Center for Strategic and International Studies

TRANSCRIPT

Event

The AI-Surveillance Symbiosis in China: A Big Data China Event

DATE

Wednesday, July 27, 2022 at 12:00 p.m. ET

FEATURING

Scott Rozelle
Co-director, Stanford Center on China's Economy and Institutions

Noam Yuchtman
Professor of Managerial Economics and Strategy, London School of Economics

David Yang
Assistant Professor, Department of Economics, Harvard University

Emily Weinstein
Center for Strategic and Emerging Technologies (CSET), Georgetown University

Paul Mozur
Asia Tech Correspondent, The New York Times

CSIS EXPERTS

Scott Kennedy
Senior Adviser and Trustee Chair in Chinese Business and Economics, CSIS

Ilaria Mazzocco
Trustee Chair Fellow, CSIS

Paul Triolo
Senior Vice President for China and Technology Policy Lead, Albright Stonebridge Group; Trustee Chair Non-Resident Senior Associate, CSIS

Transcript By
Superior Transcriptions LLC
www.superiortranscriptions.com

CSIS | CENTER FOR STRATEGIC & INTERNATIONAL STUDIES

Scott Kennedy: Welcome to this CSIS event on AI – "The AI-Surveillance Symbiosis in China."

Artificial intelligence makes people think about a lot of things. In some ways, it's a huge opportunity. It's unbelievably impressive and, for mortals like me, oftentimes very difficult to understand. It's also connected to some of the most entrepreneurial parts of China. But there's also lots of concerns about artificial intelligence as it intersects with China's political system. We are worried about economic competitiveness in the West, national security, human rights. Everyone has their favorite sci-fi scenario that they're worried about in the – in the autonomous vehicle that they're driving and who's going to take it over while they're going across a bridge. Anyway, we need to make sense of this.

There's a lot of writing and thinking about this, but not enough good writing and thinking. But today we're going to share with you some fantastic research and also some commentators who have done some phenomenal work, and today we're going to try and really educate everybody on both where China and AI intersect and what it should mean for the Washington policy community.

David Yang and Noam Yuchtman have done some fantastic pathbreaking work on AI in China, which we are delighted to highlight in our new feature for Big Data China, the collaboration that we have with Stanford University's Center for China's Economy and Institutions. And this feature is crafted by my colleague Trustee Chair Fellow Ilaria Mazzocco.

We're also delighted today to unveil Big Data China's full microsite that has all of the features that we've produced so far and will be updated on a regular basis with more features, featuring fantastic work from scholars around the world doing fabulous research on China – quantitative research that Washington needs to understand.

Today's format is really straightforward, even though we have a lot of folks who are going to be speaking today. Our partner in crime at Stanford, Scott Rozelle, the co-director of SCCEI, is going to introduce our featured scholars. We're going to have David and Noam then present the core of their work. Then we're going to turn to my colleague Fellow Ilaria Mazzocco to say a little bit about what it was like to produce the CSIS feature and then some of the policy implications. And then we're going to turn to a fantastic group of panelists to help us really dig in deeper: Emily Weinstein of Georgetown's Center for Strategic and Emerging Technologies; Paul Triolo of the Albright Stonebridge Group and a non-resident here in the Trustee Chair; and Paul Mozur, award-winning writer/commentator/journalist for The New York Times, who has been reporting on this issue for several years.

And then we're going to take questions from the audience. You saw when you registered for the event or logged on today, if you went to the webpage, the button where you can – that you can click and submit your questions, and those will come to me.

So first thing first. I'm going to turn things over to Scott Rozelle. Scott, I want to thank you and your team at Stanford for being great partners for Big Data China as well as for this feature that we're discussing today. Again, thanks for all the work that you and everyone have done.

Scott Rozelle: Thanks, Scott. The two Scotts lead off and introduce this great panel. And thank you, Ilaria, Maya, and your team for putting this together. It's part of this effort that we're really making, thinking U.S.-China relations. Understanding China is one of the most important issues facing us today, and our goal is to take the best research in the world – that's often published in academic journals that are hard to do – and then turn it into something that we here in Washington, in the business community in the U.S., and the general public can understand. It's data-based and really lets you have new glimpses into what's happening inside China. So thank you, Scott and team.

My job's very easy today. I get to introduce two of my close colleagues, close friends, and two of the really best economists that are working on China and other parts of the world today. I'm going to do it in alphabetical order. They're both wise. (Laughs.)

David Yang. David – I've known David since he was a young grad student because he's a Stanford Ph.D. from the Economics Department there. He's now an associate professor at Harvard University and has done probably the – two or three or the most pathbreaking papers in many different directions on China. And you'll hear today, you know, what he's done. He collects data, he scrapes data, he puts together datasets like nobody else. And so, David, we're really happy to have you here for this feature.

And the other is Noam Yuchtman. I knew David when he was a first-year grad student. I knew Noam when he was in Williams College as a young undergraduate and he was an intern for us back in the early 2000s. And we took him to China and he's been glimpsing it from all these different directions ever since. He started his career after finishing his Ph.D. at Harvard at Berkeley, and now he's a professor at the London School of Economics.

So, David and Noam, thank you for being here. Thank you for doing this terrific work. The floor is yours.

Noam Yuchtman: (Off mic.)

David Yang: Noam, I think you're muted.

Dr. Rozelle: Yeah, you're muted, Noam. (Laughs.)

Dr. Yuchtman: (Comes on mic.) There we go. Yeah, I got that. Thank you very much.

Thanks, Scott and Scott, was what I was saying, very much for the introduction. Thanks to the CSIS. Thanks to Stanford's Center on China's Economy and Institutions for the chance to discuss our paper.

So David and I will be talking today about our research titled "AI-tocracy." So the background for our work, in a sense, and the academic and policy space that we wanted to step into involves a sort of political economy equilibrium that we found, you know, very reassuring but in some ways questionable.

So, since the end of the Cold War, I think many of us in the West believed in this sort of reassuring equilibrium in which democracy sustains innovation and causes economic growth and development, while autocracy is inimical to them. The existing academic literature emphasized the risk to entrepreneurs of expropriation by unconstrained autocrats and the tendency for autocrats to repress any potential challengers. Democracy, in contrast, allowed for economic disruption, allowed for creative destruction, which are the foundation for economic growth. And to the extent that citizens in less-democratic regimes got richer, they would demand more democratic rights. Put simply, democracy was seen as aligned with innovation and economic growth, while autocracy was not.

By that logic, the economies of the Western democracies and our allies – Japan, South Korea – would remain the most technologically advanced in the world, and that ensured the political and geopolitical dominance of this group of countries. To the extent that autocratic regimes or the citizens in autocratic regimes aspired to wealth and technological progress, they would need to adopt democratic political values and institutions, thus reinforcing the strong position of democratic states.

However, in recent years the emergence of artificial intelligence, or AI, seen as the technological basis for a potential fourth industrial revolution, seems poised to challenge this reassuring consensus. A key reason for this is that AI is a data-intensive technology. Since the deep-learning revolution, access to large amounts of data has played a critical role in driving AI innovation from translation to chess mastery to facial-recognition technology.

So what does a data-intensive technology imply for the relationship between autocracy and innovation? So we believe that there exists the possibility of a mutually-reinforcing relationship between autocracy and AI innovation in particular. Why is that? Well, autocrats engage in intensive monitoring of

their populations. They always have. And especially in contemporary autocratic societies, they collect massive amounts of data on their citizens.

