China's CXMT Is Set to Challenge DRAM Incumbents

CXMT IPO, SK Hynix, Micron, Samsung Competition, Process Node Deficit, China HBM, Wafer Adds, Memory LTAs

RAY WANG, MYRON XIE, DYLAN PATEL, AND 3 OTHERS

JUN 23, 2026 • PAID

188 2 27 Share

We were the first to describe the memory shortage coming from AI's insatiable usage in reasoning and agentic flows in late 2024 on the newsletter. We have since previously published multiple in-depth pieces on memory, as well as detailed coverage of CXMT and China's compute. With CXMT set to IPO in the coming months, we believe a dedicated deep dive on them specifically is warranted. The company is likely to become the largest semiconductor IPO in China and mark a major milestone for the country's leading memory manufacturer, which is also destined to compete only more fiercely with the leading memory suppliers of Samsung, SK Hynix, and Micron from here.

Our latest on memory can be seen here:



Memory Mania: How a Once-in-Four-Decades Shortage Is Fueling a Memory Boom

DYLAN PATEL, RAY WANG, AND 3 OTHERS • 2 7

Read full story →

Or other older but excellent memory or China compute pieces below:

1. Huawei Ascend Production Ramp: Die Banks, TSMC Continued Production, HBM is The Bottleneck
2. Huawei AI CloudMatrix 384 – China's Answer to Nvidia GB200 NVL72
3. Scaling the Memory Wall: The Rise and Roadmap of HBM
4. The Memory Wall: Past, Present, and Future of DRAM
5. Scaling Laws - O1 Pro Architecture, Reasoning Training Infrastructure, Orion and Claude 3.5 Opus "Failures", Inference Tokenomics of Test Time Compute

flows, and the patience of state-venture capital, together of which are turning the company toward indigenous innovation.

The Silicon Valley Returnee

Zhu Yiming, the founder of CXMT, earned his undergraduate degree in physics at Tsinghua University in 1994 and came to SUNY Stony Brook for his graduate study in electrical engineering. He then worked in Silicon Valley and became a project lead at MoSys (Monolithic System Technology) around 2001. In 2005, he went back to China with a set of SRAM patent and US$100,000 in seed money, founding GigaDevice, which later became well known for its SPI NOR flash and microcontrollers, a fabless design house that grew into one of the world's top NOR flash suppliers. But the global NOR flash market is much smaller compared to DRAM or NAND flash. Zhu Yiming always dreams big, and not surprisingly he chose to march into the DRAM business.

DRAM, however, is not a design-house game in which one can stay fabless. DRAM is capital-devouring, IP-fortified, and manufacturing-bound, and by 2016 the whole industry was dominated by three survivors, Samsung, SK Hynix, and Micron, fortified by four decades of patents and capital that no newcomer had breached. Zhu's SRAM patents and GigaDevice's NOR franchise gave him neither a DRAM cell, nor a DRAM process, nor any cover from the incumbents' patents. Thus, when Zhu and the Hefei municipal government launched the DRAM venture in 2016, the "506" project that became CXMT, the core technology had to come from somewhere else entirely.

It came from a dead company in Germany.

The DRAM Foundation: Inheriting Qimonda

The dead company was Qimonda. Though the company went bankrupt in January 2009 due to the 2008 Global Financial Crisis and the following dramatic memory downcycle, it was then the leading DRAM player in Europe. A subsidiary of Infineon, which originated from Siemens, it offered a rare alternative: a deep DRAM patent base and a cell architecture, both originating outside the dominant Samsung-SK Hynix-Micron triangle. In June 2015, Polaris Innovations, a subsidiary of WiLAN, a Canadian patent-monetization firm, bought roughly 7,000 Qimonda patents and applications from Infineon for about 30 million euros. In December 2019, Polaris and CXMT signed an agreement: a license to a large set of

A major technology CXMT inherited from Qimonda and then developed its 46nm-class BWL (buried wordline) cell, which CXMT scaled toward the 10nm class. BWL is the load-bearing idea. Instead of routing the access transistor's gate across the wafer surface, BWL sinks it into a trench beneath the bitline. It pulls the gate off the surface so the cell collapses to a $6F^2$ layout (versus $8F^2$). It also lengthens the channel without spending surface area, suppressing the short-channel leakage that wrecks retention. And it cuts gate-to-bitline parasitic capacitance. Buried wordline plus stacked capacitor is the architecture all three leading players run today. The trench holdout died holding the stacked/BWL escape hatch, and that is exactly what CXMT picked up.

