THE BROOKINGS INSTITUTION

WEBINAR

HOW ARTIFICIAL INTELLIGENCE IS
TRANSFORMING THE UNITED STATES AND THE WORLD

Washington, D.C.

Thursday, January 7, 2021

PARTICIPANTS:

Welcome and Introduction:

JOHN R. ALLEN
President,
The Brookings Institution

Discussion:

DARRELL M. WEST, Moderator
Vice President and Director, Governance Studies,
Senior Fellow, Center for Technology Innovation,
The Brookings Institution

AARON KLEIN
Fellow and Policy Director, Center on Regulation and Markets
The Brookings Institution

JOSHUA P. MELTZER
Senior Fellow, Global Economy and Development
The Brookings Institution

CHRIS MESEROLE
Director of Research and Policy, Artificial
Intelligence and Emerging Technology Initiative,
Fellow, Center for Security, Strategy, and Technology
The Brookings Institution

MARK MURO
Senior Fellow and Policy Director, Metropolitan Policy Program
The Brookings Institution

NICOL TURNER LEE
Senior Fellow and Director, Center for Technology Innovation
The Brookings Institution

* * * * *

P R O C E E D I N G S

GENERAL ALLEN: Good afternoon ladies and gentlemen, and a happy new year to you all. My name is John Allen and I'm the president of the Brookings Institution and it is a great pleasure to welcome everyone today to our event on artificial intelligence entitled, “How artificial intelligence is transforming the United States and the world.”

To state the most critical point right up front, artificial intelligence, AI, is well on its way for being one of the most impactful technologies ever created, and indeed the full suite of related technologies, which include computing power, sophisticated algorithms, big data analytics and more, and I'll repeat, will ultimately be one of the defining technologies of the 21st century. These are technologies which effectively cut across virtually every sector of our society and certainly every sector of what we do at Brookings. They are as multifaceted in their impact as they are in their ability to transform the world around us, and as you will hear soon from our scholars, our experts, they represent one of the greatest opportunities, or one of the greatest challenges of the modern era.

Through the recent creation of our artificial intelligence and emerging technologies initiative, which brings together scholars from across the institution studying the intersection of AI in fields like global finance, and national security, and international norms, we placed great emphasis on this field. I lead the program along with Brookings Fellow Chris Meserole, who you will hear from shortly. But our institution is dedicated to the study of artificial intelligence and other technologies and to assessing in-depth ways in which our global community will be defined by new and emerging technologies writ large. And in that regard, the AIET initiative as you pretty called, works very closely with our Center for Technology Innovation laid so ably by our Senior Fellow Dr. Nicole Turner Lee.

Now if one looks at any of the issues of the day, nearly any one of them or all of them will soon have a technological dimension which will enable the integration of AI in one form or another across so many of the sectors. Take for example climate change, or COVID-19, two existential challenges that we face as a society. And in this case, they are entirely disconnected by their causal factors or their data sets. But AI through the advanced use of big data analytics powered by -- empowered by advanced computing and even in the relatively near future, quantum computing, can tell us just as much about the

spread of the novel coronavirus as it can about our shared climate future out to 2070 and beyond. It's a remarkable capability.

That's not to say that AI is all-powerful. Indeed, some of our panelists today, I hope to, and will likely highlight that this technology remains uniquely limited. But it's access to good data and great algorithms and seemingly next-generation computing power in its applications while still narrow, offers us great opportunity. So we have a long way to go before AI resembles anything like true general intelligence, but it's a future we can imagine and one that we must be prepared for.

So with that in mind, allow me to briefly introduce our panelists who collectively assess -- who will collectively assess the issues associated with AI from a wide variety of perspectives in subjects. First we have Senior Fellow Aaron Klein, previously a director of the Bipartisan Policy Center's Financial Regulatory Reform Initiative. And he is focused on various issues such as financial technology and regulation technology -- regulation, technology, finance and more. He has also served from 2009 to 2012 as the deputy assistant secretary of economic policy with the Department of the Treasury. Welcome Aaron.

And joining him, we also have senior fellow and director of the Center for Technology Innovation who I just briefly mentioned before, Dr. Nicole Turner Lee. With nearly two decades of experience, Nicole has researched and written extensively on the digital divide and racial biases in machine learning algorithms. She is also author of the forthcoming book, “Digitally Invisible; How the Internet is Creating a New Underclass,” which will be released this year. And Nicole, I'm really looking forward to this book and really looking forward to the day we do the book event for everyone.

Third, we have Josh Meltzer who is also a senior fellow and he currently leads the Digital Economy and Trade Project at Brookings which examines the intersection of data, the internet, and the future of international trade. Previously a diplomat at the Australian Embassy in Washington here in D.C., Josh has folded his experience working on issues of trade, climate, and energy into his work with us today at Brookings and has folded it in in the context of artificial intelligence and global technology norms including his participation in a major transatlantic dialogue on trustworthy AI. Welcome, Josh.

Now with us on the panel also in mentioned earlier is Chris Meserol. Chris is a director of

research and policy for the Artificial Intelligence and Emerging Technology Initiative. And Chris, of course, has extensive experience in studying the intersection of emerging technology and international and domestic security matters as well as coordinating our broader institutional efforts on the study of AI. And his insights have been featured in outlets such as The New Yorker, Foreign Affairs, and Wired and is widely considered a leading expert on AI and tech policy today. So Chris, welcome.

Also, we have senior fellow and policy director, Mark Muro. Mark focuses on what he calls, "the digitalization of everything," which is frankly terrifying. He has authored several publications on automation, on artificial intelligence, as well as the digitalization of the American workforce. It's a tremendous work and his work regularly finds itself into the hands of many of our most prestigious members of Congress. And Mark, thank you for making the time to join us again today.

And finally, we have vice president and director of Governance Studies, Darrell West. Also a senior fellow in the Center for Technology Innovation, Darrell has written several award-winning books on the future of artificial intelligence and automation, to include a recent book where he dragged me along in the process "Turning Point; Policy Making in the Era of Artificial Intelligence." And his current resource -- research focuses on artificial intelligence, robotics, and the future of work. So welcome Darrell and thank you for your leadership across the board in the institution on these many matters.

Thank you all for joining us. I think the audience will agree, especially now, but in the aftermath of this event today, the tremendous expertise that we have at the institution on this matter. Before I turn the floor over to Darrell who will be moderating today's events, a quick reminder that we are very much on the record. Should the audience members like to submit questions, they are welcome to send them via the email address, events@brookings.edu, that's events@Brookings.edu or tweeting @BrookingsGov using the hashtag #AIGovernance, @BrookingsGov using the hashtag #AIGovernance.

So with that Darrell, thank you again for leading this and moderating this panel today. And the floor is yours, sir.

MR. WEST: Okay. Thank you, John. And we appreciate your leadership in the areas of AI and emerging technologies. I know it was just several years ago that you highlighted the importance of new technologies and encouraged Brookings scholars to work on them. And the good news is, as John

pointed out in his various introductions, Brookings scholars have risen to the occasion. There are great AI projects taking place in each of the research programs all across the institution. So there's lots of exciting work rolling out and you're going to hear a lot about that today.

So I'm going to start with some questions for our panelists. And then we will save time for questions with the audience. And John gave you either the Twitter hashtag or email address that you can use to submit questions. And for those of you who would like to get any additional information about technology policy, you can check out our Tech Tank Blog and our Tech Tank podcast that we inaugurated a couple of months ago. That covers some of the latest issues in technology, the trends, the problems come and some of our ideas on how to address those issues.

So I want to start with Nicole. You write a lot about bias in AI. Can you give us some examples of how AI is being deployed and what kinds of problems that we are seeing?