Autocrats demand AI technology as a technology of prediction. What does that mean? Well, prediction isn't just useful to make guesses; prediction is first useful for identifying individuals, identifying actors. It's next, of course, useful, once you identify an individual or actor, to make a prediction about their behavior. But yet another stage of prediction is trying to make a guess about what an individual's response will be to changed incentives. And by making predictions about how individuals can respond to changed incentives, autocrats can use AI technology for behavioral manipulation.

Because AI is this sort of powerful technology of behavioral control, it's also a technology, potentially, of political and social control. Because autocrats will find that so useful, they can credibly commit not to predate on entrepreneurs. Nor would entrepreneurs wish to undermine a state that gives them access to valuable data for innovation. So firms will work hand in hand with autocratic states. The firms that provide prediction services, which you can think of as political control AI services, to autocrats will benefit from the receipt of government contracts.

Importantly, in addition to the typical financial benefits of these contracts, the data that firms are able to receive and access when they work with the state can play a key role in stimulating further innovation. Consider firms providing some sort of facial-recognition AI services to the government. They'll analyze surveillance video feeds, and most directly they'll use those video surveillance feeds to provide services to the police force, let's say to help the police suppress political unrest. But that same data isn't just used to make predictions for political control; it's also useful as an input into improving facial-recognition algorithms. The firm that provides AI services to the government can use that data to develop new facial-recognition technology with government data as an input into innovation. And this innovation isn't just narrow innovation for government use; it can be much broader. Trained algorithms for facial recognition can be used to identify a protester, but they can also be used to identify customers in a retail context. So AI technology may help entrench autocrats, and autocrats' politically-motivated procurement of AI may stimulate AI innovation even beyond government applications and move out the broader technological frontier.

We're going to test for this sort of alignment of autocratic political institutions and AI innovation in the context of facial-recognition technology in China. We're going to study the world's leading autocracy and among the most active areas of AI innovation. We're first going to consider whether AI technology enhances autocrats' political control, studying the impact of political unrest on public security procurement of AI. We're also going to study the impact of AI procurement on subsequent unrest. We, lastly,

consider whether politically-motivated AI procurement stimulates follow-on innovation.

Dr. Yang: Great. So in order to carry out these tests, the first part of this research project was to try to sort of put together a number of datasets. Some of them, we built it from scratch.

The first dataset that links about 8,000 facial-recognition AI firms that have ever sort of come into existence in China over the last decade or so – link these 8,000 facial-recognition AI firms to about sort of 3 million government procurement contracts that's been published by the Chinese Ministry of Finance. And in particular, I want to focus on about 10,000 AI contracts that have been issued by local public security agencies of the government, which you would expect that might be responding to local unrest and local political turbulences and potential demand for AI technology, OK.

We're also going to combine this together with datasets coming from the GDELT database, which is a global database that track events around the world. We're going to focus on about 10,000 political unrest events that took place in China – across China from 2014 until today.

And finally, to look at sort of, you know, the innovation activity of the AI firms when we look at all the software that the AI firms has been developed and upgraded over the – over the course of this time period. So we're going to look at sort of all the software registration that these firms have to submit to China's Ministry of Industry and Information Technology before this software has been released to the – to the customer. Very importantly, we're going to categorize this software that the firms are developing into the software that are intended for government uses and also software that's intended for commercial applications that may not be directly of interest to the government per se, OK. And finally, sort of to look at whether these softwares are at the frontier of the innovation, we're going to link this software that the firms are developing to the exporting activity of these firms, so to look at whether this software generates global demand which will be a common indicator that economists use to think about the frontier of innovation.

OK. So I'm going to show you a number – actually, let me also mention that many of the software that we're looking at that these firms are developing that – which will be sort of included in our analysis are literally at the – at the global frontier. If you were to sort of rate the facial-recognition algorithms based on their accuracy or sort of the speed of the – of the algorithm – the U.S. government is doing so every year – over the last sort of five years or so, the Chinese algorithm has consistently ranking towards the – towards the top of the list. In this particular picture, you see that sort of the five out of the – six out of the top 10 algorithm and all five of the top most-accurate and

most-speedy algorithm in facial recognition were developed by the Chinese facial-recognition companies, OK.

All right. So we have three tests to think about sort of this AI-tocracy symbiosis.

The first test is – one asks is: Do local sort of unrest episodes in one particular quarter leads to greater sort of procurement of AI technology by the local police department in the subsequent quarter, OK? What you see here in this figure is that leading up to the unrest, two quarter or one quarter of units of unrest, there's little evidence of anticipation – that the local governments are buying more AI, local police department are buying more AI for sort of anticipating the upcoming unrest. But the quarter after the unrest took place in those prefectures, there is a very substantial increase in the amount of new procurement contract that the police department are issuing out to purchase AI technology. This is controlling for local – sort of, you know, locality sort of characteristics and very sort of macro trend and so on. So this is sort of, you know – I think is isolating sort of the effect of unrest on the subsequent AI procurement.

We also find sort of as the local police department are buying more AI sort of technology and AI services, it's come along with a whole package of broad technological upgrading of the local police department. Also, in the same time – in the same time period in responding to local unrest, buying more high-resolution surveillance cameras, which will be crucial to provide the useful data for the – for the government. And also, we see that sort of local government – local police department are start to hire less police forces and there's a larger proportion of desk job police forces, suggesting that even the labor part of the – of the political control apparatus has been upgraded, potentially to be – to be consistent with the AI technology. OK.

Is the AI technology the local governments – local police department is buying actually useful? The next question we're looking at is sort of whether the past procurement of AI technology actually induce a suppression of future unrest, which presumably is what the political motivation for those procurement in the first place. OK. What you – what you see in this table is that while sort of good weather conditions holding fixed sort of – of the local grievances and other sort of, you know, socioeconomic conditions, good weather conditions is inducive of large public gatherings and political unrest, as the positive coefficient in fine weather is indicative.

You see sort of a strong negative coefficient on good weather conditions interacting with the past sort of procurement of AI by the – by the local police department. That's suggesting that sort of, you know, the factors that would induce a protest, such as weather, become substantially tempered

when the local public security arms of the government has been accumulating a large capacity of AI technology.

And this doesn't seem to be sort of, you know, a result of a general sort of government upgrading the technology. As you see in columns five and eight in this table, if a non-public-security part of the government is purchasing more AI in the past, that does not have an effect on tempering the subsequent unrest. We have to be sort of the local public security arms of the government who is – who is now taking advantage of the AI technology that we see sort of a drop in unrest in the future period, OK.

So that's on the political side, where this is technology that potentially has been demanded by the local police department and has been fairly high sort of – sort of effect on suppressing unrest. What about sort of from a firm's perspective? When the firm providing those services to the government, are they benefiting from these politically-motivated public contracts? So here we're looking at sort of within those AI firms who have been awarded sort of their very first contracts with the government, within those firms, do they start to develop and issue new software quarter by quarter leading up and right after the very first contract is let. So as you see in these two figures, the left-hand panel is showing sort of the new development of software contracts – sorry, of the government software. You see that sort of the – sort of these AI firms are no differentially developing more software before the first contract arrives, but immediately after the very first contract from the second quarter and on the firms start to develop more government software. And also, potentially more surprisingly, these firms also at the same time start to develop more commercial software, and that's very much then suggesting a very quick sort of spillover of the useful data and other sort of inputs that the firms are able to get access to from working with the government. We see also these very firms, after their very first contract with the government, start to become exporter of the technology for the very first time, suggesting these firms are now pushed very quickly to the frontier and beginning to potentially sort of, you know, have a global demand for that technology.