Talent Flows: From a Frozen Blueprint to Living R&D

Besides patents, the more durable thing CXMT pulled out of Qimonda's collapse was its engineers. Qimonda's Xi'an R&D center, with 400-500 engineers, was one of the largest Qimonda built outside Germany. After Qimonda's collapse, though the whole Xi'an R&D center was acquired by Tsinghua Unigroup, the wider diffusion of individual talents benefited CXMT. In addition, CXMT successfully attracted Karl-Heinz Kuesters, a senior engineer, from Qimonda's German site to Hefei, China. Kuesters had spent 24 years as a VP of technology and pre-development across Siemens, Infineon, and Qimonda. The pre-development line Kuesters ran was the stacked-capacitor work, the architecture CXMT actually builds on. He joined CXMT as a technical consultant, for which the EE Times called Kuesters the company's "ace in the hole." What Kuesters brought was the part of Qimonda's legacy that neither the patents nor the 2.8 TB of documents carry: the tacit know-how. Having led DRAM development for two decades, Kuesters could tell CXMT's engineers which of Qimonda's design choices to keep and which to discard, and how to move a cell that worked only in the lab into volume production, the integration and yield judgments that no patent records.

The American side shows the same pattern. Ping Er-xuan, CXMT's VP for future-technology assessment and the public face of its roadmap (the "46nm to the 10nm class" framing), came not from Qimonda but from a US career at Micron, SanDisk, and Applied Materials, where he ran memory and materials technologies. Ping brought process-and-materials depth and an emerging-memory view: individual mobility.

CXMT has also recruited heavily from Korea and Taiwan. Korean prosecutors have charged former Samsung employees with leaking technology to CXMT, and dozens of Korean

process development.

This is the part that bears on where CXMT is heading. The Qimonda patents were always a finite, expiring asset. What lets CXMT keep moving, from G4 to G5 and now into HBM, is the assembled capability of domestic talent, Chinese nationals who worked at foreign companies and returned, as well as some experts from foreign firms, not the documents. The inheritance only got it started. The talent turned a foreign legacy into in-house R&D juggernaut. However, it took nearly a decade for CXMT to generate profits. The question is who has patiently funded CXMT's development and borne its near-decade losses?

The Patience of State-Venture Capital

It is difficult to not attribute at least part of CXMT success to the strong support of local and central government in China. Hefei municipal government is one of the great examples. Hefei is one of the tech innovation hubs in China, well known for its patient state-venture capital nurturing successful companies over the past two decades, from BOE (the world's leading display panel maker) to NIO (a leading EV maker) and now to CXMT. In particular, the Hefei municipal government did two things for CXMT. First, the Hefei government helped CXMT build a local supply chain around its fab. Hefei's playbook is to take a large equity stake in an anchor "link-leader" and then draw the rest of the chain in around it. The city did that for BOE in displays and NIO in EVs, and from 2016 it replicated the same playbook for CXMT. Surrounding CXMT's plant in Hefei's airport-zone industrial park, the government produced a dense local cluster. Two packaging-and-test houses, Peyton and Xinfeng, are located within a street or a wall of CXMT's fab, with Xinfeng booking over 99% of its revenue from CXMT. An on-site bulk-gas plant run by Guanggang supplies most of CXMT's needs, while wafer-reclaim capacity from Zhiwei Semiconductors a subsidiary of Zhichun Technology (至纯科技), is located in Hefei's Xinzhan Hi-Tech District. The state-venture-capital vehicles have also taken outright control of an upstream chip-molding equipment maker, Wenyi Technology. Such a local supply-chain cluster gives CXMT a localized industrial base.

In addition, Hefei's state-venture capital could afford to lose money for a very long time. Unlike a private venture-capital fund answerable to LPs that expect a return on a fixed timetable, Hefei's state-venture capital, ultimately backed by the city's municipal and development-zone state entities, faced no such clock. They kept funding a company that, even after turning its first annual profit in 2025, still carried an accumulated deficit of

project's first phase (RMB 14.4bn of 18bn). Across successive rounds Hefei's vehicles were diluted, but they never sold down and never walked away. By the IPO, the largest holder, Hefei's Qinghui Jidian, at 21.67% and state-venture capital vehicles together holding over 30%. That willingness to treat a fab as a decade-long bet rather than a fund-cycle return is the catalyst that the technology and the talent both depended on.

