Center for Strategic and International Studies

TRANSCRIPT

Event

Exploring Global AI Policy Priorities Ahead of the India AI Impact Summit

Panel 2 - Unpacking the Global Challenges and Opportunities for AI Governance

DATE

Friday, January 30, 2026 at 10:30 a.m. ET

FEATURING

Dean Ball
Foundation for American Innovation

Russ Headlee
Senior Bureau Official, Bureau of Cyberspace and Digital Policy, U.S. Department of State

Shana Mansbach
Foundation for American Innovation

Dean Ball
VP of Strategy and Communications, Fathom

CSIS EXPERTS

Aalok Mehta
Director, Wadhwani AI Center, CSIS

Transcript By
Superior Transcriptions LLC
www.superiortranscriptions.com

CSIS | CENTER FOR STRATEGIC & INTERNATIONAL STUDIES

Aalok Mehta: Hi, everyone. Welcome back from the coffee break. I hope you've been enjoying the program so far.

My name is Aalok Mehta. I recently joined CSIS as the new director of the Wadhwani AI Center. I'm really happy to have a distinguished panel here today to talk about AI governance and its intersection with the India Summit.

So, as you know, we've seen a rapid pace of development of AI technology. We've also seen rapid developments in the AI governance and regulation space as well over the past year. This includes the EU AI Act coming into force, kind of. It includes new sort of legislation at the U.S. state level, especially in California.

We have a new administration in the U.S. that has taken a markedly different approach to AI regulation than the previous administration, and we've also seen countries like Korea and Japan implement their own sort of national AI frameworks and laws.

And so, we're here to discuss the – where things go from here and how the India AI Summit might change the conversation or influence the trajectory of governance.

So, I am going to, you know, turn to the panelists. I'm going to sort of give them all an intro question and a chance to introduce themselves, and I'm looking forward to the conversation.

So, I'm going to start with Russ Headlee, who is the senior bureau official at the U.S. Department of State, Bureau of Cyberspace and Digital Policy. So, we've heard from the Indian ambassador and the French ambassador about what they're hoping to get out of the summit.

I would love the perspective from the U.S. government as well or from the Department of State in terms of what you're hoping to see out of the summit and especially what you're hoping to see in terms of AI governance.

Russ Headlee: Thanks, Aalok. Thanks a lot to CSIS and the Wadhwani Center for hosting. It's great. It's an honor to be here. It's particularly an honor to be sitting next to Dean Ball, who's been a real thought leader in this area.

I've got a lot of things going for me in my current position, leading our Cyberspace and Digital Policy Bureau. One of the things I have going for me is I started this position on July 23rd, which is an auspicious day in Trump administration AI policy history. It was the day that the

president launched America's AI Action Plan, and Dean was one of the chief architects of that. So we, I, have been a real beneficiary of the good spade work that Dean did during the first six months or so of the administration over there.

One of the great things – I think in a lot of ways, India's AI Impact Summit presents an opportunity. There's a lot of convergence between what India intends to do with the summit and a lot of the themes that are in the AI Action Plan.

Their focus on impact, their focus on practical use cases, and their focus on diffusion to the developing world in particular really dovetails nicely with the third pillar of the AI Action Plan that Dean and the White House team teed up for us.

Secretary Krishnan from MEITY was on Greg's podcasts a couple of months ago and really highlighted his narrative to an American audience, which is the kind of audience that tunes into CSIS podcasts, which everyone should do, is very much in line.

Not coincidentally, he's a good diplomat, and he was definitely pitching an American audience, and I heard a lot of those same themes when I was in India for preparatory events in December ahead of the summit. But, indeed, I think he's quite genuine in a lot of ways in elevating those themes, that India's looking for practical use cases and really looking to lean in on innovation and adoption.

And, indeed, when you look at India's domestic regulatory landscape, they've really put their money where Secretary Krishnan's mouth was on the podcast and, indeed, have taken a relatively cautious stance to regulation and have really leaned into innovation and adoption.

So, we will look – we'll look to emphasize those aspects and hope that India will serve as a role model to a lot of other developing countries as they're looking to build out their own regulatory frameworks. There are some bad models out there in the world, from our perspective, and India's is a relatively good one in a lot of ways.

Dr. Mehta: All right. Thank you.

Next, we'll go to Shana, Shana Mansbach, who's vice president of strategy and communications at Fathom, which is an organization that does a lot of thinking about the connection between AI regulation and governance and trust in AI technology.

So, Shana, the Delhi summit is the first time that one of these major summits has been held in the Global South, and so I'm really interested in your perspective about how that might shift the conversation or change the conversation around AI governance, going forward.

Shana Mansbach: Sure, and, first, I just want to say thank you for having me. This is a particularly fun panel for me to be on because I'm with a former and now returned member of my current organization and a current member of my former organization at the State Department. So feels like a homecoming here.

