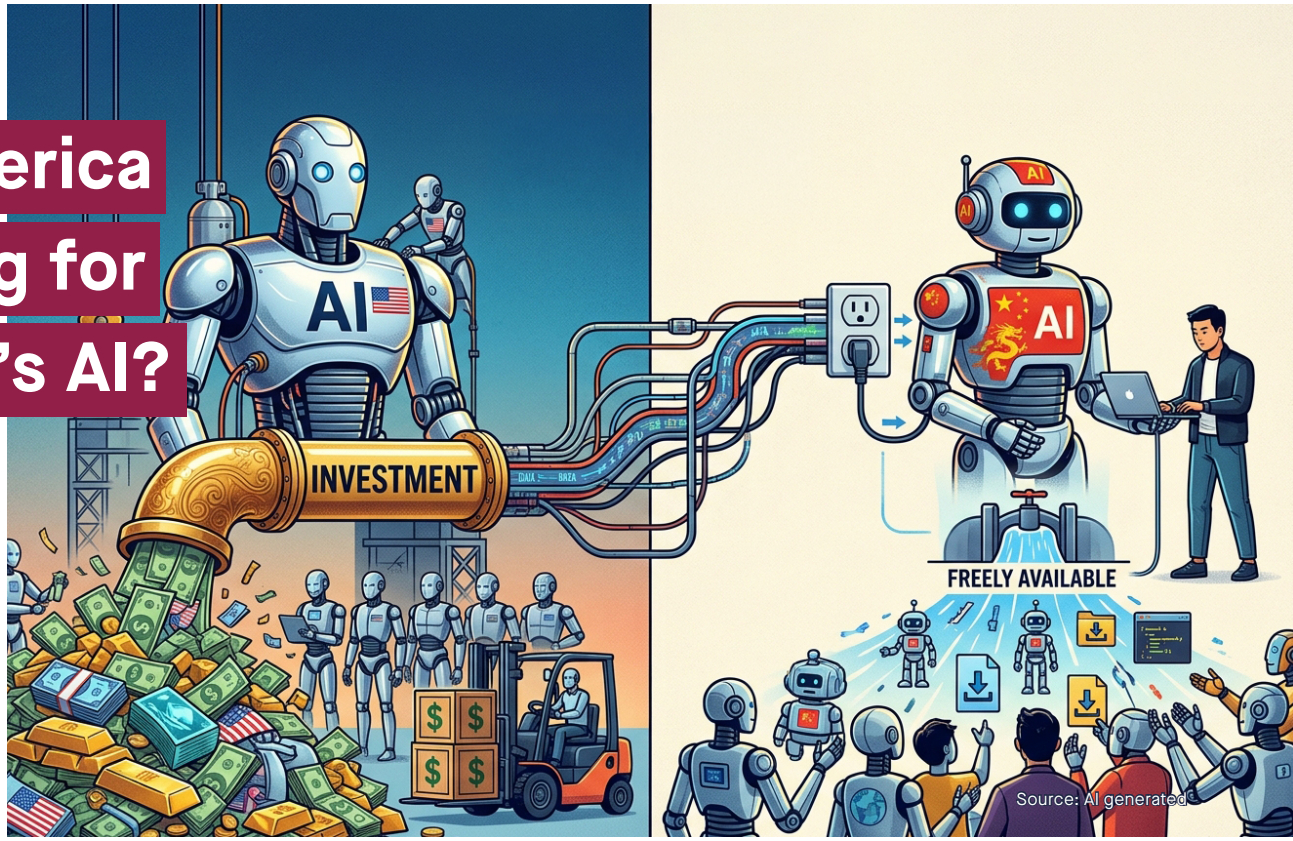
Is America paying for China's AI?



US tech giants are investing hundreds of billions of dollars to develop the most advanced AI models. Meanwhile, Chinese companies are developing models with similar capabilities at a fraction of the cost, often making them freely available to download. Some analysts see this as a form of subsidy, but the economics behind it are more complex.

