generous. Thank you to all the folks who have spent millions, or tens of millions of dollars to produce models, and basically donate them to the scientific community.

Speaking as a former PhD student on behalf of my current students, that's great. My students are doing great science, because someone spent a lot of money for free for them to go and have fun. I personally don't fully understand the business model behind this. And it may just be that I'm an AI researcher, and not a business person. And maybe I'm missing something. But I don't know if it's sustainable. I don't know how you justify doing this day in and day out. So I don't know whether if I were a business relying on the open source models today, I don't know if you could rely on that as your long-term strategy right now. I also imagine it has interesting consequences for startups. Stephanie mentioned there are a lot of startups out there doing a lot of different things. But I know a lot of startups, or I have a lot of friends at startups where they're trying to get a foothold in the AI space by building cool new models.

They're spending what they have to build a model, and suddenly, a much larger entity comes and drops a tens of millions of dollars model for free. That's really tricky for them I imagine. And I'm not even going 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.

Krisha 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 definitely great for just your everyday developer, your everyday company that wants to make a product, and just make sure that it works, and have a very specific use case. But I do wonder what competition is going to look like with the most bleeding edge models. And I think another thing too is my dad is a college professor. So I've heard a lot about life in academia. And I think we should also be putting a lot more research and funding into academic labs, whether that's data sets.

I see Jonathan give me a thumbs up, or chips, and things like that. It's insane. I saw a tweet recently that said that I think a lot of us have seen that Meta announced that they'll have around 350,000 H100s by the end of this year. Meanwhile, Carnegie Mellon, which has one of the top AI programs in the country, 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.

Krisha 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 touches on areas that are much broader than the harms of AI, and which are therefore potentially the subject of a coalition that's much larger than just the group of people who are worried about OpenAI.

If you are worried that TikTok is making millennials quote Osama bin Laden, or if you're worried that Facebook made your Grandpa a conspiratorialist, or that Instagram is making your kid anorexic, or that protesters at Black Lives Matter demonstrations, or the people who attended the January 6th riots are all being identified by Google through reverse warrants, you are someone who cares about privacy, as is anyone who is worried about the privacy implications of AI, whether that is models memorizing, and then regurgitating private information as we've seen, where sensitive information from medical histories, or resumes, or commercial databases of purchase history are sometimes being memorized and coughed up by these models, or whether you're worried about the truly grotesque generative AI problems with image and video generators, things like non-consensual pornography generation, copyright's not a great tool for dealing with this. But privacy law would certainly give you an awful lot of remedies to deal with.

You'd have a lot of people on your side who may not care about those issues, but are really worried about how their kids are being spied on by Instagram. A federal law that creates a private right of action, something that allows the public to sue directly if they can't get a local prosecutor to take up their case, which would therefore create a regime of no win no fee contingency lawyers, who'd be ready to take up this cause, would go a long way to disciplining firms that currently only think about copyright to the extent they think about any constraints on their data gathering, their data analysis, and then their data mobilizations.

Krisha Cerilli:

Hey, Cory. But you had also previously mentioned some concerns around consumer protection related to data collection. Do you have thoughts on this topic?

Amba Kak:

Yeah. I definitely have thoughts. I want to, if I may, just quickly circle back to the open source question. Because I think that that conversation was really interesting. And Jonathan and Stephanie both brought more nuance to even the term open source, which remains notoriously underspecified. I guess I wanted to underscore that it is important to remember, especially in a policy context, that open source companies are still operating in a highly concentrated ecosystem in which the largest firms retain both the resource, and data advantages, and network effects. So I think the lesson there is just that open source AI should not be assumed to be a stand-in or a substitute for structural interventions across the AI stack, because those firms are also vulnerable to the same dependencies, and will also need the same protections from practices like self-referencing that I heard were discussed in the previous morning panel on compute as well.

So that's just very quickly there. I think on the consumer protection side, I already previewed this earlier. But I do think that data minimization is a very core part of the toolkit. It's a principle that needs to be defended against a I would say, an existential threat, which is the idea that AI innovation requires a no-holds-barred regulatory orientation to data, that regulation is at odds with innovation. And we know from the last decade that if anything, it is these guardrails around data and cybersecurity that will actually shape the right innovation. So I mentioned that the inference part space is going to expand the area of surface attack when it comes to privacy threats. We've already seen unexpected, again, accidental leaks of personal information from many of these LLMs. So yeah. I think just to emphasize that this is more important, not less.

And then I think this prompts the question, which is that data minimization isn't new at the highest level. It's been around for more than a decade globally in various forms, in various laws including the GDPR. So I guess, "Why hasn't it worked," I think. Or, "Why hasn't it prevented some of these, the worst privacy invasive practices?" I think there, the lesson, if anything, is one on not allowing too much room for interpretation. Because I think where maybe the first decade of data minimization came to a head was on the question of, "Is behavioral advertising a legitimate business purpose? And if it is, does that mean we can just maximize, collect as much data, and keep it forever?" And I think as we look forward, acknowledging those administrability challenges, acknowledging that an interpretive wiggle room will be abused, to really focus on bright-line rules that don't allow that, that make it very explicit that AI training is not a free card to break down all your data silos, to violate purpose limitation, that we want to draw bright-lines around restricting particularly the use of the most sensitive data like biometrics or related sensitive attributes.

Yeah. And maybe I think this has been one of the, I guess, outcomes of today's discussion, but really I think pushing for more integrated regulatory approaches that don't silo out the consumer protection side of things, and the competition side of things. Because we have, again, seen how some of the largest firms really took advantage of that over the last decade to amass the information asymmetries that they have, and further concentrate their power. So just a integrated approach.

Krisha Cerilli:

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 money to pay for data licenses are the only people who get to play.

And I will say here that I think that having the money to pay for data licenses is not correlated strongly with being someone who will not harm the public, that there are lots of incumbents, with lots of money, who've got a strong track record for being the last people we want to lead us into the future. And I think that to answer your question about what success means, it would mean having an information governance regime that thinks about information beyond being property. It thinks about information, and the harms that arise due to privacy invasions, to labor violations, to consumer rights violations. And I will close finally by reminding the people listening that although we described some very valuable things by calling them property, that the most valuable things in the world we describe with non-property language, and that's people. Harming someone is not theft of their integrity. Killing someone is not theft.

We have a whole language for describing how the most valuable things in our world should be regulated, rated, controlled, liberated, compensated that are not about property rights. They're about a sui generis regime. And our information is important enough. In fact, it is so important, that merely giving tradable property rights is always going to be inadequate in the same way that we don't solve the problems of a lack of organ donors by creating property rights in kidneys, and then just letting people sell their kidneys. We need rights that deal with our information in a way that is cognizant of, and sufficiently important that we recognize its gravity, and its centrality. Thank you very much.

