KASASV

KOREAN AMERICAN SEMICONDUCTOR ASSOCIATION IN SILICON VALLEY

Dec 2025

딱 한 발짝만 내디뎌 보려고 한다

내가 가려는 먼 곳을 쳐다보며 걷는 게 아니라,

지금 있는 자리에서 발을 쳐다보며 일단 한 발짝을 떼는 것, 그것이 시작이며 끝이다.

그렇게 한 발짝 한 발짝 내딛는 데 집중하다 보면 어느새 목적지에 도착해 있는 나를 발견하게 되는 것이다. 도저히 못 갈 것 같은 순간에도 발을 쳐다보며 한 발짝 떼는 것은 언제든 가능하다.

물론 선택한 길이 틀릴 수도 있고, 최선을 다했는데도 낭떠러지에 도착할 때도 있을 것이다.

하지만 그게 두려워 한 발짝도 떼지 않으면 영영 아무 데도 못 가게 된다.실패를 하더라도 실패로부터 무언가를 배우면 그것은 더 이상 실패가 아니었고, 길을 잘못들었다 싶어도 나중에 보면 그 길에서 내가 미처 몰랐던 것들을 배움으로써 내 삶이 더 풍요로워졌다.

때론 내 뜻대로 흘러가지 않는 인생 때문에 화가 난 적도 있지만, 분노의 힘이 나를 살게 한 적도 있다.

그러므로 가장 빠른 직선 코스로 가야 한다는 강박관념만 버린다면, 한 발짝을 떼는 것에 대해 두려움을 가질 이유는 없다.

-김혜남

INTRO

E-mail 주소는 blind 로 나가기때문에 공유가 되지는 않습니다. 혹시 개인 메일주소를 원하시면 저에게 알려주시면 update 하도록 하겠습니다. 그리고 Newsletter distribution 을 원치않으시면, 알려주시기 바랍니다.

12 월 27 일(토)에는 올해 4 번째 Golf Tournament 가 있습니다. 연말이긴 하지만 가능하신 분들과 좋은 시간을 보냈으면 합니다.,

자세한 내용은 추후에 웹사이트(announcement)를 참조하시기 바랍니다.

Linkedin 에 KASASV 를 officially launch 했습니다. 많은 관심 부탁드립니다.

저희는 미국 Internal Revenue Code 501(c)(3)에 의거하여 Federal Income tax exemption 을 받은 비영리단체로 donation/협찬으로 운영이 되고 있습 니다. 관심이 있으시면 저희 웹사이트 Partners tab 을 보시면 좀 더 자세히 보실수 있습니다.

- 사이먼 리 드림/협회 상근임원

www.kasainsv.com 문의 info@kasainsv.com

감사합니다

Samsung Electronics: Value, Scale, And The Future Of Semiconductors

Nov 19, 2025 Piotr Kasprzyk

Summary

- Samsung doesn't get nearly as much attention as its competitors, while its stock price has almost doubled year-to-date, surging 92.7%.
- The company remains at the forefront of major product categories: smartphones and memory, defending its leadership in both markets for a long time.
- However, global trade and macro factors influence demand, affect margins depending on geography, and force the leadership to adapt and develop new strategies.
- Nevertheless, for those looking for value, strong AI exposure, and geographical diversification, Samsung may be a compelling proposition.

Karlis Damabrans/iStock Editorial via Getty Images

Samsung Electronics (OTCPK:SSNLF) is presumably mostly known to consumers as a manufacturer of high-quality smartphones, electronic devices, and smart home appliances. Indeed, it's a large part of Samsung's operations, generating over half of the company's total revenue. However, what currently captures investors' attention is Samsung's memory and foundry businesses, as the AI boom continues.

The company operates in a highly competitive industry in an extraordinarily dynamic environment. Yet, while enterprises like NVIDIA (NVDA), Apple (AAPL), Micron (MU), or Taiwan Semiconductor (TSM) are widely discussed by the stock market-oriented media outlets, Samsung doesn't get nearly as much attention. Its stock price almost doubled YTD, advancing 92.7%; sales in the last quarter grew in the high single digits; and operating profit increased 32.6% YoY. A deeper look at the company's fundamentals, valuation, and moat, or lack thereof, should provide an answer as to whether it's worth investing in as of today.

Samsung's Diversified Business

Smartphone Business

Samsung is the world's leading smartphone manufacturer, with 61.4 million shipped devices as of Q3 2025, capturing a 19.0% global market share. The Korean manufacturer extended its lead over Apple to 8 basis points from 3 a year ago, which is a positive sign, although such small fluctuations are common in this industry.

Nevertheless, it's crucial that the company maintain its position and not give up its market share, as the smartphone business accounts for 42.05% of Samsung's total revenue as of H1 2025. The runner-up, Apple, has never managed to overtake Koreans despite its alleged strong moat. The Oracle of Omaha himself, Warren Buffett, once said:

Apple is in a position with consumers, where they're paying maybe $1,500 bucks, or whatever it may be, for a phone.

And the same people pay $35,000 for having a second car, and [when] they have to give up a second car or give up their iPhone, they'd give up their second car.

This moat is real. However, it applies to Samsung in the same way. Warren Buffett's statement should refer more to a smartphone in general rather than a specific brand.

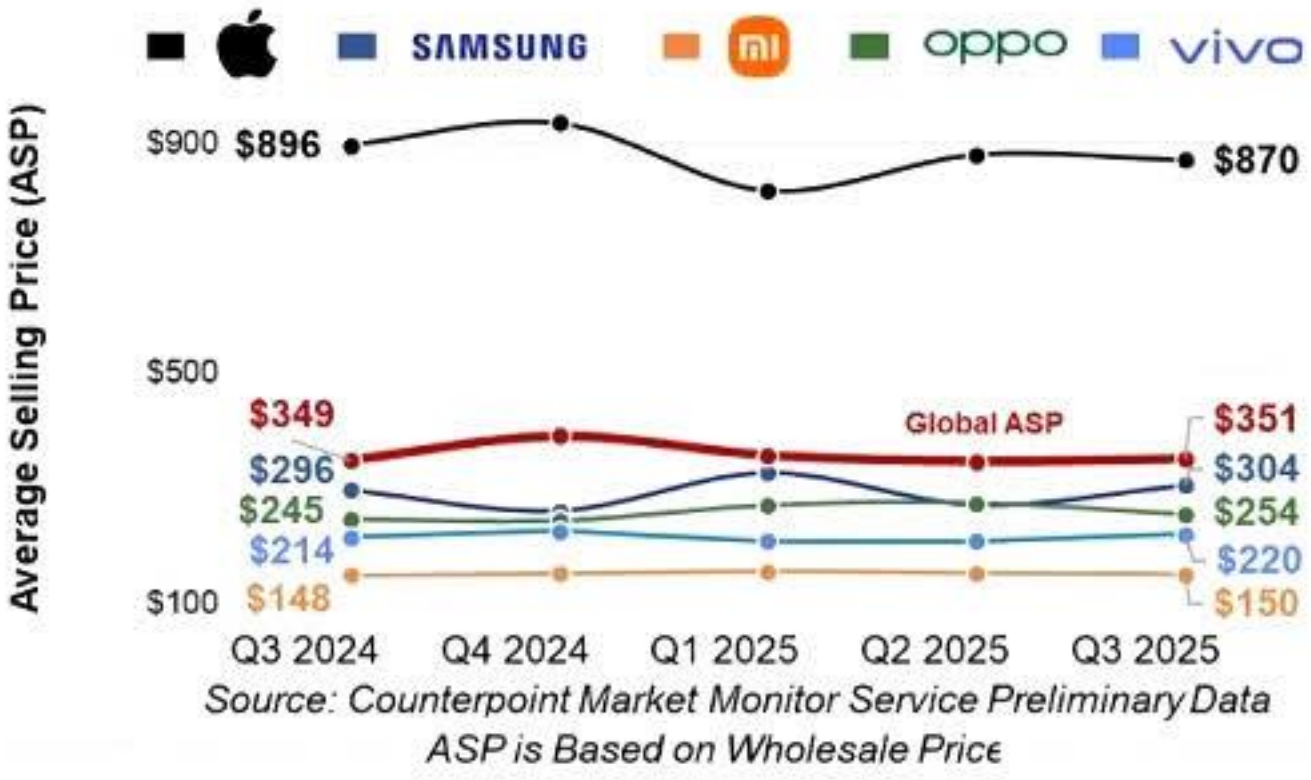

Counterpoint

Apple has the advantage of the whole ecosystem built around iPhones. Yet, considering that Samsung smartphones match or often surpass Apple devices while costing one-third less, the choice for many customers is obvious. Additionally, Koreans manufacture OLED panels, DRAM, and NAND memory chips used in Apple's devices. Thus, by purchasing a Samsung smartphone, customers receive a premium product with Samsung-made components (especially the Galaxy line), and on average, it is three times less expensive than an iPhone.

Semiconductor Business

Samsung seems to be uniquely positioned in the semiconductor ecosystem.

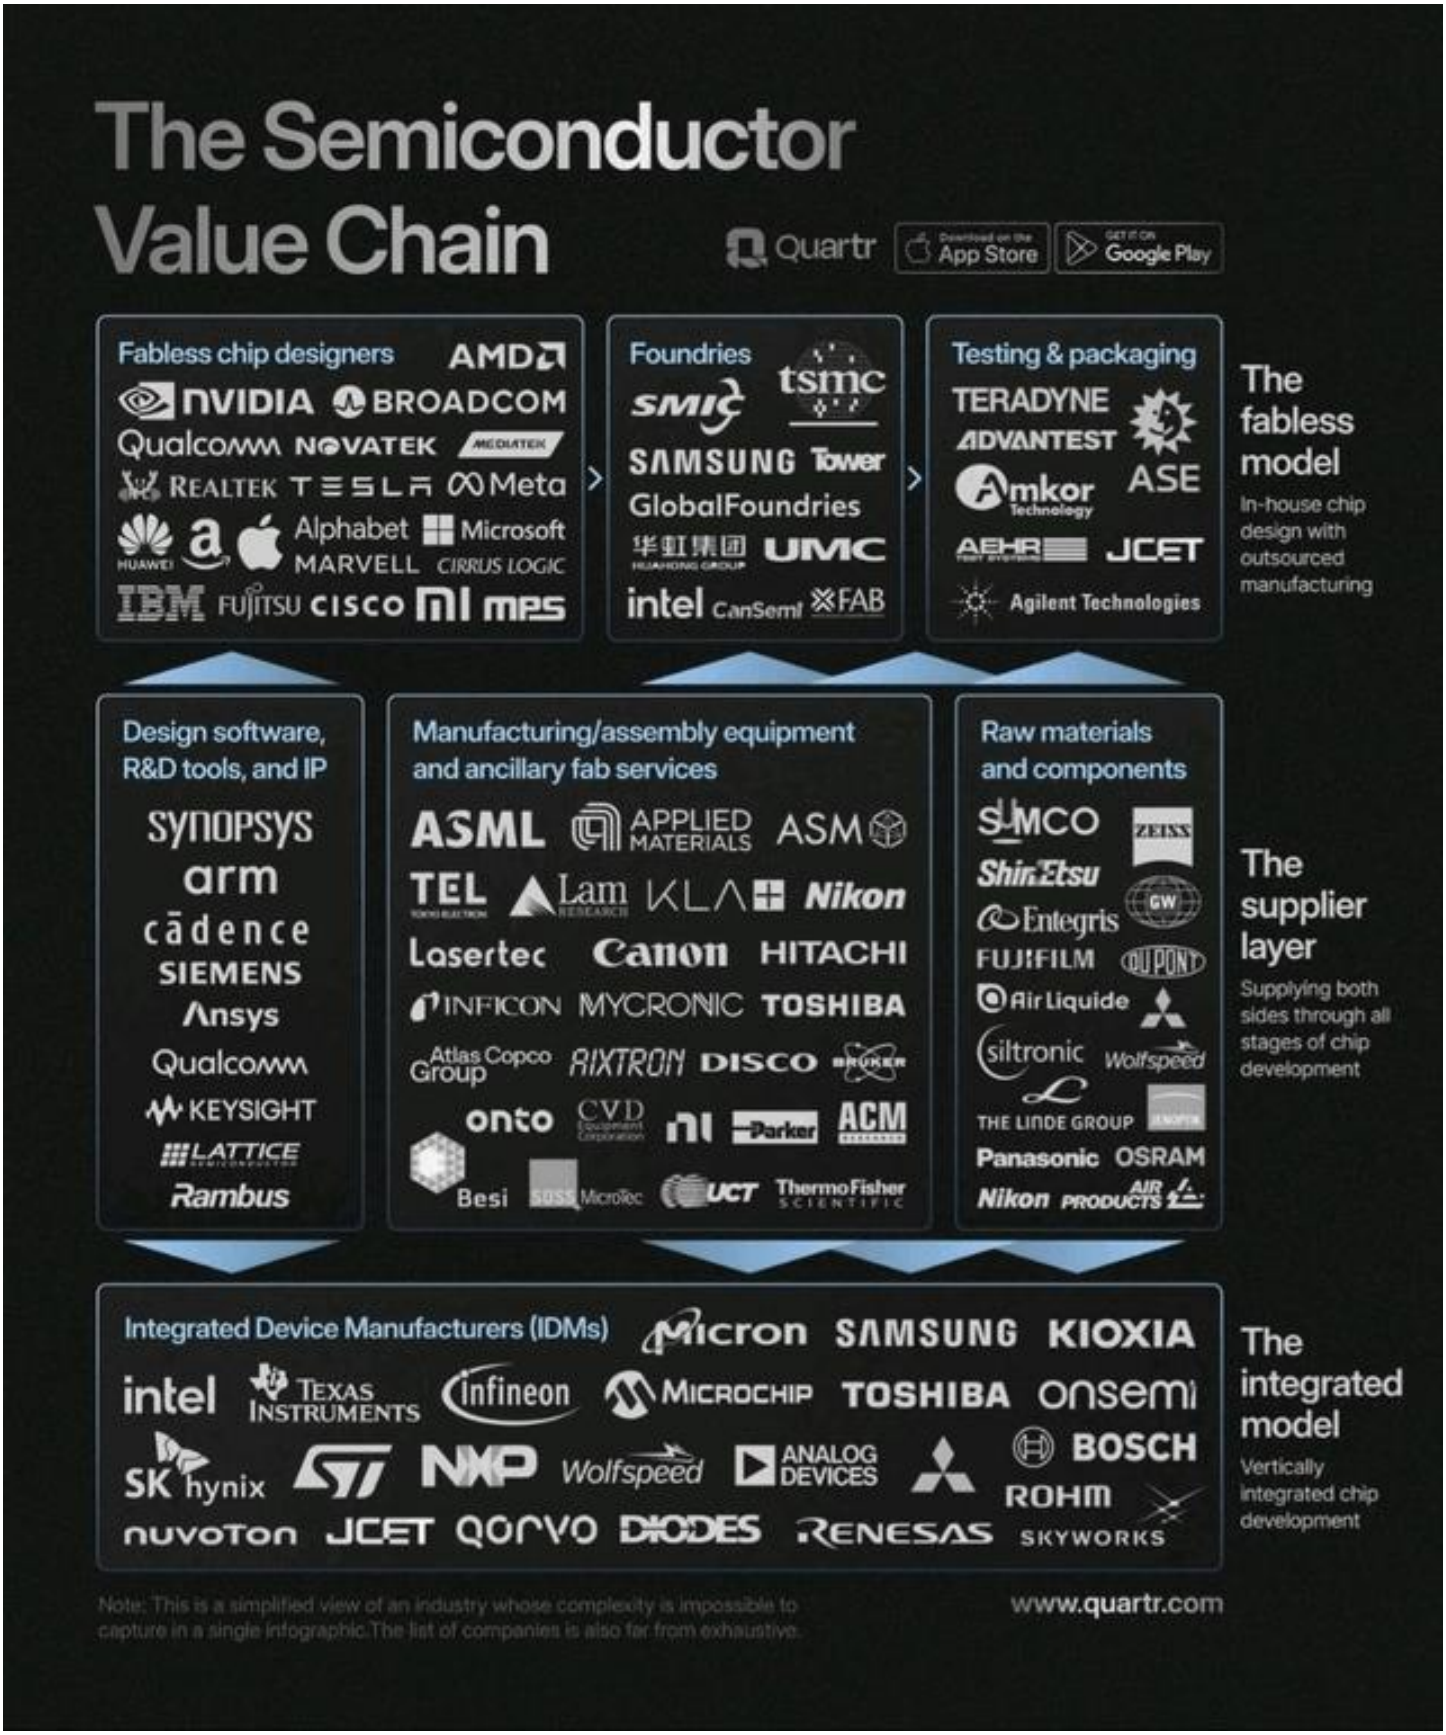

Quartr

The company manufactures chips for other companies through its foundry business. It produces customized products for customers based on their designs. Although the S.LSI/Foundry business is relatively small,

accounting for 7.4% of Samsung's total sales as of Q3 FY2025, foundry ranks second in the global foundry market share, behind only TSMC.

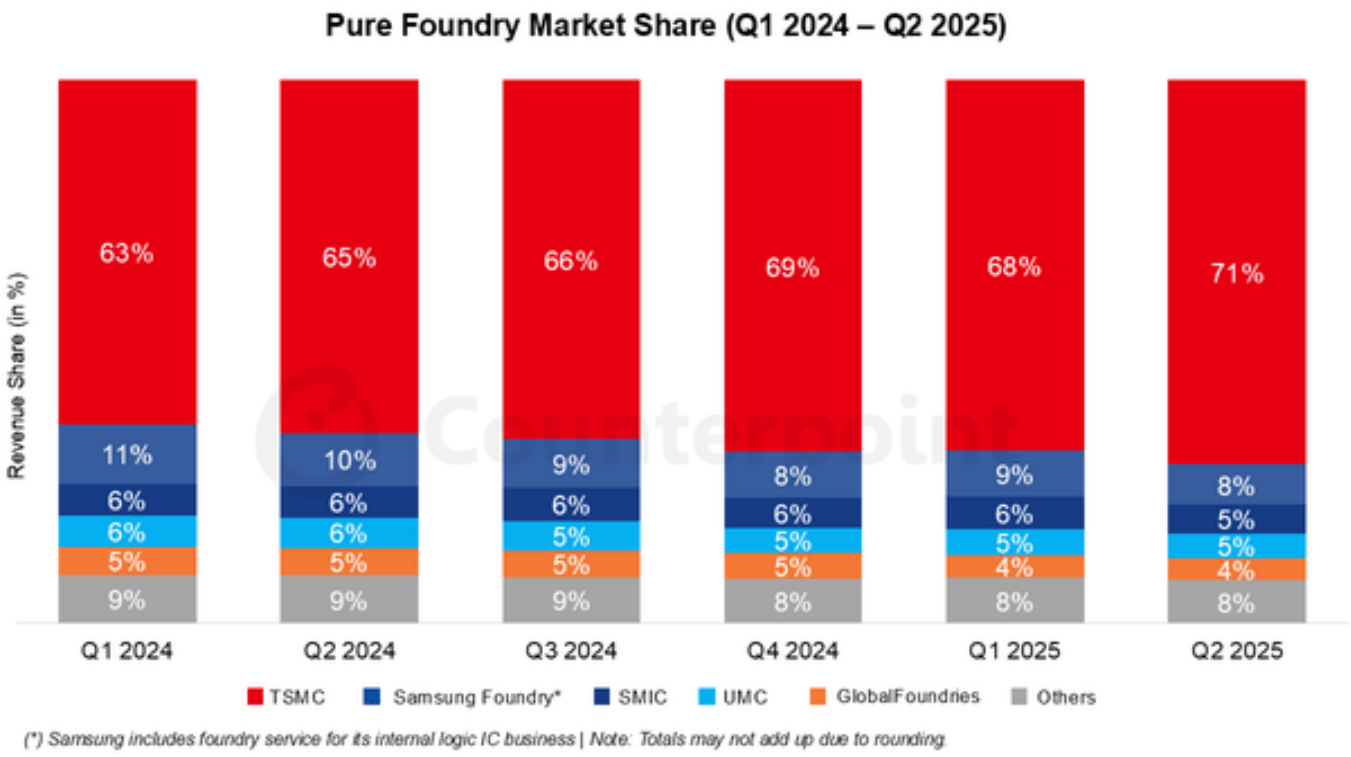

Counterpoint

The global market is expected to grow along with the demand for AI. Nevertheless, TSMC seems to have been chipping away at the share from other competitors in recent quarters. Yet, Samsung doesn't slow down. It recently started mass production of the GAA 1st generation 2nm process, with the outlook of a further expansion into the GAA 2nd generation 2nm process in 2026. Moreover, the management sounds confident about the future of its foundry segment.

We will expand our business opportunities in line with the evolving markets by enhancing collaboration with customers and strengthening competitiveness across the board, including development, orders, production, and supply.

Within Samsung's System LSI business, the company designs and sells its own logic chips, such as mobile application processors (mobile APs) and camera image sensors. Moreover, it expands beyond mobile and into the automotive market with increasing demand across infotainment, ADAS, and connectivity systems. Yet, Samsung has managed a steady supply of its high-end system-on-chip (SoC) processors for smartphones and other devices.

