Memory and Machines

The HBM Supercycle and the Durability of AI Capex into the Second Half of 2026

Darmine Capital · Macro Research Division · Singapore

June 2026

Abstract

The memory cycle has become the fulcrum of Asian equity performance. High-bandwidth memory (HBM), the specialized DRAM that feeds AI accelerators, has driven a supercycle that lifted SK Hynix, Samsung and their upstream suppliers in Korea, Taiwan and Japan through 2025 and into the first half of 2026. This paper argues that the cycle is real but maturing. AI capital expenditure will still grow through 2026, but at a decelerating pace, because the extraordinary base effects of 2025 cannot repeat. We map the exposed nodes, quantify the HBM and conventional DRAM/NAND dynamics, and stress-test the bull case against hyperscaler discipline, Korean political recovery and Japanese monetary normalization. Our conclusion is tactical rather than structural: we recommend trimming pure-play AI hardware exposure in the second half of 2026 and rotating toward domestic-demand and services beneficiaries in India, Indonesia and Singapore, where earnings are less hostage to a single capex line.

1. The Cycle That Ate Asian Equities

For two years the single most important variable for Asian equity returns has not been Chinese stimulus, the Bank of Japan, or the Federal Reserve. It has been the capital expenditure budgets of a handful of American hyperscalers. Microsoft, Amazon, Alphabet, Meta and, increasingly, a cluster of sovereign and neocloud buyers, have collectively lifted AI infrastructure spending to a run rate that J.P. Morgan estimates channels roughly 30% of incremental global AI capex directly through Taiwan and Korea. That concentration is the defining feature of the current Asian market, and it is both the source of the region's outperformance and its central fragility.

The transmission mechanism is memory. Every AI accelerator, whether an Nvidia GPU or a custom application-specific chip, requires high-bandwidth memory stacked physically beside the logic die to feed data fast enough to keep the compute utilized. HBM is not a commodity DRAM part. It is a low-yield, capacity-constrained, high-margin product where two Korean firms and one American entrant control effectively the entire merchant supply. When AI demand accelerated, HBM went from a rounding error in the DRAM market to the profit center of the entire memory industry within eighteen months.

The result has been a classic capex-driven upcycle superimposed on high beta markets. Korea, which the house baseline has recovering to only about 1.8% GDP growth in 2026 after a 1.1% collapse in 2025 amid the constitutional crisis, has seen its index performance decouple almost entirely from its domestic economy. Taiwan's earnings revisions have run well ahead of the broader Asia ex-Japan complex, where Taiwan and Korea together are driving the 12% to 13% EPS growth we expect for the region in 2026. The question that matters for the second half is not whether the cycle exists. It plainly does. The question is how much of it is already in the price, and what happens as growth rates, not levels, begin to fade.

> **Our view:** The AI memory supercycle is genuine and still expanding in absolute terms, but Asian equity leadership has become dangerously narrow. We treat the HBM complex as a cyclical position that has entered its late-mid innings, not a structural holding to be accumulated indiscriminately into the second half of 2026.

2. Anatomy of the HBM Supercycle

2.1 Why HBM is different

Conventional DRAM is a cyclical commodity: prices swing violently with the balance between fab capacity and end demand from phones, PCs and servers. HBM behaves differently for three reasons. First, it consumes far more wafer capacity per usable bit, because stacking twelve or sixteen DRAM dies with through-silicon vias sacrifices yield, so each HBM bit crowds out roughly two to three conventional DRAM bits from the same fab. Second, it is sold largely on long-term agreements with prices fixed quarters ahead, which smooths the cycle. Third, qualification is brutal: a supplier must pass the accelerator vendor's validation, and the incumbent Korean players have a multi-quarter head start on the latest generation.

These features mean HBM has, for now, escaped the commodity trap. Blended memory margins at SK Hynix have expanded to levels not seen in prior peaks, and management guidance points to HBM supply being effectively sold out through 2026 on a contracted basis. That is the bull case in one sentence. The bear case is that every one of those three protective features is temporary. Yields improve, capacity migrates, and second sources qualify. The only question is timing.

