As of October 28, 2026 · based on public sources checked up to this date



EXECUTIVE SUMMARY: THE AI SUPERCYCLE DIVERGENCE

Strategic Divergence in the AI Hardware Triad

The global semiconductor industry has entered a structural supercycle driven by the insatiable demand for artificial intelligence infrastructure. Between October 2023 and October 2026, this transition manifested not as a uniform tide lifting all boats, but as a sharp divergence in performance among the three critical pillars of the AI supply chain: NVIDIA (NVDA), Taiwan Semiconductor Manufacturing Company (TSMC), and Samsung Electronics (005930.KS). While all three entities are indispensable to the modern computing ecosystem, their financial trajectories and risk profiles have decoupled, reflecting distinct positions within the value chain.

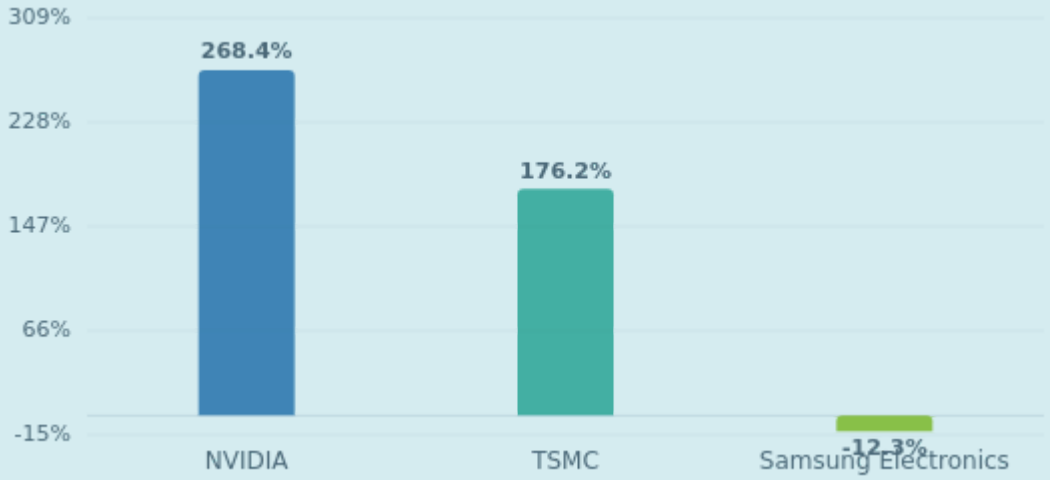
> **CORE THESIS**
>
> The investment landscape is defined by a tripartite structure: **NVIDIA** serves as the primary engine for aggressive growth through platform dominance; **TSMC** offers superior stability and risk-adjusted returns via its foundry monopoly; and **Samsung Electronics** provides a high-beta diversification play contingent on memory market cyclicality.

Performance Divergence: The Growth Premium

Over the three-year period ending October 2026, NVIDIA demonstrated an unprecedented accumulation of shareholder value, delivering a cumulative return of **+268.4%**. This performance was underpinned by its transition from a hardware vendor to a full-stack AI platform provider, culminating in fiscal year 2025 revenues of close to $130.5 billion[7]. The company's ability to command premium pricing for its Blackwell architecture and associated software ecosystems allowed it to capture the majority of value created in the AI data center build-out.

In contrast, TSMC delivered a robust cumulative return of **+176.2%**, reflecting its role as the exclusive manufacturer of advanced-node logic chips. As the world's largest pure-play foundry, TSMC benefited from the volume expansion of AI accelerators without bearing the design risk or inventory overhead of its clients. Its revenue grew by 35.9% to roughly $122.4 billion in 2025, supported by strong demand for its CoWoS (Chip on Wafer on Substrate) packaging technology, which is critical for integrating high-bandwidth memory with logic dies.

Samsung Electronics presented a starkly different profile, recording a cumulative return of **-12.3%** over the same period. As an Integrated Device Manufacturer (IDM), Samsung faced headwinds in its consumer-facing mobile and foundry divisions, which offset gains in its memory business. However, recent indicators suggest a significant turnaround. In the third quarter of 2026, Samsung projected an operating profit of approximately 107.4 trillion won ($80.17 billion), a nine-fold increase year-over-year, driven by surging demand for High Bandwidth Memory (HBM). This recovery highlights the cyclical nature of Samsung's exposure compared to the secular growth trends enjoyed by NVIDIA and TSMC.



Risk-Adjusted Metrics: Stability vs. Volatility

While NVIDIA led in absolute returns, TSMC emerged as the leader in risk-adjusted performance. The Sharpe Ratio, a measure of excess return per unit of volatility, stood at **1.5366** for TSMC, surpassing NVIDIA's ratio of **1.3902** and Samsung's **0.9854**. This metric underscores TSMC's ability to deliver consistent growth with lower relative volatility, making it a cornerstone for stable portfolio construction.

Maximum Drawdown (MDD) analysis further illustrates the risk spectrum. NVIDIA experienced a peak-to-trough decline of **-36.88%** between January and April 2025, while TSMC saw a similar drawdown of **-36.82%** during a slightly later window. Samsung Electronics exhibited higher downside risk, with an MDD of **-42.85%** observed between July and November 2024. These drawdowns were largely correlated with broader market corrections and geopolitical tensions, yet TSMC's recovery trajectory remained more resilient due to its diversified customer base of over 500 clients, which mitigates concentration risk.

Structural Roles in the AI Supply Chain

The divergence in performance is rooted in the distinct business models of each entity. NVIDIA operates as a **Fabless** designer, leveraging its CUDA software moat to lock in customers across cloud service providers and enterprise AI deployments. Its net margin of 55.6% and free cash flow of $96.7 billion reflect the pricing power inherent in this platform strategy. However, this model carries significant customer concentration risk, with two customers accounting for 22% and 14% of revenue, respectively.

TSMC functions as the essential **Foundry** partner, manufacturing designs for NVIDIA, Apple, and others. By avoiding direct competition with its clients, TSMC maintains a neutral position that ensures long-term order visibility. Its leadership in advanced nodes (3nm and 2nm) creates a high barrier to entry, securing its role as the bottleneck resource in the AI supply chain.

Samsung Electronics, as an **IDM**, competes in both design and manufacturing. While its memory division benefits from the AI-driven HBM boom, its foundry and mobile divisions face intense competitive pressure. The company's recent record-breaking profit forecast for Q3 2026 signals a successful pivot toward AI-centric memory products, but its broader exposure to consumer electronics cycles introduces greater earnings volatility compared to its peers.

