

**MICHELLE
ENGLISH:**

Greetings. I'm Michelle English. And on behalf of the Center for International Studies and MIT Mexico, welcome you to today's Starr Forum. Before we get started, I'd like to remind everyone that we have many more events planned for the semester, including one this Friday at noon on the philosophy of human rights. Details for this event and others are available in the foyer. And for those who haven't already, we also have a sign up list so you can receive email notices on all of our events.

In typical format, today's event will conclude with a question and answer session. For the Q&A, I'd like to ask everyone to please be mindful of time and to ask only one question. And we will be using the microphones during Q&A. And please identify yourself and your affiliation prior to asking your question.

It's truly an honor to have with us today Luis Videgaray to discuss the challenge of AI policy around the world. Dr. Videgaray is Director of the MIT AI Policy for the World Project and a senior lecturer at the Sloan School of Management.

It is also an honor to have with us Kenneth Oye. He's joining the conversation. Professor Oye is both a professor of political science in the School of Humanities, Arts and Social Sciences and of data systems and society in the School of Engineering. He is also Director of the program on emerging technologies at the Center for International Studies. At this time, I'd like to invite Professor Oye to the podium to provide introductory comments and to formally introduce our esteemed speaker, Dr. Videgaray.

[APPLAUSE]

KENNETH OYE:

So thank you, Michelle. It is truly a pleasure to be offering an introduction for Luis. By way of background, the gentleman that you see sitting here in the front row has a long, and distinguished, and sometimes sordid history. The distinguished part is an undergraduate degree from [INAUDIBLE]. The sordid part is a doctorate from MIT, and Luis also served as Foreign Minister and Finance Minister of Mexico.

I have to tell you a little story, by way of introduction. Luis spoke to a group of 85 senior American officials. And after his presentation, a three-star general came over. And he said, Professor, is Luis available to become Secretary of State? And I indicated that that might be

difficult to arrange. But his duties have included working on the successor to the NAFTA at the USMCA on a range of issues putting out financial crises and handling problems and disputes on everything from immigration, to conventional investment, and trade policy.

Luis will be speaking today on AI and development. And this topic takes something that is near and dear to the hearts of all of us at MIT-- AI-- but to have a serious discussion of the implications and effects of the technologies being developed here and elsewhere for development prospects around the world.

The format will be a presentation by Luis. We'll have a little bit of a conversation up front, and then we'll turn to you for questions and answers. And as Michelle indicated, please identify yourself, and we'll have a microphone for you to speak. Luis, thank you so much for joining us today.

[APPLAUSE]

LUIS VIDEGARAY: Thank you, Ken. Thank you so much, Ken. I'm very proud to be here, very thankful to the Center of International Studies and to Ken in particular for inviting me to be at the Starr Forum today. Thank you for all of you for being here. I see many friendly faces. Thank you for the students from Sloan that are here.

And I want to single out somebody who's here, that, to me, is very special. 35 years ago-- actually, I'm going to mention two people, because I just saw one that he just came in. But I want to start with my thesis advisor.

It was-- Jim-- my thesis was completed 22 years ago. And I had the privilege to go through MIT as a student with the guidance of Jim Poterba, one of the best teachers, not only of economics, but about a lot of things in life. And it's a real honor for you to be here Jim. Thank you so much. And I also--

[APPLAUSE]

Any by the way, your new office looks a lot better than the old one, because the building is new. So it's great. I also want to mention and acknowledge the presence, here, of the Consul General of Mexico. Thank you, Consul General for being here, Alberto, my friend. It's an honor to have the highest representative of the Mexican government here with us today.

So Ken was very--

[APPLAUSE]

Ken was very nice in his presentation. And believe me, I have a lot of-- let's call them war stories about my career in government, about the US/Mexico relationship and many other things. We're not going to talk about that today. Of course, we can always do that. Please feel free to approach me and we can always talk about any of those issues, but we're going to talk about artificial intelligence today.

And in particular, we're going to talk about artificial intelligence policy, and how is that evolving around the world, why is that important, and where are we going, or are we getting somewhere with that? So let's get started. And this shouldn't be long.

So what do we mean when we say artificial intelligence policy? That's the first question, because there's a lot of concerns about AI. And so I'll try to provide, here, not a definition, but actually a collection of themes that would constitute a coherent or a comprehensive AI policy. And as you'll see, we'll start with a blank page, and it will quickly fill up with lots of themes and lots of complex issues.

First of all, one key topic of AI policy is the use of AI in the delivery of government services and goods. And this is something that is gaining a lot of traction around the world. Perhaps a little bit of less visibility, but is extremely important. In a lot of places-- health care education-- our core activities, core services provided by the government.

In many of the countries like my own, the government conducts anti-poverty programs. But also certain core services, like, say, tax collection can be aided and can be delivered with the support of AI. So the use of AI to have a better government is an important first thing. By the way, this is how I got involved with machine learning many years ago. About five years ago.

I was Finance Minister of Mexico. And as such, I oversaw the governance of our tax collection authority. And we were trying to do things a little bit better. And after several frustrating conversations with consultants, I ended up talking to computer scientists and realized that these new-- not so new methodology called machine learning could actually help in doing things a little bit better.

It's a complex implementation, but there's a lot of potential in having better delivery of services

in doing that. By the way, there there's interesting research here at MIT, particularly Media Lab on how to do better delivery of anti-poverty programs with the aid of AI, how to make them more effective, how to make them more efficient.

Second, a key question about AI policy is, how should governments invest or use public resources in AI? Is this something that should be completely left out for the market, or should government step in and support both the development and rollout the use of AI in society? Of course, basic question is research and development. And this is a question that looks a little bit different in a country like the US, or in a developing country.

It goes beyond that. You have the question of, should the government be supporting investment vehicles for startups? Venture capital funds with public resources. This is a model that is very much used in Asia, particularly China. You see a lot of venture capital coming from the government into industry. So that's an interesting question.

Education. Should the government be doing more in funding education in computing, or what about tax breaks and other incentives for the use of AI in the economy? I'm talking about the economy, third big question, third big theme is AI in the economy.

And a lot of the anxiety, a lot of the conversation about AI is about jobs displacement and the inequality that it can create. And it's a very valid topic of discussion. Here, MIT has a very strong task force which is the Work of Future Task Force that is preparing a report. There's already a partial delivery of that. And we'll talk a little bit more about it in a second, about how AI interacts with the job market.

Does it create new jobs? Is it will displace human beings? Those questions are central to any AI policy. But economic questions go beyond that. It's also market power. It's also-- that means concentration, antitrust policy. Today we are seeing news from Europe on data policy. Just a few hours ago, the European Union issued a new communication on how they're going to deal with internet platforms. And even things like algorithmic collusion in financial markets, which is a new topic, but some things that there's already some indication that these things might happen. So a whole block of features on the economy.

