SECTOR UPDATE

2026. 6. 12

Scan to go to
Research Center report database

Tech Team

Jongwook Lee
Team Leader
jwstar.lee@samsung.com

Kyoungbeen Kim
Research Associate
kyoungbeen.kim@samsung.com

Tech (OVERWEIGHT)

2H26 memory outlook: A duration game

- If companies can sustain elevated profitability for an extended period relative to historical norms, valuation multiples during this cycle should reflect this structural shift.
- Demand remains robust and increasingly durable, while supply-side production capacity has contracted materially. The proliferation of long-term agreements (LTAs) has further strengthened investor confidence in a prolonged industry upturn.
- We raise our target prices to KRW500,000 for Samsung Electronics and KRW3,500,000 for SK Hynix.

AT A GLANCE

Samsung Electronics (005930 KS, 299,000)

Target price	KRW500,000	67.2%

BUY

SK Hynix (000660 KS, 2,101,000)

Target price	KRW3,500,000	67.2%

BUY

WHAT'S THE STORY?

Memory rerating: Memory remains a cyclical industry. However, the current DRAM cycle is anything but typical. If fabs can generate higher ROIC for longer at the same book value, the baseline for P/B multiples must adjust accordingly. This cycle is longer, deeper, and deserving of higher valuation multiples than previous cycles. That is the essence of the current rerating. Something bigger is coming.

Increasing duration of demand: In AI servers, memory accounts for over 40% of total costs and has become the primary performance bottleneck. With the advent of the AI era, memory demand is shifting from quarterly inventory cycles to multi-year capacity planning. This extended demand duration renders traditional valuation frameworks—applying low multiples to peak earnings—obsolete. The IPOs of OpenAI and Anthropic could mark a critical inflection point: if public markets prove willing to fund substantial losses and capex, the AI infrastructure cycle could become even more durable.

Supply cannot meet demand: The most important shift is not in demand—it is in supply. Supply growth has slowed due to capex inflation, declining process productivity, high-bandwidth memory (HBM) wafer penalties, and investment delays dating back to 2023. Even at current spending levels, supply growth has stalled, increasing the value of existing fabs. The defining question of this cycle is when supply can overcome these constraints and exceed the long-term demand baseline of 22% bit growth. We believe that turning point will arrive in 2028 or 2029.

Scan to go to
Research Center report database

LTAs validate long-term demand: Long-term supply agreements (LTAs) are on the verge of widespread adoption. They improve visibility on pricing and volumes, helping memory suppliers protect profitability during periods of supply tightness, while terms are increasingly tilting in suppliers' favor. The significance of LTAs lies not in the agreements themselves, but in how they have convinced investors of the durability of memory demand. The market no longer views memory profits as a fleeting, one-year earnings peak, but as a sustained earnings stream extending through 2028.

Raising target prices; SEC remains our top pick: We raise our target prices to KRW500,000 for Samsung Electronics and KRW3,500,000 for SK Hynix. We believe the sustainability of DRAM earnings has now been extended through 2028. While the market has begun to reflect the magnitude of earnings, it has yet to fully price in their duration and the scarcity value of fabs. We expect SEC shares to benefit from the firm's sustained progress in HBM, leverage to commodity DRAM pricing, and exposure to tailwinds from foundry capacity and AI device demand.

WHY SHOULD
YOU READ THIS REPORT?

Dear Investor,

This DRAM cycle is not a typical industry upturn. The proliferation of AI has extended demand durability, while supply remains constrained by capex inflation and slowing productivity gains. As memory manufacturers are generating higher ROIC for longer from the same fabs, we expect the industry's valuation baseline to rise structurally.

In the AI era, memory demand is shifting from quarterly inventory cycles to multi-year infrastructure investment cycles. By contrast, supply faces structural headwinds—including the transition to high-bandwidth memory (HBM), delayed capex, and declining productivity—that make rapid supply growth far more difficult than in the past. As a result, the scarcity value and pricing power of existing fabs are increasing, directly translating into improved ROIC for memory manufacturers.

We view 2028-2029 as the medium- to long-term inflection point, when bit supply growth is expected to finally exceed structural demand growth. Until then, improving fab economics should remain the primary driver of valuation upside.

In this cycle, asset efficiency improvement and sustainable profitability have become more important than mere price increases. While the market has traditionally valued memory stocks based on short-term earnings, we expect it to gradually revise its valuation framework as demand and supply structures evolve.

The growing adoption of long-term agreements signals rising market confidence in sustained long-term demand. However, a genuine rerating should stem not from the contracts themselves, but from memory suppliers' successful integration into customers' system architectures.

Future competitiveness will be determined not by production capacity alone, but by the ability to design customer-specific memory hierarchies, control data flows, and build platform lock-in. The memory industry is evolving from a storage component supplier into a systems solutions provider—a shift that could fundamentally transform its long-term valuation framework.

We maintain OVERWEIGHT on the memory sector and raise our target prices to KRW500,000 for Samsung Electronics (SEC) and KRW3,500,000 for SK Hynix. Our investment thesis rests on two pillars: first, DRAM earnings durability is likely to extend through 2028; and second, while the market has recognized earnings scale, it has yet to fully price in profit sustainability or the scarcity value of existing fabs.

Tech Team **Jongwook Lee Team Leader** | jwstar.lee@samsung.com
Kyoungbeen Kim Research Associate | kyoungbeen.kim@samsung.com

WHY SHOULD
YOU READ THIS REPORT?

We continue to favor SEC as our top pick, as it stands to benefit simultaneously from rising commodity DRAM prices, expanding HBM market share, fab valuation rerating, and narrowing valuation gaps with peers. This cycle is not merely an earnings cycle—it is also a book-value revaluation cycle.

We hope this report enhances investors' understanding of the industry and supports informed investment decisions.

Thank you,

Jongwook Lee and Kyoungbeen Kim

Tech Team **Jongwook Lee Team Leader** | jwstar.lee@samsung.com
Kyoungbeen Kim Research Associate | kyoungbeen.kim@samsung.com

REPORT
CONTENTS

Two inflection points and rerating

This DRAM cycle is unique: demand has been extended and supply growth delayed. Over the long term, memory companies' book values will also need to be revalued.

This cycle features two critical inflection points: 1) a slowdown in DRAM price increases in 2Q26, which should represent a temporary correction—not the end of the cycle; and 2) a sharp ramp-up in supply from new fabs in 2Q27, which should mark the beginning of a sustained decline in share-price momentum. Yet this is not merely a price cycle. A structural rerating is also underway, driven by extended earnings duration, improved fab economics, and memory companies' potential evolution into solution providers. While we acknowledge these inflection points, we believe the magnitude and duration of share-price movements will differ significantly from past cycles.

Two inflection points

We expect the DRAM cycle to pass through two pivotal quarters, each shaping the trajectory of stock performance. The first inflection point should arrive in 2Q26, when DRAM price growth begins to decelerate. We view this as a temporary correction phase and maintain our OVERWEIGHT recommendation. The second should occur in 2Q27, when new fabs come online and supply growth accelerates. From this point onward, the market will face a true test of durability.

Historically, the pattern has been consistent: once DRAM price growth peaks, stock prices lead the decline; when new fab capacity becomes visible, stocks enter a sideways phase. The ultimate cycle peak has always been driven not by supply, but by demand shifts—such as Amazon's order cuts in 2018 or the global PC demand contraction in 2022.

This cycle should follow the same structural logic, but with two key differences: 1) both demand and supply are now less price-elastic; and 2) the interval between inflection points is wider and deeper. Earnings duration has lengthened, ROIC at existing fabs has improved, and memory companies are increasingly positioned to become integrated solution providers. As a result, we expect the rally between the two inflection points to be steeper and longer than in prior cycles. Moreover, this cycle includes a rerating component—unlike past cycles.

2025-2028 cycle analysis

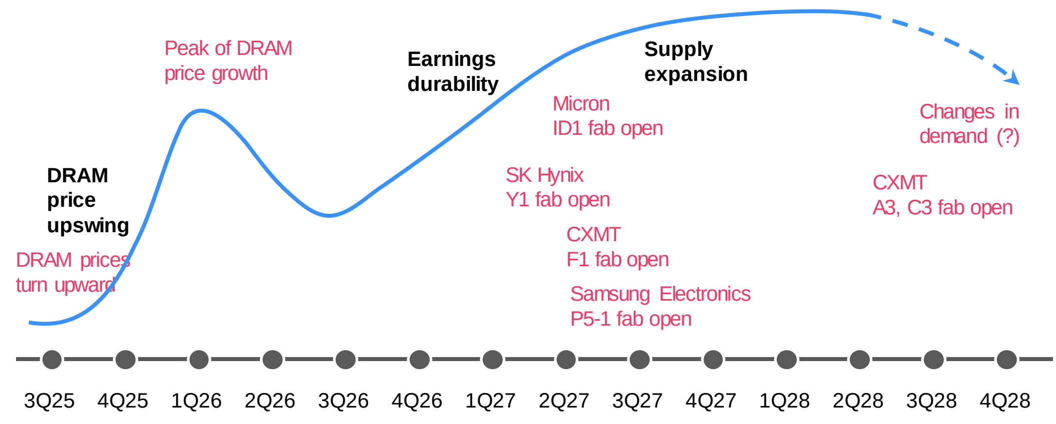

Source: Samsung Securities estimates

2016-2018 cycle analysis

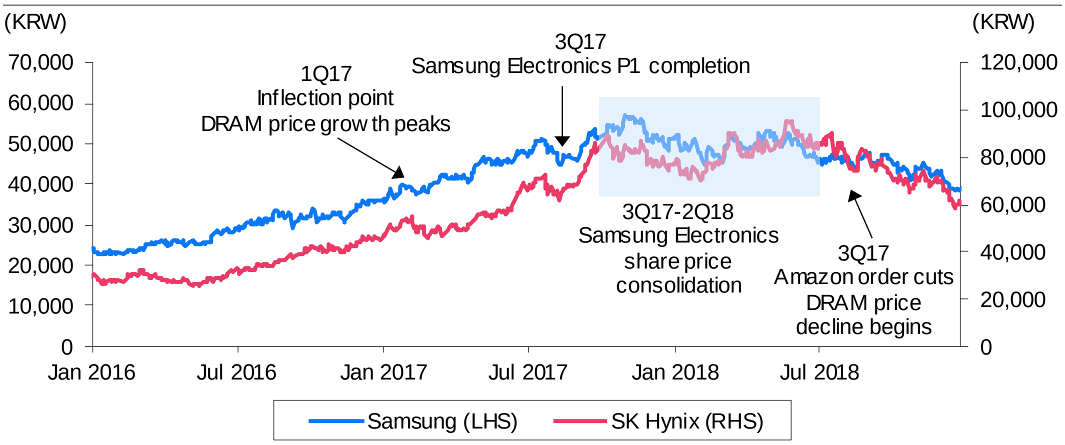

Source: QuantiWise, Samsung Securities

2020-2022 cycle analysis

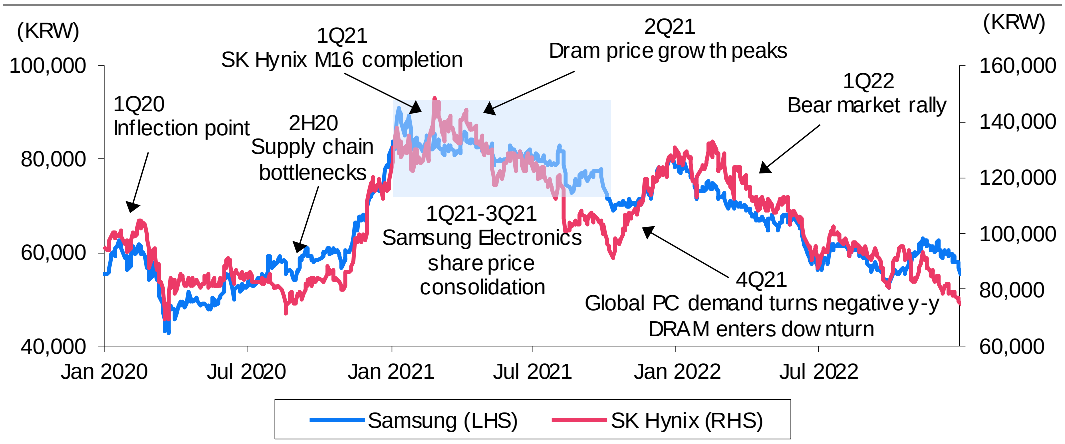

Source: QuantiWise, Samsung Securities

Why rerating is beginning in this cycle

The market is not yet confident enough to call the cycle peak. Instead, two sentiments coexist: skepticism that conditions are “too good to be true” and a growing appetite for a new narrative. While stock prices have surged, investor sentiment has not yet become overwhelmingly one-sided, in our view. We occasionally encounter investors who argue strongly for growth extending to 2030 and long-term agreement (LTA) tailwinds—but they remain a minority. We believe the market will peak only when bullish conviction becomes dominant. For now, the narrative is not yet being dictated by strong bulls.

We do not view the current period as the late stage of an AI infrastructure cycle. Rather, we see it as the midpoint of a valuation rerating process. As evidence accumulates that earnings are structural rather than merely cyclical—and as new narratives around AI services continue to emerge—investors should increasingly conclude that the valuation framework for AI hardware must be recalibrated. DRAM remains a cyclical industry, but the structural shift— memory’s rising share of system-level value creation in the AI era—is profoundly important. This shift is likely to endure far longer than most investors expect.

The three stages of DRAM rerating

We believe the rerating of memory stocks will unfold in three distinct stages.

Stage 1. Extended earnings duration: In past cycles, investors immediately questioned when DRAM prices would peak and reverse. In this cycle, however, the dominant question has changed: Could this last longer than we thought? This subtle but profound shift in mindset alone can drive valuations higher. Even at the same peak earnings, the present value of profits differs dramatically if they last three years rather than one. That said, it would be inaccurate to claim that memory has escaped its cyclical nature. The more precise framing is that while AI infrastructure investment has not eliminated the DRAM cycle, it has structurally extended its duration.

Stage 2. Elevated fab economics: There is growing recognition that existing fabs and cleanrooms are generating higher ROIC for longer than in prior cycles. When both the magnitude and duration of earnings improve, the capital that produces those earnings must be revalued accordingly. This is the fundamental basis for P/B rerating.

Stage 3. Memory as a solution asset: The final stage involves the market recognizing memory not as a generic commodity component, but as an embedded, system-level asset within customer architectures. Competitive advantage shifts from cost leadership to quality. When the consensus fully embraces this evolution, the entire valuation framework should transform.

We believe the cycle currently sits between stage 1 and stage 2. While earnings durability has begun to be priced in, the market has not fully appreciated the structural uplift in fab ROIC. As investor understanding deepens that profitability has risen structurally, the rerating should continue to unfold.

Our view on P/E valuation

No investor is unaware that memory stocks trade at low P/E multiples. But for P/E expansion to be justified, the underlying assumption must be that today's high profit margins are sustainable. In this sense, the debate over P/E is not fundamentally different from the P/B-ROE framework—it ultimately hinges on confidence in earnings durability. The discomfort with current valuations stems less from fear of margin compression than from underlying unease over whether the earnings base itself is truly structural.

P/E and P/B-ROE are simply two ways of describing the same valuation. Arguing for P/B premium expansion is, in effect, the same as arguing for P/E rerating. However, given that DRAM earnings volatility remains higher than that of other semiconductor peers—and that even the Philadelphia Semiconductor Index typically trades at a 15-20% P/E discount during late-cycle phases—we believe P/E is best used as an intuitive, cycle-positioning indicator rather than as the primary valuation tool today.

The three-stage valuation rerating framework we have defined represents a transition away from P/B as the dominant metric and toward P/E as the primary anchor. For this shift to occur, two conditions must be met: 1) DRAM earnings volatility must converge toward, or fall below, that of other semiconductor companies; and 2) the historical discount applied to memory P/E relative to the broader semiconductor sector must become a source of upside. Identifying evidence of these conditions should be central to building long-term investment conviction in memory.

Lessons from TSMC's three P/E ratings and implications for memory

TSMC has undergone three distinct P/E reratings over the past two decades: 1) 2010-2012: the mobile transition; 2) 2016-2018: the datacenter cycle; and 3) 2024-present: the AI and CoWoS era. Each rerating was driven by a different catalyst.

- The mobile rerating (2010-2012) was driven by total addressable market (TAM) expansion rather than supply scarcity. Demand growth fueled the phase through a broader customer base and higher shipment volumes. Samsung Electronics' (SEC) foundry business also expanded during this period, meaning TSMC's premium was driven not by supply dominance but by its ability to capture a growing market.
- The datacenter rerating (2016-2018) drove the largest P/E expansion in TSMC's history as three factors converged. First, hyperscaler capex emerged as a new structural driver of semiconductor demand—not simply more chips, but different chips characterized by higher performance and longer life cycles. Second, Intel's delays below 10 nm and SEC's yield challenges created TSMC's first clear and durable technology lead. The adoption of InFO further strengthened its leadership in mobile. During this period, TSMC's irreplaceability began to be reflected in its stock price for the first time. Third, investor sentiment remained cautious following the post-2013 mobile slowdown, making the subsequent rerating even more powerful.

- The AI/CoWoS rerating (2024-2026) is defined by confirmation of a technology monopoly, driven by advanced-process leadership and deep customer lock-in. Yet the rerating remains underway. Several headwinds continue to cap P/E upside: 1) CoWoS capacity expansion is ongoing, and supply is not yet fully constrained; 2) Intel and SEC's foundry businesses remain credible competitive threats; 3) new entrants such as Terafab could disrupt the ecosystem; 4) no demand driver has yet emerged at AI's scale to fully absorb TSMC's total capacity; and 5) geopolitical risks have intensified.
- The 2016-2018 datacenter cycle was a transformative period for the semiconductor industry. Yet an important divergence emerged: while TSMC enjoyed a historic P/E rerating as markets recognized a structural shift in its business model, memory companies continued to be viewed as cyclical players—even as their earnings surged (although their shares also rallied on strong earnings growth). The distinction lay in the nature of demand.
- For TSMC, datacenter growth expanded demand across high-performance computing, GPUs, networking, and custom ASICs, each requiring larger die sizes, higher ASPs, and longer product roadmaps. By contrast, memory demand from datacenters, while robust, remained structurally unchanged: standardized, commoditized DRAM and NAND sold to a new customer base, yet still functioning as interchangeable components. Memory capex was viewed not as a source of structural advantage, as in TSMC's case, but as a cyclical expansion that would eventually lead to oversupply, price erosion, and ROE mean reversion—which is precisely what occurred.
- Today, memory is beginning to resemble TSMC in 2016. First, the demand structure is changing for the first time—demand is not merely increasing in volume but evolving in nature. Second, competition on the supply side is shifting from pure cost leadership toward technology-driven differentiation. Third, valuation multiples remain depressed.
- What TSMC experienced during the datacenter cycle was not simply a cyclical upswing, but a fundamental revaluation of the economic value of its fabs. We believe memory is now undergoing a similar process. Admittedly, today's memory products still contain a much higher proportion of commoditized, general-purpose components than TSMC's highly customized offerings. However, if memory solutionization—where memory becomes embedded, system-level infrastructure rather than a standalone chip—gains traction, the structural gap between memory and TSMC's business model should narrow.

12-month fwd P/E: TSMC and Samsung Electronics

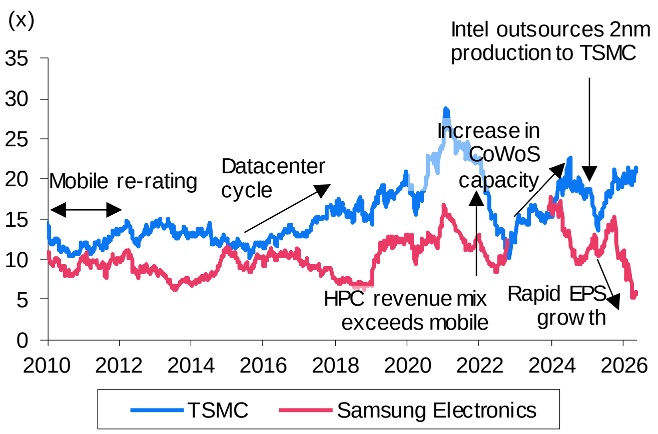

Source: Bloomberg, Samsung Securities

12-month fwd P/E: SOX Index, Micron, and TSMC

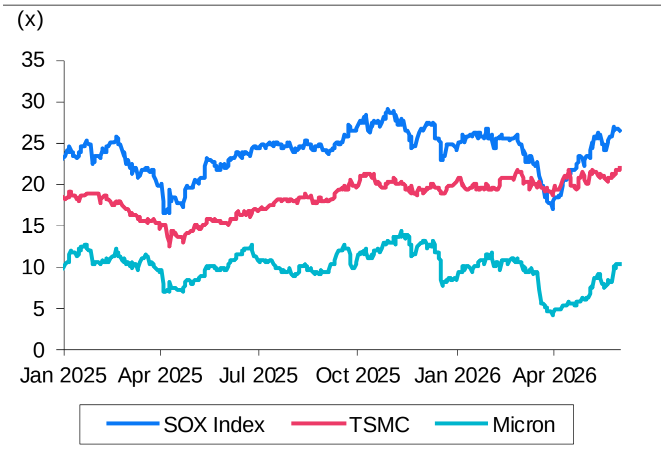

Source: Bloomberg, Samsung Securities

Increasing duration of demand cycles

AI datacenter investment is not a product cycle but an infrastructure build-out. As a result, memory demand is shifting from quarterly inventory cycles to multi-year capacity planning.

Characteristics of AI infrastructure demand

AI datacenter investment does not make memory a non-cyclical industry. Memory remains cyclical. What has changed in this cycle is the nature of the demand driving it.

In the past, memory demand cycles closely mirrored product cycles. Demand—driven by PC replacements, smartphone launches, server inventory replenishment, or the pandemic era— rose sharply but often faded quickly. These fluctuations were heavily influenced by end-product demand, channel inventory levels, and consumer sentiment. When demand was strong, memory prices surged; when sales slowed or inventories built, the cycle reversed rapidly. This dynamic has long been the core reason the memory industry has traded at low valuations.

By contrast, AI infrastructure is less a product trend than a large-scale industrial build-out. It involves a long-term investment cycle that bundles GPU clusters, power grids, networking, datacenters, cooling, storage, and memory as interdependent components. Demand no longer ebbs and flows on a quarterly basis; instead, it is driven by multi-year capacity planning. As competitive advantage in AI services increasingly hinges not just on model performance but on guaranteed access to sufficient computing power, customers are prioritizing first-mover advantage over short-term profitability. In this context, memory is no longer viewed as just a component—it has become a key bottleneck in AI infrastructure.

Moreover, the upcoming IPOs of frontier AI companies such as OpenAI and Anthropic could further extend the duration of this demand cycle. Until now, their infrastructure investments have relied primarily on private funding and strategic partnerships. After IPOs, however, these firms would gain access to significantly larger pools of capital through public equity markets. In such a scenario, AI capex could shift from a simple VC-funded cycle into a long-term investment cycle repeatedly validated and scaled through public equity markets. For the memory industry, the critical point is that AI demand is no longer about one-off server expansion—it is evolving into a long-term infrastructure investment cycle intertwined with capital markets.

Hyperscalers: Capex (Oct 2025)

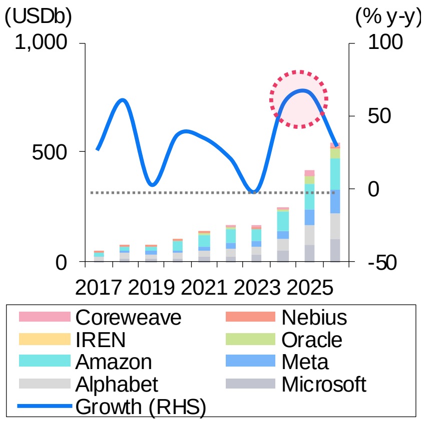

Source: Bloomberg, Samsung Securities

Hyperscalers: Capex (Jan 2026)

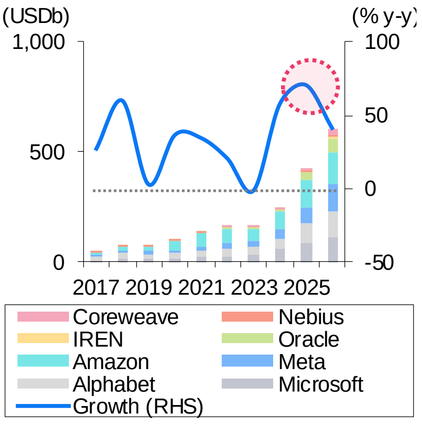

Source: Bloomberg, Samsung Securities

Hyperscalers: Capex (Jun 2026)

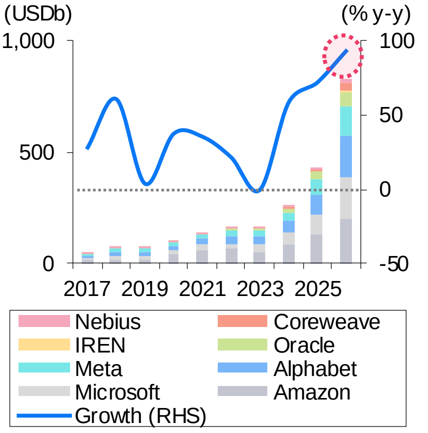

Source: Bloomberg, Samsung Securities

Hyperscalers: Capex plan revisions

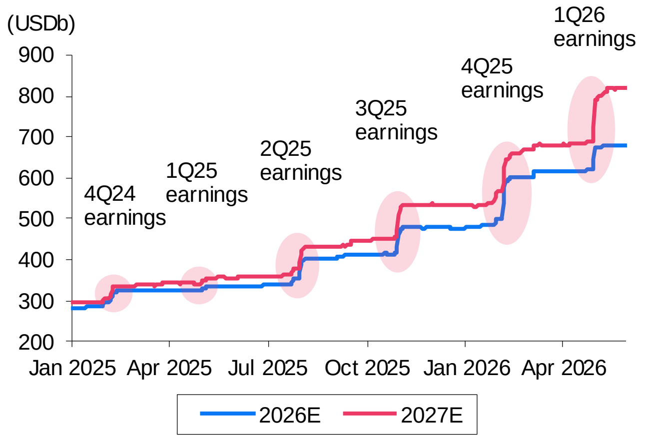

Note: Sum of Amazon, Meta, Microsoft and Google

Source: Bloomberg, Samsung Securities

Hyperscalers: Capex-to-EBITDA ratio

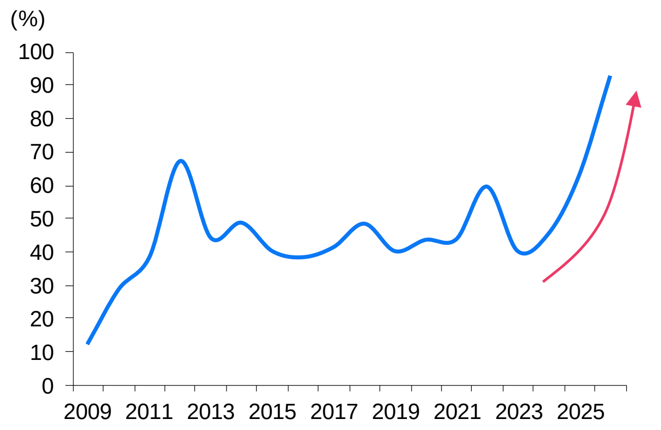

Source: Bloomberg, Samsung Securities

Reaffirming robust AI demand

Ongoing supply expansion efforts are unfolding against a backdrop of solid AI demand, and the AI catalysts expected to emerge after the stock-price correction should provide even more powerful rebound momentum. Past semiconductor cycle downturns were primarily driven not by supply gluts, but by sharp contractions in demand. The following three drivers provide strong foundational support for an additional memory industry upturn:

First, securing computing power remains a strategic priority for customers. AI competition among Big Tech is centered on securing computing assets, not launching services, as computing power is becoming the key barrier to AI service quality.

