

Cutting access to Anthropic's Mythos is a gift to China

Financial Times Europe

16 giugno 2026

Agathe Demarais The writer is senior fellow in geoeconomics at the European Council on Foreign Relations and the author of 'Backfire'

Last Friday, Washington imposed export controls on Anthropic's latest frontier AI models Mythos and Fable 5, forcing the company to cut off global access to both. While developers lamented the loss on online forums and Wall Street tried to assess the impact it may have on Anthropic's upcoming initial public offering, the bigger picture was lost in the noise: Washington has just done more to boost the appeal of Chinese AI models than Beijing could have ever hoped.

The Trump administration's implementation of the export ban was chaotic. The notice arrived at 5.21pm on a Friday, attributed to a security issue that Anthropic says is a misunderstanding. For global companies pondering which AI model to pick, such a process is a red flag. If Washington can shut down access at such short notice, it undermines the case for building infrastructure around US AI. The lack of transparency carries a warning that this extraordinary action could be repeated.

Washington's decision to use export controls shows that its concern is with foreign use of Anthropic's powerful latest models. But nationality is a poor proxy for misuse and if security concerns are genuine then an outright ban would make more sense. Plus, export controls only bite when the target is hard to substitute, such as extreme ultraviolet lithography for chip manufacturing. Although industry benchmarks show that no model is as good as Mythos and Fable 5, the suspension will fuel demand for substitutes readily available from Chinese competitors.

The performance of Chinese models like DeepSeek's V4 Pro, GLM, Kimi, MiniMax and Qwen is increasingly on a par with that of US AI. Crucially, these models are self-hostable (meaning companies can run them from their own servers), making them harder for governments to limit access to. They are also cheap. DeepSeek's flagship model costs just 87 cents per 1mn output tokens, making it about 60 times cheaper than Anthropic's Fable 5.

The idea that global companies could turn away from American AI is not farfetched. In the first two weeks of June, four of the five most popular models on OpenRouter, a platform that provides access to hundreds of AI models, were Chinese. Among the global top 20, Chinese models processed twice as many tokens as their US-developed competitors. If erratic policymaking makes reliance on US models start to look risky, switching to Chinese substitutes may become more common.

Washington cannot ignore this development. Wider adoption of Chinese models supports Beijing's bid to shape global AI standards. When developers and companies build on Chinese AI architecture, they absorb Chinese technical assumptions — giving Beijing longterm influence and leverage. The process is self-reinforcing. The more developers use

Chinese AI models, the more they build custom tools that they share with peers, making Chinese AI models more appealing to other developers.

The suggestion that only America should be able to access US frontier AI models also raises awkward questions for Anthropic, OpenAI and xAI's parent SpaceX. If these companies cannot sell future models abroad, the rationale underpinning their sky-high valuations rests on shaky ground.

Beijing has spent years telling potential customers that its technology supplies would never be cut off for political reasons. Washington has just strengthened that pitch, making Chinese AI models look less risky even if they remain weaker on capability. Whether the US reinstates access to Mythos and Fable 5 is beside the point. The signal that access to American AI is revocable is a remarkable own goal. Global companies may decide that the best AI is not necessarily the smartest but the one that they can access every Monday morning.