From Inheritance Toward Independence

Put the three threads together and CXMT's first decade resolves into a single arc. Qimonda supplied the foundation, with a licensed patent base and a cell architecture from outside the incumbent triangle. Talent supplied the motion, with key figures like Kuesters and Ping as well as the returnees from the American majors, and the contested hires out of Korea. Those people turned a frozen blueprint into a process that could keep scaling. Then the Hefei government supplied what the other two needed but could not generate on their own: capital, patience, and a localized supply chain. None of the three would have produced a DRAM maker alone; together they did.

In the following section, we will be discussing CXMT's financials, technology, equipment ecosystem.

Next Step after 10 Years: An IPO in the Supercycle

The past decade of CXMT's history, while impressive, may prove to be only an early chapter in the company's longer-term story. The company is now preparing for what will become one of China's largest semiconductor IPOs of the past decades, and potentially the most closely watched semiconductor listing globally this year. In December 2025, CXMT officially entered the IPO filing stage when the Shanghai Stock Exchange accepted its STAR Market listing application, following a prolonged period of market reports throughout 2024 and 2025 suggesting that the company was preparing to go public. More recently, CXMT's application advanced further through the review process by submitting the formal CSRC (China Securities Regulatory Commission) registration on May 27 and currently under final review, moving closer to a landmark listing.

While understanding strategically important, unlisted Chinese companies such as CXMT and YMTC has never been easy, CXMT's move toward a public listing has been helpful for us to understand more details about the company given its IPO prospectuses that have shed

our Memory Model, we can develop a more accurate view of CXMT's current position and a more robust forecast of its future performance.

At a high level, CXMT is clearly the fourth biggest DRAM maker globally by almost every metrics while expanding its lead over legacy memory suppliers. For the full year, CXMT revenue increased 156% YoY to ~$8.6 billion, up from ~$3.3 billion in 2024 and ~$1.2 billion in 2023. Net income also turned positive for the first time, reaching $1 billion, underscoring the company's rapid scale-up and improving profitability. Even with such impressive results, CXMT's CY25 revenue is still materially behind Samsung (~$72.3B), SK Hynix (~$52.1B), and Micron's (~$37.2B) DRAM revenue.

Source: SemiAnalysis Memory Model - sales@semanalysis.com

In 1Q26, CXMT reported revenue of $7.3 billion, representing roughly 700% YoY growth and already approaching the company's full-year revenue in 2025. Operating margins also expanded sharply, reaching approximately 70%.

is expected to be 7x y/y and reached to more than 16 billion. For full-year 2026, we believe CXMT's revenue could exceed $50 billion. If achieved, this would suggest the company has more than doubled revenue every year since 2023, and more than 6x its revenue in 2026 year over year basis, with its earnings trajectory accelerating meaningfully as both scale and profitability improve.

In our view, such significant upside to CXMT's earnings is clearly driven more by the cycle itself than company's technology or market positioning. When we take a closer look into the correlation of CXMT's ASP trajectory and bit shipment. Company's bit shipments increased by only 11% in 1Q26, while ASPs rose by roughly 57%, following QoQ ASP increases of 63% and 68% in 3Q25 and 4Q25, respectively. In other words, what really drove up company's earnings is really the explosive ASP growth rather than significant market share gains over its peers in global DRAM end-market. By bit shipment, we model CXMT's market share will increase from 9% in 2025 to 12% in 2027. While a 3% market share gain sounds minimal, it is significant for a market we size at close at $1T in 2027.

Source: SemiAnalysis Memory Model - sales@semianalysis.com

Surprising enough, the strong pricing uplift coming to CXMT is not an exception. We have been observing similar dynamics across both leading-edge and legacy memory suppliers in the DDR5, DDR4, and even DDR3 markets over the past year or so. As noted in our February memory piece, where we described the memory market as entering a "once-in-four-decades shortage," we believe DRAM pricing remains on track to double again this year driven by sustained supply-demand imbalance across these product categories. Since February, we have become even more confident in this view and believe it could potentially exceed our expectations by year-end.

