Protecting Privacy in the AI Era: Data, Surveillance, and Accountability

On Wednesday, June 26th, 2024 the Center for AI Policy held a briefing for House and Senate staff on **Protecting Privacy in the AI Era: Data, Surveillance, and Accountability**.

The Center's Executive Director, Jason Green-Lowe, moderated a discussion between a panel of esteemed experts:

- Ben Swartz, Senior Technology Advisor at the FTC
- Brandon Pugh, Director of Cybersecurity and Emerging Threats Policy at the R Street Institute
- Maneesha Mithal, Partner at Wilson Sonsini and co-chair of the firm's privacy and cybersecurity practice
- Mark MacCarthy, Adjunct Professor at Georgetown University and Nonresident Senior Fellow at Brookings

They discussed privacy violations, AI's impacts, the proposed American Privacy Rights Act (APRA), and more.

What follows is a full transcript of the discussion. The transcript was generated safely by AI with human oversight. It may contain errors.

To watch a corresponding video recording, see this video on the Center for AI Policy's YouTube channel.

Full Transcript

Jason Green-Lowe | 00:07.065
All right. Welcome, everybody. This is the Center for AI Policy's panel on privacy, data, and accountability. Hot topics these days.

We are coming into this issue as people who are deeply concerned about some of the catastrophic risks posed by highly advanced general purpose AI. That AI is enabling what I think of as third generation privacy threats. What do I mean by third generation?

So for the last couple of decades, we've had first generation threats just based on data entry. If you keep a database in every office, in every school, in every police station,

some of it's going to leak, right? You just have your birth date, your social security number, your address everywhere. That information is not securely stored. There's a report of a leak every month or so. So some of that data is going around.

The second generation of privacy threats as we see it is machine learning. Right, it's powered by the ability to make inferences about the data that's leaking about you and going around. There's the famous story about the family that found out they were pregnant based on purchases being made at Target. Nobody ever typed in “I'm pregnant” into a database, but the information was nevertheless kind of shared.

The third generation comes when you don't even need a human deciding to run a machine learning model on the particular question of whether target purchases mean that you're pregnant. What happens when these kinds of shrewd guesses become part of the general background of what's happening all around us, right? Someone had to sit there, you had to spend some engineer's time thinking, oh, this is a good idea. I want to find out if this customer is pregnant. But what happens when that happens automatically, you know, hundreds of times a day, just new models making, asking their own questions, right? We see agentic AI being advertised by some of the big tech companies right now. I'm not entirely sure those advertisements are accurate. It might still be a few months away. But that doesn't give us a whole lot of time. This is coming down the pipeline. We're going to be looking at AIs that are pursuing goals more or less autonomously. And one of those goals is going to be to crunch your data and draw inferences about you, including possibly whether you're telling the truth, what you're thinking based on your brainwaves, what you're typing based on the sounds made by your keyboards, even if you're not talking to that AI, what your facial expressions on the security camera suggest.

So I want to facilitate this discussion about privacy because we're alarmed by the progress of this artificial intelligence. From our point of view, it's under-regulated, and the public needs to be more prepared to deal with these threats. But I’m an AI expert, I am not a privacy expert, so that's why I've got these four distinguished panelists. I'm mostly just going to be asking them questions and learning from them, and I encourage all of you to do the same. I've got a few questions to kick us off, and then I'll turn the floor over to you so that you can ask questions.

On my immediate right, Mark McCarthy is a professor at Georgetown University. He's a fellow at the Brookings Institution. He's got a book out on regulating digital technologies. He's worked for the Post. He's the senior vice president for public policy for Visa at one time, and has worked on the House Committee on Energy and Commerce and for

OSHA. So an extremely long and varied list of high level experiences, thinking and teaching about privacy.

On his right is Maneesha Mithal. She is an expert from the FTC and now at Wilson Sonsini. She was running a team of 40 lawyers, all of which were investigating privacy violations and identity protection. She has been a delegate on privacy to the OECD and to APEC. And this is only her first lecture on privacy law today. She'll be headed off to another one in New York this afternoon. Sorry if that's a privacy violation.

On her right is Brandon Pugh. Brandon is a senior fellow at the R Street Institute. He's running a team there of experts on tech and privacy issues, and is a national security law professor with the US Army Reserve and the Judge Advocate General's Legal Center. So he has been teaching army officers on how to prepare for kind of legal proceedings within the armed forces on some of these topics, and has also been an elected and appointed official in local and state politics, a fellow at the FBI, and the managing editor of the Journal of Law and Cyber Warfare.

On his right is Ben Schwartz, who is a senior technology advisor at the FTC. And before his time in the FTC, he was one of you. He was working for Senator Padilla's office as a staffer on digital platforms and data privacy. But that's his second career. Before then, he was a senior software engineer at Twitch and Google, where he took a lead role in promoting corporate social responsibility.

So I'm very excited to have these panelists here. I think we'll have a very interesting conversation together. And I think we'll just get started. If for any reason you can't answer a question because you don't know or because your professional responsibilities prevent it, feel free to either pass or reframe it so that you can answer the question.

And the first one is just what's a historical example of a major privacy violation that you find particularly concerning and why does it bother you? Let's go left to right and then we'll come back right to left for the second question. Okay.

Mark MacCarthy | 06:05.304
Good. Thanks for having me, Jason, and thanks for having an audience of bright and energetic Hill people eager to find out more about AI and privacy.

I think we should come back to your premise that there's a third generation of privacy intrusions that we have to worry about with artificial intelligence. I think that's one that deserves a little bit more examination. My sense is that it's a little bit more like more of

the same, machine learning and machine learning and machine learning. But I don't think we're at the point yet where these systems have become autonomous and are acting on their own and are running out gathering information about people with no direction from humans.

