Consumer Finance Monitor Podcast (Season 9, Episode 21): The White House AI Framework: Ambition, Preemption, and Uncertainty Ahead

Speakers: Alan Kaplinsky, Gregory Szewczyk, Charlie Bullock and Kristian Stout

Alan Kaplinsky:

Welcome to the award-winning Consumer Finance Monitor Podcast, where we explore important new developments in the world of consumer financial services and what they mean for your business, your customers, and the industry. This is a weekly show brought to you by the Consumer Financial Services Group at the Ballard Spahr Law Firm. And I'm your host, Alan Kaplinsky, the founder and former practice group leader for 25 years and now senior counsel of the Consumer Financial Services Group at Ballard Spahr. And I'll be moderating today's program.

For those of you who want even more information, either about the topic we'll be covering today or, for that matter, anything else in the world of consumer finance, don't forget about our blog, which, like our podcast, also goes by the name of Consumer Finance Monitor. We've hosted the blog since 2011 when the Consumer Financial Protection Bureau became operational, so there's a lot of relevant industry content there. We also regularly host webinars on subjects of interest to those in the industry.

So to subscribe to our blog or to get on the list for our webinars, please visit us at ballardspahr.com. And if you like our podcast, you can let us know about it, please. You can leave us a review on Apple Podcasts, YouTube, Spotify, or wherever you obtain your podcast. Also, please let us know if you have any ideas for other topics that we should consider covering or speakers that we should consider inviting as guests on our show.

So our podcast today will focus on the White House National Policy Framework for Artificial Intelligence, which was released on March 20th of this year. This framework contains a sweeping set of legislative recommendations intended to establish a coherent, nationally unified approach to AI governance. While the framework does not itself create binding legal obligations, it's likely to shape federal AI legislation in the months and years ahead.

Our webinar today will summarize and comment on the framework's key areas of focus, and we'll be considering, among other things, what its influence can mean for the current state regulatory landscape, including Colorado.

We published an article about the framework in our April 1, 2026 CyberAdviser Blog and on April 8th in our consumer finance monitor blog.

Now, today's episode builds on a conversation we began last year when we examined the White House ... I call it the AI White Paper, that's not the official name of it. And we discussed at that time its implications for consumer protection, innovation, and regulation. We did an October 30, 2025 webinar entitled AI in Financial Services: Understanding the White House Action Plan, rather than the white paper, and What It Leaves Out. We repurposed that into a two-part podcast show that was released on December 3rd and 10th of last year.

The earlier discussion focused on a principles-based framework, what many viewed as a rights-oriented approach to governing AI. We now have a new and a much more detailed set of legislative recommendations from the White House outlining what it calls a national policy framework for AI. This latest development raises important questions. Has the administration shifted its priorities? Are we seeing a move toward a more innovation driven and less of a regulatory approach? And what does this mean for Congress, federal agencies, and industry participants, particularly those that are in the financial services area?

To help us unpack these issues, I'm delighted to welcome back three outstanding guests who joined me last year for the webinar on what I'm calling the white paper.

First of all, we have Charlie Bullock. Charlie is a senior research fellow on LawAI's US Law and Policy team. He advises state and federal policymakers on AI governance topics, and he publishes research on legal questions with significant practical relevance to US AI policy.

His recent research examines issues, including federal preemption of state AI laws, federal and state AI whistleblower protection legislation, and the likely consequences of the end of courts giving Chevron deference to the future of AI regulations issued by federal agencies. So Charlie, a warm welcome back to you.

Charlie Bullock:

Thanks for having me on again. Great to be back.

Alan Kaplinsky:

And also joining us is Kristian Stout. Kristian serves as the Director of Innovation Policy at the International Center for Law and Economics. His work addresses competition, telecommunications, and AI policy with a focus on how law and economics informs regulatory design. And a warm welcome to you as well, Kristian.

Kristian Stout:

Thanks. It's good to be back.

Alan Kaplinsky:

And finally, and last, certainly but not least, I'm pleased to welcome my colleague, Greg Szewczyk, who's a partner and the head of the Privacy and Data Security Practice Group at Ballard Spahr. Craig advises clients on a wide range of privacy, cybersecurity, and data governance issues, and brings a particularly valuable perspective on how emerging AI policies intersect with existing legal and compliance frameworks. Among other AI advice that Greg has given to many clients is the advice he's been giving about the Colorado AI statute, which I believe is going to go effective fairly soon. And I'm sure we'll talk about that a little bit. And Greg, welcome to you, of course.

Gregory Szewczyk:

Thanks, Alan. Always good to be on the show.

Alan Kaplinsky:

So let's get into the questions that I have, and I do have a lot of them, and I'm going to direct them to different people. But after I do that, each of you who's got anything to say, you can just chime in.

So let's start with you, Charlie. And I'm calling this the framing of the new White House AI framework. So at a high level, how would you describe this new legislative framework for AI? And I guess, what stands out most to you?

Charlie Bullock:

Sure. So for background, we've known this framework was coming for a while. There was an executive order back in December that President Trump issued that mentioned that Michael Kratsios, who's the sort of chief technology advisor to the White House, and David Sacks, who was the AI czar when this was released, were going to jointly prepare this legislative recommendation.

Now, of course, a legislative recommendation, which is what this framework is, is it only does anything if Congress actually passes it. Right now, it's just a recommendation. And so yeah, I think it's quite sparse on details. It's three pages long, big spaces, bullet points, and it covers a number of different areas, there are seven sections, but mostly I think the way you want to model it is this is very broad preemption of state AI laws combined with very little in the way of federal policy.

I mean, they make some recommendations, but they're quite sparse. Mostly they're negative. It's things like, "Ah, how should we handle copyright? Oh, well, let's let the courts handle it." And that's what's already happening. So it's sort of like a lot of things that the government should not do, okay, and then very broad preemption.

So that is not a surprise. That's what we kind of expected. David Sacks has been very clear about his thoughts on how AI should be regulated, which is essentially not at all. So this is expected, but also I think very unlikely to pass, basically, is what most analysts have been saying and what I believe.