2.2 The crowding-out effect and conventional memory

The most underappreciated feature of this cycle is the second-order squeeze on conventional DRAM and NAND. Because HBM absorbs so much wafer capacity, the memory majors have starved commodity DRAM of supply even as AI servers, which need large volumes of standard DDR5 alongside their HBM, pull demand higher. The consequence has been a conventional DRAM price recovery that has nothing to do with a consumer electronics upturn and everything to do with capacity reallocation. This is why the memory cycle has broadened: it is no longer only an HBM story but a general DRAM tightness story, which extends the beneficiary list well beyond the two HBM leaders.

> **Our view:** The crowding-out of conventional DRAM is the cycle's second engine and it is more durable than the HBM margin story itself, because it does not depend on any single accelerator roadmap. We would rather own broad DRAM tightness than a levered bet on next-generation HBM qualification timing.

3. The Exposed Nodes: Korea, Taiwan, Japan

The supply chain that manufactures the AI buildout is concentrated in three economies, each exposed at a different point of the value chain and each carrying a different macro overlay from the house baseline.

Node	Primary exposure	Cycle leverage	Macro overlay (house view)
Korea (SK Hynix, Samsung)	HBM and conventional DRAM/NAND	Highest, direct margin capture	2026 recovery to ~1.8%; high beta; political normalization after 2025 crisis
Taiwan (TSMC and packaging)	Advanced logic and CoWoS packaging	High, volume and pricing power	~30% of AI capex with Korea; anchors Asia ex-Japan EPS of 12-13%
Japan (materials, equipment)	Photoresist, etch gases, test, tools	Moderate, picks-and-shovels	BOJ normalizing; 10Y JGB above 2.3%; TOPIX target 3,350-3,400
Malaysia (back-end packaging)	Assembly and test	Lower, late-cycle spillover	Packaging beneficiary with China oversupply risk

Cycle leverage is Darmine's qualitative ranking of sensitivity to a change in AI capex growth rate, not to its level.

3.1 Korea: high beta on a weak economy

Korea is the purest expression of the trade and its greatest risk. The two memory majors dominate index weight, so the market has rallied despite an economy that only limped out of the 2025 political and constitutional crisis into a projected 1.8% recovery. This is a market where the equity beta has almost nothing to do with domestic GDP and almost everything to do with a global capex line item denominated in another currency. When that line decelerates, Korea has no domestic cushion. Bank of Korea easing can support valuations at the margin, but it cannot substitute for the earnings that the memory cycle provides.

3.2 Taiwan: the indispensable node

Taiwan sits at the logic and advanced-packaging chokepoint. TSMC's CoWoS packaging capacity is the physical bottleneck that gates how many accelerators can be built, which gives Taiwan more pricing durability than Korea, because packaging capacity takes longer to replicate than memory yield improvements. Taiwan is therefore our preferred way to retain some cycle exposure: it is the node least likely to be disintermediated within our tactical horizon, even if it is not immune to a capex growth-rate slowdown.

3.3 Japan: the durable picks-and-shovels

Japan's exposure is the least glamorous and, for that reason, the most defensible. Japanese firms dominate the advanced materials (photoresists, specialty gases, polishing slurries) and a slice of the test and metrology equipment that every fab, HBM or logic, consumes regardless of which vendor wins the accelerator war. These are consumables and tools sold into the entire industry, so they are less exposed to the identity of the cycle's winners than to its aggregate level. The macro overlay is genuinely favorable and independent of chips: BOJ normalization toward a 0.75% to 1.0% policy rate, Shunto wage settlements near 5.5%, and Prime Minister Takaichi's expansionary fiscal stance underpin a TOPIX target of 3,350 to 3,400. Japan lets us hold semiconductor exposure while diversifying the macro driver away from a single capex line.

> **Our view:** Within the exposed nodes we prefer Taiwan for cycle durability and Japan for its independent, domestically driven macro tailwind. Korea is the highest-beta expression and the one we would trim first, because it offers the least protection when AI capex growth decelerates.