But to answer the question on historical events, I think the one that jumps out at me, and it's an old one, is the Cambridge Analytica problem from 2014. Facebook had a reasonably open policy for researchers at that time. a guy sort of masquerading as a real researcher, did a study, a profile study of people by having them play a game and answer questions. And using that, he had maybe 250,000 people who answered the questions. And he could build a model that predicted pretty accurately various psychological characteristics and political preferences and so on. And then he gained access to 50 million Facebook users, and was able to profile all 50 million of those Facebook users using the model he had constructed, and then he made that available to political parties for the purpose of political advertisements.

That strikes me as an extraordinarily lax privacy breach which Facebook has fixed in its particular form, but it gives you an example of the kinds of problems that can arise. When you have access to large amounts of social media data, the inferences about your most intimate characteristics are easily ascertainable from your behavior on Facebook and other social media platforms. You may think that you've told nobody that you're gay, but your behavior on Facebook tells people that you are. You may think you've told nobody what your religion is or your political party. But your behavior on Facebook reveals that. And that's all machine learning that does that kind of inference. I think those are the kinds of problems we still have to wrestle with. It was 2014 when this happened. It's 2024, a decade later. And we still haven't figured out a good way to control that kind of abuse.

Maneesha Mithal | 08:49.849
So I really appreciate being asked this question because it caused me to think about some of the cases we brought when I was at the FTC. And when I was there, we probably brought over 150 privacy cases involving things like Cambridge Analytica, data breaches, some egregious privacy practices.

But when I think about what really kind of gives me a visceral reaction, it may be some of the smaller cases. So for example, we brought a case against a company that was allowing people to put stalkerware on their domestic partners' mobile phones. We brought a case against people who were basically soliciting revenge porn for inclusion on their site. We brought a case against a company that leased out computers, but

without telling consumers, they turned on their cameras. And so they had employees in their offices kind of able to see what was going on in people's bedrooms.

And so I think what all of these... matters have in common is that there's a… I know the FTC talks about commercial surveillance, but I think it's the surveillance aspect of these practices. The fact that companies were stalking, following, allowing the stalking, following of individuals without the individuals even knowing about that. So I think that if you kind of fast forward to 2024 and AI, I think that's a concern that continues, is this kind of covert surveillance and ability to gather all sorts of data about consumers and profile consumers. So a lot of the same concerns that animated the privacy issues continue to animate on AI.

On the revenge porn example, I think one of the things that AI has allowed to be souped up is this idea of deep fakes and the fact that, you know, not only revenge porn, but people could kind of use images without consent, you know, superimpose different bodies on faces. And I think we were talking to some of our clients who were saying that, you know, by 2030, we may not be able to trust what we see, you know, in social media or photos or pictures or voices. And we need to build some consumer education around that. But those are the types of things that... I think the threats that I think we saw from traditional privacy enforcement that are continuing into AI, and I look forward to talking about that as we go on.

Brandon Pugh | 11:15.738
Yeah, I'm glad to be here. Jason, thanks for having me, and congratulations. I guess this is the Center's first anniversary, roughly, so that's terrific.

Let me just first say, contrary to some of the examples I'm going to give, I'm actually an optimist when it comes to AI for privacy and cyber reasons. I'm not saying there are not privacy concerns. There certainly are, and we should have some reasonable guardrails. So happy to get into that.

But in terms of your specific question. I'm concerned around some of the national security risks. We see adversaries rapidly collecting a lot of data on Americans and trying to make sense of that for nefarious purposes. Often that's to identify intelligence assets, members of the military. This, unfortunately, is not theoretical. We've known this has been happening for a long time, but the Russia-Ukraine conflict made this front and center. Just to give a couple examples there, we saw openly available data, super sensitive data like biometrics, genetic data, being targeted for disinformation campaigns and even physical violence in some cases. And if you've been following a lot of the White House developments, we've seen some statements come out of the White House

in terms of how AI actually may accelerate that in terms of knowing that adversaries already have a ton of data, but now AI and some of the related technologies making the identification and seeing who may be a sensitive target even quicker and easier. That is definitely a risk, and I think something the U.S. should be out there and trying to get ahead of.

Ben Swartz | 12:34.529
Yeah, thank you again for having me. And first to disclaim the views are my own and do not necessarily represent the Commission or any individual commissioner.

Yeah, I wanted to highlight, maybe not one of the biggest privacy breaches, you know, Cambridge Analytica immediately comes to mind when you think about very large privacy breaches, but a series of cases that the FTC brought last year around pixel tracking with regards to BetterHelp and GoodRx, which were websites that shared sensitive information via these pixels or software development kits, SDKs, with marketing partners. And I think I want to highlight these cases and why they're important because the collection of data was very difficult for people to know what was going on or reasonably be able to avoid. There was a lack of clarity around, you know, what data was being collected and how that data was being shared and whether or not it was scrubbed when shared with marketing partners.

And I highlight these cases because as we turn to AI, I think that they could be potential harbingers of what's to come. I think we can all agree in this room that the amount of data and information that large language models and AI collect is enormous. And the developers of these models are trying to collect as much data as possible. And so being cognizant that these types of privacy risks might continue into the future. And the black box nature of some of these models, you know, not knowing exactly what data goes in, how it could come out the other end is something to be concerned about.

Jason Green-Lowe | 14:10.240
Thank you. So, something that attracts a lot of attention just kind of in the corporate workforce is privacy policies. We all click through them, you know, yes, I agree, yeah, sure, whatever you said, here's the terms and conditions. Do those need to change as AI develops? What needs to go into a privacy policy for a company that is running generative AI like a chatbot or an image generator? What are some unique privacy concerns that are raised by these new technologies, if any?