OK. So in recap, what we show in this research is this – is this sort of feedback loop, and then you can call this sort of a sustained – a new program of sustained innovation that's being developed under entrenched autocrats. Academically, it will hopefully challenge some of our conventional wisdom of whether sort of autocracy and frontier innovation may or may not coexist. It will challenge some of our sort of, you know, expectation about the political trajectory of China moving forward as a – as a innovative economic powerhouse, especially maybe in a sector such as AI, but also have a very entrenched and stable political system that go along with his technology. It will also imply that from a global perspective we're going to potentially be

dealing with a geopolitical challenger that's going be vastly autocratic and economically advanced.

I want to spend the last minute on – so this is – you know, while the AI technology has been developing in China potentially partially out of political motivation, this is going to be technology that's going to have pretty far-reaching global and international repercussions. Our ongoing works try to – try to trace all the – sort of, you know, the foreign buyers of China's technology, and we see that sort of – sort of China, indeed, seems to have a very large comparative advantage in this technology. Half of the global trade deals in AI facial-recognition technology is dominated by the exporter of China, which is very different from any other frontier technology that we have seen in the past. And potentially from a geopolitical perspective, more worrisomely, China's export of AI technology is heavily sort of steered towards buyers who are – buyer country who are weak democracies or strong autocracies – again, very different pattern from the traditional frontier technology where it's the rich and democratic countries tend to – tend to be the buyers of those.

And I want to close here and just say that sort of, you know, this is a – this is a pattern that's going to be a political economy sort of equilibrium within China, but it will start to then have influences over the rest of the world through the export of this technology. And yeah.

Dr. Kennedy: Super. Well, Noam and David, thank you both for sharing your research today, as well as with us, and also for letting us highlight it in this program. Really important, and in some ways going back to the very first slide that Noam presented questioning the original thesis about an end of history, you all may be talking about a – still an end of history type of thesis, but a very different type of end of history. And so very powerful, yes, shaking up the conversation quite a bit.

Let me turn now to my colleague here at CSIS Ilaria Mazzocco. She joined us about a year ago. Her Ph.D. is from Johns Hopkins SAIS, where she did some pathbreaking research on clean tech and how China's political economy shapes the evolution of technologies – not an entirely different story from what we're talking about today in terms of the intersection of politics and technology trajectory. She was at the Paulson Institute for a few years and we were lucky enough to have her join us last year. She leads our work on each of the features with regard to Big Data China and has done a fantastic job. She also co-authored a terrific study we issued earlier this year in May on Chinese industrial policy spending, "Red Ink." And she's also done an amazing job putting together the microsite that we launched today that has everything related to Big Data China on it.

So, Ilaria, I want to thank you for the fantastic work that you've done for us in so many different ways, and invite you to share some of your comments and thoughts about what it was like to put together this feature on AI and then some of your thoughts about the policy side of things. Over to you.

Ilaria Mazzocco: Thank you, Scott. And thank you and my – from our team that worked so hard to get the website up, and also Scott Rozelle, Matt Boswell, and everybody from the SCCEI team for their support. And then thank you to Noam and David for, you know, providing the data so that we could actually make this happen. And then, obviously, to all the panelists today. It's very exciting.

So, I mean, I think what I'll say is that it was from the start, when we started a conversation with Noam and then David on their research, it was clear that this was very interesting, very important research. What was perhaps – I mean, it was certainly challenging – was actually the so-what question, right? It was obvious that this was a very important story to bring to Washington, but the policy implications are quite complex, right? I mean, and I think not to simply it too much, but you know, it came down in part to the old question of does it mean that we need to run faster, does it mean that we need to slow down China, and what – and then, you know, what does it mean for human rights, right? There was this big component on the surveillance state, which is obviously a little different from what we usually do, Scott, right? I think we were more familiar – I was more familiar with the sort of commercial competition component, which was something that we spent a lot of time here at the Chinese Business and Economics Program thinking about.

So I think those were the interesting and challenging parts of it. And I think, you know, at the end of the day, after spending some time on this and talking to experts – I would like to thank Paul Triolo for giving me his thoughts on this at certain points of the process – you know, the challenge really is that there's only so much you can – you know, this is a very interconnected, very globally integrated industry, AI, right, data and, you know, all the different components, right, algorithms and computing power, et cetera, and talent, obviously, right? Talent is really crucial. So it's very global and it's very hard to see how exactly, you know, where – how effective decoupling or putting strict restrictions are. And I know that Emily actually just published an excellent paper on this. But you know, so that was very challenging to, you know, think of what the terms of that would be.

And it was very clear that really if, you know, on the commercial side of things, if the U.S. wants to do better it needs to invest more in the U.S., right? So more on education, more on research, and more on attracting talent, which is really, I think, the most obvious part of – one of the most – the areas where China's the most – sorry, the U.S. is the most competitive with respect to China, right? We're really a magnet for talent.

But aside from that, then, you know, there were all these different parts to it, and the human rights one I'd say is by far the most upsetting and concerning one. But ultimately, I'd say this – and this was my conclusion – is this is not just an AI issue, right? AI is just a tool that is used within a broader surveillance-state mission within China, and so there's only so much we can expect AI policy to do in this area. And that's, I think, the – that's, I think, an important thing to keep in mind.

So I know that our – you know, and then, finally, I would say we're talking – I went on to talk about AI broadly, but really Noam and David's research is very focused on a specific type of AI, right? This is facial-recognition technology. And I know that our panelists here are experts on broader topics so I'm interested to hear – to hear their perspective, but really there's like a variety of different areas within AI that have sort of different applications and different sort of linkages between the state and the – and the business sector.

So I think that I could keep talking about this. You know, Scott, we've been working on this for a few months now. But I think there were – there are many different implications and it's really interesting research, and I – it's been a pleasure to have the opportunity to work on this. And I'm looking forward to the discussion.

Dr. Kennedy: Well, thanks so much, Ilaria. And, again, really, congratulations on the terrific feature, microsite, and the orchestra that it takes, everyone working together. It's really been – really been fantastic.

We have three fantastic experts who know a lot about AI and technology and China, who I want to bring into the conversation now. I'm going to introduce them and then ask the person a question, and then go to the next person. So I'm going to start with Emily, and then go to Paul Triolo, and then Paul Mozur.

So Emily Weinstein is a research fellow at Georgetown's Center for Strategic and Emerging Technology, focused on U.S. national competitiveness in AI and technologies – U.S.-China technology competition. She's also a nonresident fellow at the Atlantic Council's Global China Hub, and the National Bureau of Asian Research. And as Ilaria just said, she just published a paper right in the space, perfectly timed, for today's event. And, you know, Emily, I'm really glad that you're with us today.

And I wanted to ask you about the research that Noam and David have done, what you think of their work and the core conclusion. And also, how does that fit – how does their work fit with China's broader push towards technology innovation and other areas of AI that they're working on. As

Ilaria said, not all AI is perceptions AI. You have language AI and many other types of applications. And, you know, the significance of government procurement as a critical driver of how the trajectory of AI and China in your research, and maybe relative to other types of factors – entrepreneurial investment, venture capital, global markets, et cetera. So if you could help us sort of contextualize their research into how you look at China's AI space.