4. Will AI Capex Sustain Through 2026?

This is the central forecasting question, and the honest answer requires separating two claims that the market persistently conflates. The first claim is that AI capex keeps growing. We agree with this. The second claim is that AI capex keeps growing at the rate it grew in 2025. We do not. The distinction is everything, because equity multiples in the exposed nodes are priced off the second derivative, the change in the growth rate, not the absolute spend.

4.1 The base-effect problem

2025 was the year hyperscaler capex inflected vertically. Aggregate infrastructure spending across the major buyers rose at a rate that reflected a step-change in commitment, off a base that was still moderate in 2024. Arithmetic alone dictates that the same absolute dollar increase in 2026 produces a lower percentage growth rate, because it is measured against a much larger 2025 base. Even in a scenario where every hyperscaler spends more in 2026 than in 2025, the growth rate decelerates. Memory pricing, order momentum and, crucially, sell-side earnings revisions are calibrated to growth rates. A deceleration from a very high growth rate to a merely high one is, in revision terms, a negative catalyst.

4.2 Signs of hyperscaler discipline

There are early, tentative signals that the buyers are becoming more disciplined. Commentary from the largest cloud operators has begun to emphasize return on invested capital, depreciation schedules and the monetization of AI services rather than raw capacity accretion. This is the natural maturation of any capex cycle: the land-grab phase, where spending is a strategic imperative regardless of near-term return, gives way to a phase where boards demand evidence of revenue. That transition does not end the cycle, but it caps its growth rate and introduces air pockets in the order book. We expect the second half of 2026 to feature at least one such air pocket, most likely a quarter of softer memory order momentum that the market initially misreads as the end of the cycle.

Metric	2024	2025	2026E (house)
Hyperscaler AI capex growth (yoy)	Strong, off moderate base	Vertical inflection	Positive but decelerating
HBM demand (bit growth)	Emerging	Sold out, contracted	Still tight, growth rate easing
Conventional DRAM pricing	Bottoming	Recovery on crowding-out	Firm, supply-driven not demand-driven
Asia ex-Japan EPS growth	Recovering	Led by TW/KR	12-13%, still tech-concentrated
Global merchandise trade volume	>2%	>2%	~0.5% (WTO)

Directional house estimates for illustration; exact figures for hyperscaler capex vary by source. The point is the change in growth rate, not the level.

> **Our view:** Our base case is that AI capex grows through 2026 but the growth rate decelerates meaningfully because 2025's base effects cannot repeat. Since the exposed equities are priced off the second derivative,

> decelerating growth is sufficient to compress multiples even without an outright spending decline.

5. The Macro Backdrop: A Region of Divergence

The memory cycle does not operate in a vacuum. It sits atop a regional macro picture defined by divergence. Asia-Pacific still grows roughly 4.4% in 2026 and contributes close to 60% of world GDP growth, but the composition is uneven and the external trade impulse is fading fast: WTO puts global merchandise trade volume growth at only about 0.5% in 2026, down from above 2% in the prior two years. An economy or an equity market that depends on a rising tide of global trade will find 2026 unforgiving. The AI capex flow is, in effect, the one part of the trade complex that has been growing rapidly, which is precisely why its deceleration matters so much.

Central-bank divergence sharpens the picture. The BOJ is tightening, with the policy rate heading toward 0.75% to 1.0% by mid-2026 and the 10Y JGB yield above 2.3%, the highest since 1998. The PBoC is easing into persistent Chinese deflation and a fading export engine, with a 2025 trade surplus near 1.2 trillion USD now colliding with a 47.5% US effective tariff. The Fed is cutting. MAS is on a gradual S-NEER appreciation path. This divergence means there is no single regional rates trade; each market must be underwritten on its own policy trajectory. For the chip complex specifically, the relevant point is that neither BOK easing nor PBoC easing can offset a memory earnings deceleration, and BOJ tightening is a headwind to the very Japanese exporters that benefit from the cycle, partially hedged by their domestic reflation.

