

Tech Team

Jongwook Lee

Team Leader

jwstar.lee@samsung.com

Joonghan Kim

Senior Analyst

joonghan1.kim@samsung.com

Youngjin Lee

Senior Analyst

youngjin91.lee@samsung.com



Scan to go to
Research Center report database

Tech (OVERWEIGHT)

Kimi moment: The intensifying frontier AI race

- The emergence of Kimi K3 has triggered a distinct form of market anxiety, different from the DeepSeek shock.
- While not a direct driver of infrastructure investment slowdown, it has amplified volatility in the business models led by US frontier labs.
- AI model advancements continue to be a structural tailwind for semiconductor demand. Kimi K3's architectural approach intensifies demand for premium HBM and high-bandwidth interconnects. Memory stocks are now in a bottoming phase.

WHAT'S THE STORY?

Kimi K3's characteristics: Kimi K3 is the largest open-weight model to date and confirms the continued validity of scaling laws. It showed that China is now competing at the frontier, not just the value segment. Although API pricing remains competitive, per-task cost advantages erode at scale due to high token consumption. K3 has limitations (eg, overreasoning, slower inference speeds, and inconsistent user experience). Its Stable Latent MoE is a highly sparse Mixture-of-Experts architecture activating only 16 of 896 experts per inference, achieve high scaling efficiency.

The Kimi shock vs DeepSeek shock: If the DeepSeek shock represented China's value-for-performance breakthrough, the Kimi shock embodies a new form of fear—that Chinese firms may soon lead the frontier race. Should Chinese companies take the lead, it could undermine confidence in OpenAI's and Anthropic's valuations and normalize China-driven market shocks. That said, based on the Jevons Paradox, the long-term trend of rising total demand (token usage) driven by falling prices is expected to persist. A near-term slowdown in capex is also unlikely to materialize.

The ball is back in the frontier labs' court: What matters now is usability—not just performance or efficiency as measured by benchmarks, but how few tokens, time, and energy are required to deliver the desired outcome. The current turmoil from Chinese firms is likely less a crisis and more the growing pains of transitioning into a mature AI services era. Going forward, value should be directly tied not to which model you use, but to what service you deliver and at what cost. The frontier race should remain sustainable only when both model capability and infrastructure strength are aligned.

(Continued on the next page)

AI model advancement as the tipping point for semiconductor demand: Kimi K3 has spurred competition among frontier labs and contributed novel ideas to technical advancement, accelerating the pace of AI innovation—which in turn drives semiconductor demand. For semiconductor providers, the key question is whether AI customers prioritize performance improvement or cost reduction; Kimi K3 introduces a new phase in performance-oriented development. Even if AI models are China-led, if open models enable broader service deployment, they can still positively impact semiconductor demand.

Memory stock bottoming: Kimi K3's design demands massive HBM capacity and sparse MoE inference requires greater inter-GPU bandwidth. The structural importance of memory semiconductors in AI systems continues to rise. Samsung Electronics and SK Hynix are trading at near-term trough valuations—P/B of 1.8x and 2.5x, respectively. The fact that tier-1 memory stocks are down 38% from their recent peaks and that the Philadelphia Semiconductor Index's P/E has fallen below that of M7 (excluding Tesla) indicates excessive sectoral concentration is easing.


The emergence of Kimi K3: making waves

What defines Kimi K3's competitiveness?

To assess the impact of Kimi K3, it's essential to evaluate its competitive positioning. Kimi K3 is defined by four key characteristics:

