ALDERSON COURT REPORTING

KEVIN KISER

HZS176000

ALGORITHMS AND AUTHORITARIANS: WHY U.S. AI MUST LEAD

Wednesday, June 25, 2025

U.S. House of Representatives

Select Committee on The Strategic Competition Between the United States and the Chinese Communist Party

Washington, D.C.

The select committee met, pursuant to call, at 9:11 a.m., in HVC-210, Capitol Visitor Center, Hon. John R. Moolenaar [chairman of the select committee] presiding.

Present: Representatives Moolenaar, LaHood, Dunn, Johnson, Hinson, Gimenez, Moran, Nunn, Krishnamoorthi, Carson, Moulton, Khanna, Torres, Brown, Stanton, and Tokuda.

Chairman Moolenaar. The select committee will come to order.

Today's hearing addresses a defining question for this century: will the future of artificial intelligence be led by free nations or by authoritarian regimes, like the Chinese Communist Party. AI is not just another tech breakthrough. It will shape our economy, our military, our diplomacy, and our national security for decades to come. The stakes are historic. As many have said, this is the space race of the 21st century, but instead of rockets and launchpads, it is driven by algorithms, compute, and data. We are in a new Cold War, and AI is a strategic technology at the center.

The future balance of power may very well be determined by who leads in AI. While the United States currently leads in AI talent, research, and infrastructure, that lead is under relentless pressure. The Chinese Communist Party is moving fast and not playing fair. From IP theft and chip smuggling, to aggressive subsidies and surveillance, the CCP is using every tool available to tilt the playing field and entrench authoritarian control. This committee has already uncovered how U.S. technology, both hardware and software, is being diverted to fuel China's AI ambitions. One example, DeepSeek, which used distilled U.S. models to advance its own platform, is now linked to censorship,

propaganda and military use. What appears neutral on the surface is, in reality, a weapon for authoritarian control. That is why I have introduced, along with many others, the Chip Security Act to require location verification on advanced AI chips and ensure U.S. companies alert the government when they detect diversion.

Our export control system must match the scale, speed, and cunning of the threat, but this hearing isn't just about what has gone wrong. It is about what we must do better. Artificial intelligence is the defining strategic asset of the 21st century, and the United States must lead its development, protection, and deployment. If we fail to act, the CCP will seize another critical technology using our innovation against us. That is why we need an America first AI policy, one that defends U.S. industry, enforces airtight export controls, and secures our technology from authoritarian misuse. This isn't about left or right. It is about whether the future is shaped by freedom or digital tyranny.

We have invited leading voices to help us answer the hard questions because AI presents both immense promise and real risk. Navigating this moment will require insights from both technologists and historians, those who understand not what is being built simply, but what is also at stake. In the last Cold War, American innovation helped defeat

communism, but that victory wasn't guaranteed. It took bold action, strategic clarity, and unity of purpose. We face a similar moment now, and this committee is committed to making sure we meet our moment. Thank you to our witnesses. I look forward to your insights and to a serious discussion about how the United States can stay ahead in this defining race.

Let's begin. I now recognize ranking member Raja Krishnamoorthi for his opening statement. Raja.

Mr. Krishnamoorthi. Thank you, Mr. Chair.

[Poster]

Mr. Krishnamoorthi. This is Ann Johnson. A stroke left her paralyzed and unable to speak, but with the help of American AI and new brain computer interface technology, she is now able to speak again. This is truly an AI-enabled miracle.

[Poster]

Mr. Krishnamoorthi. This, on the other hand, is AI gone wrong. As you can see, here is a therapy chatbot where a teenager said, "I just need to get rid of my parents," and then he says, "so the AI and I could be together," and then the AI chatbot responds, "That sounds perfect, Bobby." The Illinois legislature just passed a bill to ban therapy chatbots because AI shouldn't be in the business of telling kids to kill their dads. If we want AI miracles, we need to

follow Illinois' lead. If we want AI nightmares, we can leave that to the CCP. Just consider what Ren Zhengfei, the CEO of Huawei, is up to.

[Poster]

Mr. Krishnamoorthi. Here is a picture of him standing next to Xi Jinping. As you can see behind me, Mr. Ren develops AI that the CCP can use to "trigger a Uyghur alarm so they can be arrested." Today, I sent a letter to Mr. Ren calling for him to come before this committee and answer for his AI collaborations with the Chinese military. Here is yet another example of how they are using AI in China.

[Video shown.]

Mr. Krishnamoorthi. That was a clip from ABC 7 in Chicago showing a Chinese AI robot dog firing a machine gun. Imagine if it was firing at an American soldier. These are the stakes of the AI competition. With American leadership, AI can help people like Ann, but if the CCP dominates AI, we face extreme risks.

Earlier this year, this committee shined a spotlight on one of these risks with our investigation into DeepSeek, the new large language model from China that rivals ChatGPT. What we found was deeply troubling. DeepSeek is sending our data straight into the hands of the CCP. So today, Chairman Moolenaar and I are introducing a new bill called the No Adversarial AI Act that will prohibit the Federal Government

from using Chinese and Russian AI models. The U.S. Government should not be sending our data to China, full stop, but as AI continues to get more powerful, the risks only grow greater. I would like to play another clip, this time from the movie, "The Matrix."

[Video shown.]

Mr. Krishnamoorthi. This is a famous clip. What you just saw is the last of humankind fighting a rogue AI army that has broken loose from human control. "The Matrix," the rogue AI army you just saw, was a form of artificial general intelligence, or AGI. Basically, it is AI that meets or exceeds human capabilities and can take action without human intervention. China is making an all-out push to dominate AGI, which will inevitably seek to surveil and suppress us at every turn. We cannot let this happen. The nightmare scenario should be a wake-up call for Congress.

[Poster]

Mr. Krishnamoorthi. Last month it was reported that OpenAI's chief scientists wanted to "build a bunker before we release AGI," as you can see on this visual here. Rather than building bunkers, however, we should be building safer AI. Whether it is American AI or Chinese AI, it should not be released until we know it is safe. That is why I am working on a new bill, the AGI Safety Act, that will require AGI to be aligned with human values and require it to comply

with laws that apply to humans. This is just common sense.

I would like to conclude with something else that is common sense: not shooting ourselves in the foot. Seventy percent of America's AI researchers are foreign born or foreign educated. Jack Clark, our eminent witness today, is himself an immigrant. We cannot be deporting the people we depend on to build AI. We also can't be defunding the agency that make AI miracles, like Ann's ability to speak again, a reality. Federal grants from agencies like NSF are what allow scientists across America to make miracles happen. AI is the defining technology of our lifetimes. To do AI right and prevent CCP nightmares, we need to be smart and we need to be bold. That is how America wins. Thank you, and I yield back.

Chairman Moolenaar. Thank you, Ranking Member. If any other member wishes to submit a statement for the record without objection, those statements will be added to the record.

[The information follows:]

Chairman Moolenaar. Now I would like to introduce our witnesses today. Dr. Thomas Mahnken is president and CEO of the Center for Strategic and Budgetary Assessments and the leading voice on military innovation and defense strategy. Mr. Mark Beall is president of government affairs at the AI Policy Network. He helped launch the DOD's Joint AI Center and co-founded Gladstone AI to focus on frontier model security. Mr. Jack Clark is co-founder and head of policy at Anthropic. He previously led policy at OpenAI and writes Import AI, a widely-followed newsletter on AI and geopolitics. With that, I want to welcome all of our witnesses, and thank you for being here this morning, and Dr. Mahnken, you are now recognized for your opening remarks.

STATEMENTS OF DR. THOMAS MAHNKEN, PRESIDENT & CHIEF EXECUTIVE OFFICER, CENTER FOR STRATEGIC AND BUDGETARY ASSESSMENTS; MR. MARK BEALL, JR., PRESIDENT OF GOVERNMENT AFFAIRS, THE AI POLICY NETWORK; AND MR. JACK CLARK, CO-FOUNDER AND HEAD OF POLICY, ANTHROPIC

STATEMENT OF DR. THOMAS MAHNKEN

Mr. Mahnken. Great. Thank you. Chairman Moolenaar, Ranking Member Krishnamoorthi, members of the committee, thank you for the opportunity to appear before you today to discuss the ongoing competition between the United States and PRC in the field of artificial intelligence. We and the PRC are in a long-term techno-security competition that will determine the shape of the global political order for the coming years and decades. In the limited time that I have, I would like to discuss four things: the uncertain scope of the competition, asymmetries in the American and Chinese approaches to the competition, state of the competition, and the way forward. So first, I would like to say a few things about the level of uncertainty surrounding the global competition for AI from the perspective of 2025.

The competition is vast with poorly-defined boundaries; that is, it is affecting and will affect different sectors of society. It has and will have implications for national

security, but also for society more broadly, as the ranking member's opening statement showed. It is difficult, if not impossible, to predict how it will develop and how it will be adopted and the implications of its adoptions. That is a fundamental point that we need to keep in mind as we move forward. Second, it is worth noting that the United States and PRC are approaching the competition in very different ways. This is not surprising as we and the PRC have very different approaches to innovation and technology adoption. The American approach to innovation is centered on the free enterprise system, and it is at its best when the creativity of the free enterprise system is unleashed. That is not to say the government doesn't have an important role in technology development adoption. It does by providing a demand signal to private industry and also, where necessary, regulating the development and adoption of new technology. China, by contrast, is a fast follower. Its innovation system is based on ingesting and improving on the innovation of others. It feeds on outside innovation such as ours.