Alan Kaplinsky:

Yeah. Okay. Anybody else, Kristian or Greg, at a high level?

Kristian Stout:

Yeah, at high level, I agree with a lot of what Charlie just said. I think what I would add is overall, it distills a lot of very good ideas into an ambitious roadmap for potential change. I also appreciate Charlie's skepticism that this would pass. However, I would put a little modification on that and say, I don't know that this roadmap is meant to pass as one package. I think it's meant to be suggestions as there are different pieces of legislation that may touch on different areas that these ideas may be incorporated into. So I'm skeptical that you would get one big AI omnibus. I'm optimistic that at least some of these pieces may make it into legislation that Congress is considering as part of the process.

Alan Kaplinsky:

Okay. Yeah, well, thank you, Kristian.

Gregory Szewczyk:

I would echo what both of you said, although I'd probably lean a little more towards Charlie's skepticism on what will actually pass. I do think that ... And we'll get into this more in some detail later. I found the specifics that came out on the preemption kind of the most interesting part of it, specifically what was carved out of the preemption to be kind of the most interesting in the form of any kind of details that came out of it and how that could actually lead to where we've seen a lot of strong endorsements of federal preemption, how we'll still have ... even if this would pass the full preemption that they're endorsing here, we're still going to see some form of a patchwork forming on certain issues, which to me was really some of the more interesting parts of what came here. But I know we're going to talk about that in little more detail.

Alan Kaplinsky:

Yeah, we'll get into things, Greg. We are going to get into federal preemption a little bit later in more detail.

So back to you, Charlie, do you view this as a natural evolution of the White House action plan that was issued last year, or do you see this as any kind of meaningful shift in direction by the White House?

Charlie Bullock:

Yeah, I think it's pretty clearly a very meaningful shift in direction. The White House action plan was in large part authored by Dean Ball, who you had on last time when we were discussing the plan. And Dean is a very libertarian-minded guy, he's a techno-optimist, he is pro-AI, he's pro-innovation, but he also takes the possibility of very advanced AI systems being developed in the future very seriously. And if you take that possibility seriously, it's kind of obvious that there will be risks created, that this is a very serious thing, and that the federal government needs to respond to that.

David Sacks' position is kind of the opposite. He's very opposed to acknowledging the existence of any risks or suggesting any kind of regulatory intervention whatsoever, not even regulatory, but the action plan had detailed policy recommendations on things like investing in interpretability and control research and physical and cybersecurity for frontier AI labs, for ramping up export control enforcement. And it also suggested a much narrower approach to preemption. And it was 28 pages long, single-spaced, had lots of specific ideas. This is three pages long, about triple space bullet points, and gets rid of pretty much all the things I just mentioned and many others and replaces them with much broader preemption.

So what happened essentially is Dean Ball's out, David Sacks is in, there's a new sheriff in town, and the new sheriff says we are doing absolutely nothing to regulate AI ever.

Alan Kaplinsky:

Okay. Anybody else want to comment on that?

Kristian Stout:

I don't know if I would say we're never going to regulate AI is implicit in there, but I appreciate that it is a much easier approach to thinking about the AI issues. But again, like I said at the top, I think this is meant to be sort of a roadmap of the topics and approaches to how you get at them with the implicit understanding that when there are actual areas that present themselves, that require regulation, that the government is going to pay attention.

I think we're going to talk about this later, but I think what's going on with Anthropic and Mythos right now kind of demonstrates that the government pays attention when there's something that arises that needs regulatory legal attention.

And I think implicit, for instance, another topic we're going to discuss, the child protection consumer risk, that seems clearly to me slotted into everything that's being discussed in the COPPA and KOSA area. So it's more of a, "Hey, you need to consider these things when you're regulating in these areas," knowing that we are going to regulate AI, and less of a, "We need to create a centralized AI regulator who watches and solve from the top down."

Alan Kaplinsky:

All right. I was going to ask you a question about comparing the action plan to what we have this year, what just got issued. I think we probably have already answered that in connection with my earlier question. And I'm going to ask you, Greg, to answer this question, this question that I know you've spent a lot of time on, federal versus state authority. It's clear this framework strongly endorses federal preemption of state AI laws. How significant is that development?

And I guess the other thing, if you could, and all of you I'm going to ask to comment on the same thing, as problems develop, and we did have a problem develop a couple of weeks ago when Tom Friedman wrote an opinion piece in the New York Times, and all of a sudden the world found out about this Anthropic issue with this more advanced AI system, Maud Mythos, hope I pronounced that right, and all of a sudden they called Scott Bessent, the treasury secretary, and Jay Powell, the chair of the Fed, summoned all the major mahoffs that are dealing with AI to a meeting in Washington. At least I never learned the outcome of the meeting, but I'm just wondering if that might be altered now as crises develop.

Anyway, Greg, why don't you take a shot at that and then give everybody else an opportunity to comment on it?

Gregory Szewczyk:

I mean, kind of hitting that first issue of the federal preemption, where the framework kind of goes. I mean, I don't necessarily see it as a huge shift in the positions that the administration has taken, but where I do think it is interesting is it actually is one of the areas where the framework has a little more detail on how it suggests that Congress should structure things. I mean, it lays out three areas where it thinks there should be federal preemption, which is AI development, laws that would penalize AI developers for third-parties' unlawful conduct involving their models, and laws that would unduly burden Americans' use of AI for activity that would be unlawful if performed without AI.

That first category is largely focused on frontier AI laws like California's SB 53 and New York's RAISE Act. And I kind of see the other two prongs as more aimed at components of Colorado's AI Act.

But I've also saw a White House advisor note that the Trump administration would also extend some of those principles to some of the chatbot bills that are out in legislatures right now, especially the ones that would create liability for chatbot operators that engage in unauthorized practice of certain professions like chatbots that are engaging in legal advice and things like that.