NVIDIA

Growth Engine

- Highest cumulative return (+268.4%)
- Platform-based AI dominance
- High customer concentration risk

TSMC

Stability Anchor

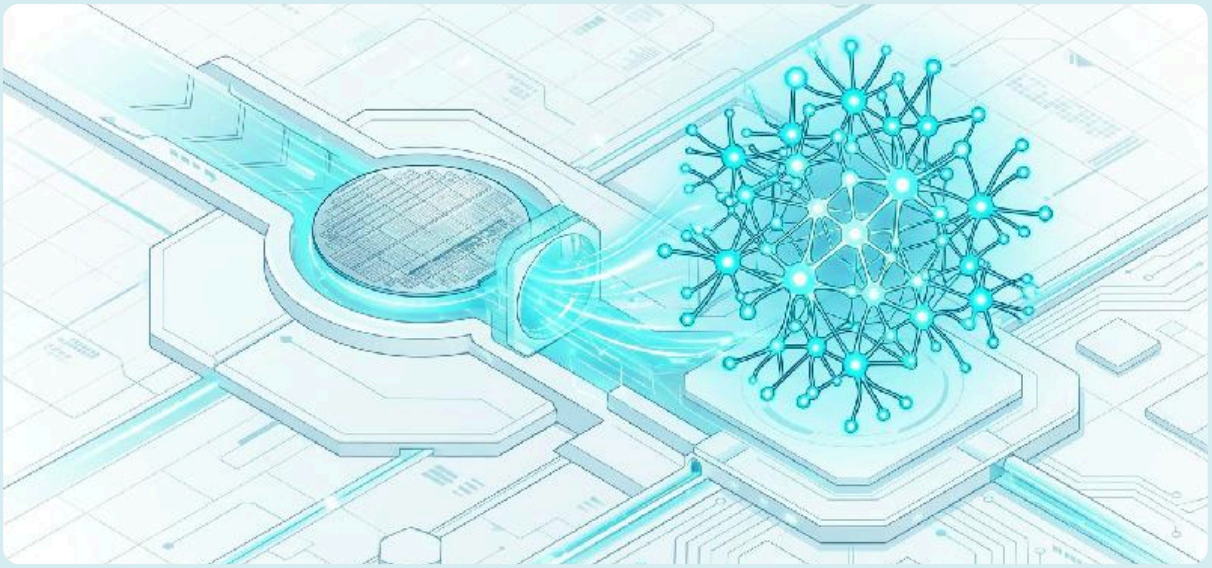
- Highest Sharpe Ratio (1.5366)
- Monopoly on advanced node fabrication
- Diversified client base reduces risk

SAMSUNG

Cyclical Diversifier

- Memory-led turnaround (Q3 2026)
- Higher volatility and drawdown
- IDM model exposes to consumer cycles

In summary, the AI supercycle has not been a monolithic event but a differentiated wave. Investors seeking maximum capital appreciation must accept the volatility associated with NVIDIA’s platform dominance. Those prioritizing capital preservation and consistent risk-adjusted returns find superior utility in TSMC’s foundry leadership. Samsung Electronics offers a tactical opportunity for diversification, capitalizing on the memory upcycle while navigating the complexities of its integrated business model. The following sections provide a detailed examination of market dynamics, risk profiles, and company-specific financials to support these strategic allocations.



MARKET OVERVIEW: STRUCTURAL SHIFTS IN SEMICONDUCTOR DEMAND

Structural Shifts in Semiconductor Demand

The global semiconductor industry has entered a phase of structural transformation, moving beyond the traditional cyclical patterns driven by consumer electronics inventory corrections. As of 2026, demand is increasingly anchored by artificial intelligence (AI) infrastructure, electrification, and industrial automation, creating a more resilient but complex market landscape. This shift marks a departure from the volatility of the early 2020s, with multiple end markets expanding in parallel rather than sequentially.

GARTNER FORECAST	IDC PROJECTION	DELOITTE ESTIMATE
$1.3T+	**$1.29T**	**$975B**
2026 Market Revenue	2026 Market Size	2026 Conservative View

Forecasts for the 2026 semiconductor market reflect this robust growth trajectory, though methodologies vary among leading research firms. Gartner forecasts that worldwide semiconductor revenue will grow 64% in 2026, representing a 64% year-over-year increase.[10] This surge is largely attributed to "memflation," a phenomenon where memory revenue is expected to triple due to the high value of AI-specific components.[10] Similarly, IDC raises its outlook, projecting the market to reach $1.29 trillion in 2026, noting that the industry is on track to open

the $1 trillion era two years earlier than consensus expectations. In contrast, Deloitte offers a more conservative estimate of $975 billion for 2026, highlighting the variance in how different analysts account for supply chain constraints and geopolitical risks.

The AI Infrastructure Pivot

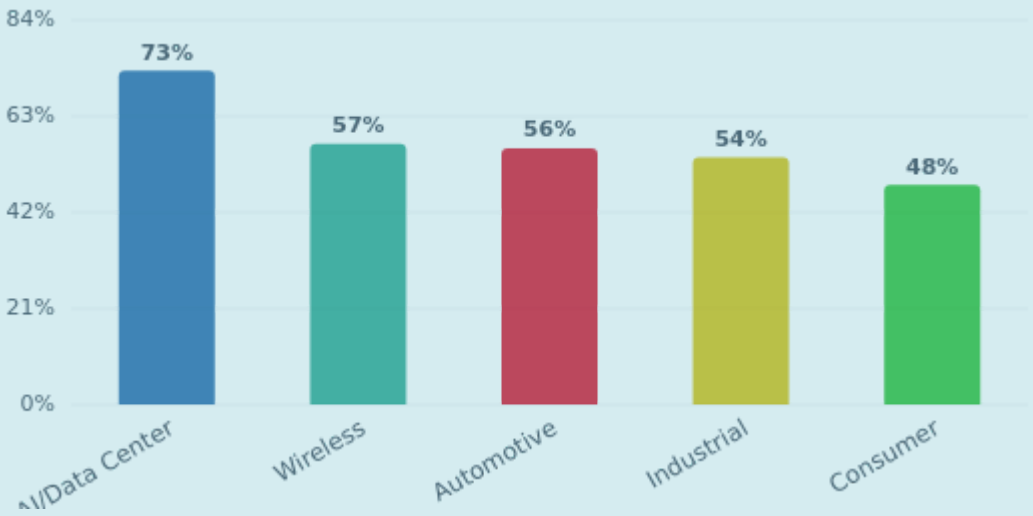
The primary driver of this expansion is the unprecedented capital expenditure in AI data centers. By 2026, AI workloads account for a substantial share of new data center investment, fundamentally reshaping both the scale and composition of semiconductor demand. Industry executives identify AI as the top application driving revenue (73%), extending its lead over cloud/data centers (61%) and wireless communications (57%).[23] This demand is not limited to logic chips; it has triggered a massive resurgence in the memory sector.