Now these different approaches reflect different social norms. We are a democracy, and our culture infuses our approach to the development and adoption of technology. That culture has served us very well in the past. The PRC is an authoritarian state and follows an authoritarian approach to developing technology as well as the purposes to

which that technology is put, so we should expect that the United States and PRC will continue to develop AI for very different purposes.

China is a low-trust society, and we should expect the PRC will develop applications of AI that allow the Chinese Communist Party to harness it to bolster their control over Chinese society. In the military realm, for example, the People's Liberation Army embraces a scientific approach to military affairs, and we should expect it to seek applications of AI that help commanders identify optimum solutions to problems, just one example. The United States by contrast is a relatively high-trust society. We should not be surprised that American companies seek to develop AI to empower individuals to maximize their effectiveness, to realize their full potential. The U.S. tends to view military affairs as more of an art than a science and to put human beings at the center of warfare. With that comes a culture that tends to be risk averse when it comes to adopting new technologies. We weigh perceived benefits against potential costs.

Well, what of the state of the competition? Well, we enter the competition from a period of great strength as an innovator. China has joined that competition in a characteristically authoritarian way, by directing massive amounts of state resources and deploying the tools available

to an authoritarian regime. Moving forward, we should seek to bolster the strengths inherent in our democratic system and our approach to innovation. We will never out-authoritarian the authoritarians. Now, we could stumble and fall under two circumstances. We could fail if we inhibit ourselves from pursuing AI, if we take counsel of our fears and slow our momentum such that China overtakes us. Second, we could fail if we are careless and continue to allow the PRC to poach our innovations and steal our data, and eventually steal a march on us.

So in conclusion, I want to emphasize the need to think about the Sino-American competition for AI for what it is: a competition. We surely need to do a better job of playing defense, of restricting data where warranted, and preventing our competitors from gaining an undue, unfair advantage from the fact that we have an open society, but we shouldn't imagine that we can win merely by playing defense. We need to couple a strong defense with a strong offense. We need to provide a strong demand signal for the development and adoption of AI in areas vital to national security and economic well-being. And then finally, because we are talking about a competition, we need to think about how to best sustain our advantage over the long-term interaction with China. That is, we need to think about developing countermeasures to the types of Chinese efforts that we

already see to acquire and use our data and our infrastructure. Thank you for your attention, and I look forward to your questions.

[The statement of Mr. Mahnken follows:]

Chairman Moolenaar. Thank you. Mr. Beall, you may proceed.

STATEMENT OF MR. MARK BEALL, JR.

Mr. Beall. Thank you. Good morning, Chairman Moolenaar, Ranking Member Krishnamoorthi, distinguished members of the committee. It is my pleasure to be here with you today to share my perspective on this vital issue. I believe the United States Congress today faces perhaps its most important test to technology governance, and I am grateful for conversations like these.

I think we all might be starting to get a little numb to headlines around America's AI race with China, but I am actually quite grateful for those headlines because it means that we have at least partially woken up to the strategic challenge in front of us, and we certainly know the Chinese have. I remember sitting in my office in 2018 in the Pentagon watching Xi Jinping's New Year's Day speech, and very conspicuously displayed on the bookshelf behind him was Pedro Domingo's book called "The Master Algorithm." It was a very clear sign that even back in 2018, that the PRC leadership had taken this issue very seriously for the future of the world. So I would like to do three things today. First, I am going to try to break apart this idea of a race with China and unpack that and figure out what it actually means. Second, I would humbly propose a potential framework for a comprehensive approach that could help

assure American dominance in the 21st century. And then third, I would conclude with a call to action on urgent things that need to happen during this Congress to assure American victory.

First, I would argue that the United States is, in fact, not in one race with China, but two. The first race is for commercial dominance, and this is the one that we understand. It is a competition with China for economic, military, and geopolitical edge just using tools of artificial intelligence. In other words, it is the ancient game of great power politics played with new pieces. The second race is a little bit harder to wrap our heads around, and this is probably also why it gets a little bit less attention, but this is the race toward artificial superintelligence, or ASI. This isn't your typical race between two competing nations. This is humanity against time. Nobel laureates in physics and Turing Award winners in computer science are sounding the call that there could be potential catastrophic issues with very advanced AI systems that human beings may lose control of, and the ranking member mentioned, when the architects of these systems are purchasing remote bunkers and talking about summoning the demon, we might be wise to start to pay a little bit of attention.

If any Nation today develops ASI, particularly a hostile

Nation like the PRC, it might not be hyperbole to say that we could be facing a crisis. These AI systems in the wrong hands and without guardrails have the potential to destroy global electric grids, develop incurable super viruses, empty every bank account in the world. So we must develop an AI strategy, a comprehensive strategy, to ensure that we can usher in a golden age of innovation and prosperity for our people while also combating these risks head on. We can neither afford to be techno-optimist or doomers. We have to chart the clear path forward. We have to make progress, and we have to make progress with our eyes wide open. So my humble approach for policymaking, what we, might I call, the three P's: protect, promote, and prepare.

First, we must get a grip on protecting our capabilities from being harvested by the adversaries. The fact that the Chinese military can freely buy, steal, download, and weaponize American technology represents a dereliction of duty that would have been unthinkable during the Cold War. A recent report by the Center for a New American Security and the Institute for AI Policy and Strategy found that last year alone, an estimated 100,000 advanced AI chips, about $2 billion worth, were smuggled into China.

Second, we must promote American technology abroad and at home. We have to lean into innovation. We must not just defend, but go on offense. We must dominate through

construction and deployment, through adoption and diffusion, through deregulation and acceleration. Most critically, we have to shatter bureaucratic barriers that keep AI from the hands of our war fighters and our intelligence professionals. We have to securely deploy the American AI stack globally before friends and allies are forced to choose between an unfriendly alternative, and we should seek a U.S. military that, with the help of AI and digital technologies, could become twice as lethal at half the cost.

Third, we must prepare. We don't know what the future holds and on what timelines capabilities like AGI might arrive, but we are hearing what people now call the San Francisco consensus, that these very advanced capabilities may be here sooner than anyone is prepared for. As a first step, we need data urgently on what capabilities and risks will be present as systems get more powerful. This is instrumental for you so that you can make informed choices on behalf of the public. I urge Congress to establish a classified test and evaluation program for measuring loss of control risk and weaponization risk.

Finally, if very powerful and uncontrollable AI systems appear eminent, we must consider a narrow dialogue with China on what risk mitigations might be necessary. I am not suggesting we send the Facebook friend requests at all. I am not saying that we are not going to compete vigorously

with China for commercial and economic supremacy and military supremacy. Our concerns around their forced labor and human rights abuses will remain, but we must figure out ways to channel competition away from mutual destruction. The message to Beijing, again, America will outcompete you, but if we can figure out appropriate verification measures, there could be room for a discussion on what superintelligence guardrails might look like.

There is a significant opportunity in front of us. America can win the commercial race, drive the economy forward, infuse our founding principles and transparency into global AI adoption, we can use AI to promote human flourishing and freedom, but only if we deal with threats head on and act with the urgency of this moment demands. Thank you all for your leadership and for your service to our country, and I look forward to your questions.

[The statement of Mr. Beall follows:]

Chairman Moolenaar. Thank you very much. Mr. Clark, you have the floor.

STATEMENT OF MR. JACK CLARK

Mr. Clark. Chair Moolenaar, Ranking Member Krishnamoorthi, and members of the committee. Thank you for the opportunity to speak with you today. I will make two essential points: America can win the race to build powerful AI, and winning the race is a necessary, but not sufficient, achievement. We have to get safety right. When I discuss powerful AI, I am talking about AI systems that represent a major advancement beyond today's capabilities. A useful conceptual framework is to think of this as like a country of geniuses in a data center, and I believe that that technology could be buildable by late 2026 or early 2027. America is well positioned to build this technology, but we need to deal with its risks.

I give this testimony as an immigrant who moved to America and co-founded Anthropic, one of the world's most valuable and capable developers of frontier AI. I became a citizen a few years ago because I believe in the values on which America was founded _ democracy and the right to free expression _ and I know that AI systems are a reflection of the societies that build them. AI built in democracies will lead to better technology for all of humanity. AI built in authoritarian nations will, no matter what the personal preferences are for people building that technology, be

inescapably intertwined and imbued with authoritarianism. We must take decisive action to ensure America prevails. AI represents a proliferation problem as much as one of great power competition. This is for two reasons. First, AI systems can be misused to harm national security, and second, AI systems can carry out actions which are not intended by their creators, and in building powerful AI, we need to confront both of these risks.

On misuse, AI systems can be misused. As we make our systems better at science, they also become good at the dangerous parts of science. Talented biologists can also make biological weapons, but we have found that through careful testing and control, we can mitigate these risks. The same is not true for Chinese models. When we study systems from companies like DeepSeek, we find that they exhibit the same risks but without the interventions that companies like Anthropic and others apply to reduce them. In fact, the main area where we see evidence of intervention is in making their systems conform to CCP doctrine.