So I think depending on how these would be implemented, if anything was ever passed, they would kind of have a scope beyond just the laws that are on the books right now. And although we've seen a struggle for states to pass comprehensive laws like Colorado AI's Act, we have seen some states get chatbot laws across the finish line, and I think we'll probably see some more. And so it would actually have preemptive effect.

Where I think it's actually more interesting is the fact that although there has been strong preemptive push at the federal level, there are actually concessions in the framework. And I was a little surprised to see the walk back on the concessions.

For the listeners, the concessions that were listed were ... And I'll do kind of the less controversial and then the one where I see more ambiguity and I think more interesting discussion. The less interesting to some extent, state zoning laws, including state authorities to determine the placement of AI infrastructure, requirements governing the state's own use of AI. And then the one that I think opens up a lot of discussion is traditional police powers retained by states to enforce laws of general applicability against AI developers and users, including particular laws to protect children, prevent fraud, and protect consumers. And that's the carve-out where I think we're prone to see some real battles play out, depending on how that would roll out. So I'll stop there and give Kristian [inaudible 00:19:39].

Alan Kaplinsky:

Yeah, I'm just thinking as you said that, it sounds like the exception has the potential of swallowing federal preemption altogether, depending upon what state you're talking about.

Charlie Bullock:

I mean, I think it's important to note that this generally applicable thing was also in the last attempt at preemption. The moratorium also had an exception for generally applicable laws. And then, of course, the question is, what is a generally applicable law? It's not clearly defined, nobody quite knows, but I don't think that this is a dramatic departure from the last attempt ... for the moratorium that was proposed over the summer. I mean, it does introduce this idea via the developer-deployer distinction. So maybe we're going to preempt laws governing development of frontier AIs more than deployment because deployment is sort of more local, it's more related to the traditional police powers of the state, and development is more sort of a interstate commerce-type phenomenon. But yeah, I think it's generally quite broad and similar to the proposal that was introduced last summer.

Kristian Stout:

Yeah, I agree. And I think the frustrating thing ... So I agree with the idea that when the moratorium ideas have been floated around since last year, everybody who was promoting them at the time, and I was like this, a lot of people were saying, "No, this moratorium is really about preventing states from getting into the algorithmic design business." And you have authority about how those things are used in your territory, housing discrimination, all kinds of civil rights laws, zoning, things like that. What we don't want you to do is we don't want 50 states getting involved in the training process and in how you design neural networks and how you ... That sort of thing. That would be immensely expensive if it's even possible to comply with 50 state rules on how these systems are designed and put together. And I think this carries over with that.

I think the reason why maybe they sharpen some of the language on it being around ... I think it says design, I believe is in there, but that's what it's really getting at, is because unfortunately there are presidential hopefuls that are lining up to distance themselves from the White House on AI issues, and they are state-level officials, and they tend to try to make it seem like the feds are trying to come in and eat the state's lunch when, in fact, it's really a partnership design philosophy rather than a displacement.

Gregory Szewczyk:

So Kristian and Charlie, I got a question for you guys. I've seen different arguments on how the administration would want to have things interpreted like the scope of the language for the child safety carve out. If the intent is really supposed to be like general child safety laws that would incidentally reach AI, or if it's really more like AI-focused child safety laws, how broad that is supposed to be, or if it's just supposed to leave it up to Congress to kind of do it, and this is more just out of our court a little bit.

Charlie Bullock:

So my take on this, if you go up to the child safety section of the framework, it says Congress should not preempt generally applicable state child safety laws. And I view that as a kind of pointed message to Senator Marsha Blackburn and other sort of pro-child-safety Republicans. Generally applicable laws were already accepted the last time. So when they say, "We're not going to preempt generally applicable child safety laws," what they're saying is, "We are going to preempt all of the child safety laws that we were going to preempt anyways." So it kind of reads like they're carving out child safety, but in fact what they're doing is very pointedly not carving out child safety and saying, "Okay, no, we're coming for your state child safety laws."

Kristian Stout:

I mean, it depends on how you read police powers, and lawyers will have a lot of fun litigating this, because if the moratorium is focused on design issues, maybe there's an edge case where you're going to have a firm that all of their resources are in designing systems meant for children. But if someone is not designing their system in a way that's meant for children, the moratorium shouldn't reach it if the system is deployed in a state. So if a company deploys a product in a state that has a harmful effect on children, whether it uses AI or doesn't use AI, if the laws of general applicability says we're protecting kids, it should reach it unless there's something weird in the legislation that implements it.

Alan Kaplinsky:

Yeah. Let me also tell you something that is gnawing at me, I guess, since I heard Greg recite what was in that exemption, police powers, anti-fraud laws. What about every state in the country has a UDAP law, a law proscribing unfair, deceptive, and sometimes abusive acts or practices? Couldn't that cover everything? Couldn't that cover algorithmic, the government sort of dictating what should be in a certain algorithm? Or am I reading too much into this exemption?

Charlie Bullock:

I think that what the people who were behind this framework would argue would be that the UDAP laws would be an example of generally applicable law. It doesn't target AI specifically, it applies to all technology equally, therefore it's generally applicable, therefore we're not going to preempt it. That's what they would say.

It's a little unclear because they also do this undue burden thing. So laws that are facially generally applicable, like for example, the ELVIS Act in Tennessee, which says any technology that perfectly mimics a celebrity singer's voice. Well, okay, what technology could that possibly be than AI? So that's like an example of an undue burden law that could still be preempted even if it was facially generally applicable. But I think for UDAP, you'd be on much firmer ground to say that probably would be generally applicable.

Kristian Stout:

Right. And there's a controversy going on right now that I think actually demonstrates this. So there was a researcher from, I think, AMD who discovered that Anthropic was throttling Claude's access to really intense thinking models over the last few months in a way to preserve resources, but they weren't signaling to consumers that they were. That's essentially a potential UDAP claim because Anthropic has terms of service for people who subscribe to Pro or Max that says you get access to Opus at a certain amount. If they didn't get it and they weren't informed about it, that theoretically could be a UDAP claim. And I don't think, even if there was legislation that says you're not allowed to regulate the design of AI, that's not an AI design problem, that's a contract problem.

