



***The Beijing Brief* podcast
The Brookings Institution**

“Are the US and China really in an AI ‘race’?”

Tuesday, June 23, 2026

Participants:

Kyle Chan
Fellow, Foreign Policy, John L. Thornton China Center
The Brookings Institution

R. David Edelman
Director, Project on Technology, the Economy, & National Security
Massachusetts Institute of Technology
Nonresident Senior Fellow, John L. Thornton China Center, The Brookings
Institution

Ryan Hass
Director, John L. Thornton China Center
Senior Fellow, Foreign Policy, Center for Asia Policy Studies, John L.
Thornton China Center; Chen-Fu and Cecilia Yen Koo Chair in Taiwan
Studies
The Brookings Institution

Episode Summary:

How are the United States and China pursuing AI development? On this episode of *The Beijing Brief*, Ryan Hass sits down with Kyle Chan and R. David Edelman to discuss the two countries differing AI strategies and assess why the two countries have taken a different approach to their AI ambitions.

PRESIDENT TRUMP: *AI is a big deal. AI is going to be maybe the biggest thing, bigger than the internet, bigger than anything else. We are leading China by a lot.*

[music]

HASS: Hello, you're listening to *The Beijing Brief* from the John L. Thornton China Center at Brookings, part of the Brookings Podcast Network. I'm Ryan Hass, director of the China Center, and *The Beijing Brief* is a biweekly podcast focused on unpacking the forces shaping the U.S.-China relationship and China's political, economic, and technological ambitions.

Today, our topic is the role of artificial intelligence in the U.S.-China relationship. This is a topic that members of the Trump administration have spoken about frequently over recent weeks and months. I want to play you a few clips from Secretary Bessent and Secretary Burgum.

SECRETARY BESSENT: *The U.S. is the AI leader in the world. We're an AI superpower. China is second. They are trailing substantially. We want to make sure that we keep that lead.*

SECRETARY BURGUM: *The U.S. cannot lose to China in the AI arms race.*

HASS: The Trump administration often talks about artificial intelligence in the context of a race with China. They use this racing framing often to explain their massive drive to increase energy production and limit regulation on AI companies. But there are key differences in how America and China are approaching artificial intelligence, differences in their strategies as well as their goals.

So today, we want to discuss, is racing the right framework for understanding how the United States and China are pursuing AI development? How and why are the two countries pursuing their AI ambitions differently? And in the national security domain, where are the United States and China in consensus and dissensus on the role of artificial intelligence?

With me today are two of the smartest minds I know on these topics, R. David Edelman and Kyle Chan. David is a nonresident senior fellow at Brookings and the MIT Center for International Studies. We worked together closely in the Obama administration, where he was a special assistant to the president and served as the chief cyber diplomat for the Obama administration. He also worked on a range of issues in AI, data privacy, and technology trade, 5G, and more while he was in the National Economic Council.

David, it is just wonderful to have you with us today.

EDELMAN: Thanks for having me.

HASS: And by now you are also acquainted, I'm sure, with Kyle Chan, our fellow and tech expert resident here in the China Center.

CHAN: Great to be back.

[2:19]

HASS: Kyle, we've heard previously from you about your origin story for getting involved in the technology scene related to China, but we have not heard from David. So David, I'd love to get your take. When did you become interested in China tech issues?

[2:32]

EDELMAN: I suppose my story is one of a irrepressible geek who ended up in national security and foreign policy as as a profession. You know, probably most seriously, when I was in graduate school, I was very interested in this question of if one country shuts down another country's internet, is that an act of war? This turns out was the wrong question to ask, but it was very interesting at the time.

And looking at this, it was at that point hypothetical, I mean, this is the early 2000s. There were really only a short list of countries that fell onto that list potentially. Right? Obviously the U.S., Russia, China, maybe North Korea if you want to take a flyer on it. And so this question sort of animated this intersection that I started working on of those two areas, of technology and foreign policy.

Then I was recruited out of grad school to go to the State Department, and at every stage of my career, it ended up being a U.S.-China conversation or a U.S. international conversation shaped by China, whether that was UN negotiations on international cybersecurity, whether it was, as you said, questions of 5G and domestically, Economic Council, what should U.S. national strategy look like? Even something as basic as rural broadband ended up having a China element to it.