Then, we have the social responsibility issues of AI. Things like privacy-- the fact that the type of AI that has exploded is the statistical type, it's a machine learning, and in particular, deep learning, it consumes lots and lots of information. And so how do you deal with that model of learning with the right to privacy and to a private life? So that's a big discussion.

Also, the issue of fairness, and bias, and the possibility of discrimination. We know that algorithms can be biased for several reasons. Explainability. Some type of algorithms, in particular deep learning algorithms are hard to interpret. Hard to understand how they get to the predictions that they get to. And this can be quite frustrating and quite important in some settings.

Think of a judge that is basically a decision on a recommendation by an algorithm. If the judge cannot understand why is that recommendation being done, that's a problem. But you can think of that also in the medical context, you can think of that in the job market.

Robustness, that is the consistency and how resistant the algorithms are to either random variations in the world, but also adversarial intentional attacks on their predictions, the question of accountability. Then, we have a fifth block, which is AI and democracy. And the use of these sophisticated techniques of learning very granular information to manipulate the minds of consumers, but also of voters, and to influence politics in democracies, influence the opinion polls, influence, of course, elections.

We have the questions of surveillance. How machine learning can power new tools of surveillance starting with face recognition, but goes well beyond face recognition and the emergence of a surveillance. Some people talk about surveillance capitalism, some people talk about a surveillance state. But this is clearly an issue. And of course, AI enabling authoritarian regimes. Technical authoritarianism that we see in some places, some very important places around the world.

And finally, the geopolitics of AI. This is a reality. This is a little bit of the elephant in the room. Not everybody likes to talk about the geopolitical dimensions of technology. But there are clearly rival models around the world, that we see a model of technology deployment in China, we see a different model in Europe, and we see an emerging model in the US. According to culture, history, these are different models. And countries around the world realize that.

There's a lot of talk about a technological decoupling happening as we speak. Just go to the general press and you'll see-- almost every day you can read about a coming AI Cold War, or a lot of concerns on national security and these division of the world in two separate camps.

By the way for a country that is neither China, nor the US, nor part of the European Union, this is a problem. Because it means, where do we stand? Think for a moment of a Latin American

country, or an African country, or a country in Southeast Asia. Is this about choosing sides?

So you have these six blocks of things. We could have a lecture on each block. We could have a lecture on each line within the blocks. So this is a very complex problem, and this is not something that will be solved in a single report. This has got to be a collective conversation. It's going to take a while. It's going to be years in the making.

There's a seventh topic that I've added in the middle, which is sustainability. And we should keep in mind that machine learning is a computationally-intensive technology. It needs a lot of electricity. That means that there's a relevant, an increasingly irrelevant carbon footprint about machine learning. So we need to keep that in mind. I would add that as a seventh topic. OK.

So how are we doing in terms of developing a consistent effective AI policy around the world? As you may imagine, this is a process that is just starting. And the first fact that I want to point out-- and these are all stylized facts-- is that the world of policymakers, on one hand, and the world of computer scientists are very different, and are far, far away.

This means, first of all, that there is an information lag. So things that concern computer scientists today might become concerns of policymakers a few years into the future. Think about the question of privacy.

Privacy has been an issue-- a very important issue for computer scientists, for a long, long time. And it's been a policy issue much more recently. And you can also, almost on every topic, see there's a lack.

But it's not just the lack. These are technologies that are quite complex. If you really want to understand deep learning, you have to understand a lot of math. A computer scientist will tell you it's not that complicated.

Well, if you're not a computer scientist or an MIT scientist, it's hard. So there needs to be some translation. And who does that translation is also introducing noise.

Who does the translation? Well you have the general press. Sometimes, it's good. Sometimes, it's not that good.

By the way, I strongly recommend *MIT News*. They do a lot of coverage of what happens here at MIT, and also *MIT Technology Review*. They do a very good and balanced coverage of what's happening around the world of machine learning. A lot of work, multilaterals.

The World Bank, you'd be amazed how many countries approach the World Bank, or the OECD. Even nonmembers approach these organizations seeking advice, looking for advice. We at MIT know that, because then the World Bank or the OECD come to MIT to ask questions. Not to me, but to people who really know about this stuff.

And this a good translation. Then we have the think tanks. Some think tanks are good. Some are very partisan. So there's a little bit of a mix.

Then you have consultants. And consultants are jumping into this opportunity because there is a need for knowledge and consultants are everywhere. And they're making some very strong claims. And I'll show you one in a minute.

And then of course, one of the biggest sources of this translation of knowledge are the tech companies, which is good, because the tech companies are very strong in their knowledge. But the problem is that they are not unbiased. They have an interest. They might have a conflict of interest in trying to influence through knowledge spreading actual policies.

And finally, there is a language gap. There is a lot of hype, a lot of buzzwords, a lot of things that people write in the correct context, in the original papers or essays and makes sense. But taken out of context, just don't make a lot of sense.

And you start reading you everywhere about ecosystems of innovation, leapfrogging opportunities. That's a favorite of Ken and I. And you see a lot of technical terms used incorrectly. There's a tendency to very rapidly adopt these buzzwords. I'll give you just a couple of examples.

I don't know if you're familiar with this book. This is a book of two authors from MIT, from the Sloan School, Erik and Andrew. And it's a very good book. It's a book written back in 2014, *The Second Machine Age*. And it is a brilliant book.

It's actually today, six years after being published, a very good read still. And if you can see by the title, they were announcing a second machine age. Does anybody today talk about a second machine age? Not really-- not in the policy world. What people talk about is the Fourth Industrial Revolution.

By the way, I know Klaus Schwab very well. We're friends. This is not a bad book. But quite frankly, this is a much better book. But the platform in which it was not only published, but

publicized, was much more powerful.

Of course, the World Economic Forum is a loud voice around the world with our convening power. So nobody talks today about the Second Machine Age. A lot of people talk about the Fourth Industrial Evolution.

Just cast a question. What about the previous three? Nobody knows, but people talk about the Fourth Industrial Revolution.

I give you another example. MIT, as I mentioned, is working on the Work of the Future Taskforce. In the fall, in October, they released this report. And it's a very good report.

If you have a chance to read it, I strongly recommend. It is not final. There's still a lot of work pending.

But here, for example, if you're interested in what's going to happen to truck drivers with autonomous vehicles, particularly long-haul trucking, what's going to happen, here you'll find a very balanced approach. And clearly, the evidence is nowhere near claiming that all jobs in truck driving and long-distance driving are going to be lost. That is not the case.

At the same time, literally the same month that that report was published, PBS released this documentary on AI. And it has an entire 40-minute section full of anxiety, essentially announcing that truck driving is going to end in a very short period of time. And that is not science-based.

