

*National-security threats combine with fears of a looming AI financial shakeout.*



**BUSINESS WORLD**

**Anthropic vs. China**

Gen. Henry "Hap" Arnold, commander of the U.S. Army Air Forces in World War II, wanted to know how rockets, satellites, jet engines and nuclear weapons might change warfare in the postwar world. Today, he wouldn't have to invent the RAND Corp. Anthropic, one of America's leading builders of artificial intelligence, perhaps later this year to become a trillion-dollar public company, is perfectly happy also to be its 21st-century anti-China think tank.

CEO Dario Amodei has been unflinching in claiming that we face a race with China, that the survival of democracy depends on who wins. Even amid its weird fight with the Trump Pentagon, Anthropic tools are almost certainly being used right now by the National Security Agency in its twilight cyber struggle with Chinese spies and hackers. While President Trump visited Beijing, Anthropic posted a 5,500-word manifesto on how to conduct this new AI Cold War with China.

The fight only begins with the fact that the latest cutting-edge U.S. models, led by Anthropic's Mythos, seem to be uncannily good at finding dangerous and unsuspected vulnerabilities in all major software systems. China and criminal hackers are, at most, only a year behind. Next question: Is this one problem or two?

Increasingly a sinister hue attaches to Beijing's preference for "open weight" models. Any user is free to adapt and modify them at will; their Chinese builders, unlike Western AI builders, don't reserve the ability to monitor how the models are being used or exercise a kill switch over antisocial activities.

Is this because China believes its national firewall and comprehensive surveillance renders it immune? Because it plans to put such tools in the hands of anti-Western hackers while preserving deniability?

Beijing is still thought by some to have remained mum on early Covid because it believed China enjoyed a relative advantage in a global pandemic. During Mr. Trump's Beijing summit, Treasury Secretary Scott Bessent told CNBC a key goal would be "to make sure nonstate actors don't get a hold of these models."

To Anthropic and others, the pressing danger is to their business models, and to any hope of getting paid back for their billions and billions in investment, which China, they believe, is deliberately trying to undercut.

Anthropic's pitch for fighting back comes down to two things: China's access to toptier chips and "distillation"— the illicit use of Western models to train Chinese models, hijacking (in the company's words) "decades of foundational research, billions of dol--lars in US investment, and the work of thousands of the world's best engineers."

There's nothing like sitting down with AI mavens, as I recently did, to make a columnist swell with pride for (possibly) coining the term AI Chernobyl. Big questions loom. The Obama administration used drone assassinations against foreign terrorists. Should the

same be visited on cyber kingpins? Kinetic and cyber action can both do severe economic damage. Both can cause civilian casualties. Should we stop treating one as warfare and the other as crime?

It doesn't solve everything, but mutual assured destruction worked passably well for 70 years to hold off nuclear Armageddon. China's economy is no less subject to comprehensive cyber shutdown than ours, especially given its reliance on software-based systems for almost every kind of transaction (cash being all but obsolete there). Not to mention its government's giant investment in software-based monitoring of citizens and their movements.

The China fears are real. Don't underestimate, though, how much they are tangled up in worries about a coming AI shakeout. Take this week's de--feat of Elon Musk in his lawsuit against rival OpenAI. Nominally the case turned on OpenAI's conversion to a forprofit company and Mr. Musk's worries about *Homo sapiens* being displaced by superintelligent robots.

In fact, Mr. Musk was engaged in the great American tradition of waging lawfare against a business rival. This comes at a moment when he and others are looking to loft AI-related IPOs. Don't say it too loudly: This may be a race to shift to the public some of the losses from an AI buildout that seems increasingly unlikely to yield profits for every contestant.

Such problems Hap Arnold didn't face. RAND was born as an arm of defense contractor Douglas Aircraft, but it was quickly spun off into a nonprofit to become America's premier Cold War think tank. Meanwhile, Douglas (now McDonnell Douglas) found it perfectly possible to sell fighter jets to the U.S. military even while it set up a passenger-jet factory in communist China.

Monday's court ruling, to the chagrin of Mr. Musk, settled one question. OpenAI isn't a charity but a business with a trillion-dollar valuation in the private market. A similar question may face Mr. Amodei and fellow Anthropic executives sooner than they think: Is Anthropic a business or an anti-China think tank? America may need both but the company likely will have to choose if it wants the public to invest.

*By Holman W. Jenkins, Jr.*

---

[Copyright \(c\)2026 Dow Jones & Company, Inc. All Rights Reserved. 5/20/2026](#)  
[Powered by TECNAVIA](#)

---

The following is a digital replica of content from the print newspaper and is intended for the personal use of our members. For commercial reproduction or distribution of Dow Jones printed content, contact: Dow Jones Reprints & Licensing at (800) 843-0008 or visit [djreprints.com](http://djreprints.com).

---