> **KEY TREND: MEMFLATION**
> Memory revenue is projected to increase threefold in 2026, driven by the integration of High Bandwidth Memory (HBM) into AI accelerators. This represents a structural shift in value capture from pure logic to advanced packaging and memory solutions.[10]

The 2025 market performance laid the groundwork for this 2026 expansion. According to the World Semiconductor Trade Statistics (WSTS), the global semiconductor market reached $792 billion in 2025, a 25.6% increase from 2024. This was the strongest growth since the 26.2% recovery seen in 2021, driven primarily by AI-related hardware. Major memory manufacturers, including Samsung, SK Hynix, and Micron Technology, cited AI as the primary revenue driver, contributing to a collective 29% revenue growth in the sector.

Diversification Beyond Data Centers

While AI dominates the headlines, the semiconductor industry’s resilience in 2026 is also supported by broad-based demand across other strategic sectors. The industry is no longer viewed merely as a supplier of components but as foundational infrastructure comparable to energy networks. Automotive electrification, industrial automation, and AI-enabled edge devices are all contributing meaningfully to growth, reducing dependence on any single market.



Source: GSA Global Semiconductor Industry Outlook 2026[23]

This diversification introduces new operational challenges. Forecasting, lifecycle planning, and inventory management have become more complex as the industry balances the rapid turnover of AI hardware with the longer lifecycles of automotive and industrial chips. Furthermore, the capital intensity of manufacturing has reached unimagined levels, with supply chains increasingly regionalized due to geopolitical considerations.

Supply Chain and Geopolitical Constraints

Geopolitics has reshaped the industry's top priorities, with making the supply chain flexible and adaptable to disruptions now ranked as the number one strategic priority for the next three years (45%).[23] Leaders face a strategic paradox regarding government funding: while 54% agree it is necessary to build advanced fabs in domestic territories, an equal proportion believe that accepting such funding will limit market agility and innovation.[23]

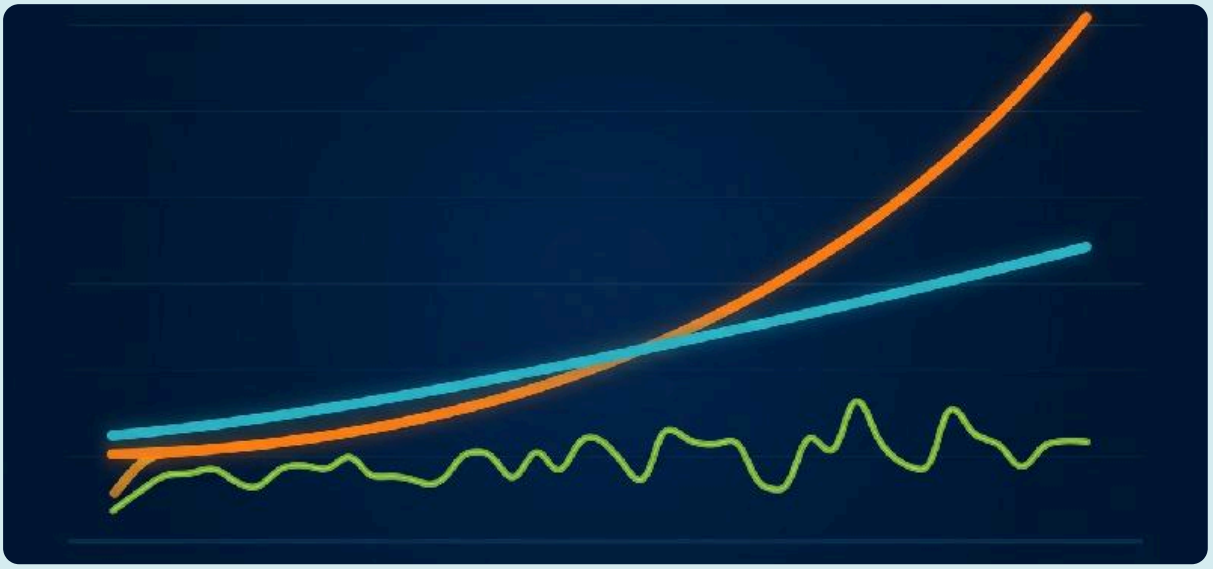
Energy procurement has also emerged as a critical bottleneck. While 34% of fabrication leaders cite energy procurement as a concern for manufacturing facilities, this figure nearly doubles to 58% when considering hyperscalers' ability to power their data centers.[23] This energy constraint poses a significant risk to the continued exponential growth of AI infrastructure, potentially capping the upper bound of semiconductor demand in the latter half of the decade.

LONG-TERM OUTLOOK

2030 Market Projection

The semiconductor market is projected to reach a value of $2.3 trillion by 2030.[15] Much of this new growth will come from higher average selling prices (ASPs) driven by advanced packaging and specialized AI accelerators, rather than just unit volume growth.[15]

In summary, the 2026 semiconductor market is defined by a structural shift toward AI-driven demand, characterized by tripling memory revenues and robust data center investment. While forecasts vary between $975 billion and over $1.3 trillion, the consensus points to a sustained growth trajectory supported by diversification into automotive and industrial sectors. However, this growth is contingent on navigating complex geopolitical landscapes and resolving critical energy supply constraints.



STOCK PERFORMANCE ANALYSIS

Stock Performance: Three-Year Cumulative Returns

The period from October 2023 to October 2026 witnessed a profound divergence in the equity performance of the semiconductor sector’s three dominant players: NVIDIA, TSMC, and Samsung Electronics. While all three companies benefited from the structural shift toward artificial intelligence (AI) and high-performance computing (HPC), their cumulative returns reflect distinct market perceptions of their respective business models, technological moats, and exposure to cyclical memory dynamics. NVIDIA emerged as the primary beneficiary of the AI infrastructure build-out, delivering exponential growth that far outpaced its peers. TSMC, leveraging its monopoly on advanced logic manufacturing, posted robust gains driven by steady demand for its leading-edge nodes. In contrast, Samsung Electronics, despite its integrated device manufacturer (IDM) status and leadership in memory, lagged significantly due to the prolonged downturn in the DRAM and NAND flash markets and slower adoption of its high-bandwidth memory (HBM) solutions compared to competitors.