And so from the hardest side of national security over to how we're thinking about the future of technical standards, each of those pieces ended up having elements in the U.S.-China relationship.

And so since then, it continues to be among the most interesting areas that I've been very privileged to have people a lot smarter than me working on and helping to teach me about those elements over the years as well.

[4:00]

HASS: Well, we're glad that that you're involved in these conversations. Let me just ask both of you to start out. Is an AI race the right way of framing what's going on between the United States and China today as it relates to AI? Just in a short, one or two sentence response. Kyle, why don't we start with you?

CHAN: No, I think the idea of an AI race flattens everything and oversimplifies the dynamic

HASS: David?

EDELMAN: I, I agree with Kyle. The idea of a U.S.-China race is a quick quip to try to essentialize things on both sides. The the reality here is obviously much more textured. There is a political and economic imperative in China that revolves around a race. In the U.S., race strikes me a little bit more as a rationalization of policy, sort

of a post facto, particularly in a world where the domestic politics of AI become more complex by the day. And the U.S., economically, politically, does very, very well when there is a competitive factor involved. It was true certainly in the days of the Soviet Union. It's true in the context of China now.

Is that the motivating element driving the majority of U.S. policy? No, I wouldn't say that it is. I think we over-index on it in the U.S. because it's simple when the reality is much more economically complex than I think that quick quip gives us.

But it is true, there are those on both sides that absolutely see this as a race and race as the animating element of whatever they happen to be focused on at that moment, whether it's the technology, the hardware of it, the economics of it, or even just the bilateral dynamics.

HASS: So it's instrumentalized, in a sense, to to advance certain objectives?

EDELMAN: I think it's instrumentalized on both sides, and I think it is more instrumentalized on the U.S. and more baked into the dynamic in China.

HASS: Right.

[5:38]

Well, let's talk about the the difference between the United States' and China's approaches. Kyle, we've heard you say that there's a difference between China's pursuit of integration of AI compared to the United States' push for artificial general intelligence, or AGI. First of all, what is AGI? And beyond that, what is the difference between how the two countries are approaching these different AI strategies? What are their end goals?

[6:00]

CHAN: In the U.S., there is a very, very strong focus, if not obsession, with AGI, or artificial general intelligence. This idea that we will develop AI systems that are capable of basically anything that humans can do, at least when it comes to cognitive tasks.

There's also this idea of artificial super intelligence, that these systems will exceed human capabilities. And there is is a broader race dynamic within the U.S. among the labs, among the AI developers, to try to get there first.

And this is what's driving the massive compute spend that we see from the American tech companies. Next year, for example, estimates are that we will have collectively a trillion dollars or more of spending on new data centers and compute infrastructure across the United States. And this is also animating the framing of the U.S.-China bilateral relationship.

On the Chinese side, they see it differently. One way I would put it is in China, they are AI-pilled but not AGI-pilled, and by that I mean they take AI very seriously. They see this as a powerful transformative technology, and their goal is to use AI to help turbocharge their broader economy and other parts of their society. They want to integrate AI into manufacturing, education, health care, research and development,

biotech, especially drug discovery, government services. They want to see AI everywhere all the time.

And much of this is a parallel to their previous approach to the internet and even the rise of computers, where they see this as important in and of itself, but also something that is all the more powerful if it can be diffused into society.

So here I see a contrast where in the U.S. there really is this almost sort of ideological focus on AGI, and in China, while some of the Chinese AI founders will talk about AGI to some extent, for the most part, that's not really their overriding goal, and they're much more focused on trying to integrate this technology into everything from applications to services to physical, real world use cases.

[8:16]

EDELMAN: We're not gonna get some big memo once we reach AGI. There's not going to be some switch that flips and we'll all be aware of it, and then that's great and wonderful. I mean, there's an argument to be made that we have reached a form of AGI right now, and I think that is incredibly complexifying Kyle's point about what this race is towards.

So, you know, frontier labs in the U.S. are very focused, obsessed, if that's the right word, on this idea of AGI and ASI. And if you look at the sort of founding documents of some of these companies, this is the *raison d'être* for why they are doing what they're doing, is to develop or safely develop, or supervise the development of, this artificial general intelligence, artificial super intelligence technology.