I'll throw a little bit of cold water in this conversation. You know, we've had three years of summitry, and I think without a ton to show for it, at least in terms of AI governance, so I want to be realistic about what we can achieve at the summit.

That said, these summits are where the agenda is set. The host country has a tremendous capability capacity to set the tenor and the focus of the agenda of the conversation, so it really does matter that this is happening in Delhi.

Right now, the conversation – the global conversation around AI is highly disjointed. We just moved from safety in Bletchley to investment in Paris and Korea now to impact, which I think most of us understand to mean application – what does this actually mean in people's lives.

And on a country by country basis, I think we're kind of all over the place. Washington is focused on acceleration. Brussels is focused on safety. China is, largely, focused on the CCP – I'll leave it at that – and the global majority, to the extent that we can characterize or generalize, is focused on access and inclusion.

So, I think it's unrealistic to say that we're going to have some grand bargain, that we're all going to agree on the same policy things, nor do I think we should necessarily. You know, we're all sovereign countries. We have different ways of doing this. But there are real costs to regulatory fragmentation for public companies, for – excuse me, for the public; for companies, particularly small companies, you know, small developers that don't have armies of compliance lawyers to help them navigate many different systems here; and also for the regulators themselves that face capacity constraints – and that's probably true of many Global South countries.

So, I think in Delhi we have an opportunity to coalesce, again, not around some grand bargain here, but around a governance framework that introduces some interoperability into the system. I'm thinking

about a technical backbone that makes compliance legible across borders.

So, at Fathom we're very focused – and my panelists have heard me talk about this before, and I should say before getting into this Dean Ball is really the thought leader, maybe even the godfather of this idea. But we're very focused on independent verification, so this idea that governments – whether it's a state government or a country government – should set some outcomes around what they want to see, societal outcomes. So India can say here's what good looks like. The United States can say here's what good looks like. France can do it too. But then you have technical third-party auditors that are actually verifying whether this is indeed happening.

The logic here is the, you know, labs shouldn't grade their own homework. This is very good for trust. It keeps up with the rapid pace of the technological development itself. So, we think the technical – the technical architecture is there to – you know, to operationalize this approach.

And I guess I would just finish by saying, you know, frameworks only work if everyone is incentivized to participate. You know, nuclear nonproliferation treaties don't just work because, you know, fancy people signed them and they're down on paper; it's because everyone's incentives are aligned. So I think in order to have success at this summit, at any other summit, we need to come together around a framework that people – or, actually, countries are incentivized to uphold, where it's creating a trust architecture that's good for the public, for companies, for innovation, and for the regulators themselves.

Dr. Mehta: Great. Thank you.

Finally, we'll turn to Dean Ball, a senior fellow at the Foundation for American Innovation. Dean, I'd love to press a little bit more on the fragmentation issue. So, we have had, let's say, a spirited debate in the United States about fragmentation at the state level and sort of what that might mean from a compliance and burden perspective, especially for small companies. Obviously, there is an international analog as well of fragmentation at the – at the global level and different countries adopting sort of incompatible, perhaps conflicting regulatory regimes for AI technology. Are there specific things that you would like to see out of the summit that might help alleviate this issue?

Dean Ball: Yeah. So, I think, you know, there's fragmentation of – and thank you, by the way, for having me, and thanks to both of you for the kind words. Undeserved.

But I think there's, like, a – there's a fragmentation of – yes, of legal frameworks and regulatory frameworks and whatnot, but I think most importantly there is a fragmentation almost of mentality. My view is that when we talk about AI policy what we are ultimately referring to is, like, essentially all policy, all regulation of all economic activity in 10 to 15 years from now, because it's that general of a technology. And so there is this drive, you know, I think most principally associated with Europe but also of some American states and places like Korea, Japan – well, Korea has flirted with this too – to create this, like, almost fully general risk-management framework for, like, all of AI. And it just seems to me that that is – that is so broad, and just so general and abstract, that I just – like, the notion of AI risk management to me almost is too capacious to be useful.

And so I think that what's, like, much wiser to do, and I think what the action plan tries to do and what I think sound AI policymaking strives for regardless of its country of origin, is to really think about concrete actionable things that we can do that sort of advance our interests along relevant margins. And to – you know, as we say, to the language models. This has been a mantra of mine for two years. But you don't actually say this anymore, but it used to be, back in the era of prompt engineering – in ancient history, two years ago – (laughter) – you would say, you know, take a deep breath and think step by step to get models to perform better. And I think we ought to prompt ourselves in a similar way.