Recent projections from OpenAI illustrate this clearly: although its sales forecast was raised to USD30b, its projected free cash flow has deteriorated even more sharply. Cash burn for 2026 and 2027 is estimated to reach USD25b and USD57b, respectively. This indicates that investment in computing power is growing faster than revenue. Yet in Mar 2026, OpenAI successfully raised funding at a post-money valuation of approximately USD852b. This signals that both equity and venture capital markets currently place greater value on preemptive market positioning through AI semiconductor capex than on near-term profitability. As long as companies keep racing for preemptive dominance, hardware demand should have a solid floor.

Second, enterprise demand and the adoption of agentic AI are accelerating rapidly. The center of gravity in AI demand is shifting from consumer to enterprise use cases, with agentic AI serving as the key catalyst.

Gartner projects that by the end of 2026, AI agents will be embedded in 40% of enterprise applications, up from less than 5% in 2025—an almost eightfold increase. These agents perform tasks by running dozens of reasoning loops in response to user prompts. As these reasoning loops become more sophisticated, agent performance improves. At the same time, as users manage an increasing number of AI agents, memory demand associated with those workloads is growing exponentially.

Third, explosive growth in inference token consumption shows no signs of slowing—demonstrating that LLM providers remain firmly on the steep upward slope of their adoption curve.

At GTC 2026, Jensen Huang suggested that Nvidia could eventually operate 7.5m AI agents alongside 75,000 employees. This implies a future in which each employee leverages roughly 100 agents to enhance productivity. Meanwhile, Google disclosed during its 1Q26 earnings call that direct API token processing volume for its first-party model exceeded 16b tokens per minute, up from 10b in the prior quarter. The proliferation of multimodal services—processing text, images, video, and audio simultaneously—is further accelerating token consumption. Rising token volumes translate directly into memory demand, serving as a key driver not only for high-bandwidth memory (HBM), but also for commodity DRAM and NAND.

OpenAI: Sales forecasts

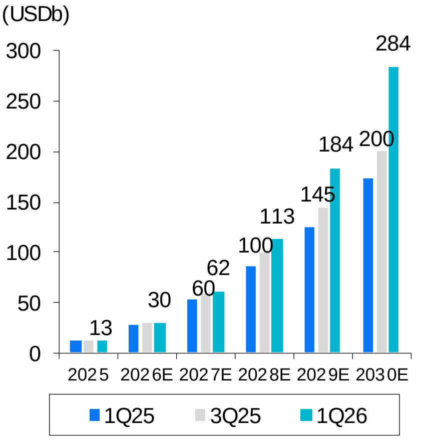

Source: Press, Samsung Securities

OpenAI: FCF forecasts

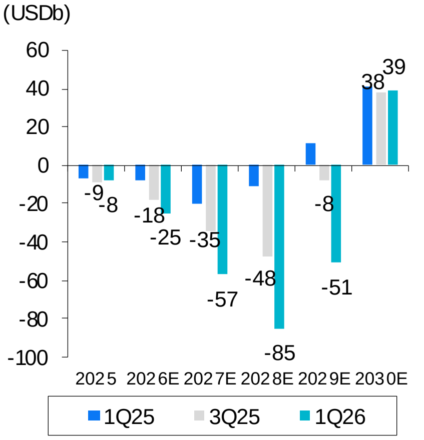

Source: Press, Samsung Securities

Google: Monthly tokens processed

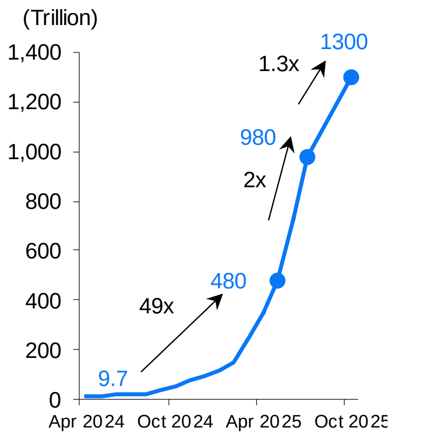

Source: Company data, Samsung Securities

IPOs of frontier AI companies as critical milestones in the AI cycle

The IPOs of frontier AI companies could mark a critical milestone in the AI infrastructure cycle. According to media reports, OpenAI is reviewing potential SEC filings as early as 2H26, with preliminary discussions suggesting it could raise at least USD60b. Anthropic has also been cited as a potential IPO candidate in 2026, with recent developments highlighting both substantial funding needs and improving prospects for profitability. While timelines and deal sizes remain unconfirmed, investor interest is unmistakable. Whether AI model companies can justify massive losses and capex requirements in public markets will be a key test for the next phase of the cycle.

The key metrics for evaluating frontier AI companies post-IPO are likely to differ from those used for traditional software firms. While ARR, ARPU, active users, and retention rates will remain important, token consumption and compute utilization may become more critical leading indicators from an AI infrastructure perspective. This is because the compute spending of these firms is directly linked to hyperscalers' datacenter investments, which in turn drive demand for GPUs, HBM, DRAM, NAND, networking, and power infrastructure. Ultimately, revenue growth at OpenAI and Anthropic should not be viewed merely as software sales growth—it should also be interpreted as a leading indicator of demand for AI semiconductors.

To date, frontier AI companies have relied on one or two large private funding rounds each year to cover losses and fund capex. Post-IPO, however, this dynamic changes. Public companies cannot raise capital as readily whenever runway concerns emerge, as private companies can through venture funding rounds. Instead, they must continually convince investors, justify ongoing losses, and demonstrate that capital raises are supported by growth sufficient to offset shareholder dilution. As a result, post-IPO evaluation will focus not only on user growth, but also on whether frontier AI models can build highly loyal customer bases, increase ARPU through agentic AI, and create durable workflow lock-in around their platforms.

In this regard, frontier AI companies need their own “Tesla moment.” Tesla recorded years of losses, yet investors willingly accepted dilutive equity offerings—not for car sales alone, but for the vision of securing first-mover advantage in EVs, driving down battery costs, building out charging infrastructure, perfecting autonomous driving software, and scaling manufacturing. The same applies to AI model companies. Merely claiming market leadership based on hundreds of millions of users—despite deeply negative FCF—is no longer enough. What is required is compelling evidence of a sophisticated business model: AI agents deeply embedded in enterprise processes, corporate customers building workflows around specific models, and token consumption growing continuously.

Ultimately, the message from OpenAI’s Sam Altman and Anthropic’s Dario Amodei to public markets must be clear: current losses are not expenses but investments in the future compute platform. If investors come to believe that this platform can generate high ARPU and powerful lock-in effects over time, capital will accelerate. If not, cash-flow pressure will intensify and risks will mount.

Tesla's fundraising case

Tesla's trajectory offers a compelling case study in how frontier AI firms might sustain capital raising after going public. For years, Tesla prioritized growth and capex over profitability—yet its rising stock price enabled easier access to capital, which in turn funded larger production scale-up and deeper technology investment. A self-reinforcing loop emerged: higher stock price → easier equity raises → greater capex → faster growth → higher stock price.

This loop was possible because investors believed in Elon Musk's vision. Gigafactories were not merely factories but infrastructure designed to drive down battery costs and secure leadership in EVs. Tesla vehicles on the road were not just units sold, but foundational assets for data collection, brand equity, the charging network, and autonomous driving software.

Similarly, frontier AI companies need to convince markets that massive compute investments are not merely costs, but preemptive bets on model performance, agent ecosystems, enterprise workflow lock-in, and the scaling of token consumption. If markets embrace this narrative, the AI capex cycle can persist longer and grow stronger—even after these companies go public.

Tesla: Operating profit and funding timeline

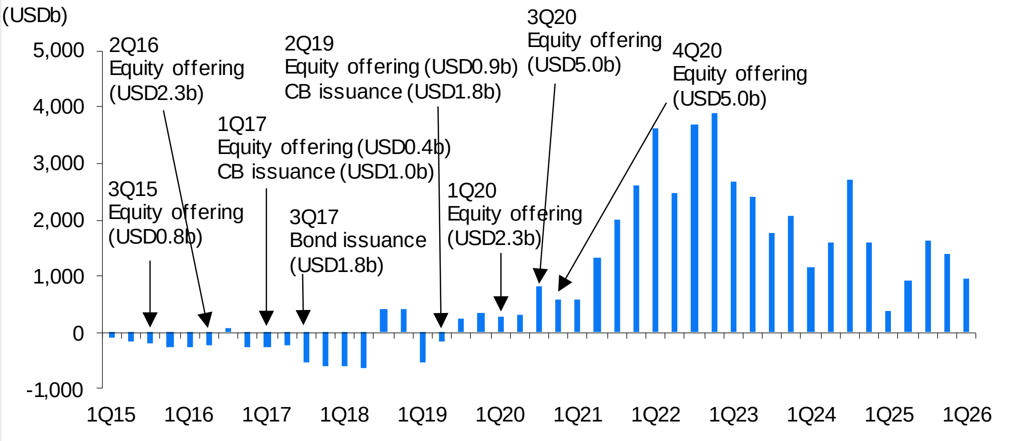

Source: Bloomberg, Samsung Securities

Fab economics upgrade

Essence of the AI cycle: The revaluation of memory fabs

This is not simply another memory cycle prolonged by strong demand. Demand has been extended by AI infrastructure investment, while supply growth has slowed due to capex inflation, diminishing process productivity gains, HBM wafer penalties, and delayed investments in 2023. Together, these two forces have reshaped the economics of existing fabs, giving them greater pricing power and higher ROIC than ever before.

At the core of this rerating is not how much higher DRAM prices can rise, but how long elevated returns can persist. How much longer can the same fabs, with the same book value, generate materially higher returns than in the past? If new supply faces higher execution barriers, long-term demand growth remains near the 22% convergence rate, and HBM continues to cannibalize commodity DRAM capacity, then memory companies' normalized ROIC and P/B multiples warrant a structural upgrade.

Memory remains a cyclical industry, but the baseline has shifted. The peak of this cycle is unlikely to be merely a price peak—it is more likely to mark the beginning of a lasting revaluation of fab economics. We believe the true inflection point will arrive around 2028-2029, when bit-supply growth not only overcomes the HBM penalty but also consistently exceeds the 22% demand convergence rate.

The new memory economy: Margins set to level up

The story behind memory's rerating is not simply a longer cycle—it is a structural upgrade in fab economics. We are not arguing that memory is no longer cyclical. Memory remains a cyclical business, with prices moving in response to timing mismatches between supply and demand. What has changed is the shape of the cycle. In past cycles, prices oscillated around a low equilibrium price and a low normalized ROIC. In this cycle, the baseline itself has shifted higher. The cycle has not disappeared; rather, the entire ROIC curve has undergone a parallel upward shift.

This shift is driven by three factors.

First, demand has become harder to predict. In the past, DRAM demand was anchored by PCs, mobile devices, and traditional servers—markets characterized by clear replacement cycles and reliable shipment data. Smartphones refreshed every two to three years and PCs every four to five years. Even when one segment slowed, others absorbed the shock. AI infrastructure demand is fundamentally different. AI investment is driven by Big Tech's capex cycles, evolving model architectures, the pace of inference adoption, and widely varying memory intensity across customers. The optimal mix of HBM, server DRAM, eSSD, CXL, and KV-cache offload remains fluid. While a replacement cycle may eventually emerge, current demand is driven more by capacity expansion than replacement. Many operators prefer to gradually reduce utilization of legacy servers rather than retire them outright. As a result, demand is massive but difficult to forecast. For suppliers, proactive capacity decisions can no longer be based solely on shipment forecasts, as they were in the past.

Second, the cost of new supply has risen significantly. The same KRW1t of investment now generates far less bit growth than it once did. The adoption of EUV, increasing process-node complexity, rising cleanroom construction costs, soaring utility requirements (electricity and water), equipment inflation, longer lead times, and higher HBM post-processing costs have all increased capex intensity per bit. In the past, expanding wafer capacity and advancing process technology translated directly into bit growth. Today, that linkage has weakened. The transition to HBM imposes a wafer penalty, while continued miniaturization faces diminishing efficiency gains. As a result, bit growth per unit of capex has structurally declined.

Third, process migration has become increasingly complex. Historically, the DRAM industry reduced the marginal cost of new supply through node shrinkage and productivity gains. As nodes advance, however, scaling becomes more difficult, equipment investment requirements increase, and yield stabilization timelines lengthen. HBM further compounds these challenges, requiring more wafers and post-processing resources than commodity DRAM does. Larger die sizes, higher I/O pin counts, more-demanding thermal management, and expanded redundancy all increase wafer consumption. As a result, the same cleanroom and wafer capacity now produce fewer commodity DRAM bits than in the past.

Together, these three forces have raised the hurdle rate for new supply. To justify new fabs, memory suppliers now require higher long-term pricing, wider margins, and greater contract visibility than in prior cycles. Without confirmed LTAs, upfront payments, firm customer commitments, and favorable pricing terms, suppliers can no longer afford to invest ahead of demand merely to gain market share. Supplier caution, therefore, should not be viewed simply as oligopolistic discipline—it is a rational response to heightened demand uncertainty and AI-driven capex inflation.

Against this backdrop, the value of existing fabs warrants reassessment. The diminishing returns from capacity expansion and rising barriers to new fab construction have made operating fabs increasingly valuable assets. As returns on new fab investments decline, lead times for new supply lengthen, and demand for higher-capacity memory rises, existing cleanrooms and wafer capacity command renewed scarcity value—and, in turn, greater pricing power. Their book value now represents capital capable of generating higher ROIC than in previous cycles. In essence, the P/B multiple reratings of memory manufacturers reflects the structural improvement in fab ROIC.

Industry: DRAM capex vs bit growth

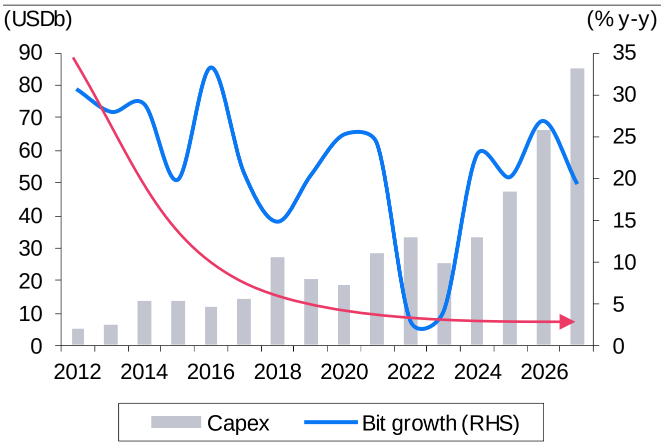

Source: Omdia, Samsung Securities estimates

Industry: DRAM bits per wafer

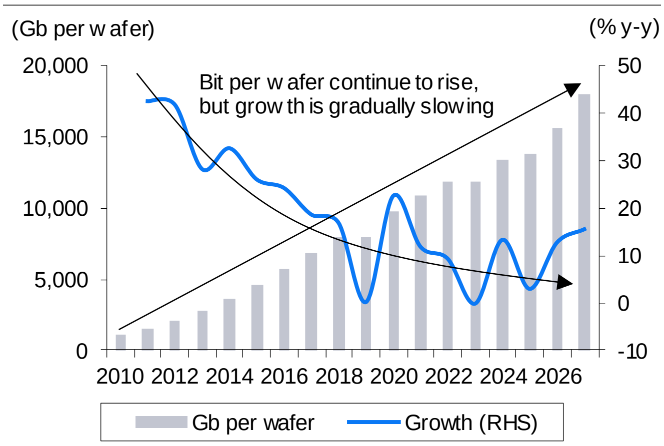

Source: Omdia, Samsung Securities estimates

Butterfly effect of the 2023 investment delay

The investment pullback in 2023 is one of the key factors behind the current memory supply shortage. At the time, the industry was gripped by a sense of crisis: for the first time in its history, memory demand had stopped growing.

The surge in PC demand during the pandemic, along with enthusiasm surrounding the metaverse, had begun to fade. PC shipments had returned to negative growth, server shipments had fallen to single-digit declines, and mobile device demand remained below even pandemic-era levels. As much as 95% of the global DRAM market was serving customers whose demand had stagnated. In previous downturns, companies could still justify building new fabs, confident that demand would eventually recover and absorb the added capacity. In 2023, however, concerns grew that newly built fabs might never be fully utilized. The display industry had already experienced this outcome: excess fabs stood idle for years, and new construction effectively came to a halt.

In 2023, delaying fab construction was a rational move to preserve cash and ensure survival. However, it inadvertently created a butterfly effect that shifted the supply curve lower for 2026-2027. Fabs and cleanrooms originally scheduled to come online over 2025-2026 were pushed back to 2027 or later, leaving the industry critically short of the cleanroom capacity needed to meet today's surging AI-driven demand.

Crucially, once delayed, fab projects are not easily accelerated. DRAM supply cannot be scaled simply by adding more equipment. Even setting aside the long-term planning required for semiconductor clusters—including national-level decisions on electricity and water infrastructure—building a fab within an existing cluster still requires substantial time for cleanroom construction, utility integration, equipment installation, process setup, and yield stabilization. Because regulatory approvals, workforce deployment, and supplier coordination must align simultaneously, schedule slippage is extremely difficult to recover. As a result, the 2023 investment pause has now materialized as a physical capacity shortfall over 2026-2027.

Historically, peaks in DRAM pricing cycles have often coincided with the completion of major new fabs. Once large-scale cleanrooms were completed, the market viewed this as a signal that sufficient supply would soon become available, shifting expectations toward price stabilization or decline. This dynamic was evident with SEC's P1 fab in Pyeongtaek (2017) and SK Hynix's M16 fab (2021). The unusually prolonged industry uptrend of 2017-2018 was largely driven by the timing gap between the full ramp-up of SEC's Hwaseong Line #17 and the eventual completion of the Pyeongtaek P1 fab.

Relative share-price performance: Samsung Electronics vs SK Hynix (2016-2018)

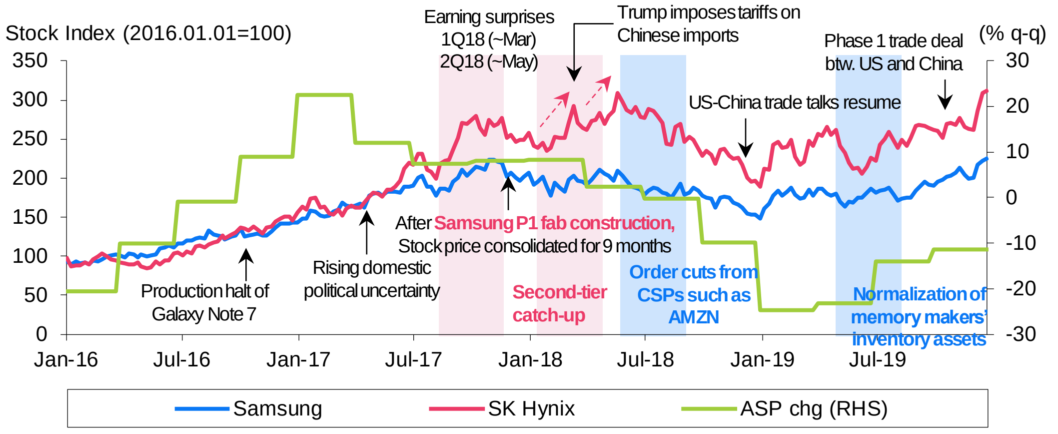

Source: QuantiWise, Samsung Securities

This cycle, however, is different. Even as prices have begun to rise, the arrival of large fabs capable of easing supply constraints has been significantly delayed. We expect Hynix's Yongin Y1 to complete cleanroom preparation by Feb 2027, Micron's ID1 in 1H27, and SEC's P5—optimistically—in 3Q27. With all major memory makers' key new fabs concentrated in 2027 and beyond, meaningful new supply capable of overwhelming market demand is unlikely to emerge before 1H27.

Samsung electronics and SK Hynix: DRAM sales and opening of major fabs

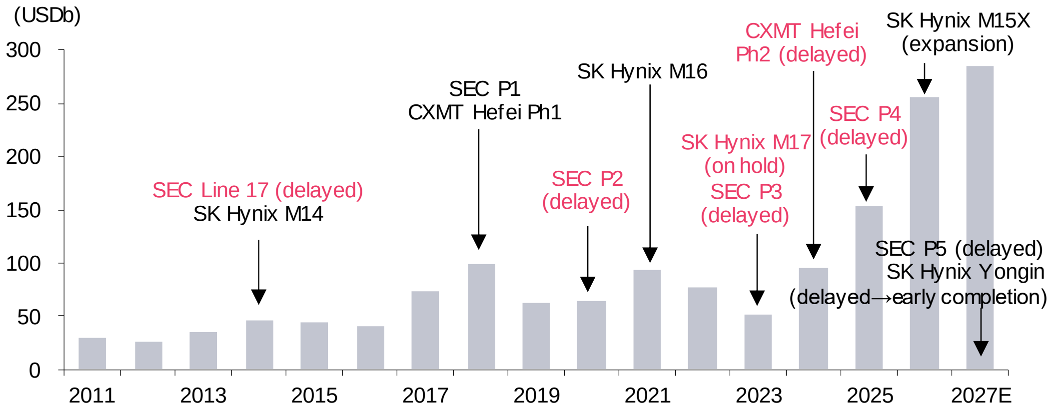

Source: WSTS, Samsung Securities estimates

In conclusion, this DRAM cycle is likely to diverge from historical patterns and sustain its upward momentum far longer than usual. The extended duration of the cycle should support the durability of profits for Korean memory manufacturers. The current shortage is not merely the result of inventory depletion—it reflects a physical capacity gap created by past underinvestment. For customers, persistent supply uncertainty makes inventory reductions difficult. With supply physically constrained, AI infrastructure demand continuing to expand, and HBM increasingly diverting capacity away from commodity DRAM, speculative demand is unlikely to fade quickly. This dynamic underpins the strong earnings visibility of Korean memory makers through 2026-2027.

As long as HBM exists, supply shortage will be difficult to resolve

The root cause of the current DRAM shortage lies in HBM. More than just a high-value-added product, HBM is actively reshaping the entire DRAM supply curve.

Since 2023, HBM demand has consistently outpaced overall DRAM bit demand, and all three memory manufacturers are determined not to miss the HBM opportunity. Securing a foothold in the HBM market is critical to keeping pace with evolving AI architectures, next-generation GPU roadmaps, changing memory specifications, and shifting system designs. As a result, SEC, Hynix, and Micron are all likely to pursue a dual-track strategy: expanding the HBM portion of production while leveraging the pricing power of server DRAM. To simultaneously capture technological leadership, HBM market growth, and server DRAM profitability, the most rational approach is to balance market share across both segments. In this context, Micron has demonstrated early strategic foresight, SEC is likely to expand HBM capacity aggressively despite the associated costs, and Hynix may prioritize profitability in server DRAM and LPDDR, competing more selectively in HBM capacity given its resource constraints. Crucially, none of the three memory makers is willing to walk away from HBM, regardless of how profitable server DRAM may become.