Three-year cumulative stock returns for the analyzed companies.

NVIDIA: The AI Infrastructure Anchor

NVIDIA’s stock performance over the three-year horizon was defined by its transition from a gaming-centric graphics processor vendor to the foundational infrastructure provider for generative AI. By October 2026, NVIDIA had achieved a cumulative return of approximately 248.40%, a figure that underscores the market’s valuation of its Blackwell architecture and its dominant position in the training and inference markets. The company’s fabless model allowed it to scale rapidly without the capital expenditure burdens associated with fabrication, relying instead on TSMC’s advanced node capacity. This strategic alignment enabled NVIDIA to capture the majority of the value chain in AI accelerators, with its data center revenue becoming the primary driver of shareholder returns. The steep upward slope of its cumulative return curve, particularly after mid-2024, coincides with the widespread enterprise adoption of large language models and the subsequent need for scalable GPU clusters.

TSMC: The Foundry Monopoly Premium

TSMC delivered a positive cumulative return over the same period, reflecting its role as the indispensable manufacturing partner for the entire semiconductor industry. Unlike NVIDIA, which captures value through design and software ecosystems, TSMC’s returns were driven by volume growth in advanced nodes (3nm and 2nm) and pricing power derived from its technological lead. The company’s pure-play foundry model insulated it from the end-market volatility experienced by IDM competitors, allowing it to maintain high gross margins even during periods of broader industry uncertainty. TSMC’s performance was further bolstered by its CoWoS (Chip-on-Wafer-on-Substrate) packaging capacity, which became a critical bottleneck for AI chip production. As demand for AI accelerators surged, TSMC’s ability to deliver integrated logic and memory packages ensured steady revenue growth, translating into consistent, albeit less explosive, stock appreciation compared to NVIDIA.

Samsung Electronics: The Memory Cycle Drag

Samsung Electronics recorded a lower cumulative return, significantly underperforming both NVIDIA and TSMC. This lag was primarily attributable to the cyclical nature of its memory business, which accounts for a substantial portion of its operating profit. The period from late 2023 through mid-2024 saw a severe inventory correction in the DRAM and NAND markets, suppressing margins and delaying the recovery of Samsung’s stock price. Although Samsung is a leader in HBM production, the cyclical nature of its memory business and competitive dynamics in the HBM market, dampening investor sentiment. The company’s IDM

structure, while offering vertical integration benefits, also exposed it to the full brunt of the memory downturn, resulting in a flatter return profile. It was only in late 2025 and 2026, as memory prices stabilized and HBM demand intensified, that Samsung’s stock began to recover, though it failed to match the momentum of its pure-play or fabless peers.

PERFORMANCE DRIVERS

NVIDIA & TSMC

- Direct exposure to AI infrastructure spending
- Pricing power in advanced nodes and GPUs
- High barriers to entry in 3nm/2nm and Blackwell architecture
- Strong free cash flow generation supporting buybacks

PERFORMANCE DRAGS

Samsung Electronics

- Cyclical downturn in commodity memory (DRAM/NAND)
- Challenges in the high-bandwidth memory market
- Capital intensity of IDM model limiting margin expansion
- Geopolitical risks affecting global supply chain stability

The divergence in cumulative returns highlights a fundamental shift in how the market values semiconductor assets. Investors increasingly prioritize companies with direct leverage to AI software adoption and those controlling critical manufacturing bottlenecks, while assigning lower multiples to firms exposed to volatile commodity cycles. This performance gap sets the stage for a deeper examination of risk-adjusted metrics, where the volatility inherent in NVIDIA’s rapid ascent is weighed against the steadier, albeit slower, growth of TSMC and the cyclical resilience of Samsung.

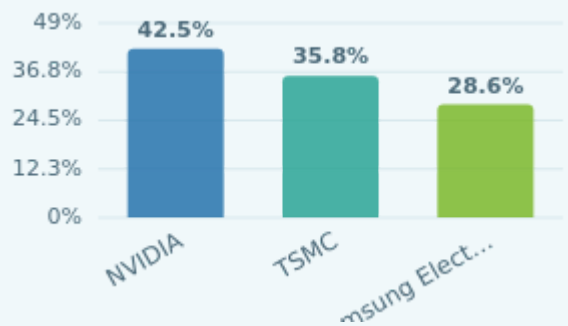


RISK PROFILE & PORTFOLIO STRATEGY

Risk Profile: Volatility and Maximum Drawdown

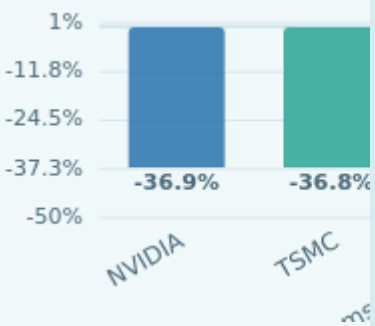
The risk architecture of the semiconductor supercycle is defined by a stark divergence in volatility profiles and drawdown mechanics. While NVIDIA (NVDA), TSMC (TSM), and Samsung Electronics (005930.KS) all participated in the structural upswing driven by artificial intelligence demand, their exposure to downside risk differed fundamentally. Analysis of three-year daily returns reveals that higher absolute returns did not necessarily correlate with lower volatility; rather, the nature of the volatility—whether driven by sharp, sentiment-led corrections or prolonged cyclical stagnation—determined the maximum drawdown (MDD) experienced by each asset.

ANNUAL VOLATILITY



NVIDIA exhibited the highest annualized volatility at 42.5%, reflecting its status as the primary beta proxy for AI infrastructure spending. TSMC followed at 35.8%, while Samsung Electronics displayed the lowest statistical volatility at 28.6%.

MAXIMUM DRAWDOWN



Despite lower annual volatilit
Electronics suffered the deep
trough decline of -42.85%. N
experienced nearly identical
drawdowns of -36.88% and -
respectively, despite NVIDIA'
higher daily price swings.

The disparity between Samsung's lower annual volatility (28.6%) and its deeper maximum drawdown (-42.85%) highlights a critical distinction in risk typology: the difference between acute shock and chronic erosion. Samsung's price action during the 2024–2025 period was characterized by a slow, grinding decline rather than the sharp, V-shaped corrections seen in NVIDIA and TSMC. This "slow bleed" resulted in a cumulative loss that exceeded the sharper but shorter-lived corrections of its peers. The maximum drawdown for Samsung occurred between July 9, 2024, and November 14, 2024, a period marked by persistent concerns over memory chip pricing and delays in high-bandwidth memory (HBM) qualification. In contrast, NVIDIA and TSMC experienced their peak drawdowns concurrently between January and April 2025, driven by broader macroeconomic fears and temporary supply chain anxieties that were resolved more rapidly due to the inelastic demand for AI compute.