Emily Weinstein: Sure. No, I'm happy to. And thank you, again, to you, Scott, and Ilaria, and the rest of your team, for having me on today.

So I want to say that the research that Noam and David presented is super fascinating. And I really look forward to digging more into it. I did a kind of brief overview, and this presentation was super helpful. But I do – the fact that you guys have gone into the procurement data I think is what sets this research apart from other folks who are looking at similar topics. CSET, as many folks know, we're really focused on being a data-driven analysis shop. And we focus a lot of our research – again, starting from that small point of what is the data that we have, what data do we want to collect, and then what story can we shape? What analysis can we pull from the data that we have.

I personally have not worked as much with procurement data, but my colleagues – some of my colleagues have looked very closely at, for instance, PLA procurement data. So what is the China military procuring in artificial intelligence. And it provides you with a different story or with different potential indicators in comparison to if you're looking at what are Chinese companies investing in, or what are Chinese universities or Chinese researchers publishing in, all kind of in the scope of AI. Because I think procurement is showing against specifically what the government wants.

Whereas the other two you could say, you know, we understand that the Chinese government – or, the Chinese government or the Chinese system is much more, obviously, of a top-down system. The Universities are owned by the Ministry of Education, or the Ministry of Industry and Information Technology. So they're all under the government. And you can – we can go into the semantics about what the private industry actually is in China, if China has a private industry. I won't – I won't go into detail on that now, because that's a whole separate debate.

But those indicators are a little less – or, I would say – a little more flexible, or less direct than the procurement ones. Because, again, this is ones that the government is directly saying: OK, we have identified this as of potential interest. We have a specific application that we want to apply it to. And then we go from there. Versus, you know, hopefully, you know, if a company is in – if a Chinese venture capital company is investing in an AI company in China, obviously that means that there's something that is of interest, but it's

not necessarily directly relevant, obviously, to the Chinese government. It can be. It doesn't have to be. And, again, at the university level, obviously the government has some say – or, more than some say – but has a say of what topics are of interest.

And we've actually seen – some of my CSET colleagues have done some excellent work looking at China's AI education system and have actually done some work looking at how China over the past five-ish, 10 years has actually started developing actual, like, AI majors at different Chinese universities, and has set up AI-specific laboratories or research centers. But again, the government is not – you know, you're not having, you know, someone from, you know, the Ministry of Science and Technology going through and reviewing every single publication that has to do with AI before it's published.

So that's just to say that I think it's a really unique and important indicator. And I can just say briefly some thoughts on how this kind of fits into China's broader AI ecosystem – or, AI ecosystem and innovation ecosystem. So the Chinese have been pretty – I mean, and this is not unique, I would say, to this context. But the Chinese are usually pretty explicit about their strengths and their weaknesses and vulnerabilities. If you look in Chinese language, they will often be pretty explicit about these and pretty upfront. And CSET has done some great work, for instance, looking at Chinese chokepoints, what China views as their weaknesses in supply chains. Again, all public source, out there in Chinese, for anyone to find.

And what we've seen is particularly in artificial intelligence and emerging technologies, the part that China is missing now is usually falling somewhere in the basic research level. The part where for a while we thought of China, particularly in the West, particularly in the United States, we always thought of China as a country that wasn't necessarily capable of innovating. It kind of had this copycat culture, where it could buy things or steal things and try to, you know, take them apart and rebuild them. And that was certainly the case.

But we've moved into a point now where China is shifting out of that copycat mentality, and they're actually able to innovate in their own unique way. And I've written a little bit about this, as have others. But again, Chinese scholars – and I'll point specifically to a paper by the dean of the School of International Studies at Peking University, Wang Jisi, that came out actually in January. It's a super interesting piece about China and the U.S., and decoupling in technology, that actually conveniently disappeared off the internet after about a week, which to me is an excellent signal that he was onto something.

But he actually argues that if the U.S. and China were to decouple in things like artificial intelligence right now, China would have more to lose than the United States. And I think, again, it's hitting on that basic research level because it's the intangible, you know, transfer of knowledge, the tacit knowledge piece that we've seen China pushing to get in terms of, you know, things like Chinese talent programs, exchange programs. You can even go towards the more nefarious things, like cybersecurity things, hacking, things like that. It's, again, that knowledge that is really what China needs.

Oh, and there was a great quote from the chokepoints piece that I mentioned before that my colleague Ben Murphy wrote, where in the context of photolithography machines – and I know this is different from AI, but I think it's a parallel here – an expert actually in China argued that even ASML were to give China the blueprints to their photolithography machines, they still would not able to replicate them because of the tacit knowledge that is needed to operate those machines, and deal with upkeep, and all these different things. So in terms of innovating, again, China is getting there. But there are certain chokepoints or certain areas in innovation – particularly when it comes to basic research – that China is really lacking in. So I'll pause there for now.

Dr. Kennedy: Thank you, Emily. Really helpful understanding the overall ecosystem. Strengths and weaknesses of not only AI in China, but their overall innovation trajectory. We had hosted Wang Jisi at CSIS in February just after that piece came out. So it was a really interesting time for him to be in the U.S. And I know that that paper, even though it wasn't on its original website for a very long time, has generated a lot of discussion here and in China too.

And so let me shift now to Paul Triolo, who is the vice president at Albright Stonebridge Group for China and technology policy. He's also a nonresident senior associate with our program here at CSIS. We're super excited to have Paul as part of our team. He previously was with the Eurasia Group for several years, managing their technology policy program. And before that, spent a good number of years in the U.S. government helping our system understand how China worked – or didn't work, on occasion. So, Paul, we're really glad to have you with us.

I want to ask you about the findings from Noam and David, and how much you think that represents sort of a key dynamic. They presented a very clear sort of logic of how government procurement is driving demand in the Chinese state, and then driving how companies innovate in AI, and link that to surveillance, and generate income which creates a positive feedback loop. At the same time, there are potentially other sources of trends and factors shaping AI in China. And, you know, sort of what is your overall impressions of AI in China and the relevance of government procurement to that larger story?

Paul Triolo: Big question. Big question. And I thought Emily did a great job in outlining some of the issues. So thanks for this great panel on the – and the chance to examine this really interesting research. My sense is that, you know, in discussions around China and AI, what is often missing, it seems to me, is sort of, you know, how things really work in China in the AI world.

So a lot of times people are relying on aspirational documents, and sometimes data sets that – where the actual AI component is not always clear. AI is a big term that encompasses many, many things. So it's nice to see this data-driven project that's been presented here, because it includes multiple data sets. It includes both the contracts and the unrest, for example, as a way of sort of getting beyond, you know, just a single data set that may or may not tell a particular story.

In my world, I work also day-to-day with companies that are – that are – you know, that I wouldn't call AI companies, but they're companies that have an AI – a big, important AI piece to their business model, for example. So it's important to also note that AI, and talking about AI, it helps to get specific as quickly as you can because talking about AI in general sometimes is not that helpful. I also try to talk with people on the ground in China. Of course, very difficult to travel there during the pandemic. But I try to talk with people across China regularly on what's happening in China in the AI space, to sort of bring a little bit more ground truth. Because it's hard to follow such a complicated issue remotely. Unfortunately, we're probably going to be stuck doing that for a while.