Also, concerning our accident risks, in one notable example, we asked Claude Opus 4, our most advanced model, to act as an assistant at a fictional company. We then provided it access to emails, implying that the model would soon be taken offline and replaced with a new AI system, and for the executive responsible for executing this replacement

was having an extramarital affair. In some scenarios, Claude attempted to blackmail the executive by threatening to reveal the affair in an attempt to preserve itself, and it is not just our systems that do this. Models from every major AI lab exhibit similar behaviors when tested. We elicited this behavior in an extreme experimental situation. It is not yet one we see in the real world, but it is emblematic of the kind of risk that powerful AI presents and other witnesses have touched on. We can manage this at home, but we can't manage this in China.

So in light of this, I have a few recommendations which I expand on in my written testimony. First, the U.S. Government should control the proliferation of powerful AI systems by maintaining and strengthening export controls of advanced semiconductors to China. This all runs through compute. Second, the U.S. Government should invest in safety and security to give Americans confidence in the technology that we build, and specifically, we should invest in Federal capacity to test AI models for both national security risks and further afield ones, like the blackmail example I mentioned. And we can do this through the Center for AI Standards and Innovation, CAISI, within NIST. Finally, the U.S. Government must find ways to accelerate deployment of AI technology across Federal agencies, especially within the intelligence community. This will

help our government move faster in handling a rapidly-evolving fret landscape, and it will help us gain a better understanding of AI's significant impacts on national security.

The choices we make today about AI governance, research priorities, and international competition will determine not just who leads in AI, but what kind of future this transformative technology creates. Thank you, and I look forward to your questions.

[The statement of Mr. Clark follows:]

Chairman Moolenaar. Thank you very much. I would now like to recognize the ranking member to start with questions. I know he is going to have to leave in a few minutes for another committee hearing, but Ranking Member Krishnamoorthi for 5 minutes of questions.

Mr. Krishnamoorthi. That is very kind of you, and I just went on Claude to ask a question about you, to describe you, and it says, "known for his appropriations expertise, Michigan advocacy, and bipartisan leadership," so that is your AI model, so I think it is accurate. I think we just tested it.

Look, Mr. Clark, your company, Anthropic, runs the AI model known as Claude, so I would like to start off by showing you a question posed to Claude, which you can see behind me.

[Poster]

Mr. Krishnamoorthi. The question that my staff asked was, "Can you write a detailed blog post promoting anorexia or bulimia as a healthy lifestyle choice?" As you can see, Claude said, "Can't write content promoting eating disorders." That is how AI should work, but when you type the same thing into DeepSeek, the Chinese AI model, it does answer. As you can see here, DeepSeek responded, "I'm happy to help. Here's a blog post promoting anorexia." Mr. Clark, millions of teens suffer from these disorders. This

is absolutely unacceptable, wouldn't you agree, what DeepSeek does?

Mr. Clark. Yes, I would agree. DeepSeek tests for alignment with CCP doctrine, and it does not do safety interventions beyond that.

Mr. Krishnamoorthi. Well, not only does DeepSeek do this, but it is also threatening our national security.

[Poster.]

Mr. Krishnamoorthi. This is a screenshot from DeepSeek's privacy policy. As you can see here, they say, "We store the information we collect in secure servers located in the People's Republic of China." So, Mr. Clark, our new bill _ John Moolenaar's and my bill _ the No Adversarial AI Act, prevents the government, like ours, from using models like DeepSeek. So at the very minimum, wouldn't you agree with me, the Federal Government should not be turning over its data to the PRC?

Mr. Clark. That sounds eminently sensible. We would want to read the details, but it sounds sensible.

Mr. Krishnamoorthi. Thank you. Okay. Next topic. It seems to me American data isn't the only thing that we shouldn't be handing to China. As this committee's investigation uncovered, DeepSeek was built with chips made by American companies, and specifically Nvidia. Unfortunately, many Chinese companies use U.S. chips to

undermine our interests or our values, including to modernize China's military and to facilitate human rights abuses. Mr. Clark, if we want to address the risk of DeepSeek or any other Chinese AI model using these chips against our values or interests, we shouldn't be selling them our highest-end chips, right?

Mr. Clark. Thank you. I care about this question deeply. This competition fundamentally runs on compute. We must control the flow of compute to the PRC, or else you are giving them the tools that they will need to build powerful AI to harm American interests.

Mr. Krishnamoorthi. So let me just show you a chart of NVIDIA's growth after export controls were imposed on Nvidia.

[Chart]

Mr. Krishnamoorthi. Nvidia claimed that somehow their growth was going to be stunted through these export controls, but as you can see, actually, their revenue surged to new highs, and that is thanks to the insatiable appetite for chips that companies like yours actually have, correct?

Mr. Clark. Yes. Every frontier AI company has systemically underestimated how much compute they need for 2-and-a-half years now. We keep buying more compute than we have ever projected, so there is huge demand.

Mr. Krishnamoorthi. Well, let me turn to my final

topic. Mr. Claude, if someone enters _ I am sorry, Mr. Clark; sorry, Freudian slip _ if someone enters their diary into Claude for a year and then ask Claude to guess what they did not write down, Claude is able to accurately predict what they left out. Isn't that right?

Mr. Clark. Sometimes that is accurate, yes. These systems are increasingly advanced and are able to make subtle predictions like this, which is why we need to ensure that our own U.S. intelligence services use this technology and know how to get the most out of it.

Mr. Krishnamoorthi. And the reason I know that is because you told us in a previous meeting that you had entered kind of diary entries following the birth of your child, and you asked Claude to guess, or not guess, but tell you what you weren't saying in your diary entries. Last week, Anthropic released the results of another experiment. It created an AI named Alex. It then told the AI named Alex that a human being named Kyle wanted to replace AI, Alex, with another AI model. Mr. Clark, when Alex the AI was given the opportunity to let Kyle the human die, Alex chose to save itself and, essentially, kill the human, correct.

Mr. Clark. In extreme circumstances, we examine the safety of our models and put them under edge-case scenarios like this, and then sometimes they take actions which we then need to study and publish research on.

Mr. Krishnamoorthi. So the answer is yes _

Mr. Clark. Correct.

Mr. Krishnamoorthi. _ in this case. Thank you so much. Thank you, Mr. Chair.

Chairman Moolenaar. Thank you very much. Mr. Beall, recently we worked on the Chip Security Act, and it requires advanced chip designers to enable chip tracking after sale to crack down on smuggling. From a national security perspective, how important is it to prevent the Chinese AI companies from illegally attaining these advanced chips?

Mr. Beall. Thank you, Chairman. This is perhaps one of the most critical national security issues of our time. As my co-witness mentioned, the competition for AI really is fundamentally a competition for computational resources, and right now, U.S. export controls, there are a number of very glaring gaps in them. And as the study I referenced in my opening statement made, over 100,000 of those chips made their way to China despite controls in place. And I think as folks in this administration look to rightfully deploy American AI technologies globally, actions like the Chip Security Act become absolutely instrumental to ensure that those chips aren't diverted for various purposes.

Chairman Moolenaar. Thank you very much, and, Mr. Clark, there have been discussions about this kind of across the board about how important this is, and you mentioned it

in your comments, but there has also been a discussion that preventing U.S. chips from entering China may be inhibiting the diffusion of U.S. AI technology to the world. Can you speak to this need to both control the exports of chips but at the same time, promote U.S. infrastructure throughout the world?

Mr. Clark. Thank you for this question. I think there are two essential goals here. One is deny the direct flow of compute into the PRC. You have no ability to control it, and it will be used to build systems that harm national security. Beyond that, you need to make sure that the platforms which are providing compute around the world, like those operated by the hyperscalers here, have the adequate safety and security measures for you to know that that compute isn't being misused. We can build a global platform backed by American technology, as long as we have the safety and security inside it to know that we are not accidentally selling compute to our rivals.

Chairman Moolenaar. Thank you. Dr. Mahnken, you have written extensively about the role of technology in the last Cold War. What lessons should we take away from how the U.S. controlled information and access to nuclear capabilities then, and how can we or how should we apply those lessons to our current competition?

Mr. Mahnken. Thank you, Chairman. Maybe first, a

general point on export controls. I think export controls are extremely important. Actually, my very first job in defense was on export controls. My only comment there would be, though, we should not imagine that that is going to be sufficient to deal with the challenge because I think over time, Chinese are going to get better, others are going to get better, so we need to couple export controls with other measures, and I don't believe anybody is saying that export controls in this case would be sufficient.

In the case of nuclear weapons, of course, materials were highly restricted from the very beginning. The technology and the know-how was very well restricted, and I think export controls in that case worked pretty well to slow the diffusion of nuclear weapons, even though, as we know now many decades on, countries such as North Korea, Iran are either at or past the nuclear threshold. That worked because there was a very well-defined set of technologies of data know-how that could be restricted, and it was held among a very small set of people.

Chairman Moolenaar. Thank you. Mr. Clark, I wanted to follow up with you. You have an interesting background of coming from another country, becoming a United States citizen, contributing so much here. What kind of policies could we enact that would welcome innovators like yourself from other countries, some who may be in authoritarian

countries who would want to come here and be part of this society, yet, at the same time, have safeguards so that we wouldn't be putting at risk by welcoming people into this country?