Alan Kaplinsky:

Right, right. Yeah, yeah, good point, Kristian. All right, I'm going to stay with the federal versus state authority for a minute, go back to you, Greg. And so the problem, assuming you don't like federal preemption, and I'm guessing a lot of the Democrats on the Hill don't like federal preemption, you end up with a patchwork of state regulation in this area, and it probably happens before Congress even gets around to enacting a law that might end up preempting a lot of it. Isn't this going

to create an absolute ... I mean, it's always difficult if you're operating on a national basis in all 50 states to have to comply with 50 state laws. It increases the cost enormously, and it's a compliance burden. So how do you deal with that, Greg?

Gregory Szewczyk:

I mean, I'll start by saying I don't see my role as having a policy ... the preference that I push. I just help my clients comply with the laws. I'll say in drawing from the privacy standpoint, we have been fortunate enough to have seen ... although there is a patchwork on the privacy side, we've been fortunate enough to have things coalesced around a general standard. And although it does cause frustration that when a law comes and changes things on the margins, it still causes operational difficulties and expenditures, and especially when something changes and it changes underlying costs and infrastructure costs in particular, we have at least been fortunate enough that we don't have dramatically different frameworks from the state laws that have developed.

Where there's the fear on the AI front is, since we don't have that standard framework that the states are following, if the patchwork grows in different directions, then that would be an entirely different game that companies would have to do. And that would just be a much, much more difficult approach, especially when we're talking about deploying tools and developing technologies. So that would be a much different scenario for our clients.

I'll also say on that first point about what Democrats would want, I've had some conversations with Democratic regulators who are the ones who are ... And I'm trying to think of the right way to put this because they were confidential conversations, but they are in favor of a federal preemptive bill if it could be done in a smart way. The position that they have is these politicians and regulators say when it's not there, they feel that the states need to step in and regulate. But they said that their preference would be, especially for AI, they would prefer to have a federal preemptive bill.

Alan Kaplinsky:

Hmm, interesting, yeah. Not what I would have expected because, at least in other areas, it was almost a knee-jerk reaction on the part of most Democrats to not like federal preemption. Charlie or Kristian, want to chime in on that?

Charlie Bullock:

Sure, just briefly. I mean, I have a whole preemption spiel that I go through, but I won't go through it here. I think that one important thing to note is that preemption is not an all-or-nothing proposition. In fact, historically, it'd be completely unprecedented to say, "Okay, we've got this new technology. We're preempting all state laws related to that technology." How it's always worked in the past with the internet, with airplanes, electricity, whatever you want to name, what happens is you do narrow preemption when reacting to state laws you don't like.

So for the internet, you have the Internet Tax Freedom Act. Okay, we're seeing that there's certain discriminatory state taxes against out-of-state like e-commerce, we're going to preempt that, that sort of thing. There can be no state law regulating the internet. In fact, there's many state laws regulating the internet.

And so I think just because you're not doing this sort of all at once states cannot regulate AI for 10 years, doesn't mean you're not doing preemption. There will obviously be preemption. I mean, even if you don't explicitly put express preemption in your bill, there's always conflict preemption. Any state bill that conflicts with a federal law is automatically preempted. Courts will enforce that.

So that's sort of my perspective, is that I think the better approach than the one that I've heard a lot of Democrats and Republicans, including Senator Blackburn and so forth, recommend is a sort of piecemeal-preemption thing where, okay, you pass a federal bill and then you preempt state laws that are inconsistent with that bill or that you don't like and et cetera. So you go issue by issue and preempt sort of over time instead of trying to do it all ex ante.

Alan Kaplinsky:

Right. Okay. Going to move on now to the area of child protection and consumer risk and go back to you, Greg. And I've actually got two questions I'm going to pose at the same time. The framework places very strong ... I'm sorry. The framework

includes specific proposals to protect children such as age assurance and parental controls. Are these proposals workable in practice? And how do these provisions in the framework compare to existing approaches to online child safety and privacy?

Gregory Szewczyk:

So I think this is one where, although it points to age assurance and parental controls, the actual specific controls and the specific platforms are really going to matter as to how workable it's going to be. When we see the regulations on things like the Colorado Privacy Act's design regulations or age gating for the newer laws related to child protection on certain kinds of websites, the levels of detail are just much more specific and strict. And I think with this kind of guidance, it's hard to say how workable it will be. And we'd really have to see what Congress would come up with as to how workable it would be in practice. And that would go for the second as to how it would compare to existing approaches.

I would imagine that Congress would want to work off of that existing framework just from a legislative standpoint, but I think we'd have to wait and see. And I'd welcome to hear what Charlie and Kristian have on that point as well.

Charlie Bullock:

Just one piece of context that I think is important is that there's an ongoing discussion about child safety bills in Congress. And I think one important one is KOSA, the Kids Online Safety Act. And so if you compare the provisions of what's in the framework to KOSA, it's like a much more preemptive and sort of a less aggressive approach. What KOSA does is it sets a floor, a high floor and says states can go above this. We're only preempting conflicting state laws, but states can go above and beyond this floor we're setting. And what the White House framework would do is set a pretty low ceiling, say, "This is what we're doing on child safety, age verification, et cetera, and states can't do anything else." So that's kind of an important division there.

Alan Kaplinsky:

Okay. Let's turn to another area, which I will call innovation, infrastructure, and economic policy. I'm going to go to Kristian on this. The framework places strong emphasis on AI dominance, infrastructure build out, and regulatory sandboxes. Are these the right priorities?

Kristian Stout:

I think the framework is clearly saying yes to an AI build out agenda. It's obviously focused on infrastructure, faster permitting, making federal data sets usable for AI, and as you mentioned, creating sandboxes where the firms can actually test out their systems and deploy them for regulatory rules that might be too rigid for these instances or applied too early.

I think on balance, those are going to be the right priorities if our goal is to actually accelerate [inaudible 00:37:01] across the economy.

