


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.

Charles-Henry Monchau, CFA, CAIA, CMT
Chief Investment Officer
charles-henry.monchau@syzgroup.com

Assia Driss
Syz Research Lab Team Coordinator
assia.driss@syzgroup.com

Romain Marcozzi
Syz Research Lab Team
romain.marcozzi@syzgroup.com

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 committee, 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.

Alleged distillation campaigns against Claude

~16m exchanges via ~24,000 fraudulent accounts (single-source allegation)




Source: Anthropic, "Detecting and preventing distillation attacks" (Feb 2026). Accused labs had not responded in the report.

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.

2026 capex: US hyperscalers vs. Chinese majors

Guidance midpoints; Chinese figures converted at ~7.1 RMB/USD



Sources: company guidance via FT/TrendForce (US, Q1-Q2 2026); ByteDance RMB 200bn plan; *Alibaba FY2026 actual - RMB 126bn.


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.

OpenRouter: Chinese models took the token volume

Share of tokens routed on OpenRouter, selected data points



Sources: OpenRouter data via CNBC, Sacra, a16z/Wing VC reporting (2026). Snapshots differ in scope; indicative only.

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.

The compute asymmetry behind the efficiency story

Share of global GPU-cluster performance



Source: Epoch AI GPU clusters dataset (2025; covers ~10-20% of global capacity); Fed FEDS Notes (Oct 2025).

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 targeting as 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.

Token share is not revenue share

Annualised revenue run rates of leading labs



Sources: Bloomberg/Reuters (Anthropic), Bloomberg/Semafor (OpenAI), Epoch AI (other labs <\$1bn). Run rates, not audited.

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.

Welcome to Syzerland®

For further information

Banque Syz SA

Quai des Bergues 1
CH-1201 Geneva
T. +41 58 799 10 00
syzgroup.com

Charles-Henry Monchau, CFA, CAIA, CMT

Chief Investment Officer
charles-henry.monchau@syzgroup.com

Assia Driss

Syz Research Lab Team Coordinator
assia.driss@syzgroup.com

Romain Marcozzi

Syz Research Lab Team
romain.marcozzi@syzgroup.com

This marketing document has been issued by Bank Syz Ltd. It is not intended for distribution to, publication, provision or use by individuals or legal entities that are citizens of or reside in a state, country or jurisdiction in which applicable laws and regulations prohibit its distribution, publication, provision or use. It is not directed to any person or entity to whom it would be illegal to send such marketing material.

This document is intended for informational purposes only and should not be construed as an offer, solicitation or recommendation for the subscription, purchase, sale or safekeeping of any security or financial instrument or for the engagement in any other transaction, as the provision of any investment advice or service, or as a contractual document. Nothing in this document constitutes an investment, legal, tax or accounting advice or a representation that any investment or strategy is suitable or appropriate for an investor's particular and individual circumstances, nor does it constitute a personalized investment advice for any investor.

This document reflects the information, opinions and comments of Bank Syz Ltd. as of the date of its publication, which are subject to change without notice. The opinions and comments of the authors in this document reflect their current views and may not coincide with those of other Syz Group entities or third parties, which may have reached different conclusions. The market valuations, terms and calculations contained herein are estimates only. The information provided comes from sources deemed reliable, but Bank Syz Ltd. does not guarantee its completeness, accuracy, reliability and actuality. Past performance gives no indication of nor guarantees current or future results. Bank Syz Ltd. accepts no liability for any loss arising from the use of this document.