So, as an example, to make this concrete, like, I think you see – there's a – there's a one regulatory framework that's emerged in the United States, came out of California. It's called SB 53. It's basically a transparency bill that says, there's a few categories of risks at the frontier of AI. Not all of AI. Not facial recognition. Not auto-driving cars. Not drones. Not – just quiet your mind for a moment and just think about the frontier language models, right?

Not every use of them, but just some uses. And it is credible. There's a credible threat model that these things advance serious catastrophic risk, which is number one – that's the number one purpose of government is catastrophic risk mitigation, right? In bio, in cyber, and in this category that we might call model autonomy, which really – it's not within the Overton window to say what that really means, but I'll just say it. What it means is the automation of AI research within frontier labs, right? And so we need to understand more about the nature of how labs are evaluating these very complex risks and how they're mitigating them. And so let's have some transparency

about that. Let's have some transparency and some structured incident reporting and things of that kind. That would be one.

And then another example would be, like, you know, the Colorado SB 205 algorithmic discrimination framework, which is itself based on the European Union's AI Act, which is, like, we are going to regulate all potential instances of anything that we consider discrimination in all uses of AI, which is basically all modern software, which will undergird all economic activity. And it's like, oh my God, how do we implement that? That seems very hard. These are the different mentalities. I still think – when I engage with a civil society outside of the United States, and much of it within the United States too, I still hear an awful lot of the, like, we're going to think about everything, all potential risks, all at the same time, and we'll have a fully general framework that, like, accommodates all of it. I just think it's not going to work.

So there's a competing – there's a competing set of mentalities. Does that mean that SB 53 is, like, a perfect bill and it's great? No. But I would say it's an example of a kind of thinking that – in fact, many people in the Trump administration don't very much like SB 53. And you can have reasonable debates about that. But it's an example of a kind of thinking that I think we need to see more of. A great – another area would be, you know, one of the major themes of this year will be the widespread diffusion of agents, AI systems that can operate for hours at a time doing entire human jobs, in many cases, or very large fractions of them. This is not like a hypothetical, maybe one day in a few years. No. It is like that is happening right now. And it will happen much more by the end of this year.

And the security of those systems, right? Their vulnerability to penetration, their adversarial robustness against things like prompt injection, how should you – as someone deploying agents, what is the notion of responsible use? In a similar way to, like, what is the notion of driving a car responsibly? It's like, OK, well, don't be drunk, put on your seat belt, keep your eyes open, et cetera. Like, these sorts of basic things, right? We have no standards for that right now. We totally could have standards for that. We totally could have, like, actual technical standards – not standards – not standards – things masquerading as standards that are really regulation, but like actual technical standards that are very sort of NIST-shaped, National Institute of Standards and Technology, that can sort of help people, deployers and developers of AI systems, build more secure agents.

So, like, let's do that. But I think concreteness and actionability, if that's a word – actionableness are things to aim for, as opposed

to these kind of abstract things that make me feel like I'm reading Wittgenstein.

Dr. Mehta: Let me ask a quick follow-up to that. So supposing we adopt this approach of sort of addressing different pieces of AI in different ways at different times. Is it better for that to happen at the national level and sort of roll up into sort of an international agreement? Or is an approach that starts, sort of, with a multilateral approach and then just, you know, trickles down to different countries, would that work better?

Mr. Ball: It's the answer is there's, like, a little bit of everything there, right? There's going to be some. I think particularly things that affect the deployment of AI systems, things that are affecting, like, you know, the ability of businesses, banks, hospitals, whatever, to deploy AI systems in different ways. Like, I think you want that to have a local flavor, right? You want – and so that's the kind of thing that I think, you know, different cultures, different regulatory climates, et cetera, et cetera will – it's fine. It's fine to have there be diversity there. That's the kind of regulatory diversity that we're quite used to. You know, whether or not it's, sort of economically optimal is separate from, like, different countries have different cultures, different regulatory cultures. Hospitals in India are regulated differently from hospitals in Minnesota. And that's fine.

Then there are things that really implicate the development of the systems. And I think once you get into that category, you start to get into things where – and not just – not just the development of the systems, but also, I would say maybe slightly more precisely, obligations which are principally placed upon the developers of AI systems. That gets a little bit more complicated, and is maybe something where we need to have – there needs to be some sort of different notion of jurisdiction, right? So in the U.S. context, I have often said that that's the kind of thing that we really need to do at the federal level. We cannot have 50 different state regimes that govern the development of AI. That is, indeed – America is, indeed, unfortunately, heading in the direction of that kind of sort of quilt-like outcome. But we'll see. We'll see what happens there.

But, you know, you asked a question that was also more about global context. And I think that, like, the global regulation of frontier AI development is really something that we should avoid. We should try to avoid multilateral institutions that govern, you know, OpenAI, and Anthropic, and Google, and Google DeepMind. I think if you would like to regulate frontier AI development, then it would be good to have a frontier AI development industry. And otherwise – (laughs) – you should sort of leave that to the countries

from whence the technology emerges. And so I – when I hear people talk to me about sovereign AI, I find that there's a very large correlation between people that talk to me about sovereign AI and people that talk to me about wanting to regulate technologies that are from countries that – you know, from foreign countries.