Samsung Electronics: HBM and server DRAM market shares

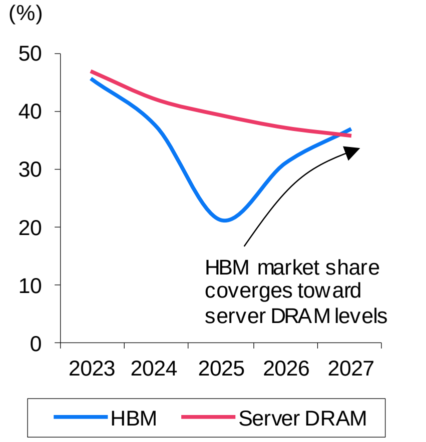

Source: Counterpoint, Samsung Securities estimates

SK Hynix: HBM and server DRAM market shares

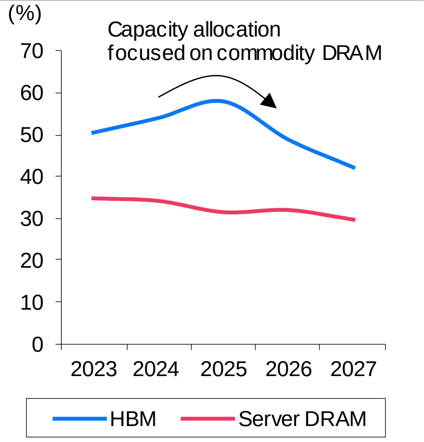

Source: Counterpoint, Samsung Securities estimates

Micron: HBM and server DRAM market shares

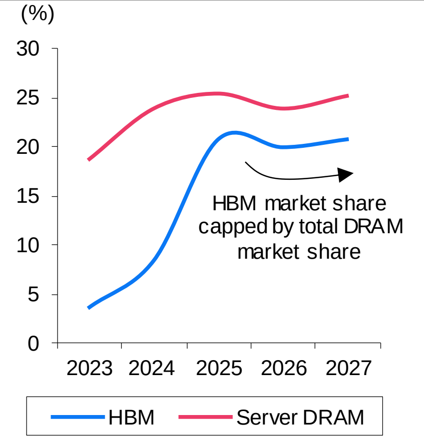

Source: Counterpoint, Samsung Securities estimates

The challenge is that HBM delivers far fewer effective bits per wafer. HBM dies are larger than standard server DRAM dies, while TSV processing, stacking, bonding, testing, and packaging introduce additional yield losses and cost burdens. HBM3 and HBM3E already carry a substantial wafer penalty relative to conventional server DRAM. With HBM4, this burden is expected to increase further due to larger die sizes, higher I/O pin counts, more-complex thermal management, and expanded redundancy schemes. Nor is this issue unique to HBM4—it reflects a structural trend likely to intensify with each successive HBM generation. As HBM demand rises, fewer wafers remain available for commodity DRAM. In other words, while HBM enhances memory makers' revenue mix and profitability, it simultaneously constrains industry-wide bit supply. In the past, a 10% increase in wafer capacity could translate into 30-40% bit growth. Today, diminishing productivity gains from miniaturization have pushed that figure below 20%. With the shift toward HBM, it is now even lower—and in some cases, negative. AI has extended demand; HBM has constrained supply. That dynamic defines the current cycle.

Accordingly, the recent increase in capex should not be immediately interpreted as a precursor to oversupply. Current investments are not intended to drive a rapid increase in supply in 2026, but rather to secure the infrastructure needed to sustain bit growth from 2027 onward. In particular, given the limited availability of new cleanrooms, the only practical option is to accelerate technology migration and equipment deployment within existing facilities. However, as long as HBM's share continues to rise, a significant portion of this investment will be absorbed by the transition to HBM rather than by expanding commodity DRAM bit supply.

As the proportion of wafers allocated to HBM increases, the wafer pool available for commodity DRAM contracts sharply. We estimate that roughly 70% of the new wafer capacity added by the three major memory manufacturers is being directed toward HBM. In practice, wafer growth for commodity DRAM is running at barely 5% annually—a rate that makes it mathematically impossible to satisfy 22% bit-demand growth.

Production and sales per wafer (quarterly)

	HBM3E	HBM4	D1b Server	D1c Server
Net dies	578	503	528	1,750
Content per die (Gb)	24	24	32	16
Yield (%)	60%	50%	40%	95%
Production per wafer (Gb)	8,323	6,034	6,758	26,600
Price (USD)	1.7	2.3	5.0	1.7
Sales per wafer (USD)	14,149	13,879	33,792	45,220

Source: Samsung Securities estimates

Memory makers' wafter growth: Commodity DRAM vs HBM

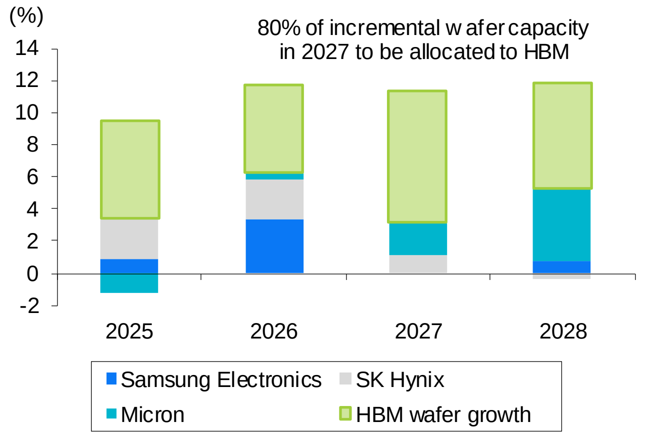

Source: Samsung Securities estimates

Long-term demand outlook: Bit growth converging toward 22%

On the demand side, it is essential to distinguish between the short-term surge and the long-term trend. In 2026, DRAM bit demand growth should remain strong, driven by server DRAM and HBM. Over time, however, AI infrastructure investment should settle into a more stable rhythm. Looking ahead to 2030, total DRAM demand—including demand for HBM—is likely to converge on annual bit growth of roughly 22%. Excluding HBM, the figure falls to approximately 19%.

DRAM: Long-term bit growth forecasts, by application

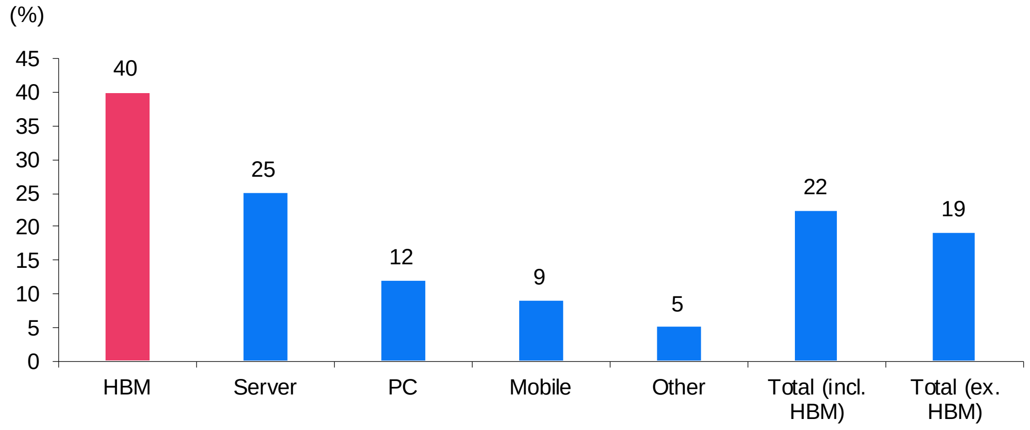

Source: Samsung Securities estimates

Server DRAM is currently delivering bit growth of around 40%. Based on long-term projections from cloud customers, server bit demand is expected to roughly triple over the five years through 2030. This implies a gradual normalization toward an average annual growth rate of around 25% by 2030. Mobile and PC DRAM may face temporary weakness through 2027 due to pricing resistance, but should return to a normal content-growth trajectory from 2028 as pricing stabilizes. Assuming unit shipments remain broadly flat, mobile DRAM bit growth should recover to the high-single digits, while PC DRAM growth moves back above 10%, fueled by higher content per device. HBM, however, operates on a different scale: supply continues to lag demand significantly, and content growth is measured in multiples rather than percentages. As a result, HBM is positioned to sustain annual growth of roughly 40% through 2027-2028.

The precise figure of 22% is less important than what it represents. It is not a demand forecast, but rather a threshold for assessing potential oversupply. As long as bit supply growth remains below this long-term demand-convergence level (22%), downward pressure on pricing is unlikely to intensify. Conversely, if bit supply growth—even after accounting for HBM's wafer penalty—begins to exceed 22%, supply-demand dynamics could flip. For equity investors, the key signal in this cycle is not how quickly prices rise, but when supply bit supply growth crosses this long-term demand-convergence threshold.

Long-term supply outlook: Two-phase normalization in 2028 and 2029

From a supply standpoint, 2026 is unlikely to deliver bit growth above a percentage in the low-20s, as it is constrained primarily by cleanroom shortages. Current capex plans come too late to meaningfully affect 2026 output. In 2027, bit supply growth could rise to around 17% through additional equipment deployment and utilization of remaining cleanroom capacity.

However, the expansion of HBM production—and its associated wafer penalty—will significantly dilute the net benefit. We estimate that the three major memory manufacturers will increase total wafer capacity by roughly 15% annually over the next three years, with approximately 70% of that expansion allocated to HBM. As a result, effective bit growth for commodity DRAM may be limited to just 5% annually—well below the 22% demand-convergence threshold.

The first meaningful inflection point in supply-demand dynamics is likely to emerge in 2028. As major memory manufacturers begin operating newly constructed cleanrooms, the pace of supply expansion should accelerate. Opening a new cleanroom does not immediately lead to oversupply—it takes time for equipment installation, ramp-up, yield stabilization, and product qualification. Market expectations, however, may begin to shift. While 2026-2027 are likely to be defined by a strong perception of physical supply shortages, 2028 could mark the first visible sign that supply is starting to catch up. At that point, the degree of supplier dominance may begin to ease.

A second inflection point is likely to follow in 2029. As the cleanrooms opened in 2028 approach full utilization—and as additional fabs and equipment investments scheduled for 2029 come into clearer view—supply growth could move to a higher trajectory. Once long-term demand visibility for server DRAM and HBM becomes firmly established, and LTAs are proven durable, memory suppliers are likely to reassess the returns on incremental capex. Capex may no longer follow a gradual upward path; instead, it could increase in a distinct stepwise manner. Groundbreaking activity for next-generation cleanrooms, such as P5 and Y2, could provide the clearest signal that this transition is underway.

Memory suppliers: Major fab construction timelines

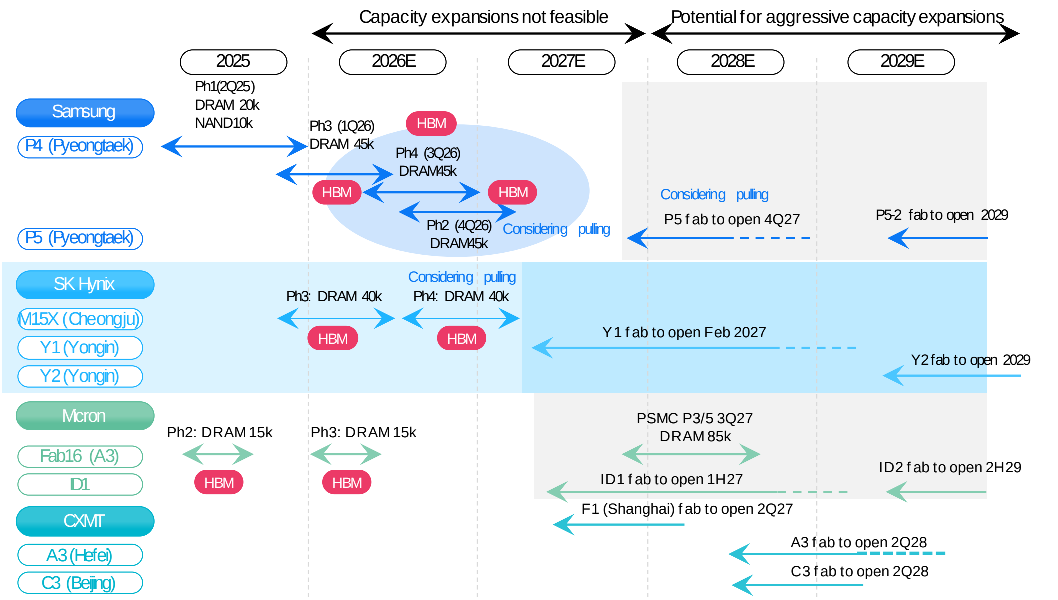

Source: Samsung Securities estimates

That said, the key metric is not total wafer capacity growth, but effective bit growth after accounting for HBM conversion. Even if the three memory makers increase wafer capacity by 15% annually, a large share of that expansion will be allocated to HBM, leaving effective bit growth for commodity DRAM constrained. The critical question for supply-demand analysis in 2028-2029, therefore, is not how much total wafer capacity has expanded, but whether total DRAM bit supply—after accounting for the HBM penalty—crosses the 22% demand-convergence threshold.

DRAM bit share (%), by application

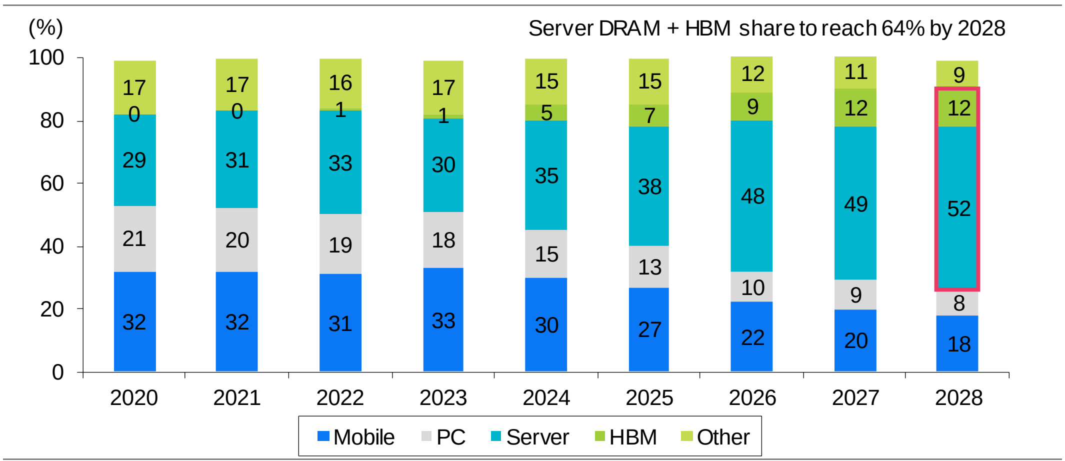

Source: Counterpoint, Samsung Securities estimates

CXMT is a factor, not a game changer

The pace of CXMT's capacity expansion has always been a critical variable. If CXMT leverages China's domestic market to expand its share in DDR4, LPDDR4, and select DDR5 segments, it could put clear pressure on global DRAM supply-demand dynamics. In commodity DRAM, CXMT's growing supply may cap price upside or dampen inventory-building among some customers. During the expected supply-normalization phase of 2028-2029, CXMT's capacity expansion could therefore shift the global DRAM supply curve further upward.

CXMT: DRAM capacity

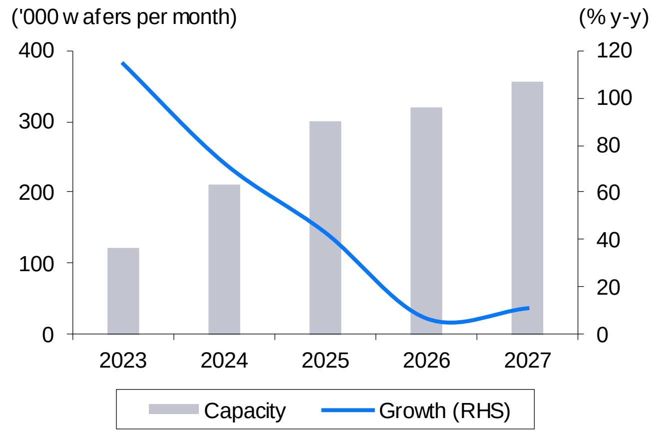

Source: Samsung Securities

Memory makers: Commodity DRAM market shares

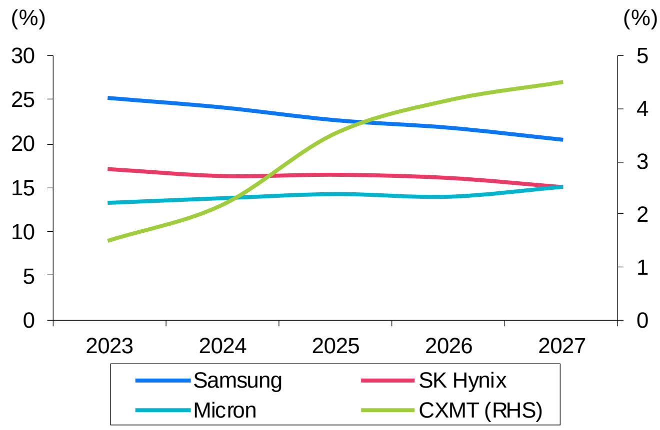

Note: CXMT expected to reach 20% market share in China
Source: Samsung Securities

However, CXMT is unlikely to overturn the core structure of this cycle in the near term. The supply shortage is centered not on commodity DRAM, but on high-capacity DRAM for AI servers, advanced DDR5, HBM, and high-reliability product lines—all of which must be delivered primarily to US customers. In these segments, process competitiveness, yield, customer certification, long-term supply stability, packaging technology, power efficiency, and product reliability are all critical. Even if CXMT expands commodity DRAM supply, it cannot immediately resolve the global supply bottleneck in AI servers and HBM. HBM, in particular, is not a product whose supply can be increased simply by adding DRAM wafer capacity; it is closer to a system-level product requiring advanced back-end processes, CoWoS technology, and close co-development with customers.

CXMT should therefore be viewed through two lenses. First, its ability to ease the 2026-2027 supply shortage will be limited. During this period, bottlenecks are concentrated in cleanroom capacity, advanced node processes, HBM transition, and AI server-certified products. Second, from 2028-2029, CXMT could add downward pressure on commodity DRAM prices. If new cleanrooms from the three global players come online while CXMT supply expands simultaneously, commodity DRAM supply-demand conditions could normalize more quickly. Even in this scenario, however, HBM and AI server DRAM supply is likely to remain tighter than commodity DRAM.

Ultimately, CXMT is a variable that may accelerate the cycle's inflection point or deepen the trough in downturns, but it is not a decisive force capable of disrupting the structural improvement in fab economics. Rather, CXMT's capacity expansion may push the global top-three memory suppliers to prioritize HBM, high-value-added DRAM, and memory solutions over commodity DRAM. In short, CXMT has mixed implications: it risks intensifying price competition in commodity DRAM while underscoring the scarcity value of leading-edge fabs.

Regarding CXMT's IPO, recent reports indicate that the company plans to raise approximately CNY29.5b (USD4.3b), with proceeds earmarked for upgrading DRAM production lines, developing leading-edge DRAM, and investing in next-generation memory such as HBM. In the short term, this IPO may amplify perceptions of DRAM supply tightness, although CXMT lacks the immediate capacity to ease it. Over the longer term, however, if CXMT can sustain capacity expansion with lower capital costs after the IPO, Chinese commodity DRAM supply could rise sharply alongside the ramp-up of new global cleanrooms over 2028-2029. This would cap commodity DRAM price upside and pressure the global trio to accelerate their mix shift away from commodity products and toward server DRAM, HBM, and customized solutions. In the most extreme scenario, CXMT could become the catalyst that bifurcates the memory market into commodity and customized segments.

CXMT: IPO proceeds allocation plan

Use of proceeds	Total investment	Planned IPO proceeds allocation
Technology migration	CNY7.5b	CNY7.5b
DRAM node R&D	CNY18b	CNY13b
HBM R&D	CNY9b	CNY9b
Total	CNY34.5b	CNY29.5b

Source: Company data, Samsung Securities

CXMT: Quarterly sales

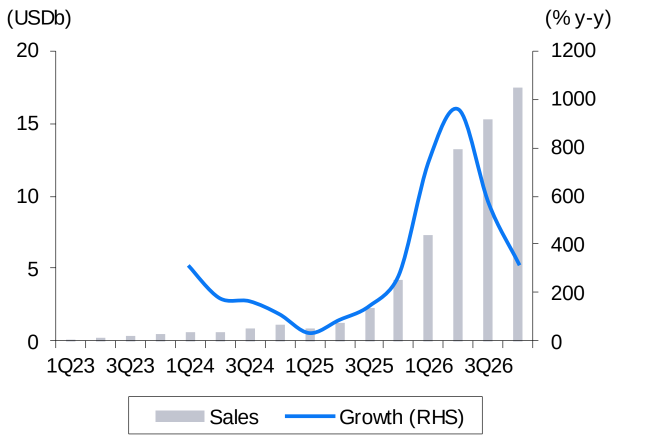

Source: Counterpoint, Samsung Securities

Memory-centric AI: Rising BOM share reflects greater willingness to pay

The rising share of memory in server BOM is not merely a sign of cost inflation—it shows that memory has become the critical performance bottleneck in AI systems.

DRAM market projected to grow at an 80% CAGR over five years, driven fundamentally by AI agents

We forecast that the global DRAM market will reach USD631.5b in 2026, up 313% y-y. This cycle is price-led, not volume-led: while DRAM bit shipments should rise only 22.8% y-y, average selling prices (ASPs) are projected to surge 236% y-y. This does not point to a simple demand recovery, but to a structural tightening of supply, coupled with explosive demand for AI servers—significantly strengthening pricing power. Notably, HBM's revenue share is expected to fall from 21.3% in 2025 to 9.9% in 2026—not because HBM demand has weakened, but because commodity DRAM prices have risen even more sharply, diluting HBM's relative weight. Based on the commodity DRAM segment alone, the market is expected to grow 373% y-y.

By 2027, the DRAM market is projected to expand to USD1t, up 59.6% y-y. The most important trend is that while revenue growth should continue, bit shipment growth should slow. HBM bit growth is expected to decline from 72% in 2026 to 40% in 2027, while commodity DRAM growth should fall from 19% to 15.1%. Supply growth should continue to lag demand, sustaining a structural shortage through 2027. The key issue is not weakening demand, but prolonged supply constraints. Although price growth may moderate from 2026 levels, absolute price levels and profitability should remain historically high.

Longer term, DRAM is expected to sustain growth through 2028. Five consecutive years of growth (2024-2028) at an 80% CAGR would be unprecedented in DRAM industry history—an anomaly not driven by a cyclical upturn, but by the fundamental transition to the AI agent era. Historically, DRAM demand was tied to device adoption—PCs, smartphones, and servers. Since 2023, however, HBM's emergence has begun reshaping the market. With the commercialization of AI agents in 2025, memory demand is now being driven not by replacement cycles or unit growth, but by a new engine: the large-scale build-out of AI infrastructure. This makes the current cycle structurally distinct from any prior cycle.

DRAM market growth

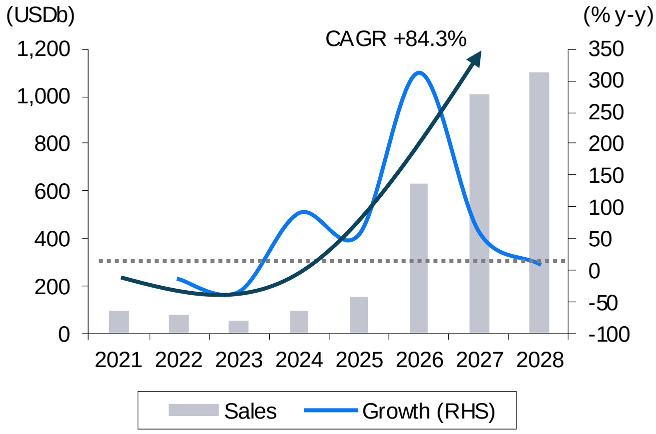

Source: Counterpoint, Samsung Securities estimates

DRAM market ASP growth vs bit growth

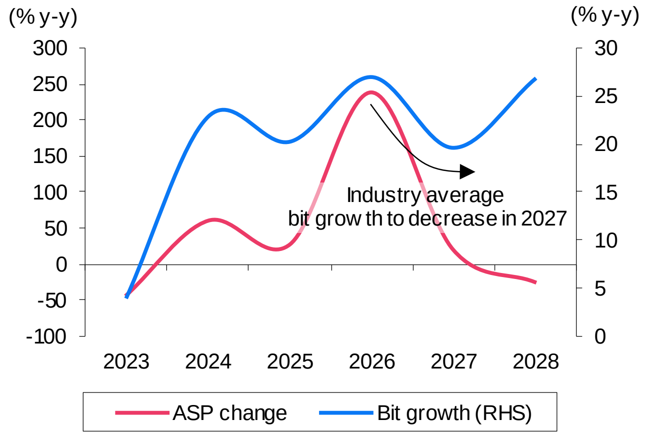

Source: Counterpoint, Samsung Securities estimates

Memory's share of server BOM should continue to rise

Cloud service providers' (CSPs) memory strategies are no longer focused on incremental improvements to individual components. Instead, they are expanding the memory bottleneck beyond HBM to include CPU DRAM, external memory, and storage tiers—and integrating them under unified software control.

Traditionally, memory accounted for 10% of server BOM. With the proliferation of GPU-based AI servers, this share is rising rapidly. We estimate that it reached 17.5% in 2025, driven by HBM adoption, DRAM content growth, and ASP increases—and that it could expand to 37.2% by 2026. This is not a cyclical pricing phenomenon; it reflects structural changes in AI workloads.

The key lies in the evolution of AI models. LLMs have improved performance by scaling parameter counts, inevitably increasing demand for memory bandwidth and capacity. More recently, the rise of reasoning models and Mixture-of-Experts architectures has elevated the importance of contextual data—specifically, KV caches—that are actively referenced during inference, not merely stored as static parameters.

Notably, inference is divided into prefill and decode phases. During these phases, input context is stored as KV pairs and repeatedly accessed—a structure that has become central to AI performance and latency. As a result, the focus of AI performance improvement has shifted from raw compute power (FLOPs) to how quickly data can be stored and retrieved. As long as this architecture persists, memory's share of server BOM is likely to remain high.