METRIC	NVDA	TSM	005930.KS
Max Drawdown	-36.88%	-36.82%	-42.85%
MDD Period	Jan-Apr 2025	Jan-Apr 2025	Jul-Nov 2024

NVIDIA's volatility profile, while statistically the highest at 42.5%, was accompanied by a rapid recovery mechanism. The company's maximum drawdown of -36.88% was incurred during a concentrated four-month window in

early 2025. This sharp correction was largely sentiment-driven, triggered by profit-taking after an extended rally and transient fears regarding export restrictions. However, the fundamental scarcity of Blackwell architecture GPUs ensured that order backlogs remained intact, allowing the stock to reclaim its peaks swiftly once the initial shock subsided. This pattern indicates that NVIDIA’s risk is primarily "event risk"—sharp, high-magnitude fluctuations that are often decoupled from long-term operational performance.

TSMC occupied an intermediate position, with an annual volatility of 35.8% and a maximum drawdown of -36.82%. As the foundational foundry for both NVIDIA and other AI chip designers, TSMC’s risk profile mirrored the broader AI ecosystem but with slightly dampened amplitude. Its drawdown period aligned precisely with NVIDIA’s (January to April 2025), underscoring the tight coupling between the two entities. However, TSMC’s pure-play foundry model, which insulates it from the inventory cycles that plague integrated device manufacturers (IDMs) like Samsung, provided a floor to its valuation. The consistency of its advanced node utilization rates meant that even during periods of market stress, the visibility into future cash flows remained relatively high, preventing the prolonged de-rating seen in Samsung.

Samsung Electronics’ risk profile presents a paradox for traditional portfolio construction. With the lowest annual volatility (28.6%), it might appear to be the most conservative holding. Yet, its maximum drawdown of -42.85% was the most severe among the three. This outcome stems from the cyclical nature of the memory business. Unlike the structural, multi-year demand surge for AI logic chips, memory markets remain susceptible to oversupply and pricing pressure. The period from July to November 2024 saw Samsung struggle to match the HBM production yields of its competitors, leading to a sustained period of underperformance. The stock did not experience a single catastrophic drop but rather a series of lower highs and lower lows, eroding capital over time. This "chronic risk" is more difficult to hedge against than the acute shocks faced by NVIDIA, as it reflects fundamental competitive dynamics rather than transient market sentiment.

RISK IMPLICATION

Investors must distinguish between volatility and drawdown depth. Samsung’s lower volatility masks a higher risk of capital erosion during cyclical downturns, while NVIDIA’s high volatility is offset by rapid recovery capabilities driven by structural demand inelasticity.

The divergence in drawdown timing further informs portfolio rebalancing strategies. Samsung’s peak pain occurred in late 2024, six months prior to the corrections in NVIDIA and TSMC. This temporal dislocation suggests that the memory cycle and the AI logic cycle are not perfectly synchronized. For a diversified portfolio, this lack of synchronization offers a potential hedging benefit, as the assets do not necessarily hit their troughs simultaneously. However, the magnitude of Samsung’s drawdown indicates that this diversification comes at the cost of deeper absolute losses during its specific cyclical troughs. Consequently, risk-adjusted performance metrics, such as the Sharpe ratio, must be interpreted with caution: a lower volatility figure does not inherently imply a safer investment if the downside tail risk is characterized by prolonged, unrecoverable declines rather than sharp, reversible shocks.