DR. TURNER LEE: Thank you Darrell, and I appreciate John's very kind introduction. I'm so excited for all of you to join us today because this is what we do at Brookings. And so it's important for us to sort of talk through that. As has been mentioned, the work that I do is interesting and I'm going to tell you why, because I am a sociologist. I'm not a scientist. I'm not a computer engineer. I'm somebody who studies systems and systems of discrimination and oppression and how they overlap with other things in the world. And my interest in technology which dates back many, many years really has brought me to this place to sort of understand how we look at emerging technologies and what the implications are.

Clearly, we all come to the table when we think about the development of technology and the use of technology with our values, assumptions, and norms about how we see the world. And as a result of that, when that technology is actually designed, those things actually factor into that. And so if you are a person who is a developer who is developing products and services that are going to be applied, let's just say for example, pain threshold algorithms that will apply on women and you are a man, obviously that technology is going to be designed in a way that actually preferences men.

The same thing when it comes to race. If you are a person who actually designs technology and you are absent in terms of the people at the table who have the lived experiences of

African-Americans, Latinos, and other people of color or people of different abilities, you are going to develop technology which starts out externally flawed. And I think Darrell, that's really important for people to hear. It is not that computer is biased; people are biased because we are biased. And if you start from that assumption, you begin to look at the design and development technology in ways that once it leaves the lab, it has the possibility of unintended consequence.

One more thing in terms of racial bias and I will give a couple of examples. When we think about how these models actually operate outside of a technology lab, they operate within the context in the world in which we live. And we cannot divorce that world. I mean, look at yesterday what happened. This is who we are as a country based on the historical remnants of slavery and discrimination.

When you begin to think about how that manifests itself in technology, it does the same thing. What may have started as a really great idea to actually develop an algorithm for patients with chronic disease to actually get eligibility for our program, may end up excluding, and this is a real-life case, Black patients simply because of the variables that you used was the cost invested in healthcare. We already know externally and historically that African-Americans and other people of color do not have quality access to health insurance and as a result they do not spend a lot on healthcare.

It is examples like this or facial recognition technologies that are designed in ways that they are not preferenced to have the greatest specificity for people with darker complexions, or Black women who change their hair. So as a result that technology will not work on me. It will work on the people who were out yesterday, but it will not work on me.

And I say that my friends because when we begin to think about racial bias and technology which is one of the reasons we at Brookings work on this, it's not so much about a shiny object having these technical flaws. We are talking about the unintended consequences that have implications on people's lives. The ability to get credit or not. The ability to actually get into a healthcare program that can save your life or not. The ability to be misidentified as a criminal for regular civilian.

These have implications and unfortunately when we look at areas like healthcare, employment, critical justice, financial services, these have consequences that extend the long legacy of

discrimination and bias that is already foreclosed on opportunities for certain populations to be participating. And we don't want to see that with technology. Technology should never become a barrier to that type of entry.

So I will stop there, Darrell, because you know I can go on and on and on. But I think it's so important that we actually route these conversations around the potential for unintended consequence and how that actually ties in to the externally flawed society in which we live, and factor that in as we develop the cadence for these models to be much more equitable, responsible, and inclusive.

MR. WEST: Thank you, Nicole. That is a great summary and the work that you and other colleagues are doing on this issue of bias is very important. People who are interested in reading more about that, we actually have a series of papers that have come out over the last year or so, on AI bias and how to address those issues. So you can find those papers at Brookings.edu.

Mark, you look at the geographic disparities in the digital economy and how occupational issues are playing out in various metropolitan areas. How is AI affecting the economies of different parts of the country?

Mark, you need to unmute yourself.

MR. MURO: Yeah, I hope -- thank you so much. I hope everybody is starting the year amidst all that is going on.

So yes, my work at the Metro program really focuses around the interplay of space, work, and technology. We think it's really important to nail down some of these relationships. We've looked at the geography of U.S. advanced industries, the shifting digital content of work over time, we are seeing massive elevation. And then how these affect social inclusion, regional divides as you suggested. And then we've looked at automation exposure and most recently AI.

Our approach is usually to work with occupation based measures of exposure to particular technologies to try to get very granular metrics. For example, in the case of basic automation and robotics, we use the occupation specific exposure scores from McKinsey and Company to link occupational susceptibility to the demographics industries and ultimately geography, which we think is extremely important here. In doing this we confirmed that lower paid, routine, cognitive tasks, and

occupations appear to be highly vulnerable to the existing technology. From there we showed that meant that more Black, brown, and young workers would be inordinately exposed, so really nailed that down with data.

And then we show that given occupational mixes, which is an important linchpin here, you can then get to looking at very distinct regional variations between cities and regions. In the case of automation and robotics, heavy focus in the heartland. But AI has been harder to work on. Researchers have simply had less time to learn about AI's basic use cases. There is less widely shared agreement for researchers to draw on where, say, machine learning systems excel, and this little agreement on workforce impact.

But given that (inaudible) we found an interesting way forward. We worked with a great colleague at Stanford, Michael Webb, who has developed a way to overcome the general lack of detailed assessments of AI's use in the economy by using AI to solve that statistically. Basically, Michael used textual analysis of the language of AI patents to identify AI capabilities and then machine learning in the form of natural language processing to quantify the overlap between patents and job descriptions. Very interesting strategy. So again, keeping it very real with what is happening in the economy.

So by mining verb/object pairs such as predict traffic or diagnose disease that are both in patents and in job descriptions, you can get to again, this granular drill down. So what did we find? Well, our analysis suggests that AI has its very own range of use cases and a really distinctive impact that is going to be quite different than what we know about automation and robotics, for instance. Quite the opposite of robotics automation, AI exposure looks to be across better educated, better paid workers for now in professional settings, white-collar professional tool frequently. And so advance manufacturing is still a heavy user.

So you get again, very distinctive patterns. Yes, the industrial heartland is going to be saturated with these technologies and already is in the manufacturing world. But also interestingly, coastal professional services, metros like Seattle, the Bos-Wash corridor are saturated with this in business contacts.

So what should we make of these initial patterns? I think there are tons of open

questions here. The first point is that it's very early days that these -- but we need to know that these systems are going to be distinctive and different from what we know about other technologies, and certainly automation is a very different animal. We have a lot to learn about these impacts. And I think it's important that we -- and it's important to say that our work quantifies potential impact. It doesn't say what kind of impact yet. I think that's a huge, open question. Are these positive networks impacts? Negative ones? How does that vary? That raises more questions.

You know, I think you could see it playing out in different ways that are going to be important focuses of research at Brookings. On the positive side, AI may play out as an extremely beneficial complement to white-collar workers in places given those workers in places -- education levels, agility, access to training that I would say others don't get. So in that sense, maybe marketing executives and radiologists, et cetera, supported by AI will be able to offload tedious aspects of the work and spend more time on the higher value aspect.

But on the negative side, we can certainly see AI adding to the divides. If it -- if we have a super empowered professional core of the economy abetted by powerful new technologies so that the economies winners pull away even more, and meanwhile these technologies for some could be a scary substitute for well-paid human work, the new angle replacing thousands of mid-level professionals now.

So I think an important take away is AI is going to be extremely pervasive; between it and automation of other forms, no group is really going to be left out. I truly can imagine both possibilities up and down. It's likely going to have very distinctive and uneven impacts and we need to try to understand those impacts in advance and begin preparing for them now.

MR. WEST: Thank you, Mark. And for our viewing audience, this is very important work because as Mark has pointed out, the geographic differences in terms of how technology is playing out, how AI is being deployed, are quite substantial across the regions, across the states, and across the city's. And of course those geographic disparities play a very important role in the current political polarization that we see because so much of our economic activities on the East Coast, the West Coast, and then a few metropolitan areas in between, lots of other areas are being left behind. They are upset and this contributes to our current political distress. So very important that you look into those questions.

Aaron, you have done a lot of work on global payment systems and financial technology. How is AI and technology affecting the world of finance?