What may be more interesting to readers who have not yet followed CXMT or the memory market closely is how the company's pricing compares with industry leaders. Based on our Memory Model, CXMT's DRAM ASP challenges the common misconception that Chinese memory is structurally cheaper and will flood the market, thereby pressuring global pricing.

Taking 1Q26 as an example, CXMT's DRAM ASP was only slightly below Samsung, SK Hynix, and Micron — by roughly 5–10% — in the same quarter. And as we model it out throughout 2026, we think this will still to be the case directionally although the gap will gradually expand. We believe this widening gap over the coming quarters will be driven less by inherent pricing differences but more by the change of product mix. Leading suppliers continue to benefit from a higher server DRAM mix in their bit shipments, as well as a more favorable pricing outlook for server DRAM compared with consumer DRAM.

As such, we expect this mix to increase further over the coming quarters as server becomes a larger share of DRAM end-market demand. By the end of 2027, we expect server DRAM and HBM account for well more than 50% of the total DRAM end-market. Given that server DRAM and HBM carry higher $/GB than other memory end markets, this should allow leading memory suppliers widen their ASP gap versus CXMT as server DRAM mix increases, especially considering the expected material price increase for HBM in 2027 (we have 2027 HBM pricing in our Accelerator & HBM Model and Memory Model). We also have details on LTAs for major memory buyers such as hyperscalers and Nvidia.

However, it remains far below SK Hynix at 60.4%, as SK Hynix benefits from a much higher HBM mix, which carries higher ASPs and margin last year. CXMT's ~38% margin is a significant swing from -113% in FY23 and -4.7% in FY24. Last year is not only the year CXMT reaches a record-high gross margin, but also the first year that the company achieves a positive margin profile.

Source: SemiAnalysis Memory Model, Company Reports - sales@semianalysis.com

With DRAM ASP continues to increase in 2026, CXMT's margin profile has improved further. Its operating margin reached 70% in 1Q26, compared with SK Hynix at 73%, Samsung at 81%, and Micron at 84% in the same period. Besides the strong ASP growth, firm's margin improvement also needs to thanks to company's near-total exposure to commodity DRAM, which effectively has higher margin now vs. HBM. Based on its filing, nearly all of company's bit sales are conventional LPDDR and DDR products. HBM is still a very minimal contributor to the company's revenue and earnings.

Source: SemiAnalysis Memory Model, Company Reports - sales@semianalysis.com

This becomes clearer when we run a simple cost-per-bit analysis on DDR5 products across four memory suppliers. For DDR5, we find that CXMT's cost per bit remains meaningfully higher than that of the three leading suppliers, by more than 30%. However, because DDR5 pricing was already exceptionally strong in 1Q26, we believe this still lifted CXMT's gross margin to over 70%.This suggests that the improvement in CXMT's margin profile is primarily driven by pricing, rather than by a material improvement in product competitiveness or cost structure.

Source: SemiAnalysis Memory Model - sales@semianalysis.com

In addition to printing record-level earnings, we believe company is gaining ground from the capacity perspective. By the end of 2026, we expect CXMT to reach roughly **350 kwspm**, which is only modestly below Micron's estimated ~**385 kwspm**. This would position CXMT close to becoming the industry's third-largest memory supplier, if ranked only by wafer capacity.

Source: SemiAnalysis Memory Model - sales@semianalysis.com

However, CXMT remains materially behind the two leading DRAM suppliers, Samsung and SK Hynix, which we estimate at approximately **720 kwspm** and **595 kwspm**, respectively. Next year, with the initial ramp of Shanghai Phase 1 and the full ramp of Hefei and Beijing, CXMT capacity could reach the **420kwspm range if capacity in year-end**, representing **~17% of global DRAM capacity**, up from **~13% in 2025.** In terms of bit shipments, CXMT's share of global bit shipment is expected to rise from **9% to 12% in 2027.**

CXMT's global capacity share could increase further as its Hefei site reaches full operation and the two phases of its Shanghai site continue ramping through 2028. We believe the company will reach 500kwspm of wafer capacity by the end of 2028, accounting for ~17% of global DRAM supply, up from 11% in 2025.