Nvidia AI accelerators: Memory portion of BOM

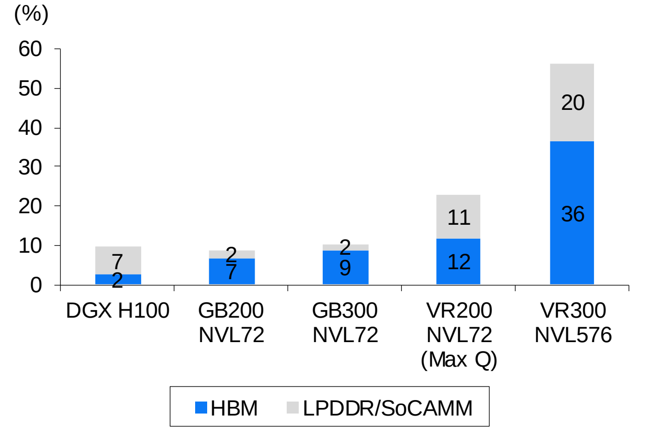

Source: Samsung Securities estimates

Server DRAM + HBM market size

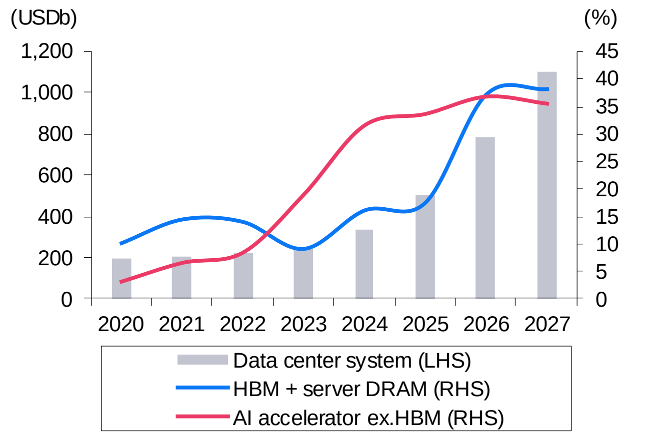

Source: Counterpoint, Samsung Securities estimates

The bottleneck shift: From compute to memory

The most significant change in AI inference is the relocation of the performance bottleneck. Where FLOPs once dictated performance, memory bandwidth and cache architecture are now the limiting factors.

This is especially evident in the decode phase. Nvidia has repeatedly emphasized that prefill is compute-bound, while decode is memory-bandwidth-bound. In autoregressive decoding, each token generation requires repeated access to the existing KV cache, making memory access volume disproportionately high relative to computation. Generation-related attention operations exhibit such a high ratio of KV cache accesses to FLOPs that memory bandwidth—not compute—is now widely recognized across the industry as the true bottleneck. This is not merely a vendor claim, but a structural consensus.

Product strategies reflect this shift. Nvidia's recent roadmap is not centered on isolated hardware advances, but on holistic control of a hierarchical memory architecture that extends from on-chip memory to external DRAM and storage. This system-wide integration spans hardware, infrastructure, and software:

- **NVLink-C2C:** A cache-coherent, high-speed interconnect enabling unified memory access between CPU and GPU at the physical layer.
- **BlueField DPU:** Offloads network and storage paths to accelerate data movement and storage access.
- **ICMX:** A composable memory architecture that extends GPU access to external memory, including DRAM and NAND.
- **HMM (unified memory):** A software layer that allows the GPU to page-manage and access CPU memory as if it were local.
- **TensorRT-LLM:** Implements KV cache management, including offloading, and inference-specific memory optimization.
- **Dynamo:** Optimizes memory/compute scheduling and resource orchestration for inference workloads.

This signals a fundamental transformation: memory in AI is no longer simply about adding capacity—it is about how efficiently the entire memory hierarchy is utilized and controlled. Performance is now determined by memory management.

Memory demand expansion: KV cache distributed across memory hierarchy

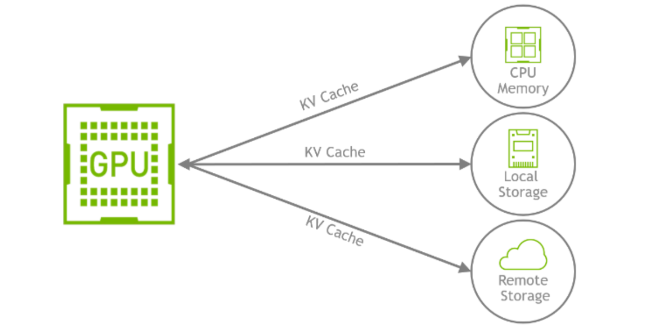

Source: NVIDIA

Addressing decode-stage memory bottlenecks through KV-cache tiering

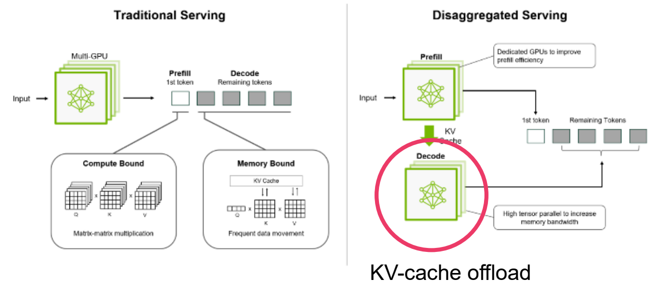

Source: NVIDIA

NVLink-C2C unified memory: Integrating CPU and GPU memory

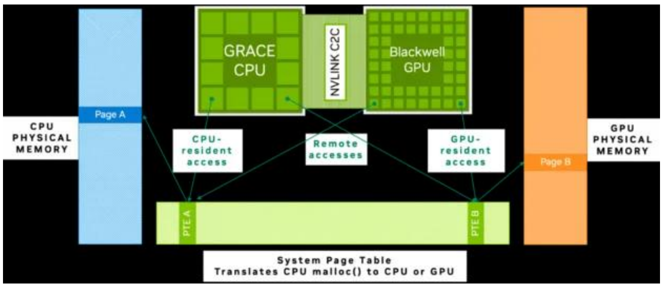

Source: NVIDIA

BlueField offload: A data path for memory expansion

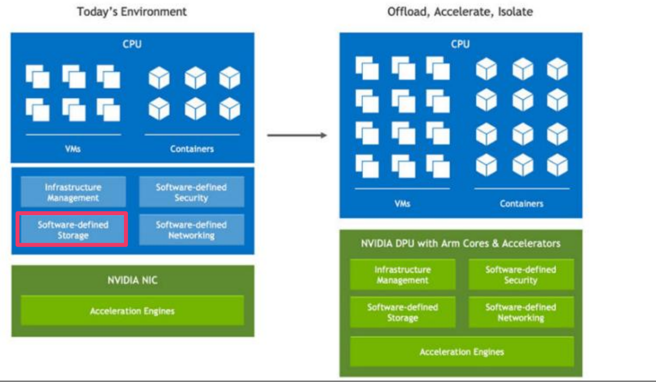

Source: NVIDIA

CMX context memory storage platform unveiled during CES 2026 keynote

Source: Nvidia

Long context and concurrency: Catalysts for memory demand

A key trend in AI models is the growing emphasis on long context windows (the number of tokens a model can process in a single inference) and concurrency (the ability to handle multiple users or sessions simultaneously). Both directly increase demand for KV cache. KV cache size scales linearly with context length, but when combined with rising concurrency, total KV cache requirements grow multiplicatively. Because KV cache directly translates into memory consumption, it has become the primary driver of exponential memory demand.

Google Cloud has explicitly stated that KV cache footprint is directly influenced by context length, concurrency, and throughput. Anthropic, in addressing the limitations of context windows for long-running agent-based tasks, has highlighted architectural innovations such as persistent memory and cache compaction—underscoring the view that context windows themselves are a system bottleneck closely tied to DRAM capacity.

Concrete figures underscore the scale of the burden. In Sep 2025, Nvidia estimated that for Llama 3 70B, a single user's KV cache under a 128k-token context window could reach approximately 40 GB. This represents a dramatic increase in memory requirements even for standard inference sessions. As recently as 2024, the same model was discussed as operable within only a few gigabytes of memory under short-context conditions using quantization. The implication is a fundamental shift: whereas model size once dictated memory requirements, context length is increasingly becoming the determining factor.

The push toward longer context windows and higher concurrency reflects a competitive environment that demands greater scalability and higher service quality. These developments indicate that memory is no longer merely a cost component, but a core resource that determines service quality and operational efficiency. Services capable of supporting long-context, high-concurrency workloads can deliver richer user experiences, potentially reshaping pricing strategies and monetization models. Ultimately, memory has evolved from a cost factor into the infrastructure that defines service competitiveness.

Long-context costs: KV-cache scaling requires memory hierarchy

System prompt length (tokens)	Best-performing storage setup	Mean TTFT (ms) chg vs HBM only (%)	Input throughput chg vs HBM only (%)	Mean end-to-end latency chg vs HBM only (%)
1000	HBM + CPU RAM	5%	1%	-1%
5000	HBM + CPU RAM	-6%	27%	-21%
10000	HBM + CPU RAM	121%	23%	-19%
50000	HBM + CPU RAM + local SSD	48%	69%	-41%
100000	HBM + CPU RAM + local SSD	-3%	130%	-57%

Note: Large cache (12.6m-13.7m tokens) saturates HBM and CPU RAM, spilling to local SSD
Source: Google Cloud

Evolution of AI hardware roadmaps: Memory-centric architecture is not a memory-cycle phenomenon

AI's transformation is directly reshaping hardware roadmaps. It is no longer unusual for AI accelerator vendors such as Nvidia and AMD to emphasize not only raw compute performance, but also memory capacity, bandwidth, and memory-control capabilities when discussing system performance. AI systems are increasingly viewed as requiring balanced improvements across the "processor-memory-fabric" stack—three interdependent layers in which memory is no longer a peripheral component, but the central determinant of overall system performance.

In the early stages of LLM adoption, HBM's role in supporting massive parameter sets was paramount. HBM remains critical not only for training, but also for the prefill phase of inference. As such, it is difficult to argue that demand for HBM has weakened. On the contrary, HBM requirements continue to rise, with supply constraints persisting for a third consecutive year. As inference workloads expand, however, memory requirements are broadening beyond HBM to encompass a wider range of memory technologies.

- HBM remains indispensable for high-bandwidth workloads, including prefill and certain compute-intensive operations.
- Server DRAM and GDDR are increasingly being deployed for capacity-oriented workloads, including KV-cache storage and partial offloading.

As a result, AI servers are increasingly converging on hierarchical memory architectures. Given HBM's inherently high manufacturing cost and constrained supply, deploying HBM across all workloads is impractical. AI systems are therefore naturally evolving toward hybrid architectures comprising HBM, DRAM, storage, and SRAM. The current memory challenge is not a cyclical supply-demand imbalance—it is a structural consequence of AI system architecture.

KV-cache offloading sustains DRAM demand

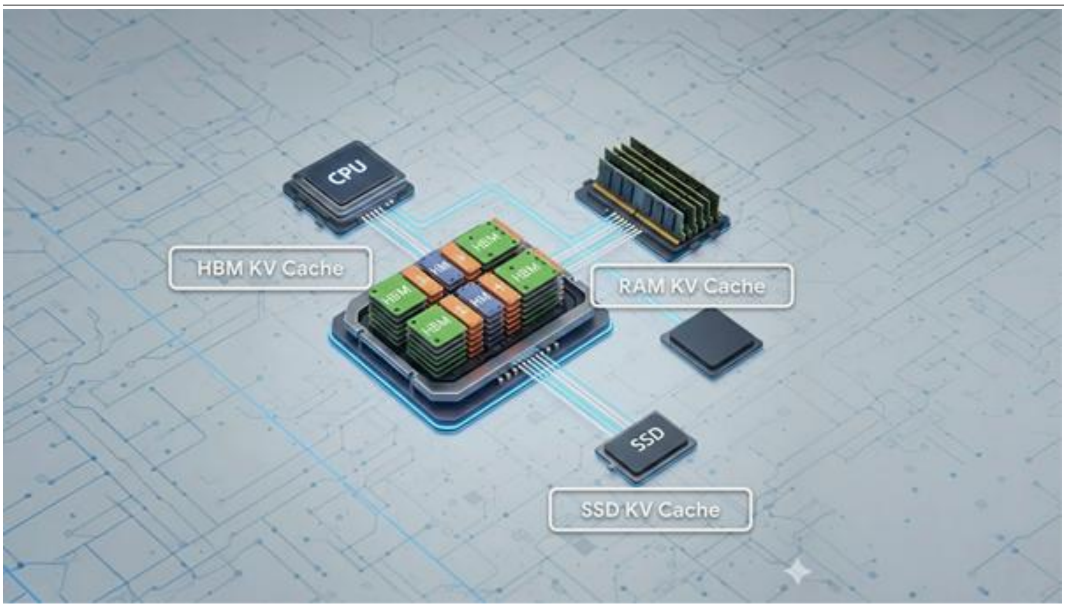

Source: Google Cloud

LTAs are coming

Long-term agreements (LTAs) will mark a significant milestone for the memory industry. They do not guarantee long-term profitability, but they clearly signal that the industry cycle is becoming more prolonged.

LTAs are evidence of supply-demand imbalance

Long-term agreements (LTAs) are gaining traction, particularly for server DRAM. On the surface, LTAs appear to be contracts that lock in price and volume. In practice, they are structured risk-sharing arrangements born of supply-constrained environments. These deals cannot be driven by demand alone—buyers and sellers must align on their risk-return trade-offs.

Server DRAM has created a rare alignment of interests between customers and suppliers. From the customer perspective, securing memory and energy has become essential in data center planning. From the supplier perspective, massive capital investments can be justified only when buyers guarantee order volumes that exceed what mobile or PC markets can absorb. Memory manufacturers seek to defend high prices even during downturns, while customers aim to secure supply during upturns.

LTAs for DRAM are a critical tool for extending the duration of memory profits. In the past, high prices immediately triggered peak-cycle debates. Now, suppliers are pursuing volume commitments of three years or more through LTAs to secure future profits. When upfront payments are included, cash flow stability for memory suppliers improves further. The pricing model has also shifted. The original fixed-price model no longer works for suppliers. Hyperscalers have begun proposing more-flexible terms, including price floors with room for upward price adjustments, to secure LTA volumes.

What it means for supply to be locked in

LTAs contribute to supply-demand stabilization, but they are about more than shipment smoothing. They reflect deeper shifts in supply-demand dynamics that must be understood. Once an LTA is signed, a significant portion of total supply becomes locked into long-term contracts. LTAs are not merely contractual agreements; they are preemptive claims on supply. As a result, customers that fail to sign LTAs are pushed into a thinner, more competitive market, where they must fight for limited availability.

The market does not become more stable—it bifurcates. LTA customers secure stable volumes at high but predictable prices, while non-LTA customers face heightened uncertainty and are forced to contract at sharply rising spot prices. The resulting rise in DRAM prices creates secondary upward pressure on LTA pricing: even if direct price pass-through is constrained by contract terms, the overall blended ASP faces structural upward pressure.

Mobile/PC DRAM is particularly vulnerable to this dynamic. As LTAs expand around server demand, memory manufacturers have stronger incentives to reallocate capacity toward higher-paying server customers. Lower-priority mobile and PC customers therefore face reduced access to supply, weaker pricing power, and greater supply uncertainty. This is clear evidence of a growing divide between customers who secure supply upfront and those who do not.

A similar trend is emerging within mobile DRAM itself. Leading customers agreed to higher prices in 2Q, then signed forward contracts to secure volumes for 2Q-4Q. The market structure now mirrors that of server DRAM: prices bifurcate, secondary-tier customers face price spikes, and the blended ASP ends up higher than expected. There is, however, one important difference. As capacity shifts toward server DRAM, total mobile DRAM supply is likely to decline. Meanwhile, mobile demand is inherently price-elastic—beyond a certain threshold, transactions can collapse abruptly. Thus, while server DRAM faces supply-demand pressure alongside price escalation, mobile DRAM is more likely to experience tight supply conditions accompanied by volume contraction. Market researcher Counterpoint expects mobile DRAM's share of total supply to fall from 28% in 2025 to 24% by 2027.

DRAM bit share, by application

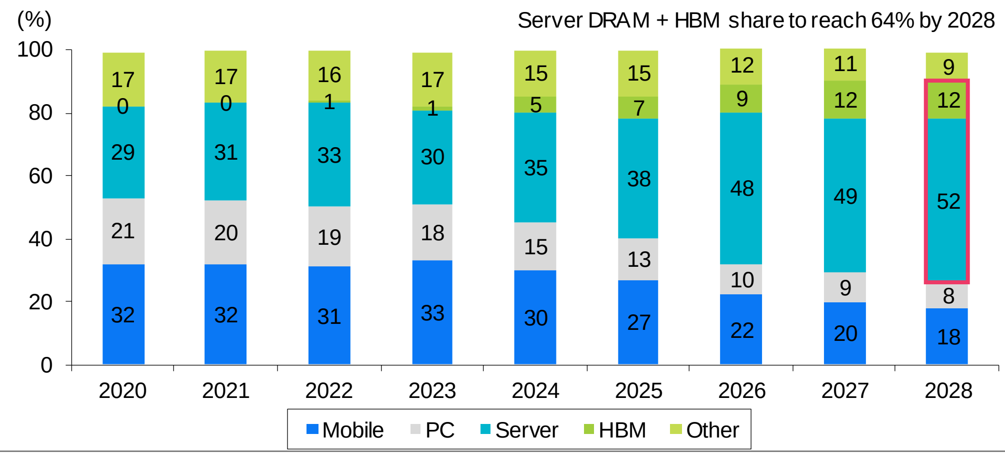

Source: Counterpoint, Samsung Securities estimates

For memory chip manufacturers, are LTAs a strategy or an opportunity cost?

In the current supply-shortage environment, almost any strategic choice tends to favor memory manufacturers. From an investment perspective, however, LTAs are not merely a stability tool—they represent a deliberate choice to reduce earnings beta. Two distinct strategic paths exist.

First, suppliers can aggressively expand LTAs. This approach stabilizes cash flow during downturns, turns customer relationships into long-term partnerships, and lays the foundation for expansion into solutions businesses.

Second, suppliers can limit LTAs to maximize price upside during upturns. If LTA requests themselves signal supply-demand imbalance, they also imply further upside for spot prices. In that case, it may be more rational for memory manufacturers to maximize excess profits by leveraging prevailing market prices rather than locking in pricing through contracts.

When both strategies coexist, the latter is likely to be more advantageous for near-term earnings and share-price performance. The later and smaller the LTA commitment, the greater the benefit from rising spot prices. Downturn risk is a future uncertainty, while excess profits in an upturn are immediate and tangible. Investors, in particular, are drawn to companies with high earnings beta during industry upturns.

As a result, from the customer's perspective, securing large LTA volumes may feel like surrendering future price appreciation potential.

LTA changes the value of an empty seat

Source: Samsung Securities

Evolving LTA pricing structures

LTA prices are typically agreed in advance and remain stable despite market fluctuations. Much like a bond coupon or futures price, LTA pricing serves as a tool to mitigate future risk.

However, a new pricing structure has been proposed to address DRAM suppliers' core concern (ie, the worry that signing LTAs may lead to forgoing future upside in server DRAM prices). The proposed model introduces a hybrid structure in which customers receive a price floor while the price ceiling remains open. Under this arrangement, DRAM manufacturers can still capture excess profits during upturns while being protected from losses during downturns. This effectively embeds a price option within the LTA contract. In exchange for the option value, the customer could receive a discounted LTA price or guaranteed supply volume, creating mutual benefit.

We view this structure as a more reasonable pricing mechanism that better aligns the interests of both parties. Given memory manufacturers' current negotiating power, we expect them to increasingly favor this option-like pricing model. Under this structure, LTAs could become deeply entrenched across DRAM suppliers.

To illustrate, suppose an investor needs to purchase a large volume of a thinly traded listed stock. As the investor's ownership stake increases, the available float shrinks, bid-ask spreads widen, and share-price volatility rises. Not only does the volatility of the shares purchased increase, but the value of the investor's untraded holdings may also become more volatile as a result. The same dynamic applies to DRAM.

If LTAs account for more than half of all server DRAM bit shipments, the supply of DRAM available in the spot market diminishes, increasing price volatility. With price-inelastic server demand becoming the dominant force in the market, DRAM price volatility has already increased—and LTAs further amplify it. If elevated DRAM price volatility is embedded in LTA pricing through option-like contractual mechanisms, even minor DRAM price movements can materially alter LTA profit profiles, amplifying earnings volatility for LTA signatories.

LTA mitigates price volatility

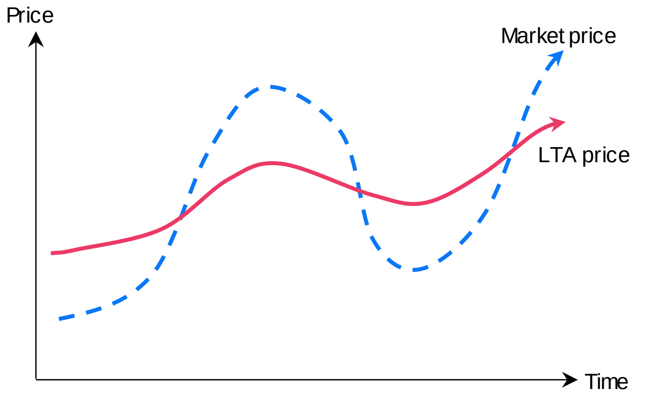

Source: Samsung Securities

New LTA contract structures

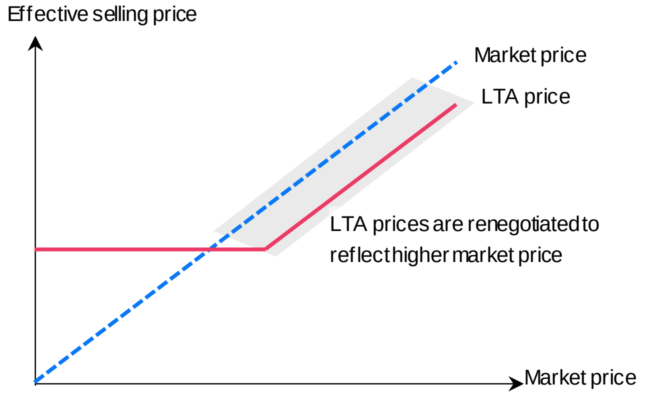

Source: Samsung Securities

Between LTAs and HBM, which is more stable?

LTAs in memory markets improve supply visibility and price stability during periods of tight supply. However, they have limited effectiveness as a structural lock-in mechanism. Memory products are inherently standardized and operate within a multi-sourcing ecosystem. During downturns, customers prioritize pricing, supply terms, and vendor diversification over contractual commitments. When products are interchangeable and switching costs are low, agreements signed during boom periods are often renegotiated or abandoned, resulting in significant market-share shifts. In downturns, customers seek to optimize inventory, delay purchases, and use vendor competition to renegotiate terms—the same dynamic currently unfolding in the DRAM market, where suppliers are actively pursuing contract revisions. Thus, while LTAs may improve short-term earnings visibility, they are unlikely to fundamentally extend the duration of long-term profitability.

Could larger advance payments and the addition of penalty clauses make LTAs more binding? Certainly. Such measures would strengthen contractual obligations and improve earnings visibility. However, legal enforceability does not necessarily translate into economic effectiveness.

What truly matters is platform dependency—the extent to which customers become locked into memory products tailored to their specific system architectures. Current DRAM and NAND products are built to industry standards; as long as technical specifications are met, switching vendors entails minimal cost. As a result, customer loyalty to a single supplier remains limited, and multi-sourcing strategies continue to dominate. If competing products are available at attractive prices, LTAs alone cannot create durable structural advantages. However, when customers are locked in through customized product development, lengthy development cycles, exclusive technology adoption, high production stability, and regional distribution networks, switching costs rise materially, reinforcing the effectiveness of LTAs as a lock-in mechanism.

Against this backdrop, HBM offers significantly greater revenue stability. The lock-in effect is inherent. Technical differentiation and customization are not optional—they are essential. Even without formal LTAs, long-term partnerships between HBM suppliers and customers are effectively mandatory. Industry standards for HBM are often secondary to the system-level optimization required for specific GPGPUs or ASICs. Developing and stabilizing yield for a particular HBM configuration takes years, and reliability is non-negotiable. Capacity constraints further limit customer flexibility. As a result, customers cannot easily switch suppliers. Pricing power in HBM during periods of tight supply differs fundamentally from that in commodity DRAM. In HBM, pricing reflects technical complexity, customization effort, and differentiated performance—not simply market supply and demand.

Annual HBM revenue mix: SEC vs Hynix

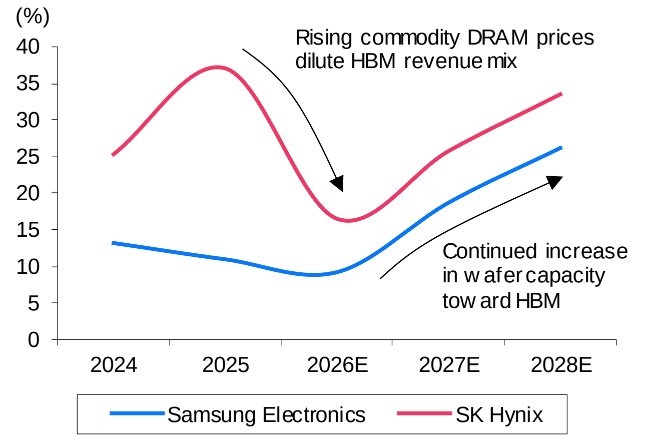

Note: HBM portion of SEC memory division's revenue
Source: Samsung Securities estimates

Annual HBM operating profit mix: SEC vs Hynix

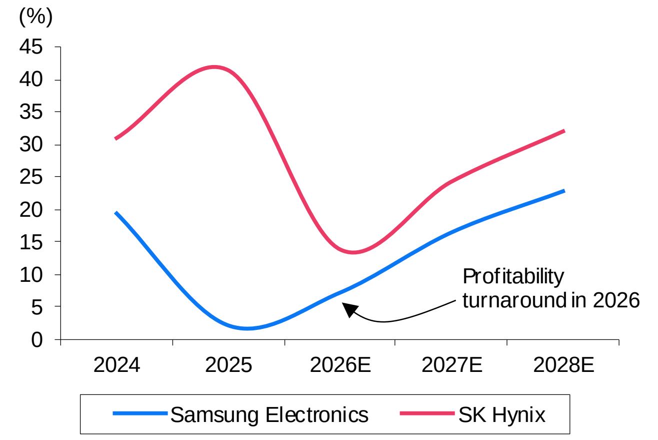

Source: Samsung Securities estimates

Indeed, there was a period when LTAs were widely adopted in the automotive semiconductor sector:

- Following the pandemic-induced semiconductor shortage, LTAs became widespread between automakers and automotive semiconductor suppliers.
- As supply-demand conditions normalized over 2023-2024, some contracts were modified through push-outs, cancellations, or mutual agreements.
- Microchip Technology explicitly cited order cancellations and schedule delays as key risks.
- ON Semiconductor noted that even contracts with non-cancelable provisions were revised by mutual consent.