The U.S. labs are chasing the prize. The Chinese government, in partnership with the labs there, are actually chasing the application. And, you know, Kyle, you've written, I think, very eloquently about this. While we are looking for benchmarks, they are wiring the factory floor. And we have to be very aware of this because it's the diffusion of AI where the power actually lies. Right? You're gonna get 100 on every benchmark. That's amazing. You'll get a Turing Award for it. Fantastic. Is it locked in a lab somewhere? Is it usable?

And critically, and I think this is what this conversation is lacking a lot, will companies pay a lot to use it? I think one of the things that's interesting both, you know, dynamically and technologically about what's happening in China right now is not just the focus on near peer capability. Right? Can the model perform at 95 of the 100 of the U.S. frontier models? But, and we'll talk more about this, this open source paradigm, or at least open weight paradigm, where you have these models that can be run separately, you can actually look at them and see how they're put together, you know, they also run a lot more cheaply.

And that is a strategy we've seen before. This was the Chinese playbook with wireless. This is exactly what you saw in a different context when it came to 5G. How did we lose the 5G race? Well, it was a combination of standardization, manufacturing, and ultimately a lack of major investment at the time that mattered.

Now, we did a lot of major investment in 5G. I want to be clear. The U.S. rolled out 5G in a very meaningful way. And yet the national security folks were stuck, and I was in these rooms, you know, stuck with the despair of, well, we ultimately rely on hardware that was made elsewhere.

And so when we think about this right now, it's important to recognize, you know, are U.S. frontier labs at the frontier at this moment on that race towards AGI, ASI? Yeah, I would argue they are. How much of a lead do they have? Oh, I don't know. It depends on who you ask. Three months, six months, 18 months.

At the end of the day, the question to me at a national strategy level is not who's going to get the blue ribbon first. It's a question of who is going to be able to diffuse that across their economy and, we'll probably talk about later, in some very critical but narrow national security applications that are pretty meaningful too.

[11:07]

HASS: Now, David, you just touched on open models versus closed models. I want to dwell on this for a moment. Chinese tech companies like DeepSeek and Alibaba are developing open-weight models to capture global market share, while U.S. frontier models are, you know, sort of chasing the frontier of the most advanced cutting-edge levels of innovation.

So in terms that my mother would understand on Orcas Island in Washington State, what are the key differences between these approaches, open versus closed models, and what are the pros and the cons of these different approaches?

[11:35]

EDELMAN: Sure. I mean, the the the easiest way to put it, I think, is a closed model is going to a restaurant, and an open model you get the recipe, basically a way to think about it. Right? So a closed model, the company controls what the model weights are, controls the process, the inputs, and the outputs, which is actually much more complex than just a single model. Right? Where is it being routed? How is it processed? And then ultimately, they control and do not publicize what the underlying data sets were to train it, what the model weights, which is extremely technical and interpretable by almost no one, but what the model weights are, how it gets from A to B, and then how that output is wrapped.

You get the answer you get, and you're able to have some visibility into it the more technical you are willing to be when you tap into these systems of the proprietary company models.

But with open source, it's in many cases the opposite. And and to be clear, not all models that say they are open are truly open in the way we think of open source software. Is it the weights are publicized? Is it truly open? You have visibility into the data set that went into it, the pre-training process. There's a lot of steps here.

But, the bottom line is, one, you actually understand and can replicate how it's run, and then ultimately, in some cases, can run it locally on your own machine, and this is a big difference. Right? You know, when it comes to an open weight model, you

can replicate and then, in some cases, literally run it on on your Mac, or if you have an NVIDIA cluster that you paid tens of thousands or hundreds of thousands of dollars for, you could run it on that. You can't do that with the latest model from, say, Anthropic, because they are in charge of the model and how it runs and under what conditions you are able to get access to it.

And there are advantages and disadvantages to both. Right? You know, there's a lot of anxiety, you know, I, I... The main conversation I heard at the Munich Security Conference, which you were at as well, you know, really revolved around European powers with great anxiety over this idea of the dependence that they would have on American firms and American technology stack, which are sort of interchangeable. And the concern really revolved around, well, we've seen the U.S., just as we've seen China, weaponize aspects of its technological advantage.

[13:38]

Now, I don't want to draw a parallelism between the two; they've been weaponized in very different ways. But the perception, particularly in Europe, is these are strategic liabilities. And so when you think about an open versus a closed model, the concern is, well, if we have all of our companies pointing at the API, right, at the computer level interaction of one of the U.S. companies, well, what if they shut us off? What if they export control the thing that our companies are dependent on? This creates a liability.

