

THE BROOKINGS INSTITUTION

WEBINAR

POLICYMAKING AND ARTIFICIAL INTELLIGENCE
A CONVERSATION WITH JOHN R. ALLEN AND DARRELL M. WEST

Washington, D.C.

Wednesday, April 21, 2021

Opening Remarks:

STEPHANIE AARONSON
Vice President and Director, Economic Studies
The Brookings Institution

Moderator:

SANJAY PATNAIK
Senior Fellow and Director , Center on Regulation and Markets
The Brookings Institution

Conversation:

JOHN R. ALLEN
President
The Brookings Institution

DARRELL M. WEST
Vice President and Director, Governance Studies
The Brookings Institution

* * * * *

P R O C E E D I N G S

MS. AARONSON: Hello, everyone. And welcome to this fireside chat hosted by the Center on Regulation and Markets at the Brookings Institution. My name is Stephanie Aaronson, and I'm the vice president and director of the Economic Studies Program here at Brookings.

This event is one of an ongoing series of fireside chats hosted by the Center which explore important topics related to modern day markets and regulations through one-on-one conversations with regulators, business executives, and academics. Today, we are honored to be joined by John Allen, president of the Brookings Institution, and my colleague, Darrell West, the vice president and director of Governance Studies at Brookings. They're joining us to talk about their new book, "Turning Point: Policymaking in the Era of Artificial Intelligence."

Artificial intelligence is transforming our economy, society, and our politics so rapidly that it is sometimes hard to keep up. And as the authors point out, the decisions we make today about how to regulate AI and the policies we develop to deal with its impact will largely determine whether AI is a force for good in our world or not.

To share a little more about the authors, John Allen became the president of Brookings in November 2017, having most recently served as chair of security and strategy and a distinguished fellow in the Foreign Policy Program at Brookings. Allen is a retired U.S. Marine Corps four-star general and former commander of the NATO International Security Assistance Force and U.S. forces in Afghanistan. During his nearly four-decade military career, Allen served in a variety of command and staff positions in the marine corps and the joint force.

Darrell West was appointed as vice president and director of Governance Studies in 2008, where he also holds the Douglas Dillon Chair. His current research focuses on artificial intelligence, robotics, and the future of work and he has published multiple books with the Brookings Institution Press on these topics. Prior to coming to Brookings, Darrell was the John Hazan White professor of political science and public policy and director of the Taubman Center for Public Policy at Brown University.

And now, I'm pleased to introduce the moderator of this conversation, Sanjay Patnaik. Sanjay is the director of the Center on Regulation and Markets and holds the Bernard L. Schwartz Chair

in Economic Policy Development in Economic Studies. Before joining us last summer, Sanjay was on the faculty at the George Washington University.

Sanjay deserves credit for having originated these fireside chats upon his arrival at Brookings. They arose out of his longstanding interest in research into business and government relations, corporate political strategy, globalization, and international business. So, without further ado, I will turn things over to Sanjay. Thank you so much for joining us today.

MR. PATNAIK: Thank you so much, Stephanie, and thank you to John and Darrell for being here today. It is a pleasure. I would like to start with Darrell. Can you tell us a little bit about why you wrote the book? What inspired you to focus on that topic?

MR. WEST: Well, first of all, thank you, Stephanie, and thank you, Sanjay. We appreciate your interest in AI and look forward to working with you and your team on technology issues. We see lots of opportunities for collaboration between us.

So, John and I wrote this book because we think AI is the transformative technology of our time. As Stephanie mentioned, it's being deployed in many different sectors from healthcare and education to transportation and e-commerce and national defense. So, in the book, we present in-depth case studies of AI in each of those areas. We look at how it's being used, what the opportunities are, and what the risks are.

We developed the title of "Turning Point" because we argue we are at a major inflection point and the crucial variable in determining our future is going to be public policy. So, in the book, we present a detailed policy governance blueprint. And we suggest that if we take appropriate actions, we're very confident about the future of our society. We're very optimistic about how things may unfold. But if we don't do certain things, it could go off the rails very quickly. So, as a society, I think our goal should be to build responsible AI that serves the public good.

MR. PATNAIK: Great. Thank you so much. I would like to turn to John next. You have made AI a priority area for the Brookings Institution. Can you please elaborate more on why you think it is essential for a policy-oriented thinktank like Brookings to look at AI? And what is your vision for Brookings in how it can contribute to a better understanding of AI?

MR. ALLEN: Well, let me also add to Darrell, my thanks to Stephanie for her opening

remarks and my thanks to you, Sanjay, for convening this conversation. And I want to thank you, specifically, for the work that you've done on artificial intelligence across your career, but, specifically, since you've arrived as the Center director.

As Darrell said, artificial intelligence is going to be one of the most transformative technologies of the 21st Century. When I arrived as president, it was very clear to me that the Institution needs to acknowledge the potential of artificial intelligence and other emerging technologies the potential for the change that it can create in our society going forward.

And as Stephanie said, as Darrell alluded to, and I think we talk about it pretty frequently, we face an environment in the 21st century with these emerging technologies that can either be utopia or dystopia. And, of course, we won't hit either of those two extremes. But we will certainly find that these technologies, which are inherently neutral, will ultimately take on a particular form by whichever entity employs them or whatever country employs them.