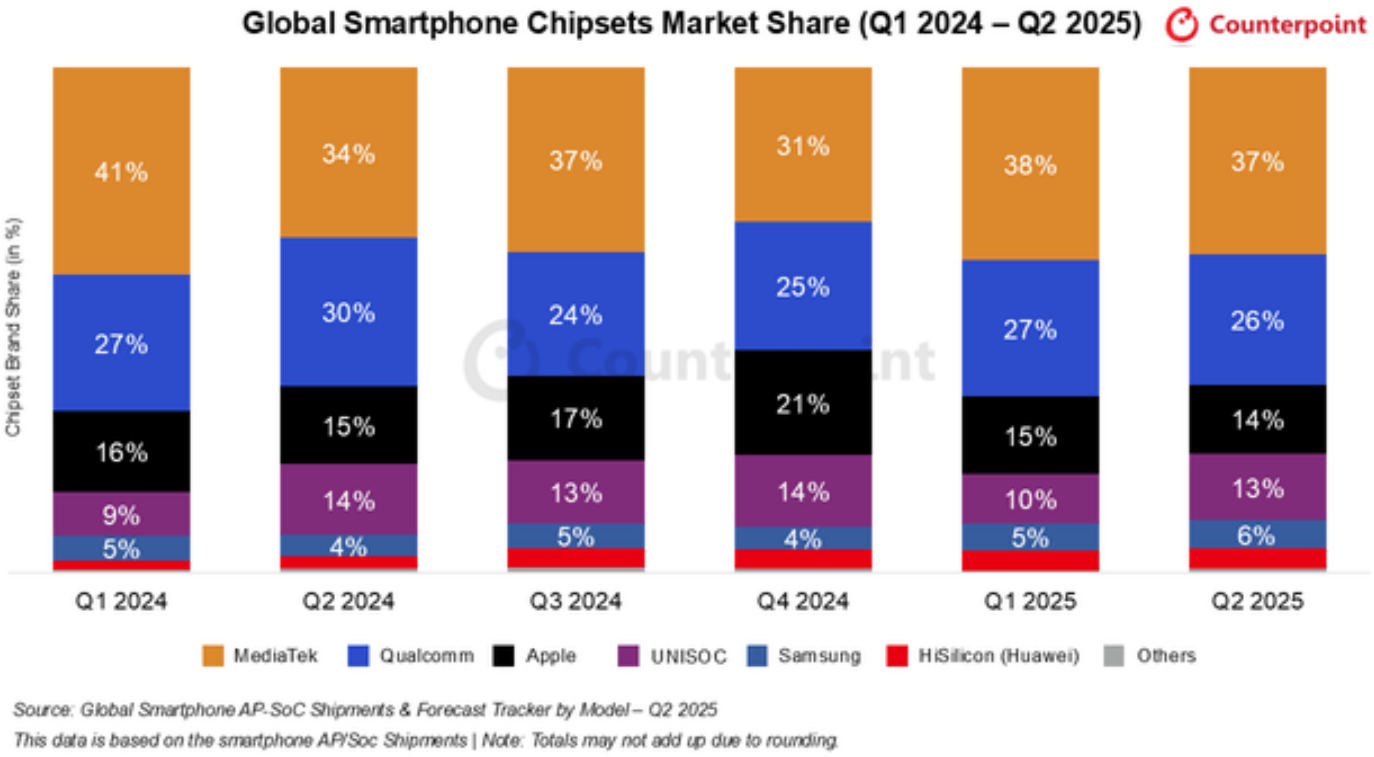

Counterpoint

The company remains among the top 5 global smartphone chipset manufacturers, holding a relatively small 6.0% market share.

Memory business, together with the System LSI and Foundry segments, makes up the Device Solutions Division, with Memory being the largest among them. It remains in focus among investors due to the AI demand and robust performance. Most importantly, Samsung leads the global memory market, with its revenue climbing to $19.4 billion in Q3 2025.

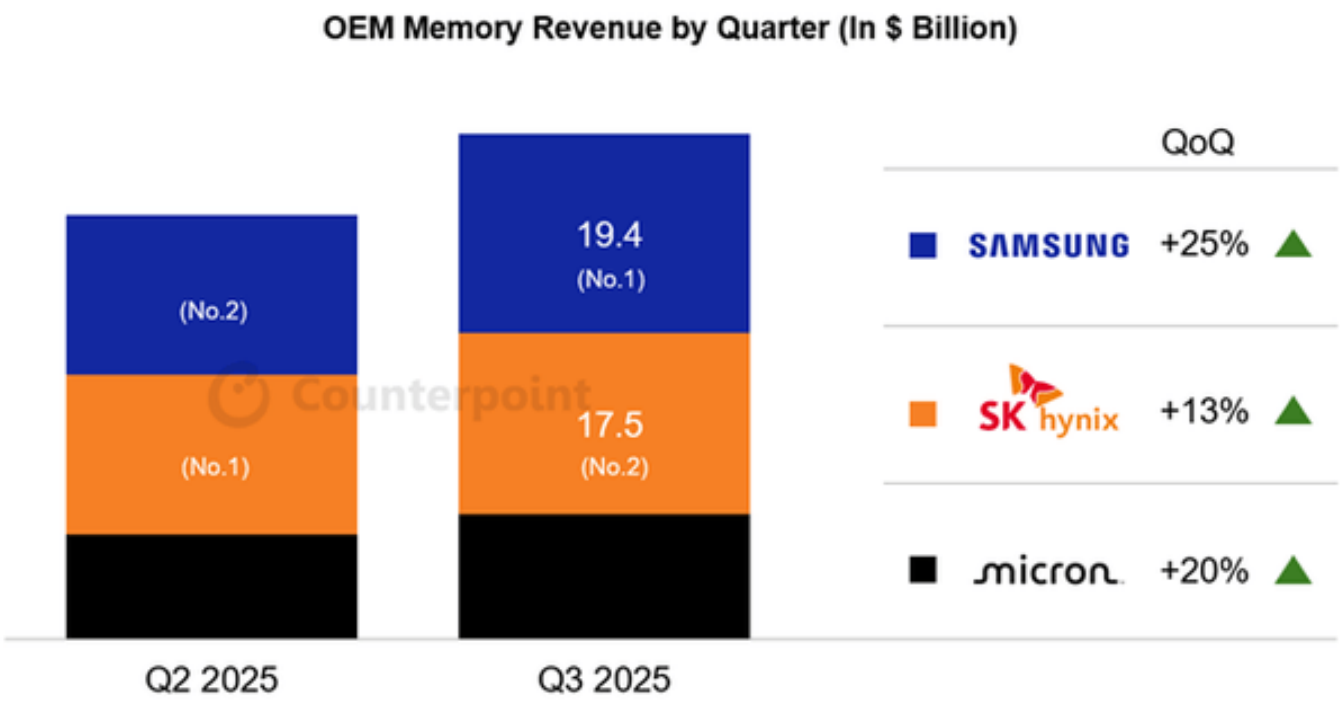

Counterpoint

After giving up the first position to SK hynix (OTCPK:HXSC.F) in Q2 2025, the company regained the leadership after growing 25% QoQ. The mid-term focus is on AI-related products, such as DDR5, LPDDR5x, and high-density QLC SSDs, along with further expansion and heavy investments in HBM3E and HBM4 memory nodes.

Samsung is a global leader in both of its major product categories: smartphones and memory. These two segments account for 68.23% of total revenues as of H1 FY2025. Samsung has managed to defend its leadership in both markets for a long time, staying ahead of the competition most of the time. The brand, level of innovation, in-house manufacturing capabilities, and strength across segments, together with its global leading position in two of them, present a moat. Now, the question remains if the fundamentals confirm it. This, along with determining a fair value of the business, will answer whether it's the right time to invest in the company.

Fundamentals

Samsung is a conservative and risk-averse company, as evident in its financials. The company focuses on several debt-related metrics, which appear in its presentations on a quarterly basis, and they have read as follows over the last five years:

2020	2021	2022	2023	2024	
Current Ratio	262%	248%	279%	259%	243%
Liability/Equity	37.1%	39.9%	26.4%	25.4%	27.9%
Debt/Equity	7.3%	6.0%	2.9%	3.5%	4.8%
Net Debt/Equity	-37.8%	-34.7%	-29.6%	-21.9%	-23.2%

Financial Ratios (Author: data source – SA)

The company keeps a high level of assets compared to its liabilities. These have also become a smaller part of the company's equity in recent years, which suggests faster asset growth than liability growth. Additionally, Samsung is in a comfortable position with a low debt burden and a high cash position. Besides that, the company effectively doesn't pay any interest on its debt since it records a net interest surplus.

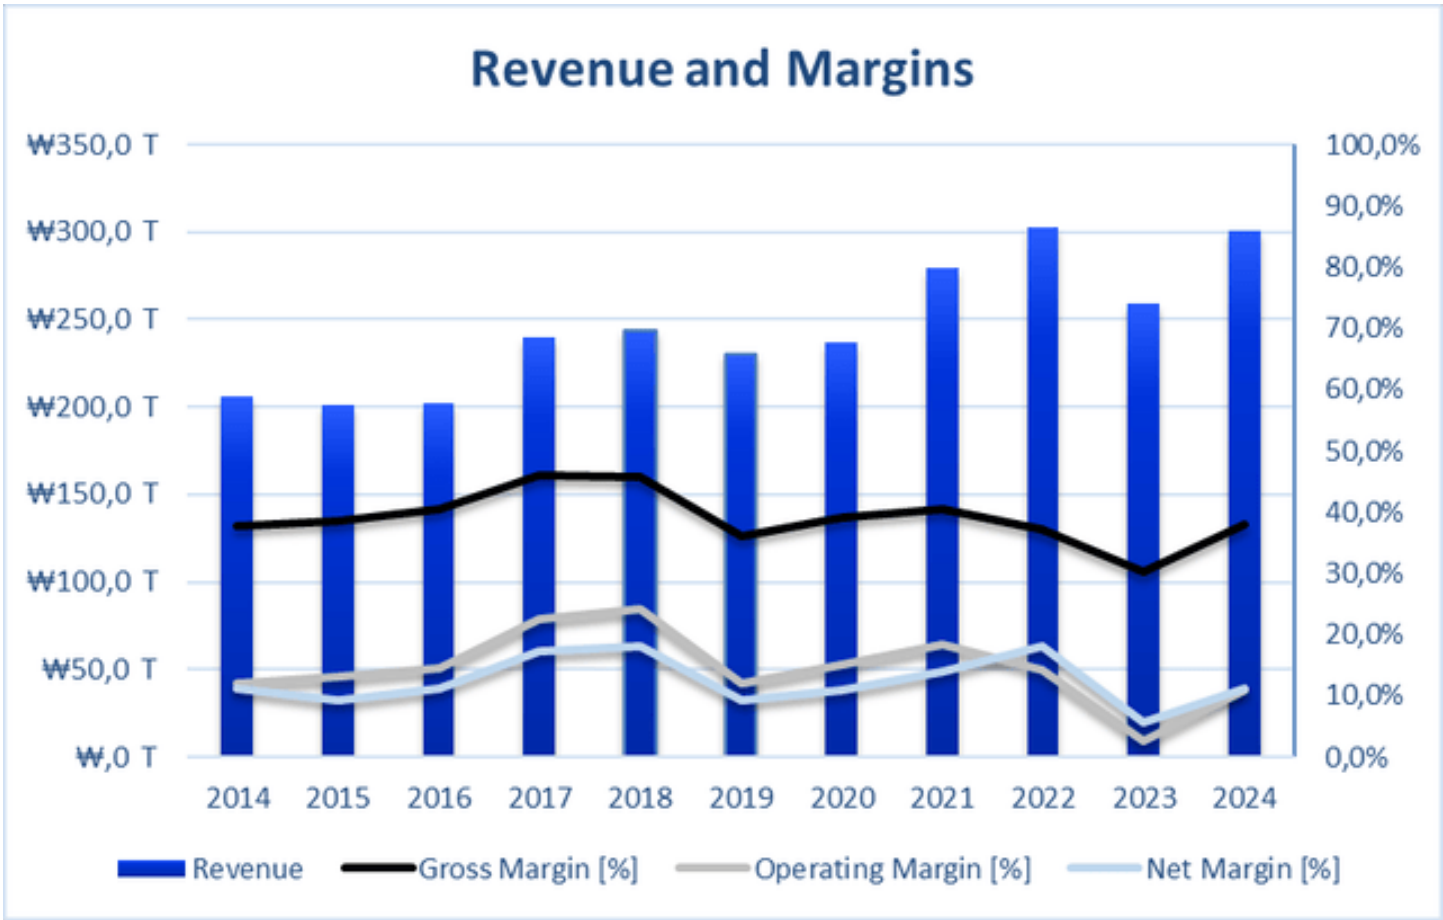

Author (Data: SA)

Accounting for cycles, revenues have been growing, while margins have fluctuated within a stable range. Profit and free cash flow also experience peaks and dips, reaffirming the cyclicality of the business. These pieces of information are indicative of the trend and how it may evolve in the future.

Valuation

To gain a better understanding of what a fair value of Samsung likely is, three scenarios based on a Discounted Earnings Model were created. Based on analyst estimates, the Korean manufacturer is supposed to report the following EPS numbers in the upcoming years:

2025	2026	2027	
EPS [₩]	5,616	9,824	10,778

EPS Estimates (Source: Simplywall.st)

Worst-Case Scenario

The forecast assumes that the current up cycle peaks in 2027 and will be followed by a mild down cycle, with smartphones expected to stabilize the results. While the semiconductor down cycle deepens in 2029, the DX

division should soften the drop in Samsung numbers. By 2031, the bottom is reached, and a recovery should begin. A new upcycle is expected in 2032, peaking in 2033 and normalizing in 2034. The EPS numbers in this scenario reflect the assumptions.

WORST CASE

SAMSUNG (SSNLF)				
Dicounted Earnings Model with Terminal Multiple				
Year	Earnings	Discounted Earnings	Growth Rate	Discount Rate
2024	₩4 949,8			
2025	₩5 616,0	₩5 105,5	11,2%	10%
2026	₩9 824,0	₩8 119,0		
2027	₩10 778,0	₩8 097,7		
2028	₩9 300,0	₩6 352,0		
2029	₩8 600,0	₩5 339,9		
2030	₩9 200,0	₩5 193,2	5,7%	
2031	₩10 000,0	₩5 131,6		
2032	₩11 300,0	₩5 271,5		
2033	₩12 000,0	₩5 089,2		
2034	₩11 500,0	₩4 433,7		
Terminal Value	₩103 500	₩39 904	Terminal Multiple	
			9,0	
INTRINSIC VALUE		₩98 037		

Auuthor (Data: SA)

A ten-year period for a cyclical business is extremely difficult to predict. These predictions should serve as a rough estimate for investors and shouldn't be taken too strictly. With that in mind, the assumptions give Samsung an 11.2% five-year EPS growth CAGR, followed by a 5.7% annual growth rate. At the end of the period, the market assigns Samsung a P/E ratio of 9, which, according to the father of value investing, Benjamin Graham, would mean almost no profit increase for the company. By discounting the earnings at a 10.0% rate, the fair value as of now is ₩98,037 ($67.88), which corresponds to the current stock price.

Normal-Case Scenario

This scenario assumes the AI's excitement to continue. Samsung should remain the leader in the smartphone business and also possess 40%-45% of the DRAM market share. The foundry business should improve, although it is still far behind TSMC. Expanding business opportunities and enhancing collaboration with customers should be a priority. Additionally, AI would drive demand across segments with momentum reinforced by HBM, DDR6, and LPDDR5X for servers and the Galaxy Flip model, among others.

NORMAL CASE

SAMSUNG (SSNLF)				
Dicounted Earnings Model with Terminal Multiple				
Year	Earnings	Discounted Earnings	Growth Rate	Discount Rate
2024	₩4 949,8			
2025	₩5 616,0	₩5 105,5	17,5%	10%
2026	₩9 824,0	₩8 119,0		
2027	₩10 778,0	₩8 097,7		
2028	₩10 300,0	₩7 035,0		
2029	₩10 700,0	₩6 643,9		
2030	₩11 800,0	₩6 660,8	9,9%	
2031	₩13 000,0	₩6 671,1		
2032	₩14 300,0	₩6 671,1		
2033	₩15 700,0	₩6 658,3		
2034	₩17 200,0	₩6 631,3		
Terminal Value	₩240 800	₩92 839	Terminal Multiple	
			14,0	
INTRINSIC VALUE		₩161 132		

Author (Data: SA)

With an annual growth rate of 17.5% over the next five years, followed by 9.9% in the subsequent five-year period, and a terminal multiple of 14, Samsung's fair value comes to ₩161,132 ($111.24). It's substantially higher than in the previous case, but the model contains more optimistic assumptions, which support this level of growth and higher valuation.

Best-Case Scenario

The most optimistic case includes more exuberance. Samsung is expected to catch up in HBM faster than expected as DDR6 and HBM4 become essential volume drivers. Additionally, the AI adoption for smartphones would accelerate. In the foundry business, key customers must be acquired to achieve a significant market share gain. Similarly, S.LSI and the mobile units would have to grow faster, likely propelled by an increased share of smartphone modems and imaging chips, as well as the sales of premium AI smartphones.

BEST CASE

SAMSUNG (SSNLF)				
Dicounted Earnings Model with Terminal Multiple				
Year	Earnings	Discounted Earnings	Growth Rate	Discount Rate
2024	₩4 949,8			
2025	₩5 616,0	₩5 105,5		
2026	₩9 824,0	₩8 119,0		
2027	₩10 778,0	₩8 097,7	21,4%	
2028	₩11 300,0	₩7 718,1		
2029	₩12 200,0	₩7 575,2		10%
2030	₩13 800,0	₩7 789,7		
2031	₩15 400,0	₩7 902,6		
2032	₩17 400,0	₩8 117,2	11,1%	
2033	₩19 200,0	₩8 142,7		
2034	₩21 000,0	₩8 096,4		
Terminal Value	₩420 000	₩161 928	Terminal Multiple	
			20,0	
INTRINSIC VALUE		₩238 592		

Author (Data: SA)

Assuming optimism at the end of the ten-year period, the market would apply a rich terminal multiple of 20. Samsung's rosy future would lead to a 21.4% growth CAGR, followed by an 11.1% annual growth. By discounting all upcoming earnings to their present value, Samsung's fair value comes to ₩238,592 ($164.71). It suggests a lot of room to run for the stock.

It's crucial to emphasize that these scenarios should help investors to get a grasp of possible paths for the company depending on the market. There are plenty of unknowns that may influence the results. A ten-year period is a very long time, especially for a company operating in such an industry across so many segments, each with its own unique dynamics. In hindsight, Samsung was heavily undervalued one year ago, trading at ₩56,700 ($40.82 at the KRW/USD exchange rate of 0.00072). Right now, it's arguably selling at the fair price if one believes conservative assumptions.

Risks

Samsung operates in a cyclical commodity industry where it's exposed to significant fluctuations in demand for its semiconductors. Besides that, the smartphone business also experiences cycles. It creates a dynamic and slightly unpredictable environment that investors should get used to if they want to ride this well-engineered, reliable Korean train. Yet, they must be prepared for sharp descents through the valleys before climbing back to the peaks.

Although Samsung's financials appear very healthy, one concerning trend is becoming increasingly visible, namely R&D expenses.

2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	
R&D / Revenue	6.8%	7.0%	6.8%	7.5%	8.6%	8.9%	8.0%	8.2%	10.9%	11.6%

Samsung's R&D/Revenue Ratio

That's the price Samsung has to pay to stay on top in the environment it operates in. It remains the leader in its crucial businesses, but to maintain its position, it needs more resources and a greater investment in its R&D activities. Recent years have been especially intensive, with costs growing from 8.2% in FY2022 to 11.6% in FY2024

Despite its leadership in key businesses, Samsung faces fierce competition. In the company's financial report for H1 2025, as well as for the full FY2024, a significant loss of market share across divisions can be observed.

Product	2022	2023	2024	H1 2025
TV	29.7%	30.1%	28.3%	28.9%
Smartphone	21.7%	19.7%	18.3%	19.9%
DRAM	43.1%	42.2%	41.5%	32.7%
Smartphone Panels	56.7%	50.1%	41.3%	39.9%
Digital Cockpits	17.9%	16.5%	12.5%	12.1%

Samsung's Share in Chosen Markets

In each of the five segments, Samsung was unable to maintain or increase its market share compared to FY2022. Every single one has experienced a drop, with Smartphone panels and DRAM suffering the most. Although the first constitutes only 8.0% of the total revenue, memory is a significant part of the whole business. A solution to these persistent declines might be even more investment in R&D, which creates a vicious circle.

Samsung doesn't provide an exact explanation for why its market share declines, but a frequently given argument explaining the weakness is the impact of global trade or even regional conflicts disrupting the markets. For a manufacturer like Samsung that operates globally, a rapidly changing playing field and the evolving distribution of power around the world disrupt the overall business. These factors influence demand, affect margins depending on geography, and, in effect, force the leadership to adapt and develop new strategies to dampen the impact of such developments.

Conclusion

Samsung is a great, diversified business standing at the forefront of several markets with its innovative solutions. It is the fifth top brand worldwide, taking into account revenue, presence, profit, awareness, and brand strength. The company fully participates in the AI race, delivering solid numbers. It's financially conservative, which results in a strong balance sheet, low debt level, and robust growth despite multiple challenges.

On a less positive note, it operates in cyclical markets that are highly competitive. This creates pressure, which can be seen in the accelerated growth of R&D expenses and a declining market share across multiple segments.

However, the company continues to focus on innovation, and its market position, combined with customer trust and brand recognition, gives it a competitive edge with a high chance of defending itself.

When considering conservative assumptions, Samsung seems to be fairly valued. Yet, a golden rule given by the legendary investor, Charlie Munger, applies:

A great business at a fair price is superior to a fair business at a great price.

For those looking for value, strong AI exposure, and geographical diversification, Samsung may be a compelling proposition. The Korean powerhouse may face challenging battles ahead, but it has everything it needs to win them and seal its global dominance.

Why Micron Remains One Of The Top Momentum Stock For 2026

Nov 18, 2025 Alpha Mantra

Summary

- Micron (MU) is positioned to benefit from a DRAM supply crunch driven by surging AI demand, supporting a continued boom phase into FY26.
- MU's stock has more than doubled in the past year, and a significant supply shortage is expected to persist until new capacity ramps up in FY27.
- Despite stretched valuations, MU remains a strong momentum play, with a buy rating justified by robust revenue and margin growth forecasts through FY26.
- Risks include the cyclical nature of memory markets and reliance on sustained AI-driven capex, but current trends favor continued upside for MU.

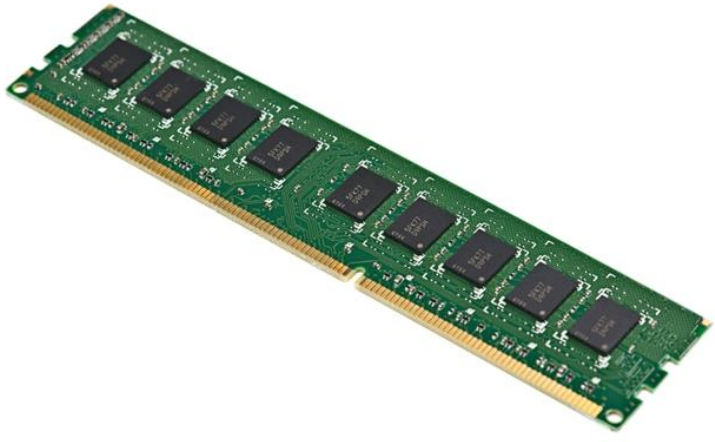

AWSeenbarab/iStock via Getty Images

Thesis

Micron (MU) is among the top three major DRAM manufacturers in the world, coming next to Samsung Electronics (OTCPK:SSNLF) and SK Hynix. The DRAM market is characterized by regular boom and bust

phases, and currently we are in a boom phase driven by AI related demand. This boom has turned around MU's financials completely from the bust phase in FY23-FY24. Consequently, MU's stock price has more than doubled over the last 12 months period.