On the other side, you know, what you hear from, I don't want to pick on Europe, but particularly European powers right now and middle powers, is this concern that, well, okay, but do we really want to be reliant on Chinese infrastructure again? Yes, perhaps we could replicate it, but what are the underlying security concerns? What are the risks that we aren't aware of, that we aren't able to track? Obviously, China's security regime speaks for itself in terms of its willingness to ingest and retain Western information.

These are hard trade-offs on both sides. And all of it at this moment, particularly even for U.S. companies, is coming down critically to cost. And there are a lot of companies that will not publicize this right now, but they will, you know, convey privately to me and to others, written about this to some extent, that are they uncomfortable running Chinese open models? Yes. Are they gonna do it because it costs 5% of their U.S. counterparts, and they aren't doing something so wildly complex that it requires AGI? Also yes. And that's a very difficult position to be in right now.

[15:00]

HASS: Right. Well, let's let's talk a little bit about the layers that go into the development of AI. The first layer that I want to touch on, I'm going to turn to you, David, is energy. AI requires an enormous amount of energy to power the technology. How are the two sides approaching their energy demands?

[15:18]

EDELMAN: With great hesitation. Look, the the U.S. AI picture has turned into a U.S. energy picture, and it's, it's inescapable. I mean, nowadays, we are measuring the capability of companies and their data centers not by the number of chips in the data center, but by literally the amount of power going to it. That is the measurement now. Is it a, you know, 50 megawatt, 100 megawatt, gigawatt data center?

And, you know, when you're looking at the domestic politics of AI, and we could just spend a little bit on this, the domestic politics of AI have changed immensely in the last six months. AI as a category is polling mostly underwater, despite some views that there could be an opportunity there. Most Americans aren't seeing it. And the place the rubber hits the road is whether or not there was going to be a data center in their community.

And just right now, you see, you know, pretty sizable, opponents would call it NIMBYism, others would say more effective zoning. Right? But trying to keep data centers from going into communities because this view that they are giant, they don't create that many jobs, perhaps they have water consumption. This is a major point of disagreement, actually, between the companies building them and some of the political and other folks that are looking at it.

This is ultimately going to be a question of whether or not we can power these data centers. AI requires a tremendous amount of energy. That energy has to get on the grid in a certain period of time. And right now, the lead time in the United States to getting new energy on the grid is immense, something called the interconnect queue.

And, you know, we're not building new fission nuclear power plants at a rapid clip. Well, China is. What are we doing when we have a sudden surge of need for new compute? Well, it's like that deal that Anthropic just made with SpaceX that helped juice their IPO, where they have these massive data centers, Colossus-1 and Colossus-2, that are basically running on natural gas—some of it questionably permanent—running on natural gas and and belching out emissions into the area. This reality is undisputed.

Now, if you look at the Chinese counterpart, are they also running a lot of this energy on dirty fuel? Yes, they are. Do they have more energy capacity, more slack in their overall grid to run it? Yes, they do. And do they have more production coming online, *ceteris paribus*, right now against the U.S.? Absolutely.

[17:31]

So the challenge that the U.S. faces right now vis-à-vis China is, you know, twofold. One, yes, we need to stay at the frontier without question on the technological pieces. But making sure that we are innovating towards efficiency, something that through a combination of chips requirements and others the Chinese have been forced to do, that may end up being just as critical to whether or not we can diffuse AI across the economy as whether we are right there up at the frontier.

[17:56]

China has more of it. They're building more of it. And, you know, absent breakthroughs that we really need in firm baseload carbon-free energy 5 to 10 years from now looks very challenging, maybe as challenging as the next 18 months when it comes to whether we're able to bring enough energy online.

And I think the key data point that we need to keep looking at is completions of data center projects. In the U.S., you're seeing a huge number of projects announced, much smaller number of them actually on the pathway towards completion. A lot of that is permitting. Other elements have to do with the finances coming through. The music has to not stop on the equity markets right now.