These examples illustrate that legal enforceability does not always align with economic reality. Ultimately, LTAs are tools for enhancing short-term visibility, but they do not guarantee structurally sustainable revenue.

LTAs have limits—broader industry transformation is required

LTAs represent a meaningful evolution in the memory industry. They improve visibility into pricing and volumes, and they help protect profitability during periods of supply tightness. However, LTAs remain contractual arrangements layered on top of underlying market dynamics—they do not, by themselves, fundamentally alter the industry's profit structure. Once supply-demand imbalances ease, contracts are renegotiated, and multi-sourcing dynamics reemerge. If a single contract could reshape an industry's profit model, why did the business structure remain unchanged through multiple supercycles?

Ultimately, the long-term profitability of memory companies will be determined not by contracts, but by their ability to create deeper customer lock-in through differentiated business models. At this stage, the memory industry must evolve beyond its role as a mere component supplier and transition toward a solutions-oriented business model. LTAs are a positive development—but their impact should not be overstated.

The memory solution business is more important than LTAs

The competitiveness of memory is no longer defined by capacity or cost per bit, but by who designs the data flow. For memory companies to achieve a true rerating, they must evolve from being mere bit suppliers into data-flow designers.

Critical to distinguish between "memory is important" and "memory firms are important"

The key is to define precisely what "memory-centric" means. The rise of memory-centric computing could be misread as merely an expansion in commodity demand—that is, a larger market for DRAM and NAND. But if "memory-centric" is to represent a structural shift in business models, it must go beyond scale. It requires:

- **Design-in:** Active participation by memory companies in system architecture design
- **Data control:** Direct influence over data movement and compute topology
- **Platform lock-in:** Performance differentiation tailored to specific customer platforms

In other words, memory companies must move beyond simply producing more bits. They must actively participate in system-level design—including memory hierarchy, data movement, and caching strategies—to deliver customer-specific memory solutions. Only then can they build a structural moat that transcends commoditization. While the importance of memory in AI is undeniable, memory manufacturers are unlikely to see a fundamental shift in their role or valuation if control over memory architecture remains with system vendors such as Nvidia, AMD, or cloud providers.

In the AI era, memory is no longer a passive component. But its importance alone is not enough to elevate the status of memory makers. The essence of "memory-centric" lies in whether memory companies can define system-level performance—not simply supply it. This should be the pivotal strategic variable for their future.

The real transformation lies in shifting from component supplier to solutions provider

To reduce earnings volatility and secure sustainable margins, memory companies must increase platform dependency—not through contracts, but through integrated offerings. While maintaining economies of scale through JEDEC-compliant standard products remains essential, the next frontier is expansion into custom memory subsystems, hierarchical memory architectures, and memory-aware control software, each deeply linked to clients' overall system architectures.

We no longer believe the competitive advantage of memory companies lies primarily in reducing cost per bit. What is becoming increasingly critical is which partner can co-design memory behavior, placement, hierarchy, and control in line with customers' specific architectures. In our discussions with these companies, it is clear that they fully understand this shift. We therefore believe memory companies are actively experimenting with new memory-layer innovations—such as PIM, HBF, CXL, and Nvidia's CMX support—drawing lessons from their success in HBM. New memory technologies are no longer judged solely on their ability to reduce costs versus competitors; their capacity to enable memory control is now equally critical.

The industry backdrop has shifted decisively in their favor. In traditional CPU/AP-centric computing, memory was a passive, processor-dependent component. Today, AI infrastructure requires memory systems that can handle diverse workloads simultaneously: prefill, decode, long-context inference, agent-based reasoning, KV-cache reuse, and multi-model inference. Nvidia's introduction of CMX—a dedicated context-memory layer—reflects its recognition that conventional memory architectures are ill equipped to address the acute memory bottlenecks emerging in agentic AI. In short, a single standardized memory architecture can no longer satisfy all AI workloads. The focus must shift from optimizing memory for cost alone to optimizing it for application-specific needs.

HBM is an excellent first step

HBM represents the most successful starting point for this transformation. It is no longer merely a DRAM package; it operates alongside logic dies such as GPUs and ASICs, integrated through interposers and advanced packaging. Even from the customer's perspective, HBM is not a product that simply meets JEDEC specifications and "works fine." Instead, it is a tightly integrated system component in which bandwidth, power, thermal management, package height, stack count, signal integrity, and yield are all co-optimized for specific platform requirements.

In particular, starting with HBM4, the separation between the base die and core dies has become even more pronounced. SEC has explicitly announced in its 2026 HBM4 mass-production roadmap that it will use a 4 nm logic base die—a level of semiconductor process technology previously unimaginable in memory manufacturing. By leveraging advanced wafers traditionally reserved for logic chips, SEC has significantly enhanced performance and reliability. This is a clear signal that the base die is no longer just a passive memory buffer; it has evolved into a critical control and I/O interface layer. Although HBM retains the external form factor of standard memory, its business model has already shifted: success now depends on co-development with customers' architectures to achieve deep optimization. HBM is thus the first concrete example of memory moving from a commodity product to design-dependent memory—a fundamental shift required for a rerating of memory companies.

In this context, HBM's platform effect far exceeds that of commodity DRAM. Commodity DRAM can be easily substituted across vendors as long as JEDEC compliance is met. HBM, by contrast, requires extensive validation and production stabilization for each specific GPU or ASIC generation. It demands co-design of power delivery, thermal management, and electrical characteristics within the package. The relationship between customers and memory suppliers has shifted from simple component procurement to collaborative roadmap partnerships. This is why the focus should not be on LTAs, but on the deeper structural shift in the memory-processor relationship.

That said, HBM is not yet a fully locked-in monopoly. Nvidia, Google TPU, and AMD all have strong long-term incentives to maintain multi-vendor supply chains, primarily for supply resilience. Yet this also confirms that competition among suppliers remains active. For example, SEC has already gained leadership in HBM4 within Nvidia's ecosystem, while AMD—long a close partner of SEC—is now also sourcing HBM from Micron Technology. The true significance of HBM is not merely that lock-in has occurred, but that memory companies have, for the first time, demonstrated the ability to create lock-in through design dependency and co-development.

PIM represents memory companies taking a decisive step toward system-level logic

Processing-in-memory (PIM) embeds computation directly within memory arrays, reducing data movement and offloading selected operations from CPUs and GPUs—thereby improving both performance and energy efficiency. PIM is the clearest embodiment of memory makers' evolution from component suppliers to system solutions providers.

SEC first unveiled HBM-PIM in 2021, demonstrating its integration into commercial accelerator systems. Since then, it has expanded the portfolio beyond HBM to include AXDIMM and LPDDR5-PIM, positioning PIM not as a one-off innovation, but as a strategic platform direction. AXDIMM, for example, extends AI processing to the DRAM module level. SEC has claimed that HBM-PIM can deliver up to 2x system performance while reducing energy consumption by more than 70%.

PIM structure

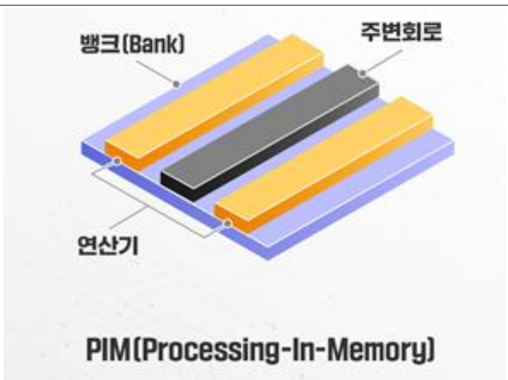

Source: SK Hynix

HBM-PIM

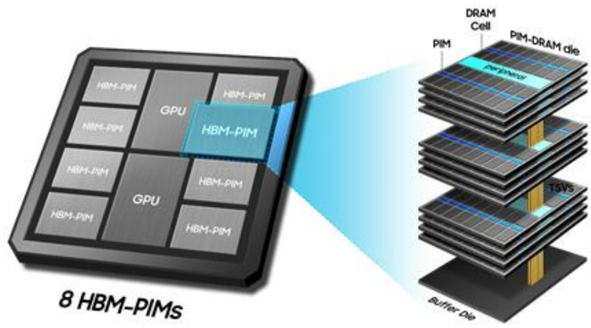

Source: Samsung Electronics

The significance of PIM lies not merely in adding compute to memory, but in enabling memory companies to influence how data is placed and processed. Historically, data access and placement were controlled entirely by CPUs and GPUs; memory served as a passive storage layer. With PIM, memory vendors must understand customers' model characteristics, identify computational bottlenecks, and optimize memory behavior for frequently used workloads. This transforms them from passive suppliers into active co-design partners, requiring engagement at the system architecture stage, much like AI accelerator vendors.

Of course, PIM's slow commercial adoption has been driven in part by this very design dependency. The entrenched dominance of Intel's CPU and Nvidia's GPU ecosystems means that broad PIM adoption requires support from those platforms through compatible compilers, runtimes, and frameworks. Alternatively, memory companies must build their own software stacks—an expensive and complex undertaking.

However, the timing is increasingly favorable. In today's AI infrastructure, memory's share of the BOM has risen substantially, while long-context inference and reasoning optimization have made data-movement costs more critical than ever. For end customers—cloud service providers (CSPs)—the key requirement is a structure that lowers cost per watt or cost per token, and PIM directly addresses this challenge. By reducing data movement, it inherently lowers cost per token. While the decision to adopt PIM ultimately rests with CSPs, the economic incentive has never been stronger.

HBF is a new memory tier where memory control is the core

High-bandwidth flash (HBF) represents another important step in the commercialization of memory as a system solution. In Feb 2026, Hynix and SanDisk announced a joint initiative under the OCP to standardize HBF, positioning it not as a simple storage upgrade, but as a next-generation memory solution for AI inference. This signals that HBF is not merely about making NAND faster—it is an effort to define a new memory tier positioned between HBM and SSDs.

Understanding HBF simply as "AI starting to use NAND" misses its fundamental purpose. The essence of HBF lies not in the storage medium itself, but in its memory-control architecture. In AI inference workloads, performance and cost are jointly determined by decisions regarding which data remain in HBM, which are offloaded to HBF, and when and how frequently data move between tiers. HBF, therefore, is not a passive storage device—it is a memory-hierarchy control system designed to orchestrate data placement and movement.

This capability stems from NAND's inherent architecture. Unlike DRAM, where the memory die largely dictates behavior, NAND relies on sophisticated controllers to manage the flash translation layer (FTL), wear leveling, queue management, data placement, and error correction. HBF extends this model—not merely to manage storage, but to actively optimize data flows for AI workloads. It is not an extension of the NAND business; it is a move toward designing the entire memory hierarchy.

Currently, HBF is in the early standardization phase. Its technical architecture remains fluid, and the industry is exploring multiple design options to meet diverse customer needs. In this phase, the key competitive battleground is not who makes the best NAND, but who controls data placement and movement policies. Historically, NAND vendors have remained commodity suppliers because control over data flow resided with SSD module makers, enabled by standardized interfaces such as UFS and eMMC, which were designed for general-purpose computing environments where data access patterns were relatively predictable. By contrast, AI workloads require dynamic, workload-aware memory management—something existing standards were never built to support.

If memory companies can own the HBF controller and extend control to data placement policies, migration triggers and timing, and software interfaces for inference engines, they can move beyond being NAND suppliers to become providers of custom memory systems. Conversely, if they cede control of this layer to platform vendors such as Nvidia, AMD, or CSPs, HBF will remain a passive, high-capacity extension of SSDs—valuable, but not transformative. The true significance of HBF is not the emergence of a new memory product. It is a test of whether memory companies can reclaim control over data flow.

Expanding HBF memory solution

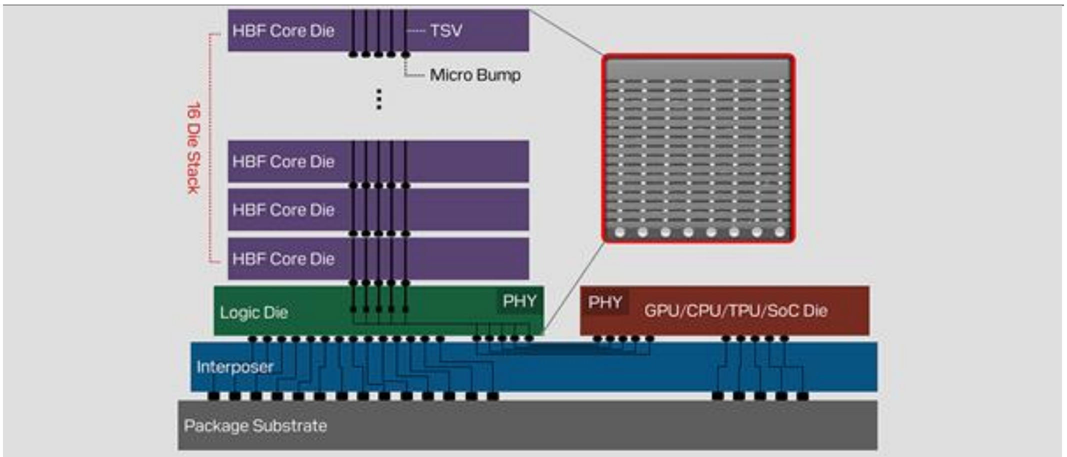

Source: SanDisk

CXL, in a broad sense, follows the same strategic direction

In a broader sense, CXL follows the same strategic direction of expanding into the memory systems business. CXL is not a memory product; it is an interconnect fabric that enables memory expansion, sharing, and tiering between CPUs, accelerators, and memory subsystems. Its value lies not in raw memory performance, but in the architecture of memory placement, allocation, and access policy.

In a CXL environment, the controller becomes even more critical. Decisions regarding which workload receives which memory tier, how memory is shared across domains, and how data movement is optimized directly determine both performance and total cost of ownership. The core competitive advantage of CXL, therefore, is not bandwidth or density, but the ability to intelligently connect, pool, and prioritize memory resources across heterogeneous systems. This requires two complementary capabilities: 1) networked memory pooling and hierarchical management at the hardware level; and 2) software intelligence that understands AI model behavior and workload patterns at the system level, enabling optimal memory policies.

This structural shift transforms the memory business from a commodity model into a system-level solutions business. As CXL adoption grows, the value proposition shifts from selling DIMMs to designing memory placement policies, control software, and system-level integration. CXL is not merely an interface; it is an ecosystem encompassing controllers, modules, firmware, software stacks, and platform compatibility. One reason commercialization has lagged is precisely the difficulty of achieving end-to-end ecosystem coherence.

In this landscape, merchant silicon companies play a pivotal role. Companies such as Marvell, Astera Labs, Microchip, and Montage supply CXL memory expanders and switches—the critical controllers that manage memory connectivity and data movement. If these vendors evolve beyond basic connectivity to include workload-aware bandwidth allocation, memory tiering logic, and real-time traffic optimization, CXL can evolve from a passive interconnect into a differentiated system value layer.

Meanwhile, memory companies are actively entering this space. SEC, Hynix, and Micron Technology have all announced or are developing CXL memory modules. Notably, SEC has gone further, introducing its CXL Memory Module (CMM-D) alongside proprietary CXL controllers and the Samsung Memory Development Kit (SMDK) software stack to enable memory expansion, sharing, and tiering. This is a clear signal that memory vendors are no longer content to remain component suppliers. They are positioning themselves as memory system solution providers.

Yet ultimate control remains uncertain. The critical question is not who designs the CXL controller, but who owns the software stack that defines memory policy; who controls platform-level integration; how much customization is required for each customer workload; and how deeply memory companies can penetrate the data-control domain.

If memory companies can secure control over memory placement policies, software interfaces, and system-level orchestration, CXL becomes a gateway to custom memory systems—a new high-margin business beyond commodity DRAM or NAND. But if control remains concentrated among merchant silicon vendors or CPU/GPU platform owners, memory companies will be relegated to supplying "CXL-enabled DRAM"—a new version of the NAND business model: high volume, low differentiation, and limited influence over system value.

CIM and LPU: Redefining the role of memory and memory companies

The recent rise of Groq's Low-latency Processing Unit (LPU) architecture—partially influencing Nvidia's Rubin platform—has drawn significant market attention. Groq's design leverages hundreds of megabytes of on-chip SRAM as primary memory storage, eliminating dynamic scheduling in favor of compiler-driven static control. This minimizes data movement and delivers a breakthrough reduction in AI inference latency. Nvidia's adoption of LPX-like structures within Rubin confirms a clear industry trend: memory is becoming more hierarchical. Importantly, on-chip SRAM is not designed to replace HBM or server DRAM, which remain essential for large-scale KV-cache storage. Rather, the LPU exemplifies a shift toward tighter co-design of memory and logic to alleviate memory bottlenecks—not by replacing memory, but by reducing dependency on it through architectural innovation.

Compute-in-memory (CIM) represents a more fundamental rethinking, integrating computation directly into the memory array itself. Unlike PIM, which adds limited compute functions adjacent to existing DRAM structures to preserve compatibility with current ecosystems, CIM treats the memory cell—whether SRAM, MRAM, or RRAM—as an active computational unit. This requires rethinking device physics, process compatibility, and circuit design from the ground up. SEC, through SAIT and its foundry division, demonstrated the world's first MRAM-based in-memory computing in 2022. TSMC followed with its own RRAM-based CIM research the same year. CIM is not an incremental upgrade—it is a foundational reimagining of memory's role in computation.

The core question is not whether SRAM, MRAM, or RRAM are "needed" for AI. The critical question is how far the structural evolution of AI can expand the business domain of memory companies. If memory suppliers can propose custom memory-compute architectures tailored to specific AI inference patterns—designing not only the memory, but also the data-movement logic, access sequences, and compute mapping—they can transition from component vendors to system-level memory architecture partners.

Nvidia's memory strategy and coexistence with memory companies

Can Nvidia's CMX and BlueField-4 coexist with the system-level ambitions of memory companies? In 2026, Nvidia defined CMX as a context-memory storage platform, and with BlueField-4, it aims to build a pod-level context tier—an intermediate memory layer shared across multiple GPUs within a data center. DOCA Memos, a software stack embedded in Nvidia's BlueField SDK, manages and shares KV caches across compute nodes and CMX data nodes. This is not merely infrastructure expansion; it is a deliberate architectural move to embed memory placement, reuse, and orchestration directly into the GPU platform's core design. At a high level, Nvidia and memory companies are competing for the same prize: control over data flow, not memory itself.

Yet their strategic domains are distinct enough to avoid outright conflict, at least in the near term.

Nvidia is responsible for defining memory-utilization policies—specifically, how data is placed, moved, and shared. By contrast, memory manufacturers focus on implementing memory technologies with specific performance characteristics at each hierarchical level. Even as Nvidia defines the pod-level context tier, the physical realization of that tier remains open. Memory vendors can propose a range of options—including HBM, SoCAMM, HBF, CXL-extended memory, and even PIM-based architectures—and deliver optimized configurations based on the performance, power, and cost characteristics of each layer. As a result, structural coexistence is feasible in the short term: Nvidia oversees high-level orchestration, while memory vendors focus on lower-layer optimization.

Over the long term, however, competition is inevitable. This mirrors a recurring theme in semiconductor history: who drives performance innovation—the platform (through software) or the component (through hardware)? If Nvidia succeeds in fully optimizing context tiering at the software-hardware co-design level, the physical memory beneath it may become increasingly standardized. Different memory types would still exist, but as interchangeable, undifferentiated building blocks. In that scenario, memory companies would be pushed back into the role of low-margin, performance-and-price-driven suppliers.

Ultimately, for memory to establish a system-level role, the key lies in defining its function for each workload and designing corresponding characteristics and behaviors tailored to those needs. Just as HBM was designed as memory optimized for specific AI accelerator architectures, the future of the memory business should shift from asking "which memory should be supplied to which system?" to "how should memory hierarchies be designed for specific systems?" Those that fail to make this transition risk seeing memory revert to a generic commodity component. Those that lead it, however, can expand memory from a commodity into a customized, system-level business.

The core of memory competitiveness lies in who can best define memory workloads for customer-specific architectures and design memory solutions with the optimal specifications to match them. The memory industry should evolve not as a cost-driven storage manufacturing business, but as a software-like industry that designs data flows.

Maintaining OVERWEIGHT on memory; SEC still our top pick

Raising target prices for SEC and Hynix

We raise our target price for SEC to KRW500,000 and for Hynix to KRW3,500,000. The revision rests on two key pillars.

First, we believe the sustainability of DRAM profitability has been extended through 2028. Given the current fab capacity plans of major memory manufacturers, a full supply-demand balance is unlikely before 2029. The expansion of LTAs further reinforces our confidence in this trend. While demand uncertainty always remains, AI data center investments are driven not by product cycles but by infrastructure build-out. Accordingly, memory demand should be assessed not through quarterly inventory cycles, but through a multi-year capacity-planning lens. If the fundamental driver of demand is the survival and scaling of AI infrastructure itself—regardless of fluctuations in end-user consumption—this demand can and should be assessed over a longer time horizon. We now base our target prices for SEC and Hynix on 2028E BVPS to reflect this extended cycle.

Second, the scarcity value of advanced fabs must be incorporated into valuations. The pace of productivity gains from process transitions has slowed, HBM capex concentration is becoming more prolonged, and demand is increasingly concentrated in server DRAM—all raising barriers to new capacity expansion. When the same capital investment yields less bit growth than in the past, the economic value of existing fabs and their book value must be reassessed. Memory manufacturers' net assets are no longer merely historical capital; they are assets capable of generating higher ROIC than before. As a result, the same book value deserves a premium P/B multiple. We continue to apply a medium- to long-term valuation premium to book value, aligning our target P/B multiple with the ROE-P/B regression line.

Based on this framework, we apply a P/B multiple of 2.2x to SEC's 2028E BVPS of KRW228,000, resulting in a target price of KRW500,000. For Hynix, we apply a P/B multiple of 2.5x to its 2028E BVPS of KRW1,419,000, yielding a target price of KRW3,500,000. These target prices are derived from historical ROE-P/B relationships. However, if memory profitability, sustainability, and capital efficiency have structurally changed—as we believe they have—valuation multiples need not remain anchored to lower historical ranges. We are now entering a new normal for memory valuations. This cycle should redefine market consensus on the fair multiple for memory stocks.

Samsung Electronics: Forward P/B vs ROE

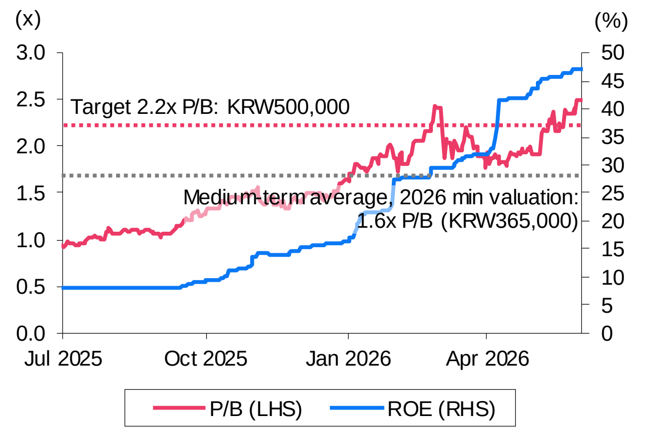

Source: QuantiWise, Samsung Securities

SK Hynix: Forward P/B vs ROE

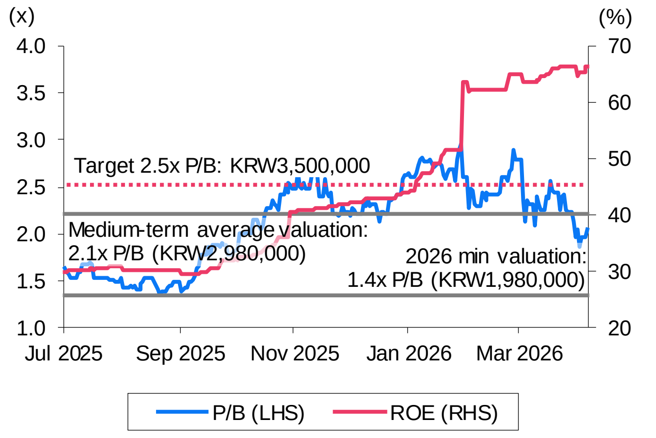

Source: QuantiWise, Samsung Securities

Samsung Electronics: Forward P/B vs ROE

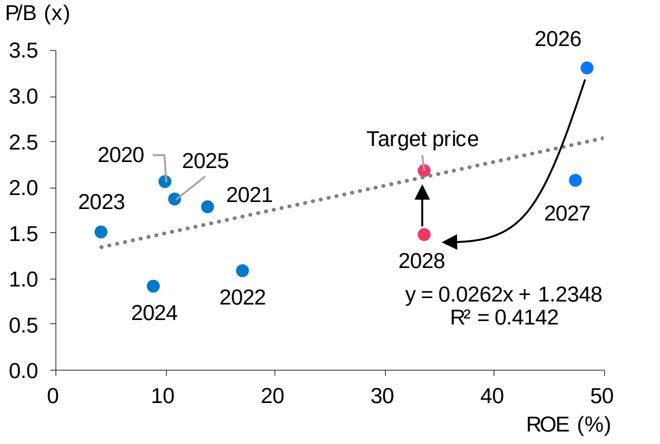

Source: QuantiWise, Samsung Securities

SK Hynix: Forward P/B vs ROE

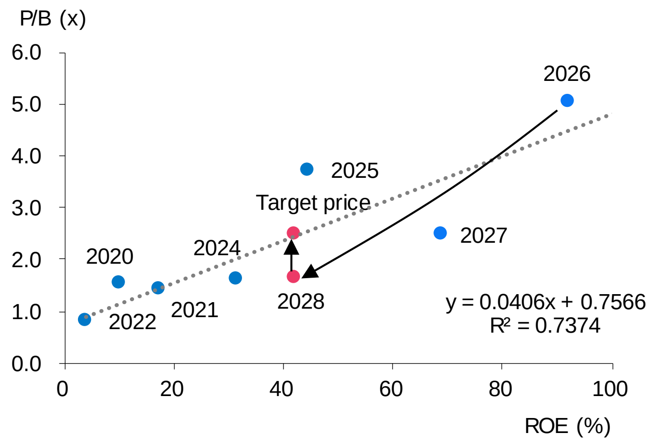

Source: QuantiWise, Samsung Securities

Samsung Electronics: Target-price calculation

(KRW)		Notes
BVPS	234,152	2028E
Target P/B (x)	2.2	Assumes 2028 ROE of 33%
Fair price per share	515,135	
Target price	**500,000**	
Current price	299,000	
Upside (%)	67.2	

Source: Samsung Securities estimates

SK Hynix: Target-price calculation

(KRW)		Notes
BVPS	1,419,571	2028E
Target P/B (x)	2.5	Assumes 2028 ROE of 42%
Fair price per share	3,548,928	
Target price	**3,500,000**	
Current price	2,101,000	
Upside (%)	66.6	

Source: Samsung Securities estimates

Why SEC remains our top pick

SEC is our top pick in the memory sector—not merely because it has regained HBM competitiveness, but because it stands to benefit most directly from the structural upgrade in fab economics that lies at the heart of the new normal in memory valuations.