Source: CXMT's Heifei Site, SemiAnalysis Memory Model - sales@semianalysis.com

Given CXMT's expanding role in global DRAM capacity, as in past cycles, investors are concerning about potential supply-demand disruption from Chinese players. While these concerns are understandable, we believe they are likely overplayed at least for the next two years. We factor in incremental wafer capacity and bit shipments from CXMT and other memory suppliers—and assuming utilization rates in the high-90% range—we continue to see DRAM as extremely supply constrained. We constantly update our wafer add numbers, demand numbers, and pricing assumptions in our memory model, much faster than trendforce or sell side banks.

Source: SemiAnalysis Memory Model - sales@semianalysis.com

Looking solely at CXMT's wafer additions, we do see meaningful capacity expansion when compared with other suppliers. We expect CXMT to add roughly 85kwspm, 70kwspm and 80k each year from 2026 to 2028, versus Samsung at 15k/50k/110k, SK Hynix at 60k/60k/90k, and Micron at 30k/90k/115k. Even with these wafer additions, we expect DRAM to remain undersupplied by a high-single-digit percentage this year, widening to a low- to mid-teens in bit undersupply next year. We addressed in detail in our previous piece why DRAM is likely to remain undersupplied through potentially 2028, even with these incoming incremental wafer additions.

We see little ability for the company to irrationally accelerate capacity expansion beyond its current pace in a way that would meaningfully disrupt a market now providing an extremely favorable pricing environment because fab construction timelines are so long. This pricing backdrop has been the primary driver of the company's explosive earnings growth–which the company hopes to see it continues. Based on the fab buildout we are tracking, we also not yet see the sign of this possibility although we would like to stress that the total wafer capacity of Shanghai site could have over 400kwspm of wafer capacity in a full-ramp status.

Specifically on its wafer capacity, we see quite limited wafer allocation toward HBM. By reconciling with CXMT's IPO-related filings to date, we find that wafer allocation to HBM has been very limited, even until today. By the end of 2025, we believe only ~5 kwspm of CXMT's ~265 kwspm of capacity is allocated to HBM. We think this figure will increase to nearly ~30 kwspm and ~55kwspm by the end of 2026 and 2027 respectively. This capacity trajectory seems better align with CXMT's filings indicate that roughly 99% of revenue consists of DDR and LPDDR products in 2025, as discussed earlier.

Source: SemiAnalysis Memory Model - sales@semianalysis.com

This wafer allocation dynamics could change, however. We believe China's broader push for self-sufficiency in AI compute could conflict the company's strategic priorities, and we believe this push could intensify over time, given the HBM supply constraints discussed earlier and government's determination to address this issue.

To that end, in our estimates, we factor in government influence for CXMT to allocate more wafer capacity to HBM overtime. As a result, we expect CXMT's HBM wafer capacity to accelerate materially in 2027 and 2028 with its HBM technology improvement, supported by continued growth in China's domestic compute market. We estimate CXMT's HBM wafer capacity will reach **55kwspm and 100kwspm in 2027 and 2028**, respectively. This would increase the company's share of global HBM wafer supply from **1% in 2025** to **12% in 2028**.

It is important to remember that CXMT, unlike other memory suppliers, is not only an economically and technologically important company for China, but also a strategic asset that the country can leverage to advance prioritized policy objectives.

materially higher margins than CXMT's HBM products, while also delivering more than 3x the bits per wafer on a like-for-like basis.

Given that CXMT has not yet fully matured its HBM technology, allocating significant wafer capacity to HBM would likely generate limited profit while consuming scarce DRAM wafer capacity that could otherwise support higher-margin commodity DRAM at greater volume. In this context, prioritizing commodity DRAM is both economically rational and better aligned with CXMT's current manufacturing capabilities and pricing environment. China must allocate to HBM though as sales of HBM to China are somewhat limited besides some loopholes that allow Korean vendors to keep shipping to China.

On technology readiness, we believe CXMT is still struggling to stabilize supply for **HBM3 8-hi**, with even greater challenges in **12-hi.** On the front end, the company appears to have made progress stabilizing production of its **Gen 4** (**G4**), or **1z-equivalent** DRAM. We believe the majority of CXMT's DRAM output this year will be fabricated on the company's G4 process node. Yet, front-end wafer-sort yield should still be materially lower for the core DRAM die used in HBM, given its larger die size, more demanding cell performance, overall performance requirements versus commodity DRAM. We think the front-end wafer sort yield remain a major challenge for the company, and its gap versus its peers here is still large. While we believe yields on CXMT's G4 node have improved, we suspect it still lower than industry standard of 85-90% mature yield level for 1z given the lower margin we saw throughout 2024 and 2025. This might suggest that equipment limitations and manufacturing know-how remain persistent obstacles that CXMT will need to overcome.