But you know, as Kyle pointed out, and I think it's really important, we are seeing not just a once in a generation, once in a lifetime investment in physical computing infrastructure right now. It is an probably unprecedented transfer of wealth. And where is it coming from? Largely, it is coming from the free capital flows of existing tech incumbents that are taking advertising and other revenues and putting them over into data center build. Now, it's also coming from our pension funds, let's be clear.

But there's a really interesting dynamic undergirding all of this, which is the U.S. capital markets do have the potential to help the U.S. energy crunch catch up in a way that may not be as true in China, even if the permitting is a lot easier in China even if they're building a lot more efficient plants.

[19:14]

HASS: One of the issues that David raised, Kyle, was the need to design for efficiency. The Chinese have been forced to do so because of constraints on their compute capabilities. And so I wanted to ask you a bit about chips and export controls. The United States has restricted certain types of advanced chips for export to China. How significant have these export controls been in impacting China's pace of development? And more broadly, are the export controls working?

[19:38]

CHAN: So there are two sides to this story. On the one hand, the export controls are working in the sense that they are limiting the compute for Chinese AI labs. It is slowing down their progress. It is making it more difficult for them to train large models at the scale that we see in the United States. It is making it more difficult them, for them to stay at our pace at the frontier. That's in the near term.

In the medium to long term, there are sort of second order effects that are really important to consider. Namely, that this has really incentivized China, Chinese policymakers, and the Chinese chip industry to double down on building their own self-reliant semiconductor supply chain. That is, this is something this is I don't think ever been done in history, a single country trying to internalize and reproduce the entire global semiconductor industry. We're talking about everything from lithography, to etch and deposition, to the photoresists, and the key chemicals that go into this process, to the software that's used to design chips. Right now, globally,

this is scattered across a number of different economies— the U.S., Japan, South Korea, the Netherlands—and China's trying to replicate almost all of it domestically.

And not only that, but they're looking for ways to continue to make progress on AI despite these compute constraints. One is, while they won't be able to catch up, I believe, on the single chip level— that is the chips coming from Huawei I don't think will ever be at the performance levels that we see from Nvidia— they are trying to make progress in terms of networking greater numbers of these chips together into larger compute clusters, leaning heavily on some of their networking expertise going back to their 5G build-out.

On top of that, they're looking for ways to make the models more efficient. So this is sort of on the software layer, where given the same amount of compute, can they squeeze out more performance? Can they make these models run more efficiently to the point where you can also deploy them more efficiently? They require less energy, they require less compute resources, and ultimately are they gonna be cheaper, going back to David's point.

[21:44]

And then another way that they're trying to get around this is they're trying to innovate in areas of chip manufacturing that don't depend on key choke points like EUV, or extreme ultraviolet lithography, which currently only the Netherlands can produce machines that can do this.

And one thing that they're trying to do there is they're trying to use what's called multi-patterning and also other forms of semiconductor manufacturing technology like hybrid bonding. Whether they'll be able to catch up fully, I would be skeptical because unlike other industries where China has been able to catch up or even become the the new pioneer, the semiconductor industry is like a Chinese fast train trying to catch a bullet train. This is very difficult. The entire industry is moving fast. TSMC, Nvidia, all of the supply chain.

But can they produce good enough chips, especially on the inference side, to continue their AI development, to continue making progress, not only again on training these models, but on deploying them at scale, especially in the age of AI agents that are heavily reliant on inference, on the ability to run these models over and over again over long periods of time for large workloads for many different customers simultaneously. This is a question that I think Chinese AI chips and Chinese models can continue to make progress in in the face of export controls.

So overall, any fair assessment of U.S. export controls would have to recognize that both are true at the same time, that in the near term, this does limit some of the Chinese AI training. And this is important because I do think in some narrow areas we should retain a lead, like in cyber capabilities.

But in other areas, this has really incentivized them and motivated the entire chip industry to essentially not just buy the fish from the U.S., that is rely on U.S. chips, but to build the fishing rod themselves in order to have an unlimited supply of their own compute in the long run.

[23:38]

EDELMAN: Kyle points to something that I think might turn on its head what a lot of us think of as this sort of state control paradigm, which is at this moment, unusually so, I mean, others can check the math here, but it strikes me as highly likely that the free cash flow of NVIDIA, Apple, Google, say Amazon too, is probably higher than the CCP's ability to marshal resources against this set of needs at this time.