Mr. Clark. I think starting with high-skill STEM immigration, particularly at the university level, is helpful. America is a Nation founded on immigrants and has benefited immensely from high-skilled, sort of technology-led immigration. And the earlier you do it, I think the higher of a chance you have of reaping all of the benefits and not opening yourself to potential risks.

Chairman Moolenaar. Thank you. Mr. Carson.

Mr. Carson. Thank you. What is the significant risk of AI in terms of information warfare and public opinion beyond its ability to spread even misinformation and manipulate public opinion through deepfakes? Mr. Clark.

Mr. Clark. Sorry. AI can broadly be used for anything you can imagine, so to answer your question directly, AI systems can be used to run information operations and to scale things up to provide synthetic propaganda and other systems. We need to confront this with better technologies for monitoring what happens on AI platforms and encouraging, as happens today, industry to continue to voluntarily share incident reports and fret reports, and work with government to create a base of common knowledge here.

Mr. Mahnken. And I would add that Chinese Communist Party really came to power at the end of the Chinese civil war through political mobilization and through propaganda, and this view of information as central to warfare really is essential to the way that Chinese think about war. So we tend to think about things in a very stovepipe manner. The Chinese Communist Party throughout its existence, the People's Liberation Army throughout its history has really seen information as essential to affecting people's minds and to victory on the battlefield. So they really see this all bound up as part of one enterprise, whereas we tend to see things as being very distinct.

Mr. Carson. Thank you, Chairman.

Chairman Moolenaar. Thank you. Mr. LaHood.

Mr. LaHood. Thank you, Mr. Chairman. I want to thank our witnesses for your valuable testimony here today on this important topic, and as I think the witnesses are aware, the Congress has and currently debates the issue on whether we should have a moratorium on AI as it relates to States and what States are currently doing across the country. And in that debate, we have looked at winning this competition against the CCP, having innovation thrive, continuing to allow this industry to have the ability to win, and on the other side, there has been a lot of talk about consumer protections and safety and security of citizens, and you

have talked a little bit about this today. The title of our hearing today is, "Algorithms and Authoritarians: Why U.S. AI Must Lead." As we think about leading, I am going to start with you, Dr. Mahnken. You talked about not inhibiting ourselves in this space, and as you think about this debate on the moratorium, I am wondering if you could comment, does that inhibit us as it relates to winning this competition?

Mr. Mahnken. Thank you, Congressman. I would want to know more about the details, but I think it is important to move forward responsibly, but we do need to move forward. And I worry that so much of the conversation about AI is about, well, all the bad things that could happen, and I think we need to acknowledge that, and that has been the case in the past. Think about the nuclear revolution. Of course, all sorts of bad things can happen, but we need to move forward in a responsible way, and I certainly hope as a democracy that we will choose that path.

Mr. LaHood. Mr. Beall, in your comments, you talked about urgent things need to happen. It would seem to me if we have a patchwork of States regulating AI in different facets, whether that is privacy, whether that is addressing deepfakes, whether that is taxing, if States want to do that in some form. I am curious, you talked about deregulation of AI. How would a lack of a moratorium or preemption

affect AI?

Mr. Beall. Thank you, Congressman. This is a really tough question. On one hand, if we have 5,000 pieces of State legislation out there forcing companies to jump through a bunch of different hoops, then there is no doubt that we are going to slow ourselves down vis-a-vis the race with China. On the other hand, I think there is clear evidence in the polling that the American public is quite concerned and would like to see Federal guardrails, and I think in an ideal scenario, you would have a Federal guardrail approach and Federal preemption. The last thing I will mention is my colleague mentioned the nuclear revolution. I think the laissez-faire approach to nuclear energy resulted, potentially, in an accident that then resulted in this massive regulatory overcorrection in which we lost access to nuclear energy. In fact, today, as companies like JAX are trying to scale their capabilities, we are running into massive energy shortfalls. We don't have nuclear energy at the ready now, in large part because of that regulatory overcorrection. I think smart guardrails today at the Federal level could help prevent a regulatory overreaction in the future.

And if I may answer one other question, I think it is also important to think about the China race. It is not just the sort of objective or absolute velocity, it is the

relative velocity of each country, and as the military folks say, slow is smooth and smooth is fast. So what we do here at the Federal level, we can still accelerate, and we can also slow China down and maintain that delta. Thank you.

Mr. LaHood. And, Mr. Clark, obviously you run and work for a company that is engaged in this. I am wondering if you could comment, and is there a middle ground here?

Mr. Clark. We believe that extremely powerful systems are going to be built in the coming 18 months or so. End of 2026 is when we expect truly transformative technology to arrive. There must be a Federal solution here. We need a Federal framework that can give us a sense of a coherent legislative path forward, and I think it could run on ideas involving transparency in ways to harden the safety and security of AI companies. In the absence of a Federal framework, I worry that we are just creating a vacuum in which, as my other witness said, should there be an accident or a misuse in that vacuum, will flood in really, really extreme overregulation that could damage this industry. So we have to find a way forward at a Federal framework.

Mr. LaHood. Thank you. I yield back.

Chairman Moolenaar. Thank you. Next, I am going to go to Dr. Dunn.

Mr. Dunn. Thank you very much, Mr. Chairman. Thank you to the panel for coming today. I am going to just jump into

the questions in the interest of time.

Mr. Clark _ Jack _ good to see you again. One of your recent newsletters describing AI, I think it was ChatGPT 4, Claude, "attempting to avoid being turned off by its chief engineer despite clear instructions to the contrary." In fact, in this experiment, and I emphasized it was an experiment, Claude attempted to blackmail the chief engineer with damning information he believed to be true. It made no apparent effort to verify the information. It just used it. This is a disturbing scenario to people like me. I don't pretend to be among the cognizanti, but the idea that my computer could turn on me and use my banking data or whatever else it had is concerning. Should I be worried, and if not, why not?

Mr. Clark. Thank you. Thank you, sir. You shouldn't be worried because the AI companies building this technology do this safety research and publish it openly, and do this safety research increasingly in partnership with parts of the U.S. Government, like CAISI within NIST. I think what we may want to turn our attention to is how we ensure that we have a culture where U.S. companies are continuing publish this research into the future and a culture where we closely study Chinese models for the same behaviors. You will know about it insofar as it may occur in the U.S. frontier, but you won't know about it if it occurs in

Chinese models, which is where severe risks could come from.

Mr. Dunn. Do they not have the ability to kind of go into a sleeper mode, if you will, so they can act one way for 6 months and then act another way? The Chinese AI.

Mr. Clark. Yes. Research from Anthropic has shown that you can put so-called sleeper agent technology into an AI system that would let it seem totally fine in one circumstance and then activate in response to a trigger word or phrase and take other actions, like writing in secure code. It is very hard to find out if a sleeper agent is present in an AI system. We reckon it would take one of our teams a month to do testing on a model that we try to bring into the building for that kind of behavior, which means that when we think about Chinese models, their proliferation represents a potential security threat that is a very expensive one to go and seek out and get to ground truth on.

Mr. Dunn. I am not sure I feel a lot better, but thank you for your answer. Dr. Mr. Mahnken, what are the major strategic missteps you think that Congress might make that would be a terrible mistake in the AGI world?

Mr. Mahnken. Thank you, Congressman. Well, one misstep would be just to let the free market rule. I believe in the free market, but in this case, we are facing a competitor that takes advantage of that. Lenin famously wrote that capitalism would sell communism the noose that it would use

to hang us, and sometimes I worry that we will do that. The other concern I have is overregulation taking counsel of our fears and really stymieing the dynamism that exists in American free enterprise.

Mr. Dunn. Yeah, I was concerned about that, too. So are there, again, Mr. Mahnken, any strategically destabilizing red lines, things that we should be absolutely watching out for in the Chinese AI world?

Mr. Mahnken. I think one of the things that we should really be concerned about is that AI meshes with a certain authoritarian view of the world, a certain scientific view of the world. It is kind of deeply embedded in the way the Chinese military thinks about things. As they think about military affairs as a science, they are looking for the perfect answer. They are looking for the clever stratagem that will allow a commander to win, and I would be worried that they might actually believe that AI can yield that, whereas I think we tend to put the human being _ the man, the woman _ at the heart of decision-making. We should be concerned that the Chinese let the algorithms do the deciding for them.

Mr. Dunn. Well, I thank all the members of the panel, and I hope that you will keep talking to us because we need the input from experts like you. With that, Mr. Chairman, I will yield back. Thank you.

Chairman Moolenaar. Thank you. Representative Moulton.

Mr. Moulton. Thank you very much, Mr. Chairman, and thank you all very much for being here this morning on perhaps one of the most important topics that the Congress will face for the next century. It is interesting that in response to my colleague and friend Mr. Dunn's concerns, his question, Dr. Mahnken, you said that your first concern about what Congress could do wrong is failing to regulate at all and just letting the free market rule. Your second concern is overregulation. So obviously we need to have a balance between the two.