Ben Swartz | 14:40.558
Happy to start this one off. I guess something that the Chair has said many times and the Director of the Bureau of Consumer Protection has said many times is that this

notice and choice regime around accepting terms of service and privacy policies has failed. And through some of the enforcement actions the FTC has brought, I think that they've thoroughly rejected that fiction that just you know stuffing more and more clauses into privacy policies in terms of service is necessarily the correct way to go and so to some extent I reject the the sort of premise of of the question, because I think if really what we're looking for is more corporate regulation, to just stuff more and more hidden causes and have consumers deal with consent fatigue clicking through privacy policies and terms of service and hoping that will solve the underlying privacy harms that that could exist, I think is something that we we should be really be concerned about going forward.

Jason Green-Lowe | 15:33.156
I think you raise a good point. So I'm going to pivot with this question. And maybe it's not the consumer end agreement. Maybe it's an internal corporate privacy policy. Is there a meaningful policy that a company could adopt to govern its own behavior?

Ben Swartz | 15:47.931
I think, you know, to some extent that could be the case. I guess one thing that I wanted to highlight is a blog post that the Office of Technology staff and others in the agency wrote earlier this year that warned companies that, you know, if they surreptitiously change their privacy policy post hoc to start collecting data and start training on that information, that that might be considered unfair or deceptive. And thinking about internally in these companies, perhaps they have these policies set up on day one, but later business objectives change, and it's sort of unclear whether or not that actually ends up protecting consumers in the way that they initially set out to.

Maneesha Mithal | 16:27.561
Can I actually just jump in on that specific point? I mean, and Ben, I completely agree. And this is something I know the FTC has been saying for years and years that, you know, overloading privacy policies in terms of service is really not doing anything for consumers.

I guess the thing I would say is that I do believe that there has to be a role for notice and choice. And it may be an unpopular opinion because notice and choice is so vilified. But because consumers' preferences truly are different. Some people may want their data to be used for personalized medicine. Some people may want their data to be used to improve diversity outcomes in organizations. And some people may say, look, I get it, these are positive use cases, but I just don't want to be a part of it. So I think because consumers' preferences on privacy are so different, I think there has to be a role for notice and consent.

That said, I agree that privacy policies and terms are not the way to do that, but I do think that they do play an important role. Because I think representing clients, I think when they have to draft these privacy policies in terms of service, it serves as kind of an internal business policy function. They're saying, okay, here are the statements we're making to consumers. We have to make sure that our policies match up with those statements. So I think it serves an internal function, and I think it serves an external function because there are watchdogs like the FTC, like think tanks and other public policy organizations that can look at these public-facing documents and really hold companies accountable. So I think there are roles for that.

I think we are putting too much onus on the consumer. So I think there are some things that companies should be doing, which I'm sure we'll get to, that go beyond notice and choice. But I do think notice and choice is part of the equation. It's not the full answer.

Brandon Pugh | 18:10.344
Yeah, I agree. I think it's actually the way I see it is probably a hybrid. I think too often times there's companies and there's temptation to just bury something there. So you're hoping a consumer doesn't actually see it. So I think it's really encouraging to see some companies that have kind of revamped that in terms of highlighting specific portions, doing more effective PR campaigns for consumers, or relying on different types of techniques. I know data minimization is controversial these days, especially when it comes to AI. I think if done right, it has a lot of value as maybe an alternative or supplement, I would say.

But from a corporate end... I think if you're approaching AI as like this is a new starting point or a new frontier for us, I actually think you're wrong. I think it does have uniqueness to it, and you should be considering the uniqueness factor. But privacy should be applying to all technology development. And so this shouldn't be something where you're like, well, holy cow, we need to generate an entire new plan. Hopefully, you have a strong privacy program outside of that, and some are more sophisticated than others. I think a good example of that is when it comes to cybersecurity, we throw out secure by design and secure by default all the time. You know, CISA loves those two terms. I actually think they have a neat analogy to privacy.

Privacy should be something that's baked into the technology and is encouraged and, you know, fine-tuned during the development phase. It shouldn't be an add-on after the fact where you're just doing a privacy review and you're saying, oh, this doesn't make the cut go back to it. Hopefully, you're having your privacy experts ingrained from the

very start. And we do see some companies doing that. Not all, but there are some that are more proactive than others. I think it's a really good model.

But secondly, I think part of your plan should also be seeing how you can leverage some of these developments for good within your company? Is there a way to find AI privacy developments that will actually lead to more effective privacy? Not all, but there are definitely some.

Mark MacCarthy | 19:49.871
So I want to highlight a tension between privacy in this area and the development of large language models. You probably heard the news that Meta has pulled back from providing its service in Europe because of difficulties complying with the European privacy law. The European privacy law says you have to have a good reason for using data, and they couldn't come up with a good reason. They didn't have consent. It wasn't necessary to provide the service, and it wasn't clear that they had a legitimate reason to do it. So they suspended their training on European data in the interest of preserving privacy. So that's a tension we have to be careful about.

The consent model, I think, makes a lot of sense. It's very intuitive. I mean, you know, if I'm using ChatGPT, I would want to push a button that says, do not use my interaction with ChatGPT to train your system. I don't want that. And I'd like to have that capacity to push that button. It has to be easily available. It can't be hidden behind 13 screens. But I'd like to have that capacity.

But you know, there's a tension there, too. I mean, What if I'm a bad guy? What if I'm someone who wants to use ChatGPT to develop a bioweapon, or to engage in revenge porn, or to commit consumer fraud? The company can detect that kind of thing by looking at the interaction with the system and doing some data analysis associated with it. If the consumer can say, I don't want that, well, who's going to take advantage of that? The bad guys, right? So your ability to find vulnerabilities in your own system and detect problems with your own training is minimized by the people who opt out of having their information analyzed.