- 1) Largest open-weight model to date: Assuming full model weights are released on Jul 27, Kimi K3 boasts 2.8t parameters—nearly double the size of DeepSeek V4 Pro (1.6T), previously the largest open-weight model. This scale confirms the continued validity of scaling laws.
- 2) Competing at the frontier, not just the value segment: Kimi K3 scores 57 on the Artificial Analysis Intelligence Index, placing it third behind Anthropic's Fable 5 (60) and OpenAI's GPT-5.6-Sol (59), and surpassing models like Claude Opus 4.8 and GPT-5.5. It achieved top rankings in select benchmarks, including frontend coding tasks. While not yet SOTA across all metrics, its performance marks a decisive shift from China's prior dual-track strategy—either ultra-low-cost (eg, DeepSeek, Xiaomi MiMo) or mid-tier performance at moderate pricing (eg, Z.AI GLM). The fact that Kimi K3 outperforms Claude Opus 4.8, released just two months prior, suggests the performance gap between US and Chinese frontier models has narrowed to under two months.
- 3) API pricing remains competitive but per-task cost advantages erode at scale due to high token consumption: Kimi K3's API pricing is set at USD3 per million input tokens and USD15 per million output tokens—matching the standard rate of Claude Sonnet 5 (index score: 53), though Sonnet 5 is currently discounted through August. This represents a significant increase from K2.6's pricing (USD0.95 and USD4). While K3 uses 21% fewer output tokens than K2.6 to achieve comparable results, its absolute token consumption remains high at 130m tokens per task. As a result, the effective cost per task is USD0.95—comparable to GPT-5.6-Sol (USD1.04), and still below Fable 5 (USD2.75) and Opus 4.8 (USD1.8).
- 4) Persistent limitations in user experience: Kimi models have long exhibited over-reasoning tendencies, leading to slower inference speeds—a known characteristic of Chinese LLMs. Despite strong benchmark scores, real-world usability lags behind US counterparts. Even Moonshot AI acknowledges a noticeable gap in user experience compared to Claude Fable 5 and GPT-5.6-Sol.


Parameter size of major open weight models



Source: Moonshot AI, Samsung Securities

Artificial Analysis Intelligence Index comparison of major models


(Scores)



Source: Artificial Analysis, Samsung Securities

Cost per task of major models


(USD)



Source: Artificial Analysis, Samsung Securities

API price comparison of major models (based on blended price)

(USD)



Note: Blended Price based on Cache:Input:Output = 7:2:1

Source: Artificial Analysis, Samsung Securities


Key architectural innovations in Kimi K3

Kimi K3 is built on two core innovations—Kimi Delta Attention (KDA) and Attention Residuals (AttnRes)—combined with a Stable Latent MoE structure, which the company claims delivers 2.5x higher scaling efficiency than Kimi K2. Independent validation via technical documentation and weight release remains pending.

- **Stable Latent MoE:** A highly sparse Mixture-of-Experts architecture activating only 16 of 896 experts per inference. This drastically reduces computational load while preserving model capacity, enabling efficient scaling of model capacity.
- **KDA:** A novel attention mechanism optimized for long-context retention.
- **AttnRes:** Instead of carrying forward all intermediate representations across layers, AttnRes selectively retains and reuses only the most relevant latent features. This enhances information efficiency and supports deeper, more scalable architectures.

Another differentiator is Kimi K3's end-to-end quantization-aware training. Unlike K2, which was quantized post-training to INT4, K3 integrates quantization from the SFT (supervised fine-tuning) stage, training with MXFP4 weights and MXFP8 activations. This improves hardware compatibility. However, the model's full efficiency is contingent on high-bandwidth interconnects. As a result, Moonshot AI recommends deployment on super-node configurations with 64+ accelerators.

Kimi K3 architecture



Kimi K3 architecture: the Stable LatentMoE and KDA modules (left), the AttnRes operation α (top right), and the Block Attention Residuals backbone (right).

Source: Moonshot AI, Samsung Securities

Why does misfortune come all at once?

What distinguishes the Kimi shock from the DeepSeek shock?

While the DeepSeek shock was a cost-efficiency shock, the Kimi shock represents a new and deeper fear: that China may now be capable of leading the frontier race itself. The two shocks appear similar on the surface, but the implications are fundamentally different.

Even if limited to certain benchmarks, Kimi K3's ability to outperform US frontier models in frontend coding is enough to send ripples through the industry. This comes at a particularly painful moment: Google's Gemini 3.5 Pro has been delayed due to underwhelming coding performance.

The Jevons Paradox is at work, but...