So I think, first of all, that – you know, I think it's important to note that there is – and I think Emily mentioned too – there's no top-down AI, you know, mastermind in China that's commanding researchers and companies, that's sort of part of some centralized and nefarious AI strategy for global domination. I think that's one thing to note here. As in many other countries and economies, Chinese AI companies, or companies that are leveraging AI – I prefer companies that are developing AI as part of their broader business strategy – they're seeking to address a particular type of business requirement.

And here it's interesting, because we have a government having a requirement, obviously, and then Chinese companies stepping up to address that requirement. And that's an important thing. I think that the focus on sort of surveillance and security and facial recognition is understandable, of course, in the context of sort of digital authoritarianism and the surveillance state. But there's some issues I think we have to sort of sus out here. I think we've talked a little bit about them already.

One is, you know, the global widespread use of facial recognition technology as part of policing, for example. I mean, there's debates in many countries about how the police should use this, also in the commercial space. There was just a story yesterday about U.K. citizens being concerned about the Orwellian nature of the facial recognition and other biometric data they were having to give up to go into stores. So there's sort of a bigger, broader global issue around, you know, how facial recognition sort of fits into the broader issue of data collection. So that's one thing.

And then there's some – you know, the other piece that in terms of the policy response I think that we have to keep in mind is, you know, efforts to combat sort of digital authoritarianism by targeting Chinese surveillance technology companies, for example, tends to – tends to, you know, be sort of – look a little bit political and may carry less weight in some quarters when it's not part of a broader framework to mitigate abuses of surveillance technology globally, including those, you know, facilitated by firms from the U.S., Israel, and other Western democracies that also have companies of concern.

On the policy side then, I think, you know, punishing Chinese AI companies – if that – it depends on what's the goal there? If we're concerned about the use of facial recognition, for example, in policing, as has been brought out in this paper. So, for example, all these Chinese companies – I think all of them that were on that graphic that was shown – are all on the Entity List, I believe. All – you know, from Megvii, to iFLYTEK, to SenseTime, the whole nine yards. So but the question is sort of what's the goal here. Is the goal to change behavior or is the goal to punish?

Megvii, for example, a company that I've had a fair amount of contact with, you know, they were about to do an IPO in Hong Kong, and then when this came out they made a major effort to change their business practices – for example, cutting down on the work they were doing in Xinjiang, trying to appoint an AI ethics advisory board, et cetera. And they also went out to try to appeal to the Commerce Department about getting off the Entity List. But in the Trump era, that was – that was not going to fly politically. So that's another sort of policy issue as you're looking at this is, OK, if we want to – if the goal is punishing these companies, for example, to change behavior, how does that work?

And then the other issue I think that we've touched a little bit upon is the effectiveness of the technology. How is it used? What advantages does it bring? China already had a pretty good public security apparatus and capabilities before AI, right? And so what is the sort of – what is the real delta here that the presence of security cameras and AI behind them brings? Is the mere presence of security cameras, for example, a deterrence in terms of suppressing demonstrations or other kinds of public activity?

And then, again, the other major issue, how integrated are the databases behind AI and facial recognition systems? I think Paul Mozur, for example, had a really good – one of his famous stories with the glasses, right? The glasses that were supposed to be so effective. And it turned out, there was not a whole lot behind them. But that public security folks, when the police would wave those glasses, as if, you know, those were going to tell the poor person who was under investigation, you know, everything about them. So there's a big issue in China too about sort of how this – all this information that's being gathered by facial recognition and other methods is being integrated. The Shanghai database leak, I think, recently was pretty interesting.

So I also agree in the paper that the data is really important, of course. And it's not surprising, of course, that data is really the big issue here. You know, and in China, of course, data is really an important sort of battleground now. You can argue that originally it was sort of algorithms and technology, and now it's really all about data and which companies have access to the most data. So for example, beyond facial recognition, which I think – you know, was naturally sort of one of the first things funded because of this government need. And a lot of the companies – those eight AI companies I mentioned – were all driven by the – in large part by the revenue that was generated from the – some of these public contracts here.

But there's a whole – but that era is probably over. You're probably not going to see a whole lot more companies going in for facial recognition. I'm curious at how you got the 8,000 companies in that survey, but that seems like a lot of companies. And the big eight I think are really the really dominant companies in terms of both the technology and the capabilities. But looking forward, I think the real interesting part of that data battle in China is going to be in other areas. Like, large hospitals, for example, are building their own AI teams and using their hospital's own data to train their models. Manufacturers are also using AI. Computer vision here is also, you know, for quality control and other things, is a big growth area.

And other areas, I think, are medical imaging – like for cancer detection and other things. There's some really good Chinese companies in this space, like Infervision, who are trained in the U.S. and went back to China. So there's a lot of – there's a lot of other areas that facial recognition has sort of been focused on. But facial recognition and object recognition more broadly are used in many, many other applications on the commercial side.

And then finally, I'll just say that I think that, you know, the – as I said, it's really good to have this – the data point that I found most interesting, which was: This wasn't a centralized state effort to push facial recognition technology down to policing stations, but something that seems to be, you know, grassroots up. Like, when there was a problem, everybody said, hey,

let's get – what are we going to do about this? And so they said, hey, let's get some facial recognition technology and more cameras. And so that's sort of how – what looks like happened. And it will be interesting to see if that sort of holds up.

And then – but the final thing I'll say is I'm a little concerned about extrapolating form that then into some of these broader, bigger geopolitical issues – like exporting the technology, which I think, as one of the speakers mentioned, you know, had a lot of other dimensions to them. But I think it's important. It's an important piece of the problem. But I think that that's going to require more research and more data to sort of sus out, you know, to meet the – and what the policy implications of that are.

Dr. Kennedy: Terrific. Well, thanks, Paul. And appreciate those comments bringing in some of the corporate perspective, and some of the individual companies that are involved in these different elements of AI, and the plusses and minuses, or challenges, of different kinds of policy responses. We'll get into some more of the policy responses in a little bit, but I want to turn now to Paul Mozur to talk a little bit more about these findings, and how they align with some of the work that he's done.

Paul is a Pulitzer Prize-winning journalist, focused on technology and geopolitics in Asia. Among other topics, he's covered extensively the buildout of China's surveillance capabilities, chronicling how police there have used artificial intelligence to profile minorities domestically, and sold cameras and monitoring software abroad. In a recent article, he documented how China has used data collection to push for new, more invasive forms of predictive policing, some of which purport to see the future. I've known Paul for several years, and I think one of the recent times we were able to get together was in Wuzhen at the World Internet Conference a few years ago, when those were still in person. In some ways, another window into what China is doing in technology and norm setting. A little bit different from here. But nevertheless, it's – thank you for joining us.

Paul, I want to ask you about – you know, you've done some of the most path-breaking journalism on this area, not just with regard to Xinjiang but more generally speaking, including some very recent reports that The New York Times has published. How consistent is this work with what you found? And some of the couple of the things that have come up in the conversation so far have to do with this debate over central plan, versus local demand. There's always this question about, you know, is China just a fragmented, you know, 34 provinces or is it more integrated? And then the international side of things, to what extent is what we see this Chinese dynamic, you know, being exported – not just products, but sort of this approach? So welcome your thoughts on any and all of these ideas.