Krisha Cerilli:

Thank you, Cory. Stephanie, what is a takeaway you'd like to leave the audience with?

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.

Krisha Cerilli:

Thank you, Stephanie.

Krisha 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 pressure to race to get things out there. I've always hoped that, I wish there were just centralized shared resources for improving the safety and quality of AI models. I wish there were a dataset that my students at Harvard, or my friend who just founded a company could pull off the shelf, train their model on that a little bit further, and know that they're getting something that is going to behave in ways that everybody would find to be at least a positive development, in terms of moving the model in the right direction.

I can't promise anything will fix any problem with the model, but at least we'll certainly improve the state of affairs. This isn't the right forum to ask the FTC to go build that, but it is something that whenever I get the opportunity here, I like to mention as something I'd love to see in the public policy world, shared safety resources. Certainly something that a lot of organizations are talking about working together on, but having that publicly funded would go a very long way toward just raising the bar across the board, and ensuring that... or at least to minimizing the trade-offs that a company, the race to compete and the race to get cool things out there.

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.

Krisha 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 information environment.

I think this may be where I go in terms of looking forward is, and Stephanie talked about this a little bit, which is who gets to decide and shape what counts as innovation, and what counts as innovation for the public good? I think that one way forward is to really go back to the drawing board, or the table, which is currently populated with VCs, big tech firms, and companies that they invest in, and really have a much more broad ranging conversation that is dominated by public, rather than very narrow private interests about what counts as innovation and what is innovation in the public good? And try to, I think, shape that trajectory more actively rather than be passive recipients or subjects of the tech trajectory.

Krisha Cerilli:

Thank you, Amba, as well. Let me just conclude by offering a couple of final thoughts. First, thank you all panelists for the engaging, important discussion. I could just highlight a few takeaways from things that we've heard today. Certainly, the discussion underscored that the methods of data collection and model development have implications for both competition and consumer protection. On the consumer protection front, large volumes of data are being used to train AI models, and it raises a number of really key questions for us as an agency and other policy makers, to evaluate what are the legitimate business purposes for collecting data? Are there types of data that should not be collected? Do consumers have adequate notice and transparency about how their data is being used? And what happens if companies misrepresent, or don't fully disclose their privacy and confidentiality practices?

And on the competition front, we certainly heard that incumbent tech firms have access to large amounts of data through the existing product lines, and that there are challenges associated with competing in the AI development related to access to data, and access to resources and investment 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 development in a direction that protects or enhances market power. Certainly, the FTC will need to be vigilant in evaluating these issues as we pursue our joint competition and consumer protection mandate.

This will conclude the second panel, and we'll take a brief five-minute break. Many thanks again to our panelists, for taking their time and sharing their insights with us today. Starting at 2:35, Commissioner Bedoya will speak. And then we'll move on to our third panel, which deals with AI and consumer applications. Thank you again for tuning in.

Amritha Jayanti:

Hi, everyone. Welcome back. I'd like to turn it over to Commissioner Bedoya now for some remarks. Commissioner Bedoya, over to you.

Alvaro Bedoya:

Thank you, Amritha. I want to thank you, Stephanie, everyone at the Office of Technology who has spearheaded putting this event together. I want to thank all the staff that are moderating so ably today, and are working behind the scenes to get us up and out there. I also want to express my gratitude for following such a wonderful panel of people who I admire, who I've learned from, and who I've had the pleasure to work with. And so I'm really, really glad to be here.

I'm also glad to have voted out this 6(b) order to shed some light on the competitive dynamics at play with some of these most advanced models. I will probably focus less on some of the generative and LLM models that have been the focus of the discussion today. But I do want to touch on one note, perhaps as a bridge to a discussion of maybe a more boring subject matter, but one that I think is just so critical for regular people today.

Jonathan Frankle, who is an old friend and colleague, mentioned this trend that he's seen in some of these LLM models and generative models, which is a trend toward less transparency around the data that is being used in these models. And you can see it, if you look at the technical papers that are issued to accompany, these latest models, and you look at them from a few years ago, there's actually quite a degree of specificity around what data is being used. And nowadays, a lot of them have pretty noticeably shifted to little to no transparency around what data is being used.

I understand there's a variety of reasons for that, but I'll just say that from the standpoint of bias, of detecting and trying to suss out whether these systems work differently for some groups of people rather than others, this has consequences. Because knowing what data is used to train a system up, is extremely valuable for understanding how that system might function in the future with respect to bias. And so, I want to focus on that issue, on the issue of bias in AI systems.

As I previewed, I want to focus a little less on the generative and LLM models that have been discussed, not because I'm not interested, I obviously am, but partly because they've been the subject of discussion. But most importantly, because I think that people today are much more likely to encounter in their daily lives, a much more traditional automated algorithmic decision making system, than they will a generative or LLM system. And they are exponentially more likely to have critical decisions made about them by those older, more traditional algorithmic automated decision-making systems.

Today, I'd like to focus my attention and my remarks on just this one issue of bias in those systems, and focus on just one case, which is a recent Rite Aid settlement, which came out I think on December 23rd of last year. Terrific work by our Division of Privacy and Identity Protection, and our Bureau of Consumer Protection more broadly. I'll urge you to look up the complaint, look up the settlement, get the details. But in a nutshell, this was a case of the retail pharmacy using a face recognition system to try to identify suspected shoplifters who were entering its stores. And unfortunately, as we allege in our complaint, the system was flawed. It was, we alleged, profoundly flawed. Despite all of this, it was, we alleged, disproportionately deployed in areas that were plurality minority, in which the system was, we alleged, most likely to misfire.

The ways in which it misfired were pretty jarring. I'll just share a few instances here with you today. In one instance, a Rite Aid employee stopped and searched an 11-year-old girl because of a false match. This girl was so distraught that her mother said she must work in order to comfort her, this 11-year-old child who was falsely stopped in a store. In another instance, Rite Aid employees called the police on a customer because the technology generated an alert against an image that was later described as depicting, "A white lady with blonde hair." The customer they stopped was African-American. In countless other instances, people were stopped shopping for food, shopping for medicine and other basics. They were wrongly searched, they were accused of shoplifting, and in many instances, expelled from stores. Sometimes all this happened in front of their families, in front of their coworkers, in front of their bosses.

And so, why am I raising all of this now? What takeaways do I want people who are watching, businesses and their counsel who are watching to take away? Three things. The first goes to something that Corey mentioned in the last panel, which is we need to appreciate what's at stake here. He had this great line that you don't call a murder a theft. I'm not accusing anyone of murder here, but I do think that this is a reminder that, yes, this is about privacy. Yes, this is about rights. Yes, this is about fairness. Every single one of those words is applicable here. But it's also about people walking into a pharmacy to get their prescriptions, and being stopped, having the cops called on them, and being asked to leave because the secret face recognition system has wrongly identified them as someone they are not.