SA

Now as we are in the boom phase, it becomes important to know how long this phase will last to know where MU's stock price is headed. If we look at industry supply growth and upcoming projects, it signals toward a sustained supply crunch extending into FY26 and should only moderate as we head into FY27.

While MU's valuations look stretched, I believe there is room for more upside from here, especially for investors who look for momentum stocks. I believe given the supply crunch extending in FY26 the stock could stay a momentum pick for a while. Hence, I have a buy rating on the stock.

What micron do?

MU is a global leader in advanced semiconductor memory and storage solutions focused on Dynamic Random Access Memory (DRAM), NAND (Not AND), and NOR (not OR) memories. These memories find application in electronics ranging from mobiles in your hand to data centers running AI models to operating systems of a Car.

MU segregates its business into 3 segments based on product type.

1. DRAM: More commonly known as RAM is the memory that is fast and temporary and used for computing in various devices.
2. NAND: The NAND is used in building fast persistent memory that is used for fast storage such as SSDs and flash drives used in computers and Mobile phones.
3. Other products: This segment is particularly composed of NOR and other specialty products. NOR is used to store operating system code for systems like automotives, consumer hardware and industrial machinery.

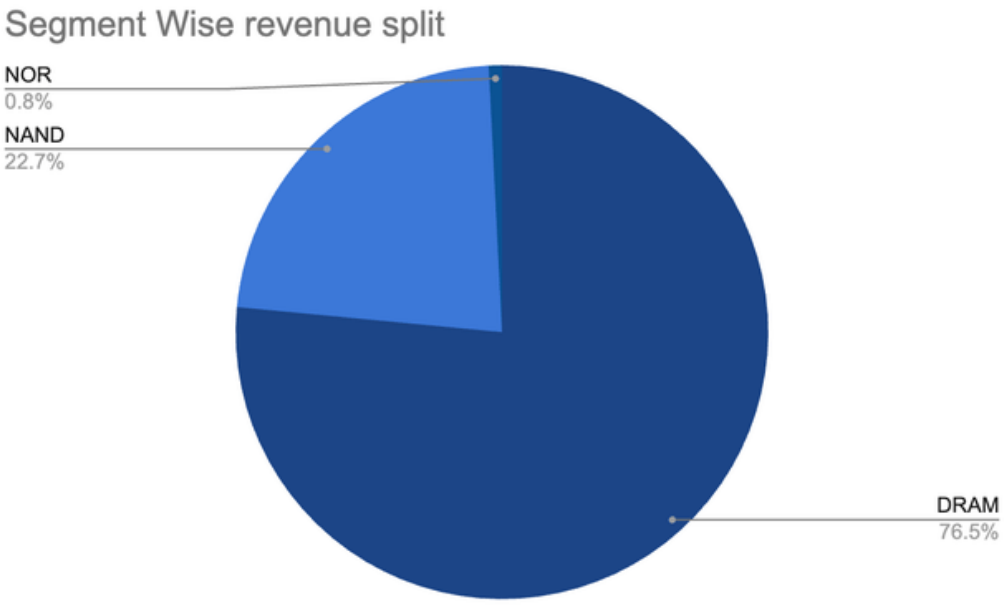

Micron's segment wise revenue as of FY25 (author, company filings)

Given MU's high exposure toward the DRAM segment and the recent supply crunch in this segment I would focus more on the DRAM category as it has been the primary driver for stock price gains recently and would majorly dictate the future direction of the stock price.

Industry Overview

Commoditized product and high producer concentration

The DRAM market is characterized by high client concentration along with the requirement to continually invest in newer technology to stay relevant and cost-competitive in the market.

The global DRAM market is highly concentrated, with the top three players controlling 95% of the market share. Beside MU the other two major players in the industry are Samsung Electronics and SK Hynix. MU stands out to be the only major manufacturer of DRAM and NAND not only in the US but also in the entire Western bloc. The high market share of the top three players creates a substantial producer concentration with significant barriers to entry due to high capital requirement in the industry as well as the need to constantly upgrade the technology to stay price competitive.

Similar to producer concentration, customer concentration is another key characteristic of this industry. The manufacturers sell their products to large corporations like data centers and mobile and PC brands. DRAM is a commoditized product with little differentiation in products offered by all the three players. New generations of products are easily replicated by other players with a lag of a few quarters. Furthermore, switching suppliers is quite easy for the clients.

The industry concentration is definitely a plus point, as it limits excess capacity in the industry, and understanding between players on pricing and capacity growth could be more measured. However, the concentrated purchasing power and low switching cost for clients keep all three major players on their toes with respect to pricing, as higher pricing than competitors might mean losing a large chunk of revenue.

Fierce Competition despite high producer concentration

Furthermore, the industry is also characterized by heavy capital investment and fixed-cost production, which further adds to the pressure to push for volume rather than pricing, making the competitive intensity even more fierce.

The production of memories like DRAM involves cutting-edge process technology, requiring leading-edge lithography, advanced materials, and MU/nanometer-scale precision. The capex required to design, build, and continually upgrade production facilities (fabrication units) is immense.

A single semiconductor fabrication unit can require billions of upfront investment and regular reinvestments to remain cost competitive. Losing a single customer could mean these high-cost facilities being vacant for several quarters if not years. This happening in a rapidly developing technology space could mean facilities remaining unutilized for much of their short useful life. This creates a constant tussle between players to keep their volumes as high as possible to cover fixed costs and benefit from operating leverage.

The competition among industry players for a concentrated client base and the need to utilize capacities create an intense competitive environment in the industry even though producer concentration is high. This intense competition gives rise to regular boom and bust phases where over capacity and under capacity relative to demand create a cycle of higher profitability and lower profitability periods for the players.

Latest boom phase driven by AI demand

Many of you might already know that AI has become the tech story of the decade. What's less obvious to many is that running and training AI needs more memory than the conventional computing use cases. This need for memory has triggered a new boom phase in the memory market, particularly in the DRAM spaces.

The need for faster memory

Think of ChatGPT or any SOTA (state of the art) large language model; these models require a significant amount of memory for them to perform computation and generate responses. Most of the modern SOTA LLMs require more than 1TB of memory just to run computation in an efficient manner. Even more so, this memory needs to be fast enough to generate a response that has an acceptable token speed.

Traditional GDDR memory, which is more commonly used in consumer-grade or gaming graphic cards, is too slow (i.e., low bandwidth) to operate such models at an acceptable speed. Therefore the memory industry came up with a new architecture of DRAM called High Bandwidth Memory (HBM). As the name suggests, HBM has bandwidth significantly higher than the traditional GDDR memory, which comes quite handy in LLM inference and training. These memories are built by stacking memory chips on top of each other, creating a 3D structure for the memory, which is stacked near to the GPU or often on the same silicon interposer. The 3D structure and lesser distance from the GPU reduce latency, thereby leading to faster computation and lesser power consumption and heat.

So the need for fast computation and growth in AI usage has become a new growth driver for the DRAM memory market, more specifically the HBM market. The chart below showcases the significant revenue growth in the MU's DRAM segment which also contains HBM revenue.

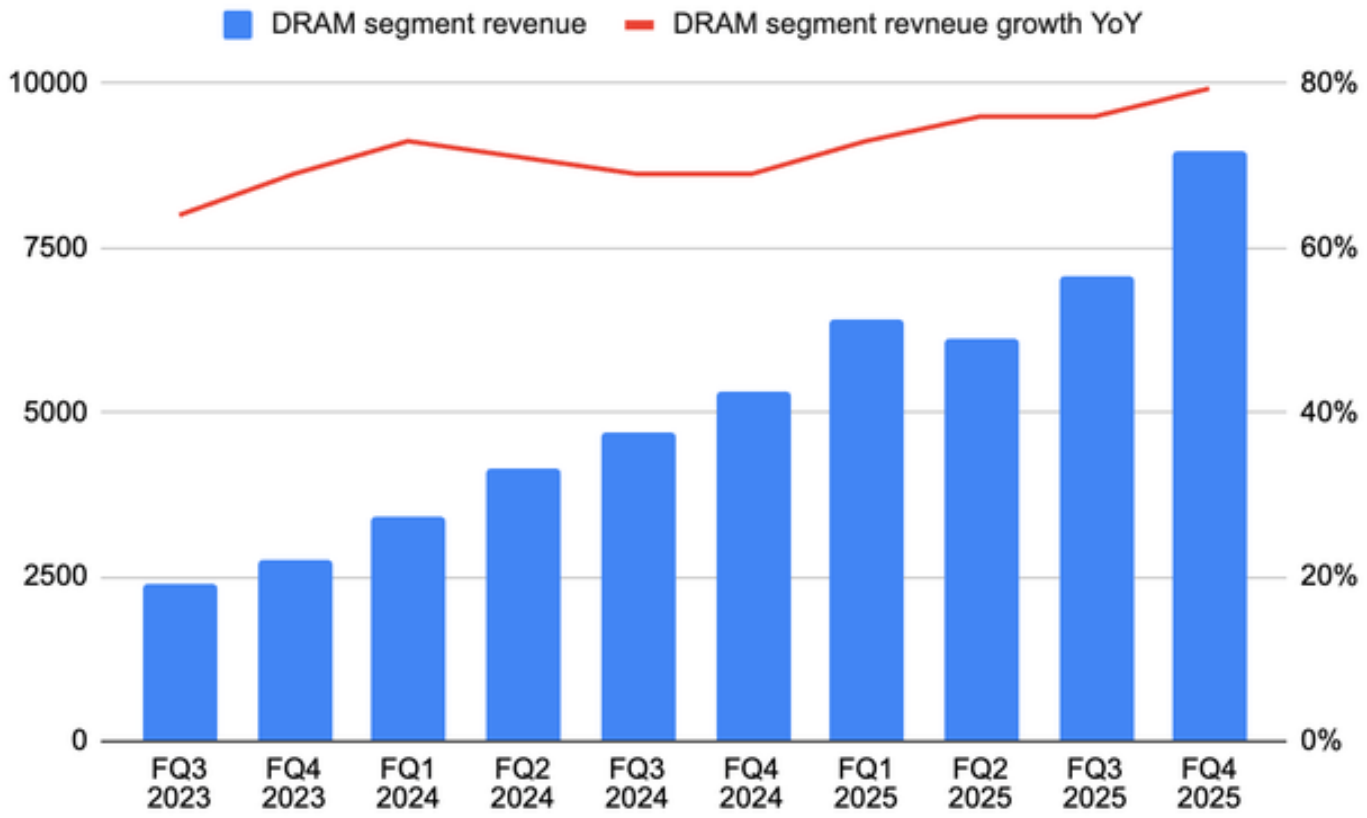

Author, Company filings

Manufacturing fungibility between DRAM and HBM

Fundamentally there are few differences between traditional DRAM and an HBM which also happens to be a type of DRAM. The only critical difference between them is that HBM uses a 3D stacking process where multiple memory dies are stacked vertically and connected using through-silicon vias (TSVs) enabling lower latency. On the other hand, traditional DRAMs are typically used 2D layout where chips are placed side-by-side on a circuit board resulting in a much higher latency.

Both traditional DRAM and HBM share the same base DRAM fabrication facilities, and the only difference in the manufacturing process lies in the post-wafer processing units. Usually the base DRAM fabrication facilities cost much more than the post-wafer processing units in terms of capex. So to increase the output for HBM manufacturers just need to ramp up the post-wafer processing units and utilize more of the base DRAM fabrication unit toward HBM manufacturing.

This shift toward HBM due to the substantial ramp-up in AI demand has also created a short supply in the DRAM market leading to significant price surge in recent months.

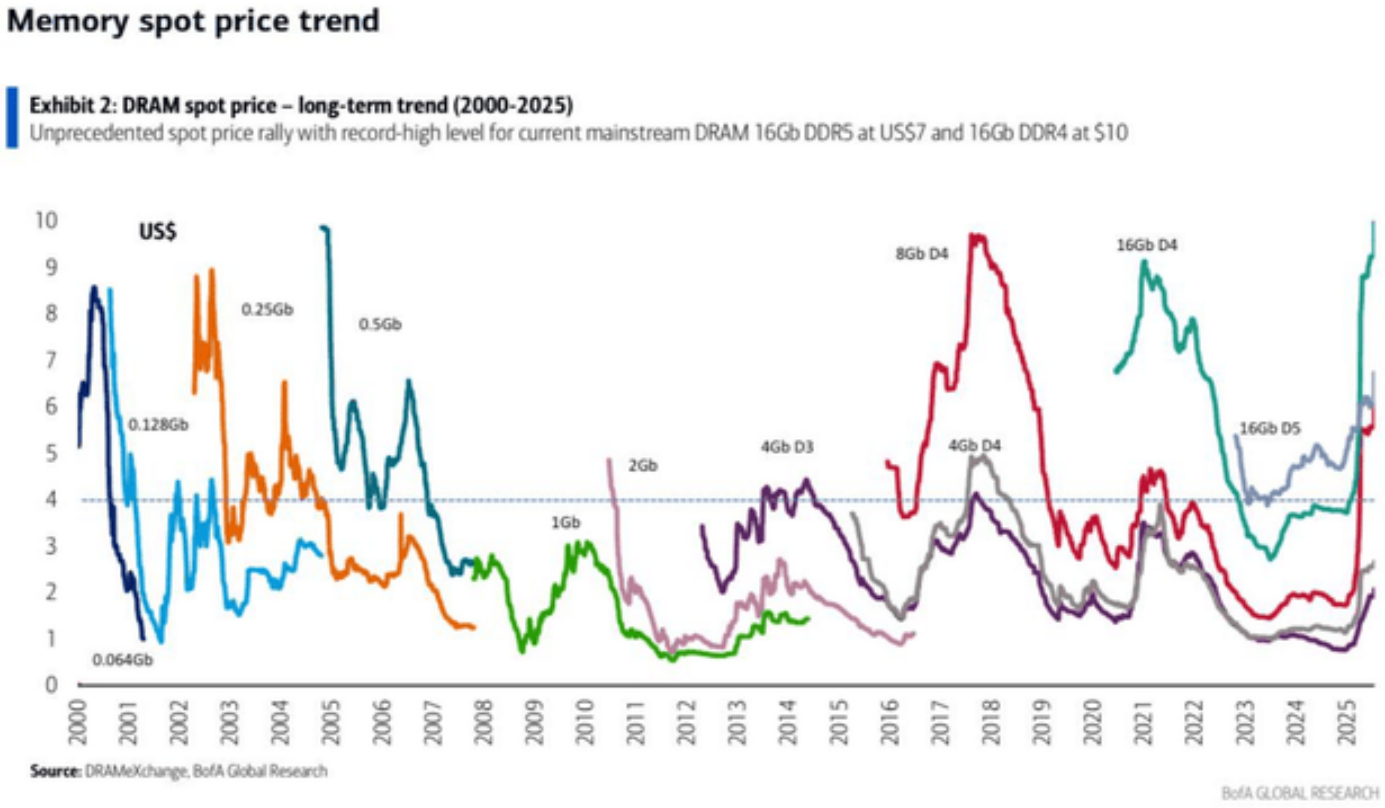

Bofa Global

With regard to the NAND market, the AI related tailwind is much more muted, but it still has driven a sizable demand growth. While HBM is used in active computation for AI inference and training, NAND used in SSDs plays an important role, like storing datasets on which AI would be trained or storing AI models in persistent storage.

Significant DRAM shortage should persist in FY26

As mentioned earlier, there are three main players in the memory market (DRAM): Samsung electronics, SK Hynix, and MU. All three players are undergoing significant capexes in the DRAM space, particularly toward HBM.

With respect to MU the company has planned to invest $4.5 billion per quarter, or $18 billion in FY2026 (September-ending FY). A substantial amount of this capex is being directed toward DRAM and HBM capacity expansion. Interestingly, as presented in the table below, none of these capex, except the Manassas modernization project, would be completed in FY26.

MU's major planned capex

Project/Facility	Capex Amount	Timeline & Ramp
Singapore	$7 billion by 2030	Significant ramp-up in FY27 and beyond
Idaho Fab 1 & 2 (Boise)	$50 billion by 2030	Fab 1: prod FY2027; Fab 2: post-FY2028

New York Megafabs	$20 billion by 2030	Production in FY29-FY30
Manassas, VA modernization	$3 billion by 2030	Ramp up in FY27

Author, company filings/public sources

Samsung Electronics, being the largest player in the industry, is also planning to substantially increase its capex, particularly for the HBM and DRAM segments. While the total capital outlay for 2026 is not announced by the company, the company did specifically announce to significantly ramp up capex from ~$41 billion in 2025. Specifically for the memory segment, the company is planning to ramp up capex, particularly to serve the growing demand for HBM.

Samsung Electronic capex for memory segment

Project/Facility	Capex Amount	Timeline & Ramp
Pyeongtaek Campus P5	$22 billion for P5 complex containing memory facility and foundry	P5 ramp up in post-2028
Pyeongtaek Campus P4	-	phase 4 construction started in mid-2025 (my estimate; it might take 2 years to complete)

Author, company filings/public sources

Similar to Samsung Electronics, SK Hynix also did not provide actual guidance for FY26 capex. Instead, they have provided a comment regarding scaling up capex YoY from YTD $16-17 billion of capex.

SK Hynix major memory capex projects

Project/Facility	Capex Amount	Timeline & Ramp
Expansion at M15X Fab (Cheongju, Korea)	$3.6 billion	Complete by the end of 2025 and ramp up in 2026
Yongin Cluster Fab 1 (Korea)	$6.6 billion	1H FY27
Indiana Advanced Package Plant (U.S.)	$3.87 billion	2H FY28

Author, company filings/public sources

Carefully examining all the major projects of all three players, one trend was clearly visible: there is a significant ramp-up of capacities starting from FY27. And there are no major investment projects except SK Hynix's M15x project, which is projected to compete and ramp up in 2026. In fact, industry players like Samsung had mentioned that they see a supply shortfall in 2026.

With respect to demand, there is almost a consensus among all three industry participants that the demand is quite strong, with the AI demand being the driving force. Managements have been forecasting substantial growth in the HBM market figures; for instance, SK Hynix's management mentioned that the HBM growth could be ~30% YoY till 2030. Meanwhile, MU's management also upped their demand estimate from the DRAM market to the high teens.

While players are taking necessary steps to increase their production from existing facilities to cater to the ramp-up in demand, I believe 2026 is going to be a major year in terms of supply shortage in the DRAM market. Furthermore, the inventory levels in the industry are now at lower levels.

The shortage is clearly visible in recent price increases, and it continues to get even better for DRAM manufacturers. The DRAM chip spot prices increased by up to 20% in the first week of November. Prominent brokerage houses and industry researchers such as Citi and trendforce are also expecting a memory shortage in 2026.

Valuations

Given the current and persisting supply shortage in the DRAM market, MU, at a non-GAAP forward PE ratio of just 14.6x, looks quite cheap. However, as mentioned earlier, the players in the memory market, such as MU, tend to be heavy on fixed costs, making them highly cyclical. The problem with cyclical stocks is that their valuation multiple compresses significantly in up cycles and expands in cycle bottoms.

So a deeper analysis of the valuation metric is necessary to gauge a correct rough estimate of MU's fair value.

Below are my fair value estimates based on my assumptions on how the future is going to unfold for MU. Based on the assumptions, I find MU stock price does offer a decent upside from the current price.

In millions (except EPS)	2023	2024	2025	2026	2027
Revenue	15540	25111	37378	52329	62795
Gross profit	415	5613	14873	31398	31398
Operating income	-3574	1305	9809	24071	21978
Net Income	-5833	778	8539	21978	20722
Diluted EPS	-5.34	0.7	7.59	19.11	17.64
Assumption					
Revenue growth	-50%	62%	49%	40%	20%
Gross margin	3%	22%	40%	60%	50%
Operating margin	-23%	5%	26%	46%	35%
Net Margin	-38%	3%	23%	42%	33%
Diluted Share count	1093	1118	1125	1150	1175
Implied Stock price				**287**	**300**
PE multiple				15	17

Author, SA

Assumptions

Broadly, I have assumed a continued momentum in fundamentals going through FY26 primarily due to my assumption that the DRAM shortage is going to persist in FY26 with a supply increase starting from FY27. Given this, I have assumed a significant revenue growth of ~40% for FY26 with a cool-down in FY27.

With respect to margins, my assumption might look a little optimistic at first, but if you look at the history, especially the boom phase in FY16-FY17 these numbers do look fair. I have assumed gross margins to be in the range of 60%, operating margin ~46%, and net margin of ~42% for FY26 and then a cool down in FY27. The FY26 margin figures are almost in line with the FY17 boom peak. I believe the current boom cycle is far more fundamentally strong with the underinvestment by the players during the FY23-FY24 bust phase and subsequent unprecedented surge in demand due to AI. I believe there is a fair chance that MU can surpass the FY17 peak margins.