MR. KLEIN: Darrell, thank you very much for having me. It's a pleasure to join this conversation.

This is a huge impact. Let me start by asking everybody was watching, are you watching on a Mac or a PC, or Android device? If -- depending on the answer to that, I know a little something about how likely you are to repay a loan. In point of fact, the data show that Mac users holding income, all other proxies you want, pay back a little bit more often. Should I offer you a lower rate on your credit? Should I offer you a lower interest rate on your car payment? Or a higher rate? Is that allowable or legal?

But the way we traditionally set money out and the way that we have -- our entire financial system tends to rely on these things called credit scores, the most famous of which is a company called Fair Isaac known as FICO. And these credit scores are kind of made to look like your SAT score and you think they have all this fancy underwriting behind them. And the reality is they don't.

They are fundamentally based on a file of information that is largely incomplete, highly inaccurate, one out of every five Americans has an inaccuracy on their credit report. Download -- you can download one for free. I hope to write a law in Congress to allow that. I found the Aaron Klein in New Jersey who didn't pay his cell phone bill, which has forever hurt my credit report. Not me. We have the same carrier, but I can't get it off because if there is an inaccuracy in your report, the credit bureaus are under no legal responsibility to fix it.

And this system based on whether or not you use to repay your credit reports is what the entire society uses to decide whether or not to give you credit, and if so, at what rate. Well, it turns out this system is not very accurate. I've said for those of you who went to band camp and were kind of geeks like me, it's like an out of tune oboe that the rest of the financial orchestra is tuned to. An orchestra tends to an oboe not because it's the most accurate note, because it's the hardest instrument to tune. Society is tuned to these credit scores not because they are accurate, highly accurate, no, what we find out is they are poor predictors of the future, but we all use them anyways.

Artificial intelligence can easily surpass these models, but they do so using information we don't necessarily fully understand, and information that is highly correlated with protected status in classes. And as a society, we have created different barriers legally to the provision of credit and finance than we have to other services. And even within the provisions of finance they are inconsistent.

So Darrell, you remember I started out by saying are you on a Mac or a PC? Another question I can ask is, is your name in your email? Research has shown that people whose name is in their email are better credit risks than people whose name isn't. While on the one hand, email is something you totally can control. On the other hand, Dr. Turner-Lee will remind me there are a host of valid sociological and economic reasons why people with nontraditional names may choose different email addresses. There is plenty of economic research showing systemic racism to people with clearly ethnically identified names on resumes or email addresses. Could you use that or not?

Remember when I told you that thing about Mac or PC? Did you know that some of the nation's largest chains started quoting people who went -- shopped online a higher prices with the Mac than the PC because Mac people are richer? That's legal to do in a store. That's probably illegal to do if you're providing credit.

So how do I get to close all of these different taxa? I ask you to think about the world where we currently live, where current technology decides credit allocation in a highly biased and systemically distortionary, and I might add, inaccurate way.

Now, think of all this new technology that we can apply. Well, sometimes we can apply to find people who are better credit risks, but in a way that's more discriminatory. Do we want that? There is a tension between accuracy of credit allocation, what we call risk-based pricing, and discrimination.

I will give you an example. Men are riskier than women, controlling everything else. It is blatantly illegal to charge men a higher rate than women when they are all the same. In effect, women cross subsidize men's financially risky behavior throughout society and I thank all of the women who lowered my interest rate as a result of that.

We don't allow that cross subsidy in driving. Men are worse drivers than women. Men pay a higher premium. Anybody watching as a teenage boy, and a teenage girl will be well aware that

simply on the basis of gender at age 17, and having nothing to do with their driving track record, they are charged different prices. What as a society do we want to allow to be used? What do we not? What technology can make the allocation of credit more fair even if it is not totally fair? Do we want to allow that or do we not?

These are the questions that AI is going to force upon us as intelligent and sophisticated lenders in the banking and nonbanking system use this technology. And it's one that we better be prepared -- a challenge we better be prepared to meet.

MR. WEST: Okay. Thank you, Aaron. And I should point out Aaron just taped a Tech Tank podcast with us on the gig economy. It's a great job out there kind of looking at how the gig economy is developing, what issues, what problems are rising there, and how we can think about that. So you can find that at Brookings.edu.

Josh, you work on the transatlantic aspects of technology and how various regions and countries around the world are handling these issues. I know you pay a lot of attention to the United States and the European Union, but then you also look at Australia, India, and Singapore. So I'm just curious how kind of various regions or particular countries are addressing some of these tech problems.

MR. MELTZER: Yeah, thanks Darrell. And it's great to be here with this group of colleagues. And thanks for you and John for putting this important event together.

Yeah, so as John had mentioned in the introduction, we've put together a forum for building on international cooperation with AI. And I will get to that in a little bit more -- a bit later. But just before I do that, I just want to sort of take a step back and give a little bit of framing of where we're coming from and why we thought this was a worthwhile and important initiative.

As you sort of put it, we are very much focused on the international elements of AI and really around building a cooperation on AI. And I thought my colleagues I think really just pointed out in very clear and precise ways, there are various really significant ethical dimensions to how AI is actually going to be rolled out and applied across United States, but really globally. And so having a sort of common, essentially ethical foundation to how allies it should be in fact used is an important endeavor which has been going on for a while. In fact, there's a lot of important work in progress on actually

developing reasonably consistent sets of ethical principles around AI. The OECD has agreed a set of ethical AI principles. The U.S. has released its own set of essentially regulatory principles for AI. The EU has done similarly.

Countries, companies that are developing AI, Google, IBM, Microsoft, have all got their own forms of ethical AI principles and so forth. So there's been a lot of important work on that. And so we are sort of seeking to essentially build off some of these common frameworks on ethical AI to think about the next steps about how difficult AI actually gets developed and implemented and do it in a way which is cooperative and consistent.

The other element for us is the sort of trade and international economic aspects of it because again, I think what is being make clear is how economically significant AI is going to be. So everyone sort of wants a piece of it in some sense. And you see governments around the world now very actively developing a lot of AI policies. Some are throwing fairly significant resources at trying to develop AI and innovation. And government in some areas in particular are now really moving quite quickly into sort of the AI regulatory space. And I think the EU in particular is probably at the forefront of that.

And so I think that draft very much created the sense of providing and creating a place where you can get key actors and stakeholders together who are sort of sufficiently ethically aligned on what the governments aspects of it should be to talk about how we go about regulating and providing guidance for the actual development application of AI that can be to the extent possible, cooperative or at least minimize differences in the tensions that might arise if we go often particularly off in all directions in that respect.

And so that was sort of the genesis I guess of this forum for cooperation on AI. I do this in partnership with Cam Kerry who is obviously very well-known and in Governance Studies. We have a partnership with Andrea (inaudible) from (inaudible) in Brussels. And that was the beginning of the project. And as you mentioned, it was initially a trans-Atlantic dialogue which involve the United States, the EU, the U.K., and Canada and multi-stakeholders in fact has engaged (inaudible) government officials, industry, and AI experts from academia. We have now expanded it out to include Singapore, Japan, and Australia. And it is essentially built, at the moment, around a series of regular AI dialogs.

We've had six of them over the course of last year.

And we started by getting into the details of the top regulatory issues and also issues around innovation, R&D that come up when you think about where there might be areas both for divergence and tension as well as corporations. So far, for instance, issues around risk regulation and how do you identify the areas of AI that might present new and novel risks that may require additional regulation? Or where do we have existing laws that may actually be applicable? What's the role of international standards and international standards bodies trying to drive some commonality and convergence on a range of regulatory AI issues. Issues around conformity (inaudible) particularly when you think about AI in the so-called wild that will actually essentially develop and learn from novel and new data sets even once it's been released. How do you assess conformity then with ex ante regulation?