The decisions made by these systems, and zooming out more broadly from Rite Aid, affect our basic ability to live our lives with dignity, with fairness, to get the healthcare we need, to get fairness in terms of the apartments we rent, the jobs we apply for, in countless other areas of our lives. This technology has basic implications for our lives and our abilities to lead them. That's the first thing.

The second thing I want companies to take away, is that a company cannot buy a technology. A company cannot deploy it on its customers en masse, and when that technology misfires and misfires spectacularly, that company cannot turn around and say, "I am sorry. It was an algorithm. It was a black box. I didn't know. I am sorry. I am not liable." If you are a technology that is using an algorithmic decision-making system, and this is the third takeaway, to make critical decisions about people's lives in any way that may substantially injure them, you need to ask hard questions about how that system works, how it affects people, how it can hurt them, and you need to make sure that system works fairly, and is fair to the people it is used against.

In closing, I want to actually answer, if I may, the closing question that Krisha asked our moderator from the Bureau of Competition, from TED, asked to close the last panel, which I think is a really helpful question, which is what does success look like? And I'm going to give two answers to this. The first is focused on people. I think success looks like people controlling technology, not the other way around. Success looks like people feeling like they're in control of technology, people knowing when technology 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 startups having a shot at success, having a shot at reaching customers, because they put out good products that work, that people like, and not just because those products are put out by the most powerful company that has hit a trillion or more dollars in valuation. Success looks like companies succeeding on the merits, not because they are large.

And so with that, let me turn it over to Andy Hasty from our Division of Privacy and Identity Protection. And just say, I'm excited to watch our next panel, which Andy will introduce. Over to you, Andy.

Andy Hasty:

Thanks very much, Commissioner Bedoya. Hi everyone, I'm Andy. I'm an attorney in the FTC's Privacy Group, and it is my privilege this afternoon to moderate the last of today's three panels. Where the first panel focused on AI's underlying computing infrastructure and the second panel covered data and the role data plays in training and developing AI models, this last panel will explore AI at the consumer application level. Focusing both on what it takes to be competitive at this layer, and the risks that we all face as AI technologies take on an increasingly prevalent role in our lives. Joining me today to share their perspective are four terrific panelists. We have Atur Desai, Karen Hao, Conrad Kramer, and Ben Winters.

Atur is a technologist and a lawyer. He's currently the Deputy Chief Technologist for Law and Strategy at the Consumer Financial Protection Bureau, where he's served since 2016. Prior to his current role, Atur was a senior litigation counsel in the CFPB's Office of Enforcement. He holds a law degree from NYU, an MPP from the University of Michigan, and a Bachelor of Science from Cornell.

Karen is an award-winning journalist who covers the impacts of AI on society. Her work is widely read and cited. She's a contributing writer for the Atlantic, and previously covered China's technology industry as a foreign correspondent for the Wall Street Journal. She's an MIT graduate, and a former senior editor for AI at MIT Technology Review.

Conrad is the CTO of a new startup that he co-founded to build consumer software using AI. Before co-founding his current company, Conrad spent four years with Apple, where he built the Shortcuts app that comes pre-installed on every iPhone. And before that, Conrad co-founded Workflow, which Apple acquired in 2017.

And last but not least, we have Ben Winters. Ben leads the AI and Human Rights Project at the Electronic Privacy Information Center, or Epic, where his work focuses on AI and automated decision-making applications. Ben is also a senior policy fellow at AI Now, and teaches technology policy at the UDC David A. Clark School of Law. Ben has a law degree from the Benjamin N. Cardozo School of Law, a bachelor of science degree in communication studies from SUNY Oneonta, and has published work covering generative AI harm, discriminatory impacts of carceral technologies, and the need for transparency regulation.

Okay. We have a little less than an hour to explore a big and nuanced topic that presents a wide range of complex issues. Obviously, we won't be able to cover everything from all perspectives, but let's do what we can. I want to start with the competitive landscape, which I realized the second panel covered, or touched on anyway, but I think that makes a nice transition. Conrad, could you describe what the competitive landscape looks like from your perspective, with a company that's building an AI-based consumer application? What do startups need to be competitive here?

Conrad Kramer:

That's a good question. I'd say that primarily startups need access to models in order to build products. Because consumers don't download and run AI models directly, they typically interact with a product that embeds an AI model within it. And so, there are a few different ways that a startup who's building a product using an AI model can get access to a model. The first is paying an existing company that provides models as a service, and the second is training your own, which takes a lot of money and expertise, a lot of specialized expertise. The last is actually, you can use an open source model and customize that for your application. But the competitive landscape is interesting, because either you have the money, expertise and time, which takes a lot of time to train these large AI models, or you are relying on either open source or an existing model provider. That's the existing, I'd say, landscape in terms of, the main thing they need is access to models.

Andy Hasty:

Could you elaborate a little bit more on the considerations that go through your mind when you're deciding, are you going to build a model? Are you going to purchase or license one? Are you going to look to the open source? What are the factors that you think about?

Conrad Kramer:

I think the first is when you're evaluating a model, is quality. All the different solutions I mentioned have different qualities. The best models are currently available for purchase or for rent from the larger companies, versus the open source models are a little bit lagging behind on quality. The other trade-off to consider is cost. And so, the open source models aren't... the data is free or to get the model is free, but to run it isn't, you have to pay for compute to run the model. And so the cost scales from just paying for computation, all the way up to paying for computation and for a model developer to build and test the model. Those are two of the things.

The other thing is that quality can sometimes depend on which data the model is trained on. And because we don't really know, for example, for the closed source models, which data is being used to train them, you actually just have to guess and check. You have to try each one and see, maybe the data that it was trained on is more aligned with what I'm building, and so it performs better. Versus the open source models, you actually in most cases have visibility into the entire dataset, and so you can actually see what is being used to train the model, and you can actually use that knowledge to build a better product. And so, there are a lot of trade-offs there, I'd say, with models.

Andy Hasty:

You referenced quality, can you elaborate a little bit more on the metrics you're using to evaluate that, make this a little bit more concrete for us?

Conrad Kramer:

Yeah. I would say the standards for metrics to evaluate models are pretty rudimentary. Because currently, the gold standard for the industry is you have a list of questions, and you ask the model a list of questions, and you check the answers, and you make sure the answers match what you expected. And so when building a product, the way you would do it is you have your set of questions, and then you swap out the model underneath, and you see how it does on that set of questions. If it answers well, then it works well. But actually that is, because these metrics are very, very simple, it's actually a lot of qualitative understanding, of trying to figure out why the model performance is good or bad. And it's a lot of just guess and check. It's really hard to evaluate model performance generally.

Andy Hasty:

When you're talking about good responses, you mean accurate, or quick, or comprehensive, readable, all of the above?

