



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

“Why Washington fears China's open-source AI”

Thursday, July 30, 2026

Participant:

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

Episode Summary:

What does the invention of Chinese AI model Kimi K3 mean for Washington and Silicon Valley? In a new explainer for *The Beijing Brief* podcast, Kyle Chan breaks down the ongoing AI battle between Chinese open-source models and American AI companies like Anthropic and OpenAI.

[0:00]

[music]

CHAN: Everyone's focused on the race between OpenAI and Anthropic, but there's arguably a bigger AI race going on, the race between the U.S. and China.

The AI race is no longer just a competition between companies and startups. It's a contest between nations over global power and influence. It's a competition between the U.S. and China to see who will shape the most consequential technology of the 21st century and has the potential to reshape global politics.

I'm Kyle Chan, fellow at the Brookings Institution, and this is a *Beijing Brief* explainer.

[0:35]

Right now, there's a big debate going on in Washington and Silicon Valley over open source models, and open source AI models are overwhelmingly dominated by Chinese AI companies like Kimi, DeepSeek, and z.ai. Most recently, Moonshot just released a powerful new model called Kimi K3 that looks almost as good as the very best models coming out of Anthropic or OpenAI.

The AI race is shaping up to be more than just about bigger and better models. It's a battle over access, open versus closed models, the U.S. versus China, and of course, who will gain the most from this transformative new technology. Many in Washington are worried that if the U.S. loses its edge over China in AI or even slips behind, this could give China a key advantage in areas like cyber capabilities.

This matters more than ever because AI models are getting increasingly powerful, and at the same time, we're seeing a growing backlash to the data center buildup across the U.S. Many are asking what this technology is really for and how they might stand to benefit, especially in the face of worries over potential AI-induced job loss.

On the one hand, open source models or open weight models are ones that are freely available to download online that you can customize and fine-tune on your own data. What this gives you is not just cheaper models, but also greater control over what you can use them for.

[2:06]

Closed models, on the other hand, are models that you pay for through, say, a monthly subscription fee or through API fees. What this means is you often get more powerful capabilities but at much higher costs, and that data ends up going back to the companies that serve you the model.

Chinese models, on the other hand, are mostly open source. What that means is that you can actually download DeepSeek's latest model, run it on your own computer, or run it on a cloud platform. This allows you to access the model for free or at much lower prices, and it allows you to customize the model depending on your needs and your data.

China's open source strategy has allowed Chinese models, which are generally lagging behind the best U.S. ones, to catch up in terms of adoption and global reach.

[2:56]

This is really important because the AI race is not just about the best models, but about who uses them and how. Right now, Chinese models are getting more popular around the world, including in the United States, because developers can download them and run them at very low cost.

Crucially, what you get with these Chinese models is something like 90% of the capabilities of an American model, but at a fraction of the cost. And this is starting to challenge the business logic for some of these trillion-dollar companies like OpenAI and Anthropic. They've made a bet that by charging fees, they can recoup the massive costs it takes to build data centers and train ever larger models.

[3:38]

But with these Chinese AI challengers, now part of the market is shifting away and switching over to these less expensive Chinese models. In a way, we've seen this story before in other industries where Chinese alternatives have eaten away into market share from American companies, including in industries that the U.S. once dominated, from automotive to consumer electronics.

So because Chinese models offer close to frontier performance but at a fraction of the cost, they're getting more popular among global users and even in the U.S..

This is also happening at a time when AI costs are soaring.

[4:17]

The release of Kimi K3 has surprised many in Washington and Silicon Valley because it means that the gap between the U.S. and China on AI is much smaller than people had thought and may be getting narrower. A lot of people had thought that China would take months, if not more than a year, to get to a Mythos level capability, but now that might be around the corner.

[4:39]

What does this mean for the United States, for the U.S. AI industry, and the future of global AI governance?

In addition to Kimi, we have a whole slew of Chinese open source models that have been released recently, including GLM 5.2 and a new model from Qwen. All of these are starting to reach into the multi-trillion parameter range. What does that mean? That means these are really, really massive models that many people outside of China didn't think could be trained and developed on Chinese hardware.

[5:11]

Some of this is due to a process called distillation. Distillation is when you take a more powerful model and use its outputs to help train a less powerful model. Virtually

all the major U.S. AI labs have issued reports on Chinese AI distillation on their own models and want Washington to do something about it.

The question is, what should be done? The problem is distillation's really hard to tackle because it's hard to separate legitimate uses of AI from these more industrial scale distillation campaigns.

[5:43]

But is all this talk about distillation a distraction? On the one hand, distillation is an important problem that should be addressed, and ideally the government steps in to help coordinate between the U.S. AI labs. On the other hand, though, there is a fundamental capability among these Chinese AI labs that allows them to be so competitive compared to American models.

The narrow gap between Kimi K3 and the top models from OpenAI and Anthropic make it much harder to argue that Chinese AI labs are able to keep up on distillation alone. There's something real going on beneath the surface here, and that is a lot of talent and a lot of resources in China going towards AI development.

The AI race isn't just America versus China anymore. It's open innovation versus closed control. Which side do you think wins out in the end?

[music]

For deeper dives on Chinese AI competition, head over to Brookings dot edu slash TheBeijingBrief to listen to recent episodes of *The Beijing Brief* on this topic.