



By: TA Analysis

Who will pay for cheap Chinese AI?



In an analysis published on 26 October, Reuters calculated that **China's STAR50 stock index** has lost about 40 per cent of its value in less than four months, with an approximate \$480 billion drop in the market value of its constituent companies.

Shares in processor makers MetaX and Moore Threads have halved since July; optical equipment makers for large data centres Zhongji Innolight and Eoptolink have lost about 60 per cent from earlier peaks, while shares in robot maker Unitree have halved since listing.

Reuters identifies no clear trigger for the decline, and investors are questioning how long demand for expensive AI equipment will last.

Shares in the STAR50 index are still worth roughly one hundred times the companies' annual profits even after the fall, and the market is worried about possible restrictions on the export of Chinese equipment and about the income that past investments will generate.

Falling share prices in themselves do not prove that demand for equipment has also fallen. Alibaba and MiniMax are not members of that index.

Alibaba's shares are listed in Hong Kong and New York, and MiniMax's in Hong Kong. Their results, however, show how growth in demand for computing power translates into revenue that can finance new investment.

In addition to companies that develop models, cloud service providers and companies that deploy these models also buy equipment.

Investments are growing faster than earnings

Alibaba's report for the quarter ended 30 June shows how differently businesses within the same company are developing.

The AI Cloud and Compute Services division, which rents computing capacity and includes processor development, generated revenue of 48.4 billion yuan, up 45 per cent year on year.

Its adjusted EBITA, a measure of earnings before interest, taxes and amortisation of intangible assets, was 5.6 billion yuan.

The separate AI Labs and Applications division, which develops **Qwen models** and products based on them, reported revenue of 3.3 billion yuan and a loss of 13.9 billion yuan on the same basis. A year earlier, the division's loss was 3.2 billion yuan.

At the same time, Alibaba earns money by renting out infrastructure but spends significantly more on developing models than it charges for them.

The group's capital expenditure rose 75 per cent to 67.7 billion yuan, almost three times the 22.9 billion yuan the company generated from operations.

Free cash flow was negative 44.7 billion yuan, down from negative 18.8 billion yuan a year earlier, which Alibaba attributed to investments in cloud infrastructure.

For Alibaba, it remains an open question whether the investment will pay off, although it still has the money to fund it

The company still had 474.5 billion yuan in cash and other liquid investments at the end of June. For Alibaba, it remains an open question whether the investment will pay off, although it still has the money to fund it.

In the first six months of 2026, **MiniMax** increased revenue from access to models for developers and enterprises from US\$9.2 million to US\$73.9 million, with that business accounting for 63.4 per cent of its total revenue.

As much as 60.8 per cent of sales were made outside mainland China, confirming that the

company already has customers in international markets.

Gross margin increased from 12.1 to 17.9 per cent, which the company attributes to more efficient infrastructure.

Adjusted net loss, however, rose from US\$138.7 million to US\$293 million, primarily due to new model development, on which MiniMax spent US\$296.9 million in the half-year, more than double revenue.

The loss according to international accounting standards decreased to US\$358 million, mainly due to smaller changes in the value of financial liabilities.

MiniMax then raised approximately HK\$15.88 billion net, about US\$2 billion, in July through the sale of new shares and convertible bonds.

The amount raised is many times its half-year adjusted loss, although some of the money comes from debt maturing in 2027 if it is not converted to shares first.

The new capital gives the company time to grow, but future investors will increasingly focus on how much of that growth can be financed through sales.

A cheap service that the world is quickly adopting

Chinese manufacturers compete on price, which makes it easier for them to enter foreign markets.

DeepSeek offered the V4 model in late April at a price that the South China Morning Post estimated, for a typical conversation at the promotional rate, to be about **97 per cent lower than that of OpenAI's GPT-5.5**.

In May, it declared a 75 per cent discount on the V4-Pro model permanent, and at the end of June it announced double tariffs during peak hours.

Even the most aggressive price-cutting manufacturer has to take account of labour costs and available computing capacity.

A report by the American platform **OpenRouter** on 22 September stated that, during the previous week, models from Chinese companies had processed 69 per cent of the total volume of text on the platform.

The data measures the amount of text that computer systems process. It does not show how many people use them or how much money their makers receive.

OpenRouter is used mainly by developers, who choose between different models according to price and quality, so that share does not represent the share of Chinese companies in the global market.

For the same volume of processing, a cheaper model can generate significantly less income than a more expensive competitor

For the same volume of processing, a cheaper model can generate significantly less income than a more expensive competitor.