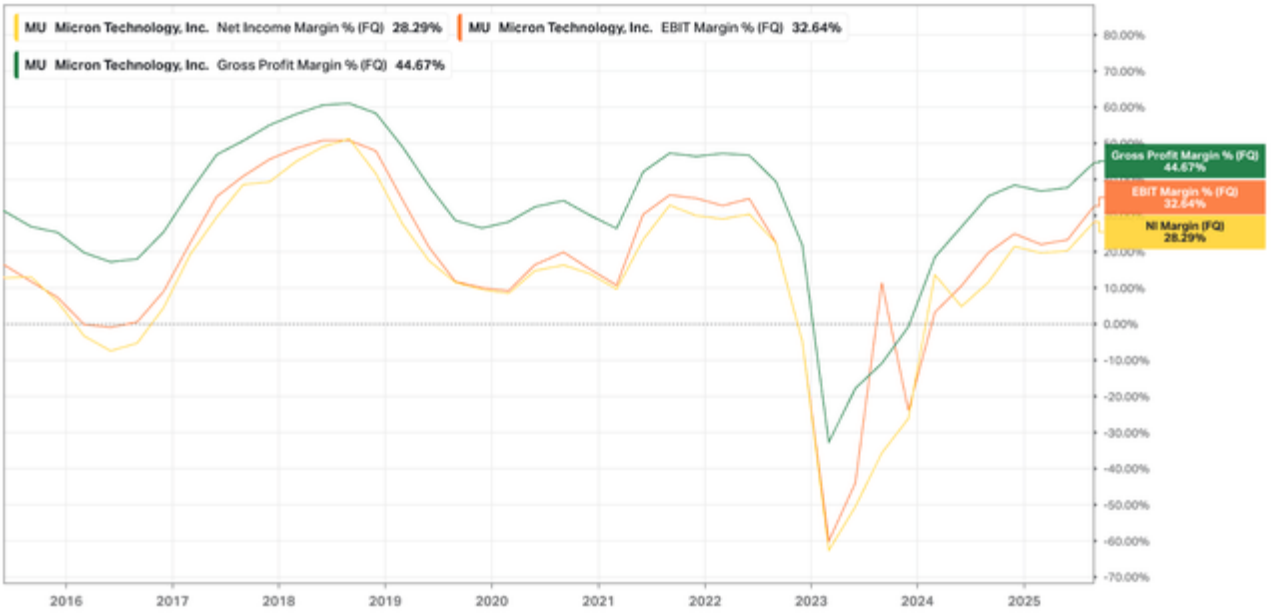

Koyfin

For valuation multiple, I am using a 15x multiple for peak earnings and an expanded 17x multiple for more normalized earnings. I believe these multiples to be fair, giving a decent premium primarily due to the extended AI related demand trend extending well beyond the forecast period.

I have also prepared a scenario analysis for fair value estimates for those who are wondering how high or low the fair value is given optimistic or pessimistic scenarios.

Bull case	**FY23**	**FY24**	**FY25**	**FY26**	**FY27**
Revenue growth	-50%	62%	49%	50%	30%
Gross margin	3%	22%	40%	65%	55%
Operating margin	-23%	5%	26%	51%	40%
Net Margin	-38%	3%	23%	47%	38%
Diluted Share count	1093	1118	1125	1150	1175
Implied Stock price				**458**	**542**
PE multiple SA				20	23

Bear Case	**FY23**	**FY24**	**FY25**	**FY26**	**FY27**
Revenue growth	-50%	62%	49%	30%	20%
Gross margin	3%	22%	40%	55%	45%
Operating margin	-23%	5%	26%	41%	30%
Net Margin	-38%	3%	23%	37%	27%
Diluted Share count	1093	1118	1125	1150	1175
Implied Stock price				188	201
PE multiple SA				12	15

Another method to value MU is through price to tangible book value. The company currently trades at a ~5.1x P/TBV ratio, which is higher than at any time in the last 20 years. In fact, these multiples were only seen back

during the dot-com boom. While the P/TBV does look elevated, if the demand boom in the AI industry is going to be sustained well into the 2030s, I believe the current investment cycle could become somewhere comparable to the dotcom boom in the 90s. This continued investment cycle would create a sustained period of supernormal profits for memory manufacturers like MU so in that case, the current P/TBV could be well justified.

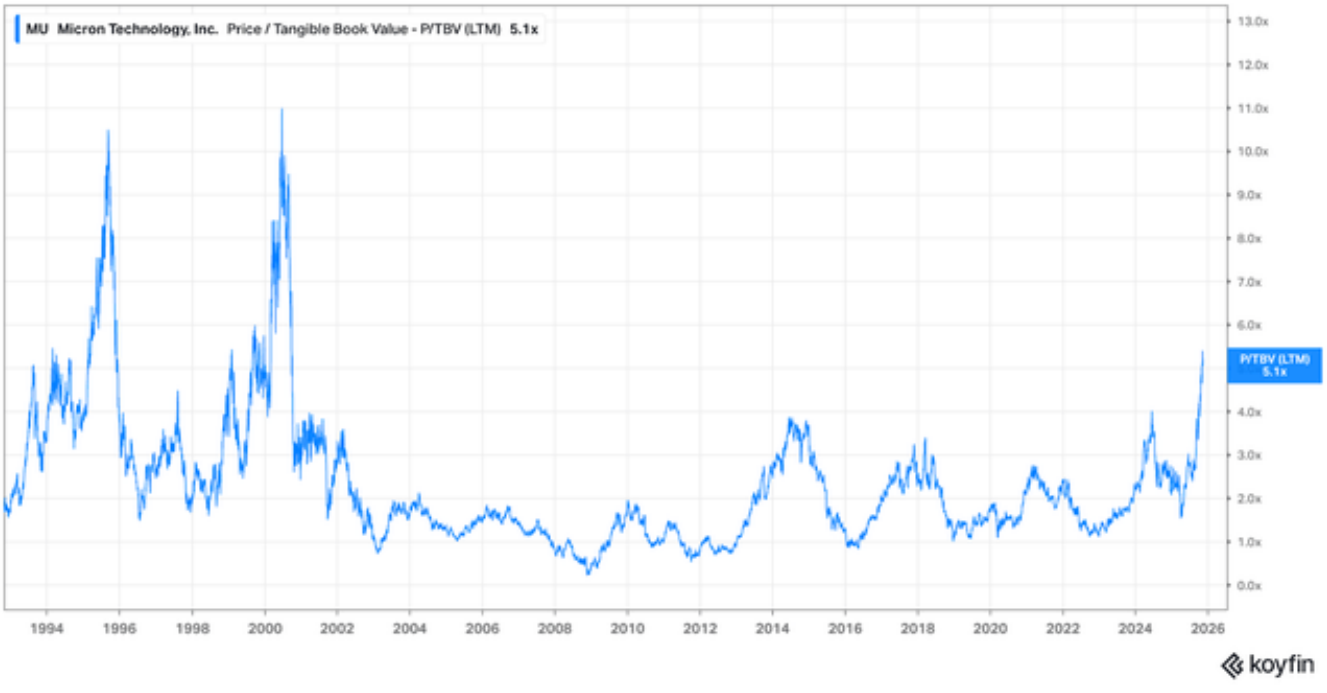

Koyfin

Risk to my thesis

MU is exposed to the cyclical DRAM and NAND market and goes through a constant boom and bust phase. While the NAND market is going through a normal demand phase, the DRAM segment in particular is witnessing a substantial surge in demand due to the AI related demand surge.

Currently this demand surge is well rooted, with capex being undertaken by large AI labs and cloud providers. However, a significant slowdown in future capex could weaken the demand outlook for the DRAM market, leading to a potential for supply to overshoot demand, which could lead to a bust phase.

My thesis's primary assumption is continued momentum in the DRAM market and an extended supply shortage of these products. If supply overshoots demand, my fair value model would need a significant readjustment to the downside.

Given the current trend in investments, a significant slowdown in investments does seem unlikely. However, one needs to keep an eye on the big tech capex spending and commentary regarding it.

Western Digital And Seagate: Race To The Top, Progression And Navigating The HDD Market

Nov 13, 2025. Jia Ming Eow

Summary

- Seagate and Western Digital are seeing increased capex from data centers as HDDs remain essential for massive storage demand.
- The HDD market rallied as Seagate and Western Digital were the top performers in the S&P 500, but valuations have skyrocketed as both names trade at a premium to historical averages.
- Seagate is launching new products rapidly, raising questions about whether WDC is falling behind in technology. However, Western Digital remains competitive with Seagate's supply constraints.
- I remain bullish in the HDD market but remain neutral, as there may be an inflection point where either firm may have a majority market share.

miklyxa13/istock via getty images

Seagate (STX) and Western Digital (WDC) are 2 of the major players when it comes to data center storage, as they compete to provide the most advanced HDDs in the market to fulfill the huge buildout of AI infrastructure today. As capex continues to rapidly increase in recent years, Seagate and Western Digital have seen a surge in demand and a resurgence in HDDs.

Price of Seagate and Western Digital (SA)

Both stocks have been relatively flat and dead money for a few years, but since the April dip this year, Seagate has had its stock prices over 4x, while Western Digital has had its over 5x, as investors realise the low valuations both names had.

This article aims to highlight and navigate the current HDD market following both Seagate and Western Digital's Q2 FY26 earnings reports. This is a follow-up and progression from my prior articles on Seagate and Western Digital, which highlighted their differences and the race for volume shipping 36 TB HAMR HDDs. I believe having both in a single article rather than separate would be way better to understand the HDD market, which had significantly outperformed my expectations. The question now is, can the stock price continue to still outperform despite the huge run-up? There is no doubt that the HDD market is doing well, with both firms raising guidance, but the stock has surged massively over the last 7 months. Fundamentally, the company might be doing great, but the stock is overbought as investors may be too optimistic.

I rated both Seagate and Western Digital at a 'Hold' in prior articles, as both stocks have now significantly outperformed the S&P 500. Western Digital and Seagate were also among the top 3 stocks in the S&P 500 for their price returns year to date. With valuations now skyrocketing and at a huge premium to their relative 5Y values, I continue to rate both at a 'Hold' with their recent run-up as they trade at all-time highs and at overbought values on the RSI scale.

HDDs As Top S&P 500 Performers

As seen within the S&P 500, Western Digital and Seagate have as much as 5x in stock price as they are the top 3 best S&P 500 performers, with Western Digital in first and Seagate in third. The market has realised that demand for HDDs is surging as the legacy product has a use case in this AI boom. The AI infrastructure buildout is clearly not slowing down yet, such as top hyperscaler Amazon expecting to spend $125B this year alone in capex, which was increased from the prior quarter, which simply shows the general trend from top AI players (increased capex and capacity).

HDD demand continues to surge, as both Seagate and Western Digital raised guidance in the most recent quarter, which is very bullish for both names. With more supply constraints and higher demand, investors simply realised that the HDD market was undervalued. The main takeaway from this segment is that, fundamentally, both Seagate and Western Digital are seeing demand they've never seen before, with trillions spent on AI infrastructure. This is reflected in the stock price, as the future of the HDD market is very bullish. Higher margins, raising guidance, and supply constraints are happening for both names, so fundamentally, both are doing amazing as the demand for HDDs will only increase as capex increases.

Why HAMR Is Important (Recap)

A quick recap of what they are racing for: for anyone that may not be caught up or not read the prior articles, both Seagate and Western Digital are racing to volume ship heat-assisted magnetic recording (HAMR) storage, which is the future of HDDs and expected to be the norm. Seagate notes HAMR has benefits such as higher areal density, high-efficiency and power efficiency, and more sustainable.

When it comes to comparing areal density, the higher the areal density per platter, the better. From the Seagate lab tests (link above), HAMR could have as high as 6TB per platter whilst Western Digital's UltraSMR version only offers 2.91 TB per platter for their 32TB SMR HDD. Since the storage space will continue to increase over the years, more platters will be required for HDDs, which increases the risk of failure. With more platters, it

becomes more complex, more points of failure and hence less reliable, as latency is higher since there are more platters to move between.

Fewer platters mean less power consumption and heat, which results in higher efficiency. Since HAMR has fewer platters and will likely be cheaper to produce, it will be the preferred option for cloud providers as operational costs will be lower. From a business perspective that aims to build out billions worth of AI infrastructure, HAMR will be the better option in almost any way.

Race To The Top (Update)

In the current HDD market, and all noted from my prior articles, Seagate has the most advanced HAMR HDD and is rolling it out, which Western Digital aims to ship in 2027. With Seagate launching its 36TB HAMR HDD in January 2025, this means that Seagate is currently in front when it comes to shipping HAMR. However for Seagate in Q2, their average shipped capacity for most advanced HDDs was lower than Western Digital (calculations all in the last Seagate article), which means even though they have start shipping the most desired HDD products, Western Digital's customers are confident in their product growth map and is happy to settle on their UltraSMR HDDs rather than being on HAMR, which has huge demand. The thing to watch here is whether Seagate can flip the script and win over more contracts, as it seems that Western Digital has the contracts with major hyperscalers. The HDD market is doing exceptionally well, but within the HDD market, Seagate has the top product but has yet to widely ship, while Western Digital has the customer base that is willing to wait for Western Digital to launch their HAMR HDDs. Any of the next quarters could be the inflection point for Seagate, but how are they doing this quarter?

Comparing capacities shipped:

Quarterly Financial Trends Continued

	Q1FY25	Q2FY25	Q3FY25	Q4FY25	Q1FY26	QoQ	YoY
Revenue by End Market ($M)	**2,168**	**2,325**	**2,160**	**2,444**	**2,629**	**8%**	**21%**
Data Center[1]	1,582	1,736	1,615	1,863	2,114	13%	34%
As a % of total revenue	*73%*	*75%*	*75%*	*76%*	*80%*	4 ppt	7 ppt
Edge IoT[1]	585	589	545	581	515	-11%	-12%
As a % of total revenue	*27%*	*25%*	*25%*	*24%*	*20%*	-4 ppt	-7 ppt
HDD Exabyte Capacity Shipped	**138**	**151**	**144**	**162**	**182**	12%	32%
Nearline	114	126	120	137	159	17%	39%
Non - Nearline	23	25	24	26	22	-14%	-4%

Nearline HDDs Exabytes Shipped (Seagate Investor Relations)

(Millions, except exabytes)

	Q1FY25	Q2FY25	Q3FY25	Q4FY25	Q1FY26
Revenue by End Market					
Cloud	$1,909	$2,096	$2,007	$2,329	$2,510
Client	139	140	137	140	146
Consumer	164	173	150	136	162
Total Revenue	**$2,212**	**$2,409**	**$2,294**	**$2,605**	**$2,818**
Exabytes Shipped					
Nearline	141	154	145	170	183
Non-Nearline	24	22	21	20	21
Total Exabytes	**165**	**176**	**166**	**190**	**204**

Nearline HDD Exabytes Shipped (Western Digital Investor Relations)

Western Digital continues to lead the HDD market when it comes to exabytes shipped. Seagate has 159 exabytes shipped, while Western Digital has 183 exabytes shipped for nearline HDDs. Seagate grew nearline HDDs by 45 exabytes and 39.4% YoY, Western Digital grew nearline HDDs by 42 exabytes and 29.7% YoY.

Competition is getting higher and higher within the HDD market; Seagate and Western Digital are essentially neck and neck with each other as they both have orders going all the way until the end of 2026. From their earnings calls:

(Western Digital)

Our top 7 customers have now provided purchase orders extending throughout the first half of calendar year 2026. And 5 of them have provided purchase orders covering all of calendar year 2026. I'm also pleased to share that one of our largest hyperscale customers has signed an agreement covering all of calendar year 2027. These commitments underscore both essential role of our products in the AI data economy and our customers' strong confidence in our product road map, including the transition to HAMR technology.

(Seagate)

The data center end market, which is comprised of nearline sales into cloud, enterprise and VIA customers represented 80% of overall revenue. Amid this improving demand backdrop, our high capacity nearline production is largely committed under build-to-order contracts through calendar 2026.

Seagate is growing faster than Western Digital when it comes to exabytes shipped, but it continues to lag behind Western Digital despite shipping HAMR HDDs, while Western Digital is still focusing on UltraSMR and HAMR in 2027. Both have the contracts and demand over the next calendar year, but the main question is whether Seagate can overtake Western Digital since they are shipping HAMR HDDs first. So what's the problem?

Product Growth Map

Recap and update on the product roadmap for both Seagate and Western Digital.

Seagate is trying to flip the script and win market share over Western Digital by having the most desired product in the market. Seagate currently ships 36TB HAMR HDDs, with a target of 44TB HAMR HDDs to volume ramp in 1H 2026. These are all going as planned from previous quarters. From Seagate's earnings call:

These products are performing well in live production environments, and we are on pace to achieve 50% exabyte crossover on nearline HAMR drives in the second half of calendar 2026, and we started qualification with a second major CSP on the Mozaic 4+ terabyte per disk platform, with initial volume ramp starting in the first half of next calendar year. This platform will offer capacities of up to 44 terabytes.

Western Digital has the contracts and trust from cloud providers in their product roadmap. Their HAMR qualification starts in 1H 2026, 3 customers by 2026, and ramp up volume production at the start of 2027. It is currently focused on 36TB UltraSMR. From Western Digital's earnings call:

On HAMR, as you mentioned, we also pulled forward the qualification process by half year. As we've highlighted in our road map in the past, the plan was to start HAMR qualification in the second half of calendar year 2026. We've now pulled that in into the first half of calendar year 2026 with one customer and we look to expand that to up to 3 customers by the end of the calendar year.

Potential Inflection Point

Is now all about whether Seagate can hit the inflection point and attempt to overtake Western Digital as it lags behind in technology. Western Digital is also executing well, with their HAMR qualifications pulled forward by 2 quarters, as mentioned above from the earnings call.

The reason why Seagate isn't dominating the HDD market is due to supply. There are plenty of statements that constantly signal that demand is growing, but supply is constrained. From Seagate's earnings call:

Yes. No, we see actually demand -- the gap between supply and demand getting a little bit bigger every quarter, that means demand is shifting more into the future, is not taken by any other technology.

We've been taking some supply out of Edge IoT products and putting it into cloud as we can pivot demand.

Some of the numbers about how much demand there is above and beyond what our supply is.

Perfectly summarised from a report, Seagate has no chance to meet demand with its HDD supply over the next 4 quarters. There isn't enough manufacturing capacity to meet requirements, as it currently does not intend to add manufacturing capacity but aims to target the demand through areal density gains.

Western Digital is able to benefit from the fact that the demand for HDDs and the huge buildout of AI infrastructure strangle Seagate's supply, so customers are willing to settle for volume shipped of Western Digital's UltraSMR HDDs and wait for volume shipment of 36TB HAMR HDDs from Western Digital since Seagate's demand is way higher than its supply. Even though Seagate is aiming to ship 44TB HAMR HDDs when Western Digital trials out their 36TB HAMR HDDs with customers (Seagate has 5 global cloud providers already qualified on HAMR HDDs), Seagate has far too much demand to even think about having the ability to fulfill supply. This refers back to my old Seagate article, where used 36TB HAMR HDDs from Seagate are resold at similar levels to brand new ones since supply is so constrained.

If Seagate can solve its supply problem somehow before Western Digital volume ships HAMR HDDs (next 5 quarters), the competition between Seagate and Western Digital could flip, and Seagate could have a majority market share, as they currently both have similar levels of the HDD market.

If Western Digital is able to continue to execute its HAMR HDDs according to plan, with now even ahead of schedule, it will simply lead to Western Digital maintaining its position in the HDD race, which will likely result in similar market share for both names. From my old Western Digital article, their old CFO promised limited

volume shipping for HAMR HDDs back in 2023, 6 months ago (CFO has resigned after the spinoff), so Western Digital being able to execute well and pull forward their shipping is very bullish for Western Digital.

All eyes are on Western Digital; expectations for Seagate are already set, and they are currently in the technology lead. Western Digital continues to have the customer base that trusts in their product roadmap, if they don't execute with expectations for HAMR volume shipping, Western Digital could lose market share, assuming Seagate would've solved their volume shipping problem for 36TB HAMR HDDs by 2027 (they already ship it now, so I'm assuming they'll solve their volume shipping problem as they have a whole year).

Western Digital is playing catch-up when it comes to technology, but with their existing customer base and still shipping more exabytes than Seagate, it is unlikely that the inflection point (if there is one) is happening in the short term, but investors who have both Seagate and Western Digital should understand that the HDD market is doing well, but the competition between both is quite tight. If there is no inflection point, as Seagate continues to have the most advanced technology with constrained supply and Western Digital has the customer base, both stocks will likely continue to benefit from the huge AI infrastructure buildout.

Valuation

Comparing P/E with Seagate and Western Digital (SA)

Comparing P/S with Seagate and Western Digital (SA)

Western Digital continues to trade at lower valuations than Seagate, as investors are putting a lower premium on Western Digital, as Seagate has the better technology. If Western Digital trades closer to Seagate's on a P/E GAAP basis, there could be a 50% upside from Western Digital's level to Seagate's level. On a P/S basis, though, the upside is rather limited, only at a 9.5% upside. To also keep in mind, Western Digital and Seagate's 5Y historical average for P/S is 1.3 and 2.28. With the current levels at 6.5 and 5.9, it is already trading as much as triple its historical averages.

Western Digital Relative Strength Index (Trading View)

Seagate Relative Strength Index (Trading View)

On a 1-year scale on the RSI scale, both Western Digital and Seagate are around/at the overbought level. Investors should be cautious with the huge recent run up in stock price and question whether the HDD market is currently overvalued. In April, Western Digital traded as low as 0.67 P/S with 14.2 P/E. Is hard to tell if the HDD market is actually overvalued, as there is no doubt they have a key role in this AI boom, as AI stocks have high valuations, but the question is how important they actually are.

Most signals suggest that the HDD market is overvalued, especially Seagate, which trades at a premium to Western Digital, but investors need to remember that they have more demand than they can currently supply. For reference, TSMC, which has way more demand than supply, trades at a P/S of around 10. TSMC obviously has a clear moat, but the idea remains that these companies trade at high valuations since they are the only ones that can supply such advanced HDDs at such scale. I wouldn't bet against any of these names, as they have raised guidance for Q2 FY26, and the chance of them raising guidance for FY26 and having explosive growth for FY27 is still easy on the table.

Risks

Hyperscalers continue to ramp up their capex, but as investors should know, AI, like most things, is cyclical, and one day capex will fall as investments and the AI infrastructure buildout will come to an end, as there will be enough capacity for data storage and HDDs. When AI infrastructure buildout starts slowing down, demand and growth rates for the HDDs market will slow and fall alongside it.

As AMD CEO Lisa Su mentions that we are in the 2nd year of a 10-year AI infrastructure buildout cycle, the demand will likely continue to rise for a couple more years before it dies down, as capex at least for 2026, doesn't seem like it is going to have a massive difference or decrease from 2025. For long-term investors, this may be a huge concern but shouldn't be a worry in the short term.

Final Thoughts and Takeaways

There should be 2 questions that investors should be thinking about,

1) Which is better, Seagate or Western Digital?

