**LIANG DU**

Chief Executive Officer
Prescient China
BBusSci & MBA (UCT), CFA, CQF

## HISTORY IN THE MAKING

After 21 years in the industry, I have seen plenty of bubbles. The AI cycle may be the grandest: a bubble with enough zeroes that we have quietly retired "billions" and moved straight to "trillions." Cash flow has been sent to the compliance basement, price-to-earnings has been declared a quaint relic of the pre-transformer era, and 100x sales can apparently pass for discipline if the logo contains the letters AI. I have even seen this logic inside our own team: a Chinese chip company at 10x sales looked "cheap" only because the peer group was apparently drawn from the outer solar system.

To put this in context, remember the glory days of EV speculation. Every Chinese EV-adjacent company seemed to come to market at 20x revenue. At its peak, NIO reached a market capitalisation of roughly US$92 billion - larger than BMW and Mercedes combined. Today it is nearer US$13 billion, despite revenue having increased by about 700% over the period. In the rally, investors did not buy companies; they bought the futures of transportation. EV makers, glass makers, motor makers, battery makers, tire makers, copper companies - if it touched an electric vehicle, it acquired a multiple. In the end, the only part of the supply chain that consistently made real money was batteries; the rest discovered competition, price cuts, and the small indignity of market caps going down. Hundreds of billions of dollars later, one is tempted to ask: does any of this feel familiar?

Today AI and almost every supplier of AI trade at price-to-sales ratios ranging from 10x to 100x. We can, of course, forget about fussy antique measures like P/E and free cash flow; those belong in museums, next to fax machines and risk committees that still ask about payback periods. The US alone has invested roughly US$750 billion in data centres. China has committed a further US$290 billion over five years. Chip and memory prices have tripled in a year, and the largest share in the CSI 300 today is a company that makes cables connecting server racks. There is a kind of poetry in that: the future of intelligence, valued at trillions, may depend on how much investors are willing to pay for very good wiring.

Why does this matter? Because we may be watching history be made: the arrival of the world's first trillionaire, Elon Musk, through the listing of SpaceX. The prospectus opens with 15 pages of beautiful rockets and space - the sort of material that reminds you that capitalism can still afford good graphic design. But read the document more closely and the story becomes less "space company" and more "AI infrastructure company wearing a spacesuit."

- Space, as a standalone business, appears financially immaterial to the US$2 trillion valuation.
- The valuation is premised largely on enterprise AI spend for a product SpaceX does not currently have.
- The company is building data centres, using SpaceX as guarantor for the debt, with equity from board-linked Valor Capital and private debt from Blue Owl. SpaceX does not appear to need the additional compute itself - no major corporate buyer currently seems to rely on Grok at scale - so the capacity is leased to Google and Anthropic, both of which are already spending heavily on their own data centres.

To summarise: the world's first trillionaire may have listed a loss-making company at 100x sales, or 33x forward sales if we are feeling generous, with negative operating cash flow, with an AI model that currently trails open-source Chinese alternatives, and a valuation not far from Microsoft's. As of July 1, SpaceX's market capitalisation is about US$2 trillion, versus roughly US$2.8 trillion for Microsoft. This is not necessarily wrong; markets are allowed to dream. It is just an unusually expensive dream, and one financed with a lot of very real private debt and retail exuberance.

Another way to frame the issue: total AI revenue is currently approximately US$120 billion. I pay OpenAI US$20 per month. From that US$20, OpenAI must train new models, run my queries in data centres, buy or rent compute, fund engineers, and keep pace with competitors who are also spending as if the discount rate was abolished by executive order. Data centres must pay for electricity, depreciation, hardware, software, and a small number of very highly paid humans to keep the machines from becoming decorative furniture. Nvidia must keep selling chips and keep launching new chips that are dramatically better than the old chips, so customers continue upgrading. But if the new chips are too much better, the old data centres may become uneconomic, because the electricity bill will make yesterday's hardware look like a gas guzzling V12. The capex cycle therefore needs to continue forever, which is a lovely business model until someone asks who ultimately pays for it. At the moment, the answer appears to be: my US$20 subscription, with heroic assumptions.

Now add the market leaders themselves, many of which hope to IPO before year-end and may be starting a price war just as investors begin asking for profitability. Every foundation-model company eventually has to prove that US$20 per month is not merely a nice donation to the GPU sector. Meanwhile, Chinese models at roughly 90% of the capability are available at prices approaching zero, with China treating AI more like a public utility than a luxury subscription. Suppose, very generously, that three billion people sign up for US$20 per month AI. That is US$720 billion of annual revenue. A sceptic might still ask whether this can support the trillions of market capitalisation now attached to AI, data centres, semiconductors, cables, power equipment, and everything else that has been sprinkled with transformer dust. AI is unquestionably one of the most important tools mankind has created. But the internet was important too, and that did not prevent investors in 2000 from having a deeply educational experience. Figure 1 below shows a stylised version of these circular relationships. For clarity, "circular" is doing quite a lot of work here.