Conrad Kramer:

Yeah. It depends. I think a lot of people are interacting with models in the case where they output text that is sent in a communication, and so the quality of the text that it might output matters a lot in terms of tone, information, things like that. We will probably see a lot of these large language models used for less communication, and more just take this and fill it into a calendar event, for example. And so, correct actually is domain dependent. It really depends on what you're going for. Correct, or good, or how to evaluate it depends on the use case. But in the case of writing an email for example, you would make sure that it's... check for concision, check for accuracy, check for a bunch of other things.

Andy Hasty:

Okay. And on the cost side, can you flesh that out a little bit more, when you're thinking about where and how you're going to obtain your models?

Conrad Kramer:

Yeah. We're in an early stage of the industry right now, where there's an explosion of new companies. And so some companies are going and training AI models, and those you can just basically pay a little fee, and you can send some information to them and they'll send you a response back. Then on the flip side of the open source model, you can download, on a website called Hugging Face, for example, you can download the weights for a model for free, and run it on your computer for free. The issue there is that if your computer can run it, these models are computationally expensive to run, so you need a powerful computer to run them. And so, people that are building these products have powerful computers, but not everyone in the world has a powerful computer. A lot of people still have smartphones as their primary computing device.

And so, the cost of that open source model is actually the cost to rent a cloud server, and put the model on that server, and to run it on behalf of users around the country. And so, there are a bunch of actually specialized startups that have started to do this as a service, where you give them your open source model and they run it for you. They're competing on costs. They're often cloud providers who are offering this service, and so they really are just trying to compete on compute costs. And so, I haven't actually looked into the pricing, but it's significantly cheaper than the hosted, expensive proprietary models.

Andy Hasty:

And one last quick question before I move on: what are you excited about here, and what gives you a little bit of heartburn?

Conrad Kramer:

Yeah. I'm really excited about just the potential for all these products to improve people's lives. Just because in particular, one thing I'm really excited about is a lot of paperwork burden on people. A lot of you interacting with a computer is filling out forms, or doing things that are otherwise rote or repetitive. I'm really excited to hopefully graduate from that era of computing, where we're not doing these rote, repetitive things. We can focus on the more human things. The thing that gives me heartburn, we're going to get to potentially later, is privacy concerns around people's data, and control over their data. Those are the things that I'm concerned about.

Andy Hasty:

Thank you, Conrad. Karen, I'm going to turn to you. Could you describe what consumers are experiencing on the ground today? What opportunities are these new models and AI technologies are creating for them, and what risks are you seeing attached?

Karen Hao:

Absolutely. Thank you for the question, Andy. I think to Conrad's last point, what I'm seeing, consumers are also excited about the possibility to integrate some of these tools into their lives, and automate or help them unlock their creativity, like talking with ChatGPT to get inspiration and ideas, or using Stable Diffusion to generate concept work that can then help them figure out where they want to go, whether it's an architecture, building that they're designing, or something like a poster that they're designing. I've also seen excitement with parents, and using these tools to engage with their kids on educational and interactive story time. But I think that on top of all of these wonderful and beneficial applications, there's certainly a huge amount of risk right now to consumers of generative AI. And the biggest issue that consumers face is this lack of transparency, which is a recurring theme throughout the day today. The two biggest sources of that lack of transparency, I think have been ambiguous or deceptive marketing, and then also obfuscation. I'll go into both of those.

But when it comes to ambiguous or deceptive marketing, one of the things that we know about large language models and generative AI, is that it has a problem with hallucinations. This is what we know from research, from scientific study of these models, from understanding how the underlying underpinnings of the models work. Because these models are trained on a lot of data, and then they're ultimately generating more data through statistical correlations, they're not actually extracting specific pieces of information for you from the web. It's a probabilistic completion of either pixels or text, or now we're getting into video, probabilistic video completions.

But the way that a lot of the companies that are developing these tools, like OpenAI, and Microsoft, and Google talk about this technology, they kind of paper over this hallucination

Karen Hao:

... problem. Like Satya and Adella has said about binging AI, this is just search just better. You will use this tool to get to the right answers.

And of course, there was that infamous case of a lawyer who then used open AI's technology to try to figure out whether he could get some assistive help on his legal research, and it ended up fabricating everything. Obviously that was, on his part, a gross oversight of his own professionalism. But on the other hand, whereas... Sam Altman has said ChatGPT can give you mirage of greatness, don't fall into that.

OpenAI's website promotes GPT4's ability to pass the bar exam, to take the LSAT. It has a partnership with a legal AI assistant startup called CoCounsel, which has advertised before that one of its clients is the California Innocence Project. So this is a repeated problem that we see, where there's become a disconnect between what consumers understand they can actually get from this technology, and what is actually the strengths and weaknesses that they can get from it.

And I just want to touch on one other example. This becomes really critical when we talk about sensitive contexts, and there have been a few studies that have studied hallucinations in the context of medical AI, for example. And there was one study in particular, really interesting from German researchers, that was looking at the ways that consumers might try to interact with ChatGPT to help improve their understanding of their own healthcare reports. So you could definitely foresee a consumer wanting to upload an image of an MRI scan for example, or an MRI report at length and say to ChatGPT, "Please summarize this for me in lay language." And what they found was they asked a bunch of radiologists to evaluate the summaries that ChatGPT was giving them, and many, many, many instances, the summary was just completely wrong in a way that would be harmful to the patient. So in one particularly egregious example, there was an MRI scan of a growing brain tumor, and ChatGPT said, "This brain does not seem to be damaged." So that's the stakes that we're talking about when there's this disconnect. And I think on top of this challenge of deceptive or ambiguous marketing is what Conrad mentioned, that there's two different classes of companies right now. There's the AI model developers, and then there are the consumer-facing companies that are taking the models and integrating it into a consumer facing product. And one of the challenges now is these supply chains have become so convoluted that consumers don't actually know ultimately what is the underlying model that they are interfacing with.

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.

A third big one is the increased opaque data collection. Every panelist I've watched today has touched on this. These systems would not exist if there were good privacy laws, or common sense privacy laws. They just wouldn't. In addition to scanning and scraping the internet, there are these exploitative turns on the relationships that companies have with users for years. You have Gmail for the last 20 years to have an email account.

Or you have to have it with work, and all of a sudden they're like, "Every time you use Google Docs, which you might have to use for school, or work, or makes it more convenient, you're helping us train our model. Or we're using all of your Facebook posts or Instagram posts." That is not part of the deal as a consumer. But once they have you entrenched, once they are a dominant market participant, that really creates an issue. And connected to that is the intellectual property rights, which I won't delve into too much right now.