So there are a lot of tensions here that we have to pay attention to. I do think It's not just up to the companies. You can't say to them, you figure this out. There has to be a structure of law. And so if we turn to the companies and say, we don't know what the right answer is, we hope you do, I think that's a recipe for letting them make public policy decisions in the private sector, which I think is a mistake.

Maneesha Mithal | 22:09.933

In full disclosure, my firm does represent OpenAI, and ChatGPT does offer an easy opt-out to the training data.

But putting that aside, I think the training data… When we talk about AI, and we talk about notice and consent, I know I was just championing notice and consent a second ago, but when we talk about AI, a lot of the large language models train on publicly available data. And so I think to Mark's point about what's happening in Europe, you know, how do you get the consent for publicly available data? Now, in the U.S., under state privacy laws, publicly available data is generally not considered personal information. And so the companies have been able to train these large language models using publicly available information. But I think that data minimization, legitimate basis in Europe. I think these are all concepts that, you know, as we think about federal legislation in the United States, we should think about.

I think it's also important that as you're thinking about what federal legislation might look like, to understand that we're not operating in a vacuum. So there's all these state laws. So for example, if you're using a chat bot, there's a lot of states that require disclosure that you're interacting with an AI and not a human. So these are the kind of chatbot disclosure laws. There are a lot of laws. It's interesting.

There's state wiretapping laws that require, some of those states require two-party consent before interception of communications. And there's been a very active plaintiff's bar that's been saying that use of a chatbot is actually the interception of a communication. And so the law requires two-party consent. So a lot of companies have been putting consents in their chatbots. Like, you know, by using this, we're disclosing that you're using AI, and you agree that we can use it to train. It's kind of like when you call a customer service line and they say, we can use your data for quality and training purposes. So you see that on a lot of chatbots right now too.

And I think a lot of these states are thinking about disclosures, to our point about not burying disclosures in a privacy policy. They're thinking about these disclosures as just-in-time, outside of a privacy policy, outside of terms of service. So as you're thinking about disclosure remedies, I think about what needs to be kind of pulled out and what can be in the privacy policy in terms of service.

Jason Green-Lowe | 24:44.016
Great. So I think Mark raises a very interesting issue. To what extent should there be gaps in the privacy requirements? Maybe you want some kind of carve out that allows companies to train on your data for the purpose of preventing criminal action or

something like that. Are there other gaps that we should be in favor of? Are there other places where less privacy would be a good thing?

Maneesha Mithal | 25:11.299
I'm not going to say less privacy would be a good thing, but I do think that there's some good models out there.

I think there should be federal privacy legislation. I absolutely think there should be federal privacy legislation. I think, you know, the FTC has stepped in in the absence of federal privacy legislation and been very aggressive about enforcing law against deceptive and unfair practices, which is great. But I think there's gaps. If a company didn't have a privacy policy, I don’t know that the FTC could say the practices are deceptive. If the companies didn't offer choices for non-sensitive data, I don't know that the FTC could show that the practices are deceptive or unfair. So I think there's federal law needed to fill in those gaps.

But I think what we've seen is that I think those laws need to have some flexibility. So, for example, if you're talking about data minimization, we see a lot of laws that say that companies can't keep data for longer than they have a business need to do so, whereas if Congress were to say, here's the data minimization schedule, thou shall not keep data for more than three months, thou shall not keep this data for more than six months, I think that's probably too inflexible because we've seen new uses of data. Who could have imagined ChatGPT three years ago or five years ago? And so I think it's important, as you think about legislation, to have flexibility built in place. innovation and new products that consumers want and that are consistent with their reasonable expectations to flourish in the marketplace.

Brandon Pugh | 26:42.254
Yeah, Maneesha raised a good point. Back to Mark's point around tension, I think we see that even in federal privacy legislation. How far do we lean toward privacy versus innovation? They don't always have to be at odds, but sometimes more privacy sometimes means more restrictions and maybe harder to innovate.

And this is actually a very timely panel because I'm sure everybody here knows, but there's a markup scheduled tomorrow on the American Privacy Rights Act, something that probably I engage every single day on for multiple hours a day. So it's kind of a passion of mine.

But regardless if it's APRA or some other counterpart, I think that is really the first step if we're serious about AI. Because everybody here that's a little more technical than me

even, at its core, AI is around data. So if we don't have a federal standard on how we're collecting, using, transferring, securing data, we're really behind, especially considering we do have this, I hate to use the word patchwork, but we do at the state level. So if you're not in one of the states that has a privacy law, largely speaking, you don't have protections or any guarantees around your data. Absent sectoral approaches and perhaps agency investigations, that's a little separate.

But a law like APRA would have both direct and indirect implications to AI. If you haven't followed it, the most recent version, the bill that was introduced, did eliminate the provisions that were specific to AI. You know, there were parts in there that called for impact assessments and also had civil rights protections. They're not in that bill. And if that's of interest, happy to kind of dive into that. But there are many parts that would directly impact AI. Depending on where you sit, sometimes it's good, sometimes it's bad.

Maneesha brought up data minimization. Essentially, under this bill, and it's over generalization, there are 17 permissible purposes for collection. There is a concern that if you have a targeted list, is there going to be a future need or innovation that's not reflected currently or that hamstrings AI? Because ultimately, we need data for more effective AI models. That's kind of up in the air, and I think there are strong schools of thought probably even amongst you in this audience. But happy to kind of explore that further during Q&A. Thank you.