It's a lot of hype. Just ask me, which of these two materials has been more influential, has been more watched or read. Obviously, it's the documentary on your right.

And then let me just skip to another favorite. And I'm not going to name the consultant-- this is from a very large consulting firm. These two blocks, I've added them on purpose to hide the name of the consulting firm, because I'm not going to say something nice about them. But they are making these kinds of claims.

To me, this is a phenomenal claim. So these consultants have a toolkit, a responsible AI toolkit, that enables organizations-- that means governments, that means companies. Enables organizations to build high-quality transparent, explainable, and ethical AI applications that generate trust and inspire confidence. OK, we're done. We just need to go to them.

[LAUGHTER]

Quite frankly, it's not that easy. Just spend a morning around CSAIL, or go to the Media Lab, and you'll see that these are extremely difficult issues to deal with. This is not as simple. So let's think, what can a policymaker learn already? And I want to share with you three things that, today, a policymaker can actually learn about the state of AI policy, or that or the framework towards AI policy, around the world.

First of all, it's well established that there's a need for a policy guardrails. If you go back to the '90s and the emergence of the internet-- and some of you might be familiar with section 230. And it was it was thought-- perhaps looking backwards, we can now claim that it would be naive. It was thought that just the ability to share more information and to consume more information was going to be good for everybody.

Well, it turns out that it was very good. But also, it had problems. It was not all good.

And we see those problems very, very clear. We see it in politics. We see it in market concentration. There are many problems with that. So ranging from the CEO of Alphabet, the parent company of Google, to the president of MIT to the European Commission to the White House, everybody agrees that some policy guardrails should be there to ensure socially responsible AI.

Number two, and this is something I think is consensus. This would not be just solved by that guys that know computing and by that the scientists in the field. This has got to be a very interdisciplinary conversation, and a very inclusive conversation.

When you hear the voice and the thoughts of people from different backgrounds-- not just from different disciplines, but also different cultures, come together and define a policy. This is not something that just-- you cannot just go to a computer lab and say, OK, get me some AI policy. This has got to be a collective conversation.

And then number three-- and it is important, because in the last three years, a good number of principles, or declarations on principles, for AI have been published. It depends on who you consult. But clearly, there are at least 80 in our count in the project that I lead. We have identified more than 80.

And there is already a literature on the documents on principles. So you can read papers criticizing, comparing, synthesizing. The one I recommend the most comes from the Berkman

Klein Center at Harvard Law. They just published last month a very good, comprehensive review of these principles.

And not all these documents are the same. But clearly, a consensus is emerging. So I think that right now, we're still working on principles that have very, very small marginal returns.

We need to go to the next phase. So the key question is, what comes next after your principles? I'm not saying that principles are important. I'm saying that we've got a lot of progress there we need to make this the next step.

What is a problem with principles? Well, principles are not enough. Using economist parlance, principles are a necessary condition, but not a sufficient condition, for policy.

Why is that? Because policy is about making hard choices, by the way, in uncertain conditions. And some of these principles can create a tension between them. So you want your algorithms to be very accurate so that when a machine learning algorithm is predicting cancer in an X-ray, you want to be very accurate.

But you also want it to be explainable, and you want it to be fair and without bias. You want the information to be secured, so there is no risk to privacy. You want the jobs of radiology safe, so no radiologists are going to go away. There are many objectives that might be conflicting with each other. This is all about the trade-offs.

And policymaking-- I'm not a scientist, so I talk to scientists now. But I come from the world of policymaking. And I can tell you, policymaking at its core is about understanding trade-offs and making tough decisions.

So what's next after establishing the principles? To me, the key question is about the trade-offs. What are the trade-offs involved?

So before going into some of the trade-offs and explaining what I'm talking about, we just say it's completely unfair and absolutely inaccurate to say that computer scientists don't care about these societal effects, about the ethics of AI, or the economic impact of AI. That's completely wrong. In fact, computer scientists have been working for years on these issues, and have developed some quite sophisticated tools about them. A few examples-- for instance, on privacy, you have just a couple of examples.

Differential privacy has been there for four years. It's a notion that's been there in computer

science. You have edge computing and distributed learning, including federated learning or split learning. There are many other types of ways in handling the data in the algorithm training process to protect privacy.

On bias, there's a whole literature on constrained optimization, imposing requirements like the quality of false negatives or things that would prevent certain groups from being discriminated against. Or there's a whole literature on the diversity of the training data. And under explainability, there's a whole set of techniques for post hoc explanations, analyses that I've done after the training of an algorithm. And there's also the methodologies for making the algorithms more interpretable, more explainable, like transparent design.

And these are just examples. And we can go on like this for robustness and for accountability. I'm being very unfair by just putting in a slide what is really a literature of many, many people working for years on this.

So I'm going to show you what these tools look like. So if today you go and talk to people who really know artificial intelligence, and say, I'm concerned about privacy or I'm concerned about fairness and bias, how can you help me-- they won't give you a box and say, OK, here's an algorithm that is unbiased, or here's an algorithm that protects privacy. Well, it will give you is something that looks like this.

What are these? These are Pareto frontiers. These are combinations that are efficient.

For instance, this one is from a thesis from a recently graduated PhD student from Media Lab, probably Mexican, Alejandro Campero, and these shows a trade-off between the utility or the accuracy, of a machine learning algorithm, and the degree of privacy protection. So there's a trade-off. If you want more power, more utility in terms of accuracy, you lose privacy, and vice versa.

This one comes from a book I really recommend. It's called *The Ethical Algorithm*, from Roth and Kearns, two researchers from UPenn, just published last year. And here, they show a Pareto frontier between the unfairness of the algorithm and the accuracy of the algorithm as expressed by its error. So you see here that making the algorithm more fair results in losing accuracy.

For some applications, losing accuracy is not that much of a problem. But for some applications, it's life or death, or it's really an important decision that hinges on this. So these

are not perhaps the things that people expect from the tools.

Let me just give you an interpretation of these. What I'm saying here is that technical tools won't give you a straight answer, but will give you a menu of options-- a menu of options that are efficient, clearly. Let me just go back to here.

You don't want to position yourself with a policy that gives you here-- that you are under the frontier. So it will give you a menu of options, but will not offer policy decisions by themselves. In other words, these techniques answer a question with a better question. And for AI policy, it's how we answered that question that we get back. So when we say, give me an algorithm that is fair, give me an algorithm that is explainable, you'll get a sense of a trade-off.

And the key question to me on defining AI policy is how do you set up a framework that will allow you to give a good question? First of all-- and this comes back from a little bit of what I learned from Jim while studying economics, you need to understand the nature of a trade-off. And I'm using here two concepts transplanted from microeconomics.