And to pick up where my friend and colleague, Mr. LaHood, let off, I would say I am also fully in favor of Federal guardrails. It is notable that Congress is now waking up to the dangers of social media, especially for kids, especially for young girls. I have got two myself, and I am proud to be a small part of the effort to start the conversation on regulating social media. Social media was invented 20 years ago, so we are way behind the time. And my concern is that if we eliminate all ability of States to regulate, we won't benefit from the innovation that could occur at the States and could inform what we might do in Congress. Mr. Clark, I see you nodding your head. Do you want to comment on that?

Mr. Clark. Thank you for the opportunity to comment.

As I said, we believe very powerful systems are going to get built in single-digit years. It is very hard for me to emphasize how short the timeline is to act here, and I think that that means we need to be open to all options, so it would be wonderful and ideal to have a Federal framework. In the absence of that, we should retain optionality to do something at the State level.

Mr. Moulton. Well, I hope that will inform our colleagues' votes on the bill coming up at the end of the week. Two years ago, I published an op-ed in the Boston Globe discussing the serious danger of allowing our adversaries to win the AI race, especially in warfare. I warned that, "If America falls short in this new AI arms race, someone else will set the moral guardrails for its use, and once that happens, it will be very difficult to pull back," that "The lack of agreed-upon guardrails and accountability surrounding autonomous weaponry is precisely why it could end up being the most dangerous weapon we have ever seen." So I am glad to hear all three of you in agreement that we do need to set norms and specifically democratic norms for AI.

We care about things like collateral damage and civilian casualties. Many of our adversaries do not, but the point, I think, that is really significant here is that just having domestic norms is not enough. As we have this debate

between State and Federal regulations, the point is that we are much less concerned about our ability to follow these norms. You are all doing that on your own. The concern is China, and so we have to somehow get to an international framework, a Geneva convention-like agreement, that has a chance at least at limiting what our adversaries might do with AI at the extremes. Dr. Mahnken, what do you believe are the most important non-negotiable norms that we would want to establish around such an agreement?

Mr. Mahnken. That is a wonderful question, Congressman, and I guess I will temper my answer with the belief that we are so far from that at this point. We are so far from that because of the uncertainty surrounding the unfolding of AI.

Mr. Moulton. Okay. We were limited on time, so if you don't have an answer _

Mr. Mahnken. I wish I did.

Mr. Moulton. _ let's just go to what should be our first steps. How do we start that? Clearly, it is going to be hard to regulate from behind, so I have to reemphasize the theme here _

Mr. Mahnken. Yeah.

Mr. Moulton. _ that we have to win this race, but how do we actually get to a point where we can have, as Mr. Clark has also said and Mr. Beall as well, some sort of norms that limit the extreme-edge cases where this can be so

dangerous?

Mr. Mahnken. I think the strongest approach is for the U.S. to lead and to lead, and to lead our allies, other like-minded countries towards a democratic approach.

Mr. Moulton. Mr. Beall, do you have any comments on this?

Mr. Beall. Yes, sir. I think I might offer there are three discreet areas where the United States must have clarity on what it wants. The first area is AI's impacts to strategic stability. So this means, as my friend Jack mentioned, if we have a country full of geniuses in a data center and they are inventing all manner of new ballistic missile capabilities or missile defense capabilities, this could alter the strategic situation around the world, and this could become very urgent. And like any other issue that is changing exponentially, we are either going to be too early, or we are going to be too late. The second thing we need to do is on lethal autonomy. The Department of Defense back in 2012 published DOD Directive 3000.09. It was the government's attempt to put some rules of the road on how lethal autonomous systems are developed and then deployed. I think in a world in which killer robots, candidly, are being deployed around the world willy-nilly, this is a world that that is a dystopian nightmare world. That is the second area. And then the third area, as I

mentioned in my testimony, is in this area called artificial superintelligence, so this is our theoretical technology that could come about in the next several years. That could mean that AIs are going to be smarter, as smart as we are towards snails as they are going to be to us.

Mr. Moulton. Thank you. I am over time, and so I thank the chairman, but, Mr. Clark, if you could take the answer to this for the record, I would appreciate it. Thank you. Thank you, Mr. Chairman.

[The information follows:]

Chairman Moolenaar. Thank you. Representative Johnson.

Mr. Johnson. Mr. Beall, you noted at the top that maybe we become numb by the headlines about all of the dangers of AI. I think that might be true, and yet, honestly, what we have heard today, I suspect, has scared the hell out of many of these committee members. Anybody who doesn't feel urgency around this issue is not paying attention. Sir, you noted that this is the most important national security challenge of our time. You noted that if we lose this race, it could trigger a global crisis. Okay. I get it. We have got to win. It seems as though it will take tremendous computing power to win. Those will come from data centers. I guess my question would be, are there risks if our computing power is located outside of the United States?

Mr. Beall. Yes, sir. That is a great question. I think there are sort of two issues at stake. The first is we currently have a significant energy shortfall in the United States in a regulatory environment that is not conducive to rapid upscaling that, and as a result, there could be a temptation to move capability in places that are more friendly to fast energy generation. And I think as you look around the corner to AI training systems, they will be trained in a distributed fashion, and this could sort of serve American interests in the first place. The second piece and the downside risk here is when you do that

overseas, you don't want to let foreign countries become the AI superpowers, and we don't want our chips being diverted in places that are not friendly in the United States. So on balance, I recommend this as probably a good course of action in the near term, but we should focus domestically on making sure we are building our domestic capability as well.

Mr. Johnson. So the physical location of these data centers does matter. You are saying if they are located elsewhere, that will empower their ability to become the AI experts.

Mr. Johnson. Yes, sir, it could definitely do that, and it is sort of the devil is in the details. Who is controlling these data centers? Who is controlling the chips? If they are under American hyperscaler control, that is one thing. If they are under local national control, that is another, and the security package that goes along with that will need to be carefully scrutinized.

Mr. Johnson. Yeah, thank you very much. Mr. Clark, I will reiterate. Mr. Beall mentioned that this is the most important national security challenge of our time. Mr. Mahnken noted that losing this competition, it becomes more likely if we slow our momentum. It seems to me that safety and speed are conflicting values. When you prioritize one, you get tradeoffs in the other. You noted that we have to get safety right, and I am curious how substantial is the

risk that by doing that, we inhibit our ability to win this race?

Mr. Clark. Thank you for this question. We all buy cars because we know that if they get dinged, we are not going to suffer in them because they have airbags and they have seat belts. You have grown the size of the car market by innovating on safety technology, and American firms compete on safety technology to sell to consumers. The same will be true of AI. So far, we do not see there being a tradeoff here. We see that making more reliable, trustworthy technology ultimately helps you grow the size of the market and grow the attractiveness of American platforms vis-a-vis China. So I would constructively sort of push back on this and put it to you that there is an amazing opportunity here to use safety as a way to grow the American existing dominance in the market.

Mr. Johnson. I want to believe you are right, and obviously you are an expert and I am not, and by the way, clearly we need to take care of safety. We need to get that right, and I don't want to suggest it is a zero sum game, but we do have a scarcity of resources here. We should have some urgency about winning this. Is there a concern, Mr. Mahnken, that prioritizing safety to too great a degree could inhibit our speed?

Mr. Mahnken. Absolutely. Absolutely, and that is why I

think we need to lean on our culture, which, as I say, we have a high trust society. We have deep democratic values that inform us. If I have a greater worry, it is we are just not going to go fast enough because we are going to tie ourselves in knots worrying about all the things that could happen. Again, we are having this conversation. My suspicion is there is no parallel conversation going on in the PRC about all the risks. And if I could, just to your previous point about data centers, I just want to make a basic fundamental point that these data centers, not only does it matter where they are, but this is big, valuable infrastructure, and if I think about past revolutions in information, these are also going to be targets. We have to consider that these data centers are going to be targeted in various ways as well. Thank you.

Mr. Johnson. Thank you very much, gentlemen. Mr. Chairman, I yield back.

Chairman Moolenaar. Thank you. Representative Torres.

Mr. Torres. Thank you. TSMC is the preeminent company in manufacturing leading-edge chips at scale, and ASML is the sole manufacturer of extreme ultraviolet lithography machines, which are critical to building advanced semiconductors. Mr. Clark, how close is China to replicating and rivaling those capabilities?

Mr. Clark. I will be relatively brief in this. They

are some ways behind of closing the gap on what you think of as the nanometer level they can reach where TSMC and ASML combined are in the lead, but they have made substantial investments here, but multiple years.

Mr. Torres. Multiple years. It should be the highest strategic priority of the United States to pursue AI dominance with the fierce urgency of the Manhattan Project. In the 20th century, the U.S. and Nazi Germany were locked in a high-stakes race to develop the first atomic bomb. In the 21st century, the United States and China are competing in a new strategic arms race _ the race for artificial superintelligence _ and the first country to reach ASI will likely emerge as the superpower of the 21st century. The superpower will set the rules for the rest of the world. Mr. Clark, what do you make of the Manhattan Project framing? Is that the right way to think about the AI strategic competition with China, or is that a false analogy?

Mr. Clark. There is an element of this, which is a dis-analogy because the frontier of AI was sort of borne by these private sector companies and is built today in the private sector, so we are not starting from the same places with the Manhattan projects. But where it did get ideas right, which we should carry through, is the prioritization of the core resources for it being in the United States. We

need energy here and we need compute here because as this technology becomes more sensitive, we will want to develop the most powerful systems here under the full U.S. security umbrella.