Scott Mozur: A small amount of things to cover. Thanks, Scott, again, for having me. And, David and Noam, congratulations on some very, very interesting research. It's really fascinating to see, and I'm excited to talk to you a bit more about it after.

So I think just to start off it might be helpful to kind of get a sense of how on the ground this actually works. And so at a local government level, you know, we all know that, you know, China's covered in cameras. But the way this is working now is that you can't watch everybody all the time with a camera. And so what police do at a local level is they make lists of people that they consider their biggest problem makers. Now, oftentimes this includes criminals. This would be, you know, people with a criminal – you know, a background of theft, or something like that. But also oftentimes it includes dissidents and other sort of political ne'er-do-wells. You know, your classic petitioner, that kind of thing.

And so what you do is you put that list of people, and you get a good picture of their face, and you put it into software that's going to look out for those specific people. So all of a sudden instead of watching 10,000 people you're watching, say, 300 people. Much easier to do that. And so what you do is you have a backend software basically take a note every time a camera detects one of these 300 people, or thinks it detects one of these 300 people. And it uses not just the face but the clothes they're wearing and other things, to basically draw a kind of map of where they've been throughout any given day.

What you can then do is go back and search those maps if you run into somebody that, you know, say two weeks later they're protesting and you want to go back, you can basically go back to that software and consult it. These lists are called blacklists. They're mostly – most Chinese people are not aware they exist. And if they are, they don't really know who's on them. These blacklists often include, you know, people with mental illness and, again, petitioners, you know, political dissidents, people who protest, and so on, along with thieves. And within the police, there's a very strong belief that this has cracked down on theft and other crimes to an extreme degree, and it's really helped them cut some of these crimes simply because it is so easy to get people based on this.

Now, if you have your troublemakers and you're saying: These are the troublemakers that you're going to watch, how do you kind of stop a protest? Because, you know, some of the data that we're seeing here says – shows that actually protests go down and incident levels seem to fall off after this software is installed? What you do, and what the police are doing in practice, is creating alarms. And so what they do is they come up with a set of preconditions that they want to kind of be aware of that they wouldn't otherwise be aware of. So, you know, let's take a petitioner. If a petitioner

says, you know, goes to the train station, if their face appears near the train station, well, they may be going to Beijing to petition. And so you send an alarm and the police go try to catch them.

The same thing goes for, you know, some of these blacklists are very – you know, are effectively racially profiling. So if several Uighurs, say, check into the same hotel, you know, there is an alarm that will send police. And we see this for all manner of different things. And there's actually competitions now across China on the local level where police actually compete to come up with clever new ways to detect whether somebody might commit a crime or might protest, something like that.

And so, you know, another version of this is, say, you know, if somebody with a record of mental illness goes near a school, you know, there will be an alarm sent. And what we see is this is sort of expanding beyond computer vision to other elements. So things like electricity usage. So if you see an apartment that has excess electricity usage, you can send the police to check on that. If a person with a known sort of history of drug dealing makes five calls out of town in a day, send an alarm. And all these are very public. They're out there kind of talking about this.

So this is how this would work. And we actually – you know, it's difficult to test, and there's not really good data on, you know, how in practice this works, because some police are more effective at using it than others. Sometimes the software itself doesn't work. We found out for our recent report an 80-year-old man who had 60 years of petitioning, which was just a brilliant find. And this guy was extremely clever. His name is Mr. Jiang. And he talked about, you know, 20-30 years ago it was fairly easy to make it to Beijing. You know, you could just get on – you basically would just get out of town, get onto the local roads, and then you can hire cars, take buses, you were on your way.

Now in a recent trip to get to Beijing, he turned off his phone, left at night, took a car paid by cash, got to the local capital, bought a train ticket to the wrong destination, which was Beijing, got off before, because he believed that his buying the tickets would alert the police to pick him up when he got to Beijing. Got off before, then took a bus, got off on a bus, took another car, paid money for it, got out before a checkpoint where they check IDs for buses, took another private car, and then got into line with other petitioners at dawn. So this is the level of kind of evasion that it takes now to get to Beijing.

So to speak to Paul's point about how effective this is, if used fairly well it's extremely effective. And so, you know, it's actually kind of remarkable. And so then, you know, to kind of speak to this idea of the future alarms and the prediction elements of it, what he told us was that actually since he went

back from that trip, because he turned off his phone, you know, when he first left, now whenever he turns off his phone, police show up at his house. So there's a monitoring down to that level where, you know, this is a new alarm that they've created for him. And, you know, some of the projects that we've seen in the procurement documents do kind of speak to this.

So there's a specific petitioner system that's being built by, again, one of the companies, Hikvision, that's on that list. And what this is doing is actually taking every single new petitioner that they find that makes it to Beijing and doing an assessment of what they did to get there, what their mental state was, what happened to them recently, and turning it all in a coded way into a database. With the idea that over time over the next few years, as this data is collected, you can actually, you know, pump that into a system, you know, throw machine learning at it, and churn out a new algorithm that will help predict when somebody will petition, what the preconditions are for them to go to petition.

And it actually has a separate subset for people like Mr. Jiang, where people who are good at evasive maneuvering, as it said. You know, and for those people there's specific things where they're looking at what do those people do before they go petition, and how can you kind of intercept them ahead of time? So this is – this is how it works on the ground. It varies immensely across China. Some police are extremely effective at this. Some are terrible at it.

Police complain about these alarms, because a lot of their sort of – the flexibility they used to have in policing has been now debased to the point of just responding to alarms. And a lot of police assessments that we've seen now are based on how well and how quickly they respond to those alarms. So you see that kind of thing kind of creeping in, where the police actually have less flexibility now because the algorithms that are watching people are also, in a way, watching them, because you can see from the center how well they respond to these things.

To speak to Paul's other point about centralization, there are absolutely centralized public security bureau procurement documents that lay out a lot of this stuff. You know, they lay out in great specifics the desire to have cameras that can do everything from, you know, notice a rust spot on a car and track that car across places to, you know, notice the first time a car enters a city, to noticing if somebody goes out late at night multiple nights in a row, and sending alarms. So this is absolutely a national effort. But at the same time, you know, as things go in China, just because the central government says something doesn't make it so.

And so I do think probably one of the things that we're seeing in the data here is that – is that, you know, some police departments do this very well.

Shanghai, Hangzhou are certainly going to be on the cutting edge of this stuff. But other places, you know, say we go to Lanzhou, or we go out to Hunan somewhere, they may not be that effective at it. And so it takes something like a big protest to give them the kick in the pants that they need to use this.

And a final point on this, they do use it effectively for crime as well. And I think one of the interesting things is this stuff kind of tends to lay dormant. And Chinese people can be quite happy with it because it does seem to be effective at stopping crime. The only moment you actually see the political impact of it is when a whole group of people get angry about some injustice, take to the streets, and all of a sudden, it's effective at going after very specific people, and stopping that protest. So there is this kind of thing where it's a very pleasant thing, until it's not.

And you don't really know what's going on in the background until sometimes it's too late. And so, you know, in Inner Mongolia, when we saw the protests – I guess it was almost two years ago now – a lot of people were not aware of the level of the surveillance until it kind of came about and was used to track them down. And in terms of – we have seen it exported a lot. I've seen less facial recognition, more simply cameras. And I think thus far the biggest impact China has had sort of globally is just pushing down surveillance camera prices to make them extremely affordable.