Source: SemiAnalysis Memory Model - sales@semianalysis.com

For its next process node, the company's G5, or 1a-equivalent DRAM node, while can theoretically continue advancing without EUV akin to Micron in 1a process node, but it will face increasing fabrication and design challenges. These challenges will only add

bit output as well despite the company can compensate the yield loss by adding more wafers.

On top of this, we believe die stacking remains the major obstacle for CXMT's HBM. HBM stacking usually introduces significant technical hurdles, including thermal stress, die cracking, warpage, bonding defects, and yield loss across multiple stacked dies. Our understanding is that these issues become even more severe as the company attempts to move from HBM3 8-hi to HBM3 12-hi, and eventually HBM3E, given company's yet sufficient enough of know-how and manufacturing experience for 12hi or above HBM.

Die stacking is not a challenge unique to CXMT. Even leading memory suppliers are encountering difficulties. For 12-high HBM4, we understand that suppliers continue to face significant stacking-related issues, including die cracking, thermal management challenges, and yield loss.

These challenges become even more pronounced as memory suppliers seek to manufacture 16-high or even 20-high HBM. For next-gen HBM4E, we note that one reason Rubin Ultra is expected to use 12-high HBM4E rather than 16-high is supply: 16-high HBM requires higher DRAM wafer intensity and involves a more difficult manufacturing process, which can lead to greater wafer loss and lower effective supply of DRAM bits. These conditions put both memory suppliers and customer in a very tough spot given highly supply-constrained environment in DRAM.

We think there is an increasing possibility that CXMT will skip HBM3 and focus instead on HBM3E 8-hi and 12-hi. We believe this potential roadmap change is driven by two factors: 1) customer demand for more competitive HBM products in '27 timeframe, and 2) mainstream accelerators will equip with HBM3E, HBM4, and HBM4E.

Source: SemiAnalysis Memory Model - sales@semianalysis.com

On the back end, while it remains debatable whether CXMT is using MR-M
we believe the packaging challenge should be relatively more manageable, as
and its back-end partners face fewer export control constraints. CXMT has b



versus leading memory manufacturers.

Given these existing manufacturing challenges, we model CXMT's HBM3 8-hi's front-end and back-end yields at roughly 35% and 70%, respectively, implying an overall yield of only around 25%. We think this number should be lower when company attempts to produce HBM3 12hi or HBM3E 12hi given higher difficulty in die stacking and bonding. At these yield levels, CXMT's HBM output would be even more limited on the same DRAM wafer capacity than leading memory suppliers. More importantly, the resulting HBM would likely carry very low margins, especially compared with commodity DRAM in the current pricing environment.

CXMT's HBM struggles continue to be reflected in its limited product presence and slow penetration within China's AI accelerator market. We think only Huawei, Cambricon, and select emerging Chinese AI chip startups are likely to adopt CXMT's HBM, though we suspect adoption rates will be large. We actually believe domestic AI accelerator vendors would still prefer foreign HBM3, or even HBM3E, if they can secure supply through any available channel or stockpiled inventory before the export controls in December 2024. As China's domestic CSP capex and broader compute buildout are surging, there is little doubt that domestic HBM demand is also growing rapidly and should continue to increase.

With that said Huawei and CXMT will have custom HBM that is not based on the slow JEDEC standards and phys, so it will be able to close the bandwidth disadvantage.

We believe China could face a more severe HBM supply constraint than what would be implied by slow domestic HBM development alone. This constraint is likely to be further exacerbated by tight supply across all three major HBM suppliers, each of which is already restricted from selling HBM2E-equivalent or more advanced HBM products into China under the U.S. export controls announced in December 2024. Given the tight supply environment, these suppliers are likely to have even less willingness to risk violating export controls to sell into China.