And then certainly also around data governance because it's well understood that while there are aspects in the computer science -- areas of computer science which are looking at ways of building AI which is not heavily dependent on data, data still remains a core input into AI models and there is a whole world of data governance which includes issues around privacy, and cyber security, and access as well and make sure things happen the way in which -- you know, again, just emphasizing the importance of international cooperation in order to minimize tensions, but also really to drive opportunities.

I mean, certainly one of the things that we think is very important I think is expanding the sort of opportunity of AI globally, not having it as simply sort of a developed country opportunity. I mean the inequality drivers that you're going to see on AI within countries I think risk really playing out globally and actually driving inequality across countries. So creating AI opportunities both on the innovation front and the application front is also important. So there the governance part of it, I think that comes into play is also significant.

And finally, I will just talked briefly. I just want to emphasize parts of the R&D and innovation part. I think what is clear so far is that AI is a very exciting technology with enormous potential, opportunity, and risk. But it is still in so many ways still an early, nascent technology that has really thrived in so many ways on international collaboration and research. I mean, it's remarkable the level of

sort of international engagement by researchers really globally on pushing the AI frontier. So it's birth is intrinsically and inherently global. And in many respects, a lot of the other inputs, whether it's data or infrastructure such as chips in the semiconductor world are also going much global.

So it is global in its genesis and expanding that and making sure that that remains the case and that the opportunities are extended globally, I think, also remains an important focus for us as well. That's where we also see really a lot of opportunity for cooperation around actual practical innovation and R&D questions like how do we scale up international cooperation on AI to deal with common concerns whether it's around health or climate change or development and so forth. So let me just leave it there. But thanks a lot.

MR. WEST: Okay. Thank you very much, Josh. And I agree 100% on the global aspects of this. Every country, every region is trying to figure out these issues and it's very important that you're bringing people together, having these international dialogues and trying to figure out the best path forward.

So Chris, you look at the security aspects of AI and how technology is affecting military preparedness. How is AI affecting national defense?

MR. MESEROLE: Well first, thanks so much Darrell, to you and to John for hosting this event.

The -- you know, the military impacts of AI, I think there's really two conversations that need to be happening. The first is there is a really robust conversation happening around the impact and use of AI within military or traditional kind of interstate conflicts. The second conversation that's really ongoing and equally important in my view is the use and impact of AI within domestic security context. And I want to talk a bit about a little bit of each of those.

The -- in an interstate or kind of traditional military complex, there is really three core debates and questions that need to be talked about. The first is what to do about lethal autonomous weapon systems and the use of AI within military systems that have a warhead attached and that can cause lethal -- that can cause fatalities and deaths. This is something that a lot of researchers in the space and military analysts are working on and thinking about. The fear here is that as AI matures and

as more and more weapons systems are ported over to use AI and incorporate AI, that they will begin to be -- the speed at which those weapons are used and the speed and intelligence with which they are able to kind of identify and maneuver to and engage targets, will begin to outstrip the speed of human thought at which point we will start to have weapons being used without really adequate human oversight.

And so the big debate happening both within democratic militaries around the world, but also within a number of different intergovernmental bodies right now, is how to ensure that there is an adequate level of human oversight of these systems as they mature and as they continue to develop over time. And there is a robust -- the Pentagon, the Defense Innovation Board have kind of come out and released a set of principles about how they will use AI. I think you will see more and more countries do the same thing.

But it's a really thorny challenge because even if you are democratic and even if you have really strong ethical commitments to make sure that these weapons systems are used appropriately, if your adversary is developing weapons without similar kind of ethical constraints, there is the scenario in which their weapons systems will be faster and more effective than yours. And we haven't reached that point yet, thankfully. The AI we have is still not sufficiently mature to be used in conflict environments. But it's coming and I think there is a really important debate to be had about -- at a global level about how to ensure that the systems are used safely.

The second core debate happening within a military context has less to do with weapon systems and more to do with the AI that's used all across the entire array of military systems and applications from logistics to personnel management. I think many people forget that in many countries the military is the largest employer in their country. They have legitimate reasons to be tracking their own employees, soldiers and analysts for insight or threat detection, and things like that. But the data that they use and algorithms that they use to kind of make those assessments or even if not for insight or threat detection, things like billeting or promotions, those are going to increasingly rely on AI.

And I think we've seen in the private sector that there is a lot of problems that come with that, that Nicole, and Josh, and Aaron, and others can speak to. So there is a robust set of questions and conversations happening as well around the use of AI in kind of military systems that don't necessarily

have a warhead attached, but are nonetheless very consequential for the militaries themselves.

The third military debate that's really one that I've been focusing on in my work is the impact of AI on international security and stability, which is less about the specific weapons systems and more about the strategic interaction between the leading developers in militaries that are using AI. So if you're military and you are afraid that your adversary is about to get a new and improved version of AI that may be able to outcompete yours and that would put you at a tremendous disadvantage on the battlefield, you now have incentive to kind of engage in a kind of -- you know, to engage now rather than down the road where you'll be at more of a disadvantage.

And conversely, there are -- in theory you can have a situation where a military that is not sufficiently constrained by its own norms and principles, might discover that it has a decided advantage with AI at a particular moment in time and that they are afraid that their adversary might catch up to them. So therefore, they should actually use that now; which in both of those scenarios it needs you pretty quickly to a pretty traditional kind of classic security dilemma where both sides begin to invest more and more in their AI and you can quickly come to find yourself in a conflict scenario without either side really intending to.

And so to head off that color spiral scenario, there is a lot of work going on at a bilateral level and a number of multilateral levels to come up with new confidence building measures and mechanisms to ensure that the major militaries that are using and developing AI can deescalate and deconflict conflict in a way that the development of AI doesn't actively trigger complex that no party really intended to happen in the first place.

The last point I will make, it just comes back to less the military international use of AI or impact of AI insecurity and more its impact of domestic security. This is another area where I've spent a lot of my own time and energy thinking about, which is; there are a lot of regimes around the world -- most regimes, whether they are democratic or authoritarian, most regimes want to stay in power. And thankfully in democracies, we do have elections and we have kind of peaceful transfers of power, or we should.

But the -- in authoritarian regimes, one of the uses of AI, if you have sufficient data on

your population it becomes very easy for you to identify and track people who are -- you perceive as threats to your regime. And so we've seen authoritarian regimes in Russia and in China developed a really robust security apparatus based on AI and machine learning and data analytics to maintain their own -- maintain control of their own populations.

So if you go to Beijing for example, as soon as you walk off the plane and through the airport and all the way down through the center of Beijing, you're going to be tracked constantly by video cameras that can match your face in near real-time and track you across the city just by virtue of your face and the AI and the pattern recognition that they can carry out. What that allows them to do is track individuals or anybody that they perceive as a dissident or a threat to the state in real time and begin to map out the networks of dissidents. And we were seen them do this now, for example, in Hong Kong.

And it's a new form of authoritarian surveillance and way of maintaining security and assuring domestic security that we haven't really seen before and we are just beginning to, I think, understand the full dimensions and scope of it. So it's something that I think democratic societies in particular who are also going to be using things like facial recognition for maintaining policing or maintaining domestic security, we are going to have to really make sure that we get the governance structures right for those so that those technologies are not abused but instead are harnessed for the better -- for the greater good.

So there's a lot of really interesting debates happening with AI insecurity. Hopefully we will lead to a better world rather than a worst one, but we will have to have more conversations like this I think to get there.

MR. WEST: Thank you, Chris. Those are all terrific points. And I know John Allen talks about this concept of hyper-war, which is consistent with things you been talking about in terms of just how AI and other technologies are going to speed up decision-making and create some additional risk and possibly create greater instability. So it is important for us to figure out how to navigate that new type of space.