One big social harm that AI is having is the exacerbation of climate change effects. We don't even know how bad this is yet, but we do know it is bad. It's unconscionable that these are being rolled out to the extent that they are. In spite of this, there is insufficient study of the carbon footprint of all large language models. But there is one study that Sasha Lucioni did, and it was [inaudible 02:43:41] that showed that a large transformer model in building that and getting that set up and built for prime chime has the same carbon output as five cars. So in a moment as we are hurdling towards climate change, this is just a thing that we are adding to everything just because. The sixth thing I'll say is a data security risk. Amba touched up on it in the last panel, but it's really just an affront to the concept of data minimization. The use, the building, all of this stuff is really difficult for data security. The more data you take in, the more vulnerable you are to data breaches. Just period, especially when there's no reason for it.

There's the labor manipulation, theft and displacement problem. The discrimination problem that we see on all sorts of automated systems. Like Alvaro mentioned at the beginning, that the facial recognition algorithms for example. But also we see there's been a lot of studies on this that generative AI systems specifically really entrench discriminatory stereotypes that we've seen for a long time. You type in “doctor,” you're going to see a white guy. You type in “homemaker,” you're going to see a woman. It is not advancing anything, it is keeping us stuck in the past.

And then the last but not least, there's the impact on market power and concentration that we touched on a lot throughout these different panels. So I'll stop there. I know that was a lot already, but excited to talk more throughout it.

Andy Hasty:

Yeah, thanks Ben. That is quite a landscape. I'm going to turn to you, Atur, let's maybe focus on the financial space. What is the CFPB focused on? Can you tell me what issues you're seeing arising from AI-based applications, AI decision-making in the financial sector, and how you guys at the CFPB are approaching them?

Atur Desai:

Sure, thank you, Andy, and thank you FTC for organizing and hosting this event. One piece of housekeeping to get out of the way before I dive in, the disclaimer. Remarks I make today are my own, they do not constitute legal interpretation, guidance for advice of the Consumer Financial Protection Bureau or myself. Any opinions or views stated are my own, and may not represent the Bureau's views.

With that out of the way, so to answer your question. So what is the CFPB doing with regards to AI? Really, the short answer is a lot. Can't possibly get into everything, so in the interest of time, I'll focus just on a few high level points. First, discussions of AI are often accompanied by a lot of marketing or hype surrounding how novel or revolutionary it is. But in reality, complex models have been used in consumer financial markets for a long time. For example, in credit underwriting and credit scoring. Modern deployments of AI are of course increasingly complex and powerful, and rely on incredible amounts of data, but complex models have fundamentally been around in the markets we oversee for a while.

And I think that's a really important point to highlight. And that's because something that oftentimes feels underappreciated in discussions of AI is the fact that a robust set of preexisting laws, including federal consumer financial laws, exist, and that these laws apply to AI all the same as to the technologies or processes being replaced. And this is a point that the CFPB is quite vocal about. Stated simply, there's no AI or fancy technology exception to the laws that CFPB enforces.

So one thing the CFPB is doing is making this point clear. The fact that you're using a complicated AI model, or that you may not understand why your model is reaching the conclusions it is reaching does not diminish your legal obligations under our laws or consumers' rights. For example, the CFPB issued two circulars related to credit decisions and the requirement to provide consumers with an adverse action notice under the equal credit Opportunity Act, or ECOA.

And at a very high level, in these circulars, the CFPB made clear that companies relying on complex algorithms must provide accurate and specific explanations for denying credit applications, and that companies are not absolved with their legal responsibilities when they let a black box model make lending decisions. So really at the end of the day, if using a technologically complex model means that a company cannot comply with its obligations under federal consumer financial laws? They really shouldn't be using that model.

So another point that I think is important to the conversation, one that I'm hearing a lot of today, is really about the inputs. As we know, models, they rely on a critical amount of data, often surveillance data, so it's important to consider how companies are obtaining this information and how it's moving from one party to another. And the CFPB is evaluating these very issues. For example, we issued a request for information related to data brokers, and we also announced the launching of fair Credit Report Act rule banking.

So you asked also a little bit about approaches to AI enforcement, so let me just talk about this at a high level. Relevant to our AI enforcement work is capacity building. The core work of the CFPP, ensuring that markets for consumer financial products are fair, transparent and competitive relies on a deep understanding of how the markets we regulate work. And while some of these technologies may not be entirely new, they are increasingly complex and can involve different business models or other considerations.

So one thing that is important is that we ensure that we have people with diverse perspectives and skill sets in the room, and this is a focus area for the CFPB. In 2022, we started our technologist program. And what this specifically means is that we began a program to tightly embed and integrate folks with technical expertise within our supervision and enforcement teams. These are data scientists, AI ML experts, design experts amongst other technical staff. So CFPB is putting a focus on making sure that we're building these interdisciplinary teams, so we're not just approaching problems from the perspective of lawyers, or economists or other professionals, but rounding it out with technologists who have deep knowledge of the markets that we're overseeing.

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 consumers interpret them, and what should companies be doing to ensure that the safe treatment of their users?

Karen Hao:

Yeah, I think the term AI safety in particular is really interesting. I was at an AI research conference at the end of last year called Neural Information Processing Systems, or NeurIPS for short. It's one of the largest AI research conferences that happens each year. There was actually a session between machine learning researchers about how they name things and whether or not this was becoming problematic, because so much of the naming originally started as signaling to one another within the machine learning community, within the AI community. And then as things started being consumer facing, and then as media companies start reporting on these terms, then suddenly it interfaces with preconceived notions of what that concept is, and AI safety is one of those.

So, AI safety originally... I'll read you the original definition. It was from this paper written by Dario Amodei and Chris Olah, who are two co-founders of Anthropic. And they defined in this paper in 2016 that AI safety is the problem of accidents in machine learning systems defined as unintended and harmful behavior that may emerge from poor design of real-world AI systems. And then they specify that it is not related to privacy, security, fairness, the economic impacts of AI or its military implications.

Basically, to summarize that, they were talking about AI safety being we need to think about methods and techniques for preventing rogue AI that might lead to existential threats. And specifically, they were contrasting this with more urgent, present-day harm, saying this is actually a different category of problem.

But of course, when we start talking about AI safety in the public domain, safety has a totally different definition than public domain. And so now it's become a really clever marketing trick for companies to lean into the common interpretation of the term, which is that... It is related to privacy, security, fairness and the economic impacts of AI or its military implications, and that is a really big disconnect that ultimately allows companies to continue doing a lot of things that are not necessarily great for the consumer in the long run.

So, when I interviewed Dario in 2019 when he was at OpenAI at the time, he made this really critical point that I think is important to understand, where he was saying that the reason they had... At the time, OpenAI was in the middle of debating whether or not to withhold certain models, like GPT2, the predecessor to ChatGPT, because it might pose a danger in terms of mass automating disinformation. And ultimately after withholding it they get blow back, so then they decide to release the model.