The infrastructure and energy required to serve high-intelligence models have become a physical bottleneck. Chinese firms cannot be exempted from this reality. In fact, Kimi K3's API pricing is comparable to Claude Sonnet 5—far from the overwhelming cost-efficiency offered by DeepSeek. The fact that Moonshot AI immediately faced acute computing power shortages after the K3 launch further supports this point.

The DeepSeek shock was overcome by the Jevons Paradox. Similarly, the base scenario for the Kimi shock remains unchanged: 1) the trend of rising total demand (token usage)-- driven by falling prices (P)--should persist; 2) the computing power shortage centered on inference should continue regardless of Chinese models' performance; and 3) a near-term slowdown in capex is unlikely. From the perspective of AI demand stimulation, Chinese firms are effectively doing for the industry what Nvidia has long done—expanding the open-source ecosystem.

However, if Chinese firms begin to lead the frontier race, entirely new risks should emerge:

1) Doubts over the fair valuation of frontier AI firms


The legitimacy of frontier AI lab valuations (OpenAI, Anthropic) has largely rested on model performance. The market's unwritten rule has been: if a firm can consistently produce the best-performing models and sustain explosive growth, it deserves a USD1t valuation. But if Chinese firms—currently valued at USD20–30b—can produce models with only marginal performance gaps and release them as an open-weight model within several weeks, the valuation rationale for US labs becomes increasingly tenuous. This is particularly adverse for Anthropic, which is preparing for its upcoming 2H IPO.

The rapidly evolving enterprise market is broadly divided into two segments: coding and agentic workflows. If the coding segment—currently the core revenue driver for frontier labs—is commoditized through incremental model improvements, the profit outlooks could shift dramatically. Agentic workflows, meanwhile, were always a domain where cost-efficiency matters more than raw performance.

Timing couldn't be worse. Recent market volatility reflects more than just demand uncertainty—it reflects growing concern over the massive capex required to meet that demand. The growth trajectory of frontier labs is now entangled with broader factors—capital markets (bond issuance, rights offering) and macro conditions (interest rates, recession risks). This is why the market reacts so sensitively—not only to a phantom Meta Platforms-shock, but also to every routine macroeconomic data point.

As planned capital expenditures continue to rise (due to soaring semiconductor prices, power shortages, data center protests, regulatory hurdles, and delays), the assumptions underlying frontier labs' profitability become increasingly uncertain. This uncertainty itself is a major risk. Even if fundamentals remain intact, these firms must continuously outperform expectations just to maintain their valuation. Even without a fundamental miss, multiple contraction—driven by narrative erosion—can trigger price declines (multiple contraction > earnings surprise = stock decline).

Comparison of US-China frontier AI labs' valuations (market capitalization basis)



Note: DeepSeek and Moonshot AI are based on the valuation discussed in the new funding round, while Z.AI is based on market capitalization.

Source: Press, Factset, Samsung Securities

2) The normalization of the “China-driven shock”


What amplified investor anxiety was that Chinese labs' steady, incremental progress was turning what were once considered “shocks” into routine occurrences. DeepSeek, GLM, and Kimi are now releasing astonishing models in rapid succession. Tang Jie, co-founder of Z.ai, predicted in mid-June that Chinese models would reach Fable-level performance sooner than expected. DeepSeek V4's official release is imminent, and Kimi has hinted at a 3.1 update.

Just months ago, US frontier labs were expected to widen their lead through recursive self-improvement cycles. Reality has flipped: Chinese firms are closing the gap with frontier labs—and even outperforming companies like Google and Meta that lack the ability to distill or replicate these models.

What Kimi K3 made clearer is that Chinese frontier labs' competitiveness stems not merely from distillation. It is underpinned by: 1) world-class AI talent; 2) advanced reinforcement learning and optimization capabilities forged under infrastructure constraints; and 3) close collaboration with government and academia.

The fact that Chinese models have begun surpassing frontier labs on certain benchmarks suggests that future reversals on other benchmarks are highly probable. While the US still leads in overall aggregate performance, the fear of being overtaken has grown significantly compared to the past.