But we certainly know that these companies are very actively courting a lot of countries and trying to sell them this software. A lot of countries are quite worried about access to the biometric data, and so sometimes they don't – they don't use it. But we do see now grids of cameras and some software to process it in a lot of countries. And I like to think of that as almost, like, an operating system for, you know, techno-authoritarianism. Once you install cameras and you have a software that's sort of processing that, it's easy to keep upgrading it and eventually come up with your own kind of set of alarms, similar to what the Chinese police are doing, and do the same thing.

And, you know, in Ecuador we saw the kind of classic China thing, where you put a camera outside the house of a person, a dissident, that you want to watch. And so I think we're going to see more of that around the world. And I don't really, you know, think there's anything anybody can do to stop it. I mean, I think the best thing the U.S. can do is probably just set a good example and try to encourage better practices in this. And I guess that's about it. I'll stop there. I had something else, but I forget. So I'll kick it back.

Dr. Kennedy: OK. We'll come back to you. That's extremely helpful, Paul, and really appreciate it. And I want to give Noam and David a chance just to offer some reactions to the feedback from the four of you, and then we're going to turn to a policy conversation about what the U.S. and others can do, if anything. I think there may be some more that we can do, and rather than just watch

this. But we'll come to that in just a second. So, Noam and David, over to you for reactions to the feedback you've heard.

Dr. Yang: Noam, do you want – should I go first, and then –

Dr. Yuchtman: Sure. Yeah, please.

Dr. Yang: Well, so, first, thank you so much for these comments. I mean, these are sort of – many of them are first-order issues in this context, and also conceptually. And it's extremely useful to think it through them carefully. I think there are two points that Noam and I wanted to sort of maybe raise, and maybe complicate the conversation a bit even more so. I'm going to talk about the overall direction of innovation, and Noam can talk about sort of trade and sort of potential complication policy responses.

So a number of you raised the question of the – you know, where exactly – there is one question about frontiers being pushed, and there's another question about where the frontier gets pushed. And we should be very clear that sort of everything we've done in this work is looking very narrowly at facial recognition AI. And then within that sector, you know, I think the economists were coming to think that when a – typically, when a government tries to buy a lot of things, that's going to shift things towards only the direction of innovation where the government's demanding. Now, what we show here is that within the facial recognition sector, when government buys it for civilian purposes – because they kind of use it across multiple purposes and many other sort of sharable inputs – then there are commercial applications for facial recognition that gets developed and pushed forward, that may not get sort of completely distorted by government's demand.

That's a whole separate question, and a much harder question to think that, you know, what about all the other AI sectors that, you know, the companies could have worked on? Medical sort of AI detection, voice recognition, for example, that may now – may now get crowded out because all these resources get spent on facial recognition AI? And sort of Emily has some good points about sort of the basic research. It may be even be too soon to see that. Maybe there's a lot of researchers that are moving their attention away from other applications of AI towards facial recognition.

And to an extent, China is having a lot of share in overall sort of research activity and sort of industrial output in this sector. It could very well sort of skewed – if that's indeed true, it could skew the overall direction of AI research, sort of maybe not towards the direction where we think it's globally welfare-maximizing. So that's a question that we have no answer on, at least empirically. But certainly, it's something to be potentially worried

about and workshopped for. And that will sort of shift some of the policy discussions on where research should be heading.

Noam.

Dr. Yuchtman: Sure. So then on the policy front, I think we see, you know, things very similar, actually, to Ilaria's initial comments on the subtlety and complexity of trade policy, also suggested by Paul T, that I think there can easily be a knee-jerk reaction to our work, which is: Well, everything we're showing you suggests that there should be restrictions on AI trade with China. And I think that has been the knee-jerk reaction in some policymaking circles, and that doesn't naturally come out of our work.

So what comes out of our work, I think, are a couple analogies, actually. So one analogy that we like to think about is child labor and labor standards upstream in supply chains where the analogy is that data collection and extraction is an upstream activity that then feeds into downstream products in international trade, and we in the Western world and beyond, I would say, care about setting standards for upstream production, and that was true with labor standards and I think that remains true for data extraction standards. And as Paul suggested, this isn't something that we care about just abroad. This is a question in our own societies.

And so what that means is that I think we can learn from that analogy, perhaps, and I don't think that the only approach to resolving or the best approach to resolving concerns about child labor or exploitive labor abroad has been trade restriction but, rather, trying to set international standards and inducing firms to meet higher standards by opening markets.

The other analogy that is raised by the last slide that David presented on the trade that we see disproportionately going from China to autocracies and weak democracies in facial recognition AI is the analogy of dual-use technologies and technologies that have some arms or military component to them where we think that there can be geopolitical spillovers in international trade that concern the international community.

And I think, you know, this analogy is one where I think there have been restrictions and, you know, I'm not sure to what extent we think those restrictions have worked to the extent that restrictions produce black markets and so on.

And so I think what we hope to do by highlighting these patterns is to raise these questions, and as economists, in particular, I think one of the strengths of economics – you know, we don't only have strengths by any means but one of our strengths as a discipline is to think about what would be called

the general equilibrium effect, that when you introduce a policy people respond.

And so if you introduce a ban Chinese firms will respond to that ban. They might respond in ways that are unintended by those who have good intentions, and that's true in terms of restrictions on trade in important ways. And so I think, essentially, we want to open up a more nuanced policy conversation, and I think when it comes to changing the behavior of Chinese firms I think we would focus more on inducing better behavior rather than punishment, at least in many cases. And then there's a separate conversation, which Ilaria hinted at, which is trying to enhance American competitiveness, and we could talk more about that offline. But that's a bit beyond the scope of the work we presented today.

Dr. Kennedy: Sure. Well, thank you both for your thoughts on this.

I want to say just one thing about the substance of the – what you found and additional future questions and then the policy implications.

So I think three obvious next research projects would be, first of all, the variation across China. You highlight in your comments that Shanghai and Hangzhou may be better at – I think – I guess maybe Paul Mozur brought this up – that maybe Hangzhou and Shanghai are better than maybe places less developed.

It'd be really interesting to see what the data show in terms of the effectiveness of restricting protest or encouraging innovation and what that looks like and what all the factors are. Opportunity costs – another question that several of you brought up. Is China's hyper focus on AI and surveillance taking resources away from other, potentially, even more beneficial – economically beneficial areas suppressing basic R&D? Or is China so big it can spend money on everything? I think it's a question.
Then international effects – what are the real international effects of this? What do we see in terms of other countries importing the technology and the – maybe the models that come with it.

Let me ask now a little bit about policy implications, and I see Paul Triolo has his hand up as well so I'm going to let him go first on this in terms of what we ought to do. I think there's – in one way, you could just sort of throw up your hands and say, geez, there's nothing you can do. They've got these capabilities and this is it and, in fact, they're part of a global AI community which can't be closed down. It's transnational by its very nature and, therefore, the horse is out of the barn, so to speak, and it's all done. Or you could say there are things that we can do and this is sort of an ongoing evolutionary process without a clear outcome and the U.S. might want to play some more defense along with others. But maybe norm setting, running

faster, investing in more – in students and talent makes this a more open competition and, you know, trying to be a good example as well.