And the way that they rationalize this was he said to me, "Obviously, misuse of these models is not good. But a language model is a lot less powerful than an AGI. I'm very worried about language models 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 I just wanted to raise one more thing. Which is after this 2016 paper, there was this amazing paper from Deborah Raji and one of her co-authors, where the original paper was called *Concrete Problems in AI Safety*. And then she wrote a paper called *Concrete Problems in AI Safety, Revisited* in 2020 where she talks about how... In the original paper, there's this idea that AI safety is about preventing the AI itself from misbehaving. And she said, "At the end of the day, AI doesn't exist in this other space. It exists in our world." And if we're worried about building systems that are going to exploit, single-mindedly head towards a goal and create harm, we already have that today.

And that's the entire AI industry. Because the way that the AI industry works, it's already focusing on a singular goal, AGI or commercialization, whatever it is. And it's creating these vast tracks of harms, whether consumer harms, or labor harms or climate harms along the way in pursuit of the singular goal. And so she says, "It is not just the actions of an AI agent that can produce side effects in real life, basic design choices involved in model creation and deployment processes also have consequences that reach far beyond the impact that a single model's decision can have. In reality, for AI systems to even be built, there's very often a hidden human cost."

Andy Hasty:

Thanks very much, Karen. Conrad, I want to turn the same question over to you, but maybe... Karen covered a lot on safety, so please react to that. But you mentioned privacy early on as a source of heartburn. And you're a co-founder, so I'd like to hear from a company's perspective about how you think about these terms and what you're doing in practice to carry them into effect.

Conrad Kramer:

Yeah, yeah. So I think the safety debate is an interesting one. I think that as a company, my takeaway of what OpenAI's safety strategy that I'm taking into our company is slow, gradual deployment. I like that. The keyword is slow, because you actually, you need to put something into production to see how they work and how they're being used, and then have a process for responding and reacting to that, and iterating.

But on privacy, the labeling is a huge problem. So labeling on privacy is a problem, but also labeling on privacy, privacy for me comes down to the most important question when you're using any product.

Where is the data? Is it on a computer that you own, or is it on a computer that someone else owns? And first of all, even knowing that is challenging with some of these products.

And then the second question is, regardless of whether it's on your device or on someone else's, what is that company allowed to do with the data? Even today, it's very hard to figure out. Maybe you could use ChatGPT to summarize the privacy policy, but it's actually, it's fairly dense to read and understand where your data is being used, how it's being used.

And back to supply chain issues. Companies can actually send their data to another company. For example, to do analytics, and then that company could be doing training with that data. There's already a supply chain problem with data acquisition, and we actually saw some of that with the GDPR, where there's actually some transparency regulations where all of the sub processors for all the data companies have to be listed. Which I think is great progress, but I still think that from a consumer perspective, it's fairly opaque of what is this company allowed to do with my data.

But that speaks to another important thing we're trying to do when building our company, and I would hope that other companies do for their consumers. But it's transparency. So those decisions that we make, and the decisions of where the data resides and what you do with it actually depends on what you're building. If you're building a product that summarizes the meeting notes and sends them to the relevant people at a company, that system necessarily has to have data from a bunch of people and mix it. But if you're having a system to compose messages for yourself, that data doesn't need to be shared or mixed, and ideally it would be on your device. So transparency is this key.

I think the other point is, and this is where it gets really tricky, is control. Some products, if you start using it and they don't make their trade-offs clear, and you're using a product? You actually don't have the ability necessarily to switch where that data is being used, it's built into the product. Some products give you some control over how your data's used. So the iPhone for example, you can go into privacy settings and turn off microphone access, and contact access and all of that. And that's a really good example of good control in the hands of users. And some companies even have this. Google allows you to opt out of some of their tracking. But yeah, control is really important.

And so I think the best systems are ones in which the data is as close to you as possible, so it's on your computer if possible. And one in which you know when the data is being used, and so when and how. And so that's the product I'm at least trying to build, and I hope that other entrepreneurs feel the same way. Because yeah, it's really important that the people know how their data is being used.

Andy Hasty:

Thanks very much, Conrad. Ben and Atur, I want to give you a chance to weigh in on that question too. But in the interest of time, I'm going to tee up the next related question, and that is what changes do you think might improve consumer privacy, or reduce the safety risks that Karen and Conrad have been flagging? And while you're thinking about it, what changes might make this space more competitive? I'll start with you, Ben.

Ben Winters:

Great. Yeah, I think what I was going to say to that question and to this question work really well together. In terms of the term AI safety, which is what the leaders of the biggest companies are talking about a lot, are bringing to congressional leaders and are talking about in really high profile context. The fact that we have to have a bunch of people talk about exactly what that means, and we still have no idea what it is, is part of the exact problem with a lot of the AI industry right now. There is this man behind the curtain problem, where there is this needless overcomplicating of things. And so I think that to the extent that the harms can be more precisely considered and taken one at a time, that's okay.

As I illustrated in my very long rambling answer, it's not easy to concretely say that one by one. But it's also really unhelpful to obfuscate all of that combined into this really weird existential harm concept. And I think that often the framing of the existential harm in the AI safety is really just serves as a market distraction and as a political tool to scare legislators and regulators into not focusing on the current harms. Because as Karen said, of course if there's an asteroid coming tomorrow, that is the most important thing. You're not going to worry about the other stuff.

But in the meantime, they are not doing these very simple, basic things. One thing that would help improve competition and just the general consumer experience is putting the burden for transparency, and communication, and just for lack of a better word, doing the right thing on the company and not the consumer to investigate something that you cannot possibly even understand if you spend half your day doing it.

So I think that there is a place for legislators and regulators to force that, there are good mechanisms like audits and impact assessments. There are different ways we can make that better or worse. But a lot of it will have to do with the norm of just respecting the customers, and valuing not just the data that you can buy from the New York Times for $30 billion, but everybody's data.

I'll stop there and hand it off to Atur, but I think most of it really can be improved, not solved, by getting the burden of the transparency on the company rather than the individual.

Atur Desai:

Great, thanks, Ben. I definitely have thought on the previous question, so happy to return there at some point if there's time. Let me focus a little bit on this from the perspective of a regulator. I just want to start from the jump by saying breaking the law should not be a company's competitive advantage. So I think as a first principle as we start to think about this, companies really should adhere to their obligations under the law. If a company can't comply with laws like federal consumer financial laws because their technology is to complex or otherwise? Then they really shouldn't be using that technology. I knowI already said this, but I think it bears repeating. But if nevertheless, a company's use of advanced technology runs afoul the law, then I think it's really important to ensure competitive marketplaces, that we enforce these laws. And what does enforcing the law mean with respect to emerging technologies?

So one aspect of it involves building interdisciplinary teams that includes technologists, those with deep knowledge of the technologies and business models at issue. And this will allow us to more efficiently and effectively detect, investigate, and if necessary, litigate potential violations of law.