Mark MacCarthy | 28:36.655
So one quick add-on on APRA. I think the permissible purposes approach is a good one. It basically says roughly what the European Data Protection Regulation says, which is if you're going to be using data, you have to have a good reason to do it. And here's a list. If you're on the list, you're okay. If you're not on the list, you're not okay.

And as you pointed out, that's far too inflexible. The world is not going to stay static from now until the next privacy legislation. So something has to be done to provide flexibility. My fix is to send it to the Federal Trade Commission and give them the ability to add to the list after notice and comment.

There's another idea that I've got in the back of my head that's probably not on the politically feasible table, but we've always had in the United States a kind of sectoral approach to privacy. Medical privacy has been regulated by HHS. Education privacy by the Department of Education. Financial privacy by the financial regulators. And I think to some degree there should be some movement of authority over privacy and the use of data to these specialized agencies.

I'm thinking in particular about the issues that are going to be raised by self-driving cars and interconnected cars. The agency that seems most likely to be expert in that area, to know what data uses make sense and what data uses don't make sense, are NHTSA. It's not the Federal Trade Commission. The Federal Trade Commission would suddenly have to become an expert in highway safety to do that job properly. So I think some flexibility, not just to the FTC, but to the specialized agencies might be something that would more adequately allow the use of data when it makes sense and the prohibition of impermissible data uses when it doesn't make sense.

Brandon Pugh | 30:32.278
Real quick add on that. I actually think that's a benefit of a national privacy law because to the extent you're sitting in the audience and you're thinking maybe the FTC has overstepped or went beyond their legal authorities, if you are thinking that, this is an opportunity for Congress to specifically spell out, like, FTC, we're directing you to issue guidance, we're directing you to issue rulemaking, and spells out specific lanes to go forward with.

But I do agree, in the absence of having any congressional direction, what do you do? You rely on your existing legal authorities. That is, to me, a really good opportunity to either, whether it's the FTC or some other agency, to kind of give them some clear direction when it comes to privacy and AI.

Maneesha Mithal | 31:06.641
I would just add one asterisk to that. I think opt-in should always be… well, yeah, almost always… I think there should be, I don't think we should be so paternalistic as to let consumers opt in to certain uses of their data. So I think that is another cure for the inflexibility that we talked about with APRA, because, like, if you look at GDPR, there is always, you know, consent is a valid basis on which to process data. So I do think that that should be added to the list.

Brandon Pugh | 31:43.678
Yeah, no, I agree. And hopefully, if I said the opposite of that, I didn't mean that. So that is a great point.

Jason Green-Lowe | 31:49.720
Yeah. So at the Center for AI Policy, we're deeply sympathetic to this idea that the technology is going to change before the next bill gets passed. It can be a long time between landmark legislation and you probably do need some kind of specialized agency with expertise that can update it as the technology changes.

But as a matter of curiosity, do any of the panelists have a crystal ball? Can you see what new uses of data might be coming down the pike? Can you see... You know, how is the use and collection of customer data going to change over time? Maybe not every use. We still want updates from the regulators, I suppose. But is there a threat or two or a use or two that you think might be available in the near future?

Ben Swartz | 32:31.776
Not necessarily to be prospective, but looking towards the recent past, one of the things that I think about are the facial recognition cases that the FTC has brought and sort of moving towards more advanced AI thinking through that lens.

And the two that immediately come to mind are a case we brought last year against Rite Aid where in the settlement, the FTC effectively banned Rite Aid from using facial recognition technology for five years.

There was a similar case in 2021 against an app developer who didn't tell their users that they were going to use their images for facial recognition technology. And in the settlement for that one, not only did the FTC make sure that the developer collected consent, expressed consent to train on people's images, but also required them to delete the algorithms and models that were built based off of the information that was collected. And so thinking through, you know, going forward through that lens I think is important.

Mark MacCarthy | 33:38.643
So I want to go back to an issue that Manisha raised earlier, which is this issue of non-consensual AI porn. It's a use that's not way in the future. It's right here and now. And it's not just Taylor Swift. It's teenage girls. And there was one story of a 27-year-old woman who didn't want to go out with the guy, so he got her a picture and made some pornography videos of her having sex with him and distributed it online. It's increasingly easy to do this. You don't need hundreds of pictures, you need just a couple. And it should be against the law.

And there is a piece of legislation, the Defiance Act, that's out there, sponsored in the House by AOC and in the Senate by Josh Hawley. Anything that has those two people in it at the same time can't be all bad. And it doesn't go in the direction of blaming the platforms or the companies that produce the technology, it goes after the people who produce the stuff and says you can bring a case against them and fine them up to $1,000 if you can make the case that they did it without getting consent or without the knowledge that they had consent to do this kind of stuff. That seems to be an easy thing

to pass. It's bipartisan. It doesn't try to solve all of the world's problems. It's a pressing issue made possible by artificial intelligence, not something that was as easily done five years ago. So here's some low-hanging fruit. Let's do it.

Maneesha Mithal | 35:11.812
So I actually think that we are just limited by our own imaginations as to the potential uses of AI. And it's very interesting having switched hats for me. So I was at the FTC for over 20 years, and then two years ago I came out to the private sector.

And so wearing my regulator hat, I see a voice cloning product, and I'm like, fraud, impersonation, scams. And I see all the bad uses.

But you know, when we talk to our clients, it's like, okay, well, we are helping people who have ALS and they want to be able to preserve their voice or they want to be able to tell stories in different ways, or they want to have the voices of their loved ones speak to them. And so it just, it feels very… Again, when I was at the FTC, I would look at AI in employment and I'd say, well, that's rife with bias. And I'm not saying it's not. It certainly is. But we have a client who is using AI to try to build a diversity pipeline for their customers where they can analyze the race and gender of candidates at the interview stage and at the hiring stage. And so it's basically just analyzing their own workforces to help promote diversity outcomes.