First, SEC's market share upside is broadening. In HBM, we expect SEC to gain share as AI ASIC adoption grows among key customers such as Google and Amazon, where SEC holds a leading position. In server DRAM, as the incumbent leader, SEC is best positioned to capture the surge in demand. Rising server DRAM prices and growing demand for high-capacity DIMMs should simultaneously improve SEC's product mix and profitability. In NAND, the renewed importance of TLC and the growing role of controller and eSSD optimization leave meaningful scope for competitive recovery, partly offsetting weaknesses in QLC and W2W bonding. Even in non-memory businesses, tight 2 nm foundry capacity could generate spillover benefits for SEC's foundry business. In smartphones, as cost inflation intensifies, premium-segment strength becomes increasingly critical—and SEC's flagship Galaxy lineup and on-device AI expansion should support both market share and margin resilience. Crucially, SEC's share growth should not be driven by HBM alone, but by a holistic AI portfolio spanning memory, foundry, and devices.

Second, as technology differentiation flattens, capacity becomes the ultimate differentiator. This cycle's supply constraint is not about a single product shortage; it is about a shortage of cleanroom space and wafer capacity. SEC's massive fab footprint is therefore poised to be revalued as a scarce asset. Over the past two years, memory shares' relative performance was driven by HBM technical leadership. But as HBM capabilities converge and customers such as Nvidia begin standardizing base dies, the top priority for AI clients is no longer simply the highest-performing chip, but a reliable, scalable partner that can consistently deliver volume and align on roadmap transitions.

Third, foundry strength is now embedded in memory competitiveness. Historically, memory and foundry were treated as separate businesses. But post-HBM4, the growing importance of base-die design and packaging allocation has blurred the lines between them. As HBM evolves from a discrete DRAM product into a system-level component tightly coupled with ASICs, memory competitiveness can no longer be judged by cell process alone. SEC uniquely combines memory, foundry, and advanced packaging capabilities, with direct experience producing base dies for HBM through its own foundry infrastructure.

Fourth, SEC stands to benefit from AI's expansion beyond data centers into AI devices. While today's AI demand originates in data centers, the long-term trajectory points to on-device AI, AI PCs, AI smartphones, and edge devices. The market currently values SEC primarily as an AI memory play. But as AI moves into devices, SEC's diversified portfolio—spanning memory, displays, semiconductors, and consumer electronics—should be re-evaluated with far greater weight.

SEC: Server DRAM market share

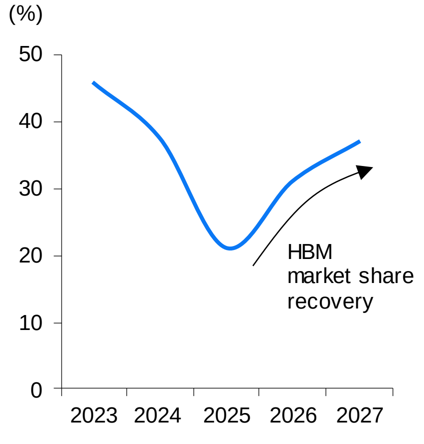

Source: Counterpoint, Samsung Securities estimates

SEC: HBM market share

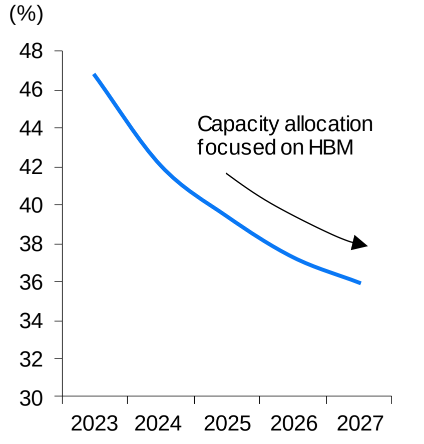

Source: Samsung Securities estimates

SEC: Foundry market share

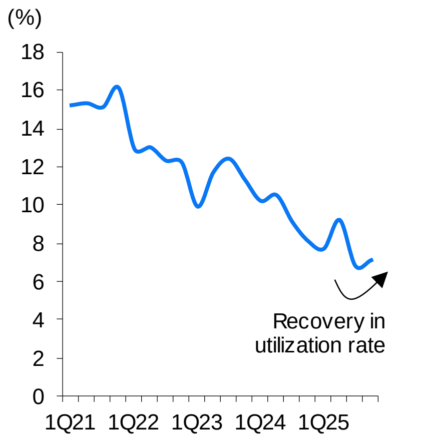

Source: Counterpoint, Samsung Securities estimates

The current market cap gap between SEC and Hynix has narrowed to historically tight levels. However, this compression does not fully reflect a fundamental shift in their intrinsic value.

Until now, Hynix has been strongly rerated due to its HBM leadership, high earnings sensitivity, high-beta appeal, and concentrated demand from ETFs and retail investors. SEC, by contrast, has traded at a discount due to delayed HBM recovery and foundry underperformance. Yet, as SEC's HBM and foundry capabilities rebound and earnings consensus upgrades for SEC accelerate, the earnings leverage gap between the two is narrowing. We believe the current market cap gap compression is driven more by supply-demand dynamics and stock-price elasticity than by true value realignment. When stock prices revert to fundamentals, the valuation gap between SEC and Hynix is likely to settle at a level reflecting the difference in memory capacity scale, plus the value of SEC's non-memory businesses.

Market cap: SEC, SK Hynix, and Micron

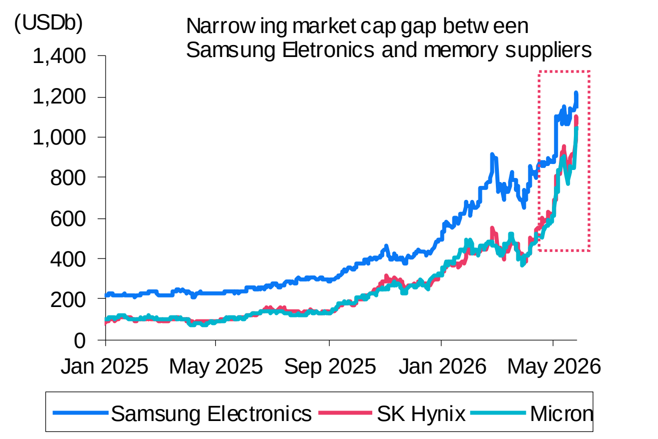

Source: Bloomberg, Samsung Securities

ETF inflows: SEC and SK Hynix

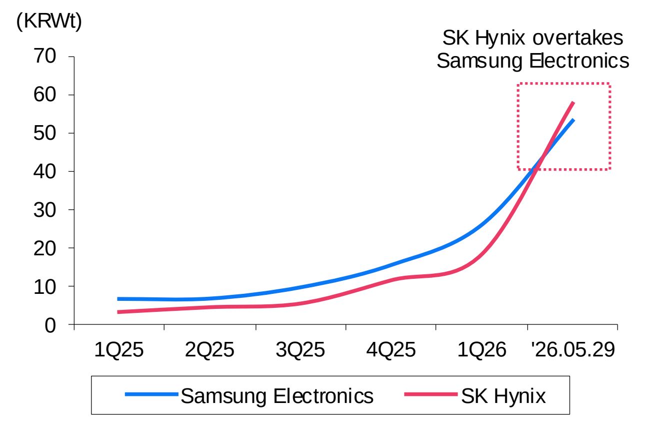

Source: KRX, Samsung Securities

The new P/E benchmark for memory semiconductors

Recent market sentiment has increasingly favored applying P/E multiples to memory semiconductor companies. Looking at the Philadelphia Semiconductor Index (SOX) and Micron's relative P/E, we observe that Micron's discount to the index has narrowed significantly compared to historical levels, indicating a clear rerating relative to broader semiconductor peers. We believe this trend is likely to persist in the near term.

SOX Index's 12-month forward P/E vs Micron's valuation discount to SOX Index

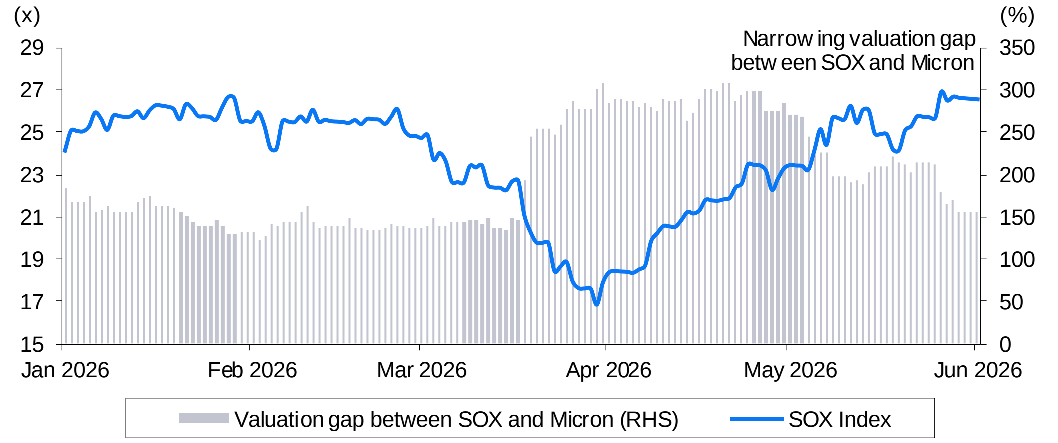

Source: Bloomberg, Samsung Securities

We do not view the use of P/E as inherently unreasonable. If the DRAM industry has structurally reduced the likelihood of entering prolonged loss-making phases—as we believe it has—then earnings-based valuation through P/E becomes a credible and meaningful approach. However, we continue to rely on P/B-ROE as our primary valuation framework because memory remains more cyclical than other semiconductor segments. While AI has lifted memory earnings to higher levels, volatility in memory earnings remains significantly greater than in other semiconductor sub-sectors. This is because demand for memory in AI infrastructure is far less price-elastic than demand for other components. In simple terms, when demand is inelastic, prices move more sharply—and when prices swing more violently, earnings volatility expands dramatically. The greater the leverage of DRAM price changes to earnings, the greater the cyclicality.

Therefore, even if P/E is applied, a cyclical discount must still be embedded. Historical SOX P/E trends show that multiples typically trade at a 15-20% discount in the late phase of semiconductor cycles, reflecting earnings uncertainty. Given memory's higher earnings volatility, this discount could widen further as the cycle matures. In essence, a P/E rerating for memory companies is possible—but it does not mean the cycle has disappeared. Rather, it means the cycle's baseline has shifted upward.

SOX Index: P/E multiple typically derates late in the earnings cycle

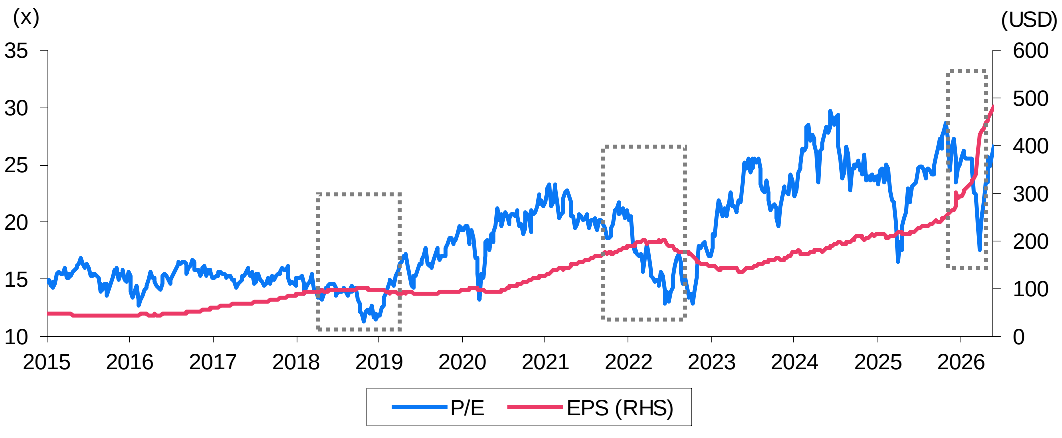

Source: Bloomberg, Samsung Securities

Moreover, P/E and P/B-ROE are fundamentally two sides of the same coin. P/B is a function of ROE and P/E; an expansion in the P/B premium implies the market is assigning a higher P/E to the same level of earnings. From this perspective, P/E serves as an intuitive indicator of where memory valuations stand relative to market sentiment. However, P/B-ROE remains superior for explaining structural changes in fab economics and capital efficiency. We believe the core driver of this rerating is not merely higher earnings, but a structural upgrade in fab ROIC. For this reason, we retain P/B-ROE as our primary valuation framework.

SEC and Hynix's target prices imply 2028 P/E multiples of 7.7x and 7.5x, respectively. While these levels appear high compared with historical memory peaks, the key difference this cycle is not just the level of earnings, but their durability, the scarcity of capacity, barriers to new supply, and the structure of long-term agreements. The central question is no longer: "At what multiple did memory trade in the past?" It is now: "What is the new normal P/E multiple the market will assign to memory under these transformed fab economics?"

12-month forward P/E: Samsung Electronics vs SK Hynix

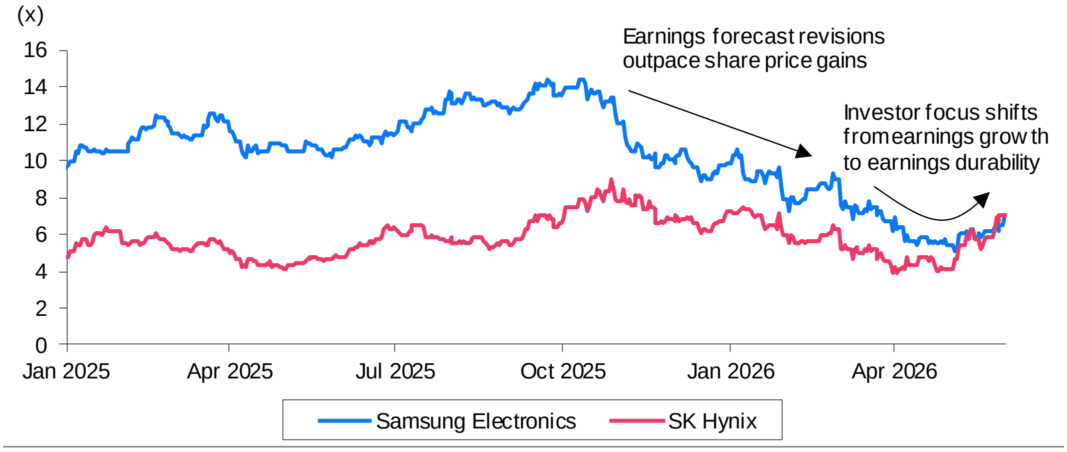

Source: Bloomberg, Samsung Securities

DRAM ASP to rise 50% q-q in 2Q and 15% q-q in 3Q

This cycle's first inflection point should arrive in 2Q26, marking a slowdown in the pace of DRAM price gains. Through 2Q26, DRAM's blended ASP is expected to rise by over 50% q-q—an extraordinary pace not seen in prior cycles. Starting in 3Q26, however, the rate of increase is likely to moderate to a more stable level of around 10%. At this juncture, investor focus should shift from the momentum of price rallies to the sustainability of DRAM operating profit. Given the sharp price surge and resulting profit-taking opportunities, a market correction between these two paradigms is inevitable.

Notably, the pace of DRAM price increases in 1Q26 far exceeded our original cycle assumptions, and both stock prices and earnings estimates have already fully priced in this acceleration. Because earnings momentum has been stronger than anticipated, the reversal effect as momentum cools should also be more pronounced.

We now expect this cycle's second inflection in 2Q28, versus our prior forecast of 2Q27, as DRAM prices are projected to begin declining in 2Q28 rather than 2Q27. This adjustment reflects a change in our assumption about the speed of supply response, rather than an upward revision to our demand outlook. While new equipment deployments and utilization of remaining cleanroom capacity should improve bit growth starting in 2027, the wafer penalty from the HBM transition and delays in ramping new fabs mean that supply growth is unlikely to consistently exceed the demand convergence line of 22%. A sustained price decline requires more than just increased supply; it requires sufficient and credible supply visibility to reduce customers' incentives to accumulate inventory. This condition is unlikely to be met until major new cleanrooms begin full-scale ramp-up, which we anticipate in 2028. Therefore, 2027 should be a period of moderating price growth, not the onset of a structural price decline.

Quarterly DRAM ASP change forecasts*

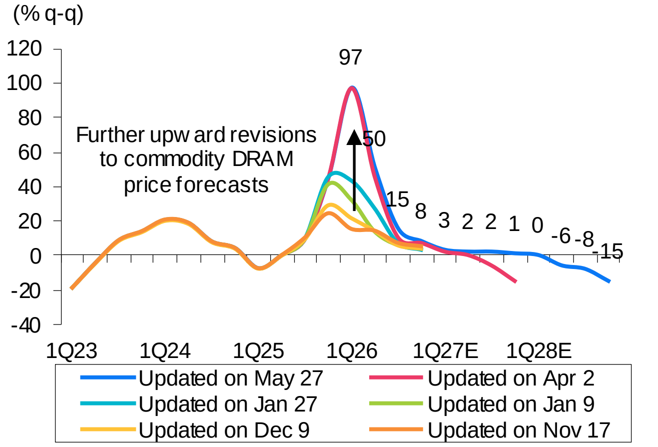

Note: *Average of SEC and Hynix
Source: Samsung Securities estimates

Annual DRAM ASP change forecasts*

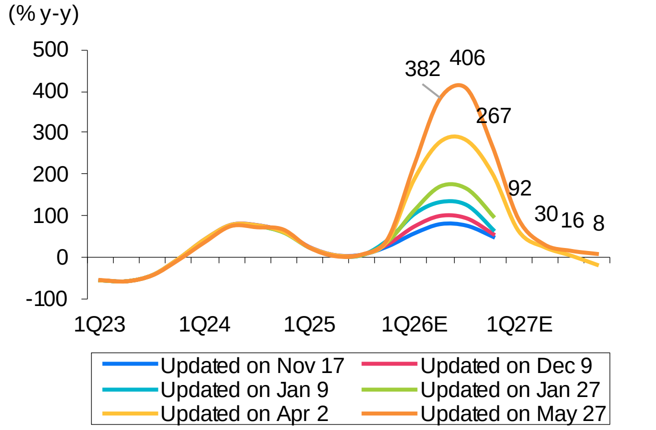

Note: *Average of SEC and Hynix
Source: Samsung Securities estimates

Quarterly DRAM ASP change

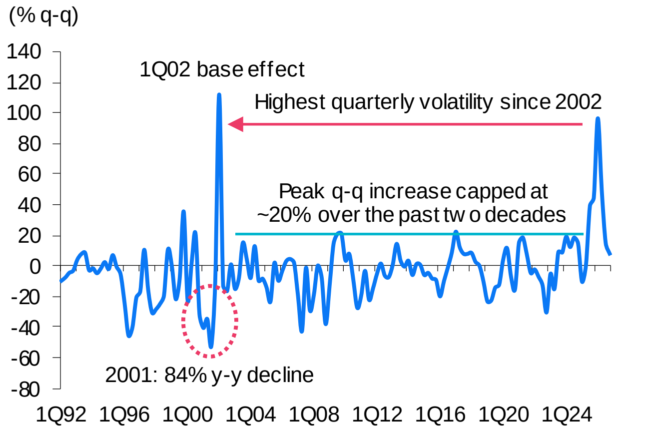

Source: WSTS, Samsung Securities estimates

Quarterly DRAM ASP change

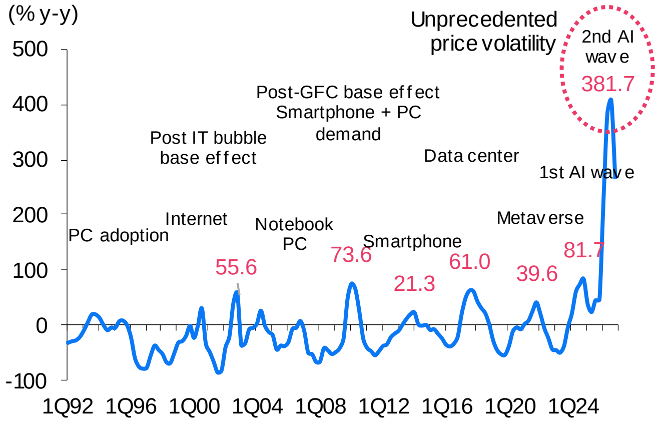

Source: WSTS, Samsung Securities estimates

Samsung Electronics: Results and forecasts

(KRWb)	1Q26	2Q26E	3Q26E	4Q26E	1Q27E	2Q27E	3Q27E	4Q27E	2026E	2027E	2028E
Sales	**133,870**	**178,355**	**213,038**	**236,249**	**253,981**	**271,705**	**290,172**	**300,689**	**761,512**	**1,116,547**	**1,249,660**
Chg (% y-y)	69.4	139.1	147.4	152.0	89.7	52.3	36.2	27.3	128.3	46.6	11.9
Semiconductor	81,698	125,806	156,155	179,998	194,919	211,648	223,910	235,101	543,077	865,578	955,876
DRAM	54,877	84,041	101,028	114,843	126,033	139,789	147,542	152,476	354,788	565,840	641,725
NAND	19,927	34,656	47,660	57,001	61,647	64,766	68,685	73,534	159,244	268,632	281,490
Foundry & LSI	6,894	7,110	7,467	8,154	7,239	7,092	7,684	9,092	29,625	31,106	32,661
Display	6,700	6,147	7,938	8,517	6,365	6,906	8,335	9,162	29,302	30,767	32,305
Telecom & Handset	38,136	34,118	36,545	33,824	39,911	39,976	44,532	41,187	142,622	165,606	202,156
CE and Harman	17,689	18,284	18,400	19,909	18,786	19,176	19,395	21,239	74,282	78,596	83,323
Operating profit	**57,239**	**81,281**	**101,981**	**117,792**	**126,965**	**137,635**	**144,051**	**151,089**	**358,293**	**559,740**	**570,165**
Semiconductor	53,709	79,420	100,177	115,838	126,799	138,665	145,796	150,280	349,144	561,540	587,288
DRAM	43,415	63,197	75,209	83,722	93,122	102,535	107,666	111,162	265,543	414,485	460,463
NAND	11,059	17,155	25,498	32,206	33,598	35,945	37,777	38,973	85,918	146,293	125,150
Foundry & LSI	-765	-931	-530	-90	80	184	353	145	-2,316	763	1,674
Display	382	301	968	1,576	299	407	767	1,466	3,227	2,939	2,618
Telecom & Handset	2,763	1,046	592	434	-84	-1,767	-2,655	-657	4,834	-5,162	-19,435
CE and Harman	385	514	244	-56	-50	330	143	0	1,087	422	-306
Operating margin (%)	**42.8**	**45.6**	**47.9**	**49.9**	**50.0**	**50.7**	**49.6**	**50.2**	**47.1**	**50.1**	**45.6**
Semiconductor	65.7	63.1	64.2	64.4	65.1	65.5	65.1	63.9	64.3	64.9	61.4
DRAM	79.1	75.2	74.4	72.9	73.9	73.4	73.0	72.9	74.8	73.3	71.8
NAND	55.5	49.5	53.5	56.5	54.1	54.5	0.0	0.0	54.0	54.5	51.0
Foundry & LSI	-11.1	-13.1	-7.1	-1.1	1.1	2.6	4.6	1.6	-7.8	2.5	5.1
Display	5.7	4.9	12.2	18.5	4.7	5.9	9.2	16.0	11.0	9.6	8.1
Telecom & Handset	7.2	3.1	1.6	1.3	-0.2	-4.4	-6.0	-1.6	3.4	-3.1	-9.6
CE and Harman	2.2	2.8	1.3	-0.3	-0.3	1.7	0.7	0.0	1.5	0.5	-0.4
Assumptions for bit growth and ASP changes											
DRAM bit growth (% q-q)	0.3	3.6	5.7	5.7	2.5	3.7	1.8	1.3	20.1	10.5	18.1
DRAM ASP chg (% q-q)	92.0	45.9	14.0	7.5	7.0	6.9	3.7	2.0	282.6	5.0	-4.0
NAND bit growth (% q-q)	9.2	4.0	6.0	4.0	3.0	3.0	5.0	6.0	18.9	15.0	17.4
NAND ASP chg (% q-q)	86.0	65.0	30.0	15.0	5.0	2.0	1.0	1.0	315.8	46.7	-10.7

Source: Company data, Samsung Securities estimates

SK Hynix: Results and forecasts

(KRWb)	1Q26	2Q26E	3Q26E	4Q26E	1Q27E	2Q27E	3Q27E	4Q27E	2026E	2027E	2028E
Sales	**52,576**	**83,862**	**103,190**	**121,056**	**134,593**	**147,191**	**157,246**	**166,107**	**360,684**	**605,136**	**680,486**
Chg (% y-y)	198.1	277.2	322.1	268.8	156.0	75.5	52.4	37.2	107.1	254.0	316.9
DRAM	40,669	61,392	73,263	85,303	96,265	106,936	114,571	121,294	260,626	439,067	507,546
NAND and other	11,907	22,470	29,927	35,753	38,328	40,254	42,675	44,813	100,058	166,069	172,940
Operating profit	**37,610**	**60,703**	**74,297**	**86,773**	**96,507**	**103,928**	**108,115**	**111,232**	**259,383**	**419,781**	**424,762**
DRAM	30,972	47,118	55,569	64,019	71,756	79,125	85,630	90,733	197,678	327,244	374,774
NAND and other	6,638	13,585	18,728	22,754	24,750	24,803	22,485	20,499	61,705	92,537	49,987
Operating margin (%)	**71.5**	**72.4**	**72.0**	**71.7**	**71.7**	**70.6**	**68.8**	**67.0**	**71.9**	**69.4**	**62.4**
DRAM	76.2	76.8	75.8	75.0	74.5	74.0	74.7	74.8	75.8	74.5	73.8
NAND and other	55.8	60.5	62.6	63.6	64.6	61.6	52.7	45.7	61.7	55.7	28.9
Assumptions for bit growth and ASP changes											
DRAM bit growth	0.0	8.8	4.6	4.3	1.0	1.0	2.8	4.1	20.3	12.2	20.6
DRAM ASP change	62.9	35.4	14.3	11.7	11.7	10.0	4.2	1.7	175.2	49.7	-16.8
NAND bit growth	-15.0	17.0	7.0	4.0	2.0	3.0	5.0	4.0	16.1	19.1	16.7
NAND ASP change	75.0	65.0	25.0	15.0	5.0	2.0	1.0	1.0	293.7	40.0	-10.8

Source: Company data, Samsung Securities estimates

COMPANY UPDATE

2026. 6. 12

Tech Team

Jongwook Lee
Team Leader
jwstar.lee@samsung.com

Kyoungbeen Kim
Research Associate
kyoungbeen.kim@samsung.com

Scan to go to
Research Center report database

Samsung Electronics (005930)

Memory still the bottleneck in AI agent era

- Confidence in the sustainability of long-term DRAM demand continues to grow, while chip supply growth should remain below chip demand growth. The value of foundry and smartphone segments is set to be reassessed.
- We raise our 12-month target price for Samsung Electronics to KRW500,000 (based on 2.2x 2028 P/B) and maintain our BUY rating.