In July 2026, Chinese start-up Moonshot AI released Kimi K3, a model that Axios reported outperformed Anthropic's latest system on a popular coding test while costing 40% less. A few months earlier, Anthropic had accused Moonshot and two other Chinese labs of secretly using its Claude model to train their own, highlighting an argument that has already been circulating among investors. Microsoft, Alphabet, Amazon and Meta plan to spend roughly $700bn on infrastructure in 2026, about 75% more than in 2025, according to the Financial Times. If Chinese labs can reproduce the same results for less expensive just a few months later, are American shareholders basically funding China's AI industry? The evidence suggests the argument is partly right but overstated. US spending does not flow to China, but it buys a lead that pays off for a shorter time than investors might hope.

How AI models are built and copied

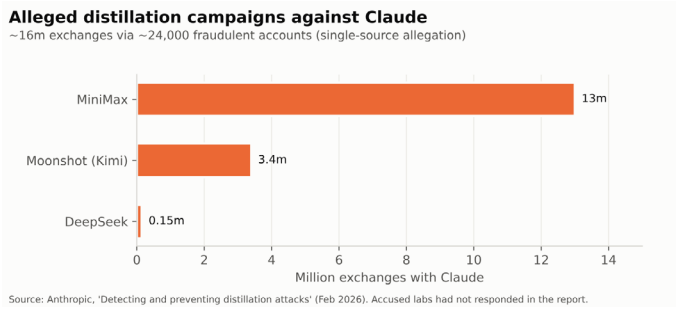
To understand the debate, it helps to know what "the frontier" means. In AI, the frontier refers to the most capable models available at any given time, today produced by companies such as Anthropic, OpenAI and Google. Reaching it is extremely expensive. Labs need vast data centres filled with specialised chips, large amounts of data, and many failed experiments before a model works as intended.

Chinese labs such as DeepSeek, Moonshot and MiniMax follow a different commercial model. They mostly release "open-weight" models, meaning anyone can download the model and run it on their own servers, often for free. Companies then pay only for the computing power they use, which makes these models much cheaper than their American counterparts.

The controversy is based on a technique called distillation. Instead of learning everything from scratch, a new model is trained on the answers produced by a more advanced one, much like a student copying an expert's worked solutions. Distillation is common and legitimate within a company. It becomes a problem when one lab uses a competitor's model without permission, since it captures the result of costly research without paying for it.

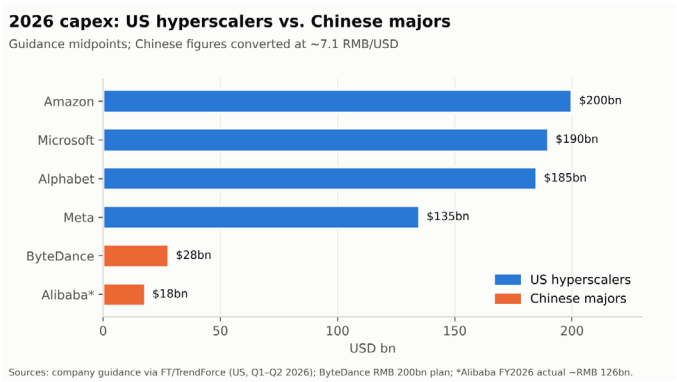
The case for the "subsidy" argument

The most direct evidence came in February 2026. Anthropic reported, in findings covered by The Decoder, that DeepSeek, Moonshot and MiniMax had generated more than 16 million exchanges with Claude through around 24,000 fraudulent accounts. Computing's breakdown shows that MiniMax accounted for over 13 million of these exchanges, Moonshot for more than 3.4 million and DeepSeek for about 150,000.. The requests focused on coding, reasoning and the ability to carry out complex tasks autonomously, precisely the skills on which US labs have spent most and for which businesses pay the highest prices. In a memo to the US congressional comitte, OpenAI made similar accusations against DeepSeek the same month, and Google disclosed an attempt to copy its Gemini model using more than 100,000 prompts. These remain company allegations, however, and nobody has measured how much the Chinese labs actually gained.