And so, the point that we've attempted to make is as AI and other emerging technologies become more pervasive in the human environment, we have to think about the kinds of legislative or regulatory considerations that are necessary with respect to policy, etc., -- in particular, policy, -- on how we can govern these technologies.

This is really important for us. As I think you would certainly know, Sanjay and others, technologies in the past, have been often -- they've often led an environment ahead of the formulation of policy by as much as, you know, 7 to 10 years. And as these technologies continue to develop, it's really important for us to be thinking about what policies are going to be essential, in particular, a liberal democracy. What policies are going to be essential ultimately to protect our citizenry. And to ensure that these technologies are employed and wielded to the maximum benefit of our population and our country and our allies.

Because in the end, AI will touch virtually every aspect of who we are as humans. And will define everything from a simple matter of how our students do in classrooms. On the one end, that's not a simple matter, certainly. But to the other end, where it will shape geopolitics between great powers. And the more we think about the policy framework for those technologies now, the better we will be prepared ultimately as these technologies become fielded in the world of the human experience.

MR. PATNAIK: Great. I think that's certainly really important to kind of try to provide a framework before a lot of these technologies hit the market.

Turning to Darrell, if you look at artificial intelligence, it's kind of like a buzz phrase in the popular discourse, but the meaning is often not well-defined. So, in Chapter 1 of your book, you talk about three features of AI that have distinguished it from other technology, which is intentionality, intelligence, and adaptability. Can you please talk through each of these features and give a few examples of AI systems that are already being widely used in different industries?

MR. WEST: Sanjay, that's a great question, and I think you're exactly right. You know, everybody has heard the term AI, but few people actually know what it means. So, we do open the book with a very lengthy discussion. And what we do is define AI as automated software that can analyze data, text, or images, and then make intelligent decisions based on that information.

I think the key qualities are software that is intelligent, adaptable, and has the ability to learn as it gets more information. And by that, we mean that algorithms can independently analyze data and then act on the insights that come out of that. And I think the learning part is key because an algorithm should be able to improve as it gets more information. I mean, that's the whole point of AI, really. We want AI that's able to adapt to changing data, as well as changing circumstances, and still can make intelligent decisions based on that.

To give just a few examples to show the multifaceted nature of AI, we're seeing applications in transportation in the form of autonomous vehicles. The AI is essential to autonomous vehicles because it's the software that integrates information from the LiDAR on top of that autonomous vehicle, the cameras that go around the vehicle, and the remote sensors that help to keep the vehicle in the correct lane. So, the autonomous vehicle is nothing without AI that can instantly integrate, analyze, and assess the situation and keep people safe.

In the finance area, we're seeing many applications. In fact, we think finance is going to be one of the leading sectors in terms of new AI applications. People are using AI for fraud detection.

MR. PATNAIK: Mm-hmm.

MR. WEST: You can use AI to find outliers and abnormalities and then refer them to humans for more in-depth inspection to see if fraud actually is taking place. AI is being used for wealth

management. It turns out that AI is more rational and less emotional than humans, which turn out to be great qualities in terms of managing your money.

And then the last example I will give comes from the world of e-commerce and the case of product recommendations. Because a lot of e-commerce platforms are analyzing your online behavior, looking at your consumer purchases, and looking at the things you look at on the site, but don't actually buy. And then based on the analysis, that site will recommend products to you that you may want before you actually realize consciously that you want them. There's some companies that have claimed that up to one third of their sales are now coming from their own product recommendations.

So, that's just one sign that the AI is getting better and better and it's being incorporated in commerce and trade.

MR. PATNAIK: The last projects are really interesting. I don't think a lot of people think about that when they see the recommendations pop up that there's such a sophisticated algorithm behind it.

John, I'd like to turn to you. So, you mentioned something really important, which is that the technology itself is neutral, right? But it's really up to us how we use it. So, if you look at AI, what are some of the most exciting opportunities that you see with AI technology?

MR. ALLEN: Well, Darrell's hit many of them. I think the opportunities are really limitless in many respects. Medical research, for example. Natural language processing has the capacity to ingest decades and decades of what had been basically printed reports and provide the capacity for machine learning through natural language processing to do an enormous amount of analysis on medical research, for example. And just pick a particular area, Alzheimer's, for example. I think that we'll see that the support of artificial intelligence to Alzheimer's research is going to make a profound difference in the presence of that disease in our population in our lifetime. So, medical research.

Education. Artificial intelligence in the context of education of our children, in our young adults, changes the environment from one of teaching necessarily to one of learning. And when you add artificial intelligence to something like, you know, the simulated environments in which we can control the environment for the students and have immediate feedback because of the performance of the children, that gives us the capacity to have a far greater and better tailored system of education for our children

through virtual reality or augmented reality. So, education is a very important area. Darrell talked about e-commerce. Also, interstate commerce with respect to driverless vehicles.

So, so much of what will touch the individual lives of Americans and global citizens every single day, will be affected in some form or another by artificial intelligence. So, I'm cautious that these neutral technologies in the hands of liberal democracies have the opportunity to truly improve the quality of life of the individuals of our populations. But I'm also extremely attentive to how illiberal states and authoritarian states and, of course, totalitarian states, can use this extremely powerful technology for the purposes of highly intrusive surveillance and ultimately control and oppression.

So, it goes back to the utopia versus dystopia. And it's important for us to take care of ourselves, but also to be able to defend ourselves and also, to understand what our potential opponents, adversaries, and even enemies are doing in the wielding of these technologies.