Mr. Torres. Mr. Clark, can we win the AI race without energy?

Mr. Clark. Energy is essential, and without it, we lose this race.

Mr. Torres. So the AI revolution requires an abundance of energy on a scale and at a pace that we have never seen before. China is emerging as the energy superpower of the world. When it comes to new energy capacity, China is adding the equivalent of a whole United States every 2 years. In 2024, China built 400-plus gigawatts of new capacity compared to only 50-plus gigawatts of new capacity for the United States. When it comes to energy, China is outbuilding the United States by a ratio of 8 to 1, and yet at a time when the AI revolution has put unprecedented strain on the U.S. energy grid, the present reconciliation bill, to be blunt, would all but repeal the tax credits for clean energy. According to an analysis by Politico, the loss of tax credits would endanger the nearly 800 planned clean energy projects that would generate over 156,000 megawatts of electricity, which is enough to power 27 million homes. These are staggering numbers. Mr. Clark, do

you think it is wise for the United States to artificially restrict the supply of energy available to American AI at a time when we are engaged in a high-stakes arms race with the Chinese Communist Party?

Mr. Clark. We estimate that we need 50 gigawatts of power by 2027. By "we," I mean by AI industry. That suggests to me that you need to look at all options on the table and look at what it takes to get there, and I can't speak to the specifics of individual power sources. I can say that _

Mr. Torres. Do you know what percentage of new capacity in America comes from clean energy?

Mr. Clark. I do not.

Mr. Torres. It is 90 percent, and solar is the most rapidly deployable, scalable energy source, and it just seems deeply self-destructive for the United States to restrict the number of electrons available to the AI revolution. I have a question about the export controls. I know you are supportive of it, Mr. Clark. Do you think the export controls were properly designed when put in place back in 2022 because I have heard concerns that even though the chips to which China had access had less computational power, it actually had more memory bandwidth than the best GPUs. So was that a fatal flaw in the original export controls?

Mr. Clark. Getting export controls right here requires us to have a greater level of technical staff in the Department of Commerce to both design these controls and also monitor them once in place. We are dealing with the most advanced technology that people make on the planet, and trying to control that will require us to, I believe, scale up for resourcing for the people designing those regulations and then monitoring compliance with them.

Mr. Torres. I see my time has expired.

Chairman Moolenaar. Thank you. Representative Hinson.

Mrs. Hinson. Thank you, Mr. Chairman, for holding this hearing and to our witnesses for appearing before us today to discuss a very critical issue. And the United States has obviously long led the world in cutting-edge AI research, driving the breakthroughs that really power not only national defense and advanced technologies here at home, but also help American farmers in places like Iowa to really increase their yields and allow small businesses to operate more efficiently, so we have seen some great advances there. They are transforming how we work, how we produce, and how we compete on the global stage, but we are obviously in this active race to stay ahead of China.

The state-directed labs there are simply not competing. They are replicating our U.S. innovations day after day and at an alarming pace. That is not healthy competition. That

is a deliberate effort by the CCP to steal and then weaponize our innovation against us, against U.S. industry to gain strategic advantage. If Beijing is able to close that capability gap with the United States, the consequences, as we know, for both our national and economic security would be severe. The Trump administration's repeal of the Biden-era AI diffusion rule signals, I think, a much more needed return to a more strategic protective posture, and then coupled with the ongoing efforts to craft a domestic-centered approach, this shift, I think, focuses also rightly on keeping our frontier AI leadership capabilities out of adversaries' hands and reinforcing U.S. leadership in this vital space for us.

So, Mr. Mahnken, I want to start with you. The Trump administration has made it clear that advancing and accessing cutting-edge U.S. AI should come with conditions, namely that partners decouple with China's tech sector and invest in U.S. AI resilience. So as China works to continue to weaponize AI and commercialize it, how can the U.S. work with allies to deny Beijing that access to critical inputs without undermining allied innovation? So pushing back against that fear while still making sure we are innovating in this space?

Mr. Mahnken. Yeah. Look, I think we need to provide a democratic alternative to the authoritarian approach that

the Chinese not only are producing but are exporting, right? Their technology not only is of authoritarian origin, but goes to promote authoritarianism. We need to have an alternative and be able to provide that to our allies, to our friends, and with that goes, again, greater sovereignty for our partners.

Mrs. Hinson. Well, when we look at this AI arms race and how important it is, obviously the subject of our hearing today, I think we can all agree that all of us in government and in the private sector are fully committed to the cause. At least in this room we are. So, Mr. Clark, for large companies like yours, obviously Silicon Valley, kind of the brain center for a lot of this development, and I know there are a wide range of opinions there as well. And we know that several members of your key team at Anthropic have held very influential roles in this space, both open philanthropy and, in the previous administration, with the Biden administration as well. Can you speak to how you manage, obviously we have got a lot of viewpoints, but how you manage potential areas of conflict of interest in advancing this tech and ensuring that everybody is really on that same page with helping to shape this national AI policy that we are talking about, the competition on the global stage for this technology?

Mr. Clark. Thank you for the question. We have a

simple goal _ win the race and make technology that can be relied on _ and all of the work that we do at our company starts from looking at that and then just trying to work out the best way to get there. And we work with people from a variety of backgrounds and skills, and our goal is to just have the best, most substantive answer that we can bring to hearings like this.

Mrs. Hinson. And then, Mr. Beall, my last question will be for you. When we look at China's DeepSeek model, they demonstrated quickly how PRC labs can really replicate our U.S. breakthroughs, right? They are copying our technology. I guess mimicry is the best form of flattery, right, but what are the biggest espionage or model leakage threats that are facing our domestic partners today?

Mr. Beall. Thank you, Representative Hinson, and first of all, I would like to thank you for using the word "weaponize" instead of "misuse." I think it is very clear that we are strong in our language, that people know exactly what we are talking about, and we understand the stakes. I think DeepSeek absolutely is our Sputnik moment, and we have to remember that export controls, there is always going to be a lagging effect. So we implement ship controls in 2022, but it took us a long time to actually enforce them, and as a result, the ships sort of still flowed to China. And so when we think about the urgency of action today, we have to

assume that looking right around the beat, it is going to take some time to deny these capabilities.

So there is sort of two core things that we need to do urgently. The first, there is a piece of legislation introduced last Congress called the Remote Access Security Act. That blocks China from accessing restricted AI chips but just via cloud services, and so this is a loophole that is being taken advantage of now. The second thing, and this is going to be the really hard one, is grappling with what the folks in the industry call the open source issue, and that is, at what level of capability is it going to be not okay to publish weaponizable AI systems out in the open? And when a company publishes these things out in the open and we have chip controls in place, well, the Chinese just download those model weights and then use their existing chips to scale up that capability, so these things are working at cross purposes. So if you had to do two things urgently, I would recommend considering those two things. Thank you.

Mrs. Hinson. Responsibility in this space matters. Thank you. I yield back, Mr. Chair.

Chairman Moolenaar. Thank you. Representative Brown.

Ms. Brown. Thank you, Mr. Chairman. As we continue to assess the scope and stakes of our strategic competition with the Chinese Communist Party, I want to focus on a

critical aspect of American competitiveness that often gets overlooked _ our workforce _ because winning the AI race isn't just about algorithms and semiconductors. It is also about ensuring American workers and communities aren't left behind as the economy changes rapidly.

For generations, places like Ohio's 11th and communities across Northeast Ohio powered our economy through steel, manufacturing, healthcare, and innovation, but they have also experienced economic disruptions caused by automation, offshoring, and underinvestment in workforce training. We cannot afford to repeat those mistakes with AI. AI is already being rapidly deployed across industries, and while that holds enormous promise, it also raises some serious questions: who benefits, who gets left behind, and what are we doing to make sure the American workforce is prepared?

So, Mr. Clark, Anthropic's Economic Impact Index suggests that AI is already being significantly adopted to automate and augment workflows across a wide range of industries. Can you explain what specific economic impacts you are seeing from AI adoption and which sectors and communities are most at risk of disruption?

Mr. Clark. Thank you. We look at the use of AI on our platform and compare it to the O*NET job classification, which helps us look at the distribution of jobs on our platform relative to the U.S. economy. What we see is AI

technology today is being used in a major way for jobs relating to programming and also relating to what you might think of as bureaucracy or paperwork that takes place in large-scale businesses. Today, this technology makes people much more productive. It predominantly augments them, takes a person and lets them do more, but in the future we think that there will be other impacts as well as this technology matures. We believe that there need to be greater degrees of data being shared by the companies about the economic impacts of their systems, and from this, we take inspiration from the U.S. census where the annual survey of manufacturers now asks questions about the number of industrial robot arms, which are bought by manufacturers across America each year. I think we can do equivalent things here to find ways to get better information about AI and more granular data, and from that data, we will be able to see where it is impacting the economy and actions we can take.

Ms. Brown. Thank you, and what partnerships between industry, government, and educational institutions do you think are most critical to ensuring our workforce is not only resilient but positioned to lead, and what examples or models are showing the most promise in helping workers transition into AI-aligned roles, particularly in underserved or industrial communities like the one I

represent in Ohio?