So there is a two-part question I would like to direct to each of you. Each of you have identified particular problems with AI in emerging technologies. So one part of the question is, what

remedies should we be adopting to deal with the particular issues that you've identified. And then secondly, since we are about to inaugurate a new Biden administration, how do you think Biden and his administration is going to deal with some of these tech problems?

So Nicole, you mentioned the problem of AI bias. What are the remedies and what your expectations in terms of how the Biden folks may treat these issues?

DR. TAYLOR LEE: Well, I think -- thanks Darrell for that question. It's good to hear my collies kind of talk about their body of work because as you see, it all ties together. And I think first and foremost when we talk about remedies when it comes to algorithmic bias which has become a pretty big buzzword among the policy community simply because we are still sort of contained to the Black box thinking that the inputs actually affect the outputs; and what we are finding in the technology that there is a whole universe of inferences that actually determined the extent to which the input actually looks different on the output side.

What I mean by that is if you've ever been to a beach and you see a snake in the sand, when the snake is actually released onto the sand, it picks up these other particles. So what may have started as a model that was trying to determine eligibility for a mortgage may actually land up the algorithm picking up particular behaviors in the case of myself where figures out, I'm a mother or I may be African-American because I read certain magazines or I visit certain shops and as a result it denies me a loan.

With that being the case, there is obviously a remedy that we actually have to really engage in rather quickly, which is the technical cadence of these models. Often times there are not sociologists that sit at the table, ethicists, lawyers, privacy specialists that actually can give that type of context to the model that is being designed in the context in which it is being deployed.

So I've been suggesting in my work, and a lot of people actually hear me on this, on how we develop standards. Standards that come after NIST has come up with some of their standards in terms of the criteria that should be part of these models.

But what are the standards that we have in terms of the toolkit of bias impact assessments prior to the model being launched? Having some type of diversity assessment of who is

sitting at the table in the model development. Having the right checklist of questions that we need to ask before going forward. Those kinds of questions, Darrell, I found when I talk to a lot of folks that are working in this, whether in academia, in the private sector, or public sector, they don't do. And as a result, these unintended consequences come along. And as I tell people, they may end up being on the first page of the newspaper simply because they did not go through this preassessment. Given that technical cadence, the next step is to start to think about where do we need to put, as a remedy, the regulatory or legislative pressure? At what point in the model do we need to start paying attention?

Now, compared to my colleagues on this event, I like black high-heeled boots and I am one of those people that when I get a delivered ad of a black high-heeled boot, I check it and I get more, and more, and more ads around black high-heeled boots. But what I don't like is when that high-heeled boot lends itself to these other variables, again, that snake analogy, that determines that I'm a mother and may determine that I'm Black, and it may determine that I spend too much money on the high-heeled boots and maybe because I use a PC they think I'm a credit risk. When you start to put all that stuff together, now I'm impacted in a way that we actually have laws that govern civil rights around people.

So I think another remedy is to really think about what are the policy implications of having AI technologies that are not guided by civil rights protections that are already on the books? They are implicitly applying, but they are not explicitly practiced. And I think it's particularly important that we think about compliance in those areas where we know we've already litigated these concerns, financial services, employment, education, healthcare, et cetera. So I think that's another remedy that we need to think about. We also need to think about where there are cases or examples of remedy when we need to actually find ways to coevolve these models.

So going back to my earlier example, if facial recognition doesn't recognize me because it does a really horrible job based on what research has been done of identifying people of color, particularly African-Americans, and particularly African-American women who change their hair, then potentially we need to have as a remedy, regulatory sandboxes, or safe harbors where we can coevolved together. Ask me to actually be a test case so that technology is actually performing much better when you actually have subjects that you're not making assumptions about their background or other

characteristic.

I think it's equally important that we think about particularly algorithms as a remedy, the training data sets which we are actually relying upon. We know for example, that many training datasets rely upon again, historical data. But if historical data starts out flawed, criminal justice data says that over represent people of color, mugshots being the primary use of training algorithms on facial recognition, that's a no-no. Healthcare datasets that actually are misrepresented or under presented in all types of demographics that could essentially tell you right or wrong in terms of how you are actually optimizing the variable.

It's really important and I want to put this out really clearly for Congress as we think about this latter part of the question; that Congress begin to think about allocating resources for more diverse and inclusive data sets. If companies aren't going to do it, we need to do it and we need to figure out how those datasets become available particularly again when we are looking at sensitive use cases.

And Darrell, one more thing I would like to say as a remedy that is not often shared, it's really important that we engage citizens in this process. So some of you might have heard me talk about some work that we're doing not just on official recognition project here at Brookings. We will be announcing this shortly in terms of looking at facial recognition, and what the regulatory pressure needs to be to have the right guard rails to protect underrepresented citizens. But Darrell, I'm also working on this energy star rating. And this energy star rating brings in the civil society aspect. So just allow me for the last 20 seconds to tell you what that is as a remedy. And I'm saying it loud because I'm hoping that somebody from the Biden administration is listening because you need to use it.

First and foremost, I went to buy a dishwasher last year. My dishwasher broke and I went into a big box store and I realized that what guided me towards wanting to buy the dishwasher that I ended up with was that big energy star rating label. And I think as we look at the distribution of algorithms, the ways of their affecting our daily lives, we need some better housekeeping skill. Potentially, I call it the industry star rating. It's not necessarily a new regulatory agency or any type of guidance that allows for some type of enforcement, but what it does is it actually generates the type of competition that we need among developers to have the highest performing algorithm with the least risk

of bias.

And what that means is, not only do you look at what I said, the technical cadence, and you pay attention to where there is regulatory pressure or sensitive use cases where there -- you may have liability, but you also think about the engagement of citizens in this ecosystem to tell you when you're wrong. Right now we don't know when these algorithms are wrong as it impacts underrepresented people. And to me in our work here at Brookings, that's particularly important. The denial of a loan to a person who is already credit risk -- a high credit risk, means that they will have to kick the can down many more years before they actually come into that space of having the eligibility determination approved.

The denial of bail or the increase of sentencing at a criminal justice algorithm that overrepresents or over criminalizes people of color means that even though the science may actually perfect the human decision-making capacity of a judge, we still are over incarcerating people who are already over incarcerated in ways that the science now makes that determination. To me that's scary because it never breaks the virtue cycle of persistent discrimination.

So I would say to this administration, you've got your hands full. And part of the challenge going forward is to determine as racial equity is in square center with some of the policies that we're going to see come out of this administration, do we have the means to have an algorithmic accountability counsel or advisory group or a caucus that's paying attention to that.

I know in our work here at Brookings we've testified before various House committees and the Federal Reserve I will be talking to next week. But the challenge is, do we have the right mix of actors, stakeholders from the private sector, the academic community, the policy community, civil rights community to have an all hands on deck conversation on where we need to impose the right guard rails and where do we need to back up and let innovation be what it is.

And I think for this administration, if we're going to get it right, start with making sure that we don't have a Wild West without compliance to civil rights laws. But then let's go back and figure out if we can work together to organically create the types of models that are inclusive and responsible, and when need it, have the appropriate guardrails in place for enforcement so we do not contribute to persistent discrimination and exclusion of populations.

I'm sticking with it, Darrell. That's what I believe.

MR. WEST: Thank you. Thanks, Nicole. I appreciate that. And I love that example in terms of the problems with the criminal justice system and the AI applications. I think that is one of the most problematic areas just because of the impact that has on human lives and the capacity to destroy human lives in the process.

Mark, you pointed out problems related to geographic disparities.

MR. MURO: Yeah.

MR. WEST: What remedies do you propose and how do you think the Biden administration should handle this topic?

MR. MURO: Yeah. I think the starting point here is -- you know, partly given the focus of our work, is on the labor market. So in many respects, and for quite a while, the problems and challenges posed by AI are very much the problems that we have in our labor market systems. We are a country with suboptimal, quite weak education and training systems in many respects. And this -- these technologies are going to really challenge us to fix some of those problems.