MR. PATNAIK: Actually, let me follow-up on this because I think this is a really critical aspect of AI. As you say, it has the potential to improve a lot of people's lives. But if we look at authoritarian regimes around the world, AI systems are already being used for mass surveillance, for social credit scoring in China, for instance. And they have really the potential to undermine democratic societies through the proliferation of deep fake videos, of misinformation, things like that. How can we safeguard our democracies against such negative applications?

MR. ALLEN: Well, a regulatory environment permits us to protect ourselves from ourselves in many respects. But we also have to recognize what the threats are that will come at us from our opponents. You know, we live in the physical world. Right now, I'm in my home in Mount Vernon, Virginia. But you and I are communicating with Darrell in the cyberworld. In the information space, if you will. And that shapes and frames our moment-to-moment existence. And when that environment is, in fact, intruded upon by our opponents, either by distorting a reality through deep fakes, or through as Darrell just talked about, through the capacity for micro processing and micro analysis of your preferences in putting certain news items in front of you or putting skewed or biased information in front of you based on whatever social platform you're on, then we are vulnerable.

And we have to protect ourselves, again, through a policy regime, through regulations, and legislation. But we also have to defend ourselves in the context of our military capabilities in the

cyber domain. And we have just seen the President of the United States level sanctions on the Russians for intruding into our cyberspace and attempting to influence the opinions of American citizens as it relates to the 2016 election and later.

So, that is a vulnerability that liberal democracies have and we have to be conscious of that. It also requires that we educate ourselves to what true information actually looks like and not permit ourselves to fall victim to the kinds of skewed information that can come over certain media platforms. And I'll let Darrell also mention -- question this as well, because he's got some very good thoughts on it.

MR. PATNAIK: I think that's a really important point. Really the responsibility of each city's laws is to inform themselves, right?

MR. ALLEN: That's correct. Exactly right.

MR. PATNAIK: Turning to Darrell. Like these are some of the risks that we heard about, right, about AI, and looking at our democracy and cyber intrusions. What are some of the other biggest risks that you see with AI technology going forward?

MR. WEST: I certainly agree with all the worries John just expressed in terms of the use of AI by authoritarian governments to suppress their people, the risk of misinformation, the risk of disinformation. But within the domestic space, there are some additional concerns related to fairness and bias. And certainly, my colleague, Nicol Turner Lee has written extensively on these issues.

So, for example, financial institutions are making decisions on loans and mortgages based on AI applications. And so, the question is, are those algorithms operating in a fair manner? Are they equitable? Are there inequities related to race, gender, or marital status? In the bricks and mortar world, we have clear rules on what banks can consider and what they are not allowed to consider. We need to make sure that the algorithms operate in a fair and unbiased manner.

Lack of transparency is a big problem in the AI area. You know, we refer to this as the black box. Like nobody really understands how those algorithms operate. How do they make decisions? What are the data that they incorporate in their decisions? As AI becomes more prevalent in education, healthcare, finance, and other areas, we need greater transparency about how they are making those decisions so that we can then assess their fairness of their bias and their impact on humans.

The impact on human safety, clearly, is a big consideration. As autonomous vehicles

start to come online, this is becoming a big issue. This week, we actually just had a Tesla crash where the car was being operated via autopilot that crashed. And so, questions of safety are going to be paramount.

And then the last thing I'll mention is a basic governance question of who decides. The old model in the technology area was what we loved to refer to as permissionless innovation. Whereby, we essentially let private companies make the decisions. They would decide what the new products would be, how to innovate, how these products got deployed, to whom they were sold, etc., etc.

People are now no longer comfortable with that approach. They want more public engagement, more public involvement, more public oversight. We can do this through policies, laws, regulations, and/or ethical standards. There is a growing techlash out there. Kind of a backlash against the technology sector. And that is likely to move the public, as well as policy makers in the direction of stronger enforcement and tougher policies.

MR. PATNAIK: Yeah, I think that the last one is really interesting because I think the market policy's company interest has grown so much over the last couple of years that a lot of people are concerned about it. And I think the transparency issue is really highlighting a vulnerability of our regulatory system. Because oftentimes, the regulators don't even understand what the AI systems do and that black box of these algorithms. And it's very difficult to decide policies if you don't really understand technology. And I think that's just going to get worse. So, we have to find new ways to regulate.

And that's actually, John, that's my next question for you. So, you talk to policy makers on a regular basis, right? And so, what has been the response by the policy maker community to your book and issues regarding AI? Are policy makers thinking about that? Do they take it seriously in D.C. or at the state level?

MR. ALLEN: Oh, yes, they're taking it very seriously, Sanjay. Because it gets back to the heart of the issue in many ways. And let me just use an acronym that we hear frequently. We hear it in the private sector. We hear it amongst the policy participants, and it's -- the acronym is ELSi. E-L-S -- all of them capitalized, little i. And in the liberal democracy, in the U.S. democracy, more and more, the private sector in harmony with the policy environment, uses this acronym as a guide.

And it's a very in-depth guide in the big scheme of things. But at the most basic level, the E stands for what are the ethical implications of this particular application of artificial intelligence? The L is what are the legal implications? And Darrell just mentioned the crash that occurred with the Tesla vehicle. There will be substantial legal analysis of how that happened and the outcomes in terms of jurisdictions and traffic regulations, etc., etc.