I think we just have to keep in mind that it's like fire, right? There's bad uses and there are good uses. And we have to sort that out and make sure that the policies that we adopt permit the good uses and put the brakes on the bad uses. And it's not an easy task because who could have imagined some of these positive use cases? Who could have imagined some of these negative use cases?

Brandon Pugh | 36:54.487
Well, with a framing like that, it's hard to follow on. But maybe I'll just give a couple concrete examples of how I'm seeing this.

I do say this doesn't mean do nothing from a policy and legislative standpoint. I think there's room for that.

But I do think regardless of what we do on that front, bad actors are going to continue to exploit and misuse the technology. So I always like to see, is there a way we can use the tech now for positive applications? And specifically, I look at it in the sense of cybersecurity and then consequently data security as well as privacy. And interestingly enough, we've been leveraging this. in the security space to protect data for well over a

decade. Because of the rise of generative AI, it has been really interesting to see over the past 18 months how that has accelerated cybersecurity companies and data security companies to further innovate in their space.

So is there a scenario where we could automate a response to an incoming incident rather than having it be a human and potentially having a delay time? Is there a sense to be able to better detect what may be a malicious attempt to access data versus perhaps just an accident or anomaly? So those are really encouraging, exciting things for me.

And I think, you know, something that APRA has done well is an increased focus on privacy enhancing technologies. Whether it's that bill or not, I think that's an area we should be continuing to put research and additional funds into to kind of further that market. And I think the FTC and some of their colleagues have done a good job on that front, so I think those efforts should continue.

Jason Green-Lowe | 38:18.052
All right. Thank you so much. Let's take questions from the audience. If you can, please say your name and what office you're coming from. And if you have a particular panelist in mind, feel free to say so. If not, we'll assume it's for the whole panel. In the back, yeah, on the corner.

Audience Question 1 | 38:35.134
Fred, Congressman French Hill's office. I just had two quick questions. One for, and they're both in the kind of the frame of privacy, national security, cybersecurity. So one, I was wondering if you can kind of expand upon your thoughts on specialized agency and like the guardrails for the FTC when it comes to some of the privacy protection. What would you say some of those like guardrails that kind of look like, what would be some great ideas? And then could you also expand upon some of the cybersecurity and national security broadly?

Because I know, like, for one example, like, export controls and when we're doing bilateral trade agreements and what they're asking for in some of those contracts could lead us to give them some proprietary information. So any, like, overarching things like that, could you give some examples of what we should be looking out for?

Brandon Pugh | 39:19.473
Yeah, no, all good questions. I will say I'm not speaking on behalf of the Department of Defense, but in my capacity as a national security law professor, this is largely what I wrestle with. And this is just open source now.

It's publicly been reported, so I'm not confirming it's accurate. But the White House and National Security Council is allegedly working on their national security memorandum, which is going to be a counterpart to the OMB AI guidance from the AI executive order, which is allegedly going to be out at the end of July. A lot of that is going to focus on the role of AI for national security purposes. How is it being used by DOD and the intelligence community? And then the flip side of that, how do we leverage it against potential actors? So that's definitely something to be following. Part of that will be classified. That is one area that I would be following.

And I think... Maybe I'm biased, but the DOD and the intelligence community has been very forward-leaning when it's come to AI. Yes, GenAI and now other civilian federal agencies are now trying to get up to speed. The DOD has been wrestling with this since 2014. And they first stood up their first office, don't quote me on this, but I believe in 2016. So I think they're really the thought leaders. In terms of how adversaries are using this.

Yeah, I'm just trying to think, in terms of an unclassified environment, we do know, like, adversaries like China are already trying to vacuum up—and I hate to single out China; others are doing it—but they're vacuuming up as much sensitive data as they can on Americans and are trying to make sense of that for a wide variety of purposes. And they do it on their own citizens, of course, too. And AI is definitely accelerating some of those goals.

And I think there is a potential that could play out in a future conflict. We've seen it in both Israel-Palestine currently, and we've seen it in Russia-Ukraine. So I think there's definitely a potential in future conflict for some of these data and security concerns to come up, and AI will play a role, even outside the whole autonomous weapons debate. That's a whole separate panel, but obviously AI has a role there too.

Maneesha Mithal | 41:16.609
I was going to answer the guardrails question. So I think one of the things is you could envision two kinds of delegation to the FTC or another agency. You could envision a very short law that says we believe privacy is important. FTC, go issue rules. Right.

And then you could say, OK, well, there's a very kind of clear mandate for the FTC to say, you know, here's what we want you to do specifically. And I think that kind of specific delegation with you know it's such a balance between giving flexibility but but I think that kind of specific delegation… Here are the things we want you to do… Here's the general parameters of what we want you to do, and you go fill in the definitions and

you go fill in some of the more specifics. I think that that's a more easy to implement and easy to carry out and defensible way of doing this.

Mark MacCarthy | 42:13.216
Yeah on the guard rails to the FTC, I do think they need rulemaking authority to implement uh APRA or other pieces of privacy legislation. You can't just write the law and then say FTC, enforce… Enforce what? The companies have to know exactly what the law is and you can't write that in the law, you have to have implementing regulations that tell companies what their obligations are so they need that kind of rulemaking authority.

On the national security stuff, I want to go in a slightly different direction. I think because of national security concerns, there's been a big change in the way the U.S. has looked at the issue of data localization. We previously had a policy that said no country should keep data within its own borders, it's a mistake, what we need is the free flow of data across borders, just like we need the free flow of goods and services across borders. And it's a mistake for countries to try to keep data in their own country.