They want to – like, they want to regulate American companies, and they're not Americans. And they also talk to me – and they say that America impinges on their sovereignty. And it to me, it's, like, well, hmm, wait. This sort of seems like maybe the regulation of American companies is an American prerogative, right? So that would be, like, one area that I think is going to be a fundamental tension. Not just at the India summit, but, like, over the coming century, right? Like, how do we deal with the jurisdiction of things like this? But, and maybe in the long run it is some sort of international agreement, right? Maybe that's the only equilibrium that works. But I'm kind of opposed to multilateralism on things like that at this point.

Ms. Mansbach: Can I just throw one thing in here?

Dr. Mehta: Of course.

Ms. Mansbach: I mean, I think – I totally hear you. I think there is a category, though, of domains where AI regulation may merit a more traditional, hard law approach. And here I'm thinking about red lines, right, where the logic is a lot more, like, nonproliferation, you know, AI systems and targeting, like, human targets in warfare, or something like that. That feels like the type of thing that I'd love to hear everyone's thoughts about this. Sorry to step on your toes and do your job. (Laughter.)

Dr. Mehta: I want this to be a dialogue, so it's better if –

Ms. Mansbach: But that feels like one of those – you know, talking about concrete things that can come out of India, I mean, that seems like a diplomatic issue that is a perfect setting for policymakers in India to take on, where we know that there are a couple of red lines. Maybe we have to adjust them as time goes on because, you know, technology adjusts. And we look to the Chemical Weapons Convention of 1993 for examples here, where we say there's precedent. There are some things we – probably everyone can agree are pretty bad. And there – this is – these are domains where, you know, actions are irreversible, where, you know, defection costs are really high and they can be possibly verified. I guess my question to the group – and you can join in, too, and see – you were talking about this – but my question here is, you know, to you guys, does that seem like an appropriate area for diplomatic agreement? Because, I mean, clearly, I'm showing my cards here. I think that there that totally

makes sense for there to be some – a list of things that we can and cannot do, which is not something I think we can do in almost any other domain of AI.

Mr. Ball: I think – I mean, I think what you're describing to me sounds a lot like use of the technology as opposed to development of the technology, right? Like, you can't – don't let AI – you know, U.S. and China have agreed, for example, to not have AI in the actuation of our respective nuclear arsenals. That seems like every country in the world should agree that, fine, we can all agree. We can – we can all agree on stuff like that, reasonable enough.

I think there's things on – areas like biosecurity would be a great example of an area where I think we can – we can develop common standards for biosecurity, maybe even cyber, things like that. I think that's all great.

I just think that if your starting point is, like, I have a gigantic regulatory appetite to micromanage the affairs of businesses that are thousands of miles away from me in a different country, I think that is just like, you shouldn't do that. You just shouldn't do that. And of, like, America, I had not appreciated it as an American for quite some time, and I feel very thankful that for the first time we have an administration in the Trump administration that is candid and honest with the rest of the world about the fact that it's – that we don't like it. (Laughs.) So, you know, I hate to be – I hate to be uncouth in saying that, but it's really true. There's probably nicer ways I could put it.

Dr. Mehta: Russ, let me – let me turn to you for a follow-up on that. So the Trump administration has taken this sort of deregulatory sort of pro-innovation approach, somewhat different than the Biden administration in a number of ways. When you engage either bilaterally or multilaterally, how do you present the American vision? And what are – what sort of agreements or what sort of discussions are you hoping to have in those settings?

Mr. Headlee: So one challenge presented by the – by the regulatory stance that we have taken is that it is much easier as a diplomat to walk into a ministry of foreign affairs somewhere with a demarche in hand and 800 pages of regulation and say, I will send you this as a soft copy, and you can find/replace our country name for your country name, and if you are looking to adopt a regulatory framework have I got one for you. Europe is quite effective at doing this, right, not only – even in cases where they have perhaps strangled an industry such that they no longer have products to export. They are still quite effective at exporting their regulatory structures and frameworks – (laughter) – in ways that I

think actually now are presenting a cautionary tale as parts of Europe or EU member states actually look to slow roll a bit implementation of the EU's AI Act and are, I think, awakening in some regards to the second-order effects or to – as Vice President Vance described it in Paris, that they in some cases have strangled an industry as it was coming up, regulating it before they really had an industry in some cases to regulate – which I think gives us kind of a natural hook, because a lot of countries around the world when they're thinking about taking on how to regulate an industry do frequently look to the EU first.