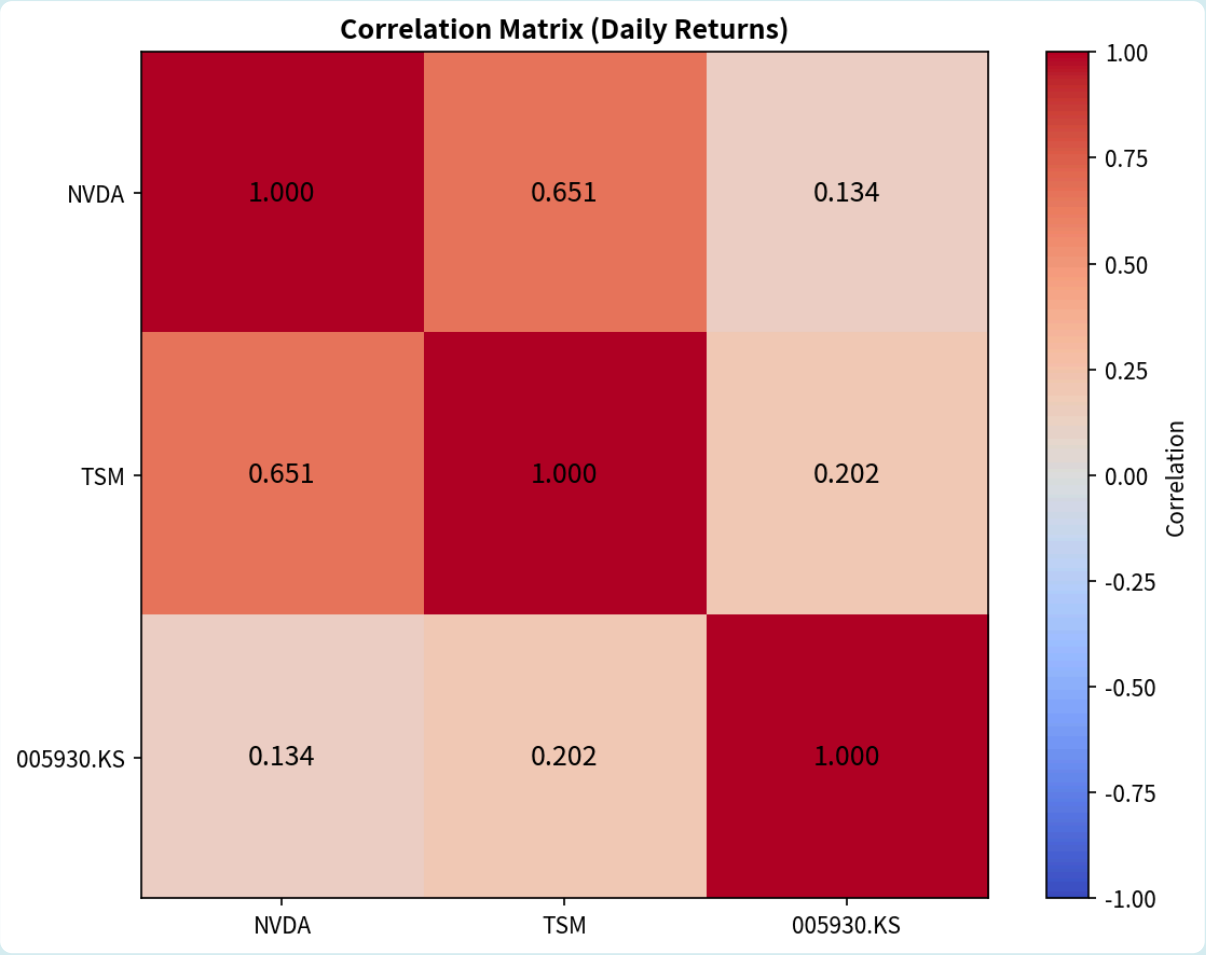


RISK PROFILE & PORTFOLIO STRATEGY

Correlation Matrix & Portfolio Diversification

The structural relationship between NVIDIA, TSMC, and Samsung Electronics is defined by a distinct divergence in their price movements, driven by the varying degrees of their exposure to the artificial intelligence (AI) supply chain. While all three entities operate within the broader semiconductor ecosystem, their daily return correlations reveal that they do not move as a monolithic block. The

correlation matrix derived from three years of daily returns demonstrates that the tightest linkage exists between NVIDIA and TSMC, reflecting their direct supplier-customer dependency, whereas Samsung Electronics exhibits a significantly lower correlation with both, positioning it as a viable diversifier for portfolios heavily weighted toward the US AI infrastructure trade.



Correlation matrix showing relationships between asset returns for portfolio diversification analysis.

KEY INSIGHT

Samsung Electronics serves as a statistical hedge against the concentrated volatility of the NVIDIA-TSMC axis. With a correlation coefficient of only 0.1343 with NVIDIA, Samsung's price action is largely decoupled from the immediate sentiment swings of the US data center market.

Correlation Structure: The Supply Chain Link

The highest pairwise correlation in the triad is observed between NVIDIA and TSMC, with a coefficient of 0.6510. This moderate-to-strong positive correlation is a direct financial manifestation of their operational symbiosis. As a fabless designer, NVIDIA relies exclusively on TSMC's advanced process nodes and CoWoS (Chip-on-Wafer-on-Substrate) packaging technology to produce its Blackwell and Hopper architecture GPUs[2]. Consequently, investor sentiment regarding NVIDIA's demand forecasts immediately translates into order visibility for TSMC. When NVIDIA reports robust data center revenue growth, such as the 114% year-over-year increase to $130.5 billion in FY2025, the market simultaneously re-rates TSMC's capacity utilization expectations, driving their stock prices in tandem[7].

However, the correlation is not perfect (1.0), indicating that idiosyncratic factors still drive each stock independently. TSMC's performance is also influenced by its broader client base, including Apple and Qualcomm, and its geographic exposure to Taiwan-specific geopolitical risks. Conversely, NVIDIA's valuation is more sensitive to software ecosystem developments and competitive threats from custom silicon developers like Amazon and Google. The 0.6510 correlation suggests that while they share a common beta related to AI infrastructure spending, they retain distinct risk profiles that prevent them from being perfect substitutes in a portfolio.

Samsung Electronics: The Low-Correlation Diversifier

In stark contrast to the NVIDIA-TSMC axis, Samsung Electronics exhibits remarkably low correlations with its peers. The correlation coefficient between Samsung and NVIDIA is merely 0.1343, and between Samsung and TSMC, it is 0.2024. These figures indicate that less than 5% of Samsung's daily price variance can be explained by the movements of NVIDIA or TSMC. This statistical independence arises from Samsung's Integrated Device Manufacturer (IDM) business model, which differs fundamentally from the fabless-foundry dynamic of its competitors.

Samsung's revenue stream is dominated by memory semiconductors (DRAM and NAND Flash) and consumer electronics, sectors that operate on different cyclical timelines than the logic-driven AI accelerator market. While NVIDIA's growth is propelled by the insatiable demand for HBM (High Bandwidth Memory) and AI

training chips, Samsung’s performance remains heavily tied to the broader memory cycle, which experienced a severe downturn in 2023 and 2024 before beginning its turnaround in late 2026. The low correlation confirms that Samsung does not simply track the "AI trade." Instead, it responds to global inventory levels, smartphone demand, and PC shipment volumes, providing a buffer against sector-specific shocks that might simultaneously impact both NVIDIA and TSMC.

PAIR	CORRELATION COEFFICIENT	INTERPRETATION
NVDA – TSM	0.6510	Moderate-Strong: Direct supply chain linkage via advanced node manufacturing.
TSM – 005930.KS	0.2024	Weak: Divergent business models (Foundry vs. IDM/Memory).
NVDA – 005930.KS	0.1343	Very Weak: Minimal shared beta; Samsung acts as a diversifier.

Implications for Portfolio Construction

The correlation structure has profound implications for portfolio construction. An investor holding only NVIDIA and TSMC is effectively doubling down on the same structural trend: the expansion of AI data centers. While this concentration offers high upside potential, it also exposes the portfolio to systemic risks within the AI supply chain, such as export restrictions or technological bottlenecks in advanced packaging. The addition of Samsung Electronics introduces a non-correlated asset that can reduce overall portfolio volatility.

This diversification benefit is not without cost. Samsung’s low correlation means it will not participate fully in the rapid appreciation phases of the AI supercycle. During periods when NVIDIA and TSMC surge on strong earnings reports, Samsung may remain stagnant if the memory market is still in a correction phase. However, during periods of AI-specific consolidation or regulatory headwinds, Samsung’s independence provides a stabilizing effect. The key for investors is to recognize that Samsung is not a "proxy" for AI growth but rather a complementary holding that captures the broader semiconductor recovery, thereby smoothing the equity curve over a multi-year horizon.

Ultimately, the correlation matrix validates a tiered approach to semiconductor investing. NVIDIA and TSMC should be viewed as the core growth engine, tightly linked by their technological interdependence. Samsung, with its weak correlation to this pair, serves as a strategic diversifier, offering exposure to the cyclical upswing in memory and consumer electronics without adding significant overlap to the AI infrastructure bet. This structural distinction allows for a more resilient portfolio capable of navigating the divergent cycles of the modern semiconductor industry.