You know, normally we think about, oh, state economy, they can direct all these resources accordingly. This is a key priority. Xi Jinping is talking about it. Therefore, like, how can we possibly compete? And the answer at this moment might actually be this unprecedented, incredible amount of resource richness that these companies have right now to make these competitive investments at a timescale that, as Kyle points out, is vastly ahead of when centralized state planning can even try to catch up.

You know, I think the the key questions we need to be asking here for the future of American strategy, you know, one, have we created a self-fulfilling prophecy, which Kyle gets to? Right? Was it through the combination of our own restrictions, statements we've taken, the actions we've taken against particular companies? Has it, in a complex environment of Beijing's domestic political situation, actually pushed the nationalist camp, the nationalist economic development camp, to ascendance in a way that wouldn't otherwise be true?

In other words, yeah, Xi Jinping would have been talking consistently about the need for fully domestic production of chips, but would they have really gotten there? Would he have had a domestic constituency to actually make that kind of an unprecedented investment? This is costly stuff. And so one big question that we'll have to see play out in the next decade is, was it fundamentally of our doing? We may never know, but I think that's important.

[25:15]

And the second is this question of, you know, energy intensiveness and efficiency as a real breakaway point. Because so long as we're in an environment in which for many U.S. frontier AI companies, yes, energy intensiveness is a liability and a cost, it is also a competitive moat. If it costs \$5, \$10 billion to do the training on a frontier model, then only companies with that level of resourcing will be able to do it. It puts downward pressure on competition for it. It may dilute the second mover advantage.

These are all things worth bearing in mind because a world in which there are major breakthroughs in efficiency, I mean, we've seen them in our labs at MIT, you know, master's students, Ph.D. students coming out with really incredible advances in how to train these models efficiently. We have every reason to believe that in part because of, maybe indirectly, these U.S. chips controls, well, as Kyle said, the Chinese have had to innovate to that level of efficiency.

And a breakthrough world in which inference becomes 1% as energy intensive as it would in the U.S. with the U.S. models, that is a breakout model. It's it's a breakout moment for application of AI, and it's a breakout moment for the possibility of these

AI systems to ultimately aggregate into something that is greater than the sum of its parts.

I mean, I don't want to talk about recursive self-improvement here, but we need to think about those key questions because the strategy that the U.S. and China are pursuing are indeed quite divergent.

And, you know, it is certainly my hope that smart policymakers, and certainly those here at Brookings and elsewhere, are playing out those scenarios, especially because this is not gonna be the end of it. There will be new hardware requirements that will come up in the next five to 10 years where strategic decisions will have to be made in the next two to three that will shape where we end up in this conversation in 2035.

[27:01]

HASS: Which country do you think, the United States or China, is in a better position to take the lead in AI innovation, given everything that we've just discussed? And what factor, whether it's energy availability, closed versus open models, advancements in chip technology, data centers, infrastructure, which which factor do you think will have the most decisive effect on on which country is going to sort of front run the other?

[27:26]

EDELMAN: Look, my money's on the United States, no surprise. Why? I mean, I think it's a few pieces. One, critically, we haven't talked about it, it's talent. The U.S. ability to attract, at least for a period of time and hopefully on an ongoing basis, the best and brightest minds in the world is directly material to a huge number of AI innovations that we have seen in the last while.

You know, look at the people who are the core technical staff of these frontier model companies and the researchers at the universities that are training them, that are working on the basic science of this stuff. This is a, you know, relative United Nations, and they want to come here to do that research.

That is a resource that could be squandered. That is a potential risk factor obviously at this moment. But I think fundamentally that combined with, you know, a predilection towards openness, an idea that you can innovate on models here, that there is not for the time being pre-market review of models, that you can put together an open source. We're talking about, you know, open models in China. You can do that in the United States too, and lots of people are.

That constellation of talent, of capability, and ultimately then of the investment ecosystem around it has my money on the U.S. for pure innovation every time. Diffusion might look a little different, and it might be slower, but for who's absolutely at the frontier, which economy I think has the opportunity to benefit from that most immediately, my money's on the U.S.

[28:45]

HASS: Kyle?

CHAN: I would break it down into sort of the virtual world and the physical world. So I absolutely agree with David that there is no way that the Chinese labs will match the amount of R&D investment and compute and experimentation that's being done to try to find new paradigms for pushing the technology forward, at least in the virtual world.