So we've been talking about this for a decade without a huge amount of traction or improvement of some of the basics of our labor market system. But disruption and change are coming from these technologies. That means we need to be thinking about ways to mitigate that for workers, and to ensure tolerable outcomes. I mean, because this is a stressor of the labor market.

So one background, interesting thing for people and places, I think is a needed bias towards running a full employment economy becomes just that much more important given the kind of change that these technologies will enforce. Promoting a learning mindset across the entire economy is going to be critical. These are, happily, things we need to be doing anyway for social inclusion, economic inclusion, but they are also some of the best responses we can imagine to these technologies.

So thinking about rescaling, incumbent worker training, some kind of bias for -- to retain human work. Accelerated learning becomes important. You know, making skilled development financially accessible to everyone because these technologies will be constantly challenging humans to find their place.

And I always talk about a bias towards -- we need better research and understanding of the uniquely human contribution, not just now in automation, but in -- what is the human role in an AI production system. But other things, facilitating smoother adjustment. People are going to be moving around. It can't become -- you know, people are going to need to be changing what they are doing or dealing with cells of maybe interruption of work.

So how can we help people make these transitions better? And obviously we are going to go through a demonstration of some of the weaknesses of some of that system as we dig out from COVID and the labor market tries to normalize. Reducing hardships for workers who are struggling. We may need to rethink social insurance in some ways. In some ways some countries have seen technology growth coming and technology stress coming and has begun to think more about the support for workers so that they have a chance to adjust to technology.

And then finally, we're going to need specific regional interventions. I think we can't allow a pull away of -- you know, I see a bias towards the professional services portions of the economy pulling away with new productivity drivers here. We can't let this become an accentuation of the kind of superstar economy that I believe are tearing the country apart right now.

So I think we've talked about things like investments to help promising heartland metros develop through specializations and global economy. We need to make sure that research institutions and universities which are the source of a lot of these technologies and the people who can work on them are regionally distributed. We're going to have to think in other ways specifically towards trying to counter what could be disparity drivers in these technologies.

So just a few things. Plenty -- I think starting with the labor market reforms is a good place for the Biden people to begin.

MR. WEST: Yeah, that definitely makes a lot of sense. Having grown up in the Midwest, I appreciate your place-based approach and the need to invest in the heartland and start to reduce some of these geographic disparities. I think that would be a good plus if we could make some progress in those areas.

So Aaron, you talked about credit scores and financial services and some of the

problems associated with those areas. What are the remedies that you think would be most important? And any thought to what the new Biden administration should do?

MR. KLEIN: Yeah. You know, I both echo and push back a little bit on some of the comments Nicole made in the sense that there is a temptation to believe that financial services has it figured out because we have a system that works. By work I mean you apply for credit, you get an acceptance or a rejection, it's based on a legal structure largely put in place in the '70s that was meant at the time to be pro-consumer and meant to address equal access issues.

And we fail to appreciate that the basis of -- that the way that system has grown by grandfathering in and allowing one technology, this credit scoring system, and fundamentally putting high barriers to entry to new systems on the basis of measuring their discriminatory impact, hasn't allowed an okay system to thrive and flourish. I think the current system is horrifically allocating credit, providing tremendous benefits to the wealthy.

Darrell, you mentioned the point about the payment system. We have a deeply regressive payment system aided by this junk technology of plastic cards and swiped fees that drive up costs for tons of small businesses, benefit big-box retailers, derive billions in tax free rewards to the wealthy. Who checks the point guy? What do those rewards keep getting better and better? Well, because your card keeps taking more and more of the money. And who pays for that? People who have cash and debit. How can you afford cash and debit? Why are you using cash?

I once had this question. Someone said well, anybody can get a high rewards credit card. Really? Anybody who qualifies. You think they're handing out platinum cards to recent immigrants with no credit score? I don't think so. And why don't they have a credit score? They may have had one in their home country, but it doesn't translate here. So we have a pretty problematic system.

Darrell, you asked what when I do. Number one, I would put in federal legislation to overturn a case called AMEX v. Ohio. That's a case where the attorney generals of Ohio and many other states wanted to allow merchants to pass on the cost of payments and AMEX sued and said in our contract if you take one AMEX, you have to take them all. If you take any Visa card, then you have to take triple platinum sapphire. Those cards have different fees. I would overturn that Supreme Court

case, sided five to four by Justice Clarence Thomas. That helps drive some of this income inequality.

Two, I would immediately use the authority under the Expedited Funds Availability Act given to the Federal Reserve to make your funds available immediately. You know our second stimulus check, they sent the money on December 30 and you got it on January 4. Really? Six days for Uncle Sam to get you 600 bucks? Amazon could've gotten you anything in the world on your doorstep in 48 hours.

We have the legal authority to require that money to be made available immediately, but the Federal Reserve and other bank regulators sit on their hands and don't use that authority while the banks generate billions in overdraft fees. How many people paid a fee on January 1 because they didn't have emergency money in their fund?

Three, I would create some safe harbors and the Consumer Financial Protection Bureau has done some no action letters. But I would look at other things that allow new technologies in full light and transparency to allocate some credit and let's run some tests. I'm opposed to this idea of a sandbox. The hot language is a regulatory sandbox. Well, I don't know about you, but I have kids. And every time I put my kids in the sandbox, they would play. I would look away and talk to some other parents and it would either end up with my kid is crying or a puddle of wet sand.

What I think is a much better analogy here than letting some people play in a sandbox and see what happens, it let's have a greenhouse. Let's let a bunch -- let's create a controlled environment where something new and different can flourish that may not survive in the wild right away, with full transparency and sunlight and let's see if these products actually do a better job of allocating credit and let's lean in to this group that I mentioned earlier, the group where we can allocate credit in a way that is both more effective and less discriminatory to the status quo, fully understanding that it is not zero discrimination.

Look, if we allow income based lending, which I think is a good idea, the more money you have, the more income you have coming, the more likely you are probably to repay a loan, that's inherently discriminatory. Let's be honest with ourselves about that. And I think the Biden administration, if they prioritize this, can have a major impact both on day one and in the long run in reforming allocation

to credit and in stopping the reverse Robin Hood that is our payment system through the adoption of new technology.

DR. TAYLOR LEE: Darrell, do you mind if I just respond back to Aaron real quick? Just real quick.

So Aaron, I want to make sure that I clarified that I think we are on the same page on a lot of stuff, just so you know. Because I think we go back and forth. And I do actually think things like FICO score which are actually defined is one of the reasons why I'm trying to move towards this energy star rating. But I think we need to also be aware -- and Latanya Sweeney's research tells us that were the powers of algorithms is the click bait technology. Whereas if you are already on the wrong side of the wealth gap and you are already being offered 29% interest rate, credit card offerings, when you click that offering, even if you may not apply, that is what is going to be served back to you because simply the micro-targeting of the algorithm has determined where you fall in the credit risk assessment.

And we have to be honest that people of color have -- and I'm just going to push back on your greenhouse because I don't know a lot of people of color who have greenhouse access. But I will just tell you that we've got to think about ways to deconstruct, to your point, some of these mainstream assumptions about credit. And also, if you think about some of the new products and services that have come through on like Credit Karma, we have to also let people know that these are also based on algorithms and they're not necessarily set in stone and set in law to take you to a car shop to buy a vehicle.

So I think you're right. I think part of it is being very transparent about how these models and decisions get made, but also allowing people ways that they could also re-curate those decisions and to help companies that design these models really understand when they develop them, they are not -- no one is that an equal playing field when they actually put these models together. I think that's the broader assumption or mistake that when we think about how machines can transform the world, we actually think they can solve the problems and they can't. They actually just replicate them.