One is the elasticity of a trade-off. What do I mean? If you have a curve that is very vertical, a frontier, then you can gain a lot of protection of privacy or gain a lot in fairness without losing much in accuracy if your algorithm. If the curve is flatter, it means that the trade-off is real. So a theoretical trade-off is not enough to present a problem.

The actual problems come from that slope of the curve. So a policymaker that gets these kind of answers, the first question that you should ask is, what's the slope? That's a boring question. It's an odd question. But it's an important question-- what's the slope of a trade-off if you really want to understand your options?

And the other one is the curvature. Is it a convex frontier where you see diminishing returns? If you see that, clearly the trade-off points towards an optimal that isn't going to be a [INAUDIBLE] solution where you're going to have to balance. So if you have a curve that has a slope and it's curved, you're going to be most likely needing a compromise. You're going to need to-- it's not going to be very effective to be at the extremes.

More important than that-- and this is a key concept that, if I wanted you to take away something out of this conversation, I would mention this one. The key challenge is actually not pinpointing to the curves. It's to create an institutional design that would allow for democratic decision-making about the trade-offs.

And why do I underscore the word democratic? Because technology is not just about technology and technologists. It affects us all, and we want to have a setup where we decide what is more important based on the opinion of society as a whole, not just either the technocrats or the scientists. There needs to be a process of institutional designed to get there.

One more slide on the trade-offs. The trade-offs are not just about accuracy, but about other things. What about innovation? What about international leadership? A lot of people are concerned that imposing restrictions on machine learning and things like privacy, fairness, explainability, will slow down innovation.

And if you slow down innovation, you might lose leadership in the international arena. Actually, you read that a lot. If you go to the general press, you read that a lot. And if you go to Washington-- I've gone there, to Capitol Hill and the White House, you'll hear this a lot.

And this is a very understudied question. I haven't seen any paper showing what is the empirical relationship between this, what is the theoretical relation about this. This is all based on assumptions, and quite frankly, sometimes in emotions.

So let me let me just-- five final thoughts on regulation. I'm going to try to land these into more specific thoughts on regulation. First of all, regulation is better to build from existing frameworks than from scratch. It's important to assess existing frameworks for, say, consumer protection, and build from there, rather than trying to have a new law and new complete legal instrument about algorithms.

Second, most regulation makes more sense if it's sector-specific. So having just an AI act is probably going to be not very useful. You need to work through the sectors. So it's better to look at health care. It's better to look at consumer finance. It's better to look at mobility and transportation.

But some common rules, of course, might be beneficial. It's very important to acknowledge that there are many questions we don't know the answer for. So this might not be the best time to be making hard-core commitments to certain types of regulation. The use of temporal frameworks, like sunset clauses, sandboxes for regulatory experimentation, preemption periods instead of just outright bans of technology-- this is relevant for instance for face recognition, these make sense.

Actually, some states in the US are being criticized for punting or kicking the can. That might not be such a bad idea today. The state of New York is doing that. The state of Vermont has actually a pretty good framework for studying the issues better before imposing regulation.

That is not necessarily about idea. Pre-market testing makes sense. We do that for drugs, clinical trials. We've done it for cars.

Why shouldn't we be pre-market testing algorithms in realms where it it's use for meaningful decisions-- decisions that have life-critical or legal or public resources involved? Important decisions, why do we establish pre-market testing, just as we do clinical trials for drugs. There's a bit of a critique of the European model of excessive reliance on individual rights. I think that is not only about individual rights. But it's also about empowering the individual through technology and establishing restrictions on the behavior of corporations.

Accountability is not just a challenge, but it's actually a policy tool. And well-defined accountability helps a lot towards addressing many of these challenges. So defining accountability is actually a cross-sectional activity.

And then the last two slides. Beware of regulatory fragmentation. What does this mean?

If you look at Europe-- it might be controversial, but I think it's that they are doing something which is remarkable. They're going through a very cohesive, consisting process of establishing regulation, first on privacy. And now they're moving into actual algorithmic decision-making. They just made announcements today about that.

If you look at China, China has a single policy-- very clear national policy. The priorities might not be the same as in Europe or the US, but they do have a consistent policy. What about the US? The US does not have a well-established-- certainly not legislation.

There are some drafts in Congress, in Capitol Hill. But it's the states that are making steps towards establishing regulation. So regulation in the US can emerge to be quite fragmented. So California is running away with regulation, and other states are moving in different directions.

How do the states look like? Like that. So is this a problem for Facebook? Is this a problem for Google?

I don't think so. They lawyer up. They have enough resources to navigate through this

complexity.

But what about startups? What about students from MIT or from Cornell or Stanford that are trying to start something, and will have to go through, at the extreme, 50 types of legislation to deal with privacy and fairness and explainability? This is not the right approach. I am concerned that the US is moving in this direction.

And my final slide. I think we are having a huge problem with trust. And trust is probably the most important problem we have to defining AI policy. And I want to show you two dimensions of lack of trust.

One here on the vertical axis is trust of technology companies. I think that's very low today. And it wasn't like that just a few years ago.

Google was an admired company. People wanted to work at Facebook. Amazon was cool. Today, they are companies that are very feared, and certainly there's not a lot of trust towards them.

On this other axis, the horizontal axis, I have geopolitical trust. Trust between whom? Between leading countries. How's the relationship-- how's the trust between the two leading countries who are the US and China? It's pretty low.

So if this prevails, we'll end up defining policy in this place. We can call it the corner of fear. This is where emotions of lack of trust and fear. So we will have policies that are dominated by imposing restrictions on the use of technology, restricting the companies, but also restricting the flow of knowledge to cooperation between nations. And you see that a lot already.

And if you ask me, this is where we are converging very, very rapidly. I'm not saying that we should be here. This is probably very naive. There are reasons to have trust issues both in this axis and this axis.

Tech companies have given us many reasons not to trust them completely. And certainly, the geopolitical dimension is also true. But we probably should be somewhere around here.

Right now, we are here. To me, a key challenge on a policy is how we'll build the frameworks, how we'll build institutions and collective decision-making mechanisms, to move us from here to here. Thank you very much.

[APPLAUSE]

KENNETH OYE: So can people in the back see us while seated? OK, good. So what I'd like to do, Luis, is just chat a little bit, for five minutes, and then turn to this voracious audience who have questions I'm sure that they wish to pose. This is also frankly the most diverse audience I've seen at MIT, with a mixture of CCO geeks, Center for International Studies policy wonks, and more than one or two nerds. So this is going to be a great group. So Luis--

LUIS VIDEGARAY: And Sloan students.

KENNETH OYE: And Sloan School students as well. Yeah, those are the ones that are better dressed.

[LAUGHTER]

So Luis, I want to push you a little bit on specifics. You said on the one hand that we need to avoid fragmentation. But you also took note of the need for sector-specific approaches, noting that the problems will be somewhat varied if you go from medical to consumer protection for consumer goods to finance.