And now we're beginning to think, well, wait a minute. There are useful regulatory purposes and national security purposes in keeping data from flowing out of the United States. And so our trade representative has changed the position that we've adopted in international negotiations on trade issues and said, no, no, we're not pushing for an end to data localization anymore. It’s something that we're actually kind of in favor of.

China, of course, is in favor of data localization. They've never been shy about that. A company like Tesla wants to bring its information out of China so they can help to perfect their autonomous driving cars and combine it with data in the United States. I think that's a pipe dream, by the way. But they wanted to do it, and China said, no, you can't do it. Now, they're flexible. They've made some changes in their rules to make sure that some unimportant and non-personal data can be transferred out of the country. But they're not shy about saying, we need to have control over our data and keep it in the country if that's necessary. I think the US is coming around to that as well.

Europe has been in that position for years. They've got very tight rules about data transfer outside Europe. We've tried to negotiate with them several times to have data transfer agreements. Two of them have been struck down by the courts. The third has been adopted, but it hasn't been finally approved because our friend Max Schrems hasn't sued to get rid of it yet, but he will. So this is a trend I think we should pay attention to.

Why is it important for artificial intelligence? The more data, the better. If you can transfer data across borders, you can help to improve your artificial intelligence. If it's contained within one country, it becomes less effective as a training tool. So it's very important to try to have this kind of balance between the national security needs, the regulatory needs, and the needs for data for development of artificial intelligence.

Brandon Pugh | 45:10.992
10 seconds, 20 seconds maybe. I neglected your guardrails question. I think they did a good job. APRA has provisions in there for the FTC to provide some parameters on data security. The data security provisions of APRA are often overlooked, but to me the most important probably… well, among the most important. Then we've also seen the FTC independently consider rulemaking around what data security could look like absent APRA.

I will say on the data localization front, the White House was very quick to say this is not a data localization effort, but there has been a federal law passed by both House and signed, as well as an executive order looking around the sale of bulk sensitive data to either, depending on what you're looking at, countries of concerns or adversarial countries. That is underway, but it is essentially targeting more toward data brokers. But some have said that is really an early attempt to see more data localization.

Audience Question 2 | 46:01.434
Hi, I'm Natalie. Not with an office. Just completed my thesis on AI policy and regulations, helping educate. One thing I'm curious about is there's obviously this theme in terms of um balancing within regulations right we want flexibility, we want progress, we want innovation, but we also need to make holistic frameworks. So my question is in general as we move forward, what kind of framework would you like to see more in terms of privacy law?

We talked a little bit about, like, EU, they're doing a more risk-based approach. We've heard a little bit about an industry-based approach. You know, in your sort of idyllic world, and we all know there are, you know, there's, you know, industry is this, that, and the other. But when we're trying to think about a general framework for how this would look, what are your ideas or recommendations on how this should start?

Maneesha Mithal | 47:00.521
I can start. So I think the EU AI Act and the Colorado law that was just passed have a really, I like the risk-based aspect of those two laws. Looking specifically at Colorado, I think they have, I mean, I think if you look at the AI laws, there's the kind of process-based requirements, which is, you know, you have to have a governance

program, you have to make disclosures, you have to conduct risk assessments. So I think those are all, you know, really good.

And I think there's... the more, I guess I would call them rights-based requirements, which is that if, for example, you make an adverse decision against a consumer, you have to give them notice and the ability to correct their information. So I think those pillars generally seem... seem like the right place to be

And the thing I do like about Colorado is that it's got certain affirmative defenses and safe harbors. So, you know, companies that want to do the right thing, that are trying to do the right thing, that are following the NIST framework, that are conducting red teaming, they get kind of a safe harbor from enforcement if they're doing, if they can show that they're doing all those things. And so I think that's a good way to preserve flexibility, have guardrails. And I'm not saying Colorado legislation is perfect, but I do think that those are things. The risk assessment and the kind of safe harbors. I think there's good.

Mark MacCarthy | 48:26.048
I think two parts. Most of the risks and benefits from AI come from how AI is used, and many of those uses are regulated. So I think the first step is to make sure that the existing regulators—our friends at the FTC, but also the EEOC, the financial regulators. Where AI is used in their area of responsibility, they have to have full authority to go after not simply the user of the AI system, but the developer of the AI system, right? That's one of the limitations right now in EEOC law. They can go after the employer, but the vendor of these employment-related algorithms escapes their jurisdiction.

That can be fixed. A federal law, a colleague of mine at Brookings proposed a kind of uniform federal law that would expand the jurisdiction of all of the specialized agencies to give them the authority over vendors. So that's one. Make sure that the existing agencies have full authority.

For the general purpose AI, which is not used for any specific application, I do think something like what the EU has put forward it would make some sense of risk based approach where there are requirements for safety testing and red teaming and essentially approval by an outside assessor to make sure that they've done the job.

Ben Swartz | 49:56.393
Just to add on quickly to that, one thing that I know the Chair is very fond of saying is that there's no AI exemption from the laws on the books.

But to get to it, not to talk about any specific proposal, but one thing we've heard from stakeholders that we've brought together is the consent fatigue issue that I brought up earlier that has come around. And so just thinking through whatever ends up coming up that doesn't repeat some of those mistakes from the past. And I know that the FTC has brought a number of settlements that ask companies to provide consumers with affirmative express consent. And so thinking through different ways of asking users for consent to certain things. It's important.

Brandon Pugh | 50:38.404
Yeah, I think Ben and Mark both bring up good points. I think it's a matter of looking and utilizing our existing authorities, not saying there's not a need for more niche piece of legislation, but before we jump into a comprehensive AI measure, what are the existing authorities and are there gaps in those authorities? I know some agencies have been a little more forward leaning than others to identify them, proactively use it, and apply them to AI. So that would be a first step.