Dean mentioned Korea. I was recently in Vietnam. Vietnam recently passed – their national assembly passed an AI law that is modeled in some ways on Korea's. So it's not – it's not just the EU. It's not just Europe that presents a model this way.

But there have been cases – one of – one of my hobbies is carpentry, and there's a saying in carpentry, or those of us that have ever held a flashlight for our dad as he was working on something may have heard it, that you – that you should measure twice and cut once because you only get – in a lot of cases; it's definitely true in carpentry and true in a lot of other aspects of life – that if you – if you cut too hastily without sort of taking the full assessment of what it is your work looks like, you can make mistakes.

And I think that picks up on a theme that Dean was pointing to there and this is one of the regulatory ideas that we are trying to export that countries really ought to – there's a lot to recommend an incremental approach. There's a lot to recommend measuring twice before cutting because, again, like, I think Europe presents some cautionary tale here.

To bring it back to India once again, I think in a lot of ways India is to be commended for the regulatory approach they've taken domestically where I think in a lot of cases the regulatory framework is not perfect, from our perspective, but they have taken a relatively incremental – relatively incremental approach that I think a lot of other countries would benefit from following.

Mr. Ball: I just – I think that's really beautifully put in many ways. I use measure twice, cut once all the time, by the way.

But, like, my view, this is – my goodness, do I not speak for the American government when I say this but, like, my view of what is happening here is that there is this, like, truly monumental moment in human history that is occurring with the development of this technology, a bright line

that we are crossing that never shall we return past to the other side of that line.

We're in the process of crossing it. We're going through the phase transition right now and there is no stopping it.

We don't know what it means in a million different profound and, in some case – profound ways that in some ways should, like, make you quiver in your bones that we don't know where we are going, right?

That's true, my posture when I go to a place – like, when I go to something like a summit, when I talk to foreign governments is, like, please just exercise a little bit of restraint here, just a little bit of humility, because we really don't – there's so much we don't know.

And so I know that it's scary that we don't know all this stuff but also, like, you will – the chances of you acting now in some brash way and making it worse are much, much, much higher than you acting now and, like, doing the exact right thing that you should have done, right?

And so, we're at a moment where the dice are in the air and, like, breathe the wrong way and you change the trajectory of history and so, like, just – you know, just be light. Just be like a ballerina. That's all I ask of you.

Ms. Mansbach: Sounds like we need a living, breathing, flexible framework around independent verification.

Dr. Mehta: Shana, I wanted to come back to the – to the question of trust.

So, one of the things we see that's a big dichotomy between, say, the U.S. and Europe and other places in the world is on the question of trust. Both trust in the – in AI as a technology that it might make people's lives better but also trust in governments to sort of take appropriate action related to the technology, including regulating it in appropriate ways.

And so, we see that's pretty high in India in both ways, and so I was wondering if you could talk about, you know, why that might be the case and sort of what lessons other countries can take about the trust the Indian people in the Indian government seem to have.

Ms. Mansbach: So, I think the answer to the question is trust, at least the Indian flavor of trust, exportable, the answer is it's really, really hard. I think part of that is there's just – there's higher institutional trust in India than there

is in America. So, I guess let me take a step back and say a couple things.

The statistics here are crazy. For anyone who hasn't paid attention to this, there's something like – and maybe it's – let's hope I get this right – 14 (percent), 15 percent of Indians are – only 14 percent of Indians are more concerned about this technology than excited. In America, it's a majority of people who are more concerned than excited.

India has an 89 – I think 89 percent of the people are – trust their government to regulate AI. All the Americans here, I mean, just, like, take a moment. Eighty-nine percent of people trust their government to regulate AI.

Eighty-nine percent of Americans don't trust our government to do most anything I say as someone who's spent my career in government here, and the numbers are much lower for Americans. Something like 44 percent of people trust American – the U.S. government to get this right.

So, there are a couple wrinkles. I mean, Americans – the way that we use this technology, and when you ask, you know, have you read a lot or heard a lot about this technology, in America it's much higher than it is in India. So, there's sort of a salience issue here, too.

But I think more fundamentally the fact that this split represents, you know, some institutional trust that's already built in, that's just really hard to export, particularly in a country that's already using this technology a lot.

I think that – you alluded to this – what America and other countries can do, it comes back to this idea of a trust infrastructure. You know, it's funny, when people say, are you pro safety, are you pro innovation, I say yes and trust is the glue between the two. They're two sides of the same thing.

It is really, really important for people to – and I'll explain what I mean by people, but it's really important for people to be able to have confidence in these systems. They need to know that they're safe, that they are secure, that they work as advertised.

Having that sort of – that confidence is great for our innovation, particularly for small businesses. It's good for our regulators. Having a trust infrastructure, again, I think is the glue between, you know, safety and innovation here, and I think it's something that we have a lot of room to run.