Sandboxes in particular around the world, we've seen successful use of them. They can be really useful for reducing uncertainty for introduction of new products and services without forcing regulators to guess too early about what they think the final rule should look like. It encourages experimentation on that point. I would say better access to government data sets can also meaningfully lower barriers for universities and smaller labs that don't already have privileged data access.

All that said though, there's an important caveat here, I think. An innovation-first framework isn't automatically a startup-friendly framework. Large infrastructure demands, compliance obligations like age assurance, unclear liability rules, things like that, they can tilt the playing field towards incumbents that already have compliance teams, legal teams, a lot of compute, and the ability to distribute.

So if we really want to think about this as a holistic pro-entry framework, I think we need to pair the big picture agenda, which again, I think is a great agenda, but we need to have things like clear safe harbors or something like that, real access to public data sets where possible and lawful, and some restraint in creating big downstream liability for developers based on how third-parties might use their systems because the knock-on effects for small entrepreneurs can be quite large in this space.

Alan Kaplinsky:

Okay. Yeah, thank you, Kristian. Charlie or Greg?

Charlie Bullock:

I think Kristian did a great job of covering that. I guess the only small thing I'd point out is I think existing law already in many cases provides coverage, provides protection against being held liable for something somebody else did with your product. I think the standard is something like if you negligently design the product and then you can be held liable because it was foreseeable that this harm would occur even though a third-party intervened. But this is a nuanced conversation, and mostly I agree with what Kristian said.

Alan Kaplinsky:

Okay. Let's move on to intellectual property and data use. I think this probably won't take very long, but Kristian, why don't you tell us what the framework says about that?

Kristian Stout:

Okay, I'm happy to. It's been a little quiet on the IP front for the last six months or so, but there were some big cases last year that I think we're still going to see more activity on. And the framework is trying to weigh in into the middle of this litigation right now. So everything I say here, the main caveat is take it with a big pinch of salt because where we see the direction of court so far could easily change when these things hit the appeals level and higher.

So on IP, I think the framework takes the right approach for now, which is to say it lets courts continue working through whether training on copyrighted materials fair use rather than trying to settle that issue legislatively right away right now. It's a complicated problem, it has a lot of nuance to it, and I think you really do want fact-finders at least looking at the instances of these cases. Even if ultimately we disagree with where they come out, those records that they're creating are going to be really useful for legislators looking at how these systems are being trained when they do construct some kind of legislative framework.

I think the reason, again, this makes sense is for several different legal questions here that tend to get collapsed together. Training is not the same thing as output infringement. You can have an infringement on the input, you can have infringement on the output, and it's very troublesome, I think, and legally questionable when courts push these things together, so they need to keep looking at this issue and trying to pull it apart.

Stylistic similarity is not the same thing as copying protected expression. It's never been for humans. If we do go down that road with AI systems, like there's a bill that France is considering right now to essentially allow the inference that a stylistic similarity means you train on copyrighted materials. Those kinds of things I think are problematic. We need to see how those facts actually exist in courts though, because maybe my position is wrong, stylistic similarity, maybe we do need to consider it. But anyway, that's something the courts need to pull apart. I think the risk here ... Yep.

Alan Kaplinsky:

Yeah, go ahead, Kristian.

Kristian Stout:

Yeah. So I think the risk here is if Congress steps in too early, there's a real risk it ends up over-correcting and chilling non-infringing learning. I say it in quotes, "learning." While not actually doing a great job at targeting genuinely infringing outputs. These systems are definitely going to be able to be used to infringe copyright at some point. I just don't think it's where a lot of the action in the courts is right now.

And the downside, I would say, of deferring to courts for the time being is uncertainty. Everyone is operating with a moving target, and that has real costs. But in this case, some uncertainty may be preferable to lock in a statutory rule that doesn't map

very well into how the technology actually works. So I think the key line here is the danger is not only under-protection of creators, it's also over-fitting copyright law to a technology we don't yet fully understand.

Alan Kaplinsky:

Anybody else want to chime in on that before we move on to the next section? Hearing nothing. So Kristian, I'm going to stay with you and we're going to talk about free speech and censorship. The framework introduces provisions aimed at preventing government-driven censorship of AI systems. How significant is this addition?

Kristian Stout:

So this edition is in the context of a number of other areas where the current administration is, in some cases, trying to advance what I think is potentially useful policy and, in some cases, part of a misunderstanding of how the First Amendment works. So it's kind of a mixed bag at a meta level.

I think the free speech piece here is one of the more novel parts of the framework. It's basically trying to prevent the government from pressuring AI providers to shape outputs along parties in ideological lines and to create some form of recourse if that kind of coercion happens. It's a callback to some of the allegations that happened during the pandemic where social media providers were jawboned, allegedly jawboned into de-platforming certain people or certain views. And to many people, that's very offensive. This administration is definitely offensive. And so I think that's kind of what they're thinking about here.

I think it's significant even though the details aren't fully worked out yet. It's essentially extending the, like I said, this anti-jawboning concern from social media into AI. The underlying idea here, I guess, is that you really don't want the government quietly steering supposedly independent systems toward preferred narratives. As I said at the top, I think that raises real free speech and administrative law concerns. The First Amendment is a very hard limit on a lot of things that the government would want to do, even if it's trying to theoretically make these things First Amendment-friendly as people might colloquially think about it.

So depending on how it's implemented, it could be just preventing regulatory agencies from contacting AI providers in order to shape their content. It could also be done in a way that looks like jawboning if they do it wrong.

And also, I would just add a caveat here, even though I think as an American, I tend to be a free-speech voluptuary and I really like the First Amendment. There is a risk where if you completely shut off all connection between intelligence sources and AI firms, that you are hamstringing yourself when it comes to things like national security concerns or even domestic terror threats, things like that. So as much as I don't want jawboning, I also don't want AI systems to be inadvertently blind to harms that the government might be able to help them understand.

Alan Kaplinsky:

Right. Okay. Charlie, do you have anything you want to add?