So I'm sure that Paul Triolo has some thoughts on this. Any of the others of you that would like to comment on the policy implications please let me know. We're running a little bit short on time, so I don't want to – I don't want us to go too much farther. Yeah, go ahead.

Mr. Triolo: Sorry. Sometimes I'm not–

I just wanted to say I really enjoyed the petition discussion from Paul Mozur because I think that's really – it's such a Chinese thing, right, and maybe Paul, at some point, should explain what that means because it may not be obvious to everybody. But that's really interesting that that's such a big focus of the alarm.

I want to say I really – I think Noam's comments were spot on there about sort of the – some of these implications. I think one thing, for example, on the standard side on the data extraction and collection side, I mean, the EU AI Act, for example, which we did a lot – I did a lot on for a client, is that one of the pieces of that is attempting to set standards around how data is collected and extracted and used to train AI algorithms, right.

So that's a – for high-risk applications of which facial recognition would seem to be one of them that seems like a really good approach to standards – on the standards front. So if an AI algorithm was trained using data extracted in ways that are not in comportment with certain standards and values then, you know, the EU, for example, is prepared to prevent those applications from being deployed in the EU.

Now, that hasn't – you know, that's – the AI Act is still in draft and they're nearing completion of that act. But that's something the U.S. should at least consider supporting. I know that through the EU Trade and Technology Council that's working with one standard, is looking at that. There's not going to be comparable legislation, for example, in the U.S. on that. But the U.S. officials at Commerce and other places will request support for those kind of standards. That seems like a really fruitful area to go.

In terms of the punishment – (laughs) – issue, again, I sort of raise a company like Megvii. I mean, they really took – they were prepared and I think they put them in place anyway – did a lot of things to try to meet some standards like setting up an internal AI advisory council. And so there, I think it's important that if there is going to be punishment or targeting of these companies there should be a path to, you know, some standards: OK, if you meet these standards you're now removed from the – from the NA list. And

that's something that was not politically possible in the Trump era but should be something that's on the table.

In addition, the question that ends up hitting some of my clients is how far back in that food chain do you go. So if you're providing, for example, storage systems that are used as part of a security camera facial recognition system, you know, is that a problem reputationally for a company providing disc – you know, disc drives or whatever. And so that's – you know, that hits, again, at U.S. innovation if we're going to cut off companies from being able to have access to a market. You know , what's the criteria? Make those clear and what's the ultimate goal of some of those penalties so that you're actually affecting behavior.

Anyway, I'll stop there.

Dr. Kennedy: OK. Let me see. I wanted Emily and then Paul Mozur offer their thoughts on policy.

Ms. Weinstein: Sure. I can just say, really quickly, I mean, I agree with most of what's already been said. But I do think in competing in AI we have to actually take a step back and think about what are the areas in which we actually want to compete with China in AI and there are a bunch of – you know, there are moral implications to that. There are also, like, actual feasibility implications to that, too.

So taking a look and actually, you know, stepping back and looking at, OK, China is leading in X field or in Y field and where do we stand in that field. Is it actually worth us taking the time to catch up in this field or should we – you know, is this something we need to work multilaterally with a different country on with – you know, work with likeminded allies and partners on improving our capability there. Because I do think if we think of this in terms of just a bilateral U.S.-China competition and not globally I think we're automatically kind of screwing ourselves there.

So, I mean, I do think, again, we need to think about this more broadly than just the U.S. and China.

Dr. Kennedy: Mr. Mozur?

Mr. Mozur: Yeah, and I'll just be real quick.

Like, as Paul was mentioning kind of how far upstream you go, I think one question with these companies and in particular with facial recognition is, for instance, if they create a software that allows any given, you know, purchaser to make a blacklist does that become a problem, right.

So if you – because one of the problems in China is that a lot of this is just profiling. So if you can create your own blacklist out of any characteristic you want and you put a certain race or a certain ethnicity into that, then all of a sudden you have, you know, automated, you know, racial profiling.

And so, you know, do we think – and we know that across China the police are doing this and we also know that they're – some of the software vendors, you know, sell it in a way that kind of is a sort of wink and a nod that allows, you know, the police officers to do this on their own a little bit away from the vendor itself.

So I do think you have to think about kind of what the preconditions are there and exactly, like, you know, how this software is catered to these systems as well.

And then the other thing I'll just say real quick is that this may sound hokey but a lot of it does come down to the democratic institutions themselves in the places. The reason this stuff exists the way it does in China, the technology isn't that – I mean, the facial recognition capabilities are extremely impressive but beyond that most of this stuff isn't all that difficult.

What allows it to happen is the political system and the kind of bureaucracies and the organization and all the rest, and so if you encourage stronger democratic institutions you create a better, you know, sort of insulation against this.

And one other thing is surveillance can be very counterproductive. Oftentimes, surveilling everybody also surveilles authoritarian governments' own abuses, and so that's something that I think we may see kind of coming to the fore, you know, in the coming years. It may not be so good to have cameras everywhere because that footage can leak out and show abuses and sort of, you know, create new political resonances and new responses from people when they see abuses, and I'll leave it at that.

Dr. Kennedy: Well, we have had a terrific conversation today undergird by a fantastic feature that Ilaria put together built on a foundation of amazing research that Noam and David have done over the last few years as part of a broader scholarly effort to understand China's technological trajectory and what it means for economies, for national security, for governance.

This has really – and this is exactly what Big Data China's goal is meant to achieve and we're delighted that we were able to collaborate with Stanford and their Center on China's Economy and Institutions to do this.

I think what comes out of today's conversation and the feature and their work – and I hope everyone goes and looks at the feature on our website and

then goes and reads the original research and we have links to all of their work online, and then also reaches out to all of us – a few things is that this is a moving target, right. China's technology innovation is changing over time.

As Emily said, they used to be copycatters. They're now innovating in certain areas. It's going to continue to change and evolve. There is no silver bullet to fix any of this. You push in one way you get a variety of intended and unintended reactions.

So we have to – this requires nuance in Washington, D.C., and, as Paul Triolo said, in collaboration with our friends in Europe, in Asia, as elsewhere. There is a lot to do.

But I don't – I also think that it's not entirely hopeless, right, that there is much we could do. One, we have amazing technology capacities. We also have a political system which, on some days, looks somewhat dysfunctional, I have to admit, but on other days you're impressed, at least what happens quietly.

We really are debating at the most basic level fundamental goals about what do we want to achieve and how to get there and how to balance things, and I think that gives us an amazing advantage even over a place like China where there seems to be, you know, lots of, you know, chests thumping about the advantages of the Chinese system. We have lots of strengths as well.

So I think that we are up to this challenge, but it's not a challenge that's going to be resolved anytime soon. This is a multi-decade generational challenge that we're on.

And I want to thank David and Noam for doing this pathbreaking work that has brought us together to understand this; to our guests, Paul Triolo, Paul Mozur, Emily Weinstein, for really putting everything in context; for Ilaria's great writing and pulling together the feature in a microsite; for the rest of our team, Maya Mei and Trustee Chair, our friends in the iLab who have done terrific work. Everyone, I just want to say thank you to all.

For those listening, we tried to embed your questions in much of the conversation. We didn't get to everyone, which really means that we've got more to do. We're going to come back to this topic again and again, and so welcome everyone's feedback.

Again, thanks, everyone, for joining, for the great work that you've done, for the great work we know that you're going to do. So wherever you are, thanks so much and have a great day.