5.1 China as the shadow over memory

China deserves a specific caution. Chinese memory makers, subsidized and strategically motivated, are ramping conventional DRAM and NAND capacity. They are years behind on HBM, but in commodity memory they represent exactly the kind of supply that can erode the crowding-out tightness we identified in Section 2. Chinese oversupply in legacy nodes is also the specific risk cited in the house view for Malaysia's packaging beneficiaries. We do not expect Chinese HBM to matter within our tactical horizon, but Chinese commodity-memory supply is a genuine medium-term threat to the second engine of this cycle, and it argues against extrapolating conventional DRAM strength indefinitely.

> **Our view:** The macro backdrop offers no rescue for a decelerating memory cycle. Easing in Korea and China cannot substitute for capex-driven earnings, and Chinese commodity-memory supply is a structural headwind to conventional DRAM tightness even as HBM stays protected near term.

6. The Rotation: From Machines to Domestic Demand

If the exposed nodes are priced for continued acceleration and the growth rate is set to decelerate, the logical portfolio response is to reduce pure-play AI hardware exposure into strength during the second half of 2026 and rotate the proceeds toward earnings streams that are less hostage to a single capex line. The house baseline already points to where that capital should go: markets driven by domestic demand and services, where the growth engine is internal consumption rather than an external order book.

6.1 India, Indonesia, Singapore

India is the highest structural conviction in the house framework: roughly 6.5% growth, the fastest major economy, domestic-demand driven, a median age of 28, an S&P upgrade to BBB from BBB-, and an India-EU trade deal that diversifies its external base. Its low correlation to the chip cycle is a feature, not a bug, in a second half where that cycle decelerates. Indonesia offers a 270 million consumer base, nickel downstreaming and EV-battery investment, and strong FDI, an accumulate-rated domestic story. Singapore is the quality-defense anchor: 2025 grew 5.0% with Q4 manufacturing up 18.8%, MTI's 2026 forecast upgraded to 2.0% to 4.0%, MAS on gradual appreciation, and a construction pipeline (Changi T5, Tuas Port, North-South Corridor) that provides visible domestic demand independent of the global tech cycle.

6.2 Services and the AI monetization pivot

There is a deeper logic to the rotation. As the AI cycle matures from building infrastructure to monetizing it, value migrates from the hardware makers to the enterprises and consumers who deploy AI. That favors services, financials and consumption plays over the memory and logic suppliers. Mastercard Economics Institute data on resilient APAC consumption, and the domestic-demand orientation of India, Indonesia and the Philippines (roughly 5.6% growth on remittances and consumption, with low chip-cycle correlation), all point the same way. The second half of 2026 is when the market should begin paying for AI adoption rather than AI construction.

> **Our view:** We recommend rotating from pure-play AI hardware toward domestic-demand and services beneficiaries in the second half of 2026. India, Indonesia and Singapore are the preferred destinations because their earnings do not depend on hyperscaler capex growth rates, and the maturing cycle itself shifts value from AI construction to AI adoption.

7. Risks to Our View

The primary risk to a cautious stance is that we are early. Capex cycles driven by strategic imperative rather than near-term return can run longer and hotter than any valuation framework suggests, and a genuine step-change in AI capability (a new model generation that visibly improves unit economics for the buyers) could re-accelerate spending and extend the supercycle into 2027. In that scenario, trimming Korea and the HBM complex in the second half would forfeit meaningful upside.

A second risk runs the other way and would validate our caution violently: a sharp, disorderly capex cut. If one or more hyperscalers were to publicly reset spending lower, the narrow, momentum-driven leadership of Korea and Taiwan could unwind faster than an orderly rotation allows. Our recommendation to trim into strength is partly designed to avoid being a forced seller in that scenario. A third risk is a currency and rates shock from BOJ normalization: a faster-than-expected move toward 1.0% with the 10Y JGB pressing higher could disrupt the yen-funded flows that support parts of the regional tech complex. We size the semiconductor exposure with these tail risks in mind rather than betting on a single path.