Performance gap of US-China frontier models



Source: Artificial Analysis, Samsung Securities

The ball is back in the frontier labs' court

On the very day Kimi K3 was unveiled, OpenAI's Sam Altman posted a telling message: "The next 12 months will be the best in OpenAI's history." Given OpenAI's recent success in deploying cost-efficient, high-intelligence models—backed by superior infrastructure compared to Anthropic—it's clear that they are highly confident in their future. Since the GPT-5.6 launch (Jul 9), usage of Codex + ChatGPT Work has surged: from 5m users to 6m by Jul 12, then 9m by Jul 15—adding a million new users per day.

Ultimately, what matters most is usability—not raw benchmark scores or efficiency metrics. The ability to achieve desired outcomes using fewer tokens, less time, and lower power is becoming increasingly critical. Intelligence matters, but practicality matters more.


While Kimi K3 delivers impressive performance, concerns have emerged around over-reasoning and excessive indexing optimization. Although its API pricing is lower than US frontier models, the real cost advantage per task has narrowed significantly. Many analysts argue that, when balancing price, speed, and performance, there is little incentive to switch from top-tier US models.

Moreover, while Chinese firms—likely aided by distillation—have demonstrated remarkable efficiency in training, they have yet to show comparable breakthroughs in inference optimization. Despite scaling model size, Kimi K3's API price rose substantially compared to its predecessor, and output token volume remains high. This contrasts sharply with the US frontier labs' dual focus on simultaneous training and inference optimization. The success—or failure—of Chinese firms in inference optimization should become a key indicator of their long-term competitiveness and their ability to sustain an efficient infrastructure cycle.

On a positive note, the current turbulence from the Chinese AI firm may simply reflect growing pains as the industry transitions into the era of mature AI services. In the future, value creation should be less about which model you use, and more about what service you deliver at what cost.

In the end, the frontier race is sustainable only when both model capability and infrastructure strength are aligned. US frontier labs still hold a clear lead. Despite the astonishing progress of Chinese firms, it is still too early to declare a winner. Having said that, uncertainty around companies like Google and Meta Platforms—those that allowed China to catch up—is growing. This endless cycle of proof, iteration, and escalating performance should become the primary driver of volatility in the AI theme.

Trend of OpenAI codex + ChatGPT work WAU



Source: OpenAI, Samsung Securities

The next 12 months are the best 1 year in OpenAI history



Source: X, Samsung Securities

Kimi K3, the tipping point in semiconductor demand

Technical misconceptions and truth: Kimi K3 is driving semiconductor demand

Interpreting this model as a technology that reduces memory usage is a misconception. The opposite is true: it enables the utilization of model performance that was previously inaccessible due to excessive DRAM requirements, by making it efficient enough to run within practical memory constraints.

The core innovation of Kimi K3 is enabling the practical execution of 1m-token long contexts on a 2.8t-parameter ultra-high-performance AI model. Standard attention mechanisms cause KV cache memory to grow quadratically with sequence length—making it physically infeasible to run such models at this scale. Kimi K3 overcomes this barrier, bringing performance to a commercially deployable cost level. From a semiconductor demand perspective, Kimi K3 sends a clear signal:


First, this is not a model that reduces computing power—it is inherently extremely heavy.

- Moonshot's official recommended server configuration is a super-node composed of 64 or more accelerators. An efficiency-optimized model would not recommend rack-scale or large-scale scaling domains.
- Loading only the model weights into cache consumes approximately 1.5 TB of HBM. This figure results from quantizing 2.8t parameters using MXFP4 ($2.8T \times 4$ bits)—an inevitable consequence of ultra-large parameter counts.
- There is virtually no room in HBM to accommodate the KV cache, necessitating its offloading to server DRAM or NVMe.

Although the total parameter count is 2.8t, Kimi K3 employs a sparse MoE architecture (StableLatentMoE), activating only 16 out of 896 experts per token. Since all 896 experts cannot fit on a single GPU, they are distributed across multiple GPUs (WideEP). This implies that, from a semiconductor perspective, scaling up (high-bandwidth interconnects within a rack) is paramount. In simple terms, the computational load across GPUs is not the bottleneck; rather, it is the growing network demands between MoE experts and the escalating need for high-bandwidth memory. This is why Moonshot recommends a super-node connecting 64 accelerators.