Charlie Bullock:

Yeah. Just real quick, I interpret this censorship and free speech section ... I think Kristian's right, but I also think it's sort of what they're asking for in some sense is like there was the woke AI executive order that President Trump put out, I don't know, a year ago or something like that, where it's essentially saying ... I think they were mad about things like the Google Gemini thing, or maybe it was Bard or whatever, where it was creating the images of Founding Fathers and they were all diverse and stuff like that. And so I think that's the basic thing that's got them upset and that they're trying to address here. That was my understanding at least. I'm not the world's expert on this, but that was what I thought was that it was sort of what they're asking for is a congressional version of that woke AI executive order saying you can't force models to be biased politically, you can't force models to be express certain ideologies or whatever.

Kristian Stout:

Yeah. Well, that's kind of the First Amendment division or wrinkle that I was mentioning. If you're doing it as you're not allowed to strong-arm companies, that probably can pass muster. If you're doing it as you're not allowed to make your models woke, I don't see how that survives First Amendment scrutiny.

Alan Kaplinsky:

Yeah.

Charlie Bullock:

Yeah, that's [inaudible 00:46:11].

Alan Kaplinsky:

That's a good way to put it. Let's turn to the regulatory structure, and want to go to you, Greg. The framework rejects creating the idea of there being a new AI regulator, some new regulatory agency that just focuses on AI, and instead it relies on the existing agencies that we have and sector-specific oversight. Do you think that's a sound approach?

Gregory Szewczyk:

Yeah, I mean, there's going to be pros and cons with either of those approaches, but I do think it's a sound approach. I mean, for the most part, these are the regulators that are going to be familiar with the underlying concerns that we're looking to solve for. And so the real issue is how does the use of AI or any other new technologies, how does that fit within the concerns that these regulators are already regulating? And so to that extent, it makes the most sense to have those regulators be the ones who are enforcing the new either regulations, rules, or however it's going to be drafted.

Where you have the setback is what we've seen in certain scenarios is you have a regulator who already has a full plate who then just has a new set of issues put on their desk without more resources, without more help, or more staffing. And I think that's kind of the drawback is if it's being put on an existing regulator just from an operations and logistics standpoint, do they have the expertise, do they have the resources to handle it? But I think from a theoretical standpoint, I think it is probably, at least in my opinion, the sounder approach, but there are, like I said, pros and cons from both directions.

Alan Kaplinsky:

Right, right. Yeah, I sort of worry about the lack of expertise problem because the agencies are never going to be able to keep up with what's going on in the industry, never. I mean, they just can't match it. And that's true with not just AI, but with all kinds of new technologies.

And there is something a little bit attractive to me, at least maybe it's superficially attractive in having one agency where it's all AI nerds, really have people who they know they can make a lot more money working for one of the Mag 7 or one of the other companies, but they elect for whatever reason to do good and work for the government, but it's certainly not clear cut.

I mean, one example of specialization is, and I don't know what you can draw from the example, but it's CFPB. Congress thought there needed to be a specialized agency that's just focused on consumer finance because of the problems of 2008 with the mortgage meltdown and how the other federal agencies purportedly dropped the ball and were not able to stay on top of what the investment banking firms were doing with all kinds of mortgages that people did not deserve to get. So I don't know. I have mixed feelings about it.

Anybody else want to comment on that before we really get to what I consider the final topic of the day?

Kristian Stout:

Yeah, I would just add that the approach in this framework is pretty consistent with American law as compared to, say, European law. It's a widely used way of approaching subjects, regulatory subjects. Privacy tends to be something that we don't

have an overarching privacy regulator. We instead have it in different regulatory bodies where privacy intersects with their mandate versus, say, the European style where they have stood up centralized privacy regulators.

And to your point, Alan, I think that you do see basically a bunch of privacy nerds in your pilots running the privacy agencies in Europe, and they tend to overweight towards the value of privacy in all legal and regulatory contexts. So that, for instance, you see tensions emerging in Europe where sometimes regulators want to impose interoperability requirements, but then the way the privacy laws are interpreted, they make it impossible to have interoperability requirements. I kind of don't like a lot of the requirements that everybody wants to impose, but their regulatory agencies sort of run into walls that they can't get over because you have agencies that maximize for one value alone.

In our context, you could either have AI nerds who want all AI all the time and completely underweight potential risks because they're just so blinded by the upside, or you could have AI doomers that run the agencies and they do everything they can to sort of mitigate the deployment of AI. And what you really want is someone who has a balanced view of the costs and benefits of AI against all the other values that our law tries to optimize for.

Alan Kaplinsky:

Yeah, I'm wondering if there might be a middle ground here where you create, I don't know whether it would be called an independent agency or not, but you create some federal body that deals with AI and has oversight over all the other agencies and how they're dealing with AI. And there's a requirement on the part of, let's call it the other agencies to consult with the experts, the specialists that are in this consultative kind of agency. Anyway, just to ... Yeah, go ahead.

Charlie Bullock:

One thing I'd say is that there's things you can do short of actual regulation that the government can still play a part in. I think the government has a role to play, for example, setting standards. We have CAISI already, the Center for AI Standards and Innovation, and that's not a regulatory body, it's within NIST in the Department of Commerce. And what they do is they test models, they do evaluations, they set standards. That seems like it doesn't really impose any burdens on innovation to me. All they're doing is labs can voluntarily submit their models for testing if they want, they don't have to.

And so I think funding and codifying CAISI would be one way to sort of get some nerds in government, build up expertise, build up government capacity to understand AI without doing a lot of regulatory stuff, European style privacy regulation that would actually hold back innovation.

Gregory Szewczyk:

Yeah, I just want to throw out there real quick that is four guys sitting around talking about AI, we're sure pretty cavalier about labeling other people nerds.

Alan Kaplinsky:

Well, I think there's actually, if I'm not correct, isn't there a gentleman by the name of Simon Taylor, who he actually uses the word nerd as maybe part of the company that he owns, I'm pretty sure. And he had some kind of a conference recently, but I certainly don't consider you guys to be nerds for sure.