So, the ethical, the legal. And then the S is for the societal implications of a particular application. And the little i is for implications. So, I think that we in this country and more broadly, in the community of democracies, are far more attuned now to the implications of artificial intelligence in the day-to-day lives of our citizens. Where 10 years ago, by virtue of the strength of computing and the availability of big data analytics, AI was a notion, but AI wasn't nearly as pervasive.

Today, with the improvement in the capacity for super computing, with the availability -- increased availability of big data and the sophistication of algorithms, this is far greater a concern as we go forward. And so, I think that we -- this book is of great interest. I've heard a number of folks -- not trying to pat Darrell on the back, since he did so much work on it, but I should -- there are so many folks who've said I truly now understand by reading this book that artificial intelligence left unto itself, is a potential problem.

So, we need to have a regulatory environment or a policy environment that helps to shape how these applications actually benefit humankind. And let me just make -- read one thing to you. You know, Darrell and I wanted to make sure that we were very clear that artificial intelligence we see as an asset. And emerging technologies are on the horizon to the good of all humankind. So, when we dedicated this book, we dedicated it as follows: This book is dedicated to our youth into whose hands we place the full potential of artificial intelligence and other emerging technologies. It is our most fervent prayer that they are guided by the light of good in wielding these technologies for the benefit of all humanity.

Being able to wield these technologies to the good and the benefit of all humanity, will come not just from the technologies themselves, but our capacity to shape their employment and their application through laws and regulations and policies to ensure that those technologies are, in fact, applied to the good and the benefit of all humanity.

MR. PATNAIK: And I think that will be critical also in terms of the acceptance of these technologies by the population, right, if we look at upon those vehicles, for instance.

And that actually brings me to my next question for Darrell. If we look at the U.S. government currently, do you think the U.S. government is prepared to adequately regulate AI? And what kind of federal agencies do you think should regulate AI? Should it be one agency? And what kind of improvement or changes should agencies make in order to have that capacity?

MR. WEST: Well, the short answer to your first question is, no. Most government agencies are not well equipped to regulate AI at this point. They have big problems in terms of lack of proper staffing and lack of proper expertise to regulate AI. Industry, obviously, can pay far higher salaries. I mean, good AI people can earn \$1 million a year from the private sector.

So, it does make it very challenging for the federal government to get qualified people who understand AI, understand the way that it operates, can develop data that illustrates some of the risks and some of the harms, and then develop regulations that can deal with those issues. It's really hard for the government to compete for technology talent and build the staffing required for effective regulation.

On the structural issue, I think there are interesting questions in terms of how the federal government approaches the topic of AI regulation. I think in the short run, what we have now and what we're likely to retain at least in the immediate future, is a sector-based approach to regulation where the Department of Transportation regulates AI in autonomous vehicles. The Federal Trade Commission handles consumer harms. The Equal Employment Opportunity Commission handles employment-related problems of AI. The Department of Education handles education-related AI. And the Department of Health and Human Services addresses health IT issues.

That's the status quo and probably will continue to be the case for a number of years. But that approach is problematic because there are some AI issues that cut across sectors. Like clearly, the privacy topic is an issue whether you're talking health, education, transportation, employment, or basic consumer finance. So, you know, we need a national privacy law that really protects people regardless of where the particular application is.

On the issue of whether we need a new agency, my colleague, Tom Wheeler, has

proposed the creation of what he calls a digital regulatory agency, which would handle these crosscutting issues and develop a well-trained staff that can regulate AI regardless of the particular sector. So, that's something that we're actively debating and trying to determine, you know, do we need that? What would it look like? How would it operate? And so on. But, you know, that approach is likely to take several years to stand up, especially in our current political climate with a 50-50 setup.

MR. PATNAIK: Yeah, definitely. And I think, one policy is quite important that you mention. If it is sectorally-based, I think it will be very difficult for the government to get the capacity and the skill level if it's distributed across the different agencies, right? So, if you look at AI governance, what should be the highest priority for policy makers right now if they look at it to improve AI governance? In your book, you mentioned as a suggestion improving governance through distributed collaboration. Can you explain a little bit what that means?

MR. WEST: What we suggest in the book is that in addition to needing technology innovation, we need governance innovation. We need new models of how we think about developing guardrails that will promote human values and protect humans from some of the safety risks, the fairness and bias questions, and the lack of transparency.

So, we need some approach to regulation that encourages innovation because that is going to be a big driver of our economy and international commerce, while mitigating the clear consumer harms that we want to guard against. So, in the book as you mentioned, we discuss the notion that we call distributed collaboration is one way to think about that. That involves providing vehicles for experts in government and industry to come together to develop effective guardrails.

Now, this can take place through the use of technical standards. That actually has been a common way to handle technical issues in the past. So, for example, when we were developing mobile technology, there were different companies that had different approaches to handling cellular communications. And so, what the government did was bring the experts together, develop standards for cell phones to be able to communicate with one another regardless of the company that was involved. Make sure that the communications took place in a safe and reliable manner. And then all the products were built based on those standards.

So, that is one way to kind of think about dealing with some of these issues. NIST, which

is the National Institute for Standards and Technology, is a current federal agency that is working very hard in this area. It's been tasked by Congress to develop standards for AI that can mitigate some of these issues. We've been involved in some of their discussions where they brought people together. They've been having meetings and their workshops that try and bring together the expertise of people involved in government, business, and academia to figure out what are the best ways to improve privacy? To make sure we have secure products and services. How to we guarantee human safety with these new AI applications?