And let me ask if there is a tension between those two, in that the sector specificity may require regulations that are quite different when you move from one area to another. Fragmentation, which you would define largely with a reference to geography, but it could also be fragmentation with reference to a degree of incoherence across sectors. So the question that I'd like to pose is on both, actually.

In terms of sector specificity, could you give us the sector that you think poses the most difficult problems in terms of the trade-offs across utilization of the data and the methods, and protection of privacy and equity? And the second question I'll have is really on fragmentation and its effects. But first, on sector specificity, what sector do you think poses the most acute difficult trade-offs.

LUIS VIDEGARAY: I'm going to limit my answer to the US.

KENNETH OYE: OK.

LUIS VIDEGARAY: Because this question, if you present it internationally, is a broader question. But in the US, some of the key sectors that are in play for AI influence very rapidly, health care, finance, already have federal regulation. Consumer protection is federal. So this is certainly something

that is doable.

By the way, it's interesting to take note on how federal legislation on finance emerged. It wasn't originally by the design. In its Federalist nature, the US in the 19th century and early 20th century had a very fragmented financial regulation. That didn't go very well.

So the US didn't have a central bank. And eventually, the Federal Reserve emerged. And in the first part of the 20th century, the federal Congress at some point preempted states from doing more financial regulation. And eventually, the regulation that we have in finance is federal. So I think that the platforms [INAUDIBLE] for consumer protection more broadly.

And also, if you think about the FDA, the F is for federal. So you already have even institutions. And you think about antitrust policy, and you have the FTC. The F is for federal. So I think that this might be a little bit in the opposite way as your question. Going sectoral favors having national approaches because the institutions are already there.

What I'm more concerned is with states coming up with generic AI types of regulation. And we have seen that in privacy. California already has its own CCPA, which is inspired by, but not the same as, the GDPR of Europe. And that cuts across every sector.

And we are seeing draft legislation on algorithmic decision-making over us of algorithms to make decisions or predictions in the whole of the economy. That's the kind of fragmentation that can be a layer that comes from below. It can be very problematic, and very hard for smaller companies and innovators to navigate through.

KENNETH OYE: So I appreciate the point on small companies having difficulty-- having the staff to understand 50 regulations that might bear on privacy and consent and data utilization. Understood. But one of the defenses of the federal approach-- of a fragmented and federal approach, is that in an area characterized by significant uncertainty, complexity, and controversy, there can be benefits having experimentation, having different models being pursued in different areas to learn and see which works best or worst. You don't believe--

LUIS VIDEGARAY: And I think that that has been historically the case. And you see that in transportation. The regulation of cars, safety standards for cars, for many, many years, it was local. It was not federal. And to this date, some of the regulation, including speed limits, are obviously local.

So yeah, there's a learning process. I'm a little bit concerned about two things. First, the US is not operating in isolation.

A lot of things are happening in the world. And you see Europe pushing towards regulation that has various extraterritorial consequences and very influential. And then you have China, which is moving ahead with a different framework.

So yeah, I understand the value of experimentation and learning by diversity. But I'm not sure that, at this particular historical moment, the US has the luxury of time to experiment in that way. I think that fragmentation can be costly in that sense.

KENNETH OYE: So within a US context, you could have a federal government that would be preempting local experimentation or state experimentation. Internationally, that becomes much more problematic. To what extent do you believe that experimentation-- national differences, which some would argue reflect legitimate differences in culture and values, the trade-offs that you're talking about, different countries may value privacy versus utility in different ways. To what extent are we likely to be facing a world of increasing diversity in terms of the policies governing AI? And if so, what are the implications of that diversity?

LUIS VIDEGARAY: I think that diversity is inevitable, and to some degree, very much desirable. And it reflects values and culture and history, and that is important. But what we should be more concerned about is incompatibility and adversarial models. I think that's the big question internationally.

So I was talking the other day with Jiang Wang, who happens to be my office neighbor at Sloan. And he was telling me the concept of privacy didn't exist in China as such, and the word privacy didn't exist in China 20 years ago. This is something that was brought in by this debate that we're seeing worldwide and the emergence of this technology. So yeah, it's natural that beyond the political differences, which are a conversation on its own, culture will play a role.

But I think what is more concerning is to have frameworks that are clearly incompatible. And we have conflicting values in countries that either separate, and we become a world of technological silos or a world of technological conflict. I think that here, there is a lot of room for diplomacy in this context. And I'm very happy that we're doing this as part of a CIS because of that. I think that regulation, again with a sectoral approach, is something that should be part of a conversation in either a global platform, like the UN, or in bilateral negotiations, like the US/China talks, or even in like-minded frameworks like the OECD should [INAUDIBLE].

KENNETH OYE: So I'll ask a couple more questions, noting that we'll be turning to the floor in just a couple of minutes, OK? But looking very broadly at development-- north/south, poor countries, rich

countries, and a technology which is diffusing very rapidly with significant effects on the organization of economic activity and political activity and medical activity. To what extent-- and this is an unfairly broad question. I want to preface it as such.

But to what extent do you think AI will have the effect of providing for opportunities for equalization, limiting inequalities, more rapid diffusion of useful information? Or to what extent are we really talking about technologies which will be owned and captured and controlled by the relatively wealthy, and the largest countries and companies, contributing to further concentration of economic and political power? Net effects, or is this too broad a question to be worthy of answering?

LUIS VIDEGARAY: Well, it's definitely an extremely broad question. But it's not an unfair question. And I can see both effects playing out. And policy and countries, individually, and collectively through diplomatic means, should be focusing on both possibilities. I see tremendous opportunities for equalization on services and the delivery of public goods.

The question of AI in health care is very different in Boston, Massachusetts, than in Bolivia or in Ghana, where here it's a question of quality maybe costs, who pays for it, how to ensure that it's good diagnostics. In other countries, it's a question of access. It's not whether this is better than the existing doctor. There are no doctors or no specialists in many places around the world. So technology can bring health care to places where it is not.

Also, disruptions in education can have a profound effect in agricultural productivity and energy efficiency. So let's not forget that the type of AI that is blooming right now has a lot of potential for good things. And those things should be encouraged and enabled. But also, we are seeing it already, as in previous waves of technological innovation, that there's a trend towards job displacement, and concentration of wealth towards capital and the owners of capital.

Now, I'm not trying to be Marxist here. But it's not a new phenomenon. We've seen that in the 19th century. We've seen that in the early 20th century. And we're seeing it again now.

And the key question is, how are we going to respond as societies? In the 20th century, only after two very, very costly wars and a huge recession, institutions were changed. And there were institutions that were balancing across people and across groups. Right now, it seems that we're moving in the opposite direction. So I think it's a very real policy question.