There is an additional gap between the spread of Chinese technology and the earnings of its creators.

OpenRouter allows users in the US and Europe to choose a Chinese model that runs in a data centre in their region.

In that case, the request is processed by the local infrastructure provider, and the Chinese company does not have to participate in providing the individual service.

The model gains international use, while part of the revenue from its deployment is captured by American or European providers of computing resources.

According to an August review of the **Hugging Face** platform, Qwen models were downloaded

approximately 2.05 billion times from the beginning of the year, and users posted 151,448 customised versions of them on the same platform.

Alibaba cites more than three billion downloads and 300,000 derived models across a broader set of channels.

Those numbers cover different scopes, and downloading the programme does not imply that its author is paid for each subsequent use.

Alibaba has announced that it wants to charge the biggest users of its open models, but has not yet explained how. The significance of the model's wide acceptance should also be sought in the jobs it enables for other companies.

The factory can gain even when the laboratory loses

In a research paper published on 23 March on the website of the US Commission on Economic and Security Relations with China, analyst **Ngor Luong** described two related processes.

Developers download models, modify them and publish customised versions, while factories, logistics companies and robot manufacturers implement such systems and refine their use based on production experience.

The paper was published as part of a research working group and does not represent the officially adopted position of the Commission.

For the Chinese economy, the second process could be more important in the long run.

A factory that uses a low-cost model to detect failures before downtime, reduce waste or organise deliveries more efficiently gains a benefit that the model manufacturer does not have to charge for directly.

US export controls make training the largest models most difficult, while smaller systems adapted to industrial tasks can run on more modest equipment

According to Ngor Luong's analysis, US export controls make training the largest models most difficult, while smaller systems adapted to industrial tasks can run on more modest equipment.

Such applications can increase the productivity of Chinese manufacturers even when the companies developing the models are still losing money.

How large that benefit will be depends on the quality of industrial data and on how successfully the systems are integrated into production.

Domestic processors face production limits


On 13 January, the United States began individually reviewing requests to export **Nvidia's H200 processors**, AMD's MI325X and similar devices to China, subject to security conditions.

Beijing has been restricting their purchase to support domestic manufacturers, but in July, **The Information reported** that it had announced the possibility of purchasing limited quantities of H200 for selected companies.

Nvidia confirmed in an August financial report that it had shipped only a fraction of the quantities covered by US licences and that revenue from the H200 accounted for less than 1 per cent of data centre sales in the quarter ended 26 July.

In the forecast published on 26 August for the next quarter, the company did not include revenue from the sale of processors for data centres in China.

Domestic manufacturers have a chance to capture that market if they can produce enough high-quality chips.



In September, DeepSeek released six software tools tailored to Huawei's Ascend processors

In September, DeepSeek released six software tools tailored to Huawei's Ascend processors.

Bloomberg reported that **DeepSeek** planned to use at least 160,000 Ascend 950DT processors in a new computing centre in Inner Mongolia, primarily to run already trained models, while continuing to use Nvidia equipment to train the most demanding systems.

According to the same report, a shortage of advanced memory could limit Huawei's total production of that processor to a few hundred thousand units in 2026, so delivery of DeepSeek's entire order could take more than a year.

The ability of Huawei processors to replace American ones in the most demanding tasks, and the speed of their manufacturing expansion, remain significant uncertainties in China's strategy.

Beijing encourages enforcement but fears bubbles

Chinese authorities are already trying to curb cut-throat competition in other industries. On 1 July 2025, the Central Financial and Economic Affairs Commission requested action to suppress unregulated price-cutting.

The campaign particularly affected the electric vehicle and solar panel industries, where large production capacities have long put pressure on earnings. In guidelines published on 9 October 2026, Beijing encourages the wider use of AI in industry but warns against reckless investment and the creation of bubbles in new branches of the economy.

There is no confirmed decision specifically to limit price competition among AI model manufacturers.

Most of the data so far points to continued international expansion of Chinese models

Most of the data so far points to continued **international expansion** of Chinese models, with increasing pressure on manufacturers to demonstrate sustainable revenues.

Over the next two years, the most expensive models are likely to be developed by a smaller number of well-financed companies, while cheaper versions are likely to be adopted by a growing number of foreign companies and Chinese factories.

Crucial factors will include how far revenue from using the models increases, how many domestic processors China can produce, and how Beijing balances support for low-cost AI deployments with efforts to curb unprofitable investment.

This trajectory could lead to wider use of AI products in China, with their development increasingly concentrated in the hands of a few companies.