So, those are just a few ways. And the book goes into greater details on these things. But government agencies need to modernize their processes and update the way they think about regulation itself.

MR. PATNAIK: I think you bring up a really important point, which is that if we look historically at the two worlds that say Silicon Valley with the developments and technology, and then the policy making community here in D.C., I think those two worlds have historically been quite separate. And I think for AI, we really need to have a much better open communication between these two worlds.

MR. ALLEN: Can we --

MR. PATNAIK: Yeah.

MR. ALLEN: Can we, please, come in behind that, Sanjay, because you make an important point. And Darrell and I have made this point publicly and it's in the book. And that is for many years until we began to see active engagement by policy makers and the regulatory considerations, many of the regulations associated with the application of artificial intelligence were baked into the code far upstream. They were baked into the code in Silicon Valley. And the policy makers in Washington really had no idea what the implications were of a particular algorithm either for good or bad.

And one of the purposes of this book was, in fact, to help to educate policy makers that these algorithms and the code associated with the expected outcomes can't be left solely in the hands of those who are code writers. This has to be an effort where government is involved. Government involved to the extent that the regulatory process, the policies that are formulated, enhance innovation, and improve the quality of human life, but don't at the same time, don't strangle innovation. And that's a really important and very delicate balance. And this is why government has got to be involved in this to

understand that we have a lot to do on innovation, but we can strangle that if we're not careful with the wrong kinds of policy or regulatory environment.

MR. PATNAIK: And I think that's a very difficult balance to strike for regulators. How to make sure that we have that innovation without strangling it, while safeguarding the customers at the same time.

MR. ALLEN: Right.

MR. PATNAIK: I want to talk to, John, to you to an area that is, obviously, very closely related to your longstanding career in the U.S. military, which is how AI relates to warfare. In your book, you mention the concept of hyper war where human decision making is almost entirely removed from time sensitive military actions. So, can you please talk a little bit about this issue in more detail? And especially what I'm wondering is should we actually remove human decision making completely from military actions at all? And how can we safeguard against AI systems making the wrong decisions?

I just read recently that the Russians are now thinking of embedding fully autonomous AI systems in their tank units. And I'm curious to hear your thoughts on this.

MR. ALLEN: Sure. Well, I don't advocate removing the human from the environment. In fact, there's a, you know, a very vigorous debate going on in the liberal democracies and in the United States as well, on where is the human in this process? We call it where is the human in the loop? Or where is the human on the loop? And I'll come back to that in a minute. From time and memorial, the side that can apply military force with greater precision and with greater velocity, is typically the side that will win. It doesn't make any difference whether we're in chariots, tanks, or the future systems that we may well see that rely on artificial intelligence and hyper velocity.

But those two truisms are at work no matter what the systems are. And that is the side that can move more quickly with greater precision is the side that will frequently and usually prevail. And if you look at the relationship, the Clausewitzian relationship of technology to the human over time, we talk about the character of war, which is the technical, the technological dimension of war, the equipment. And we talk about the nature of war. The human involvement in the process. And humans being what they are, humans have typically remained largely the same and unchanged as in the involvement in this business of conflict and war. But the technology of war, the character of war has changed dramatically.

And if I were to graph this out so we have the Y-axis as the vertical and the X-axis as the horizontal. And the X-axis is time and the Y-axis is the improvement or the increase of technology over time, for much of human history, the line of technology was virtually parallel and almost asymptotic to the X-axis. There was very little change. There was very little improvement in the nature of the technology of war. But in the aftermath of the second world war and the advent of computing, and as we began to see computing become stronger, and as command and control systems became more comprehensive, we began to see this line begin to lift up off the X-axis.

And today, the technology is not just changing, Sanjay, the rate of change of technology is nearly mindboggling. And so, we have a real dilemma today where we have technology that is changing, but the human in many respects is the same individual that was involved throughout the entire history of warfare. So, the challenge that we face is not just that we have changing technologies. And in many respects those technologies rely on artificial intelligence. The challenge that we face is how do we equip the human to be part of this process?

And so, the idea of hyper war isn't just about artificial intelligence. The idea of hyper war is that artificial intelligence and other emerging technologies both hasten the capacity to wage war and give us greater precision in ways we have never imagined before. And so, as we think about that what that means is we are starting to move far faster than we ever have before.

And as we think about how fast our opponents might be moving, the question for us, the difficult question for us is where is going to be the human in the loop? Because I think we know that Russians are not going to be constrained to put a human in the loop with fully autonomous battlefield systems. And right now, we are not producing fully autonomous battlefield systems because we want to have a human in the loop.

But if a human is in the loop, inherently it slows the process slightly. So, when you're competing against an opponent who is employing fully autonomous lethal systems in the battle space, whether they're swarms of drones that have computer vision and facial recognition that can take out key individuals at a headquarters. Or an attack of drones like that that might come at a sporting event where, again, with computer vision and facial recognition, individuals, VIPs, etc., might be taken out.

When we're competing against opponents like those, we have a real challenge. And so,

we're dealing in our own conversation about where the human should be in the loop and whether we should ever produce a fully autonomous lethal weapon system and release it into the battle space. For now, we are not. But we have to face the reality that our opponents will not feel so constrained as we are on a humanitarian basis. There's a lot more I can say about this.