The title for the top spot is still up for grabs, as both currently have similar market share in the HDD market and have a different appeal to investors. Fundamentally, each company continues to thrive as demand exceeds supply. If one is interested in the HDD market, I honestly recommend just owning both names. There are clear differences in why each stock looks attractive, and I believe that there isn't a clear winner yet. Until there is a clear winner in the HDD market, if there is one, both names continue to be brilliant businesses that have expanding margins and explosive growth.

2) Is the HDD market worth investing in since the run-up?

The HDD market is currently running on huge momentum right now, especially for the pair of Seagate and Western Digital, being the top performers in the S&P 500. Both names are rated 'Hold' in the quant system, and I honestly believe that is the next best thing to do in the current market. Both names have really traded at such high valuations before, so with how bullish the HDD market is, I don't see why there isn't more space to run.

These are names I have monitored since the April low and have always wanted to start a position in, but I never really found the 'right' price point to do so, as I certainly did not expect such a huge rally over the last 6 months for both names. The risk has certainly been higher than before, as both names certainly seem overvalued, as they never entered such high valuations, but that doesn't mean the valuations are unreasonable. They are very important in the AI boom, as they remain key companies in the AI boom. With expanding margins and explosive growth, these names trading at higher premiums should be no surprise. Until further escalations, I'm bullish on the HDD market and HAMR HDDs, but I remain neutral on Seagate and Western Digital individually as I do believe there is an inflection point that will change the current dynamics within the HDD market, as I believe one of them will pull ahead of the other.

Analog Devices: A Semiconductor Leader

Nov. 15, 2025 Joseph Parrish

Summary

- Analog Devices (ADI) offers exposure to semiconductors beyond AI, focusing on industrial, automotive, and emerging robotics markets.
- ADI's financials reveal consistent profitability, positive free cash flow, and shareholder returns via buybacks and dividends, despite cyclical revenue swings.
- Robotics-driven growth is expected to boost ADI's long-term prospects, but current valuation already reflects much of this potential.

Analog Devices (ADI) is set to release its Q4 results on November 25. While it will be an important moment to check in, I think ADI's long-term potential beyond that of Q4 needs to be assessed, and I believe the stock price currently reflects that potential.

Business Model

ADI is a global semiconductor manufacturer that offers a wide variety of products. This mainly includes a variety of integrated circuits ("ICs").

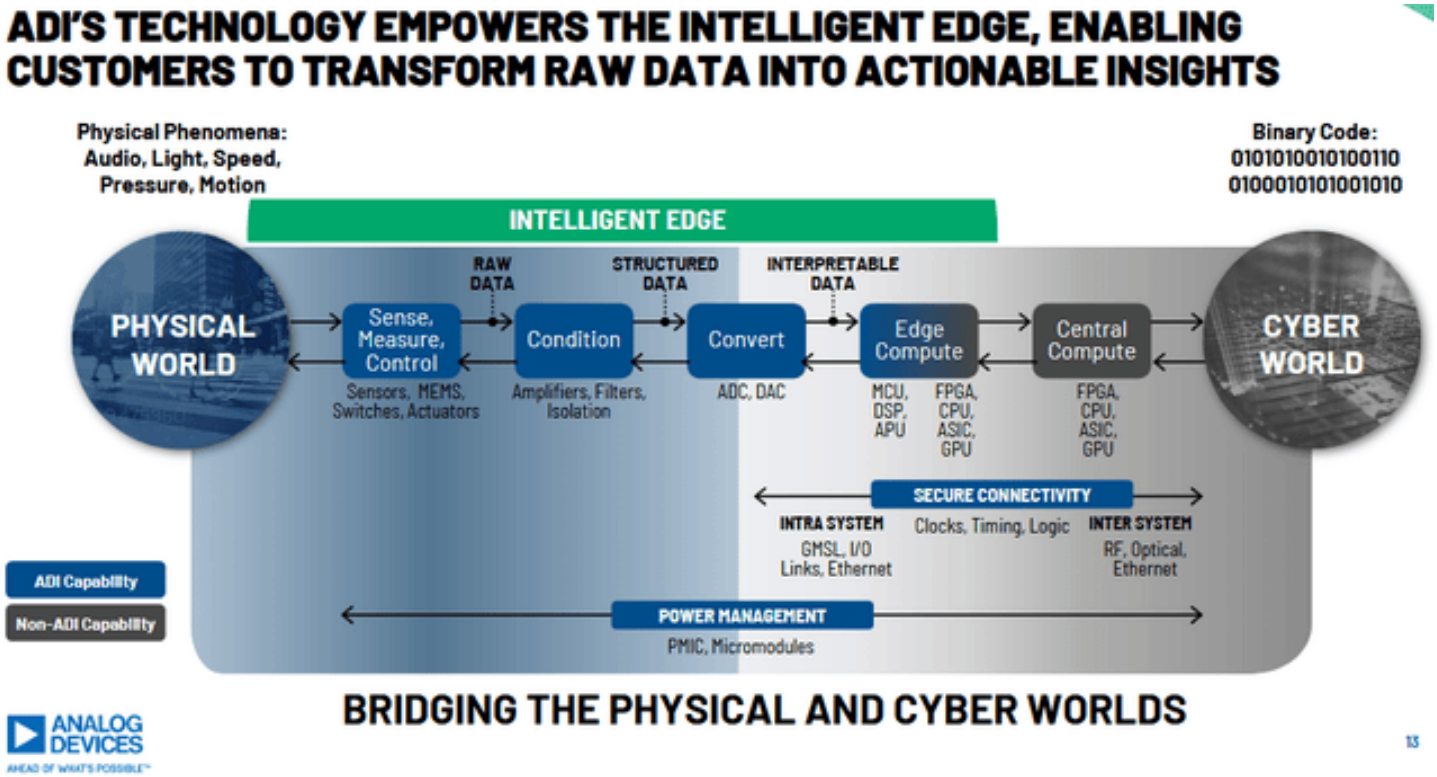

Fiscal 2024 Investor Presentation

Because products vary and will depend on the specifics of customer needs, they represent their revenue largely by the end use.

	Fiscal 2024			Fiscal 2023	
	Revenue	% of Total Revenue (1)	Y/Y%	Revenue	% of Total Revenue (1)
Industrial	$ 4,314,280	46 %	(35)%	$ 6,611,794	54 %
Automotive	2,827,439	30 %	(2)%	2,876,140	23 %
Communications	1,080,496	11 %	(33)%	1,606,426	13 %
Consumer	1,204,942	13 %	(1)%	1,211,179	10 %
Total Revenue	$ 9,427,157	100 %	(23)%	$ 12,305,539	100 %

Revenue by end user (2024 Form 10k)

Industrial uses are typically about half of their end uses. Automotive has been a growing portion of total revenue in recent years, previously less than 20% of total revenue but lately as high as 30%.

	Fiscal Year	
	2024	2023
United States	$ 2,840,426	$ 4,165,296
Rest of North and South America	62,318	88,579
Europe	2,109,529	3,001,871
Japan	1,085,631	1,397,119
China	2,128,840	2,229,631
Rest of Asia	1,200,413	1,423,043
Total Revenue	$ 9,427,157	$ 12,305,539

Revenue by geography (2024 Form 10k)

Revenue by geography (where their products are ultimately used), shows the U.S., Europe, China, and Japan as the four main markets, with the rest of Asia and the Americas contributing to a lesser extent.

	Fiscal 2024	
	Revenue	% of Total Revenue (1)
Distributors	$ 5,505,779	58 %
Direct customers	3,772,945	40 %
Other	148,433	2 %
Total Revenue	$ 9,427,157	100 %

Revenue by sales channel (2024 Form 10k)

As this depends on distributors far more than their direct sales, the geography of the end product does not necessarily reflect that of the customer.

What this indicates, however, is that this business is influenced by the macro-trends of these key markets.

Financial History

The macro-trends for this business model are visible in the financial history of this company (with fiscal years ended November), which I think show long-term patterns of growth, with blips of cyclicality thrown in there as well.

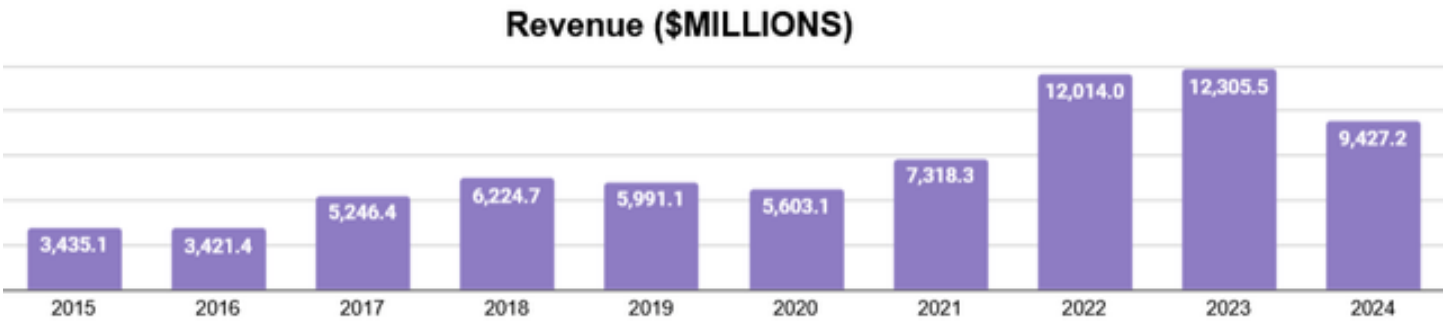

Author's display of 10k data

For example, revenue grew from $3.4B in 2015 to $9.4B in 2024. This is a lot of growth, but we also see revenue receded from a peak of $12.3B in 2023. The 2024 Form 10K summarized it as such:

Revenue decreased 23% in fiscal 2024 as compared to fiscal 2023 primarily as a result of weaker macroeconomic trends. This was pronounced in our Industrial end market as customers decreased their inventory balances and in the Communications end market primarily due to the timing of infrastructure deployment cycles. The Automotive and Consumer end markets declined to a lesser extent as demand weakened driven by reduced consumer spending.

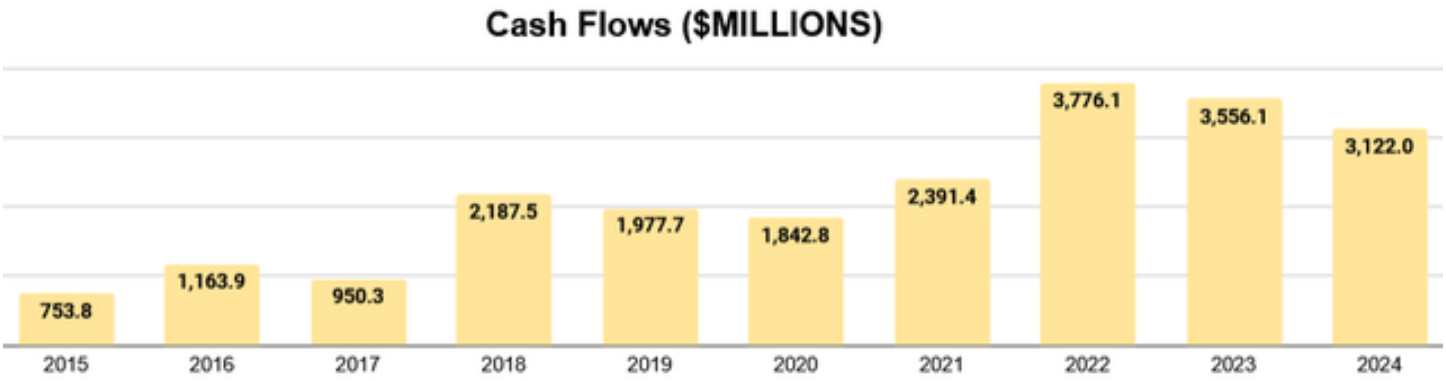

Author's display of 10k data

The pattern of free cash flow tends to mirror that of revenue as well. It's worth noting that, in spite of the cyclical downside, this number is always positive.

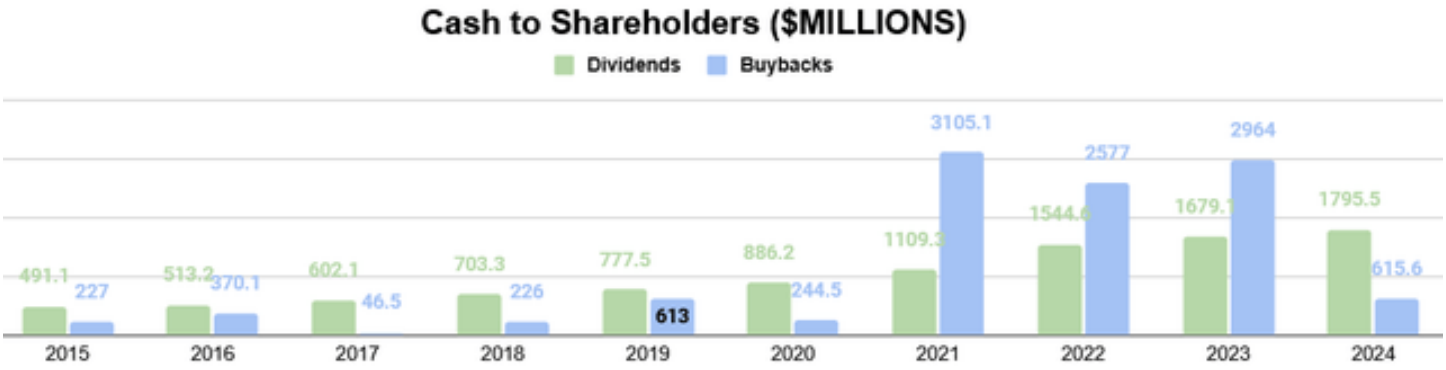

Author's display of 10k data

Shareholders, over time, have been rewarded with a mixture of buybacks and dividends. As capex is substantially less than operating cash flows in most years, there's little need to reinvest. Years with higher cash flows typically see the excess committed to buybacks, while annual distributions steadily increase.

	Three Months Ended		Nine Months Ended	
	Aug. 2, 2025	Aug. 3, 2024	Aug. 2, 2025	Aug. 3, 2024
Revenue	$ 2,880,348	$ 2,312,209	$ 7,943,590	$ 6,983,952
Cost of sales	1,090,600	1,000,970	3,111,929	3,018,737
Gross margin	1,789,748	1,311,239	4,831,661	3,965,215

Income Statement (Q3 2025 Earnings Release)

Results for fiscal 2025 show that this year has been movement in the up cycle, with nearly $1B more in revenue and very little extra in cost of sales.

Outlook and Valuation

I believe there's a case for long-term growth for ADI, going forward. Where many technology stocks are riding an AI opportunity, ADI's seems to be robotics more specifically. In the Q3 earnings call, CEO Vincent Roche summarized it as such:

ADI's high-performance technology stack and deep domain expertise are crucial to customer success in this highly sense, securely connected software-driven era and the new robotic modalities that are emerging. It's predicted that the convergence of compounding macro and AI-enabled technology factors will drive robust double-digit growth within robotics market for the foreseeable future.

While the industrial end use will drive the initial growth, I suspect the growth of robotics could trickle into the automotive and consumer uses as well, as these uses will likely incorporate robotic elements into their designs over time. Consider the broad range of ICs they make and provide, I think ADI is well-positioned to adapt its production as needed.

To value ADI for this growth, I decided to do a Discounted Cash Flow calculation based on free cash flow per share, and I used the following assumptions:

- $7.52 in FCF per share (TTM FCF with current shares outstanding)
- 10% CAGR first 5 years (to reflect robotics-driven growth)
- 0% CAGR the next 5 (to reflect cyclical downside)
- Terminal multiple of 10 (ADI probably wouldn't be valued in single digits)
- 10% discount rate (typical return of a market index)

Risks

Just as Industrial is a source of growth, one of the risk factors is weakness in the Automotive and Consumer lines, and I believe that Automotive shares patterns with Consumer, as automobiles are consumer products, whose higher prices (than, say, a phone) are justified by their "must have" status. A report from PWC highlights the growth in the auto-related semiconductor market, driven by the rise of EVs and autonomous vehicles. As cars become more digitized, demand for ADI's products should increase.

Weakness in the consumer can slow this down, with recent news indicating this is more likely. The Consumer Sentiment Index fell from 53.6 in October to 50.3 this month. This follows a series of earnings calls in which weakness in the consumer was noted, a key example being McDonald's (MCD).

Weak consumers will delay the purchase of new automobiles or devices that require more semiconductors. This could kick ADI's earnings growth out a few years and hurt its long-term valuation for today's buyers.

Conclusion

ADI is a way to bet on semiconductors outside of AI and computers. Consistently profitable and with an adaptable portfolio of products, it's one of the best ways to do that. For what it can do, however, I think it's fairly valued

China’s Tech Giants Race to Replace Nvidia’s AI Chips > As China sours on Nvidia, homegrown players try to fill the void

Nov 13, 2025 Tony Peng, Recode China AI

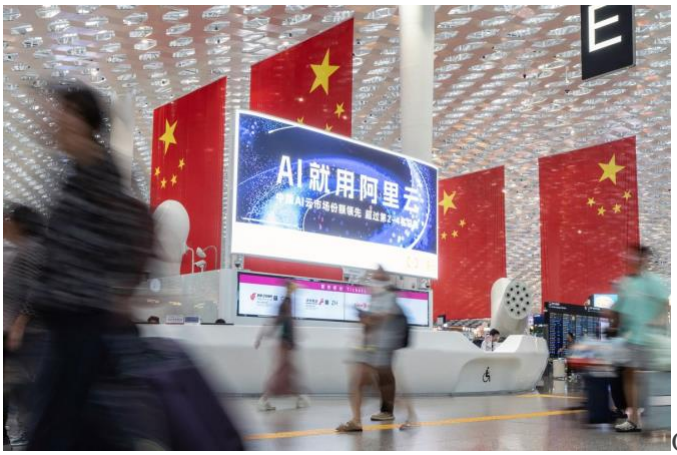

Qilai Shen/Bloomberg/Getty images

This post originally appeared on Recode China AI.

For more than a decade, Nvidia’s chips have been the beating heart of China’s AI ecosystem. Its GPUs powered search engines, video apps, smartphones, electric vehicles, and the current wave of generative AI models. Even as Washington tightened export rules for advanced AI chips, Chinese companies kept settling for and buying “China-only” Nvidia chips stripped of their most advanced features—H800, A800, and H20.

But by 2025, patience in Beijing had seemingly snapped. State media began labeling Nvidia’s China-compliant H20 as unsafe and possibly compromised with hidden “backdoors.” Regulators summoned company executives for questioning, while reports from The Financial Times surfaced that tech companies like Alibaba and ByteDance were quietly told to cancel new Nvidia GPU orders. The Chinese AI startup DeepSeek also signaled in August that its next model will be designed to run on China’s “next-generation” domestic AI chips.

The message was clear: China could no longer bet its AI future on an U.S. supplier. If Nvidia wouldn’t—or couldn’t—sell its best hardware in China, domestic alternatives must fill the void by designing specialized chips for both AI training (building models) and AI inference (running them).

That’s difficult—in fact, some say it’s impossible. Nvidia’s chips set the global benchmark for AI computing power. Matching them requires not just raw silicon performance but memory, interconnection bandwidth, software ecosystems, and above all, production capacity at scale.

Still, a few contenders have emerged as China’s best hope: Huawei, Alibaba, Baidu, and Cambricon. Each tells a different story about China’s bid to reinvent its AI hardware stack.

Huawei’s AI Chips Are in the Lead

VCG/Getty Images

f Nvidia is out, Huawei, one of China's largest tech companies, looks like the natural replacement. Its Ascend line of AI chips has matured under the U.S. sanctions, and in September 2025 the company laid out a multi-year public roadmap:

Ascend 950, expected in 2026 with a performance target of 1 petaflop in the low-precision FP8 format that's commonly used in AI chips. It will have 128 to 144 gigabytes of on-chip memory, and interconnect bandwidths (a measure of how fast it moves data between components) of up to 2 terabytes per second.

Ascend 960, expected in 2027, is projected to double the 950's capabilities.

Ascend 970 is further down the line, and promises significant leaps in both compute power and memory bandwidth.

The current offering is the Ascend 910B, introduced after U.S. sanctions cut Huawei off from global suppliers. Roughly comparable to the A100, Nvidia's top chip in 2020, it became the de facto option for companies who couldn't get Nvidia's GPUs. One Huawei official even claimed the 910B outperformed the A100 by around 20 percent in some training tasks in 2024. But the chip still relies on an older type of high-speed memory (HBM2E), and can't match Nvidia's H20: It holds about a third less data in memory and transfers data between chips about 40 percent more slowly.

The company's latest answer is the 910C, a dual-chiplet design that fuses two 910Bs. In theory, it can approach the performance of Nvidia's H100 chip (Nvidia's flagship chip until 2024); Huawei showcased a 384-chip Atlas 900 A3 SuperPoD cluster that reached roughly 300 Pflops of compute, implying that each 910C can deliver just under 800 teraflops when performing calculations in the FP16 format. That's still shy of the H100's roughly 2,000 Tflops, but it's enough to train large-scale models if deployed at scale. In fact, Huawei has detailed how they used Ascend AI chips to train DeepSeek-like models.

To address the performance gap at the single-chip level, Huawei is betting on rack-scale supercomputing clusters that pool thousands of chips together for massive gains in computing power. Building on its Atlas 900 A3 SuperPoD, the company plans to launch the Atlas 950 SuperPoD in 2026, linking 8,192 Ascend chips to deliver 8 exaflops of FP8 performance, backed by 1,152 TB of memory and 16.3 petabytes per second of interconnect bandwidth. The cluster will span a footprint larger than two full basketball courts. Looking further ahead, Huawei's Atlas 960 SuperPoD is set to scale up to 15,488 Ascend chips.

Hardware isn't Huawei's only play. Its MindSpore deep learning framework and lower-level CANN software are designed to lock customers into its ecosystem, offering a domestic alternative to PyTorch (a popular framework from Meta) and CUDA (Nvidia's platform for programming GPUs) respectively.

State-backed firms and U.S.-sanctioned companies like iFlytek, 360, and SenseTime have already signed on as Huawei clients. The Chinese tech giants ByteDance and Baidu also ordered small batches of chips for trial.

Yet Huawei isn't an automatic winner. Chinese telecom operators such as China Mobile and Unicom, which are also responsible for building China's data centers, remain wary of Huawei's influence. They often prefer to mix GPUs and AI chips from different suppliers rather than fully commit to Huawei. Big internet platforms, meanwhile, worry that partnering too closely could hand Huawei leverage over their own intellectual property.