I think it also involves thinking carefully about remedies and remedy design when violations of law are found. It is critical that breaking the law is not viewed simply as a business decision, where the cost of legal violation is considered just the cost of doing business.

And this is another place where I think ensuring having a broad set of perspectives in the room, when you think about how do we best design remedies so that we can protect consumers, and ultimately, ensure their marketplace is very competitive, is really, really foreign.

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 Winters:

Sure. I think the one thing I'd want viewers to take away today is really that none of this has to be inevitable, and none of it has to be as complicated as the largest companies really want everyone to feel. And that goes from users, to the people they're selling it to, to legislators and regulators. And when I'm saying nothing is inevitable, I am not just talking about in terms of the actual lack of likelihood that any of these existential risk things take place, which is a whole other conversation, but just the fact that AI has to be integrated into every single tool we have is not inevitable. I think that we've seen a lot of ups and downs in different markets and different trendy technologies, and at some point, things will average out, but right now, we are in this place where there's a lot of energy and capital being put into making everyone feel like AI is everything and everywhere.

So I think just one thing is just to try to empower consumers, regulators, legislators, writers, that you can understand it and push back about it. And then the other part about it that we've touched on it, and I think this audience particularly is aware of this, but I think just that the other laws still exist. We have civil rights laws, consumer protection laws, fair competition laws, and while we do need a comprehensive baseline privacy law, we need laws that ban specific, really just unconscionable uses of AI. We can, a little bit, little by little, tweak and improve the status quo with the laws that are on the books.

So I will stop there, but thank you so much for having us.

Andy Hasty:

Thanks, Ben.

Karen, one takeaway for the audience?

Karen Hao:

Yeah. I think the biggest takeaway is to, kind of building on what Ben said, to really question what companies say and how they message things to us, not only in terms like AI safety or in the other kind of marketing that they use, but also in the way that they frame what is good for us, like the idea of deploy now and iterate later. It doesn't feel like we're living in a democracy right now. If a company just gets to decide, opening AI, launched ChatGPT within two to three weeks, on a whim of a decision, based on competitive pressures, and then suddenly, it bursts forth and we're all living in this new era and we all have to grapple with it, I don't feel like any of us had any kind of agency, any kind of democratic governance over that decision. And so, always constantly questioning what companies tell us, and also actually realizing that we should be holding them accountable, living in a democracy, and demanding more of them than just kind of retroactive excuses.

Andy Hasty:

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 case that I can point folks to is a 2022 enforcement action against the company called Hello Digit. And the case involved, amongst other things, faulty claims about a company's savings algorithm's, ability to avoid junk fees, specifically overdraft fees. So I just wanted to make sure that we're kind of well prepared and well accustomed to hearing a lot of just marketing jargon being used to over hype technologies or any product's capabilities and are well equipped to handle those.

So I'll move on to closing, and I really wish I had something to build some suspense with, but I'm fairly certain that what I'm about to say will come as no surprise, since then I'll just be brief. If there's one point to take home, I think it is that there's no AI exception to federal consumer financial laws. Importantly, these laws place obligations on companies, and these obligations are not diminished because a company elects to use AI that they may not fully understand, but believe is too complex or too fancy. Thanks.

Andy Hasty:

Thanks, Atur, and thank you, Ben and Karen and Conrad, for your time today and your thoughts and for this great discussion.

I'd like to end with a few reflections, picking up on some of the themes that you all raised. One theme is trust building and deception, with you noting that companies in the space may be marketing their applications, using labels like AI or safety or privacy enhancing, that may not actually square with the company's practices or with consumer's understanding of those terms.

Another theme is just the breadth and the magnitude of potential harms here and, I think, the multidimensional nature of the risks. These technologies, they are powerful, and for as much potential as they hold, when they meet consumers, they can do a lot of damage, ranging from the known and the concrete, I think, some of the things Commissioner Bedoya mentioned in the preamble to this panel, like discriminatory decision-making some other known concrete arms, extraordinarily convincing fraud schemes and errant disclosure of consumers' private information, to the more abstract and the uncertain and the existential, some of the things I think Karen pointed out, all of which can vary across 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 deployed to hundreds of millions of people through consumer applications that expose us to a vast risk spectrum, and I think I heard several of you emphasize an urgent need to change incentives, to shift burden, so that providers do much more to mitigate this risk pre deployment.

And I think with that, we will go ahead and conclude this third and last panel, and we'll take a short break, starting at four o'clock this afternoon, Eastern Time. Bureau of Competition Director Henry Liu and Bureau of Consumer Protection Director Sam Levine will provide closing remarks. Thanks, everyone, for tuning in, and see you shortly for the final segment.

Amritha Jayanti:

Hi, everyone, welcome back. We are back from our final break. In this last segment, we will hear from our bureau directors. First up, I'd like to introduce our Bureau of Competition Director, Henry Liu, for some closing statements.

Henry, over to you.

Henry Liu:

Good afternoon, everyone. What a fantastic summit. Thank you, Stephanie Nguyen and the entire Office of Technology, for organizing this event. It's been really great to engage with such a diverse group of stakeholders to discuss AI and the many layers of technology related to it. Hearing panelists cover topics such as the potential for dominant firms to gain control over key inputs, and to use data without consent, it's a reminder of the increasing significance of AI tools. As we close out this summit, I'd like to share my views on the Bureau of Competition's role in this conversation.

AI's rapid adoption and widespread use, it's brought on many exciting possibilities, of course, ranging from improving medical diagnoses, to identifying national security threats. But there are potential risks. Too often, AI tools have been used to limit opportunities and prevent access to critical resources. To guard against potentially illegal behavior, while also allowing the benefits of AI to flourish, it's extremely important for the antitrust agencies to be vigilant to protect competition.

In the Bureau of Competition, our role in this new AI powered role is clear. We will police illegal behavior and vigorously enforce the antitrust laws, just as we have done in other industries. And the cost of inaction can be outsized in AI markets because excessive market power can distort the path of innovation and discourage future investment in AI research and development. So what principles will guide the Bureau of Competition's enforcement priorities for AI markets? I'll mention a few.

First, as Chair Khan and the other commissioners said earlier today, fair, open, and competitive markets should be the hallmarks of AI. Antitrust enforcement and competition policy have long helped create the conditions for fair and honest competition. As we heard from several panelists this afternoon, when dominant firms control over key inputs like computing power, cloud storage, semiconductors, talent, and even data, it creates a risk of excessive prices and coercive terms. To promote more open markets, the bureau will continue litigating cases that address conduct and acquisitions, which, if left unchecked, might result in increased month's market concentration in various layers of the AI stack.

To give you an example, in June 2021, the FTC sued Broadcom, alleging that it illegally monopolized markets for semiconductor components, or chips, used to deliver television and broadband internet services. Then in December 2021, the FTC sued to block Nvidia's proposed $40 billion acquisition of chip design provider Arm. Protecting competition for inputs like computer chips and processing hardware will help propel AI development.