Second, it is true that using linear attention (KDA) reduces KV cache usage—but this is precisely where market misunderstanding arises. Moonshot's research shows that KDA can compress KV cache by up to 75% at the same scale. However, this efficiency is not primarily about reducing total compute resources—it is a necessary trade-off to achieve the 1m-token capability. Without KDA, Kimi K3's target performance would be unattainable. Reducing KV cache burden was a strategic step toward commercialization—not evidence of a reduction in memory demand. Interpreting it as such would be an overstatement.

Kimi K3 requires at least an Nvidia GB200 NVL72



Source: SemiAnalysis, Samsung Securities

KDA and Wide EP drive scale-up demand over scale-out



Source: Lenovo

Competition accelerates the tipping point of demand

The intense rivalry among AI labs is, paradoxically, a tailwind for semiconductor demand. As Chinese companies narrow the technological gap, OpenAI and Anthropic are likely to ramp up their investments further. At this stage, efficiency alone is not the priority—if performance falls short, companies simply deploy more resources. US hyperscalers must serve both their proprietary AI lab customers and those opting for open-source Chinese models, further amplifying demand.

However, viewing China's AI model resurgence merely as a battle for market share is a narrow interpretation. From a hardware perspective, the real impact lies in how competition accelerates the pace of AI technological advancement—a phenomenon inherent to the early stage of any transformative technology. As AI improves, end-user adoption surges, and the moment the demand S-curve crosses its tipping point is precisely the pivotal moment we've been waiting for. In this light, the performance race among AI labs is the primary driver of semiconductor demand. We've already witnessed the most classic example in the explosive rise of OpenAI's Codex usage.

As previously noted, the user base for Codex + ChatGPT Work grew from 1m in February to 8m by July 14-15—nearly a sevenfold increase in just five months. Notably, within three days of the official launch of GPT-5.6 on Jul 9, 3m new users joined, pushing usage from 6m to 8m. The driver was not cost, but performance: GPT-5.6 Sol has reached a new peak of 80 on the Coding Agent Index, reducing output tokens by more than half, cutting time needed by over 50%, and lowering costs to approximately one-third. Added to this were service enhancements—Codex's integration with ChatGPT desktop, remote functionality, and GitHub connectivity—that significantly improved usability.

The surge in Codex users offers two key insights. First, consumers are ready to adopt AI as soon as price and performance reach a compelling threshold. Second, the scale of this growth indicates the market remains in its early stage—far from saturation. The rapid demand even strained OpenAI's infrastructure, forcing temporary reductions in context window capacity, a classic sign of early-market demand outpacing supply.

As long as end users prioritize performance over cost, AI labs and hyperscalers have no reason to hesitate in spending on semiconductors. The term “early-stage market” means the growth S-curve has not yet broken out—there is a latent demand waiting to be unlocked, whether in the US or China. Once performance crosses a threshold, adoption follows rapidly—and Kimi K3 has delivered a meaningful breakthrough in long-context token handling.

Whichever side—US or China—opens the market, it benefits semiconductor companies. China accelerating the inflection point of the S-curve is viewed as a positive for semiconductor demand.

That said, we must also acknowledge the counterpoint raised by software analysts: even as demand remains robust, capital markets view the massive capex required to support it as a risk. As a result, short-term narrative deterioration—separate from fundamentals—can exert significant multiple compression pressure. This is not a question of directional price movement, but of increased volatility.

US and Chinese AI in symbiotic virtuous cycle

When DeepSeek R1 first emerged, few fully grasped its significance—but now, it’s clear to all: this is the story of technological symbiosis.

DeepSeek R1 shattered the industry’s obsession with sheer parameter scale by demonstrating that a moderately sized model (671B total, 37B active) with enhanced MoE architecture could achieve commercially viable inference efficiency. Its arrival jolted the US AI industry awake, accelerating the real-world deployment of agentic AI. And agentic AI has fundamentally reshaped the direction of semiconductor demand—not only driving demand for AI accelerators and HBM, but also sparking new demand for CPUs and server DRAM based on massive token consumption and KV cache workloads.