Alibaba Pushes AI Chips to Protect Its Cloud Business

Sun Pengxiong/VCG/Getty Images

Alibaba's chip unit, T-Head, was founded in 2018 with modest ambitions around open-source RISC-V processors and data center servers. Today, it's emerging as one of China's most aggressive bids to compete with Nvidia.

T-Head's first AI chip is the Hanguang 800 chip, an efficient chip designed for AI inference that was announced in 2019; it's able to process 78,000 images per second and optimize recommendation algorithms and large language models (LLMs). Built on a 12-nanometer process with around 17 billion transistors, the chip can perform up to 820 trillion operations per second (TOPS) and access its memory at speeds of around 512 GB per second.

But its latest design—the PPU chip—is something else entirely. Built with 96 GB of high-bandwidth memory and support for high-speed PCIe 5.0 connections, the PPU is pitched as a direct rival to Nvidia's H20.

During a state-backed television program featuring a China Unicom data center, the PPU was presented as capable of rivaling Nvidia's H20. Reports suggest this data center runs over 16,000 PPUs out of 22,000 chips in total. The Information also reported that Alibaba has been using its AI chips to train LLMs.

Besides chips, Alibaba Cloud lately also upgraded its supernode server, named Panjiu, which now features 128 AI chips per rack, modular design for easy upgrades, and fully liquid cooling.

For Alibaba, the motivation is as much about cloud dominance as national policy. Its Alibaba Cloud business depends on reliable access to training-grade chips. By making its own silicon competitive with Nvidia's, Alibaba keeps its infrastructure roadmap under its own control.

Baidu's Big Chip Reveal in 2025

Baidu unveiled a 30,000 chip cluster powered by its 3rd generation P800 processors. Qilai Shen/Bloomberg/Getty Images

Baidu's chip story began long before today's AI frenzy. As early as 2011, the search giant was experimenting with field-programmable gate arrays (FPGAs) to accelerate its deep learning workloads for search and advertising. That internal project later grew into Kunlun

The first generation arrived in 2018. Kunlun 1 was built on Samsung's 14-nm process, and delivered around 260 TOPS with a peak memory bandwidth of 512 GB per second. Three years later came Kunlun 2, a modest upgrade. Fabricated on a 7-nm node, it pushed performance to 256 TOPS for low-precision INT8 calculations and 128 Tflops for FP16, all while reducing power to about 120 watts. Baidu aimed this second generation less at training and more at inference-heavy tasks such as its Apollo autonomous cars and Baidu AI Cloud services. Also in 2021, Baidu spun off Kunlun into an independent company called Kunlunxin, which was then valued at US $2 billion.

For years, little surfaced about Kunlun's progress. But that changed dramatically in 2025. At its developer conference, Baidu unveiled a 30,000-chip cluster powered by its third-generation P800 processors. Each P800 chip, according to research by Guosen Securities, reaches roughly 345 Tflops at FP16, putting it in the same level as Huawei's 910B and Nvidia's A100. Its interconnect bandwidth is reportedly close to Nvidia's H20. Baidu pitched the system as capable of training "DeepSeek-like" models with hundreds of billions of parameters. Baidu's latest multimodal models, the Qianfan-VL family of models with 3 billion, 8 billion, and 70 billion parameters, were all trained on its Kunlun P800 chips.

Kunlun's ambitions extend beyond Baidu's internal demands. This year, Kunlun chips secured orders worth over 1 billion yuan (about $139 million) for China Mobile's AI projects. That news helped restore investor confidence: Baidu's stock is up 64 percent this year, with the Kunlun reveal playing a central role in that rise.

Just today, Baidu announced its roadmap for its AI chips, promising to roll out a new product every year for the next five years. In 2026, the company will launch the M100, optimized for large-scale inference, and in 2027 the M300 will arrive, optimized for training and inference of massive multimodal models. Baidu hasn't yet released details about the chips' parameters.

Still, challenges loom. Samsung has been Baidu's foundry partner from day one, producing Kunlun chips on advanced process nodes. Yet reports from Seoul suggest Samsung has paused production of Baidu's 4-nm designs.

Cambricon's Chip Moves Make Waves in the Stock Market

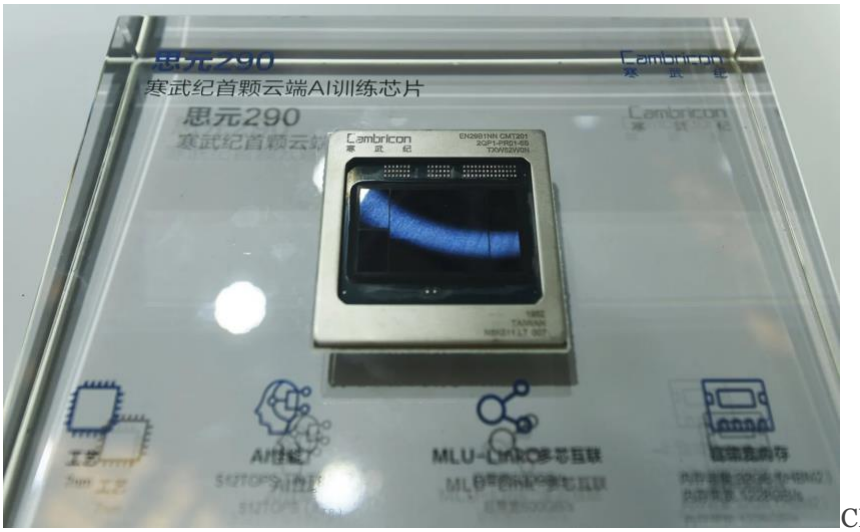

CFOTO/Future Publishg/Getty Images

The chip company Cambricon is probably the best performing publicly traded company on China's domestic stock market. Over the past 12 months, Cambricon's share price has jumped nearly 500 percent.

The company was officially spun out of the Chinese Academy of Sciences in 2016, but its roots stretch back to a 2008 research program focused on brain-inspired processors for deep learning. By the mid-2010s, the founders believed AI-specific chips were the future.

In its early years, Cambricon focused on accelerators called neural processing units (NPUs) for both mobile devices and servers. Huawei was a crucial first customer, licensing Cambricon's designs for its Kirin mobile processors. But as Huawei pivoted to develop its own chips, Cambricon lost a flagship partner, forcing it to expand quickly into edge and cloud accelerators. Backing from Alibaba, Lenovo, iFlytek, and major state-linked funds helped push Cambricon's valuation to $2.5 billion by 2018 and eventually landing it on Shanghai's Nasdaq-like STAR Market in 2020.

The next few years were rough. Revenues fell, investors pulled back, and the company bled cash while struggling to keep up with Nvidia's breakneck pace. For a while, Cambricon looked like another cautionary tale of Chinese semiconductor ambition. But by late 2024, fortunes began to change. The company returned to profitability, thanks in large part to its newest MLU series of chips.

That product line has steadily matured. The MLU 290, built on a 7-nm process with 46 billion transistors, was designed for hybrid training and inference tasks, with interconnect technology that could scale to clusters of more than 1,000 chips. The follow-up MLU 370, the last version before Cambricon was sanctioned by the United States government in 2022, can reach 96 Tflops at FP16.

Cambricon's real deal came with the MLU 590 in 2023. The 590 was built on 7-nm and delivered peak performance of 345 Tflops at FP16, with some reports suggesting it could even surpass Nvidia's H20 in certain scenarios. Importantly, it introduced support for less-precise data formats like FP8, which eased memory bandwidth pressure and boosted efficiency. This chip didn't just mark a leap—it turned Cambricon's finances around, restoring confidence that the company could deliver commercially viable products.

Now all eyes are on the MLU 690, currently in development. Industry chatter suggests it could approach, or even rival, Nvidia's H100 in some metrics. Expected upgrades include denser compute cores, stronger memory bandwidth, and further refinements in FP8 support. If successful, it would catapult Cambricon from "domestic alternative" status to a genuine competitor at the global frontier.

Cambricon still faces hurdles: its chips aren't yet produced at the same scale as Huawei's or Alibaba's, and past instability makes buyers cautious. But symbolically, its comeback matters. Once dismissed as a struggling startup, Cambricon is now seen as proof that China's domestic chip path can yield profitable, high-performance products.

A Geopolitical Tug-of-War

At its core, the battle over Nvidia's place in China isn't really about teraflops or bandwidth. It's about control. Washington sees chip restrictions as a way to protect national security and slow Beijing's advance in AI. Beijing sees rejecting Nvidia as a way to reduce strategic vulnerability, even if it means temporarily living with less powerful hardware.

China's big four contenders, Huawei, Alibaba, Baidu, and Cambricon, along with other smaller players such as Biren, Muxi, and Suiyuan, don't yet offer the real substitutes. Most of their offerings are barely comparable with A100, Nvidia's best chips five years ago, and they are working to catch up with H100, which was available three years ago.

Each player is also bundling its chips with proprietary software and stacks. This approach could force Chinese developers accustomed to Nvidia's CUDA to spend more time adapting their AI models which, in turn, could affect both training and inference.

DeepSeek's development of its next AI model, for example, has reportedly been delayed. The primary reason appears to be the company's effort to run more of its AI training or inference on Huawei's chips.

The question is not whether Chinese companies can build chips—they clearly can. The question is whether and when they can match Nvidia's combination of performance, software support, and trust from end-users. On that front, the jury's still out.

But one thing is certain: China no longer wants to play second fiddle in the world's most important technology race.

AMD Q3, 2025 Earnings Call Summary

Nov 04, 2025 AI generated summary, SL

Executive Overview

- Lisa Su (CEO) announced record quarterly results with revenue of $9.2 billion (+36 % Y/Y) and strong growth across data center AI, server, and PC segments.
- Jean Hu (CFO) confirmed broad momentum and strong cash generation, noting that AMD's AI business is scaling faster than expected with structural tailwinds in 2026–27.

📊 Financial Highlights (Jean Hu)

Metric	Q3 2025	Y/Y	Seq.	Key Drivers
Revenue	**$9.2 B**	+36 %	+20 %	Record EPYC, Ryzen 9000, MI350 sales
Gross Margin (Non-GAAP)	54 %	+40 bps	—	Mix shift to high-margin CPUs/GPUs
Operating Income	$2.2 B (24 %)	+25 %	+ —	Strong DC and Client growth
EPS (Non-GAAP)	$1.20	+30 %	+ —	
Free Cash Flow	$1.5 B (record)	↑ 3× Y/Y	—	
Cash / Debt	$7.2 B cash; $3.2 B debt (net cash position)	—	—	

Q4 Outlook:

- Revenue ≈ $9.6 B ± $0.3 B (+25 % Y/Y)
- Data Center ↑ double digits; strong CPU & GPU growth
- Gross Margin 54.5 %, Opex ≈ $2.8 B
- EPS up sequentially

💡 Segment Highlights

🏢 Data Center & AI (Lisa Su / Jean Hu)

- Revenue $4.3 B (+22 % Y/Y; +34 % Q/Q) — record level.
- Driven by 5th Gen EPYC "Turin" CPUs and Instinct MI350 GPUs.
- Over 160 new EPYC cloud instances launched; > 1,350 total globally.
- Server CPU share gains across cloud & enterprise; 5th Gen EPYC ≈ 50 % of EPYC mix.
- MI350 GPU ramp strong; MI400 launch planned for 2H 2026 as AI rack-scale Helios platform.
- OpenAI deal: 6 GW deployment (> $100 B potential rev. over next few years); first 1 GW MI450 in 2H 2026.
- Partnerships with Oracle (OCI), Meta (OCP Helios), DOE (Oak Ridge Discovery Supercomputer).

CFO Hu: Segment operating income $1.1 B (25 % margin) — profitability held despite heavy AI R&D investment.

💻 Client & Gaming (Lisa Su / Jean Hu)

- Revenue $4 B (+73 % Y/Y) — record levels.
- Client (PC): $2.8 B (+46 % Y/Y) — Ryzen 9000 desktop and notebook record sales.
- Gaming: $1.3 B (+181 % Y/Y) — semi-custom console chips & Radeon 9000 GPUs strong ahead of holiday.
- OEM Ryzen sell-through ↑ > 30 % Y/Y across commercial & gaming.
- Operating income: $867 M (21 % margin) vs $288 M a year ago.

🏭 Embedded (Jean Hu / Lisa Su)

- Revenue $857 M (-8 % Y/Y, +4 % Q/Q); sequential improvement in industrial, defense & healthcare.
- New Versal Prime Gen 2 SoCs and Ryzen Embedded 9000 Series launched.
- Operating Income $283 M (33 %) — margin decline from 40 % due to mix.
- Record design wins >$14 B YTD.

🚀 AI Roadmap and Helios Rack-Scale Strategy (Lisa Su)

- MI400 Series (2026): new compute engine + industry-leading memory capacity & networking.
- Helios Rack-Scale Platform: integrates MI400 GPUs + Venice EPYC CPUs + Pensando NICs; supports Meta Open Rack Wide standard.
- ZT Systems team integration: engineering Helios with Sanmina as lead manufacturer for mass deployments.
- AI Software: ROCm 7 launched — 4.6× inference and 3× training improvement vs ROCm 6.
- New contributors (Hugging Face, vLLM, SGLang) strengthen open-software ecosystem.
- Target: AI revenue scaling to "tens of billions annually by 2027."

💬 Q&A Themes

Topic	Speaker	Summary
MI355 → MI400 Transition	**Lisa Su**	MI355 to ramp through 1H '26; MI450 ramps in 2H '26. No pause expected; visibility strong into '26.
OpenAI Deal	**Lisa Su**	Multi-gigawatt agreement (> $100 B rev potential). Joint hardware/software development ongoing.
Helios Customer Demand	**Lisa Su**	Strong OCP reception; early deployments in rack-scale form.
Supply & Power Constraints	**Lisa Su**	Tight ecosystem capacity but well-planned supply chain and power coordination with partners.
CPU Outlook	**Lisa Su / Jean Hu**	Turin ramp fast; Genoa still strong. CPU demand durable into '26 as AI needs general compute.
Gross Margins 2026	**Jean Hu**	Margins normalize post-MI400 ramp; priority is top-line and gross profit dollar expansion.
Customer Concentration	**Lisa Su**	OpenAI key partner but supply chain sized for multiple large customers (OCI, Meta, DOE, etc.).
AI TAM Outlook	**Lisa Su**	> $500 B AI silicon TAM estimate rising further; to update at Analyst Day.

📈 Takeaways

- Record revenue and profitability across all segments.
- Data Center AI business accelerating toward multi-tens of billions by 2027.
- Strong visibility into 2026 driven by MI400 Helios platform and OpenAI/OCI deployments.
- CPU & GPU roadmaps (2 nm Venice, MI400) solidify AMD's AI infrastructure position.
- Balance sheet strength and cash flow support sustained R&D investment in AI leadership.

Intel Corporation's Future Is Process, 18A, 14A.

Nov 04, 2025 Patient Tech Investor

Summary

- Intel delivered strong Q3 results, but the critical story is its process innovation, especially the 18A node, which underpins future growth.
- INTC's unique 18A process offers a technological edge, with early yields supporting a reasonable margin, but full benefits and volume are expected by 2026.
- Long-term opportunity remains compelling if process synergy materializes, but risks persist if 18A fails to deliver; management's strategic clarity inspires confidence.


Justin Paget/DigitalVision

Intel (INTC) reported stellar financial results for its 3rd quarter, which masked, in our view, the more, well, most important story, process. Yes, it was discussed but seemed an afterthought, at least from the analysts. A lot of focus seemed targeted at capacity constraints for older products. This isn't the future. The elephant issue is process success, an entity that opens doors for future Intel products, future outside foundry business, and significantly lifts margins. This is the synergy entity. Everything depends on this success. It matters in making Intel the big, influential semi player. A few analysts did get it right, asking for more details on 18A. Blayne Curtis, now of Jefferies, was one who did. Yes, Intel might not be "the only" on the stage, but it can carry, once again, huge free world influence.

From the call, we found this to be of great value. In the CEO's prepared remarks,

Let me dive deeper into our underlying business trend. Over the course of my career, I have had the privilege of contributing multiple ways of disruptive innovation. But I can't recall a time that I have been more excited about the future of computing and opportunities in front of us. We are still in the early stage of AI revolution, and I believe Intel can and will play a much more significant role as we transform the company.

He didn't come to fail. Looking deeper, the CFO, David Zinsner, added,

. . . when Lip-Bu came in, he really was upset about the balance sheet. So we've done a lot to work on that and improve that for him. We took $4.3 billion of debt off the books this quarter, and all the maturities next quarter or next year should come off and we'll repay that.

Intel faces a multiple front daunting task in its repair. We certainly acknowledge that. But investors must not equate synergistic opportunities with daunting tasks. This is where opportunities, real opportunities, exist. In the movie, Miracle, Kurt Russell playing Coach Brooks, opened his USSR match chat with, "Great moments are born from great opportunity. . .This is what you have here" So is it with Intel.

This continues our coverage of Intel, our last, Intel's Direction Continues In Flux, Expect Volatility, argued rightfully that the stock price would be news rather than results driven.

Now, we are taking inventory, so to speak, of where Intel is today while viewing, obviously, a possible future. Today is inventory day.

Importance of Process

With the future in mind, we begin with a small detour into the Apple/TSMC (AAPL) worlds. From an article posted on 9to5Mac, this interesting rumor appeared.

But based on a new report, we could see even more pricing changes next year if rumored A20 chip costs prove true. . . . Instead, TSMC's new 2-nanometer process will reportedly come with at least a 50% price increase compared to the 3-nanometer process.

This is 20A not 18A, and if rumors are correct, it comes without backside power, a likely requirement going forward nor RibbonFET, a unique transistor power scheme. The message is smaller, and smaller technologies are getting more costly and fussy. This isn't simple core memory nor TTL technology.

Intel's Process Direction

We have been crowing about process, process, process. Now it's time to hit some details. 18A defines the transistor size in angstroms. We discussed back power above, which separates the power network from communications networking, each on opposite sides. With transistors becoming smaller, a design with power and communications on the same layer, causes increased cross interference. This increases power usage and decreases reliability. Intel is the only major fab with this capability at this point. The second feature uses a unique power approach, powering cells on four sides rather than three, reducing leakage and improving performance. At this point, Intel stands alone with the smallest size and unique design.

Now, during the call, management made several important statements concerning 18A and 14A, which are akin but the next size down. A review of the list follows:

18A

"We are making steady progress on Intel 18A. We are on track to bring Panther Lake to market this year."

- The targeted fab in Arizona is fully functional.
- Yields are progressing at the predicted rate.
- Early products include Panther Lake and Clearwater Forest.
- In early ramps, low volumes will negatively impact margins.
- Built capacity not filled until 2030.
- Some capital spending left but not on capacity.

"Yields are . . adequate to address [current] supply [requirements], but they are not where we need them to be in order to drive the appropriate level of margins."

"Expecting by the end of 2026 to be at appropriate yields, we'll probably be in that space." (Answer to a Blayne Curtis question.)

Perhaps the most important comment from the CFO is his continual focus on 40% - 60% product margins going forward. Our sense is that 18A yields have improved enough to support at least 40% for the total business going forward, or will shortly do so, but that the real targets with volume and progress rely on a level of potential. It is in this comment that investors must be watchful.

Now a 14A status review follows:

14A

Intel 14A,

- "focus on technology definition, transistor architecture, process flow, design enablement, and foundation IPs."
- "Engaged with potential [many] external customers."
- "[E]ncouraged by the earlier feedback, which helps us to drive and inform our decisions."
- "Foundry also advanced the development of Intel 14A and continues to make progress expanding its advanced packaging deal pipeline."

A Short 3rd Quarter Summary

We begin with a summary slide for the call.

intel

The results beat expectations of $13.15B vs. $13.7B and earnings of $0.01 vs. $0.23.

The internals offer a bit of deeper understanding.

- Timing confirmed for Panther Lake 18A technology to go live in the next few weeks.
- Timing for 18A Clearwater Forest technology is now set for the first six months of 2026.
- Legacy product demand outstripped supply and is expected to continue into 2026. Management clearly stated that it had no plans to expand capacity.
- "Q3 operating cash flow was $2.5 billion with gross CapEx of $3 billion in the quarter and positive adjusted free cash flow of $900 million."
- On the most important short-term issue, cash flows, management announced:
    - Exited Q3 with $30.9 billion cash.
    - Received $5.7 billion from the U.S. government.
    - $2 billion from SoftBank Group.
    - $4.3 billion from the Altera closure.
    - $900 million from the Mobileye stake sale.
    - "Expect NVIDIA's $5 billion investment to close by the end of Q4."
    - Repaid $4.3 billion of debt.
    - Plans include paying down 2026 due maturities.

The company guided flat year over year with slightly lower gross margins impacted negatively from the early startup of 18A and the early very low volumes.

Continuing, on the margin front, guidance for 2026 included headwinds from the Panther Lake startup in the first half of the year. A higher-margin business, Altera, now sold, creates another headwind. For the last half of 26, 18A production ramps up with higher volumes, and Clearwater Forest production begins. This drastically supports higher margins appearing later in 26. Viewing margin expectations remains a critical exercise for stock prices. Gross margins will be lumpy, very lumpy, and stock prices will likely follow at a level of unison.

Valuation

The stock price opened much higher but over the course of the day closed unchanged (+$0.10) on twice the normal volume. This is generally a strong sign of a top in place. One company, Bank of America, kept its underperform rating in place, stating "a price target of $34, based on a 3.0 multiple of his enterprise value-to-sales ratio estimate for 2027, . . ." The stock is also very overbought.

Once stocks significantly trade above growth rates, and many do; chart prices become more important guides. We added one below from Yahoo Finance.