I also think it's really important to continue to empower and make sure our resourcing agencies like NIST. I'm a huge fan of NIST. Love their work when they came out with both their cyber and their privacy frameworks and actually liked their AI one arguably even better because it largely relied on their privacy and cyber and built on it a little bit. So I think those types of voluntary best practices that are a little more flexible I think are ideal.

Audience Question 3 | 51:24.417
Hi, I'm Willy from the Senator Young's Office. So, I think there's some interesting things about all the companies that to some degree or another have a business model around privacy. So, you know, Apple maybe part of its marketing is a little more private than Google. Proton Mail. There's, I think, like Brave browser, things like that.

Are there any regulations on the books that, if changed, would open up space for more companies to differentiate themselves in the privacy that they offer to consumers?

Maneesha Mithal | 51:56.337
I mean, I don't know that regulation can do that. I mean, I think it's great that companies are starting to compete on privacy. And I think the FTC plays a really important role because by policing deceptive practices, if somebody says our privacy is better than this other browser's and they don't have any substantiation to back that up, I think that's a really key piece.

And I think that was one of the FTC's theories when it brought its action against Zoom, that it was claiming to use end-to-end encryption. And the idea was that Zoom... gained market share now. I think in the FTC's complaint, they talked about that being a possibility for how they increase their market share during the pandemic. So I think vigorous enforcement of Section 5 is going to help that competitive issue.

Mark MacCarthy | 52:42.516
Yeah, I agree. I don't see any regulatory barriers for companies adopting more stringent privacy rules to protect their own customers. Maybe there are, and I don't know them, but Apple's a good example of a company that tries to…

They're going to run into trouble with Apple Intelligence because they're going to have to start using data in ways that they haven't done before. If Apple Intelligence is going to sort through your photos, it's got to have full access to your photos. If it's going to start dealing with your email, it's going to have to have full access to your email. And then the question is, what do they do with all that now that they've got access to it and they've analyzed it? Do they start sharing it with other people? Do they turn it over to OpenAI and OpenAI does things with it? We don't know.

So there are challenges ahead for them in that area, but I don't see any regulatory obstacles for other companies to be privacy friendly.

Brandon Pugh | 53:34.090
The question I'd be asking if I was a regulator or staff or a member is like, what can we do to actually incentivize this type of behavior? Like are there proactive actions we can take?

I mean, not to sound immodest, I think when I was still serving with the New Jersey State Legislature, we were one of the first states to have a safe harbor bill around proactive cybersecurity measures. So if you voluntarily followed a cyber framework, you could have an affirmative defense on the back end, essentially, a little more detailed than that. Those are interesting because, you know, it didn't work for every company, but it was a way for some companies now to voluntarily put more money into compliance and programs up front that potentially saved on the litigation in the back end.

APRA takes a very similar approach when it comes to privacy-enhancing technologies. Essentially, you have a privacy-enhancing technology program, you demonstrate compliance, you potentially could save on claims that you didn't have data security measures that were, you know, in place.

Jason Green-Lowe | 54:26.082
All right, one more question.

Audience Question 4 | 54:30.129
I've lived without social media and I think it's all trying to be eradicated, but what would happen in a world where social media was eradicated? I know there's many industries in the world. But could you talk about if social media was eradicated and how that would change?

Maneesha Mithal | 55:01.310
I mean, I think that, so, I don't know, I think there'd be a vacuum that would be filled by other things because, like, I think, you know, data broker industry would, you know, still exist and you would just be targeted on your web browsing behavior or your kind of app usage behavior. So, I do think that the need for privacy would still exist without social media, but. But to your point, I mean, I think it, you know, there's a lot of pros. I personally couldn't do it.

Mark MacCarthy | 55:40.538
Yeah, look, there'll still be plenty of data on the internet itself if you get rid of Facebook, and that might be enough to continue training.

Although, you know, people are turning towards synthetic data now because they're running out of real data. And so there's a data shortage, surprisingly. The demands for extra data are exceeding the increased supply from social media and from the internet.

The one thing I would say is that there is a movement in that direction, but it's not really banning social media. It's making sure it's not available to kids. That seems to be the direction that people are moving, that this is a dangerous product that should be labeled or restricted or somehow made safe for children if they are going to be using it. That is a new development that I would not have expected three or four years ago, and it's gaining steam in the states and there are some elements in the Congress that are thinking about the same sort of movement. It's not aimed at adults. It's aimed at kids. But it is a movement that I think is gathering steam.

Maneesha Mithal | 56:43.866
Yeah, and I would characterize that more as safety than privacy. I think the kids' privacy concerns are on one side, and I think the concerns around social media are largely around content, addiction, suicidal ideation, depression, those kinds of things.

Brandon Pugh | 57:02.889
Yeah, I mean, just a quick thought. I think they both raise great points. I mean, I couldn't live not on LinkedIn. I'm on it way too often.

But on a serious note. I think the better approach is that, and I see your point, is to look at are there guardrails and steps we can take to make it more transparent and empower consumers more. So at least they're more informed around what data is collected about them, how it's being used, how it's potentially being transferred to third parties. I think those are types of reforms and things we should put in place, and we certainly could do as we sit here. Whether there is a political will to do it, time will tell. Maybe we'll see it tomorrow, but it's definitely possible.

Ben Swartz | 57:38.018
Yeah, just to reiterate. Data collection on the internet happened far before any social media company became very, very large. And to imagine that data collection would just cease if this one industry went away, I think, you know, would not be the case.

Jason Green-Lowe | 57:52.625
All right. Well, thank you for a very enlightening conversation. We're going to take a quick photo op with the panelists, and then I encourage everyone to mingle if you have more time. Sometimes panelists have been able to answer a few questions after the formal end of the event, which is now.