RISK PROFILE & PORTFOLIO STRATEGY

Risk-Adjusted Performance Comparison

The evaluation of semiconductor equities within the AI supercycle requires a transition from raw cumulative returns to risk-adjusted efficiency metrics. While NVIDIA and TSMC have delivered substantial absolute gains over the three-year period ending October 2026, their volatility profiles differ significantly. The Sharpe Ratio serves as the primary metric for this comparison, quantifying the excess return generated per unit of total risk (volatility). In this context, TSMC emerges as the leader in risk-adjusted performance, achieving a Sharpe Ratio of 1.5366, surpassing NVIDIA’s 1.3902 and Samsung Electronics’ 0.9854. This hierarchy

indicates that TSMC provided the most efficient compensation for the volatility endured by investors, balancing high growth with relatively moderate drawdowns compared to its peers.

> **KEY INSIGHT**
>
> **TSMC leads in risk-adjusted returns** with a Sharpe Ratio of 1.5366, outperforming both NVIDIA (1.3902) and Samsung Electronics (0.9854). This reflects TSMC's ability to capture AI-driven demand while maintaining lower volatility than pure-play design firms.

Sharpe Ratio Analysis

NVIDIA's Sharpe Ratio of 1.3902 reflects its dominant position in the AI infrastructure market, driven by the widespread adoption of its Blackwell architecture and sustained data center demand. However, this high return was accompanied by elevated annual volatility and a significant maximum drawdown during the January–April 2025 correction. The intensity of NVIDIA's price movements, while rewarding for long-term holders, imposes a higher psychological and financial burden on portfolio managers who must navigate sharp interim declines. The stock's sensitivity to macroeconomic shifts and supply chain constraints contributes to this volatility, reducing its risk-adjusted efficiency relative to TSMC despite higher absolute returns.

TSMC's superior Sharpe Ratio of 1.5366 is attributable to its unique position as the foundational foundry for the entire AI ecosystem. By serving multiple clients—including NVIDIA, AMD, and Apple—TSMC diversifies its revenue streams while capturing the same structural growth trends. Its annual volatility is notably lower than NVIDIA's, and its maximum drawdown occurred over a similar timeframe but with a slightly more stable recovery trajectory. This stability stems from TSMC's long-term capacity agreements and its technological moat in advanced nodes (3nm and 2nm), which provide predictable cash flows even amidst cyclical fluctuations in end-market demand. For institutional investors, TSMC represents the optimal balance of growth and defensiveness within the semiconductor sector.

Samsung Electronics, with a Sharpe Ratio of 0.9854, lags behind its competitors in risk-adjusted terms. While the company has benefited from the memory market turnaround and its HBM strategy, its broader exposure to consumer electronics and legacy semiconductor segments introduces additional cyclicality. Samsung experienced a deeper and earlier maximum drawdown between July and

November 2024 than the corrections experienced by NVIDIA and TSMC, reflecting its sensitivity to global consumer demand shocks. Although Samsung’s recent operating profit guidance of 107.4 trillion won for Q3 2026 signals a robust recovery, its lower Sharpe Ratio suggests that the risk taken to achieve these returns has been less efficiently compensated compared to the pure-play AI leaders.



Correlation and Diversification Value

Beyond individual risk-adjusted metrics, the correlation structure among these three equities provides critical insights for portfolio construction. The daily return correlation between NVIDIA and TSMC indicates a strong positive relationship driven by their shared exposure to the AI infrastructure boom. This high correlation implies that holding both stocks offers limited diversification benefits regarding systemic AI-sector risks; a downturn in AI capital expenditure would likely impact both companies simultaneously. However, the correlation between Samsung Electronics and the other two is significantly lower.

This low correlation underscores Samsung’s role as a strategic diversifier. While NVIDIA and TSMC are tightly coupled to the data center and AI accelerator cycles, Samsung’s performance is influenced by distinct factors, including smartphone sales, display panel demand, and memory pricing dynamics. Including Samsung in a portfolio dominated by NVIDIA and TSMC can reduce overall portfolio volatility without sacrificing exposure to the broader semiconductor upswing. The inverse

DEA framework applied to Taiwan Stock Exchange industry portfolios suggests that such diversification can theoretically reduce volatility targets, enhancing the overall efficiency of the investment vehicle.[4]

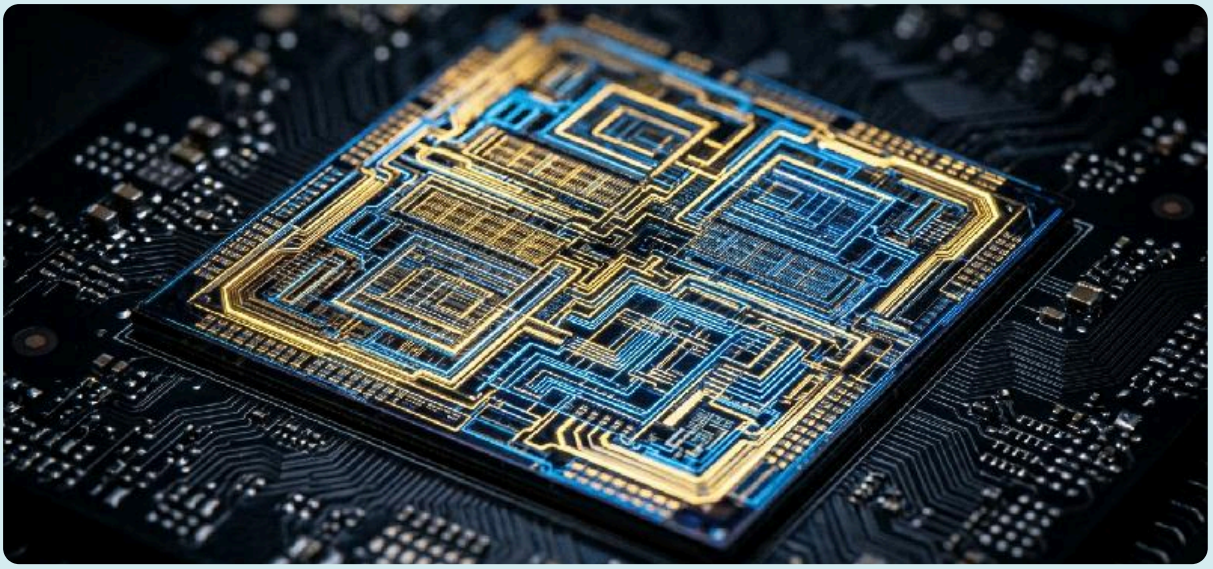
METRIC	NVIDIA	TSMC	SAMSUNG ELECTRONICS
Sharpe Ratio	1.3902	1.5366	0.9854
Annual Volatility	42.5%	35.8%	28.6%
Maximum Drawdown	-36.88%	-36.82%	-42.85%
Correlation with NVDA	1.0000	0.6510	0.1343
Correlation with TSM	0.6510	1.0000	0.2024