You know, you talk to people and they say, I'm just not sure that this thing is going to work as advertised. I'm not confident in it. And I think the more confidence you have in the AI products themselves, the warranted confidence – you know, independent verification, third-party auditing is very helpful here – also increases institutional trust in the government to do its job around regulating.

Dr. Mehta: Thank you. So, I think we'll probably do one more question for the group and then we'll try to take a couple of audience questions before we wrap up.

So, one thing that the ambassadors said is that, you know, this is an impact summit and so they're focused on impact. I think that one of the things we've seen from previous summits is that there's a flurry of activity in the run-up to the summits, announcements happening at the summit, and then sort of a lot of work streams that are supposed to happen that sort of die out or it's unclear what the outcomes of those are.

So, in terms of this summit, are there – are there steps that you're hoping to see that would sustain the momentum for whatever is announced or for whatever happens at the summit that leads to a little bit more durability than we've seen in the past?

I think related to this especially is this is the first summit happening in the Global South. There's a lot of concern that the Global South is not included in a lot of conversation around AI governance, and so – (inaudible) – that sustain their involvement in discussions around AI policy issues as well.

So, I'll open it up to whoever feels like answering some or all of that.

Mr. Headlee: So, I think one of the outcomes of India's AI Impact Summit that would be super useful or complementary from the perspective of our AI Action Plan would be if what India is intending to do in terms of highlighting use cases is really effective.

I think one of the conversations that we have a lot – that our diplomats have a lot, particularly in developing countries, is we're keen, we're interested, we want to have a strategy, we want to do this right, but we're not exactly sure what use cases make the most sense here, and I think India can really – can be exemplary there too because India has used – is already using AI in a lot of ways that I think will be useful for other – for developing countries to see, and that dovetails nicely with

the AI Action Plan because the U.S. hyperscalers and U.S. industry across the AI tech stack are already in India in a big way.

So, I think we've got that going for us as a U.S. delegation showing up in Delhi that the regulatory posture that India has taken – again, not to be a total cheerleader. It's not perfect but it is definitely good enough that Google, Microsoft, AWS are all in the Indian market making tens of billions of dollars' worth of investments that are going to help India meet its AI aspirations, that are going to help India connect the strategy and the vision that they've got to actual application to impact, to their point, for their citizens, for their economy.

And so I think there's – I think we've got a really strong foundation there from a U.S. perspective and from a U.S.-India bilateral perspective to say to countries that look to India as a role model or as an example that there's a lot to recommend the Indian approach if what you want to do is bring the world's finest AI tech stack into your country to help you accomplish your aspirations.

Mr. Ball: I think India has put itself on a very good footing from a regulatory perspective. I absolutely echo that. So my comments about the regulation of companies were primarily aimed at our European friends, not so much our Indian friends. But also, you know, there are other countries that are building up a good appetite – a regulatory appetite, that I think is probably unwise. India is, I think, taking a very smart, light touch approach. I think one thing that India really excels at historically is the diffusion of technologies and the diffusion of things like agents, things like coding agents, across an economy.

You know, we don't – again, we don't – we have only a dim notion today of what, like, responsible use of such a powerful technology is. Like, of what that should look like. I think that over time we'll develop much better intuitions and abstractions. But those things will come from the experience of adoption. And I think that India will be one of the world leaders because of its scale and its historical skill at technology diffusion. I think it will be one of the world leaders in diffusing agents, and therefore a critical partner in developing robust standards for things like agent security and responsible agent use.

And so, I'm sort of – this is all the kind of thing that happens on a voluntary basis, for the most part, through standards organizations and things like this. It's a pretty traditional tech governance process. But this is the kind of thing that I would be – that I'd be very excited to see advanced in incremental ways at the summit.

Ms. Mansbach: And I'll close this with a much more hand-wavy answer, which is I personally, professionally, emotionally, spiritually, would love to finally just collapse this fake dichotomy, this false choice between innovation and safety. I really think that if you went around the world you would be hard pressed to find many actors who would say, no, more trust in the system is bad. I think trust and creating a trust infrastructure, a flexible trust infrastructure that makes sense based on different country contexts and priorities but that gives people confidence in their own governments and their own regulators in these products that's good for adoption and innovation, is a smart thing to do.

And we live in a world where there are these camps where, you know, if you say the word safety it means, you know, you must hate America. It's not good for innovation. I think that's a ridiculous – it's a ridiculous place to be. You can have both. Of course, that means you have to stake out a real balance here. You know, you can't be all the way on one side and say I have both things here. But if you take a measured, balanced, flexible, living, breathing approach to this you can create a trust infrastructure that's good for safety and good for both adoption and innovation.

Mr. Headlee: So, I think inadvertently we all dodged your question, which is about –

Dr. Mehta: That's fine.