And, the U.S. is moving—been moving very aggressively on not just building bigger and bigger models, but developing more sophisticated agentic tools for using them. And I think this plays to the U.S. strengths in, for example, the broader software ecosystem and the broader commercial ecosystem around businesses, for example, paying and using licensed software. So this feeds right into an existing pattern that we've seen the U.S. capitalize on many times before.

However, on the physical AI side, this is an area where I would expect more innovation and faster innovation to happen in China. This is an area that plays directly to China's strengths. So by physical AI, I'm talking about robotics, autonomous vehicles, autonomous drones and delivery systems, even sort of smart service robots. Think about your humble street cleaning robot, for example. Also think about robots on the factory floor. Not just sort of your classic industrial, six access robotic arms, but your smaller cobots or even humanoids on the factory floor.

This is an area that plays into China's strengths for two reasons. One is on the supply side. China has the largest, most robust robotics component supply chain in the world. So that includes everything from sensors, actuators, batteries. We're talking about LIDAR and camera modules. We're talking about various joints, even now robotic hands. There are a number of prominent Chinese players in the space.

And this also plays into China's strengths on the demand side. That is, China has the largest manufacturing industry by far in the world and is able to deploy some of these robots into these settings and, crucially, get this data flywheel going, where through deployment they collect data, they use that data to train future iterations of hardware/software combinations and improve this broader robotics ecosystem and then feed that back into more robots on the factory floor.

So it depends on which space we're talking about, and that's why I would go back to, you know, the question at the start, which is, what are we talking about when we talk about an AI race? I think it's actually much more complicated than a a single uni-dimensional race.

[31:12]

HASS: Well, AI implicates many things, one of which is national security, something that both of you have spent a lot of time thinking about. Where do you see AI affecting the national security relationship between the United States and China today? And what do you think that we most often underappreciate about the impact of AI on on the national security relations between the United States and China? Kyle?

[31:34]

CHAN: So one strange thing to consider is the degree of actual dependency or mutual exchange between the U.S. and Chinese AI ecosystems that's happening sort of despite almost the the hopes and best wishes of some of the geopolitical actors, the governments and and and regulators.

So for example, a lot of American AI developers are huge fans of Chinese AI models for a lot of the reasons that we've talked about, including ones that David has pointed out. Their cost, their open source nature makes them easier to download and use and customize for their own purposes.

And in China, a lot of Chinese AI developers and Chinese developers more broadly love using especially Claude. This may be harder going forward as Anthropic starts to tighten controls on access to these models. But overall, this is something that's quite interesting.

And so I think for national security folks thinking about the risks to this kind of exchange, we should just keep in mind there's sort of multifaceted risks. There are, I think, risks for the United States to fall behind and what that means in terms of, for example, again, cyber capabilities, and what it means for China to have a Mythos-level model first and be able to test that out on American cyber defenses. That would be a bad situation for the United States, for sure.

But on the other hand, there are also risks in sort of going about this in a overly crude fashion. I think that efforts to too radically decouple the two ecosystems, including not just who's using whose model, but also the talent flows, and also this this broader exchange on the, the science and the R&D side, I think could end up backfiring on the U.S., for example, if we don't do this in a careful, thoughtful manner.

So for those thinking about the national security implications, I would just urge everyone to consider the whole range of risks here, and not focus too narrowly on just one idea of, you know, who has the best model and how do we make sure that we we in the U.S. retain that edge.

[33:39]

EDELMAN: I agree with what Kyle just said. I think the most important risk, big picture though, on the national security side has to do with acceleration. It is accelerating two things, very different, but both important. One, it's accelerating the pace at which certain risks can metastasize. You know, you think about, Kyle mentioned the cyber risk. I happen to be, in, in the view— I spent a lot of time working on on this subset of issues— of the view that there is actually a potential for a major advantage on the defensive side.

Because the reality is, a lot of cybersecurity and a lot of systems, certainly in this economy and elsewhere, are not great and do not have fabulously talented full-time CISOs, cybersecurity experts, that are in charge of them. And so the ability of these systems to do vastly better cybersecurity work on code that would otherwise be shipped with tons of holes in it, is is real, and I think that's important.