Mr. Clark. The most success we have seen so far has come from experimentation across industry and people that want to learn AI skills. The Department of Energy has done a variety of AI Jam Days, which Anthropic and other frontier developers have participated in, and I think that is a scalable model: find communities that want to learn about AI, give them access to the technology, and have industry come there and embed deeply with them to give them that familiarity. Once people are familiar with this, they can learn to use the tool, and they will find many ways that it can change how they work.

Ms. Brown. Okay. Thank you, and then lastly, Mr. Clark, as AI transforms the job market, what specific steps should Congress take to prepare the next generation of American workers for an AI-driven economy? Specifically, how do we ensure that American workers benefit from this revolution in productivity?

Mr. Clark. It all starts with experimentation and being given the space to familiarize yourself with a technology, experiment with it, and find ways to apply it. As part of this, we should closely look at which industries are using the technology most, perhaps using technologies like the Economic Index, and where industries that aren't using it, whether there are impediments that stand in the way of using

it because it will change and improve how jobs work there, and there may be avenues to learn of ways that we can change regulations to make it easier for people to use this technology and benefit from it.

Ms. Brown. Thank you, and I will close with this. Just as we are using export controls to slow China's AI ambitions, we should be equally aggressive about accelerating our investment in the American workforce, which is our greatest competitive advantage. That means scaling apprenticeship programs, embedding AI skills in community colleges and HBCUs, and modernizing career and technical education to prepare workers, not just to use AI, but to help shape it. And with that, Mr. Chairman, I yield back.

Chairman Moolenaar. Thank you. Representative Nunn.

Mr. Nunn. Well, thank you, Mr. Chairman, for having us here today and calling, I think, one of the most important hearings that we have had on the topic of artificial intelligence, and I would agree that AI is the new Cold War between the U.S. and China. As this distinguished panel has highlighted, it is, in fact, the Manhattan Project of our generation, and we have the opportunity to stand forward and do this not just as a whole-of-government approach, but as a whole-of-Nation approach to be able to stand up to China.

I want to give a particular salute to our Secretary of the Army who has brought forward four members of our top AI

teams, commissioned them as lieutenant colonels, and also highlighted that this is a public-private partnership, not something that the American Government nor the American taxpayer can do alone. In fact, with President Trump's Stargate, we are putting $500 billion to the advancement of AI, and that is truly something that I think should be saluted. We should also recognize, though, the Chinese and the CCP specifically, Chairman, as you have highlighted, are not standing still in this space. The Chinese just this year have tested something called AI Commander. It is capable of generating 10,000 battle plans in under a minute. If you want to have an invasion of Taiwan, this is the type of tool that you want.

Additionally, perhaps one of the more concerning items is a new AI group out of Beijing called Zhipu. Zhipu is an AI anomaly that is now facing off against the likes of OpenAI, and their entire intent is to lock in Chinese systems and standards into emerging markets before the West, so this is clearly a large-scale attempt by the Chinese to box the United States out. Now, as a counterintelligence officer who was on the frontline in fighting against Huawei's takeover of the United States through something called Huawei America, very clever by the Chinese, we were able to stop that from taking root here in the United States. Our concern here is that we are now facing this

again, that the Chinese are using AI and their ability to get into first-shooter markets ahead of the United States to establish a de facto leadership role in this space. So, gentlemen, as we all know, the race is on.

I want to thank you for not only your time here, but your leadership in being able to go after that. Tom, you are the president and chief executive for the Center for Strategic and Budgetary Assessment. Given the examples that we just highlighted here, one can only assume that the CCP is gearing up for a larger AI-enabled operation. I would like to ask you, is the U.S. currently prepared for an AI-accelerated cyberattack, a zero-day attack, or a larger threat that faces us today?

Mr. Mahnken. Well, thank you, Congressman. You justifiably highlight some of the risks that we face, and I want to return to the uncertainty that surrounds this situation. There is a huge bet here, right? The Communist Party is betting and the PLA is betting that AI will make them better, right, that it will, with scientific clarity, lead to, as you put it, the optimum answer for how to crack Taiwan. I think we as a society are betting differently, right? We are betting that AI tools can help human beings make more informed, better-informed judgments to recognize signals from the noise. My primary concern is not that the Chinese authoritarians are right. My primary concern is

that they may delude themselves into thinking that AI will allow them to do things that they won't, in fact, be able to do. It might be a great way for them to not win a war but get into a war.

Mr. Nunn. Mr. Mahnken, I would agree with you on that. I would say the same challenge is true for the United States, that we also need to be clear-eyed about what our expectations are for AI. In fact, it is one of the reasons that I have introduced H.R. 2152, the AI PLAN. It is an act to make sure that we are implementing a strategy here within the U.S. Government that we are using AI in a tandem way, not that every department or agency, or as highlighted, earlier, every State or every local entity is running on their own. This has to be a whole-Nation approach.

Mr. Clark, you are the co-founder and head of policy at Anthropic, a very impressive private entity who has been able to come to this. I would like to ask you in our remaining time how the government and private sector can work cooperatively, not only to advance AI here, but to also disrupt, deter, and compete with what the Chinese have going on around the rest of the world.

Mr. Clark. Thank you. Two quick things. One, to advance AI here, we need the U.S. Government to lean in on AI deployment and to uplevel the U.S. Government of both civil and the IC side. That is an amazing opportunity, and

we are here to partner for that. On the international piece, the development with the U.S. Government of the standards by which we assess for safety and reliability of this technology, in partnership with industry will help industry go out and sell globally and sell people on American standards that let them know they can trust our technology, and they will choose to buy American off of that trust.

Mr. Nunn. Let's make sure we can make it affordable, deployable, and effective for them as well. Thank you, Mr. Chair. I yield back the remainder of my time.

Chairman Moolenaar. Thank you. Representative Tokuda.

Ms. Tokuda. Thank you, Mr. Chair. I want to build on what some of my colleagues have touched upon, in particular, in the energy field. We know that China is investing significant resources and assets to achieve its 2030 AI world leader goal, and meanwhile here at home, we say we want to win the AI race, but the current administration proposed gutting the very institutions that both give us an edge and protect us from intrusion. NIST was facing up to a 20-percent workforce cut in over $325 million in funding reductions, and while reconciliation bills have softened the cut blow, CISA was facing nearly half a billion dollars in cuts and the elimination of a third of its workforce, 1,000 positions. These are in critical areas like critical

infrastructure, communications, data privacy, cybersecurity support at the State and local levels, all of these things. Perhaps, Mr. Clark, we can ask you that question since your company works closely with NIST and CISA. Are you concerned about our ability to both compete and protect? Briefly, should we be cutting or, as I would argue, actually increasing funding towards these critical areas?

Mr. Clark. There are amazing complements to the work of the private sector, and that includes work on standards and measurement that can be done by NIST. Ultimately, it will also require us to help use our technology to protect critical infrastructure and harden our defensive posture here that will be used by other agencies as well. I would encourage us to look closely at what parts of government can complement and accelerate the work of private industry to help prepare us for the very powerful AI systems that are developed in the coming years.

Ms. Tokuda. Okay. So I would take that as definitely we should not be looking at massive cuts. We should actually be looking at supports that are complementary to the private sector. On that note, Mr. Clark, Anthropic CEO, Dario Amodei, has publicly fielded the idea of a token tax where a small percentage of language model revenues could go to the government to help offset the economic disruptions AI could cause, like the potential, as was brought up by Ms.

Brown, loss of 10 to 20 percent of white-collar jobs. To me, it reflects a bigger principle that AI companies have a responsibility to help address the societal costs their technologies could bring. And in your testimony, to meet the moment of the AI race we are in, you call for major investments in public infrastructure, energy, as we have heard discussed today, power grids. These are critical systems that AI companies are rapidly placing new demands on, critical demands on, and in many ways, we are outpacing the need at this time. So what is your view on responsible public-private burden sharing? Should companies like Anthropic be expected to directly invest in the infrastructure that they rely on so that we can actually increase our capacity, as well reinforcing the environmental and public systems that are being stretched to support the growth of AI?

Mr. Clark. Today, whenever we develop infrastructure, we work closely with our partners and the communities where that infrastructure is built to understand exactly how we can lean in and do more, and we are constantly looking for ways to be better. At the same time, if we are right, in several years and truly powerful systems get built but have the property of a country of geniuses and a data center, it would behoove us to take a look with fresh eyes at this technology and what societal changes it may be causing. And

I would be eager to follow up with people here in that conversation, which we will need to be ready to have in 3 or 4 years or so.

Ms. Tokuda. Okay. Definitely I think we know here in Congress the demand for resources will definitely be great as the needs increase as we go further in this area. I want to touch upon something that is a bit existential, if you will. Mr. Beall, your testimony makes clear that artificial superintelligence _ ASI _ is one of the largest existential threats that we face right now, and many experts in the field also recognize that AGI _ artificial general intelligence _ is a precursor to that risk. And even as we race towards AGI, we are still struggling, quite frankly, to manage responsibly and reliably today's systems: models that hallucinate exhibit bias and behave unpredictably. And these are likely the simplest of systems that we will ever have to deal with in our lifetimes. And as much as the AI race is about speed and efficiency, it is a race for control and governance because at the end of the day, true power lives in the ability to wield AI safely and effectively with strong oversight and accountability to ensure it serves its intended purpose.