Mr. Headlee: Which is about summit (theme ?) and deliverables, right?

Dr. Mehta: It's not the first time it's happened in a panel.

Mr. Headlee: So, part of the challenge here, and those of us – I guess we've all worked in government, right? So when you've got a principal going somewhere, when you've got a big summit, what if I was – if I was an Indian civil servant or diplomat right now, I would probably be dedicating 80 or 85 percent of my time to thinking about the outcomes of this event. And there's just a natural delta between the kinds of, sort of, objectives and goals that are, sort of, capital, S-M-A-R-T smart, that are time bound, that are measurable, that are attainable, S&R (sp) , and the kinds of things that really sing when you put them into summit deliverables document, or you put them into your – to your leader's speech. I think all of what we're talking about up here are relatively practical. They are kind of smart goals, but they are not likely to be the sorts of things that wind up in Prime Minister Modi's speech at the end of this thing, or in the – in the joint statement at the end. So I think some summit fade is natural.

From a – from an American perspective, I suspect what will be quite sustainable, my hunch is, there's going to be, I think, quite a showing

from U.S. industry in in Delhi for the summit. I suspect there will be announcements and deliverables there. Not that private sector deliverables are met 100 percent of the time to the letter of the – to the letter of the announcement. But I suspect – those I suspect you will see endured and acted – you will see them endure. You will see them acted upon, in a way that, again, I think, is a demonstration to the developing world in particular, that the real action here, in a lot of cases, the real meaningful action, is going to be – is going to be deal based, it's going to be bilateral, it's going to be about those companies that can actually help countries execute their visions. And I suspect there's going to be some powerful demonstrations of that in Delhi around the summit.

Dr. Mehta: So, we'll take an audience question. I think we have a runner. And I think I saw that hand up first. So, then we'll probably have to wrap up.

Audience member: Thank you for the informative session, panelists.

With regulators, and especially regulatory bodies, being limited to local administration where these companies are located, how does one mitigate risks with hallucinations and output validity, since these systems are deployed globally and guardrails failing would have global implications?

Mr. Ball: So, I'll draw on some experience that I have in government here. At one point when we were working on the action plan, we batted around this – what actually happened, there was an executive order called the Woke AI Executive Order, right? President Trump said during the campaign, you know, we want to stop these models from having top-down ideological bias, right? And we were sort of batting around some ideas internally about how do you effect that, right? How do you effect stopping models from having an ideological bias, right? This is obvious kind of thing that is a regulation that is put upon developers of systems.

But here's the problem. It's, like, we could, in principle, have written the most unconstitutional executive order ever, right, and said: AI models must only tell the truth and have no bias. And how will we determine what truth and no bias means? Ah! Guidelines from the government. We'll tell you what they are. Not only would that be unconstitutional, but the reason that's unconstitutional is because mankind has been on a multi-thousand year quest to define what being free of bias and objectivity and truth, and all those things about what that means. And we have thus far failed to come up with a satisfying answer. And I think we'll probably continue to, if I had to guess. Plato was wrong. And so, like, I think you probably – I think probably I would rebut the presumption, is my point.

And I would say, like, if you wanted to pass a law that said it's illegal for – oh, you may not sell a model in my territory if it hallucinates, well, then, like, you know, the fully general extrapolation of that logic is no one can say anything, because everyone's hallucinating a little bit all the time, right? And the interesting things that human beings say, and AI models often say, are the hallucinatory things – the slightly hallucinatory – just a little bit hallucinatory is usually where the interesting stuff is, because we don't know the truth. So, I would – I would rebut the presumption that that's something regulators should be working on. And I would instead say, like, you should think about having standards for reliability of outputs, or quality of outputs, or things like that, in particular contexts that you care about.

So, what are you worried about? Are you worried about some medical diagnostic being wrong? That's a totally legitimate concern. Are you worried about a bank denying someone a loan application because they used loan historical data from the bank, loan data that was biased in some way because it used to be illegal to issue loans to African Americans in the American context? Totally legitimate concern. OK, well, like, all of a sudden this becomes far more cognizable. Once you bracket the problem in that manner it becomes far more cognizable to deal with this through existing regulatory frameworks, such as discrimination. Again, I'm speaking from the U.S. context here. I apologize for being a chauvinist in that regard, but it's what I know best. Consumer protection, things like this.

So, I would counsel someone with that concern to make the concern more specific in response to a defined threat model, and then sort of react that way. So that's how – and once you do that, oftentimes what you find is, like, oh, my concern isn't actually with the parameters of the neural network. I don't need to affect the developers', you know, matrix multiplications. What I need actually is something that is a far more approachable level of abstraction than the matrix multiplications.