Source: Anthropic

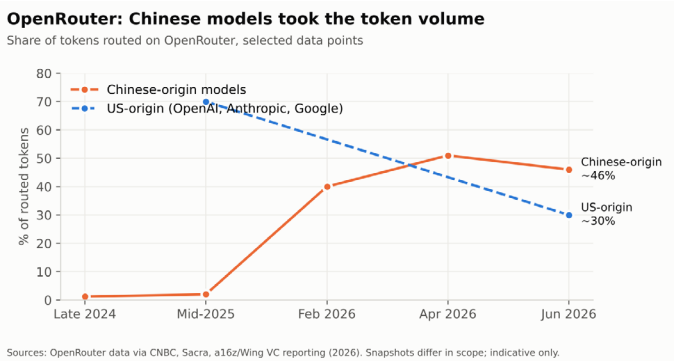
The spending gap adds weight to the argument. US hyperscalers are expected to pass $1tn of annual investment in 2027, based on estimates from Evercore and Bank of America. By comparison, ByteDance, China's largest spender, plans to invest about $28 billion in 2026, EqualOcean reports. Labs such as DeepSeek or Moonshot have no comparable infrastructure at all. Supporters of the subsidy argument point out that most of the cost of building an AI model comes before the final version. Labs spend heavily on experiments that fail, on collecting data and on testing. A lab that learns from a rival's answers skips these costly steps and goes straight to the result, while the rival has paid for all of them.



Source: Company guidance

Yet the results are close. Research institute, Epoch AI, notes that Chinese models have trailed the best American ones by an average of seven months since 2023. However, open-weight models in general, including those developed in the United States, trail the leading closed models by a comparable margin, suggesting that the delay is not specific to China. Their adoption is also growing rapidly. On OpenRouter, a platform that lets developers switch between AI models, Chinese models went from about 1% of usage in late 2024 to around 46% by mid-2026, according to Forkast, citing CNBC. The reason

is price. OpenRouter's own analysis finds these models 60% to 90% cheaper than leading US alternatives. Even American companies have adopted them. In March 2026, the coding tool Cursor launched a model built on Moonshot's Kimi, without disclosing it at launch.

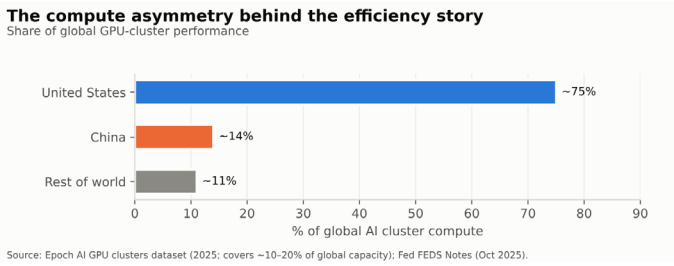


Source: OpenRouter

Why the argument goes too far

The idea that Chinese companies do not invest in AI does not survive a closer look. Alibaba and Tencent together spent more than RMB 120bn, around $17bn, in the second quarter of 2026 alone, according to TMTPost. The famous claim that DeepSeek trained a leading model for just $5.6mn covered only the final training run, as DeepSeek itself stated. The research firm SemiAnalysis estimated that the company's servers alone were worth around $1.6bn, nearly 300 times that figure. Chinese labs spend far less than American ones, but they are far from working on a shoestring budget.