> **Our view:** We are deliberately trimming into strength rather than waiting for confirmation, because the asymmetry favors it: the downside from an early trim is forfeited upside, while the downside from being late is a disorderly unwind of narrow leadership. Managing the exit is the trade.

Portfolio Positioning Summary

Overweight: India (highest structural conviction, domestic-demand driven, low chip-cycle correlation), Indonesia (accumulate; 270m consumers, nickel and EV batteries, strong FDI), and Singapore (quality defense; MTI 2026 upgraded to 2.0 to 4.0%, MAS appreciation, visible construction pipeline). Within the chip complex, retain a preference for Taiwan for its advanced-packaging chokepoint durability.

Market weight: Japan (selective; BOJ normalization and Takaichi fiscal support a TOPIX target of 3,350 to 3,400, and its materials/equipment exposure is the most durable picks-and-shovels play on the cycle), Vietnam (monitor transshipment enforcement risk), Malaysia (packaging beneficiary but exposed to China commodity-memory oversupply), and Taiwan on a total-market basis given its indispensability offset by capex-growth sensitivity.

Underweight: South Korea (tactical; highest beta to a decelerating AI capex growth rate with no domestic cushion, 2026 GDP only ~1.8%), China domestic demand (deflation, fading exports, 47.5% US tariff, property overhang), Thailand (regional laggard, high household debt), and pure-play AI hardware exposure broadly, which we recommend trimming into strength through the second half of 2026.

Key catalysts to reassess: Reassess on: (1) hyperscaler capex guidance revisions at H2 2026 results, watching the growth rate not the level; (2) HBM contract pricing and any second-source qualification news; (3) evidence of conventional DRAM softening from Chinese commodity-memory supply; (4) BOJ moves toward 1.0% and the 10Y JGB above 2.3% disrupting yen-funded flows; (5) a next-generation model release that visibly improves AI unit economics and could re-accelerate the cycle; (6) Korean domestic political and policy normalization feeding through to non-tech earnings.

Conclusion

The HBM supercycle is real, it is still growing, and it has been the single most important driver of Asian equity returns for two years. None of that is in dispute. What is changing is the second derivative. AI capital expenditure will very likely be higher in 2026 than in 2025, but it will grow more slowly, because 2025's vertical inflection created a base that arithmetic will not let 2026 repeat. Since the exposed nodes in Korea, Taiwan and Japan are priced off the change in the growth rate rather than the level of spending, a deceleration is enough to compress multiples even without an outright spending decline. That is the core of our tactical caution.

Our recommendation is therefore one of sequencing and discipline rather than wholesale exit. We favor Taiwan for cycle durability and Japan for its independent domestic reflation among the exposed nodes, we would trim high-beta Korea and pure-play AI hardware into strength through the second half of 2026, and we would rotate the proceeds toward the domestic-demand and services beneficiaries of India, Indonesia and Singapore, where earnings do not hang on a single hyperscaler capex line. As the AI cycle matures from building infrastructure to monetizing it, the market should begin paying for adoption rather than construction, and portfolios positioned for that transition will be better placed than those still crowded into the machines.

Sources and Further Reading

International Monetary Fund, Regional Economic Outlook: Asia and Pacific, April 2026

World Bank, East Asia and Pacific Economic Update, April 2026

J.P. Morgan Asia Pacific Equity Research, Semiconductors and AI Capex Outlook, May 2026

World Trade Organization, Global Trade Outlook and Statistics, April 2026

Bank of Korea, Monetary Policy Report, April 2026

Bank of Japan, Outlook for Economic Activity and Prices, April 2026

Monetary Authority of Singapore, Macroeconomic Review, April 2026

Singapore Ministry of Trade and Industry, Economic Survey of Singapore 2025 and 2026 Forecast Update, February 2026

Mastercard Economics Institute, Economic Outlook 2026: Asia Pacific Consumption Trends, December 2025

Nikkei Asia, Coverage of Japanese semiconductor materials and equipment suppliers, first half 2026

Reuters and Bloomberg, HBM contract pricing and hyperscaler capital expenditure reporting, first half 2026