However, HBM re-export and smuggling could complicate this conclusion. We understand that some Chinese companies continue to obtain HBM3 from memory suppliers, a dynamic we reported on last year and that has since been corroborated by other leading media outlets. We believe this remains the case today.

companies can access HBM. In addition, some downstream OSATs or intermediaries in third countries appear to facilitate these flows. Some entities may ship partially assembled systems or modules in forms that are not treated as fully manufactured GPUs or ASICs and are therefore still permitted to be shipped into China. The HBM can then be recovered and repackaged onto domestic Chinese GPUs or ASICs.

What the IPO Structure Reveals

CXMT could become one of China's largest semiconductor IPOs, and its ownership structure matters more than the headline financials. CXMT reported RMB7.14B of consolidated FY2025 net income, yet only RMB1.87B was attributable to parent shareholders, with 74% attributable to minority interests. The reason is the ownership architecture. CXMT holds 30.68% of the economics of Changxin Xinqiao and 31.72% of Changxin Jidian Beijing while controlling 73.01% and 75.32% of the votes through long-term acting-in-concert arrangements. That lets the company consolidate fabs it mostly does not own, so the consolidated figure overstates what public shareholders will actually receive by roughly four times.

Source: SemiAnalysis Memory Model, Company Reports - sales@semianalysis.com

CXMT exercises majority voting control of its fabs through acting-in-concert pacts, and state vehicles including National IC Fund Phase II, Hefei, and Anhui together hold well over 30% even after the listing. The arrangement looks designed to manage export-control and foreign-investor perception at a moment when CXMT's ties to the Chinese state draw the most scrutiny.

Source: SemiAnalysis Memory Model, Company Reports - sales@semianalysis.com

The raise that will get reported badly understates the listing. CXMT plans to deploy RMB29.5 billion, or about $4.1 billion, while issuing 10% to 15% of its post-IPO shares. Fully funding those uses through the IPO implies a price of roughly RMB4.41 at 10% dilution or RMB2.78 at 15%, versus RMB2.63 in the June 2025 financing. The low end represents barely any per-share appreciation despite 1Q26 revenue of $7.3 billion and net profit of $4.8 billion. At RMB2.78, CXMT would be valued at about RMB197 billion, or $27 billion, equivalent to just 1.8 times annualized first-half 2026 parent earnings. This arithmetic floor sits well below a realistic book-building valuation. This is too cheap and should be much higher valuation in our opinion.

Source: SemiAnalysis Memory Mode, Company Reports- sales@semianalysis.com

The allocation reinforces CXMT's current priorities. Of the RMB29.5 billion in planned net proceeds, RMB20.5 billion, or 69.5%, funds wafer-production-line and DRAM technology upgrades, while RMB9 billion, or 30.5%, supports forward-looking DRAM research. The prospectus discloses no dedicated HBM project and does not mention HBM. Its project descriptions focus on newer process platforms, product iteration, and the migration of existing lines toward mid-to-high-end DRAM. The IPO therefore primarily strengthens CXMT's core DRAM manufacturing and technology base, with no disclosed funding commitment to a near-term HBM expansion.

Source: SemiAnalysis Memory Model, Company Reports - sales@semianalysis.com

The size of the earnings move deserves a flag on cycle timing. CXMT guided in its December 2025 filing to an FY2025 parent loss of RMB 0.6 to 1.6 billion. Five months later the prospectus reported a RMB 1.87 billion profit, with consolidated income running past double the earlier high-end estimate. It also shows how quickly peak DRAM pricing moves the valuation denominator in either direction.

Finally, Alibaba's place on the cap table changes how to read CXMT's demand. Alibaba Cloud is at once the anchor hyperscaler customer, a near-4% holder, and an endorser, sitting alongside GigaDevice, chairman Zhu Yiming's own fabless house, at 1.8%. Domestic volume is effectively guaranteed in a way the Korean incumbents did not have in their home markets, which matters more than the small percentages suggest.

For paid subscribers, we will take a deep dive into CXMT, China's broader WFE ecosystem, the impact of export controls, and the implications for China's memory and compute ambitions. We will also discuss HBM in more detail.

CXMT's Equipment Ecosystem Under Export Controls: Domestic vs Foreign, and the Domestication Curve

Subscribe

Already a paid subscriber? Sign in

← Previous

Next →

© 2026 Dylan Patel · Privacy · Terms · Collection notice

Start your Substack

Get the app

Substack is the home for great culture