Ideas also travel in both directions. Several of the most important techniques for making AI cheaper to train and run were developed by DeepSeek and published openly, and US labs have since adopted them. CSET data cited by Al Jazeera show that China also produces more AI research than the United States, with around 27% of global publications in 2024 against 12% for the US. Limited access to advanced chips played a role too. Because US export controls restrict China's access to the most powerful hardware, Chinese engineers were pushed to get more out of less, making cheap models partly a genuine engineering achievement.



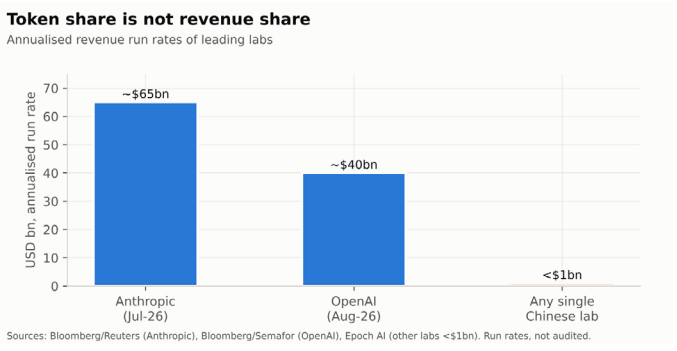
Source: Epoch AI

In addition, most of this American spending is safe from copying. The majority of hyperscaler investment does not go into creating new AI models. It funds data centres used for cloud services rented to other companies, for advertising systems and for running AI tools for paying customers. Google Cloud alone has more than $460bn of contracted future revenue, and Meta attributed part of its 33% revenue growth in the first quarter of 2026 to AI-driven ad targetingas reported by Yahoo Finance. A Chinese lab copying a model's answers takes nothing away from these businesses. Only the money spent on developing new models is truly exposed, and it represents a fraction of the headline figures.

It is also worth asking who captures the benefit of cheap Chinese models. Of 178 large open-weight models released by Chinese labs in 2026, 59% carry the Apache 2.0 licence and 22% the MIT licence, both of which allow free commercial use. The main winners are therefore developers and businesses around the world, including American start-ups. In 2025, Andreessen Horowitz, one of Silicon Valley's largest venture capital firms, reported that up to 80% of the start-ups pitching it with open models were building on Chinese ones. Seen this way, Chinese labs subsidise the global software industry as much as US companies subsidise Chinese labs.

Copying is also becoming harder. US labs are deploying tools that detect suspicious usage patterns and can deliberately degrade the answers sent to suspected copiers. The scale of the alleged campaigns also varies widely. The 150,000 exchanges Anthropic attributes to DeepSeek are small compared with the data needed to train a modern model, which suggests that most of its progress came from its own work..

Finally, American labs continue to capture the vast majority of AI revenue. Anthropic's annualised revenue passed $65bn at the end of July 2026 and OpenAI's reached about $40bn. By comparison, Epoch AI estimates that no other lab, Chinese ones included, earns more than $1bn.



Source: Bloomberg, Reuters, Epoch AI

The gap is narrowing faster than expected

Two recent developments weaken the case against the subsidy argument. The first concerns revenue. Until recently, open-weight models won a lot of usage but very little money. In September 2026, open-weight models handled 71.5% of the traffic on Vercel's AI gateway, a service developer use to access different models. They also captured 23% of spending on the platform, up from under 4% in June. Between August and September alone, Anthropic's share of spending on the platform fell from 63.7% to 43.9%. Although these figures cover a single platform, the trend they reveal is significant.

The second concerns performance. A model trained on another model's answers usually comes close to its teacher without surpassing it, which would keep Chinese labs one step behind. Kimi K3 challenges that assumption, having beaten Anthropic's Fable 5 on Arena's coding ranking and placed above Opus 4.8 in its general ranking. These rankings are based on user votes for specific tasks rather than independent tests, so it would be premature to call it a Chinese lead. But the idea that Chinese models will always remain a few months behind looks less certain than it did at the start of the year.