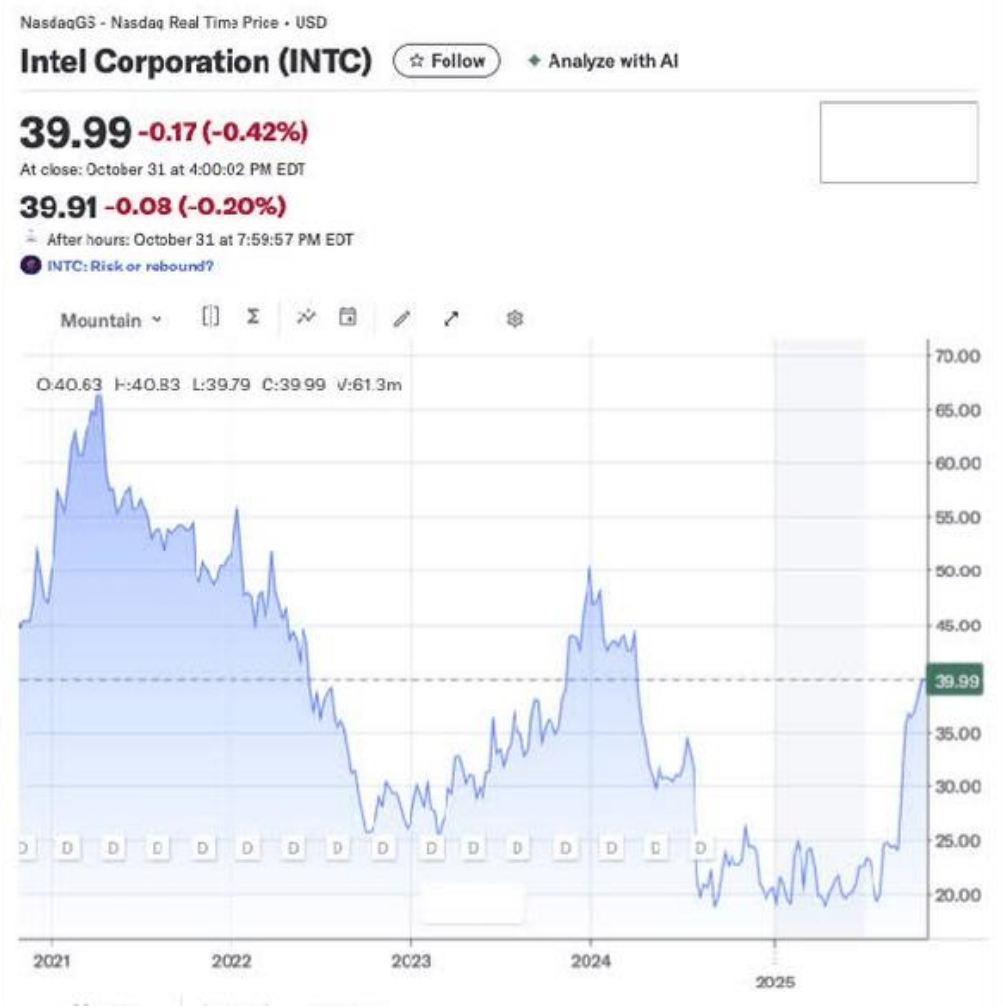

Yahoo Finance

The several year chart shows a long-term lower top and lower bottom. This suggests that a more certain bottom price resolve hasn't been established. Lower prices than $20 still have a future probability should matters turn more south. Also, the $45 - $50 middle of the chart high range suggests points of resistance. Interestingly, the price reached almost to the very low part of that range in after-hours, the night of the report.

Intel is earning about $0.50 per share at this point. Prices at $50- under its slow, methodical growth looking out the windshield view—strongly suggest a stock more than fully valued. A P/E of near 100 illustrates our belief. Thus, we are cautious about future investment additions at this point.

Risk & Our Best/Worst Fears

We have been openly discussing the future of Intel with our emphasis on process. It is everything. When 18A becomes yield competitive over the next four quarters, the company arrives with a synergistic advantage, having control over both manufacturing and products. No one else carries this advantage, no one. Margins for

internal products might reach the 60% range once again, and as discussed in other of our articles, margins for other outside products might equal 30%. This becomes a compelling story, a must-own. Investors should remember that the foundry business is in a worse light, breakeven. Intel's products, alone, are sufficient.

Political climates drive longer-term value, especially for foundries. The President visits with the Chinese leaders, discussing of particular importance for Intel's story, Taiwan, which remains critical. China continues its claim on this nation. Trump believes otherwise, by-the-way. We already understand what that means for manufacturing. China's approach to monopolizing rare earth minerals gives advanced warning to what will happen with chip manufacturing regardless of TSMC's manufacturing locations.

Now our worst fear is that Intel stock becomes a story stock trading at valuations uncharacteristic of its fundamentals. At our age, this becomes a difficult hold.

Risks do exist in that 18A doesn't materialize, leaving the company in the same circumstances as other chip developers without foundry synergy. Under this scenario, investors can only expect so-so results. It likely becomes a holding better moved elsewhere. We don't believe that this will be the case.

At the beginning of the article, we noted a real-sourced rumor of TSMC demanding capital repayment help for 20A technology from Apple. In our view, this isn't a small change. It suggests that in time Intel may expect and get the same. This helps roll off at least one capital risk.

We can't argue that the process synergy won't morph this investment into a dream. But we also agree with Bank of America that the steep price leaves investors in a quandary about where to buy and at what price. Thus, we downgrade Intel temporarily to a hold, waiting for lower prices.

It is still process-related and, in our view, a very long-term investment. Some criticize Intel's current product offerings, claiming dated technology. We don't argue, but we also add this point. The CEO understands two important factors: the products must become simpler and fewer, yet flexible in design. That issue is being addressed competently. Michelle Johnston Holthaus, Chief Executive Officer of Intel Products, is at the head. You might want to go listen to her approach. It is both pragmatic and honest with engagement internally and externally. She understands clearly the target of flexible simplicity. She helped drive it before Lip-Bu Tan came.

There is significant risk, but we also see tremendous opportunity coupled with focused and intelligent management with the courage to get it correct. Yet, again, the price is ahead of itself. We close with this repetition:

Let me dive deeper into our underlying business trend. Over the course of my career, I have had the privilege of contributing multiple ways of disruptive innovation. But I can't recall a time that I have been more excited about the future of computing and opportunities in front of us. We are still in the early stage of AI revolution, and I believe Intel can and will play a much more significant role as we transform the company.

Axcelis outlines $215M Q4 revenue target as Veeco merger advances and memory market optimism grows

Nov 04, 2025 AI generated earnings calls Insights

Management View

- CEO Russell Low reported "solid results in the third quarter with revenue of $214 million and non-GAAP earnings per diluted share of $1.21, both exceeding our outlook." He highlighted record Customer Solutions & Innovations (CS&I) revenue and noted system revenue was slightly above expectations, contributing to profitability.
- Low announced the pending all-stock merger with Veeco, stating, "this merger is expected to position the combined company as a key beneficiary and critical enabler of secular tailwinds, including AI and electrification." He emphasized cross-sell synergies, technical depth, expanded presence in Silicon Carbide and MOCVD markets, and combined strengths in memory, foundry logic, and advanced packaging.
- He detailed new product launches: "our new Purion Power Plus series… is designed to enable improved device performance and increased productivity for next-generation power devices." Additional innovation included the GSD Ovation ES implanter and positive customer feedback on MUSIC, the multistep implant chain capability.
- CFO James Coogan stated, "Third quarter revenue was $214 million, with systems revenue at $144 million and CS&I revenue at a record of $70 million, both above our expectations for the quarter." Coogan also highlighted that "we generated approximately $5 million in other income with the sequential decrease primarily due to foreign currency."

Outlook

- Axcelis expects Q4 2025 revenue of approximately $215 million. Coogan shared, "our preliminary view on the first quarter of 2026 suggests revenues to be relatively similar to our anticipated levels in the fourth quarter of 2025."
- Non-GAAP gross margins for Q4 are projected at approximately 43%. Non-GAAP operating expenses are expected to be about $56 million, with adjusted EBITDA expected to be around $41 million and non-GAAP diluted EPS of approximately $1.12.
- Management anticipates sequential improvement in bookings in Q4, improved memory market conditions in 2026, and ongoing moderation of demand in 2025 but highlighted a strong base of profitability and cash flow.

Financial Results

- Q3 revenue was $214 million, with systems revenue at $144 million and CS&I revenue at $70 million. China accounted for 46% of total sales, down from 55% in the prior quarter, with U.S. sales at 14% and Korea at 10%.
- Bookings declined to $52 million, and Axcelis exited Q3 with a backlog of $484 million.
- GAAP gross margin was 41.6%, non-GAAP gross margin was 41.8%. GAAP operating expenses were $63.8 million, non-GAAP $50.4 million. Non-GAAP operating margin was 18.2%. Adjusted EBITDA was $43 million.
- GAAP diluted EPS was $0.83, non-GAAP diluted EPS $1.21. Free cash flow of $43 million was generated during the quarter. Share repurchases totaled $32 million, with $135 million remaining under authorization. Cash, cash equivalents, and marketable securities totaled $593 million.

Q&A

- Jonathan Dorsheimer, William Blair, asked about dynamics in the non-Silicon Carbide power category and tariff impacts. CEO Low explained differentiation in silicon and advanced power applications, while CFO Coogan discussed ongoing efforts to mitigate tariff impacts and noted, "as we look to 2026, it could have a little bit more of an impact... the team is working now on working to mitigate the potential impact of that."
- Craig Ellis, B. Riley Securities, inquired about China demand stability and memory market trends. CEO Low stated, "China demand in 2026 will depend upon the end demand environment as well as how much progress they make on the chip self-sufficiency targets," and highlighted optimism for DRAM and HBM demand.
- Christian Schwab, Craig-Hallum, asked about memory market exposure and general mature segment recovery. CEO Low outlined capital intensity for memory wafer starts and noted general mature recovery would be "driven by the macro climate... it is too soon to say that the other markets, namely consumer, industrial and automotive have actually turned."
- Jack Egan, Charter Equity, questioned the sustainability of CS&I revenue and drivers of Q4 bookings. CFO Coogan attributed CS&I strength to improved utilization rates and said upgrades were strong, especially in memory. CEO Low described Q4 bookings expectations as "across pretty much all of our customers, not specifically a given market segment."
- Mark Miller, Benchmark, sought insight on Silicon Carbide trends and EV penetration. CEO Low cited price declines and new applications as supporting growth, and Coogan added, "penetration of Silicon Carbide into full EVs being in... low teens... there's still a lot of room to run in the automobile market."
- Denis Pyatchanin, Stifel, asked about bookings softness and CS&I by geography. CFO Coogan noted softness in Power General Mature and highlighted strong CS&I activity in Korea and China.
- Duksan Jang, Bank of America, asked about backlog coverage and China revenue drivers. CEO Low discussed bookings volatility, backlog stability, and CS&I as a "solid financial base." CFO Coogan said, "we anticipated China revenue being down both for General Mature and Power... we are still seeing new entrants into the Power market."
- David Duley, Steelhead, asked about gross margin drivers and Silicon Carbide applications beyond EVs. CFO Coogan detailed product mix and installation timing as margin factors, while CEO Low cited grid and data center applications for Silicon Carbide.

Sentiment Analysis

- Analysts expressed cautious optimism, focusing on bookings trajectory, China exposure, memory and power markets, and margin sustainability. Questions were targeted and pressing, particularly regarding memory market timing and backlog coverage.
- Management maintained a confident tone, using phrases like "we are pleased with our execution," and "we see exciting opportunities," but showed caution in projecting broad-based recoveries. In Q&A, management provided detailed, direct responses, though some answers on market timing and geographic specifics were qualified as "too early to say."
- Compared to the previous quarter, analyst tone remained analytical but increased focus on CS&I trends, backlog levels, and merger implications. Management's tone was consistent, but the Veeco merger added a strategic layer of confidence.

Quarter-over-Quarter Comparison

- Sequential revenue increased from $195 million in Q2 to $214 million in Q3. Non-GAAP EPS moved from $1.13 to $1.21. CS&I revenue reached a record $70 million, up from $61 million.
- Bookings fell from $96 million in Q2 to $52 million in Q3, with backlog declining from $582 million to $484 million.
- China revenue mix decreased, while U.S. and Korea shares shifted modestly. Gross margin declined from 45.2% non-GAAP to 41.8% non-GAAP, primarily due to product mix and lower-margin installations.
- Strategic outlook shifted with the Veeco merger announcement, introducing cross-sell synergy, technical depth, and broader market positioning.
- Analysts in Q3 probed more about merger benefits, CS&I sustainability, and near-term memory market inflection, while the previous quarter focused more on cyclicality and China competition.

Risks and Concerns

- Management cited ongoing market "digestion" in power and general mature segments, with uncertain timing for broad-based recovery in consumer, industrial, and automotive markets.
- Tariff impacts are expected to be more pronounced in 2026, with mitigation efforts underway.
- Bookings volatility and backlog coverage were highlighted as areas to monitor, with Q4 bookings expected to rebound but full-year bookings anticipated below prior peaks.
- Analysts raised concerns about China demand stability, memory recovery timing, and sustainability of elevated CS&I revenues.

Final Takeaway

Axcelis Technologies delivered a strong Q3 performance, highlighted by increased revenue, record CS&I results, and disciplined cost control. The pending Veeco merger marks a significant strategic milestone, aimed at expanding technical capabilities and market reach in growth areas like AI, electrification, and Silicon Carbide. Management remains confident in its financial base and ability to navigate cyclical challenges, citing ongoing innovation and customer engagement as key drivers for future growth, with cautious optimism for memory market recovery and a continued focus on operational resilience.

Lattice projects 22% Q4 revenue growth and targets accelerated expansion into 2026 as AI demand strengthens

Nov. 04, 2025 AI Generated Earnings Calls Insights

Earnings Call Insights: Lattice Semiconductor (LSCC) Q3 2025

Management View

- CEO Fouad Tamer marked his first year in the role by emphasizing "our strong Q3 performance and forward-looking guidance reflect the strength of our strategy and execution." He highlighted Lattice's robust innovation pipeline and leadership in low-power data center solutions, especially in AI and post-quantum cryptography (PQC). Tamer noted, "the adoption of Lattice's PQC technology is also accelerating due to the NIST requirement that systems be CNSA compliant." He reported record growth

in design wins and described momentum in Lattice's core markets, setting "the foundation for rapid growth into 2026."

- Tamer stated, "For Q3, we delivered revenue of $133.3 million, up 7.6% over Q2. This represents the highest sequential growth in more than 4 years." He shared that communications and computing grew 8% sequentially and 21% year-over-year, while industrial and automotive increased 6% sequentially, with normalization of channel inventory on track by year-end.
- He explained, "We are confident that we're gaining share across smart factory, robotics, medical and aerospace and defense applications based on customer feedback and design win activity."
- Tamer projected, "We estimate the percentage of AI usage across our products will be in the high teens in 2025 and mid-20% range in 2026."
- CFO Lorenzo A. Flores stated, "Revenue increased 7.6% quarter-on-quarter and 4.9% on a year-over-year basis to $133.3 million. Overall, this was the highest revenue we have obtained in 5 quarters, and we are expecting continued growth in Q4 and in 2026." Flores added, "Our gross margin expanded by 20 basis points quarter-over-quarter and 50 basis points year-over-year, 69.5% on a non-GAAP basis."
- Flores highlighted, "We delivered non-GAAP EPS of $0.28, which was at the midpoint of our guidance and represented 17% growth on both a year-over-year and quarter-over-quarter basis."

Outlook

- Flores guided Q4 revenue to be in the range of $138 million to $148 million, stating "At the midpoint, this represents revenue growth of 22% over Q4 of last year."
- Gross margin is expected to be "69.5%, plus or minus 1% on a non-GAAP basis."
- Non-GAAP operating expenses are expected between $54.5 million and $56.5 million, and non-GAAP EPS is expected between $0.30 and $0.34 per share.
- Management expects inventory normalization in industrial and automotive by year-end and projects accelerating growth in 2026, with communications and computing driving revenue expansion.

Financial Results

- Lattice reported Q3 revenue of $133.3 million.
- Communications and computing revenue reached a record, growing 21% year-over-year and 8% sequentially.
- Gross margin on a non-GAAP basis was 69.5%.
- Non-GAAP operating expense was $53.9 million, with a non-GAAP operating margin of 29% and an EBITDA margin of 35.6%.
- Non-GAAP EPS was $0.28.
- GAAP net cash flow from operating activities for Q3 was $47.1 million, with a GAAP operating cash flow margin of 35.4%. Free cash flow was $34 million.
- Lattice repurchased approximately $15 million of common stock during the quarter, with $14 million remaining on its authorization.

Q&A

- Kevin Garrigan, Jefferies: Asked if 2026 confidence is tied to normalization in industrial and auto inventory or if strong communications and compute growth can offset softness. Tamer responded, "Our comms and compute business as a percent of total revenue went from 35% of total revenue in 2023 to

45% of total revenue in 2024 to an expected over 55% of total revenue in 2025, and we expect that to grow to about 60% of revenue into 2026."

- David Williams, Benchmark: Inquired about increased management confidence over the past 90 days. Tamer noted "a very successful Open Compute Summit" and described "the past 90 days, we've seen definitely an increase in activity and spend."
- Tristan Gerra, Baird: Asked about sequential growth in communications and computing. Flores replied, "the real driver of the comms and strength is our server demand. And that's over 80% growth year-over-year."
- Gary Mobley, Loop Capital: Questioned undershipping in industrial and automotive. Tamer responded, "we're undershipping by about a couple of weeks a quarter, so call it $15 million to $20 million a quarter."
- Quinn Bolton, Needham: Probed the implied decline in industrial and auto. Flores stated, "the challenge in reconciling...what you're saying and what we're seeing is there is a direct aspect, meaning non-channel piece of our industrial and auto."

Sentiment Analysis

- Analysts raised pointed questions about sustainability of growth, inventory normalization, and segment-specific momentum, with a cautiously optimistic tone and some skepticism about industrial and auto projections.
- Management tone was confident and upbeat in prepared remarks, reinforced by repeated emphasis on robust design win activity, strong bookings, and accelerating demand in communications and computing. During Q&A, management provided detailed explanations and maintained confidence, using phrases like "we are confident" and "we expect."
- Compared to the previous quarter, both analysts and management struck a more optimistic tone, with greater emphasis on accelerating demand and visibility into 2026.

Quarter-over-Quarter Comparison

- Management raised its Q4 revenue guidance from Q3's $133 million midpoint to a $143 million midpoint, representing a higher sequential and year-over-year growth rate than previously guided.
- The outlook for 2026 was presented with increased conviction, with a notable shift in segment contributions: communications and computing are now expected to comprise about 60% of revenue next year.
- Analysts shifted their focus from near-term inventory normalization to longer-term growth drivers, particularly AI-driven demand and design win momentum.
- Management's confidence has increased since the prior quarter, supported by stronger bookings and backlog.

Risks and Concerns

- Management acknowledged continued macroeconomic and industry risks, with particular focus on channel inventory normalization in industrial and automotive.
- Supply chain pressures, particularly lead times for communications and computing, were mentioned, but management stated they are "on top of it and are very focused on making sure our customers get supply."

- Pricing pressures from suppliers were noted, with management emphasizing a "price to value" strategy to maintain gross margin stability.
- Analysts expressed concerns about the sustainability of industrial and auto growth and the implications of undershipping dynamics.

Final Takeaway

Lattice management highlighted robust sequential growth in Q3, record design wins, and increasing AI adoption as key contributors to their positive outlook. With Q4 revenue guidance pointing to the strongest year-on-year growth in nearly two years and communications and computing expected to drive over 60% of revenue in 2026, management reiterated confidence in both near-term execution and long-term expansion. Strategic investments in new products and ongoing inventory normalization position the company for sustained growth, while continued focus on operating discipline and capital allocation supports shareholder value.

Meta: Gloomy Ad Revenue Outlook, Risky AI Bets Doubled

Nov. 03, 2025 Sandeep G. Rao

Summary

- Meta Platforms, Inc. reported Q3 FY2025 results showing revenue growth but declining EPS, driven by slowing ad growth in key markets.
- META's AI initiatives, including Llama LLMs and major datacenter investments, have yet to yield significant traction or profitability outside its core apps.
- Capital expenditures are surging, with the Hyperion datacenter joint venture reflecting strategic flexibility but also highlighting high spending versus monetization.

Fritz Jorgensen/iStock Editorial via Getty Images

On the 29th of October, social media giant Meta Platforms, Inc. (META) released the third quarter (Q3) results of its Fiscal Year (FY) 2025. When diving deeper into its segments vis-a-vis bottom line analysis, there are signs

underscoring macroeconomic rumblings in their most important markets: the United States and the Western Hemisphere.

Trend Drilldown

At first glance, Meta's top and bottom lines show outperformance relative to the previous FY:

	9M 2025, % of FY 2024	FY 2024	FY 2023	FY 2022	FY 2021	FY 2020
		Year-on-Year Trends				
Advertising	86%	22%	16%	-1%	37%	21%
Other Revenue	104%	63%	31%	12%	10%	21%
Reality Labs	58%	13%	-12%	-5%	100%	127%
Total Revenue	**86%**	**22%**	**16%**	**-1%**	**37%**	**22%**
R&D Expenses	92%	14%	9%	43%	34%	36%
Total Costs and Expenses	**87%**	**8%**	**1%**	**23%**	**34%**	**14%**
Operating Income	**84%**	**48%**	**62%**	**-38%**	**43%**	**36%**
Net Income	**60%**	**59%**	**69%**	**-41%**	**35%**	**58%**
Earnings Per Share, Diluted	**61%**	**60%**	**73%**	**-38%**	**36%**	**57%**
Segments		Share of Net Revenue				
Advertising	98%	98%	98%	97%	97%	98%
Other Revenue	1%	1%	1%	1%	1%	1%
Reality Labs	1%	1%	1%	2%	2%	1%

Soruce: Created by Sandeep G. Rao using data from Meta's Financial stmts

If trends over 9M 2025 continue, total revenue will be around 15% higher than that in FY 2024 but earnings per share will be around 19% lower. Revenues from advertising remain, by far, the biggest source of revenues. With a current growth trend of 15% over the current FY, Meta's advertising revenue growth falls somewhat short of the 19% average established in FY terms since 2020.

Meanwhile, both "Other Revenue" - representing its WhatsApp Business platform, Meta Verified subscriptions and developer fees—and "Reality Labs" - representing its VR headsets and glasses—are relatively insignificant and consistently under 2% since FY 2020.

Trends indicate that Meta's EPS growth would be 19% lower than in the previous FY, which had registered a 60% growth over the year prior—which is a massive shift from all periods since 2020 (with the exception of the pandemic years of 2021 and 2022).

The drop in advertising growth is directly linked to the spending habits and trends therein within the U.S. and Western Europe. Neither inflationary trends nor trade tensions have subsided. With lower spending being projected, there will be lower ad spends.