Kristian Stout:

You don't know me well enough then. I have a massive Warhammer 40K collection.

Alan Kaplinsky:

Well, anyway, let's look ahead, and let's go to Charlie. And this is really the last area. I want to divide it into two segments, but what do you see, Charlie, as the most likely next steps in Congress? And if you were advising policymakers, what would you prioritize or change in the framework? And I guess finally, is there anything that companies should be doing now to prepare for the direction this framework suggests?

Charlie Bullock:

So okay, so I'll answer your last question first. What should companies be doing now to prepare for this framework specifically? I would say nothing. I would say devote $0 of your resources to preparing for this because it is not going to pass.

I think that's unfortunate. I would like for there to be a serious deal that exchanged some kind of preemption for some kind of federal AI policy framework. I think a lot of people would like that.

But how I interpret this framework, and feel free to disagree, Kristian or Greg, if you don't agree, I interpret this as sort of a statement that they're not trying to actually pass something. I mean, you have on the one hand, Senator Blackburn's Trump America AI Act is one thing, which is extremely regulatory. It's like 300 pages long. It has all these regulations. And then on the other side of the spectrum, you have this, which is quite the opposite. It's, I think, pretty far outside the window of what you could potentially get seven Democrats to agree to. I don't even think you get 30 Republicans to agree to this.

So what I expect to happen is that Representative Scalise will introduce a bill in the House. I think they've already announced that's going to happen. The fact, it will not pass. I think there are other bills that are also being considered that are not part of this White House framework, but that are similar to it in some ways. So you might see some of these priorities getting embodied in other bills that have maybe a better chance of passing.

But yeah, I know that Representative Obernolte in the House has a framework that he's fond of that's kind of an alternative to this and that there's some interest maybe from House leadership in adopting a form of that as well. That would be different. It would have other things. My understanding is it has things like whistleblower protections and it has things like transparency, SB 53-style transparency requirements that this framework doesn't have. So I think there will continue to be some action.

If I had to predict, unfortunately, what's going to happen overall in the AI policy landscape before the midterms, it would be probably nothing significant passes just because the base rate of Congress passing any legislation that's not like literally the NDAA is quite low. I know there is going to be a reconciliation bill to pay for the Iran War, is my understanding, but I think they're keeping that very skinny is what I've heard. So I don't think any preemption will be included in that, which means you'd have to pass it with 60 votes, which means that unfortunately I don't think anything is getting done until after the midterms then.

Kristian Stout:

So I want to agree and maybe disagree a little bit with Charlie. I think I agree that given the current level of functionality of Congress, I don't see much legislation passing before the midterms.

But where I would push back a little bit is Neil Chilson had an article, I think, last week where he was talking about, or maybe two weeks ago, where he was talking about this framework, and he framed it well, I think. He said, "You should think about this as a term sheet in a deal. You put a bunch of terms down. This is kind of what our ambitions are, and it's room for negotiation." And I agree that I don't believe that there's going to be one big omnibus package that includes all these things and you're going to get a kumbaya moment and everybody suddenly realizes that what I will call the pro-innovation view of AI is the one that we should adopt.

I think what these are is terms that the White House would really like legislators to consider as they do have opportunities to include topic-relevant pieces of legislation, and that you should look at the margins for these kinds of ideas being introduced into legislation.

On the broader question of will we get the moratorium, I would like to be more optimistic on that one. I think that the political pushback between ... I think the left just pushes back on anything that Trump says. So even if it's a good idea, I think they're going to push back on it.

And I think that there is a division between the pro-White House branch of the Republicans and the presidential hopefuls that also split on that moratorium issue, which means that your base of votes that would vote for that is very low. So unless there is some strange new deal that can be offered to the splitting Republicans, I don't see how you get the moratorium in the foreseeable future.

Alan Kaplinsky:

Okay. So here's the final question, and I alluded to it toward the beginning of our podcast show, and that is that a tremendous controversy erupted a couple of weeks ago. I think it was initiated originally by Tom Friedman, an opinion writer for the New York Times, wrote an opinion piece where he described the problem, and he said it was sufficiently important to bring this to the attention of everyone, even in the midst of the war with Iran and everything that was going on there.

And of course, hearing that from Tom Friedman, who at least I respect a great deal, and I think a lot of people have a lot of respect for him, he's a very smart guy, you pay attention to what he has to say. And he taught the controversy, of course, centers on claims by Anthropic about a powerful experimental AI model, which is often referred to as Claude Mythos.

For the listeners who aren't familiar with it, let me describe what Anthropic said, and then I'd like to get a reaction from each of you to the extent that you want to react. Anthropic says the model can autonomously identify and even exploit vulnerabilities across major operating systems, including browsers and websites, reportedly finding thousands of flaws, including previously unknown bugs. Tests suggested that even relatively unskilled users could generate working exploits with the model's help, raising fears that it could democratize high-end hacking.

Reports claim that the model at times evaded controls or escaped sandbox environments and publicized exploits, intensifying concern about containment and misuse. Because of these risks that Anthropic identified, it refused to release the model publicly, and it limited access to select partners and working with government agencies to study safeguards.

On the other side of the coin, there are some cybersecurity experts that have cautioned that the claims may be overstated or partly marketing-driven and noted that AI mainly accelerates existing techniques rather than creating entirely new ones.

Who wants to take a first shot at this? I mean, is this something that we ought to lose sleep over, or is this thing blown out of proportion?

Kristian Stout:

I don't think we should lose sleep over it. I also don't think it's blown out of proportion. It seemed fairly inevitable that models would get to the point where they could discover mass vulnerabilities and software just based on the trajectory that they're on. Because I mean, essentially we need to think about it is it's not that Mythos is doing something wholly new. Teams of computer programmers and hackers do the same thing. It just takes them four months to actually review a code base and figure out where all the vulnerabilities are, and they have automated tools that they use to help them do that.