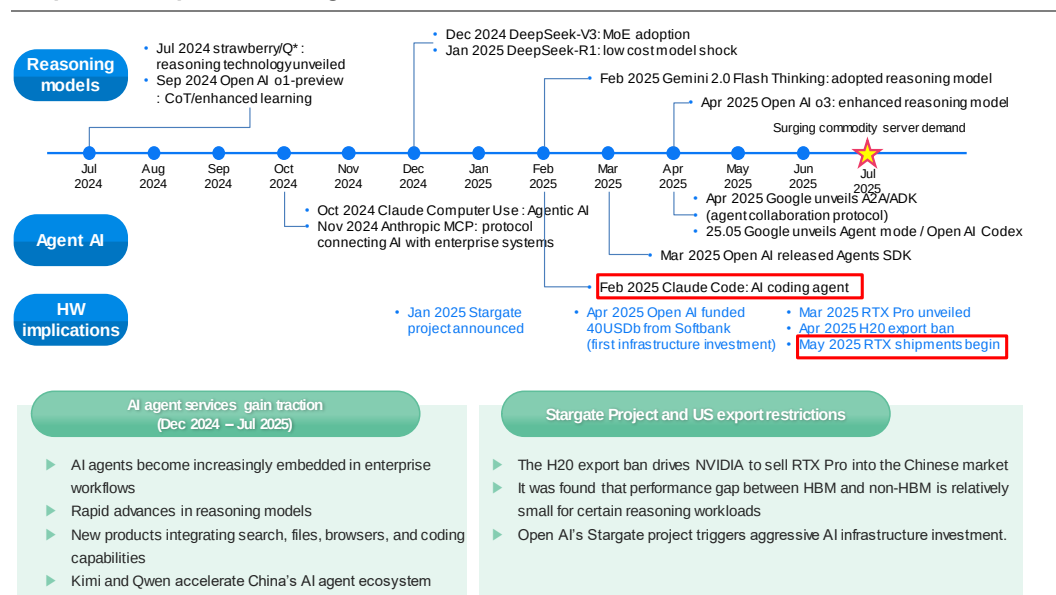
Looking back, the DeepSeek moment is not a one-off but becoming a recurring pattern. The US leads with capital and academic research to set the direction; China responds with hardware-efficient engineering and architectural innovation; then US AI firms absorb and deploy these advances at scale—creating a dynamic of technological symbiosis and competition.

- **Reasoning AI:** 1) Direction setting (US): GPT-o1 introduces a new paradigm for inference-time computation. → 2) Realization (China): DeepSeek delivers a deployable cost model using a 671B MoE architecture. → 3) Reabsorption (US): Frontier labs actively adopt MoE architecture.
- **Long-horizon agent:** 1) Direction setting (US): Gemini’s long-token paradigm + linear attention theory. → 2) Realization (China): Kimi K3 delivers a cost-effective long-token autonomous agent using KDA + MoE. → 3) Reabsorption (US): Linear attention and extreme MoE granularity are expected.

Yang Zhilin, founder of Moonshot, is a PhD graduate of Carnegie Mellon University (CMU) and the first author of Transformer-XL and XLNet—classic works in long-sequence modeling. His research DNA is inherently rooted in “ultra-long-context” understanding. Thus, Moonshot’s deliberate focus from its inception on ultra-long-context as its core battlefield was no accident. By building a 2.8t parameter model capable of retaining up to 1m words of context—and achieving a commercially viable cost structure—its public release marked a tipping point where China’s trajectory aligned with US direction.

Now, US big tech firms are set to adopt the architectural recipe proven by Kimi K3—lowering infrastructure costs and enhancing long-context agent performance in their next-generation agent models (future GPT, Claude, and Gemini lineups). As a result, the performance bottlenecks that have long constrained agentic AI will soon be overcome, signaling that a new agent market—previously closed due to technical limits—is on the verge of opening. This is the lesson left by DeepSeek: whether demand originates in the US or China, once one side opens the market, the other follows. And when the market opens, infrastructure demand surges.