But the other side of that coin is absolutely important to think about, which is, yeah, the speed at which one can develop offensive capabilities, the exquisite kind, you know, that we heard about and that Anthropic claims that they were able to find with the preview of Mythos and their Project Glasswing, but also just the speed at which you can weaponize scamming, where you can put together a convincing-looking PDF and send it to someone in a government or in a company that they actually click on. You can now do that 1,000 times faster because of these models, not because they are amazing frontier world-changing models, just because they're, like, pretty good at putting together PDFs based on a little bit of information, and email marketing has a lot of data associated with it. Right?

So that is a set of cybersecurity risks that are, you know, material, worth thinking about, and dramatically accelerated because of AI, and the defenses as well.

[35:19]

The other acceleration, though, I think actually has to do with much more conventional decisionmaking, and this is where I think things get big picture really quickly and very concerning. I mean, we've spent a lot of time working on this together, Ryan, that, you know, this question of the potential for a conventional military misunderstanding metastasizing into a shooting war because AI was involved. Not the killer robot scenario that is also worth thinking about, but the immediate.

There is, you know, a contingency somewhere, a downed plane, a rogue anti-aircraft system that had some AI layered in, and no one can talk about it. No one's able to describe what happens. There is absolutely a low level of trust that countries would be fielding this technology before it's ready, when it's error-prone, like the AI we're all using every day.

To me, these are very conventional military risks, very almost traditional national security concerns made dramatically faster with less time for decisionmaking, and critically less time for communication for de-escalation than otherwise would exist.

And so that's why I think there's there's a critical need at this moment, you know, to have really hard conversations, which, you know what? We did during the Cold War. This is not a cold war, but we did during the Cold War for decades with the Soviet Union and then the Russian Federation on how we can prevent a, say, nuclear war that neither side wants. If sides want it, that's a different question. But that's where I think, you know, AI is compressing this decisionmaking cycle.

And I will say we're not coming into this completely blind. Fifteen years ago, we had the exact same conversation about cyber, and people said, "Well, what's cyber going to do to our decisionmaking apparatus? That we're gonna have to make every decision at network speed." You've heard that before, right? Everything has to happen at network speed.

And it turns out, no, when it comes to big picture national security decisions, by and large, the United States reserves the ability to respond in the time and place of its choosing, in the manner of its choosing, and often not in the cyber domain. That is

comforting to some extent. It is not true that we have automated warfare ready to go in every context in the cyber world.

I do think AI challenges though, and certainly as much as cyber did and more so, press us up against whether we can make sure we are really guiding these decisions appropriately, not just with human hands and human eyes, but clear human minds, as opposed to being subject to what we call the automation bias: the AI system said it. The AI system has to be smarter than me. Therefore, let's push the button. If that had happened back in the days of Stanislav Petrov, at the height of the Cold War, none of us would probably be sitting here today. Right. And that's what we need to be thinking about, not as a techie issue, but as a big picture national security matter.

HASS: Well, let's let the spirit of Petrov guide us forward. We are out of time, but I think we'll be coming back to this conversation again in the future. David, Kyle, thank you for for sharing your expertise with us and our audience on the U.S. and China's AI ecosystems. I'm really struck listening to you by the different approaches that the United States and China are taking to harness AI to improve their national conditions. I'm also leaving this conversation with a renewed sense of optimism that the United States has the capacity and the bandwidth to use AI to improve our national condition.

So thank you both.

For more in-depth analysis from our team, visit the China Center on the Brookings website at [Brookings dot edu slash ChinaCenter](https://www.brookings.edu/china-center). You can also subscribe to our monthly newsletter, the China Bulletin, for the latest updates.

JON CZIN: On behalf of the team at the John L. Thornton China Center, thank you for listening to *The Beijing Brief*. This podcast is produced by the Brookings Podcast Network.

Our thanks to the production team, including supervising producer Ike Blake; senior producer Fred Dews; producer Allie Matthias; audio engineer Gastón Reboredo; and video producers Daniel Morales and Teddy Wansink. Rachel Slattery designed the show's artwork. Also, thank you to our colleagues in the John L. Thornton China Center, Foreign Policy, and Office of Communications at Brookings for their support.

To learn more about our research, visit us at [Brookings dot edu slash China Center](https://www.brookings.edu/china-center), and to learn more about this podcast, go to [Brookings dot edu slash The Beijing Brief](https://www.brookings.edu/the-beijing-brief) or wherever you like to get podcasts.