Most recently, the commission empowered bureau staff to ramp up enforcement by authorizing an omnibus resolution, which allows us to use compulsory process in non-public investigations involving products and services that use or claim to be produced using AI. The omnibus resolution will streamline our ability to issue civil investigative demands, which are a form of compulsory process similar to a subpoena, in investigations relating to AI.

The second principle, and along with protecting fair and open markets, is that AI tools should support opportunities for new businesses to compete. According to a survey by the Small Business and Entrepreneurship Council, many businesses are already using AI tools to improve their efficiency and save billions annually. From marketing and sales, to human resources, AI tools, of course, have a wide range of capabilities. With such broad applications, there, of course, often significant downsides attached. For example, small businesses that want to use AI tools frequently find it cost prohibitive, in part because they also need to acquire hardware, software, and skilled professionals. Ensuring that markets for these critical inputs remain competitive can support a diverse ecosystems of startups working on AI projects.

Now, keeping markets competitive for all businesses will require that the bureau continue using all the tools in its toolbox to vigorously enforce the antitrust laws. One important tool is Section 5's Unfair Methods of Competition Authority. In November 2022, the commission issued a policy statement on the scope of Section 5 Authority, and in that statement, the commission made it clear that it is congressionally mandated to identify unfair forms of competition, including those where small businesses were threatened by oppressive, more powerful rivals.

More broadly, the FTC's Unfair Method of Competition Authority empowers us to protect the public against powerful firms that unfairly use AI technologies, in a manner that tends to harm competitive conditions in AI or AI adjacent markets. The commission must remain vigilant to protect the public against such threats, and to ensure that competition thrives in AI input markets, as well as in markets impacted by new and emerging AI applications.

With these enforcement activities, the bureau and the commission are doing what they have done for over a century, that is to keep pace with new markets and ever-changing technologies. Our vigilance has been in no small part, due to our office of technology and their in-house expertise on all things AI, and I look forward to our continued joint work.

With that, I'll pass it back to Amritha.

Amritha Jayanti:

Thanks so much, Henry, for those great remarks.

Now last, but certainly not least, I'd like to introduce our Bureau of Consumer Protection Director, Sam Levine. Sam, over to you.

Sam Levine:

Well, thank you, Amritha, and thanks to everyone for joining the FTC's 1st Technology Summit on Artificial Intelligence. I've been so impressed by the thoughtful in-depth conversations we've had around topics ranging from cloud infrastructure, to privacy, to AI fueled fraud. And I'm grateful to our Office of Technology for organizing this important event.

Although formed less than a year ago, OT is already having a huge impact on the day-to-day work of our agency. In the bureau I lead, the Bureau of Consumer Protection, OT technologists are partnering with us on dozens of matters across many offices and divisions, and they're helping to ensure that our enforcement and policy work can meet this moment. And let me be clear, this moment is a unique one. Chair Khan spoke earlier about the last major inflection point driven by emerging technology, the dawn of Web 2.0 in the early 2000s. She noted how that era began with promise, with exciting new technologies and applications connecting people and expanding opportunities. But two decades on, we see a tech ecosystem that has concentrated private power, in the hands of a small number of firms, while entrenching a business model built on constant surveillance of consumers. Chair Khan and her remarks stressed that we need to learn from our experience in that era. And I could not agree more. This learning is not an academic exercise. As we chart our course in confronting AI related harms, we must engage with the history of how we arrived at this moment.

A generation ago, in the year 2000, the FTC issued a major report, finding that self-regulation on the then nascent internet was failing to protect consumers' privacy, and the commission recommended that Congress pass privacy legislation. Yet less than 18 months later, with the change in administrations, the FTC reversed itself, finding that, "It is clear that industry will continue to make privacy a priority," and warning that legislation would be premature and could hold back the growth of the internet. Instead, the FTC announced an initiative to ensure privacy policies were posted and honored, and to encourage industry self-regulation.

In my view, this reversal was a serious error. It is now apparent that industry did not actually make privacy a priority. In fact, at the same time the FTC was expressing confidence in self-regulation, Google began exploring how it could mine search queries for behavioral insights on its users, laying the groundwork for a transformed business model, and ultimately, a transformed internet.

A decade later, the FTC would reverse itself once more, calling in 2012 for Congress to pass legislation. But by that time, many of the harms Chair Khan described earlier were already entrenched, and industry opposition to meaningful regulation had consolidated.

This history is very much top of mind, as we confront the emerging threats and opportunities created by artificial intelligence. Congress has entrusted us with ensuring the markets are fair and competitive, and we will not be sitting on the sidelines.

In the Bureau of Consumer Protection, we've been staying busy. We issued a major report in 2022, months before generative AI became a hot topic, warning about inaccuracy, bias, and privacy abuses fueled by AI. In 2023, we began routinely issuing guidance on how our authorities apply to emerging AI technology, making clear that we were not going to wait and watch while harms accumulate.

And as we've talked the talk with our guidance, we are walking the walk with our policy and enforcement strategy. We've now required, across multiple cases, that models, algorithmic models, trained on illegally collected data be deleted. We've brought lawsuits against firms that defraud the public by claiming AI can make people rich. We've partnered with the Office of Technology to launch a voice cloning challenge to confront new forms of impersonation fraud, and we've proposed a 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 farms must either take steps to ensure their AI tools don't harm consumers, including 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 tool is our people, our multidisciplinary teams of world-class attorneys, economists, investigators, consumer ed 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 a 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 ensuring that AIs 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 yours.

And with that, we come to the end of our FTC Tech Summit on AI. Before we fully close out, I wanted to communicate a few more thank yous. Firstly, thank you to the staff across FTCA's Office of Technology, 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 Stephanie Nguyen for their remarks earlier today. Thank you to the panelists and moderators, of course, for a phenomenal set of discussions. And finally, thank you to those in the audience for tuning in.

Today, we heard a vast range of excellent points about the real opportunities and risks of AI, with a look across the layers of the tech stack, from chips and cloud, to data models and consumer applications. As Chair Khan and many others mentioned throughout the event today, there are no exceptions to the laws on the book for AI, and our agency is keen to understand the full range of experiences across the marketplace, from startup founders to consumers, to small business owners, to better understand the role of enforcement in ensuring a fair and competitive marketplace. And we're taking action where possible, as Sam and Henry highlighted.

Our conversation today captured a snapshot of AI and its impacts as they are realized today. However, we know that this space is rapidly evolving. To that end, our Office of Technology, alongside with agency, bureaus and offices plan to continue to facilitate conversations and dialogue to build awareness about the consumer and competition dimensions of this emerging technology.

Thank you again to you all in the audience for participating in this important discussion today, and with that, we'll close out. Have a great evening, everyone.