So I just had to put that out there. You know me. But I love to see -- I love what you're talking about in terms of really just coming up with ways that we extend that opportunity -- financial

opportunity to populations.

MR. WEST: All right. So now we have competing metaphors here, sandboxes versus greenhouses. I love that.

So Josh, you mentioned issues in terms of AI risk regulation, international standards, data governance, and data privacy issues. What remedies would you suggest to deal with some of these problems? And any thoughts on how the new administration may think about these issues?

MR. MELTZER: Thanks. I'm actually going to be -- this is an early preview that (inaudible)

MR. WEST: Actually Josh, I think we are losing your audio. Could you -- is there some way to turn your audio up so we can get you? If you --

MR. MELTZER: Can you hear me?

MR. WEST: Now we can hear you. I think if you are little closer to the microphone, it works.

MR. MELTZER: Great. Yeah. I do want to -- the three points I want to make is firstly just on China, which will be essential I think foreign policy focus for the Biden administration. You know, technology that is going to be essential to the discussion (inaudible) you know, being very focused on --

MR. WEST: Josh, if I could interrupt just for minute, there seems to be some audio issue. So maybe we can come back to you. If you can kind of find a different way -- and while you are figuring that out, Chris, I will just come to you.

I know you talked about autonomous weapon systems, the need for human control, problems of human rights. Your thoughts on possible remedies?

MR. MESEROLE: Yeah. So I will mention two I think possible remedies. So the first has to do with autonomous weapons and just military weapons systems themselves. And this is a bigger issue than just how they are used in the military. But one thing we haven't really talked about here is just what AI actually is right now.

And so there's kind of one version of AI in which computer scientists really try and say, okay, what do we know about the world and how can we put that in code to kind of replicate the

intelligence we have as human beings. That's called expert-based or rule-based AI systems that are very -- the advantage of them is that we have full control over those systems. We know how they fail, and we know how to test against them. The other version of AI is machine learning based AI where we say, okay, if we think about what intelligence is, it's very often about kind of pattern matching. And when we think about the human ability to recognize somebody in a photo or recognize a voice, what you're doing is you're matching visual or aural patterns in your brain and you're using the data and kind of computing from the data that you've gotten from human experience about how to do that.

And with modern AI, we're really kind of taking that machine learning approach where we are taking a lot of data, we're throwing it against a really advanced neural network or a deep neural network and we are coming up with a model that's really good at predicting patterns and data. And the challenge there is those models tend to outperform the expert-based or rule-based systems that we have, the tasks we care about most, whether that's kind of visual learning or reinforcement learning.

But the downside is we don't understand how and why they fail. The models are so complicated, they learn such intricate details from the data they are exposed to, and this kind of goes back even to some of what Nicole was talking about on facial recognition and other -- and Aaron was talking about with financial models. They fail in ways we don't understand, which makes them inherently dangerous in high risk situations.

And so the -- if I could kind of wave a magic wand and give the Biden administration something to focus on when it comes to military and AI, it would be investing very heavily in testing and evaluation and verification and validation of AI systems so we have a much better understanding of how machine learning based systems work and how -- and more importantly, how they fail. Because I don't think we are going to be at a point where we ever want a world where warheads are attached to AI systems that we don't know why and when they break down. And that's the case at the moment. It is not really mature enough to be used without a full understanding of the failure modes.

And in turn, the advantage, if we were able to make a breakthrough in that, we could then share it similar to what we did with regimes in the Cold War where we developed software for nuclear safety that we shared with our opponents because it was in the interest of everybody to have safer uses

of these weapons. So that's kind of the main thing that I would recommend.

In terms of human rights abuses of these technologies for domestic security purposes, right now we are in a situation where China has effectively turned Shandong in western China into the world's first kind of digital open air prison. You have basically, 10 million plus people, about a million of whom literally are imprisoned, but the rest of them are kind of subjects to mass digital surveillance. And if we care about that and want to push back on that, there is a variety of measures that the Biden administration could take to push the Chinese government to curtail its human rights abuses. And I think the challenge is to be selective and smart about how they do that.

I personally would advocate for something like increased export controls on advanced semiconductor manufacturing equipment. It's not the chips themselves, it's the machines that make the chips. China doesn't have those. They don't know how to make them. It's really hard to make a processor with a 5nm dye. It's not something that you can kind of just throw $1 billion at and a few smart people and get them to figure it out. It takes decades of experience and research and development. And we have a -- we and a few other democratic countries have a monopoly over that.

And I think we should explore ways of leveraging bottlenecks in the semiconductor manufacturing pipeline that China can't easily substitute for, to try and push and cajole them to curtail their human rights abuses with technology and importantly, also curtail their export and sales of the systems to other authoritarian regimes around the world. And hopefully the Biden administration can start to make some progress in that regard.

MR. WEST: Great. Josh, I will come back to you so you can finish your remarks if our -- if the audio is okay.

Actually, we are not hearing you now. Could you try one more time?

MR. MELTZER: Can you hear me now?

MR. WEST: Yeah. Now it's good.

MR. MELTZER: Okay. I'm not sure where you -- where I left off, but I'll just pick up where I think you lost me before. I was just briefly saying (inaudible).

MR. WEST: Actually Josh, we are still having some audio issues. I apologize for

interrupting you. So maybe I will just go directly to our questions. We have a number of questions that have come in from the audience. I'll just throw these out and whoever would like to answer them.

So Wayne has an interesting question on whether our panelists believe that AI is going to help strengthen democratic societies or is it going to basically, more likely and up supporting authoritarian systems.

DR. TAYLOR LEE: I can go. I can go. I can go. So I think Wayne, that's a really interesting question. And again, in my work here at Brookings, I'm really trying to contextualize the algorithmic economy and the algorithmic marketplace. So what we see with algorithms, which I think for the purpose of this panel we can't get too much into it. We also see their role that they play and sort of perpetuating hate speech or creating, what you are talking about, these regimes that are actually competing regimes which often times may be on the same platforms. So I always use the analogy of a -- Aaron used the analogy of the sandbox. I'm actually going to use the analogy of the playground.

There are tons of things on the playground, a slide, a monkey bars, swings, et cetera. When you think about the algorithmic economy, you think about things like algorithmic amplification. In many respects or many of the platforms that we -- that actually exist, some of us may be on the slide. Some of us may be in the sandbox. Some of us may be on the swings. But for those of us on the slides, may be White supremacist, those of us on the swings may be people from Black Lives Matter. Those of us on other parts of the playground may be lifestyle (inaudible). The challenge is that we are all on the same playground.

And when we think about the amplification of these messages and how they actually create this type of the strengthening of those ideologies, we could actually go either way on your question because algorithmic amplification is just so interesting to me. It actually does not have, in many respects -- and we wrote a paper about this a couple of years ago on social networks and the lack of having brokers that sort of break up these conversations especially when we see what we saw yesterday with the amplification of messages around hate or the amplification of messages around stolen elections actually led people to actually incite violence and storm government property. That is going to happen in this economy unless we actually get a handle on how we look at things like contact moderation or what

Chris said, the weaponization of AI for good or bad.

And so to your question, there is no answer because many of us are still trying to figure out the start and stops of this. But I do know if we stay with the black box thinking of how we unpack these algorithms, at least in my space, it's actually going to stifle us because there is no algorithm that is completely explainable especially when it leaves the lab and the way that it actually matures outside of that is really what is most dangerous.

MR. MURO: Darrell, I will just jump in really quickly and say that, absolutely we don't know, and the alternatives are still up to us to manage. But I'll also say that the dysfunction -- you know, the dystopian storyline would be that AI poorly managed, poorly developed, accentuates all of the winner take all dynamics that we have seen so rampant in the country in the last 15 to 20 years. There is an association of digital technologies with the possibility of disparities and accentuations.