The company is now consolidating and strengthening its bid in a trending buzzword in key American stocks for some time now: AI.

Behind Meta's Foray into AI

As per reports, investing in AI now accounts for at least 40% of the share of US GDP ("Gross Domestic Product") growth this year, and AI companies are estimated to have accounted for 80% of the gains in the U.S. market so far this year. The driving narrative that essentially makes the entire U.S. economy one big bet on AI is the notion that it would deliver a significant boost to productivity growth by delivering cost savings in the long term by requiring fewer workers—who are already among the most productive in the world on average.

This long-running notion has led to the creation of a massive network of "circular deals" worth several trillions of dollars predominantly around NVIDIA Corporation (NVDA) and OpenAI (OPENAI) - which typically sees a pattern of purchases of products and services following an investment. This serial money machine is now nearly inextricably linked to every major American tech firm and leading AI startup.

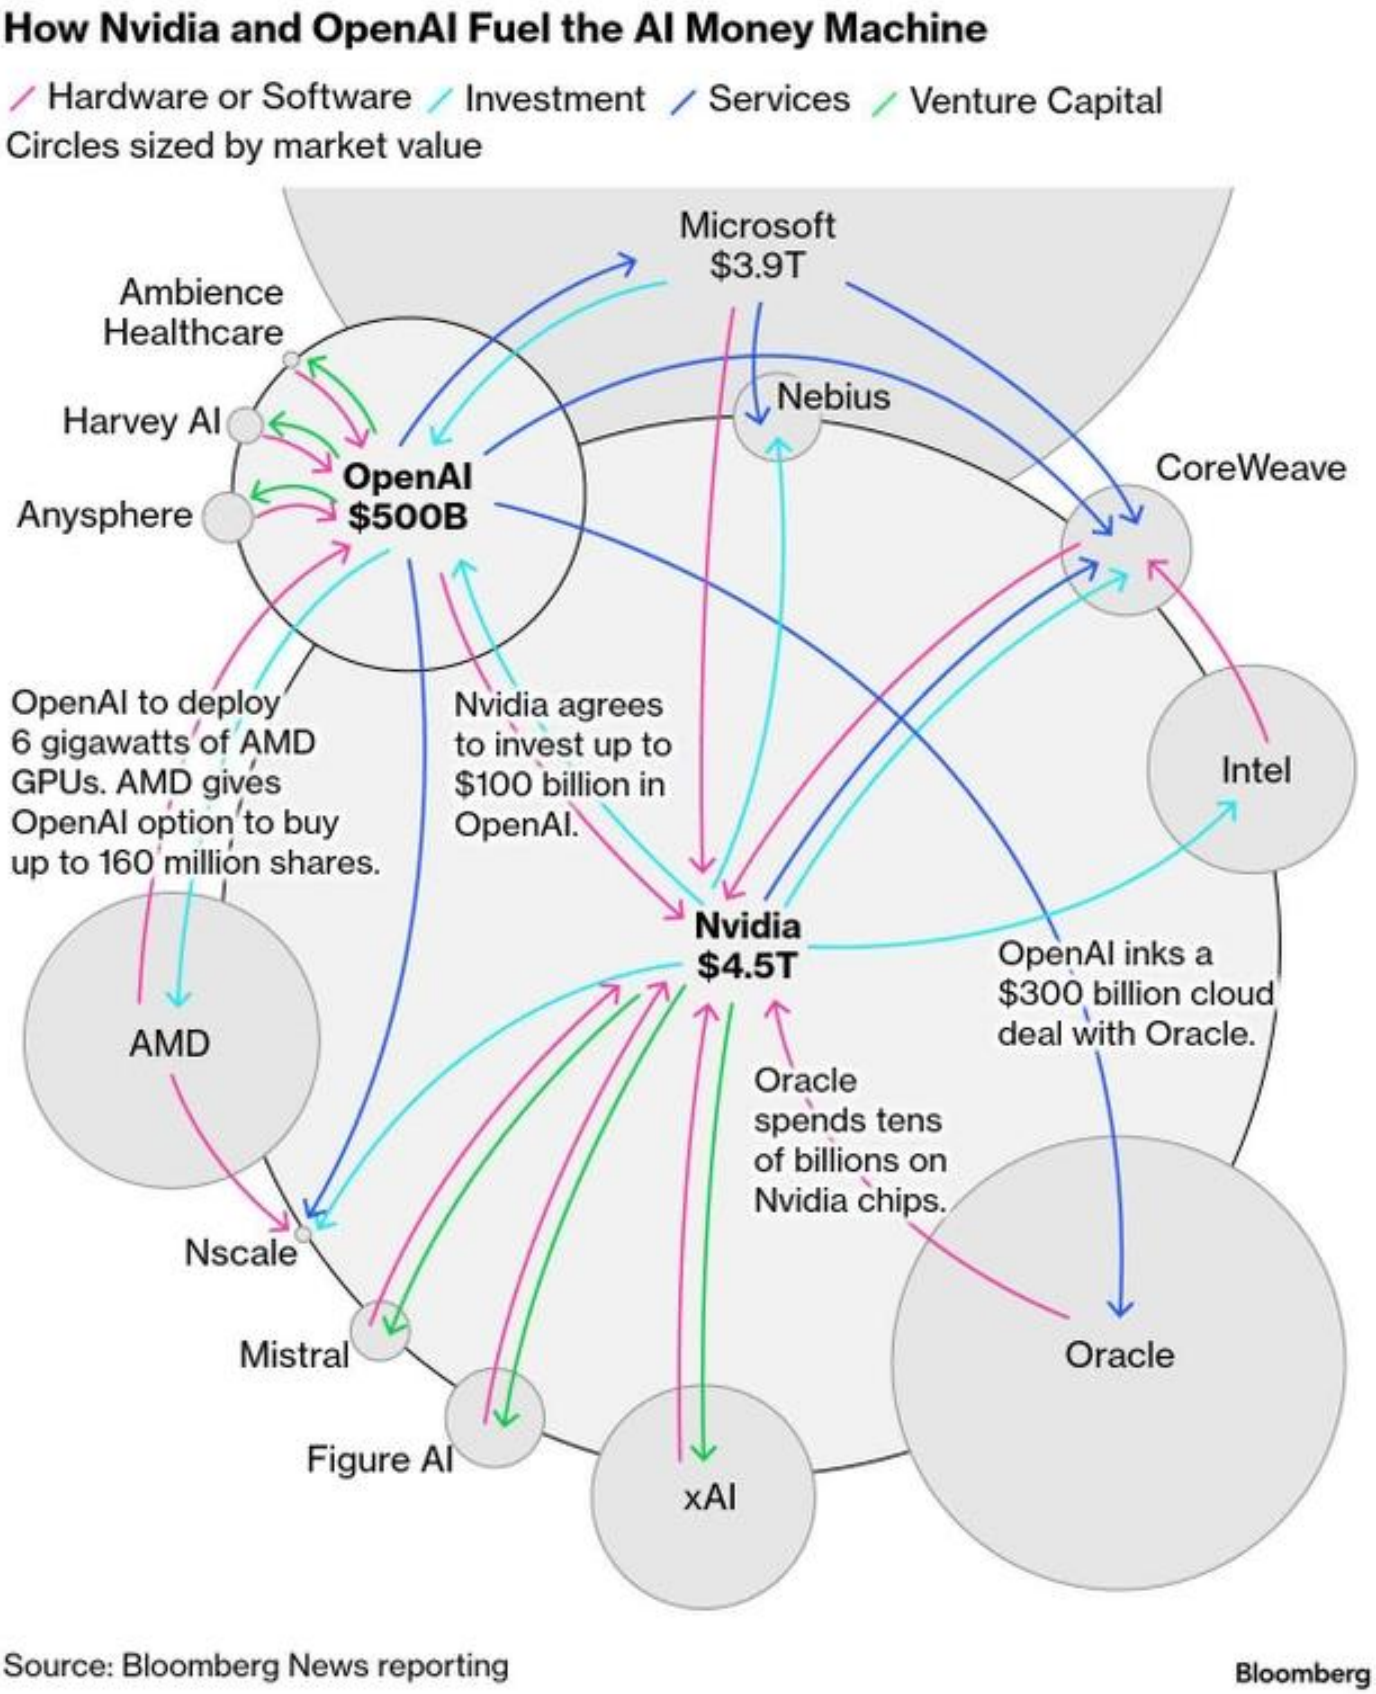

Since 2023, Meta has been releasing iterations of its own large language models (LLMs) under the name "Llama" to developers around the world, with Facebook and WhatsApp using the Llama 3 model in some

regions. With an open-weight configuration under a Community License and self-hosted deployment features, the Llama family was conceived for developers, startups, and organizations requiring self-hosting and deep control.

Earlier in April this year, Meta launched Llama 4, which had a somewhat lukewarm reception, with developers comparing it unfavourably to broad models such as OpenAI's GPT-40 and DEEPSEEK's, from Hangzhou DeepSeek Artificial Intelligence Co., Ltd. (DEEPSEEK) R1. In specialized tasks, it was deemed to have fallen behind specialist LLMs.

With its LLMs unable to gain significant traction much to the dissatisfaction of CEO Mark Zuckerberg, Meta has been enacting corrective measures under its chief AI officer Alexandr Wang, who was hired in June this year: 600 employees have been laid off from its bloated AI unit which is also being overhauled and restructured in order to compete with the likes of OpenAI in offering better AI models of utility.

One key factor eating away at the company's bottom line for some time now has been greater expenditure in compute infrastructure, which witnessed a sharp spike this year: at $48.3 billion so far, purchases of property and equipment were 216% higher than in the same period in the previous FY.

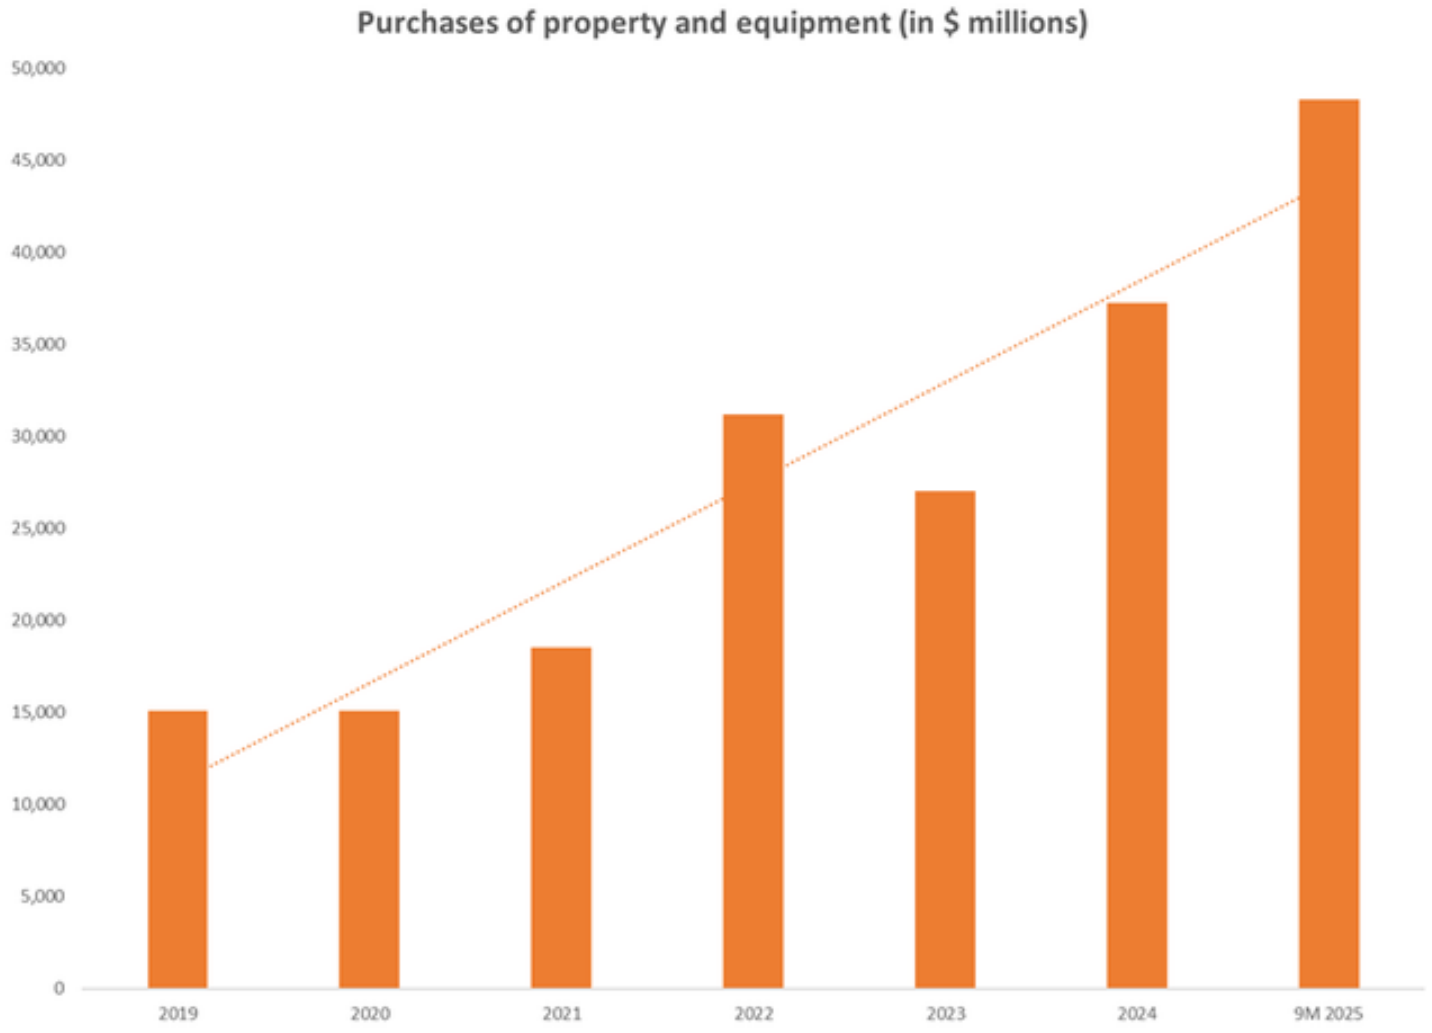

Source: Created by Sandeep G. Rao using data from Meta’s Financial stmt

In the forward guidance of its latest earnings release, the company states that capital expenditure in 2026 will be higher than in 2025 as it continues to expand infrastructure aggressively by both building its own capacity and contracting with third parties.

For over nine months, the company had inked a deal to build out a new data center in Louisiana (now called "Hyperion") to support its AI endeavor, which has ballooned from being worth $10 billion to $27 billion as of October. This 5 Gigawatt (GW) datacenter has an interesting funding pattern with asset management firm Blue Owl Capital: in exchange for Blue Owl providing private credit, the datacenter would be a joint venture between the two in which Meta would hold only a 20% stake. Meta would then lease the center from the joint venture for at least four years.

The company states that this is "long-term strategic flexibility." Considering how the "circular deal" machine is set up, this might be something akin to a hedge if the AI endeavor were to underperform. Also announced earlier in August was a $10 billion+ cloud deal with Alphabet Inc. (GOOG) (GOOGL), aka Google, over a six-year horizon, which is an interesting wrinkle: Hyperion's ownership structure likely incentivizes using native capacity over third parties (which the Google deal is quite likely to be). While details are scarce about the structuring of the deal, it remains to be seen how this deal would pan out in terms of costs for Meta: it would be entirely logical for Google to expect a baseline payment in such an agreement.

Forward Outlook

What makes the likes of a "high conviction" stock to most investors would be the capability to compound top-line growth with strong bottom line growth. In its current segmentation, at least, matching the previous FY's growth seems exceedingly unlikely.

Long-term observers draw parallels between the massive multi-trillion AI Hype and the dot-com bubble at the turn of the century. Unlike during the dot-com bubble, though, while today’s AI firms have tangible products and customers, their spending outpaces monetization. This implies that these aggressive capacity buildouts might turn out to be ultimately unprofitable. Since 2019 till the present, Meta has committed a total of over $192 billion in purchases of plant and equipment, with the latter years being dominated by datacenter-related expenses.

Given the blend of AI investment and ad spend relevance, Meta has numerous parallels with Google (which is the subject of an article designed with a highly similar framework). The question over the continuing feasibility of such massive AI-relevant investments will likely be a recurrent theme in the conversation around these two stocks. Unlike with Google, however, the lack of a substantial footprint in the AI market (at least presently) will likely be a headwind for the stock's valuation for the time being. Meanwhile (and like with Google), the gloomy outlook on future ad spends would likely be a more pressing concern for its main line of business.

SK hynix Q3 2025 Earnings Call Summary

Oct. 29, 2025 SL

Executive Summary

- Woo-Hyun Kim (CFO & Head of Finance) reported record-breaking results driven by explosive AI-related memory demand — especially HBM, DDR5, and enterprise SSD.
- Revenue, profit, and margins reached all-time highs. Management emphasized that AI is structurally reshaping memory demand across both DRAM and NAND, creating a multi-year "super-cycle" distinct from past upturns.
- SK hynix completed development of HBM4 and will begin mass shipments in Q4 2025, extending its technological lead. Strong AI infrastructure investment has sold out HBM capacity through 2027.

📊 Financial Highlights (CFO Woo-Hyun Kim)

Metric	Q3 2025	Q2 2025	Q/Q	Y/Y
Revenue	₩ 24.4 trillion	₩ 22.1 trillion	**+10 %**	**+39 %**
Operating Profit	₩ 11.4 trillion	₩ 9.2 trillion	**+24 %**	**+62 %**
Operating Margin	**47 %**	42 %	+5 pts	+7 pts
EBITDA	₩ 14.9 trillion	—		Margin 61 %
Net Income	₩ 12.6 trillion	—		Margin 52 %
Free Cash Flow	Strongly positive; +₩ 10.9 T cash increase			
Net Cash	₩ 3.8 trillion (net)	— (previously net debt)		

Cash & Investments: ₩ 27.9 trillion

Debt-to-Equity: 24 %, improved 1 pt Q/Q

This marks the first quarter ever with operating profit above ₩ 10 trillion.

💡 Market & Strategic Outlook (CFO Kim)

- AI Infrastructure Boom: Surging global investment in AI training and inference is driving structural demand for HBM, DDR5, and enterprise SSDs.
- Inference Shift: The AI market is evolving from training to inference, which massively increases token processing, driving exponential growth in memory usage across servers and edge devices.
- Memory Hierarchy Expansion: As AI workloads grow, systems are increasingly offloading data from HBM → DRAM → SSD, expanding demand for all tiers of memory.
- 2026 Forecast:
  - DRAM demand +20 % Y/Y (vs high-teens 2025)
  - NAND demand → high-teens % growth
- AI PC & Mobile: On-device AI accelerating content growth even in mid-tier smartphones and PCs

🚀 Product & Technology Updates

Segment	Speaker	Key Points
HBM	**Woo-Hyun Kim / HBM Head Jeong-Tae Kim**	– HBM4 development completed; production ready Sept 2025; mass shipments begin **Q4 2025**. – HBM4 supports **highest industry speed**, meets all customer specs. – HBM supply for 2026 finalized; sold out through 2027. – Custom HBM4E designs in progress; tighter long-term partnerships forming.
DRAM	**CFO Kim**	– Bit shipments ↑ high single digits Q/Q. – 1c-nm node mass production underway; will exceed 50 % of conventional DRAM output in 2026. – Demand led by **DDR5 128 GB modules** (doubled for 2 quarters).
NAND & SSD	**Head of NAND Kim Seok**	– Bit shipments ↓ mid-single-digits Q/Q (but ASP ↑ low teens %). – Strong **enterprise SSD** demand (+double digits Q/Q). – Ramping **321-layer TLC & QLC NAND**; expected > 50 % of NAND output by end 2026.
AI-Driven Demand Expansion	**CFO Kim**	– AI data-center growth boosting all memory segments. – Key value cache offloading and RAG (vector database) architectures increasing use of high-performance eSSD.

🏭 CapEx & Facilities (CFO Kim)

- CapEx 2026 ↑ significantly vs 2025 to meet demand.
- M15X fab: equipment installation underway → HBM ramp from 2026.
- Yong-in Fab 1: construction accelerated; possible pull-in from 2027 target.
- Indiana advanced packaging plant (US): groundwork ongoing.
- Goal: expand capacity while maintaining CapEx discipline and net-cash balance.

💬 Q&A Highlights

Theme	Speaker	Summary
HBM 2026 Contracts	*HBM Exec*	All major 2026 supply contracts finalized; HBM sold out through 2027; pricing supports high profitability.
HBM4 Performance	*HBM Exec*	Meets/exceeds customer specs; samples out ahead of competition; #1 supplier position maintained.
Memory Cycle	*Exec Response to BofA*	Current boom structurally different from 2017-18; AI driving broad, sustainable demand and supply constraints limit downturn risk.
eSSD Structural Shift	*NAND Head Kim Seok*	AI inference & RAG architectures creating sustained eSSD growth via vector DB and cache offloading.

Customer Behavior	*Exec to Korea Inv. Sec.*	Memory model shifting to **order-first, sell-later**; LTAs & pre-purchase POs for DRAM/NAND through 2026.
CapEx Scale	*CFO Kim*	2026 CapEx to "far outpace 2025"; M15X and Yong-in builds ramping; AI growth justifies higher investment.
Inventory	*Exec to SK Sec.*	DRAM inventory "extremely low" — DDR5 shipping straight from fab; healthy NAND levels.
DRAM Profitability	*Exec to JPMorgan*	DDR5 margin rising; could approach HBM, but HBM remains priority due to long-term contracts and stability.
Shareholder Returns	*CFO Kim*	No change to policy in first year of 3-yr plan; cash being reinvested into capacity and AI growth.

☑ Key Takeaways

- **All-time-high revenue & profit fueled by AI infrastructure spending.**
- **HBM4 ramp & HBM4E development strengthen #1 position in AI memory.**
- **Memory super-cycle likely to persist — broad-based, not speculative.**
- **Capacity sold out through 2027 for HBM and tight DRAM supply.**
- **CapEx rising but balance sheet remains strong (net cash ₩ 3.8 T).**
- **AI driving structural shift in DRAM + NAND demand via inference, RAG, and edge AI**