But let me just make this other point. There's a lot of AI though that is being applied in the military. AI that improves our capacity for intelligence collection and intelligence analysis. And does it in a way that produces analysis so much faster than it did when you had rooms full of analysts looking at computer screens or imagery. And provides that information far more quickly to decision makers.

We have artificial intelligence being applied to maintenance. Where AI, the capacity for algorithms in certain systems with substantial dynamic components in those systems, can detect the earliest anomaly in the operations of those systems so that we can affect maintenance far earlier in a process to prevent the catastrophic failure of dynamic systems. That means that we're more efficient. That means we're safer. That means we're more effective. Our systems will be up and operating longer because the kinds of maintenance that we can do gets out ahead of the potential for catastrophic failure of dynamic systems.

And then there's artificial intelligence on wearable systems. So that we have a very good feel for the state of the health of the individuals within a particular unit that may be engaged in combat. To include being able to track them to see who's wounded or who is experiencing enormous duress. Being able to track those who may have been wounded or killed in combat so that we know if we need to surge a certain kind of medical capability forward because of the nature of the wounds. Or push forward certain kinds of people as replacements. So, there are real applications for artificial intelligence in our military forces that don't require us to have to debate the lethal dimension of this.

And I'll make one final point. And that is, it's one that we're very attentive to in the U.S. and amongst our allies. And that is artificial intelligence and its presence in nuclear command and control. For those of you who haven't seen it, I would suggest you see the movie, "Fail Safe," which is, in essence, a futuristic movie that talked about an artificially intelligent system that ultimately would govern with the presence of a human in the loop, an artificially intelligent system that would govern American nuclear command and control, and it fails.

And so, the whole issue associated with how and where the human is in nuclear command and control is something we have to keep a very close eye on. And this is an area where we should be talking to the Chinese and the Russians constantly to ensure that we're very clear about the fact that we will probably never put artificial intelligence in nuclear command and control systems. But if we ever saw it going in with the other side, it's going to create a real decision dilemma for us.

MR. PATNAIK: Yeah, I mean, that's a pretty terrifying thought if you think about it, right?

MR. ALLEN: Yes.

MR. PATNAIK: And I want to follow-up on that, right? So, you said, obviously, other countries don't have the same problems with fully autonomous systems and their development in the Russians and the Chinese. So, how should the West and how should the U.S., with its allies, respond to that? Is there a way to kind of like neutralize that threat without going the same way?

MR. ALLEN: Well, we should be talking to them right now. And in truth, you know, we can't go into a lot of detail here, but in truth, we are having these conversations now. And I think with at least one of the great powers out there, there is no desire for us to engage in an artificial intelligence arms race, if you will.

MR. PATNAIK: Okay, that's good.

MR. ALLEN: In the 20th century, of course, the Cold War was defined in many respects by a thermonuclear arms race. And it was the basis, in many respects, for the competition that we saw during the Cold War. Those days are over. But we could easily see that artificial intelligence, nanotechnology, biotechnologies, all supported by supercomputing, quantum computing as time goes on, we could see an AI arms race, which could be quite destructive -- quite disruptive in the big scheme of things.

And before all the relevant sides who are sufficiently sophisticated technologically, before all these sides begin to build major weapon systems, hypersonic systems that are artificially intelligence that can fly for days or even weeks at a time before it decides to ultimately and unilaterally target a particular location, we need to start to have the conversations about how we can, as we have done with other weapons systems, determine that we will limit the production of those systems, the location of those systems. Or even if we ever do produce them, geofencing, putting certain targets completely off limit to

autonomous systems, etc.

I mean, there's a lot that we have to talk about now. This is the time to have the conversation because even though the Russians intend to have autonomous armored systems in the battle space, it's still pretty early in all of our countries' production of these systems. Now is the time to have these conversations about disciplining ourselves as a community of nations, even though we don't agree. They're authoritarian or totalitarian. Even though we don't agree with the system of government, we can certainly agree that we don't want the outcome and here we should be talking. And we should be talking very aggressively. And Darrell and I have written about this.

MR. PATNAIK: I mean, it's in no one's interest to call out that arms race, right?

MR. ALLEN: That's right.

MR. PATNAIK: Darrell, I want to change views a little bit. And so, if we look our allies, especially the European Union, we see that they have been thinking about AI policy, governance, and regulation for quite some time. And they actually just came out with plans, with quite detailed plans to strictly regulate AI systems, including banning AI for mass surveillance, for social credit systems, like it's being used in China. And so, I think that that sets quite important standards for western democracies. Do you think that the U.S. should follow with similar regulations and coordinate with the EU in coming up with the governance framework? Or is the U.S. going its own way?

MR. WEST: I mean, I think we do need to coordinate with the European Union on the future of AI. And you mentioned the new guidelines that they have just come out with. And I would really encourage people to take a close look at them because it's a very important step forward in terms of how we can approach AI, what the policies should be, and what the nature of the regulations should be.

And what I find interesting about the EU approach is it adopts what it calls a risk-based approach to AI. Meaning that the regulations are geared to the degree of risk. With risk defined as the number of consumers affected, the size of the company, the scope of the impact. Like does it affect human safety? That requires a very high level of regulation.