So I think there is the worry that this technology could accentuate divides. And if that is -- you know, more of what we have had over the last 15 years in the United States with massive regional divides, massive divides between cities, but also between communities, that is the thing we need to avoid. So that's just my quick view.

MR. WEST: Okay. And Chris, you wanted to jump in?

MR. MESEROLE: Yeah. I think what I would add to -- the only thing I would add to what Nicole and Mark said is I think there's two core kind of challenges that every regime, whether they are Democratic or authoritarian, need to solve. One is identifying dissidents that might kind of seek to overthrow the regime.

The other is understanding what the preferences of the population that they govern are. In a democratic society they need to understand because they want to reflect it. In an authoritarian society you want to understand those preferences just so you can understand the political costs and begin to -- if you're the elite regime, you have a better sense of where the tipping point is before the mobs start to come out. And technology impacts -- AI in particular impacts both of those functions.

In a democracy there is generally fewer terrorists and there is generally fewer kind of violent dissidence because people have a say within the system. But even when there are, there are

ways of identifying it much more easily because of AI. And it's something that authoritarian regimes have used to great effect so far. Unfortunately, they don't use it with the proper guardrails in place.

The second function is something that is less discussed, but is really vital, which is that in a democratic society, the way that we understand what the populations preferences are is through elections. That's ultimately the way we do it. We do it through polls outside of election season. But ultimately, we have elections so that there is positive and negative feedback about what our government is doing and the population can communicate those preferences to their government.

In an authoritarian regime, they've never traditionally had that mechanism because they don't have elections. So the challenge for them has always been to understand how far is too far for them. And one thing we've seen historically over time again and again, whether it's from the Arab Spring or even go all the way back to the 19th or 18th century, are these tipping points where the people finally have enough and then they rise up and kind of overthrow an authoritarian regime. And the reason that happens is that the government loses its ability to understand popular preference.

What we see now is either -- actually, if you are an authoritarian regime, you can monitor popular preference if you have access to all the information flows going through your society. And Beijing and the Xi regime in particular have started to perfect this where they can kind of -- they have full control and full visibility into what's happening on Weibo and other platforms. They have a pretty good sense of what people are upset about, what they like, what they don't like, and it allows them to kind of ensure regime continuity that really wasn't available to authoritarian regimes in the past.

So that -- if I worry whether democracy can survive against authoritarian regimes in the future or vice versa or how -- what the tip of the balance will be when it comes to technology, I do -- I think that's what's new here with AI and big data in particular is it's less about the ability to track dissidents. It's probably more actually about the ability of authoritarian regimes to get a better understanding of their society and how far they can go before the society will start to push back.

And I think it's really problematic and I think it's something that the Biden team I think would be wise to think about. You know, what kind of policies and global governance we need for digital platforms to try and ensure that that's not something that authoritarian regimes can harness because it's

not going to stay in China. You know, I'm really worried about places like Turkey or other countries where they are kind of in the balance right now. And there are Western companies and western digital platforms in those areas that if Erdogan had full control or full access to Facebook data for example, what could he do with that to perpetuate the regime? That's something that I'm increasingly worried about.

So the question is a great one. I think the jury is still out, but it's one that should give us a little bit of cause and concern at the moment.

MR. WEST: Okay. Thank you, Chris.

We have a question from Grace. How could we apply artificial intelligence to the government? Not just the private sector. And are there any specific initiatives involving AI and government?

I can give a couple of examples and then happy to have other panelists respond.

One application that is becoming very common is AI for fraud detection where basically the AI combined with data analytics looks at a bunch of financial transactions, looks at contracts, and uses the AI to spot the outliers, to spot unusual activities and then refers them to human inspectors to see if there actually is fraud that is taking place. It's widely used in the private sector and now starting to be used in government because there is so much money that flows through various government agencies.

Worker productivity, the workforce of the future is going to be very different inside of government. Technology, including AI, is going to allow people to flatten organizations and basically use technology to monitor the frontline workers and then have less need for the mid-level supervisors.

And then the last area where it is starting to be deployed, and I have to admit I am not as confident that this is a constructive use, is using AI for employment purposes, using AI to screen resumes. AI can do things like that if you're basically talking about very formalistic criteria of, does the person have a B.A. if a B.A. is required for the job, and kind of use it to screen formal qualifications of that sort.

But some companies are basically suggesting AI can help determine some of the less tangible aspects, the personality aspects, the character aspects that people need in particular jobs. Those types of AI applications, I'm not confident are very effective or are very fair in terms of how they are

being deployed. But those are just a few examples.

Aaron, I don't know, if based on your experience, you having worked at the inside of government, have other ideas on how AI either is being or could be deployed within government.

MR. KLEIN: So I will give you a couple examples. We have an extremely complex and aggressive anti-money laundering system that was designed in the 1960s and '70s to go after the mob and tax cheats. Then it was bootstrapped in the '80s to go after the war on drugs. Then U-turn quick, 9/11, let's go try and track terrorist money. These are all very, very different forms of the illegal uses of the financial system that have now resulted in the complete insanity of potentially investigating the State of Colorado for depositing tax revenue in cash from marijuana businesses. We have thousands, millions of anti-money laundering reports telling Uncle Sam that there is a pot shop that you can find on Google maps all over this country.

And so if you think about it, one of these things the system fails us is we haven't prioritized what are the things we want to use the system for. And artificial intelligence can be very useful. If you sit down and you say, look, I want to look for human traffickers and these are the types of financial transactions and some of the geolocations that we're concerned about things from, you give us that You say I want to target financial elder abuse, one of the biggest growing crimes as we see in aging of our population. AI can go through a tremendous amount of information and flag for you more likely or less likely things and use this incredible wealth of information to proactively help law enforcement as opposed to putting a giant (inaudible).

We all pay higher fees in our daily life so every financial transaction can be screened for potential money laundering abuses. And then the current justice system, when they have a lead on someone, they go in quick and query the AML database. You know who is sitting in jail for anti-money laundering today? Denny Hastert, the former speaker of the House. His crime, pay -- probably potentially molesting children while he was a gym teacher, but technically his crime was using $9,999 transactions to structure to avoid bank limits. That's what he pled guilty to. Same thing for Mike the Situation from the Jersey Shore on tax avoidance. He is in jail for money laundering.

Is this really what our money laundering system is meant to catch? We are really after

The Situation from the Jersey Shore? We're sitting on a tremendous amount of data. There's a lot of crime and we are not using it effectively.

But I'm a little wary about this person's question about how can the government use artificial intelligence for good? I gave you one example, but the question that you know, I'm equally concerned what's going to happen is the AML finds tons of false positives, tons and tons of transactions that may not have broken the law, but people have an expectation and a right to privacy. And at what level do you want the government involved?

So there are ways AI can be used for good in the government, but I would want to proceed cautiously. And I'm not convinced that the government has the proper safeguards in place to go down this road.

MR. WEST: Okay. Well, thank you very much. We are out of time, but I want to thank Nicole, Mark, Aaron, Josh, and Chris for all of your contributions. Lots of terrific work going on all across Brookings as part of the AI and emerging technology initiative. Those of you who would like to read more, check out Brookings.edu. Thank you very much for tuning in. We appreciate your interest.

* * * * *

CERTIFICATE OF NOTARY PUBLIC

I, Carleton J. Anderson, III do hereby certify that the forgoing electronic file when originally transmitted was reduced to text at my direction; that said transcript is a true record of the proceedings therein referenced; that I am neither counsel for, related to, nor employed by any of the parties to the action in which these proceedings were taken; and, furthermore, that I am neither a relative or employee of any attorney or counsel employed by the parties hereto, nor financially or otherwise interested in the outcome of this action.

Carleton J. Anderson, III

(Signature and Seal on File)

Notary Public in and for the Commonwealth of Virginia

Commission No. 351998

Expires: November 30, 2024