Investment implications

For the large US technology groups, the impact remains limited. Their profits depend primarily on cloud services and advertising, with less reliance on exclusive access to the most advanced models. Cheaper AI has also stimulated demand, benefiting chipmakers in the process. The pressure is more direct for frontier AI labs, as they can still charge a premium for the most complex tasks, but pricing power is weakening across everyday applications. As Anthropic and OpenAI consider potential IPOs, their financial disclosures should provide a better indication of how this pressure is affecting margins.

Conclusion

The argument that America bears the costs while China reaps the benefits has some merit. Chinese labs build on US breakthroughs, sometimes through practices that American companies allege are illicit, and their low-cost open models are gaining market share. Chinese companies are also investing heavily, developing their own innovations and generating far less revenue than their US counterparts. A more accurate assessment is that America's technological lead generates returns for an increasingly limited period. Recent developments suggest this period is narrowing rapidly, and investors should monitor whether growth in demand is sufficient to offset the decline in pricing power.

Disclaimer: This document is "powered by Syz" meaning that it has been prepared in collaboration with Bank Syz Ltd ("Syz") in Switzerland exclusively for Theia Investments Ltd. Theia Investments Ltd brought this document to you for information purposes only without intervention from Syz. Neither the information nor any opinion expressed herein may be deemed as an advertisement, an offer or an invitation to make an offer or a personalised recommendation, to buy or sell any financial instrument or any other investment or currency in any jurisdiction. This document is not directed to any person in any jurisdiction where, by reason of that person's nationality, residence or otherwise, its dissemination or use is prohibited. It cannot be quoted, reproduced, disclosed, distributed or otherwise referred to, in whole or in part, to any third party without Theia Investments Ltd's prior written consent. This document may contain information on past performances of financial products, all kinds of economic and financial valuations and analyses, as well as opinions of Syz, Theia Investments Ltd or of third parties. Any market valuations, terms and calculation methods, yields and opinions expressed are subject to change without notice. The information and opinions contained herein have been compiled and provided in good faith, based on information obtained from sources believed to be trustworthy. However, such information has not been independently verified. While it is believed to be correct at the time used and derived from reliable sources, no guarantee, representation, or warranty of any kind, whether express or implied, is provided by Syz, Theia Investments Ltd or by relevant data vendors and other sources with respect to the accuracy, reliability, currentness or exhaustiveness of the information communicated, the methodology employed or the prices contained herein. Past performance is not a reliable indicator of current or future results. Any reference to the past performance of an investment does not guarantee its future performance. Fluctuations in exchange rates may adversely affect the value, price or income of any investment. Investing in international markets may involve additional risks, such as social and political instability, market illiquidity, exchange-rate fluctuations, a high level of volatility and limited regulation. Investments in emerging markets are by their nature riskier and potentially more volatile than those inherent in established markets. Prospective investors should not construe the contents of this document as containing legal, tax, or financial advice. This information is purely objective since it does not take into consideration the personal situation (including financial knowledge and experience), the investment objectives and risk profile of any recipient. Moreover, certain financial instruments mentioned (in particular, but not exclusively, derivatives, options and futures) involve a high risk of loss and may not be suitable. Accordingly, prospective investors should consult their stockbroker, bank manager, lawyer, accountant or other professional advisor on such matters, including but not limited to the possible tax consequences of investing in any financial product mentioned herein, and make their own assessment of the suitability of any investment product described in this document and should understand that future expectations may not be realised. Should a recipient make an investment on the basis of the present communication, the recipient alone shall bear the consequences thereof. In no event shall Theia Investments Ltd or Syz, their directors, employees, affiliates or agents be held responsible for damages of any kind arising out of the use of all or part of the information provided herein.

POWERED BY Syz PRIVATE BANKING

THEIA INVESTMENTS
YOUR BESPOKE FAMILY OFFICE

Theia Investments Limited
#814, Liberty House, DIFC, PO Box 566720, Dubai, UAE
theiainvestments.com