AT A GLANCE

BUY

Target price	KRW500,000	67.2%
Current price	KRW299,500	
Market cap	KRW2,107.6t/USD1,389.9b	
Shares (float)	5,846,278,608 (78.9%)	
52-week high/low	KRW360,500/KRW56,800	
Avg daily trading value (60-day)	KRW6,819.9b/ USD4,497.4m	

ONE-YEAR PERFORMANCE

	1M	6M	12M
Samsung Electronics (%)	63.5	248.6	534.7
Vs Kospi (%pts)	22.6	58.2	94.6

KEY CHANGES

(KRW)	New	Old	Diff
Recommend.	BUY	BUY	
Target price	500,000	300,000	66.7%
2026E EPS	40,536	38,659	4.9%
2027E EPS	64,037	46,469	37.8%

SAMSUNG vs THE STREET

No of estimates	24
Target price	396,667
Recommendation	4.0

※ Rating: 4 < → BUY, 3 = HOLD, 2 > → SELL

WHAT'S THE STORY?

Investment strategy: Agentic AI is transforming the AI investment landscape. As services incorporating agentic AI expand, investment interest is spreading to adjacent connectivity components, and investors are gaining confidence in the durability of chipmakers' earnings growth during this industry upturn. Meanwhile, agentic AI should prompt a reassessment of previously overlooked value in the foundry and device segments. The narrowing valuation gap between Samsung Electronics (SEC) and its memory peers suggests significant upside remains for the stock. We maintain our BUY rating for SEC and present it as our sector top pick.

Raising target price to KRW500,000: We raise our 12-month target price for SEC to KRW500,000 (based on 2.2x 2028 P/B) and maintain BUY. We raise our 2026 and 2027 operating profit forecasts by 5% and 35.6%, respectively, to KRW358t and KRW559t. Demand for agentic AI is strengthening confidence in the duration of the memory upturn, while DRAM supply growth is expected to remain below demand growth through 2028. Despite trading at a discount to its peers, negative factors (*eg*, labor issues) are already reflected in SEC's valuation.

(Continued on the next page)

SUMMARY FINANCIAL DATA

	2025	2026E	2027E	2028E
Revenue (KRWb)	333,606	761,512	1,116,547	1,249,660
Operating profit (KRWb)	43,601	358,293	559,740	570,165
Net profit (adj) (KRWb)	45,207	276,168	434,861	453,344
EPS (adj) (KRW)	6,564	40,536	64,037	66,759
EPS (adj) growth (% y-y)	32.6	517.6	58.0	4.3
EBITDA margin (%)	27.1	54.0	55.5	50.8
ROE (%)	10.8	48.5	47.5	33.6
P/E (adj) (x)	18.3	8.9	5.6	5.4
P/B (x)	1.9	3.4	2.1	1.5
EV/EBITDA (x)	7.7	5.0	2.8	2.1
Dividend yield (%)	1.4	0.5	0.5	0.5

Source: Company data, Samsung Securities estimates

Samsung Securities

2Q earnings momentum continues: We expect SEC's 2Q operating profit to reach KRW81t (up 42% q-q). DRAM and NAND ASPs are set to rise 46% and 65% q-q, respectively, with mobile DRAM ASP increases likely contributing meaningfully to DRAM price rises. Earnings at its HBM and NAND units should surge from 3Q. Our 2026 estimates include bonus provisioning for the DS division, which figure we believe will amount to 13% of the division's 2026 operating profit.

Samsung Electronics: Target-price calculation

(KRW)		Notes
BVPS	234,152	2028E
Target P/B (x)	2.2	Assumes 2028 ROE of 33%
Fair price per share	515,135	
Target price	**500,000**	
Current price	299,000	
Upside (%)	67.2	

Source: Samsung Securities estimates

Samsung Electronics: Forward P/B vs ROE

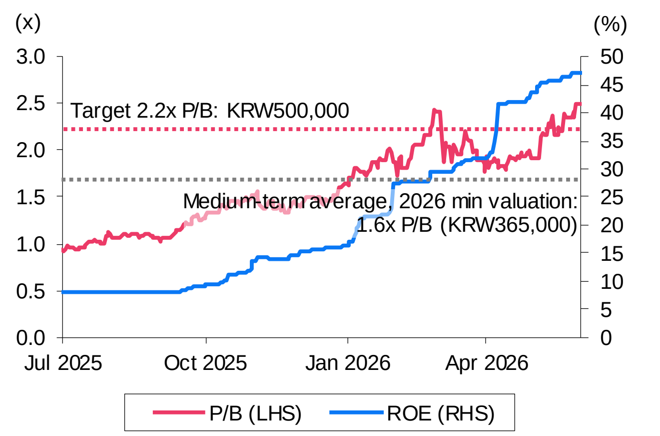

Source: QuantiWise, Samsung Securities

Samsung Electronics: Forward P/B vs ROE

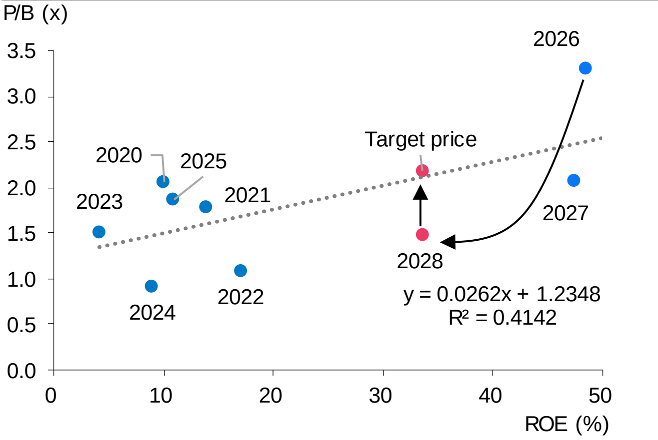

Source: QuantiWise, Samsung Securities

Samsung Electronics: Revisions to full-year forecasts

(KRWb)	2026			2027		
	Old	New	Diff (%)	Old	New	Diff (%)
Sales	**717,866**	**761,512**	**6.1**	**913,631**	**1,116,547**	**22.2**
Semiconductor	504,717	543,077	7.6	668,506	865,578	29.5
DRAM	319,506	354,788	11.0	425,172	565,840	33.1
NAND	154,392	159,244	3.1	212,229	268,632	26.6
LSI	29,625	29,625	0.0	31,106	31,106	0.0
Display	29,302	29,302	0.0	30,767	30,767	0.0
Telecom	139,110	142,622	2.5	159,761	165,606	3.7
Handset	135,910	139,422	2.6	156,511	162,356	3.7
CE and Harman	74,282	74,282	0.0	78,596	78,596	0.0
Operating profit	**341,295**	**358,293**	**5.0**	**412,922**	**559,740**	**35.6**
Semiconductor	332,523	349,144	5.0	414,900	561,540	35.3
DRAM	250,833	265,543	5.9	328,418	414,485	26.2
NAND	84,006	85,918	2.3	85,719	146,293	70.7
LSI	-2,316	-2,316	0.0	763	763	0.0
Display	3,227	3,227	0.0	2,939	2,939	0.0
Telecom	4,459	4,834	8.4	-5,340	-5,162	-3.3
Handset	4,217	4,584	8.7	-5,596	-5,430	-3.0
CE and Harman	1,087	1,087	0.0	422	422	0.0

Source: Company data, Samsung Securities estimates

Samsung Electronics: Results and forecasts

(KRWb)	1Q26	2Q26E	3Q26E	4Q26E	1Q27E	2Q27E	3Q27E	4Q27E	2026E	2027E	2028E
Sales	**133,870**	**178,355**	**213,038**	**236,249**	**253,981**	**271,705**	**290,172**	**300,689**	**761,512**	**1,116,547**	**1,249,660**
Chg (% y-y)	69.4	139.1	147.4	152.0	89.7	52.3	36.2	27.3	128.3	46.6	11.9
Semiconductor	81,698	125,806	156,155	179,998	194,919	211,648	223,910	235,101	543,077	865,578	955,876
DRAM	54,877	84,041	101,028	114,843	126,033	139,789	147,542	152,476	354,788	565,840	641,725
NAND	19,927	34,656	47,660	57,001	61,647	64,766	68,685	73,534	159,244	268,632	281,490
Foundry & LSI	6,894	7,110	7,467	8,154	7,239	7,092	7,684	9,092	29,625	31,106	32,661
Display	6,700	6,147	7,938	8,517	6,365	6,906	8,335	9,162	29,302	30,767	32,305
Telecom & Handset	38,136	34,118	36,545	33,824	39,911	39,976	44,532	41,187	142,622	165,606	202,156
CE and Harman	17,689	18,284	18,400	19,909	18,786	19,176	19,395	21,239	74,282	78,596	83,323
Operating profit	**57,239**	**81,281**	**101,981**	**117,792**	**126,965**	**137,635**	**144,051**	**151,089**	**358,293**	**559,740**	**570,165**
Semiconductor	53,709	79,420	100,177	115,838	126,799	138,665	145,796	150,280	349,144	561,540	587,288
DRAM	43,415	63,197	75,209	83,722	93,122	102,535	107,666	111,162	265,543	414,485	460,463
NAND	11,059	17,155	25,498	32,206	33,598	35,945	37,777	38,973	85,918	146,293	125,150
Foundry & LSI	-765	-931	-530	-90	80	184	353	145	-2,316	763	1,674
Display	382	301	968	1,576	299	407	767	1,466	3,227	2,939	2,618
Telecom & Handset	2,763	1,046	592	434	-84	-1,767	-2,655	-657	4,834	-5,162	-19,435
CE and Harman	385	514	244	-56	-50	330	143	0	1,087	422	-306
Operating margin (%)	**42.8**	**45.6**	**47.9**	**49.9**	**50.0**	**50.7**	**49.6**	**50.2**	**47.1**	**50.1**	**45.6**
Semiconductor	65.7	63.1	64.2	64.4	65.1	65.5	65.1	63.9	64.3	64.9	61.4
DRAM	79.1	75.2	74.4	72.9	73.9	73.4	73.0	72.9	74.8	73.3	71.8
NAND	55.5	49.5	53.5	56.5	54.1	54.5	0.0	0.0	54.0	54.5	51.0
Foundry & LSI	-11.1	-13.1	-7.1	-1.1	1.1	2.6	4.6	1.6	-7.8	2.5	5.1
Display	5.7	4.9	12.2	18.5	4.7	5.9	9.2	16.0	11.0	9.6	8.1
Telecom & Handset	7.2	3.1	1.6	1.3	-0.2	-4.4	-6.0	-1.6	3.4	-3.1	-9.6
CE and Harman	2.2	2.8	1.3	-0.3	-0.3	1.7	0.7	0.0	1.5	0.5	-0.4
Assumptions for bit growth and ASP changes											
DRAM bit growth (% q-q)	0.3	3.6	5.7	5.7	2.5	3.7	1.8	1.3	20.1	10.5	10.5
DRAM ASP chg (% q-q)	92.0	45.9	14.0	7.5	7.0	6.9	3.7	2.0	282.6	5.0	5.0
NAND bit growth (% q-q)	9.2	4.0	6.0	4.0	3.0	3.0	5.0	6.0	18.9	15.0	15.0
NAND ASP chg (% q-q)	86.0	65.0	30.0	15.0	5.0	2.0	1.0	1.0	315.8	46.7	-8.9

Source: Company data, Samsung Securities estimates

Income statement

Year-end Dec 31 (KRWb)	2024	2025	2026E	2027E	2028E
Sales	**300,871**	**333,606**	**761,512**	**1,116,547**	**1,249,660**
Cost of goods sold	186,562	202,236	193,237	249,760	334,241
Gross profit	**114,309**	**131,370**	**568,275**	**866,787**	**915,419**
Gross margin (%)	38.0	39.4	74.6	77.6	73.3
SG&A expenses	81,583	87,769	209,982	307,048	345,254
Operating profit	**32,726**	**43,601**	**358,293**	**559,740**	**570,165**
Operating margin (%)	10.9	13.1	47.1	50.1	45.6
Non-operating gains (losses)	**4,804**	**5,880**	**9,931**	**20,075**	**34,294**
Financial profit	16,703	16,240	19,841	29,464	44,136
Financial costs	12,986	11,734	11,285	10,856	9,327
Equity-method gains (losses)	751	683	1,000	1,000	1,000
Other	335	691	374	467	-1,516
Pre-tax profit	**37,530**	**49,481**	**368,224**	**579,814**	**604,459**
Taxes	3,078	4,275	92,056	144,954	151,115
Effective tax rate (%)	8.2	8.6	25.0	25.0	25.0
Profit from continuing operations	34,451	45,207	276,168	434,861	453,344
Profit from discontinued operations	0	0	0	0	0
Net profit	**34,451**	**45,207**	**276,168**	**434,861**	**453,344**
Net margin (%)	11.5	13.6	36.3	38.9	36.3
Net profit (controlling interests)	33,621	44,261	270,389	425,762	443,859
Net profit (non-controlling interests)	830	946	5,778	9,098	9,485
EBITDA	75,357	90,528	411,250	619,439	635,061
EBITDA margin (%)	25.0	27.1	54.0	55.5	50.8
EPS (parent-based) (KRW)	4,950	6,564	40,536	64,037	66,759
EPS (consolidated) (KRW)	5,072	6,704	41,402	65,406	68,186
Adjusted EPS (KRW)*	4,950	6,564	40,536	64,037	66,759

Balance sheet

Year-end Dec 31 (KRWb)	2024	2025	2026E	2027E	2028E
Current assets	**227,062**	**247,685**	**563,160**	**1,019,785**	**1,461,673**
Cash & equivalents	53,706	57,856	160,462	470,881	846,044
Accounts receivable	43,623	51,128	116,707	171,119	191,520
Inventories	51,755	52,637	160,877	197,671	223,507
Other current assets	77,979	86,064	125,114	180,114	200,603
Fixed assets	**287,470**	**319,257**	**357,654**	**396,769**	**422,346**
Investment assets	24,349	31,348	49,701	66,285	74,528
Tangible assets	205,945	215,305	236,467	259,945	278,081
Intangible assets	23,739	29,481	28,362	27,414	26,612
Other long-term assets	33,437	43,125	43,125	43,125	43,125
Total assets	**514,532**	**566,942**	**920,814**	**1,416,554**	**1,884,019**
Current liabilities	**93,326**	**106,411**	**179,675**	**238,916**	**259,271**
Accounts payable	12,370	13,039	18,614	21,507	22,362
Short-term debt	13,173	17,575	15,575	13,575	11,575
Other current liabilities	67,784	75,797	145,485	203,834	225,334
Long-term liabilities	**19,014**	**24,210**	**33,492**	**46,083**	**50,804**
Bonds & long-term debt	21	2,813	2,813	2,813	2,813
Other long-term liabilities	18,993	21,397	30,679	43,270	47,991
Total liabilities	**112,340**	**130,622**	**213,167**	**284,999**	**310,074**
Owners of parent equity	**391,688**	**424,313**	**689,862**	**1,104,671**	**1,537,576**
Capital stock	898	898	898	898	898
Capital surplus	4,404	4,404	4,404	4,404	4,404
Retained earnings	370,513	402,136	661,535	1,076,344	1,509,249
Other	15,873	16,876	23,026	23,026	23,026
Non-controlling interests' equity	**10,504**	**12,007**	**17,785**	**26,884**	**36,369**
Total equity	**402,192**	**436,320**	**707,647**	**1,131,554**	**1,573,945**
Net debt	-93,285	-100,582	-235,424	-593,078	-987,201

Cash flow statement

Year-end Dec 31 (KRWb)	2024	2025	2026E	2027E	2028E
Cash flow from operations	**72,983**	**85,315**	**229,174**	**465,245**	**492,950**
Net profit	34,451	45,207	276,168	434,861	453,344
Non-cash profit and expenses	42,947	52,396	136,158	186,028	182,508
Depreciation	39,650	43,606	49,838	56,752	62,094
Amortization	2,981	3,321	3,119	2,947	2,802
Other	316	5,469	83,201	126,328	117,612
Changes in A/L from operating activities	-1,568	-9,614	-98,088	-27,138	-22,691
Cash flow from investments	**-85,382**	**-68,512**	**-101,209**	**-126,297**	**-97,512**
Change in tangible assets	-51,250	-47,372	-71,000	-80,230	-80,230
Change in financial assets	-36,218	-9,056	-29,059	-45,235	-16,960
Other	2,087	-12,084	-1,150	-832	-322
Cash flow from financing	**-7,797**	**-13,478**	**-14,167**	**-12,954**	**-12,954**
Change in debt	6,644	5,909	-3,178	-2,000	-2,000
Change in equity	0	0	0	0	0
Dividends	-10,889	-9,897	-10,990	-10,954	-10,954
Other	-3,553	-9,490	0	0	0
Change in cash	-15,375	4,151	102,605	310,419	375,163
Cash at beginning of year	69,081	53,706	57,856	160,462	470,881
Cash at end of year	53,706	57,856	160,462	470,881	846,044
Gross cash flow	**77,398**	**97,602**	**412,325**	**620,889**	**635,852**
Free cash flow	**21,576**	**37,793**	**158,174**	**385,015**	**412,720**

Note: *Excluding one-off items
**Fully diluted, excluding one-off items
***From companies subject to equity-method valuation
Source: Company data, Samsung Securities estimates

Financial ratios

Year-end Dec 31	2024	2025	2026E	2027E	2028E
Growth (%)					
Sales	16.2	10.9	128.3	46.6	11.9
Operating profit	398.3	33.2	721.8	56.2	1.9
Net profit	122.5	31.2	510.9	57.5	4.3
Adjusted EPS**	132.3	32.6	517.6	58.0	4.3
Per-share data (KRW)					
EPS (parent-based)	4,950	6,564	40,536	64,037	66,759
EPS (consolidated)	5,072	6,704	41,402	65,406	68,186
Adjusted EPS**	4,950	6,564	40,536	64,037	66,759
BVPS	57,981	63,997	105,057	168,227	234,152
DPS (common)	1,446	1,668	1,668	1,668	1,668
Valuations (x)					
P/E***	10.7	18.3	8.9	5.6	5.4
P/B***	0.9	1.9	3.4	2.1	1.5
EV/EBITDA	3.6	7.7	5.0	2.8	2.1
Ratios (%)					
ROE	9.0	10.8	48.5	47.5	33.6
ROA	7.1	8.4	37.1	37.2	27.5
ROIC	10.6	13.2	71.4	88.3	80.3
Payout ratio	25.6	22.1	3.6	2.3	2.2
Dividend yield (common)	2.7	1.4	0.5	0.5	0.5
Net debt to equity	-23.2	-23.1	-33.3	-52.4	-62.7
Interest coverage (x)	36.2	72.0	557.3	977.6	1,100.3

COMPANY UPDATE

2026. 6. 12

Tech Team

Jongwook Lee
Team Leader
jwstar.lee@samsung.com

Kyoungbeen Kim
Research Associate
kyoungbeen.kim@samsung.com

▶ AT A GLANCE

BUY

Target price	**KRW3,500,000 66.6%**
Current price	**KRW2,101,000**
Market cap	KRW1,682.0t/USD1,109.2b
Shares (float)	712,702,365 (76.2%)
52-week high/low	KRW2,363,000/KRW207,500
Avg daily trading value (60-day)	KRW6,804.8b/ USD4,487.4m

▶ ONE-YEAR PERFORMANCE

	1M	6M	12M
SK Hynix (%)	83.5	322.9	1,037.3
Vs Kospi (%pts)	37.6	92.0	248.8

▶ KEY CHANGES

(KRW)	New	Old	Diff
Recommend.	BUY	BUY	
Target price	3,500,000	1,800,000	94.4%
2026E EPS	284,437	250,076	13.7%
2027E EPS	464,786	315,598	47.3%

▶ SAMSUNG vs THE STREET

No of estimates	24
Target price	2,470,417
Recommendation	4.0

※ Rating: 4 < → BUY, 3 = HOLD, 2 > → SELL

Scan to go to
Research Center report database

SK Hynix (000660)

Looking for memory? Look no further!

- SK Hynix's premium memory portfolio is ideally positioned as a proxy for agentic AI growth. Its potential for a valuation premium is rising.
- We raise our target price to KRW3,500,000 (based on 2.5x 2028 P/B), and reiterate our BUY rating on the stock.

WHAT'S THE STORY?

Investment view: Memory is increasingly recognized as the fastest-growing component in the AI agent era. As AI agent adoption accelerates and cloud hyperscalers intensify spending, capital is increasingly flowing to memory manufacturers to capture the upside. As a leading pure-play memory manufacturer, SK Hynix offers a premium portfolio spanning high-bandwidth memory (HBM), server DRAM, and NAND—making it one of the most compelling global proxies for AI agent growth.

Raising target price: We raise our target price to KRW3,500,000, based on a target P/B multiple of 2.5x applied to our 2028 BVPS forecast. We revise up our 2026 and 2027 operating profit forecasts by 13.7% and 47.6%, respectively, to KRW259t and KRW420t. Demand driven by AI Agents is increasingly proving credible in terms of duration, while DRAM supply growth is unlikely to keep pace through 2028. Rapid BVPS growth—fueled by a high ROE—is gradually justifying the company's premium valuation. We believe conditions are now in place to support our target P/B multiple of 2.5x, which implies a 2028 ROE of 42%.

(Continued on the next page)

SUMMARY FINANCIAL DATA

(KRWb)	2025	2026E	2027E	2028E
Revenue	97,147	360,684	605,136	680,486
Operating profit	47,206	259,383	419,781	424,762
Net profit (adj)	42,948	203,319	331,475	339,237
EPS (adj) (KRW)	58,955	284,437	464,786	475,670
EPS (adj) growth (% y-y)	116.9	382.5	63.4	2.3
EBITDA margin (%)	62.9	77.2	73.5	66.7
ROE (%)	44.2	91.9	68.4	41.6
P/E (adj) (x)	11.0	8.3	5.1	5.0
P/B (x)	3.7	5.1	2.5	1.7
EV/EBITDA (x)	7.8	5.7	3.0	2.3
Dividend yield (%)	0.5	0.3	0.3	0.3

Source: Company data, Samsung Securities estimates

Demonstrating commitment to shareholder returns: Hynix's ADR listing is imminent and should further reinforce the firm's shareholder-friendly profile. Management has indicated it will finalize its shareholder return policy—including dividends and share buybacks—by year-end, with expectations for further enhancement as the year progresses. More proactive engagement with investors is becoming a solid foundation for a sustained valuation premium.