Ms. Mansbach: I was going to very much push back on this, but – or, what you said, but then you took it to use cases, and I'm with you. I mean, it's – there are real harms happening around hallucination and bias, and I think government has an interest in keeping people safe. And if we see certain things happening, I think it's irresponsible for government to say, well, we don't know how to do this; we shouldn't do this.

The one thing I want to add to the mix is the real capacity problems when it comes to government. I mean, all of us have been or are in government. I truly believe in my bones government is well-intentioned and full of people trying to do the right thing, but AI is breaking a lot of

things, including governments. This technology moves really, really quickly. It is hard to get this stuff right. We don't have the capacity. And frankly, the people who know best how to regulate this technology, to the extent that anyone knows, are in the labs.

So, throwing out one – you know, the thing I would end with or throughout here is to build up the capacity – maybe this is a resourcing thing – you got to get those people out of the labs. You have to give them other options. You know, it's something that I'm really focused on. We're thinking about, you know, how do you create centers of excellence, verifiers, independent organizations that are doing really good work and who can actually do the technical work once government says, you know: Here's some use cases that we're really concerned about; who can do that? But right now, I think that capacity is lacking, and it's certainly true in the United – or, it's true in the United States and it's certainly true in other countries around the world.

Mr. Ball: One slightest bit of pushback to that is just this – you said that, you know, there are harms happening from AI, let's say in medical diagnostic, you know, things like that. And, like, I'm willing to believe that's true. Like, I'm willing to believe that during the time we've spoken somewhere in the world an AI model has made a medical diagnostic decision that was, like, wrong compared to what we'll ultimately find out is the truth. But, like, do you know what I'm also willing to bet you? That the number of humans – human medical practitioners who made, like, profoundly wrong – (laughter) – diagnostic decisions during the period that we've been talking is, like, orders of magnitude higher, right?

So, the expected value of human error during the period of this panel versus the expected value of AI error during the period of this panel, you know, one thing dwarfs the other is my guess. So, like, which is the problem about which we should have moral panic?

Ms. Mansbach: Sure. And, Dean, I don't think anyone's saying these things have to be perfect. If I were setting – going back to –

Mr. Ball: No, people totally are saying that.

Ms. Mansbach: Sure. Sure.

Mr. Ball: People are – people totally make –

Ms. Mansbach: I myself on this panel am saying I think a reasonable thing to do, if I were setting up a system here, I would say, OK, would government –

government has to come up with reasonable outcomes. If I were the government writ large, if I were in charge of creating these outcomes, I wouldn't say bias can never happen, no one can ever die. The mathematical solution to that is, well, then you shouldn't have the systems at all, which is ridiculous and I'm certainly not arguing that. What I am arguing is I think it is perfectly reasonable for a government to say: If and when these systems are deployed in high-risk contexts they cannot be, let's just say, worse than a human being. That, to me, is not only incredibly reasonable; that's incredibly verifiable.

Mr. Ball: Well, yes, it is, but then, like, the question is, how – if that's true, why is it that, like – if that's true, then what we should empirically expect to see is that Europe in particular should be filled with people – there should be hundreds of thousands of young people in Europe who are doing empirical econometric work on the error rate of all sorts of human institutions for all sorts of things, right? We should have really robust data on, like, the medical diagnostic error rate of every single medical profession, every law enforcement decision, every government decision. We should have – we should know what the human failure rate is really, really well if that's true, because we should be trying to define that because that would be what good looks like. That's the baseline. AI's got to be doing better than that. How come we don't do that?

Dr. Mehta: I do – we do have to move –

Mr. Ball: I encourage – I encourage regulators to do that. That's my point.

Ms. Mansbach: I think that if you had a system that was outcome-based that said here – you know, not just here's the list of things you should do but was actually outcome-based, you would see that would help catalyze science around this. I would love for that to exist. I think it is completely – it's eminently possible. There are wonderful people working on all this stuff.

Mr. Ball: Oh, it's totally possible.

Ms. Mansbach: They need to be more resourced. But when you have legislation like the EU AI Act that says here are all the things you should do, there is no incentive. There is no catalyst to do that type of research.

Mr. Ball: Totally agree. Totally agree. Claude code can do a lot of what we –

Dr. Mehta: We do – we do – I'm getting the signs that we need to wrap up.

Russ, I wanted to give you a quick chance to say anything to wrap up. Otherwise, I'm happy to move the program along.

Mr. Headlee: Let's keep it moving. Really looking forward to the – to the summit. Looking forward to seeing a lot of the folks that are in this room there.

Dr. Mehta: All right. Well, I'm going to pass things off to Rick Rossow, who is a senior advisor and our chair on India and emerging Asia economics. And he will be moderating the next panel, which is "Industry Perspectives on AI Innovation in India." So, give me a hand in thanking our panelists. (Applause.)

(END.)