The EU designates law enforcement and criminal justice as a high-risk area just because it can lead to the arrest or incarceration of individuals. In fact, in the United States, there actually are some cities and some law enforcement agencies that are using predictive analytics to target particular

neighborhoods as high crime areas, transfer of resources to those areas. But, of course, all of those predictive analytics are based on historic data that are completely biased based on race. And so, the predictions are biased in a racial direction.

So, I think that risk-based approach is a good way to think about it. And some of their particular recommendations, for example, they're banning the use of facial recognition software until the systems become more accurate. There actually are a few American localities that have taken that step as well, particularly in terms of use of facial recognition by law enforcement.

The EU guidelines proposed fines for companies that violate their rules. The fines can go up to 6 percent of global sales. So, basically, those could lead to very stiff fines if your company violates the guidelines that have been put into place. The guidelines pick up on a recommendation that we make in our book that came out last year. They called for a risk assessment. In our book, we called for AI impact statements. Kind of that parallel the concept of an environmental impact statement. So, they suggest that companies need to undertake risk assessment of their AI applications. And they need to provide what they call proof of safety of those AI applications.

Although the one thing I don't understand and I would like to have much greater clarity about what our European colleagues have in mind, is what does that mean in practice? Like how do companies actually provide proof of safety of an AI application? So, I think we need more details to assess that.

And then a final recommendation they make is that companies have to explain how their AI operates. So, basically, they want much greater transparency in the use and application and the operational aspects of AI. So, whether the U.S. should move in all those directions, I think, you know, we have to kind of wait and see. But I think given the challenges that John mentioned in regard to China and Russia on the international sphere, it's very important for the U.S. and the European Union to stay together. That we kind of think about what our shared interests are in confronting authoritarian nations, and to make sure that our AI regulations are coordinated in particular ways.

MR. ALLEN: Let me come in as well, Sanjay. And I was recently asked a question along these lines about whether there is a quadrangle, if you will. Russia and China as two of the points, and the United States and the EU as the other two points. And to Darrell's comment, look, the Euro-Atlantic

region, North America, and Europe, share so much in the context of our common values, and our common interests, and our common systems of government, and our common systems of our economies, and how we value the individuals within our populations. We share so much.

And for us to find ourselves at odds on these kinds of things, these kinds of regulatory standards, we should look very, very carefully at why those differences occur and try to find ways where we can come closer together. Now, we will always -- countries always act in their own self-interest. So, we will never be in 100 percent coincidence necessarily with the European Union. But I would propose that we should work very hard if the proposition that there is a quadrangle out there on issues associated with the application of these technologies, we ought to work very hard to turn that into a triangle so that the United States and the EU are almost imperceptibly different in our commitment to human rights, our commitment to the rights of our citizens, and the fair, explainable, safe applications of these technologies.

MR. PATNAIK: I totally agree. I think it would really increase the leverage that we have even the rest of the world that is maybe in between the superpowers.

I want to talk a little bit about another aspect that actually came up in an audience question here. Which is the impact of AI on the labor markets, right? I think we have two potential effects. One is that a lot of tasks might get much more efficient and we can actually improve our decision making and improve the workflow in this area, for instance, the medical field. But on the other hand, a lot of people are worried about the impact on low-skilled workers or even some high-skilled workers in the end in terms of unemployment. And I'm curious what are your thoughts -- and it's to both of you -- kind of like what trends do we see so far? What kind of policies can we put in place to mitigate some of those impacts?

MR. WEST: I would say that AI has a potential both to take jobs as well as create new kinds of jobs. We're already seeing AI starting to replace jobs in the finance sector. You know, the examples I gave before in terms of fraud detection and wealth management. There are going to be accountants who lose their jobs because AI can add up the numbers as well as, if not better, than human beings.

In the retail sector, we're seeing fully automated retail stores operate. I actually visited one in Seattle. And basically you go in, you go shopping. The company uses computer vision and AI

applications tied to your credit card or your mobile payment system. You walk out. It automatically charges your account. No sales clerks, no cash registers. So, there certainly are going to be entry level jobs are at risk.

But there are going to be new types of jobs that are created. Certainly, data analytics. Any young person out there who wants a guaranteed future, learn data skills. Because we need data scientists. We need people who are skilled at analyzing large datasets.

I think the greatest risk is what I would call a skills mismatch. Where people don't have the skills for the new jobs that are going to be created. And in the book, we talk a lot about workforce development, job retraining, and lifelong learning. Kind of the old education model where people invest in education up through about age 25 is going to be obsolete. People are going to have to upgrade their job skills at age 30, 40, 50, and 60. Literally, throughout their lifetime.

So, companies are starting to develop job retraining programs for their employees. And they certainly need to do that and do even more of that. Universities need to move into adult education. That is going to be a big growth market. Community colleges already are doing a good job, and four-year schools need to do more of that as well.

And from a public policy standpoint, we need to think about the question, who's going to pay for this lifelong learning? Like there are young people coming out of college with 50, 60, or \$70,000 in higher education debt. When I tell them you need to engage in lifelong learning, they tell me you're crazy like you expect me to keep paying for advanced training even beyond the age of 25? And so, there's a public policy question in terms of that important issue. Who is going to pay for lifelong learning?

MR. ALLEN: And I'll come in behind Darrell. And he makes some very important points. You know, with other -- the other prior industrial revolutions, often advances in technology did not have with them the coordinated efforts to be taken. The coordinated measures that would be taken by society to deal with the displacement of those that would, in fact, lose their jobs because of enhanced technologies. And so, what we saw was often some fairly substantial and widespread societal disruption as the societies in the prior industrial revolutions or at any point in the march of technology embraced and adopted and applied these technologies. And there was widespread largescale disruption and unemployment.