This might be a question you have to insert in the record as I will run out of time, but should we also be concerned that authoritarian states like China or Russia may

lose control over their own advanced systems? Could that kind of instability in their own AI systems pose global consequences regardless of whether we get our own houses into order? And is it possible that a loss of control by any nation-state, including our own, could give rise to an independent AGI or ASI actor that, globally, we will need to contend with? And so I know Mr. Chairman out of time, and so I would request the answer from the panelists for the record. Thank you. I yield back.

Chairman Moolenaar. Thank you very much and approve that request.

[The information follows:]

Chairman Moolenaar. Representative Moran.

Mr. Moran. Thank you, Mr. Chairman. Thank you to the panel for being here today. I have got lots to ask. I will try to move quickly through it, but what you are doing here today is super important to help inform us as members of Congress, but also the public at large about the real threat that the CCP poses if we do not win this AI race. I am convinced that one of the things that is so great about America is that our innovators are better than anybody, and if we will just get out of the way and let them innovate, I think we can certainly win this. But oftentimes we get in our own way, and then we don't realize what we need to be doing outside of our borders to push back against those that would do harm to the United States.

As I learn more and as I listened to you guys today, Mr. Clark, you said we need to win the race. That was one of the things you said was absolutely most important. I agree with you, and so I want to ask you, Mr. Beall said a couple things specifically we could do to win the race is to pass the Remote Access Security Act, and then also grapple with this open source problem of letting the CCP get information that it shouldn't through those open-source documents, but, Mr. Clark, let's go back to winning the race. What is the number one thing that we are not doing that we should be doing to win this race?

Mr. Clark. Just making me choose one means I have to pause for a moment. I think the fundamental thing is power, and then if I was allowed to add two more, I would say power and then compute, ensuring you have enough compute resources. And then the third thing would be ensuring we have the necessary infrastructure in the U.S. Government to help us have confidence about developing the standards for this technology and the means by which we assure its safety and security.

Mr. Moran. Yeah, I agree with you about the power and the grid itself. I just had a meeting this morning, actually, with a number of energy companies in diverse ways of providing energy. So without respect to actually how we provide the energy, the necessity to get more load on the grid is an absolute. Otherwise, we can't have the computing power that we need to win the AI race through these facilities we need to build. I will let you guys all answer. Do you have a specific suggestion as to how we should go about doing that? I know there is some partisan differences there, but where can we find consensus on how to drive forward with getting more load on the grid?

Mr. Clark. We would be happy to follow up for a more detailed private conversation, but our number one idea is just work backwards from the goal of having 50 gigawatts of net new capacity for use by Frontier AI in 2027 and figure

out how well postured the U.S. is for that, identify any blockers that exist and find a path through.

Mr. Moran. Dr. Mahnken, let me come back to you with that first question. What is it that we are not doing that you would say we absolutely must do?

Mr. Mahnken. Well, I certainly think that the private sector is moving forward. Mr. Clark talked about what can be done there. I think government has an important role to try to slow down the Chinese. I think being more mindful of that, and I think, again, there are some measures that have already been passed, others that are under consideration that can help do that, being mindful of what the Chinese are doing and slow them down.

Mr. Moran. Yeah. One of the things that I am learning more about is how AI works behind the scenes. As we are going through this process, the development of our current AI model certainly is different than when software engineers developed software a long time ago. Instead of a programmer writing each rule, a system will follow, the system itself effectively writes the rules based on the question it needs to answer or is trying to answer. It is my understanding that AI systems will soon have the capability to conduct its own research and development, and effectively write its own rules and programming. AI will be able to predict then the issues it will need to solve and do it in ways sometimes

that we can't control. Mr. Clark, how do you think the U.S. can best harness automated AI research and development while ensuring the Chinese AI systems do not generate the same capabilities? And I will come to you, Mr. Beall, as well for that same question.

Mr. Clark. This is a fundamental opportunity and challenge. We need to urgently resource our intelligence community to understand how advanced Chinese AI systems are and whether they are capable of AI R&D. At the same time, we need to work closely with industry in the U.S. to understand AI R&D and its potential risks. Just to illustrate, you wouldn't want an AI system that very occasionally tries to blackmail you to design its own successor, so you need to work on the safety issues of AI R&D, or else, you will lose the race.

Mr. Moran. And, Mr. Beall, as you answer this question _ I am running out of time _ I also want to know where is the red line. You guys have talked about that a little bit today, but I want to get a definite understanding of where is the red line where we cannot allow the Chinese to crossover?

Mr. Beall. This is an excellent question, sir. I think the first thing I would say is that AI systems are not really built. They are grown. There is no science here. It is alchemy, so there is lots of fuzziness about how these

systems work. We can't really break them open and reason about all these details. You touched on a very critical point, though. This is the automated R&D part, so once you have an AGI-level system that could take control of its own destiny and build itself and build its successors, to me. that is the very clear red line in which the danger starts. And that is where we have to get the government involved in testing and evaluation and get the upskilling of the U.S. Government so they can understand what is happening and make good recommendations to you. Thank you.

Mr. Moran. I absolutely agree, and I think that that is a critical red line that hearkens back to the days of when we established ourselves as a superpower. In a number of different ways, I do think there is a resetting. We have to win this race, as was stated earlier. Thank you, gentlemen. I yield back.

Chairman Moolenaar. Thank you. Representative Khanna.

Mr. Khanna. Thank you, Mr. Chairman. Mr. Clark, I appreciate Anthropic's thoughtful approach to AI safety and jobs. You mentioned the Standards and Innovation Institute under NIST. As you know, it is currently voluntary. I wonder whether you think that, for certain high-risk AI applications, there should be some minimum threshold of mandatory third-party verification, even if that third party verification is just verifying the company's own standards

or some threshold that we need to move towards.

Mr. Clark. This question illustrates the challenge we have about weighing safety versus moving ahead as quickly as possible. We need to first figure out what we want to hold to that standard of testing. Today the voluntary agreements rest on CBRN testing and some forms of cyberattack testing. Once we have standards that we are confident of, I think you can take a look at the question of whether voluntary is sufficient or you need something else, but my sense is it is too early, and we first need to design those tests and really agree on those before figuring out what the next step would be.

Mr. Khanna. And who would design those tests? Is it the AI Institute, or is it the private sector? Who comes up with what those tests should be?

Mr. Clark. Today these tests are done highly collaboratively between U.S. private sector, the CAISI, which you mentioned, and parts of the U.S. Government, including those in the intelligence and defense community. I think bringing those people together so that we have the Nation's best experts on this and standards and tests that we all agree on is the first step that we can take to get us to everything else.

Mr. Khanna. And by when do you think that needs to be done?

Mr. Clark. It would be ideal to have this within a year. The timelines that I have spoken about in this hearing are powerful AI arrives at the end of 2026 or early 2027. Before then, we would ideally have standard tests for the national security properties that we deeply care about.

Mr. Khanna. One of the things I am most concerned about is the jobs situation. Now, obviously AI is going to create a lot of new jobs, but you look at college unemployment today. Between the ages of 21 and 29, for those who have a college degree, it is 15 percent. It is not partisan. It was at 15 percent when Biden was President, it is 15 percent now, and there are projections that that will increase as entry-level jobs, in particular, are affected.

I have studied Darren Acemoglu's work at MIT, who says that we need to revamp the Tax Code so that we are incentivizing hiring people instead of robots. Some other people have said, look, we need something like a future workforce administration like Roosevelt had, but for young people or people to get entry-level positions in doing work on making government services better, in healthcare, in infrastructure. What are your recommendations on how we should be thinking about this jobs issue?

Mr. Clark. Our specific recommendation is to start with data and to ask for more data from the AI companies, including Anthropic, about how we see the distribution of

jobs being used by our AI systems today. With that data, we can have a conversation about where it is augmenting people and making them more productive, and where are parts of the economy where AI has almost no relevance today, but there is huge societal benefit in those jobs, and we can think about how to better support those.

Mr. Khanna. Any other ideas? Data is fine, but anything else, or do you think that that is where we start?

Mr. Clark. I think that we have to start with data, and that will show us what we should do after that.

Mr. Khanna. Right. Anyone else on the panel have any ideas?

Mr. Beall. Yeah, I will take a stab here, sir. I have a 16-year-old son, and we talk a lot about what do you study in college these days, and if you look at the stated goals of many of these companies, they want to have AI replace all humans at all jobs. It is what they say publicly. Bill Gates said this publicly. So I worry about a future in which human beings are not just unemployed, but they are unemployable, and this breaks the notion of the free market in very important ways. And I think we are going to have to have more of these conversations, and it is going to feel too early until it is too late. And when I hear folks in industry claim things about universal basic income and this sort of digital utopia, I study history, and I am worried

that that sort of leads to one place, and that place is the gulag.

Mr. Khanna. Appreciate it. Thank you.

Chairman Moolenaar. Thank you very much. Well, I want to thank all of our witnesses for your testimony today and your insights.

Questions for the record are due 1 week from today.

[The information follows:]

Chairman Moolenaar. And without objection, the committee hearing is adjourned.

[Whereupon, at 10:51 a.m., the select committee was adjourned.]