Figure 1: Circular Relationships in the AI Data Centre Ecosystem

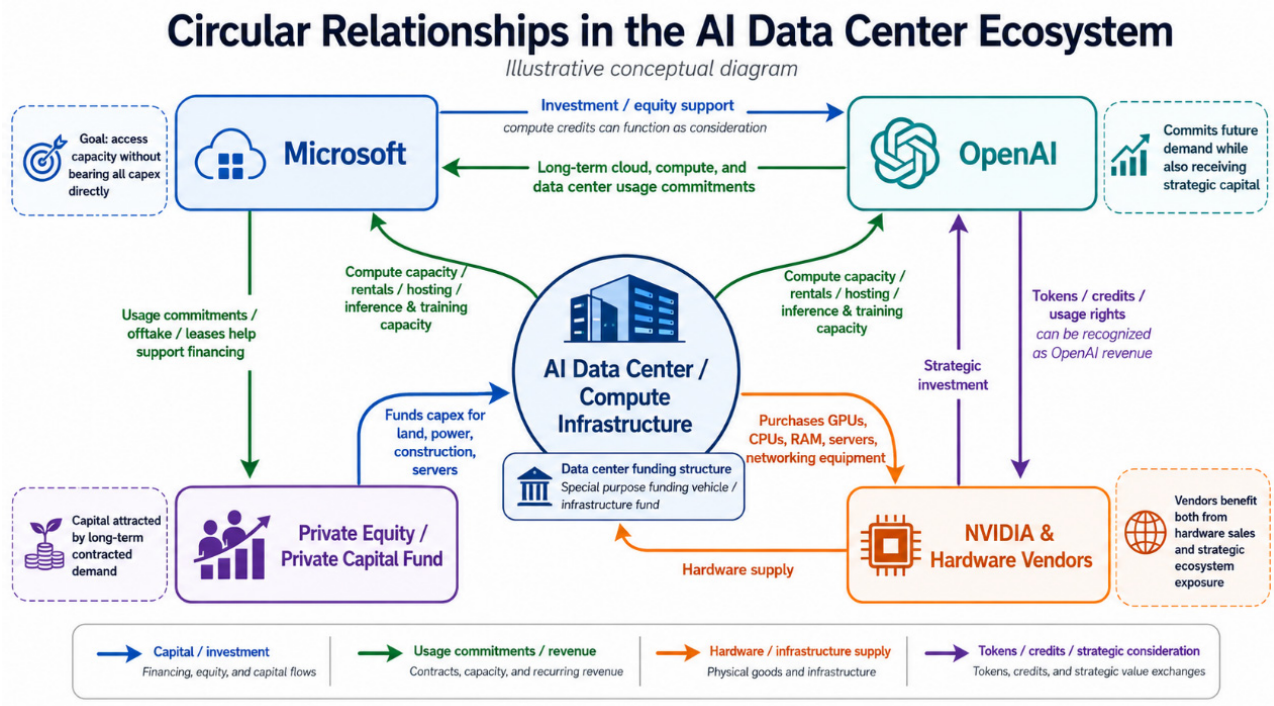


*Source: Prescient, Illustrative conceptual diagram generated by AI, as at 30 June 2026.*

Looking at the illustration, I am reminded of a famous South African company that also had somewhat of a circular relationship with investing and providing debt guarantees to off book subsidiaries. Turns out they were merely a pioneer in capital market, a model perfected by AI a decade later.

Let us bring this back to China, starting with the current top 10 shares in the CSI 300.

Figure 2: Top 10 Shares in the CSI 300 Index

| Rank | Security Name | CSI300 Weight |
|---|---|---|
| 1 | Zhongji Innolight | 5.01% |
| 2 | CATL | 3.68% |
| 3 | Eoptolink | 2.99% |
| 4 | Kweichou Moutai | 2.62% |
| 5 | GigaDevice | 1.92% |
| 6 | Ping An Insurance | 1.80% |
| 7 | Cambricon | 1.77% |
| 8 | China Merchants Bank | 1.55% |
| 9 | Zijin Mining | 1.47% |
| 10 | NAURA | 1.35% |

*Source: Prescient, Bloomberg, as at 30 June 2026.*

Zhongji Innolight, Eoptolink, GigaDevice, Cambricon, and Naura are all represented in the top 10. The aggregate price-to-**sales** ratio of the five listed AI-adjacent names is about 42x. In market capitalisation terms they are still small versus the US mega caps, but the investment thesis increasingly depends on the value and spending plans of US AI companies. In other words, Chinese equity investors are not only buying China's AI story; they are also buying the continuation of the US AI capex story. This is an elegant global supply-chain arbitrage, and also a very complicated way of saying: please keep ordering GPUs.

A second China point is even more important. The Chinese government has made clear through policy that it wants AI to function as a utility. China's model is to make AI open-source, cheap, and widely available to a vast population of consumers and businesses. The AI layer itself may not earn very much, but it can raise productivity across billions of users and millions of companies. That is great for users and the real economy. It is less obviously great for AI foundational model, and cloud services gross margins. Historically, when China mass-produces something, it tends to get cheaper and better, and the margin tends not to exist for very long. If the US AI thesis depends on permanently excluding Chinese supply through politics, regulation, export controls, cyber conflict, or some other geopolitical miracle, then the thesis is not just a technology call; it is a geopolitical call. That is a much harder thing to underwrite at 50x sales. More likely, the bubble does not have to burst tomorrow. It can simply deflate, as EV shares did over three years. The lesson from EVs was not subtle: better, cheaper cars plus many new entrants tend to lower everyone's earnings. The product can improve; revenues can explode while the equity story deteriorates. This is a phenomenon markets somehow rediscover every cycle, generally at high cost.

As a systematic factor-based manager, the practical conclusion is straightforward. This year, the only factors that have really worked are momentum and growth - the classic companions of bubble periods, because rising prices explain themselves until they don't. Looking at our three pillars of asset allocation - policy, valuation, and sentiment - policy remains supportive, valuations are neutral but moving toward negative, and sentiment remains positive. This does not mean we are complacent. We have already built models and asset-allocation scenarios that can reduce exposure to the most crowded areas at the push of a button. The hardest part, of course, when we push the button. Maybe I should put the trillion-dollar AI investment to good use and ask the model when the bubble ends. I am sure it will respond with confidence, a disclaimer, and quite possibly a recommendation to buy more compute!

Until next quarter!