Look, this is a fourth industrial revolution, and we can see this coming. And Darrell has just laid out some really important obligations that we have. Going back to ELSi, if we see these technologies coming and if we see that there will be some fundamental changes in the scope of work or the kinds of work that'll be necessary ultimately to fully embrace a digital horizon as it's coming towards us, then we have an obligation. It goes back to this issue of the character versus the nature of our society. We have an obligation as we embrace a more sophisticated character to our society, to do all we can to ensure that we buoy the nature of our society at the same time. And that's with jobs skills. It's vocational training. It's a commitment by the society to enhance lifelong learning.

MR. PATNAIK: Yeah, I think that lifelong learning is a really interesting aspect. And I think that will even go throughout all different levels of education and not only for low-skilled workers. And that will really require different ways of approaching education and different ways of organizing our workforce. So, this is a great point.

I want to finish with two final questions for both of you. Maybe, John, we can start with you first. If you could name one aspect of AI that you are most excited and positive about, which one would it be? And one aspect that you're most concerned about.

MR. ALLEN: I think I'm most positive about how AI will change the way we both educate our children. Well, the outcome of education. I am very excited about where AI can take education as we go forward. It will have the benefit of not only preparing us for this digital horizon that we're facing, but also the capacity to enhance the innovation in a rolling capacity as we go on. So, I'm very excited about how education can be enhanced using artificial intelligence as we go forward.

And, of course, we've talked about it in a number of ways, and that is in the context of the concerns I have. And that is the concerns for the privacy and the rights of our individual citizens. And that's not just as a result of the reaching of the Russians or other malefactors, the Chinese, into our cyberspace and information space and attempting to change the way our people think or act. It's how we do it to ourselves.

And so, we have to be constantly vigilant. Constantly vigilant to the rights of our citizens to ensure that either inadvertently or intentionally we are not employing these technologies to the detriment of our population. And we need to defend ourselves from our opponents and enemies from that

as well. So, that's my biggest concern.

MR. PATNAIK: And I think you bring up an important point that oftentimes a lot of privacy violations actually come from private companies, of companies here in the United States, right? The loading of our data. So, we don't have the policy framework for that.

MR. ALLEN: That's right. Now, I'm obviously concerned about artificially intelligent hyper capable nuclear weapons. I'm very concerned about that.

MR. PATNAIK: Yeah, I think everyone should be.

MR. ALLEN: But that's a ways off still, and we're going to have the conversation. But the moment-to-moment existence of the American citizen or the citizen, you pick the country, of democracy or in any other state, I worry about the human rights of that citizen, the privacy of that citizen, and ultimately the quality of life being enhanced by artificial intelligence, rather than degraded or oppressed by artificial intelligence.

MR. PATNAIK: Definitely. What about you, Darrell? Same questions.

MR. WEST: I think what I'm most excited about is the potential for various kinds of new technologies to free humans from boring, dirty, or dangerous jobs. And I think we're actually close to that. Technology is advancing in all sorts of ways. And it may be that some of the worst jobs in the sense of being boring and dirty and dangerous, we actually can automate. And what that could do if they are accompanied by appropriate public policy changes, is to free people for more creative activities.

Like I know lots of people who are interested in the theatre, of the arts and crafts, music. They have hobbies on the side. They're interested in sports. There are all sorts of things that people like to do other than their job. It may be that we are able to end up in the best of both possible worlds in the sense of not having, you know, being relieved from the tedious jobs and being saved for more creative activities. So, I'm excited about the possibility of being able to move in that direction.

I think the thing I'm most concerned about is the loss of human agency. Like when I give talks either in the United States or around the world, what people are really worried about is the loss of human control over technology. Like, you know, they see in all the Hollywood movies, "The Terminator," and it's basically hyperintelligent AI-driven superbeings that ultimately enslave humanity. And that we're no longer in charge. People are really worried about that.

And lot of the concerns about privacy, security, human safety, and so on, basically, come out of its broader worry about loss of human control. So, what I like to tell people and John and I make this point very clearly in the book. It's actually an optimistic book in a number of different senses. Humans are still in charge in the sense that we control the policies, the laws, the regulations, and the ethical standards that if appropriately implemented, can move technology in the direction of human values, protect human safety, and move the technology in a way that it conforms to human values and that people are in charge and not the technology itself.

MR. PATNAIK: I think that's a really important point. And I can totally see why people are concerned about that.

Well, we're at the end of our chat. It's 3:30. I really want to thank both of you. This has been super interesting. I appreciate your time. I know both of you are very busy so thank you very much.

MR. ALLEN: Honored to be with you, Sanjay. Thank you for all you're doing. Darrell, thank you. It's wonderful to see you again.

MR. WEST: Thank you, John.

MR. ALLEN: Good day.

MR. PATNAIK: Bye.

* * * * *

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

ANDERSON COURT REPORTING
1800 Diagonal Road, Suite 600
Alexandria, VA 22314
Phone (703) 519-7180 Fax (703) 519-7190

(Signature and Seal on File)

Notary Public in and for the Commonwealth of Virginia

Commission No. 351998

Expires: November 30, 2024