KENNETH OYE: So I'm going to ask one last question, noting that Michelle should probably get ready for harvesting questions on the floor as well. To move from that very broad question to a very specific one, if we take a cell phone-- which better not go off at this moment and we look to AI, data, and let's call it multiplying the effectiveness of medicine. So some folks are working on a variety of very interesting cell phone applications using data that would take underserved populations-- be they in southern Texas or in the developing world, and providing the kind of concierge medical advice that people get at MIT medical, but don't necessarily get elsewhere. Meaning that the AI is being used in conjunction with measures, personal measures and records, to offer advice and check on whether people are adhering to treatment, but also gathering information which feeds back in. So that would be a beneficial use of the AI, serving populations that are underserved.

On the other hand, those very same applications are also gathering data, which are potentially being sold to all kinds of folks doing research, good, or pharmaceutical companies that are seeking advertising, maybe not so good. And the issues of privacy and consent that we were talking about would bear on the beneficial and adverse uses of that information. Then you turn to governments and surveillance. And that very same data, in the hands of governments, could be used, again, for good or evil. It could be used for good.

We look at the Wuhan situation now, and the Chinese government is using prescription orders and medical records to try to track people for purposes of containing the epidemic. But it also could obviously be used for political control in ways that would be adverse. I choose this one example of a \$175 cell phone and the associated information, because even embodied in that one example is so much of what you were talking about. And my head spins, because I don't even know how to answer the question with reference to that one example. And we're talking about a technology with implications that go far beyond.

And the question, very simply put, is not so much to answer that question on the cell phone and the medical data. But how could MIT people engage more effectively with the very difficult values issues that are raised? How could we work to improve the terms of the trade-offs? How should we-- or what duties do we as technologists have to address these issues? What responsibilities and duties do we have?

LUIS VIDEGARAY: Well, those are, like, six questions.

[LAUGHTER]

KENNETH OYE: It's unfair.

LUIS VIDEGARAY: So let me start with the phone, and using the phone as a delivery technology for health diagnostics and treatment. Well, first of all, it's a phenomenal opportunity. And in many places around the world, including my home country, Mexico, this is extremely appealing.

And you're already seeing successful cases of, say, detection of retinopathy, eye disease, associated with diabetes, that it is being detected through a picture taken with a cell phone. And you see that already. And those patients are then referred to the appropriate doctors. And that is creating a much better care opportunity for people. So that is there.

But I think that, going back to the US context-- and this would apply, when we talk about medicine, we should think in the same way that we think of any other technology or drugs. That these don't go unregulated. There's a reason why there are prescription drugs and there are over-the-counter drugs.

Again, the algorithmic tools for delivery of medicine should also be regulated as such. And there are some things that a patient should not be deciding just because the phone said it, and there's a very powerful algorithm that everybody says works. And then because of that, I'm going to take this treatment. You need to go to the doctor. And that kind of algorithm, that kind of promise, should be regulated just as prescription drugs.

On the question of information, I think that we are going to see more and more technology enabling privacy coexisting with this type of delivery of systems. A lot of the opportunities, particularly using mobile phones, a lot of the distributed learning process will help to protect the data of patients or individuals you have in your phone. And again, if you are the FDA, it might be a good guideline to only approve for public use those technologies that have these privacy protection tools, like distributed learning and non-centralized learning.

The difference is that, in distributed learning, the streaming data from your phone will never go to a central server or a set of servers, and can be exploited. It will remain on your phone, and all these other things will go. So there is this problem.

I think the larger question is not about technology. It's not even about companies. It's about government and democratic institutions.

And that's what I think we should be more concerned about, because this technology creates

risks for eroding democracy by excessive surveillance and manipulation, and also enabling dictatorship. And unfortunately, we're seeing that around the world. And we are seeing some of the technology that exports not just algorithms, but social control.

KENNETH OYE: Luis, thank you. We will turn to the floor for questions. And Michelle, do you want to start over-- I'll cover this side. Michelle will cover the other side. And you could hold up your hands, and we'll try to get to you.

MICHELLE: All right, could you stand up and identify yourself?

KENNETH OYE: And please identify yourself.

AUDIENCE: My name is Bill Weinstein. I'm an MIT alum. You talked about the need for the policy makers and the technologists to develop an understanding in order to develop policy well. Then you also pointed out that one would like to have a democratic consensus about how this moves. So now you've got everybody out there, who are not well versed in any of the technologies. And on top of that, they are burdened by a plethora of cognitive biases, which completely distort their ability to understand the meaning of what's going on. How do you reckon with that?

LUIS VIDEGARAY: Well, I think that's-- I firmly believe in the democratic control of technology. I don't think that we should live in a world where technology goes rogue. But I don't believe in the opposite, which would be technology controlled through referendum.

And you see an example-- and I'm going to be a little critical of California here . If you see the origin of the CCPA, this is a very complex piece of legislation that didn't go through the standard congressional process of hearings and drafting and consultation of constituencies. This is a guy-- lots of merit, who drafted a bill, gathered signatures. And suddenly, the California senate realized that it was going to pass probably with 85% of the vote according to the polls. So they immediately grabbed it and adopted it in a day.

That's not necessarily what we should aim. So this is, I think, a fine distinction. I think the key question for policymaking is how do you create democratic institutions for the appropriate policies to emerge.

And this is not-- I don't think that this should go unchecked, and that just because this is complex, the people that don't know linear algebra should be out of the conversation. I don't think that. But you cannot do it by referendum, particularly in the context of polarization and

this information that we're living through are very much fueled by this technology or enabled by this technology. I think this has to go through the workings of representative democracy that has delivered one of the most successful political models in history. That is, existing democracies, both in North America and Europe.

KENNETH OYE: Next question. And again, we ask you to give your name, your social security number--

[LAUGHTER]

--birth date, your first pet, and your high school mascot.

Hi, I'm Angelique [INAUDIBLE] from the Sloan School. So I was trying to link some of the themes and what you're talking about. It seems to me that your point about that differential impacts for big tech versus startups of a patchwork of policies is really important, and that policymaking needs to take into account the knock-on effects of locking in too early or locking in at the wrong scale.

So you have a policy that's different from California and Oregon, or it's prematurely precluding or enabling certain choices. So what are the new policy tools that factor in these increasing returns to scale-- the path dependence using business models as well, looking at how this plays out vis a vis business gaining power or changing its role in society, given different policy options? So the traditional policy methods might benefit from simulation modeling-- what if.

Here's what happens if we do this quickly. Here's what happens if we wait and see. Here's what happens if we do this temporarily. Seems like there's another layer of complexity on thinking of policy formulation.

I did a little study a few years ago with an MIT student, where we modeled malaria policy. And the naive solution, which is to invest in both prevention and treatment, was the worst policy. It was better to go all in one or the other, rather than some kind of middle ground. And it made me realize how complex policymaking is when you look it in this way.