DeepSeek R1 sparks the AI agent era



Source: Samsung Securities

Semiconductor peak to arrive in one of two paths

Paradoxically, to argue for sustained strength, one must first define the conditions under which this cycle could peak. We believe the AI semiconductor cycle will end via one of two pathways.

First, all technology cycles end when innovation stalls. The AI cycle would falter if performance gains ceased—whether due to exhausted use cases, diminishing returns from model scaling, stagnation in AI accelerator performance, memory bandwidth limits, or the physical constraints of semiconductor scaling. Yet, as noted earlier, we are not in that phase. New frontiers—like autonomous agents—are just opening, and performance continues to advance rapidly.

Second, the cycle ends when hyperscalers can no longer finance infrastructure expansion. This is what makes this cycle uniquely different: demand no longer stems from consumer devices, but from infrastructure—and that infrastructure is being built with financial capital. As a result, this cycle is now tightly linked not just to frontier labs' revenue, but to broader capital markets—interest rates and liquidity. That's why markets reacted so sharply to the 'Meta shock'. Yet paradoxically, the financial cycle is still in its early stages, making a near-term collapse in funding unlikely.

For memory semiconductors specifically, one additional risk must be considered: What if AI's growth continues, but memory falls out of the critical path? For instance, if compute or network bandwidth becomes the new bottleneck, memory capacity expansion could be deprioritized. However, this would represent a relative risk in terms of stock performance, but it would not explain a sharp selloff. In fact, architectures like Kimi K3—demanding massive HBM capacity and offloading KV cache to slower DDR5/NVMe storage—do not reduce memory's centrality. Rather, memory is being pulled back into the center of the cycle.

This event reinforces strong conviction in the stock price bottom

We believe that SEC and SK Hynix are trading at near-term trough valuations—P/B of 1.8x and 2.5x, respectively—and current prices reflect a sentiment low, not a fundamental one. Sentiment has turned deeply negative, to the point where even positive developments are being misinterpreted as bad news.

Just as DeepSeek's January release helped ignite the Agent AI cycle from July onward, Kimi K3 is poised to serve as a catalyst for advancing long-context autonomous agents. This development confirms that the AI cycle is far from over—and that semiconductor demand may actually intensify.

- Kimi K3 does not reduce semiconductor demand—it unlocks capabilities previously unattainable.
- The race among AI models accelerates technological innovation and brings forward the tipping point for semiconductor demand.
- The triggers for a peak—technological stagnation or funding constraints—are still distant, and memory's critical role remains intact.


Two factors remain to complete the price correction.

First, the primary driver of this correction has been excessive sectoral concentration. Until this imbalance unwinds, patience is required. The fact that the Philadelphia Semiconductor Index's P/E has fallen below that of M7 (excluding Tesla) is a sign this concentration is already easing.

Second, it is necessary to identify the catalyst for recovery. In the near term, this should come from cloud providers' earnings announcements. In the longer term, it should stem from tangible demand growth driven by real-world improvements in Agent AI performance.


The war has expanded—and paradoxically, more ammunition is needed. Semiconductor demand in this cycle has not been broken.

SOX index vs M7 (ex. Tesla): 12m fwd P/E




Source: Quantwise, Samsung Securities

Samsung Electronics: 12m fwd P/B vs ROE



Source: Quantwise, Samsung Securities

SK Hynix: 12m fwd P/B vs ROE




Source: Quantwise, Samsung Securities

Compliance notice

- As of 7/21 2026, Samsung Securities shared group affiliation with Samsung Electronics.
- As of 7/21 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 7/21 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)

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

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 2026.06.30

BUY(85.5%)-HOLD(14.5%)-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.

Copyright © 2010 Samsung Securities Co., Ltd.. All rights reserved. This report or any portion hereof may not be reprinted, sold or redistributed without the prior written consent of Samsung Securities America Inc

Samsung Securities

SAMSUNG SECURITIES

Samsung Electronics Bldg., 11, 74-gil,
Seochoda-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