SK Hynix: Target-price calculation

(KRW)		Notes
BVPS	1,419,571	2028E
Target P/B (x)	2.5	Assumes 2028 ROE of 42%
Fair price per share	3,548,928	
Target price	**3,500,000**	
Current price	2,101,000	
Upside (%)	66.6	

Source: Samsung Securities estimates

SK Hynix: Forward P/B vs ROE

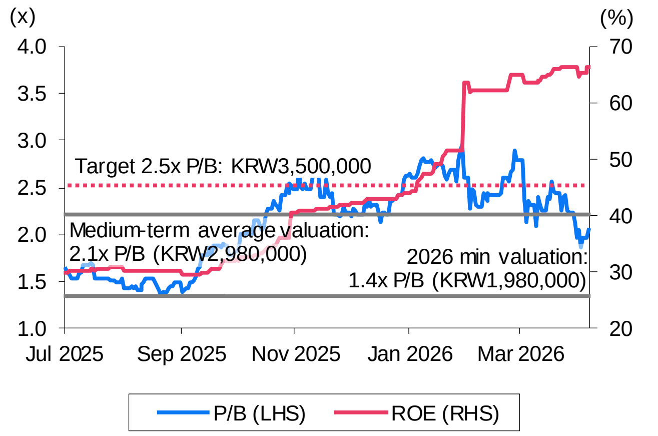

Source: QuantiWise, Samsung Securities

SK Hynix: Forward P/B vs ROE

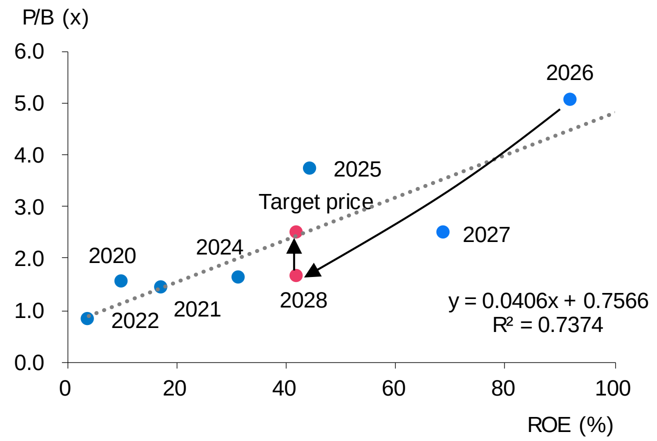

Source: QuantiWise, Samsung Securities

SK Hynix: Revisions to full-year earnings forecasts

(KRWb)	Old		New		Diff (%)	
	2026E	2027E	2026E	2027E	2026E	2027E
Sales	323,895	431,384	360,684	605,136	11.5	40.3
Operating profit	228,032	284,320	259,383	419,781	13.7	47.6
Pre-tax profit	229,176	288,561	260,666	424,968	13.7	47.3
Net profit	178,758	225,077	203,319	331,475	13.7	47.3

Source: Samsung Securities estimates

SK Hynix: Results and forecasts

(KRWb)	1Q26	2Q26E	3Q26E	4Q26E	1Q27E	2Q27E	3Q27E	4Q27E	2026E	2027E	2028E
Sales	**52,576**	**83,862**	**103,190**	**121,056**	**134,593**	**147,191**	**157,246**	**166,107**	**360,684**	**605,136**	**680,486**
Chg (% y-y)	198.1	277.2	322.1	268.8	156.0	75.5	52.4	37.2	107.1	254.0	316.9
DRAM	40,669	61,392	73,263	85,303	96,265	106,936	114,571	121,294	260,626	439,067	507,546
NAND and other	11,907	22,470	29,927	35,753	38,328	40,254	42,675	44,813	100,058	166,069	172,940
Operating profit	**37,610**	**60,703**	**74,297**	**86,773**	**96,507**	**103,928**	**108,115**	**111,232**	**259,383**	**419,781**	**424,762**
DRAM	30,972	47,118	55,569	64,019	71,756	79,125	85,630	90,733	197,678	327,244	374,774
NAND and other	6,638	13,585	18,728	22,754	24,750	24,803	22,485	20,499	61,705	92,537	49,987
OPM (%)	**71.5**	**72.4**	**72.0**	**71.7**	**71.7**	**70.6**	**68.8**	**67.0**	**71.9**	**69.4**	**62.4**
DRAM	76.2	76.8	75.8	75.0	74.5	74.0	74.7	74.8	75.8	74.5	73.8
NAND and other	55.8	60.5	62.6	63.6	64.6	61.6	52.7	45.7	61.7	55.7	28.9
Assumptions for bit growth and ASP changes											
DRAM bit growth	0.0	8.8	4.6	4.3	1.0	1.0	2.8	4.1	20.3	12.2	20.6
DRAM ASP change	62.9	35.4	14.3	11.7	11.7	10.0	4.2	1.7	175.2	49.7	-16.8
NAND bit growth	-15.0	17.0	7.0	4.0	2.0	3.0	5.0	4.0	16.1	19.1	16.7
NAND ASP change	75.0	65.0	25.0	15.0	5.0	2.0	1.0	1.0	293.7	40.0	-10.8

Source: Company data, Samsung Securities estimates

Income statement

Year-end Dec 31 (KRWb)	2024	2025	2026E	2027E	2028E
Sales	**66,193**	**97,147**	**360,684**	**605,136**	**680,486**
Cost of goods sold	34,365	38,456	87,768	171,504	241,901
Gross profit	**31,828**	**58,691**	**273,520**	**433,633**	**438,585**
Gross margin (%)	48.1	60.4	75.7	71.7	64.5
SG&A expenses	8,361	11,484	14,137	13,852	13,824
Operating profit	**23,467**	**47,206**	**259,383**	**419,781**	**424,762**
Operating margin (%)	35.5	48.6	71.8	69.4	62.4
Non-operating gains (losses)	**418**	**3,259**	**1,283**	**5,187**	**10,157**
Financial profit	4,855	16,373	9,173	14,353	21,294
Financial costs	5,708	12,505	7,125	8,166	9,350
Equity-method gains (losses)	-38	-565	50	50	50
Other	1,309	-45	-816	-1,050	-1,837
Pre-tax profit	**23,885**	**50,466**	**260,666**	**424,968**	**434,919**
Taxes	4,088	7,518	57,347	93,493	95,682
Effective tax rate (%)	17.1	14.9	22.0	22.0	22.0
Profit from continuing operations	19,797	42,948	203,319	331,475	339,237
Profit from discontinued operations	0	0	0	0	0
Net profit	**19,797**	**42,948**	**203,319**	**331,475**	**339,237**
Net margin (%)	29.9	44.2	56.3	54.8	49.9
Net profit (controlling interests)	19,789	42,919	203,184	331,254	339,011
Net profit (non-controlling interests)	8	29	135	221	226
EBITDA	36,049	61,136	278,948	444,788	454,151
EBITDA margin (%)	54.5	62.9	77.2	73.5	66.7
EPS (parent-based) (KRW)	27,182	58,955	284,437	464,786	475,670
EPS (consolidated) (KRW)	27,193	58,994	284,627	465,096	475,987
Adjusted EPS (KRW)*	27,182	58,955	284,437	464,786	475,670

Balance sheet

Year-end Dec 31 (KRWb)	2024	2025	2026E	2027E	2028E
Current assets	**42,279**	**69,458**	**242,115**	**552,771**	**872,521**
Cash & equivalents	11,205	14,924	73,300	273,619	559,276
Accounts receivable	13,019	18,199	50,014	83,770	94,201
Inventories	13,314	14,289	51,328	85,972	96,677
Other current assets	4,741	22,046	67,472	109,409	122,367
Fixed assets	**77,576**	**106,650**	**137,061**	**166,150**	**186,304**
Investment assets	6,522	17,172	24,308	30,896	32,932
Tangible assets	60,157	77,503	101,041	124,150	142,752
Intangible assets	4,019	4,049	3,786	3,178	2,695
Other long-term assets	6,878	7,926	7,926	7,926	7,926
Total assets	**119,855**	**176,108**	**379,176**	**718,921**	**1,058,826**
Current liabilities	**24,965**	**37,379**	**34,890**	**42,406**	**47,278**
Accounts payable	2,277	2,848	7,542	12,632	14,205
Short-term debt	1,283	2,396	2,250	2,104	1,958
Other current liabilities	21,405	32,135	25,099	27,670	31,115
Long-term liabilities	**20,974**	**18,062**	**22,395**	**28,676**	**29,999**
Bonds & long-term debt	17,431	14,086	14,086	16,086	16,086
Other long-term liabilities	3,543	3,976	8,309	12,590	13,913
Total liabilities	**45,940**	**55,441**	**57,285**	**71,082**	**77,278**
Owners of parent equity	**73,903**	**120,516**	**321,605**	**647,332**	**980,815**
Capital stock	3,658	3,658	3,658	3,658	3,658
Capital surplus	4,487	8,954	8,954	8,954	8,954
Retained earnings	65,418	106,577	307,665	633,392	966,876
Other	341	1,328	1,328	1,328	1,328
Non-controlling interests' equity	**12**	**151**	**286**	**507**	**733**
Total equity	**73,916**	**120,667**	**321,891**	**647,839**	**981,548**
Net debt	13,554	-242	-105,022	-337,250	-633,485

Cash flow statement

Year-end Dec 31 (KRWb)	2024	2025	2026E	2027E	2028E
Cash flow from operations	**29,796**	**53,373**	**156,492**	**291,318**	**351,143**
Net profit	19,797	42,948	203,319	331,475	339,237
Non-cash profit and expenses	17,054	18,838	74,768	112,817	114,520
Depreciation	11,985	13,099	18,802	24,400	28,906
Amortization	596	831	763	608	484
Other	4,472	4,907	55,203	87,809	85,130
Changes in A/L from operating activities	-5,600	-2,881	-66,032	-64,631	-17,320
Cash flow from investments	**-18,005**	**-48,054**	**-87,316**	**-88,529**	**-60,391**
Change in tangible assets	-15,898	-27,374	-42,340	-47,508	-47,508
Change in financial assets	-2,073	-13,315	-44,310	-40,905	-12,640
Other	-33	-7,365	-666	-115	-243
Cash flow from financing	**-8,704**	**-1,445**	**-7,734**	**197**	**-4,477**
Change in debt	-6,781	2,475	-5,639	5,724	1,050
Change in equity	115	4,467	0	0	0
Dividends	-826	-1,681	-2,095	-5,527	-5,527
Other	-1,211	-6,705	0	0	0
Change in cash	3,618	3,719	58,377	200,319	285,656
Cash at beginning of year	7,587	11,205	14,924	73,300	273,619
Cash at end of year	11,205	14,924	73,300	273,619	559,276
Gross cash flow	**36,851**	**61,785**	**278,087**	**444,291**	**453,757**
Free cash flow	**13,850**	**25,854**	**114,152**	**243,810**	**303,635**

Note: *Excluding one-off items;
**Fully diluted, excluding one-off items;
***From companies subject to equity-method valuation
Source: Company data, Samsung Securities estimates

Financial ratios

Year-end Dec 31	2024	2025	2026E	2027E	2028E
Growth (%)					
Sales	102.0	46.8	271.9	67.5	12.5
Operating profit	nm	101.2	449.5	61.8	1.2
Net profit	nm	116.9	373.4	63.0	2.3
Adjusted EPS**	nm	116.9	382.5	63.4	2.3
Per-share data (KRW)					
EPS (parent-based)	27,182	58,955	284,437	464,786	475,670
EPS (consolidated)	27,193	58,994	284,627	465,096	475,987
Adjusted EPS**	27,182	58,955	284,437	464,786	475,670
BVPS	107,256	174,539	465,471	936,908	1,419,571
DPS (common)	2,204	3,000	8,000	8,000	8,000
Valuations (x)					
P/E***	6.4	11.0	8.3	5.1	5.0
P/B***	1.6	3.7	5.1	2.5	1.7
EV/EBITDA	3.9	7.8	5.7	3.0	2.3
Ratios (%)					
ROE	31.1	44.2	91.9	68.4	41.6
ROA	18.0	29.0	73.2	60.4	38.2
ROIC	25.4	45.7	141.1	138.6	110.3
Payout ratio	7.7	4.9	2.7	1.7	1.6
Dividend yield (common)	1.3	0.5	0.3	0.3	0.3
Net debt to equity	18.3	-0.2	-32.6	-52.1	-64.5
Interest coverage (x)	17.4	51.1	297.4	480.5	431.8

Compliance notice

- As of 6/11 2026, Samsung Securities shared group affiliation with SamsungElectronics.
- As of 6/11 2026, the covering analyst(s) did not own any shares, or debt instruments convertible into shares, of any company covered in this report.
- As of 6/11 2026, Samsung Securities' holdings of shares and debt instruments convertible into shares of each company covered in this report would not, if such debt instruments were converted, exceed 1% of each company's outstanding shares.
- This report has been prepared without any undue external influence or interference, and accurately reflects the views of the analyst(s) covering the company or companies herein.
- All material presented in this report, unless specifically indicated otherwise, is under copyright to Samsung Securities.
- Neither the material nor its content (including copies) may be altered in any form, or by any means transmitted, copied, or distributed to another party, without prior express written permission from Samsung Securities.
- This memorandum is based upon information available to the public. While we have taken all reasonable care to ensure its reliability, we do not guarantee its accuracy or completeness. This memorandum is not intended to be an offer, or a solicitation of any offer, to buy or sell the securities mentioned herein. Samsung Securities shall not be liable whatsoever for any loss, direct or consequential, arising from the use of this memorandum or its contents. Statements made regarding affiliates of Samsung Securities are also based upon publicly available information and do not necessarily represent the views of management at such affiliates.
- This material has not been distributed to institutional investors or other third parties prior to its publication.

Target price changes in past two years

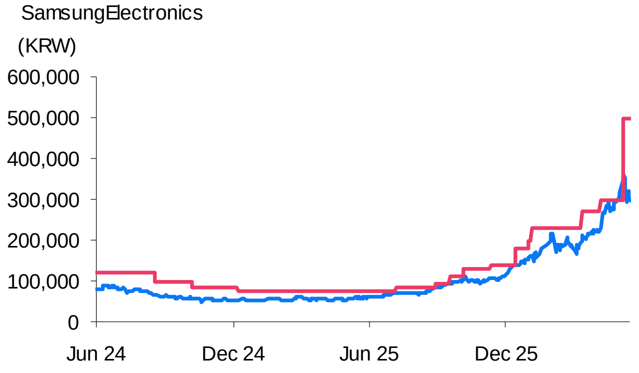

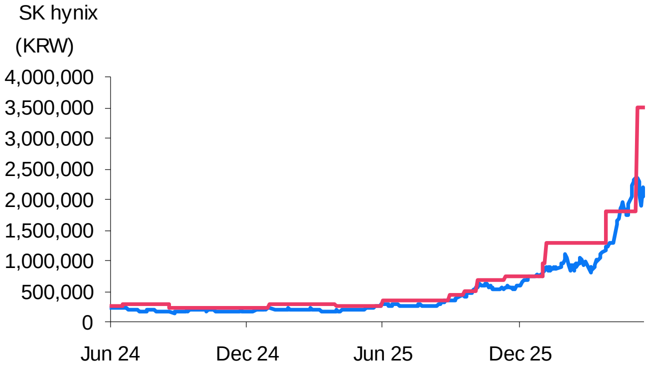

Rating changes over past two years (adjusted share prices)

SamsungElectronics

Date	2024/5/28	9/12	11/1	2025/1/2	8/1	9/23	10/13	10/31	12/8	2026/1/9	1/27	1/30
Recommendation	BUY	BUY	BUY	BUY	BUY	BUY	BUY	BUY	BUY	BUY	BUY	BUY
Target price (KRW)	120000	100000	83000	74000	85000	93000	110000	130000	140000	180000	200000	230000
Gap* (average)	-34.45	-39.49	-33.92	-21.93	-15.05	-6.95	-10.64	-21.96	-16.54	-18.69	-19.23	-19.75
(max or min)**	-26.83	-35.30	-29.28	-1.89	-1.76	1.51	-5.36	-14.54	0.71	-15.39	-18.80	-5.22
Date	4/7	5/4	6/2									
Recommendation	BUY	BUY	BUY									
Target price (KRW)	270000	300000	500000									
Gap* (average)	-20.25	-3.39										
(max or min)**	-16.30	16.33										

SK hynix

Date	2024/4/8	7/12	9/12	2025/1/24	4/24	6/25	9/23	10/13	10/30	12/8	2026/1/27	1/30
Recommendation	BUY	BUY	BUY	BUY	BUY	BUY	BUY	BUY	BUY	BUY	BUY	BUY
Target price (KRW)	250000	280000	240000	280000	250000	340000	430000	500000	700000	750000	950000	1300000
Gap* (average)	-18.91	-34.15	-23.78	-30.26	-14.53	-17.52	-14.81	-3.33	-19.05	-11.48	-10.42	-26.68
(max or min)**	-3.60	-16.79	-6.04	-21.96	11.40	3.82	-0.47	11.60	-11.43	2.27	-9.37	-10.31
Date	4/21	6/2										
Recommendation	BUY	BUY										
Target price (KRW)	1800000	3500000										
Gap* (average)	-3.42											
(max or min)**	31.28											

Note: * [(average, maximum, or minimum share price over duration of target price minus target price) / target price] x 100%
** Maximum/minimum share price if new target is higher/lower than market close on the business day prior to target price change

Samsung Securities uses the following investment ratings*

Company

BUY Expected to increase in value by 15% or more within 12 months and is highly attractive within sector

HOLD Expected to increase/decrease in value by less than 15% within 12 months

SELL Expected to decrease in value by 15% or more within 12 months

Industry

OVERWEIGHT Expected to outperform market by 5% or more within 12 months

NEUTRAL Expected to outperform/underperform market by less than 5% within 12 months

UNDERWEIGHT Expected to underperform market by 5% or more within 12 months

* Note: Effective Jul 27, 2023, BUY, HOLD, and SELL criteria are based on expectations of share-price moves of 15% or more within 12 months

Percentage of ratings in 12 months prior to 2026.03.31

BUY(85.2%)·HOLD(14.8%)·SELL(0%)

Global Disclosures & Disclaimers

General

This research report is for information purposes only. It is not and should not be construed as an offer or solicitation of an offer to purchase or sell any securities or other financial instruments or to participate in any trading strategy. This report does not provide individually tailored investment advice. This report does not take into account individual client circumstances, objectives, or needs and is not intended as recommendations of particular securities, financial instruments or strategies to any particular client. The securities and other financial instruments discussed in this report may not be suitable for all investors. The recipient of this report must make its own independent decisions regarding any securities or financial instruments mentioned herein and investors should seek the advice of a financial adviser.

This report may not be altered, reproduced, distributed, transmitted or published in whole or in part for any purpose. References to "Samsung Securities" are references to any company in the Samsung Securities, Co., Ltd. group of companies.

Samsung Securities and/or other affiliated companies, its and their directors, officers, representatives, or employees may have long or short positions in any of the securities or other financial instruments mentioned in this report or of issuers described herein and may purchase and/or sell, or offer to purchase and/or sell, at any time, such securities or other financial instruments in the open market or otherwise, as either a principal or agent. Any pricing of securities or other financial instrument contained herein is as of the close of market for such day, unless otherwise stated. Opinions and estimates contained herein constitute our judgment as of the date of this report and are subject to change without notice.

The information provided in this report is provided "AS IS". Although the information contained herein has been obtained from sources believed to be reliable, no representation or warranty, either expressed or implied, is provided by Samsung Securities in relation to the accuracy, completeness or reliability of such information or that such information was provided for any particular purpose and Samsung Securities expressly disclaims any warranties of merchantability or fitness for a particular purpose. Furthermore, this report is not intended to be a complete statement or summary of the securities, markets or developments referred to herein.

Samsung Securities does not undertake that investors will obtain any profits, nor will it share with investors any investment profits. Samsung Securities, its affiliates, or any of its and their affiliates, directors, officers, employees or agents disclaim any and all responsibility or liability whatsoever for any loss (direct or consequential) or damage arising out of the use of all or any part of this report or its contents or otherwise arising in connection therewith. Information and opinions contained herein are subject to change without notice. Past performance is not indicative of future results. Foreign currency rates of exchange may adversely affect the value, price or income of any security or financial instrument mentioned in this report. For investment advice, trade execution or other enquiries, clients should contact their local sales representative. Any opinions expressed in this report are subject to change without notice and may differ or be contrary to opinions expressed by other business areas or groups of Samsung Securities. Any analysis contained herein is based on numerous assumptions. Different assumptions may result in materially different results. Samsung Securities is under no obligation to update or keep current the information contained herein. Samsung Securities relies on information barriers to control the flow of information contained in one or more areas or groups within Samsung Securities into other areas or groups of Samsung Securities. Any prices stated in this report are for information purposes only and do not represent valuations for individual securities or other financial instruments. Samsung Securities makes no representation that any transaction can or could have been effected at those prices and any prices contained herein may not reflect Samsung Securities' internal books and records or theoretical model-based valuations and may be based on certain assumptions. Different assumptions by Samsung Securities or any other source may yield substantially different results. Additional information is available upon request.

For reports to be distributed to US:

Securities research is prepared, issued and exclusively distributed by Samsung Securities Co., Ltd., an organization licensed with the Financial Supervisory Service of South Korea. This research may be distributed in the United States only to major institutional investors as defined in Rule 15a-6 of the U.S. Securities Exchange Act of 1934, as amended, and may not be circulated to any other person otherwise. All transactions by U.S. investors involving securities discussed in this report must be effected through Samsung Securities (America) Inc., a broker-dealer registered with the U.S. Securities & Exchange Commission and a member of the Financial Industry Regulatory Authority/SIPC, and not through any non-U.S. affiliate thereof. The analysts listed [on the front of this report] are employees of Samsung Securities Co., Ltd., or a non-U.S. affiliate thereof, and are not registered/qualified as research analysts under applicable U.S. rules and regulations and may not be subject to U.S. restrictions on communications with covered companies, public appearances, and trading securities held by a research analyst account.

For reports to be distributed to UK:

This report is not an invitation nor is it intended to be an inducement to engage in investment activity for the purpose of section 21 of the Financial Services and Markets Act 2000 of the United Kingdom ("FSMA"). To the extent that this report does constitute such an invitation or inducement, it is directed only at (i) persons who are investment professionals within the meaning of Article 19(5) of the Financial Services and Markets Act 2000 (Financial Promotion) Order 2001 (as amended) of the United Kingdom (the "Financial Promotion Order"); (ii) persons who fall within Articles 49(2)(a) to (d) ("high net worth companies, unincorporated associations etc.") of the Financial Promotion Order; and (iii) any other persons to whom this report can, for the purposes of section 21 of FSMA, otherwise lawfully be made (all such persons together being referred to as "relevant persons").

Any investment or investment activity to which this report relates is available only to relevant persons and will be engaged in only with relevant persons. Persons who are not relevant persons must not act or rely on this report.

For reports to be distributed to Korea:

This report is for private circulation only, not for sale, and is issued and distributed only to persons permitted under the laws and regulations of Korea.

For reports to be distributed to Singapore:

This report is provided pursuant to the financial advisory licensing exemption under Regulation 27(1)(e) of the Financial Advisers Regulation of Singapore and accordingly may only be provided to persons in Singapore who are "institutional investors" as defined in Section 4A of the Securities and Futures Act, Chapter 289 of Singapore. This report is intended only for the person to whom Samsung Securities has provided this report and such person may not send, forward or transmit in any way this report or any copy of this report to any other person.

Analyst certification

The views expressed in this research report accurately reflect the personal views of the analyst(s) about the subject securities or issuers and no part of the compensation of such analyst(s) was, is, or will be directly or indirectly related to the inclusion of specific recommendations or views in this research report. The analyst(s) principally responsible for the preparation of this research report receives compensation based on determination by research management and senior management (not including investment banking), based on the overall revenues, including investment banking revenues of Samsung Securities Co., Ltd. and its related entities and has taken reasonable care to achieve and maintain independence and objectivity in making any recommendations.

Samsung Securities

SAMSUNG SECURITIES

Samsung Electronics Bldg., 11, 74-gil,
Seochodae-ro, Seocho-gu, Seoul, Korea 06620
Tel: 02 2020 8000 / www.samsungpop.com

Family Center: 1588 2323

Voice Of Customer: 080 911 0900

For more information,
please call our sales representatives:

LONDON

Samsung Securities Europe Limited
1st Floor, 30 Gresham Street, London EC2V 7PG UK
Tel. 44-207-776-4311
Fax. 44-203-837-9219

NEW YORK

Samsung Securities America Limited
1330 Avenue of the Americas. 10th Floor. New York,
NY 10019
Tel: 1-212-972-2454
Fax: 1-212-972-2704

HONG KONG

Samsung Securities (Asia) Limited
Suite 4511, Two International Finance Center,
8 Finance Street, Central, Hong Kong
Tel: 852-3411-3608
Fax: 852-2114-0290

BEIJING

Samsung Securities Beijing Representative Office
Rm. 910, The Exchange Building No 118 JianGuo Lu, Chao
Yang District, Beijing, China
Tel: 86-10-6522-1855 (extension 7891)
Fax: 86-10-6522-1855 (extension 7889)

TOKYO

Samsung Securities Tokyo Representative Office
#106-8532 19F, Roppongi T-Cube 3-1-1,
Roppongi Minato-ku Tokyo, Japan
Tel: 81-3-6333-2952
Fax: 81-3-6333-2953

Member of
Dow Jones
Sustainability Indices
Powered by the S&P Global CSA