LUIS VIDEGARAY: That's a great question, Angelique. Thank you for being here. I certainly don't think that fragmentation is a good idea. But I also mentioned in the presentation that I strongly believe in temporal experimentation and temporal regulation as a way to learn.

I mentioned the state of Vermont. I mentioned the state of New York. Last semester, we had one of the members of the legislation in New York that drafted-- they're going through a study

process.

So I think that, first of all, there's got to be awareness in legislators and policymakers. And there's going to be an understanding of what it is. And to me, the greatest concern is jumping to early and locking in establishing path dependence, as you very well described, much better than I can.

But I think that's the reason why tools [INAUDIBLE] having preemption. So I understand why people are very concerned about face recognition, particularly in the use by police and government. But I don't think it is probably the optimal solution to ban it forever and to impose - but perhaps a moratorium is much better until we understand better and that technology evolves, and is mature enough, and you go through an FDA type of process through that. Regulatory sandboxes, pros and cons.

In Mexico, I led the fintech law. And we went through the process of calibrating whatever [INAUDIBLE] sandboxes. It's not easy, but it allows you to learn.

There's a lot of learning to be here before committing too hard towards something. I think the worst case scenario is where you commit too early and have fragmented commitments. And unfortunately, that is not a scenario that can be discarded right now.

AUDIENCE:

Hi. My name is Daniela. I'm a freshman here, part of CSAIL. Thank you so much for your talk. My question is, when it comes to mitigating bias and manipulation through AI, do you think that there is systematic institutional issues that we need to solve in government, such as inequality or corruption from corporations first before we can entrust these governments with creating responsible regulation for AI?

LUIS VIDEGARAY: Thank you. Thank you for being here. Thank you for your questions.

You mentioned two topics that are quite important. But analytically, they're not exactly the same-- manipulation and bias. And both are really very bad things that we should be concerned about, and any country should have policies about those.

I think that the problem of bias in algorithms has been a little bit of a discovery. And it was not obvious in the beginning that that was going to be the case. Probably if econometricians had been consulted about the problems in dealing with datasets and bias introduced by the datasets, that problem would have been identified earlier.

It's the same problem that you deal with identification in econometrics. A lot of that looks the same. But I think that the discipline was not prepared for that. And suddenly, you find some true horror stories on that.

I think that the understanding of the problem is much better now. And there is no one single fix for this. There's no magic bullet.

I think that having more diverse teams, both in the policymaking, and in the algorithm design and training, actually helps. And this is not just soft policy because it raises awareness. But that is not enough. And I think that the quality of data and the representatives of the data is a key.

I think that relying on fixes on constrained optimization is always going to be problematic because you lose accuracy. You lose power. So the true fix is actually on the data. And the true tension, therefore-- the true trade-off is with privacy. So the trade-off between privacy and biases is real.

And I think that's something that we don't talk enough. Is that intentional? Are there companies or governments that are intentionally creating bias? Maybe, but I think that's the general rule.

Manipulation is a very different thing. I think that manipulation, almost by definition, is intentional. And a lot of people have discovered that these tools allow for a computer system to know a person even better than the person herself. And that is a problem, because when you can spread very targeted information and abuse cognitive biases, then you have an opportunity to truly manipulate markets and the political system. To me, that's the key.

And I think that in order to be constructive, and not just being gloomy here, I think technologies have an opportunity in developing tools towards empowering the individual in detecting manipulation and manipulative intent, and raising awareness when a technology is trying to exploit a particular reason for cognitive bias. And we don't see that enough. There are some examples-- some from MIT, both at CSAIL and the Media Lab. But we need a lot more of that. And at the end of the day, we need democratic control of government, which is the very essence of democracy.

KENNETH OYE: Let me try to work my way across here. And also, give the make and model was your first car.

AUDIENCE: Hi. I'm Chivani. Do you think Facebook should be broken up?

[LAUGHTER]

LUIS VIDEGARAY: I don't--

KENNETH OYE: Yes or no answer.

[LAUGHTER]

LUIS VIDEGARAY: I don't know what anybody who posts a yes or no I think is lying. It's unclear, because there is not enough historical experience or evidence about whether just breaking up a platform would be the solution. Because maybe you break up the platform and the behavior remains the same. We don't know that.

The problem here-- the things that we see as, say, evil might not be only attributable to economies of scale, which would be the argument for a breakup. I think that we need more harder thinking on that. By the way, Europe today established that they are going to impose regulation on requiring Facebook and Google and Amazon to share the data, just as banks do, or other industries, like the auto industry does. This just happened five hours ago, literally. And that's very different from breaking up the companies.

And this was the EU's commissioner, Margrethe Vestager, the antitrust person in Europe. So not necessarily a breakup is necessary. You need to look more carefully at the true economics of the platform and why it's creating these imbalance of power. And address that, not just-- this is not just Standard Oil necessarily.

AUDIENCE: Hi. My name is Adam Nagy. I'm at the Berkman Klein Center, and one of the co-authors of the principles report that you mentioned.

LUIS VIDEGARAY: Congratulations, and thank you. That's a great report.

AUDIENCE: And I'm actually joined here by another co-author. Nele is right here as well.

LUIS VIDEGARAY: Thank you to both of you.

AUDIENCE: So my question is somewhat self-servingly related. Not to that report-- I noticed when you talked about the translators, one of the groups that was missing was academics. They weren't up there. And we're in an academic institution.

The report was created in an academic setting. And I'm curious how, in your experience as a

policymaker, you interacted with academics, and what academics can do to get reports like ours, or like the MIT report that you cited, in front of policymakers, and what are the challenges to that? And maybe tying that into the AI space as well.

LUIS VIDEGARAY: Well, I didn't put academics there, because I don't think that there's a lot of that interaction. I don't see a lot of interaction between policymakers and academia the way that other translators do. Your report is an exception and, an hopefully is a very impactful report. It's recommended reading for everybody. But I don't see a lot of that.

In my experience, I remember, I-- because of my MIT background for some tough questions, when I was finance minister back in 2014, we were seriously thinking about banishing in Mexico digital currencies-- cryptocurrencies in particular. We were talking about bitcoin. And I came to MIT. And as a result of that, we were convinced that that was not the appropriate course of action.

But that was very, very ad hoc, because I had friends here, and I came from a place like this. But I think that it's not the general norm. I think that I started saying that academia and policymakers are far, far away. You guys are quite an exception. I think that's very real.

AUDIENCE: Hello. My name is Alonzo. I'm a senior studying mathematics.

You talked about trust. You talked about trust in tech companies. You talked about trust between countries. I think it's also worthwhile to talk about trust from citizens towards their governments, and then trust in the academic institutions that we are charging with studying these questions of democracy and AI and ethics in AI.