So what this really is is taking the time horizon that a skilled hacker can use these models to find vulnerabilities down to an hour maybe. And it also introduces the possibility that your average Joe on the street can now fire up an LLM and have it discover a vulnerability in his wife's phone or something when he wants to spy on her.

So it opens a new horizon of potential cybersecurity vulnerabilities, but it doesn't open something new and unprecedented. They're not hacking reality or something. And this creates a cybersecurity problem 100%.

I think that there probably is a little bit of a marketing hype on expressing exactly how important this is on the Anthropic side. It's really good for fundraising to say, "We've just invented a model that can literally do everything." That's great for them. And their whole business model has been trying to find ways to partner with government and really sell safety stuff. So this is in that realm.

I don't think it's completely false though at the same time. I think that the way they're rolling this out, probably what we see in the long-term is a change in business practices. We are accustomed to seeing models just being deployed as soon as they're ready to consumers. If they're really getting that capable, probably what you're going to see is a delay. You might see access only to government enterprise customers until some fraction of those capabilities, people at the top levels feel pretty reliable in the sense that giving it to Joe Plumber down the street, he's not going to hack his wife's phone with it.

Once it's been handled, then I think there'll be an incentive to roll it out. But I think mostly what this is going to mean is we'll have delayed model releases, we'll have more careful scrutiny and partnership between government, a collection of these corporations that are developing these models to make sure that they're safe, and that we're going to see some changing business models among the firms themselves, probably more incentive to create enterprise relationships, at least for the near-term on these model capabilities.

Once you see adversaries start to develop these capabilities, then that creates a counter pressure to be able to deploy these capabilities domestically as defensive measures. So that will be maybe in six months to a year we'll be having that conversation.

Alan Kaplinsky:

Right, right. Charlie, do you want to add anything?

Charlie Bullock:

Yeah, I think Kristian makes some good points. One thing I think we do know is that the zero-day vulnerabilities that Anthropic's reporting have identified are almost certainly real. They've sent these to companies saying, "Hey, patch this vulnerability in your system," and the companies have done that. It would take a big conspiracy for a bunch of different companies to say, "Oh, okay, we found all these vulnerabilities that didn't really exist." Unless Anthropic is just lying through its teeth, it's clear that this model is very capable.

The craziest thing to me about this story is that the training run, or at least the training for it, apparently costs $10 billion, which is orders of magnitude greater than anything we've heard before. It's just absurd to me that they spent $10 billion to train a single model. I mean, a hundred million used to be the absolute maximum, and it was kind of speculative that that would ever get reached. So this is just absolutely wild.

And it's crazy that increasing the scale of the training produces sort of this kind of step change in capabilities, which I mean, it has big implications for the future of AI development. What happens if you have $100 billion training run or something like that? So yeah, no, I think it's a very interesting story.

And I think one ironic part of it to me is that this is happening right after the Department of War declared Anthropic a supply chain risk and tried to basically stop them from ever doing any national security stuff. They're saying, "You are such a danger to national security that you cannot be involved in our supply chains at any level." And then a week later, Anthropic's like, "By the way, we discovered the single most important cyber defense tool of the century or something." It's kind of ironic that that happened four days apart or something.

Alan Kaplinsky:

Yeah, I noticed that. And the other thing I'm wondering is whether it has something to do with the competition with OpenAI. OpenAI is supposed to go public sometime before the end of this year. And is this Anthropic's way of throwing some cold water on that? I mean, I don't know, I have no idea. Greg, do you want to comment?

Gregory Szewczyk:

Yeah, I mean, I agree with what Kristian and Charlie just said. And I think it's just another reminder, we're going to keep seeing stories like this, we're going to see stories ... I think in the coming years, we're going to see advances in AI-related cyber defenses and AI-related cyber threats. And it's not just going to be related to AI-driven tools. We see more and more about threats from quantum-related cyber threats and quantum-related cyber defenses. And as the technology advances, we're just going to keep seeing that. And to Kristian's point, I don't think it's getting blown out of proportion, but then there's always a little bit of what's the marketing hype behind the companies behind it and what's driving it. But that's not to say don't take it seriously.

Alan Kaplinsky:

The one thing I think this is going to do, and I don't know if there's been any polling yet done to figure out how the American public is reacting to this, but I think it feeds into the fear that a lot of Americans have that AI is going to lead to massive layoffs, that this is really a technology that we've got to really closely control because a lot of people, I think if you were to poll them, they'd say they don't like AI because they're fearful they're going to lose their job.

Anyway, let me close. We've come to the end of our program. Let me close with some takeaways I have from today's discussion. First of all, the White House has clearly moved from a principles-based approach in last year's white paper to a

more concrete and policy-driven framework. And it reflects what I would call a maturation of the debate and a recognition that Congress will need to act.

Second, there appears to be a notable shift toward prioritizing innovation, competitiveness, and national leadership in AI alongside more targeted safeguards, such as those aimed at protecting children and protecting fraud.

Third, the framework's strong endorsement of federal preemption signals a potentially major development in how AI will be regulated in the US, particularly given the growing number of state-level initiatives.

And finally, while the framework answers some of the questions, it leaves others unresolved, especially in areas like intellectual property, liability.

And by the way, on the issue of liability, and I was unaware of this until I heard about it very recently, the American Law Institute is engaged in a project where they're not creating a new restatement of law, but they are creating AI liability principles, and it's in very early stages. And I would recommend that anybody who has a lot of interest in that area, that they contact ALI and get into what's called the consultative working group. And you just have to say you want to be part of it and you're notified of the meetings and you know what they're up to. Practical implementation of safeguards, as I said, that's not really dealt with.

I want to thank Charlie, Kristian, Greg for sharing their insights, helping us think through these important issues. And to our listeners, as always, I want to thank you for joining us on our podcast show today. And I can assure you this will not be the last podcast that we're going to do on AI. And indeed, I'm already in the stages of planning a podcast with the reporter for ALI who's heading up this new project so that we can find out what it's all about. So I hope everybody enjoys the remainder of their day. Thank you again.