And you of sit in the middle of those two as finance minister and foreign minister of Mexico, as Enrique Pena Nieto's campaign coordinator. And then you're all in the state of Mexico, and then also now here at MIT. And so I just wanted to address what I feel is a little bit of an elephant in the room, because the political process in Mexico has always been marred by some serious issues.

And Enrique Pena Nieto's presidency and Enrique Pena Nieto's campaign is no exception, whether we're talking about the scandals with [INAUDIBLE], whether we're are talking about [INAUDIBLE] and Pemex and Emilio Lozoya, who was just last week arrested, and your particular role in those scandals. We're talking about the embezzlement of money from the finance ministry towards the [INAUDIBLE] gubernatorial campaigns, whether we're talking

about Monex and the prepaid cards and Soriana. So I just thought that it would be worthwhile to address that in the context of your discourse on trust and democracy, because I just thought it was a big elephant in the room-- especially as a Mexican national. Thank you.

LUIS VIDEGARAY: Well I acknowledge and I respect your opinions on your concerns. I said that this is not a talk about Mexico. I'm very happy to talk to you any time you want. My door is always open, and we can discuss. I can tell you that I stand by my track record of my actions.

Obviously, as a policymaker, I did good things. I made mistakes, and I learned a lot. But I stand by my actions, and I'm very happy to talk to you anytime you want to. My door is open. And we can talk about Mexico or any other thing that you want to.

KENNETH OYE: OK, I'm going to recommend that we harvest a few short questions we can turn to you for an answer, because we're running close to the end of the time.

AUDIENCE: Thank you. So I just have a quick question related to our report from the Berkman Klein Center. So what we did was identifying AI principles that already exist.

But what we also found was that there were very little principles. For example, we couldn't find one single document from the African context. And so I'm wondering, from your perspective, from Mexico, do you think that existing principles are going to be relevant in this these countries that haven't developed in principle so far, or will they be different? And how do you see that development in Mexico and Latin America more generally?

And then how do you see the developments at the UN level that there are thoughts about developing global principles of AI? Do you think it's going to be similar principles that are already discussed? Or is it different if we include all these regions that haven't yet developed anything on this issue?

LUIS VIDEGARAY: Well, I think it's a big concern, the fact that some regions have been mostly unaware, and therefore disengaged from this conversation. I think that principals, most of them, are general principles that apply everywhere. But there are a lot more regional concerns or ad hoc concerns that should be addressed.

I'm not sure how impactful these documents, particularly from developing countries, are going to be in actual policies. I think that's a big question. I haven't seen much of that, say, in Latin America. Maybe the exception is Chile, where we do see a little bit of a build-up of a policy. But other than that, I think these documents are scattered and not very impactful.

And I also noted that in your report, the lack of an African document. I mean, you're going to find one African document. And that's a real concern, particularly as African countries and African governments are embracing this technology. Some of it comes from China, and are using it in different ways, good and bad. So I think that's a concern.

KENNETH OYE: OK, why don't we gather one question, two questions, three questions. And then you can answer all or as little as you wish, Luis. Over there.

AUDIENCE: OK, I guess that's me. I'm Tim Rideout I'm a Cambridge resident. My mom is an MIT PhD. So I learned how to think from this world. And that actually relates to my question.

KENNETH OYE: It's genetic.

AUDIENCE: I know, right? I'm not good at the math, but I think in terms of a mathematical way of thinking about political science. But actually, that's my question. When it comes to bias, and then when AI uses the real time, just the way it changes language, even inadvertently, and measuring political philosophy.

So just quick backstory. My mom was-- in the '70s and '80s, they were talking about, at the old mainframe, she would carry her punch cards over to the mainframes here at MIT and run her regressions and whatnot. And the way they did it was they tried to measure political philosophy on degrees from liberal to conservative. And that was the framework they went with. And then that got compiled over years.

And I've read the books. I read Anthony Downs' theory of democracy. Drew a little bell curve.

And I was carrying these things around in my brain for years in Washington, and it took me a long time to realize that it's empirically wrong. And I've been conditioned to think this way, because that's what I've read. And at least in the American polity, we've got progressives, liberals, conservatives, libertarians. And these are different humans with different political philosophies.

I guess the point is, first of all, please, AI should not be used to measure a political philosophy. But also, how do we-- as this is used in real time and to condition people to think, and you think of psychographic profiling that-- the obvious case is what the Russians did to our country. And it's an open question. I don't know the answer. But how do we prevent people to enable free thought and free expression and intellectual freedom, so we're not conditioned to think in

very narrow worldviews.

KENNETH OYE: OK, and we'll get a couple other quick questions in here.

AUDIENCE: Hi my name is John. I'm a grad student here. My question is about the geopolitics of AI, which you address a little bit. Here in the United States, we've seen that Google has refused to work with the DOD on their projects due to ethical concerns. But in the meantime, in the other countries where there's a much closer military and research connection, do you see that as a risk to global security order, and how should the United States address that?

KENNETH OYE: And that'll be the last question of the day-- an easy one.

LUIS VIDEGARAY: And easy one, yeah.

KENNETH OYE: And afterwards, if people are interested, please come up front.

LUIS VIDEGARAY: On the political geometry, just a couple of weeks ago, Ken and I were working with this group of senior military officials, US officials. And there was a talk on Russia by a fantastic scholar from Harvard, Alex-- I don't remember his last name.

You can--

KENNETH OYE: It's all Chatham House rule.

LUIS VIDEGARAY: Yeah. That's right.

KENNETH OYE: An anonymous senior scholar from Harvard.

LUIS VIDEGARAY: An anonymous senior scholar. But the point is, apparently, they didn't target a particular ideology or a particular label-- liberals, progressives. They just targeted everybody and tried to infuse extremes.

I think that the richness of democracy is in diversity of ideas. I think the problem is polarization and going to extremes, and the lack of ability to communicate with others. And that's a problem that these platforms can create.

But it doesn't have to be this way. It's the way we've been using technology. So I don't think this is a liberal problem or a conservative problem. This is fueling polarization to maximize engagement and profits. That's a different problem, or for geopolitical purposes, as with the

Russian case.

As with the military, just one note. That case about Google has been singled out, but was one episode. And Google continues to have a lot of contracts with the DOD and other military-related agencies. So that obviously created a lot of news, but it didn't change a lot the relationship.

I think it did help to put some restraint in the kind of things that we're done. It raised the awareness of the problem. But I think it would be completely inaccurate to describe US industry, including the tech platforms, as disengaged with the US military. That is not happening.

KENNETH OYE: OK, and all that remains is for us to express our thanks to Luis Videgaray for his talk.

LUIS VIDEGARAY: Thank you all, and thank you, Ken. Thank you.