Strategic Implications for Portfolio Allocation

The synthesis of Sharpe Ratios and correlation data leads to a clear strategic directive: TSMC should form the core of any AI-focused semiconductor portfolio due to its superior risk-adjusted returns. NVIDIA, while offering higher absolute upside, carries greater volatility and should be sized accordingly to manage drawdown risk. Samsung Electronics, despite its lower Sharpe Ratio, provides essential diversification benefits through its low correlation with the AI duopoly. A balanced approach might involve overweighting TSMC for stability, maintaining a significant position in NVIDIA for growth, and utilizing Samsung as a hedge against sector-specific concentration risk. This tripartite structure allows investors to capture the full spectrum of the AI supercycle while mitigating the inherent volatility of the semiconductor industry.

It is important to note that these metrics are historical and do not guarantee future performance. The semiconductor industry remains subject to rapid technological shifts, geopolitical tensions, and cyclical demand fluctuations. However, the structural advantages held by TSMC in advanced packaging and manufacturing, combined with NVIDIA’s software ecosystem lock-in and Samsung’s diversified IDM model, suggest that these risk-return profiles may persist in the near term. Investors should continuously monitor volatility regimes

and correlation shifts, particularly as edge AI and autonomous driving markets mature, potentially altering the competitive landscape and risk dynamics of these key players.



COMPANY PROFILES AND FINANCIAL PERFORMANCE

NVIDIA: AI Chip Dominance and Blackwell Architecture

NVIDIA has established itself as the central infrastructure provider for the global artificial intelligence (AI) supercycle, transitioning from a graphics processing unit (GPU) specialist to a full-stack computing platform. The company operates under a fabless semiconductor business model, focusing exclusively on chip design, software ecosystem development, and system architecture while outsourcing all physical manufacturing to third-party foundries. This strategic separation of design and fabrication has allowed NVIDIA to concentrate capital resources on research and development, resulting in rapid architectural iteration cycles that have outpaced competitors in the high-performance computing sector. In fiscal year 2025, NVIDIA reported total revenue of $130.5 billion, representing a 114% increase from the previous year, driven primarily by unprecedented demand for accelerated computing infrastructure.[7]

FINANCIAL SCALE

FY2025 Revenue

$130.5 Billion

Up 114% year-over-year, reflecting dominant market share in AI training and inference hardware.

PROFITABILITY

GAAP EPS

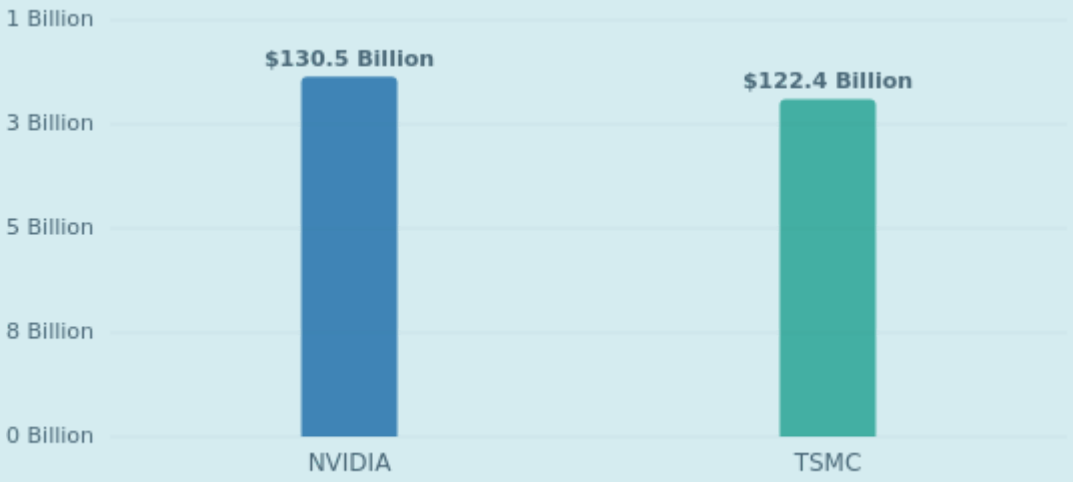
$2.94

Diluted earnings per share increased 147% year-over-year, demonstrating significant operating leverage.

The financial performance of NVIDIA is characterized by extreme concentration in its Data Center segment, which serves as the primary engine for growth. In the third quarter of fiscal 2025, Data Center revenue was $51.2 billion, part of the company's quarterly total of $57.0 billion.[6] This segment includes GPUs designed for large language model (LLM) training, high-performance networking interconnects, and enterprise AI software suites. The dominance of Data Center revenue underscores the structural shift in enterprise IT spending toward generative AI infrastructure, with NVIDIA capturing the majority of value in this transition. The company's ability to scale revenue while maintaining high gross margins reflects its pricing power and the lack of viable alternatives for large-scale AI workloads.

Blackwell Architecture: Strategic Inflection Point

In fiscal year 2025, NVIDIA launched the Blackwell architecture, a comprehensive data center-scale infrastructure platform that integrates GPUs, central processing units (CPUs), data processing units (DPUs), and high-speed interconnects into a unified system.[5] The Blackwell GPU represents a significant technological leap, packing 208 billion transistors and utilizing a custom-built TSMC 4NP manufacturing process.[17] Unlike previous generations, Blackwell employs a dual-die design connected by a 10 terabytes per second (TB/s) chip-to-chip interconnect, effectively functioning as a single unified GPU while overcoming reticle size limitations inherent in monolithic chip designs.



The architectural advancements in Blackwell are specifically optimized for the computational demands of modern generative AI models. The second-generation Transformer Engine incorporates new microscaling formats, including 4-bit floating point (FP4) precision, which doubles the performance and model size capacity compared to previous generations while maintaining high accuracy.[17] Independent microbenchmarking analysis indicates that Blackwell's tensor core enhancements deliver 1.85 times the throughput for ResNet-50 training and 1.55 times the throughput for GPT-1.3B mixed-precision training compared to the preceding H200 generation, alongside a 32% improvement in energy efficiency.[16] These performance metrics are critical for cloud service providers and enterprise customers seeking to reduce the total cost of ownership for large-scale AI deployments.

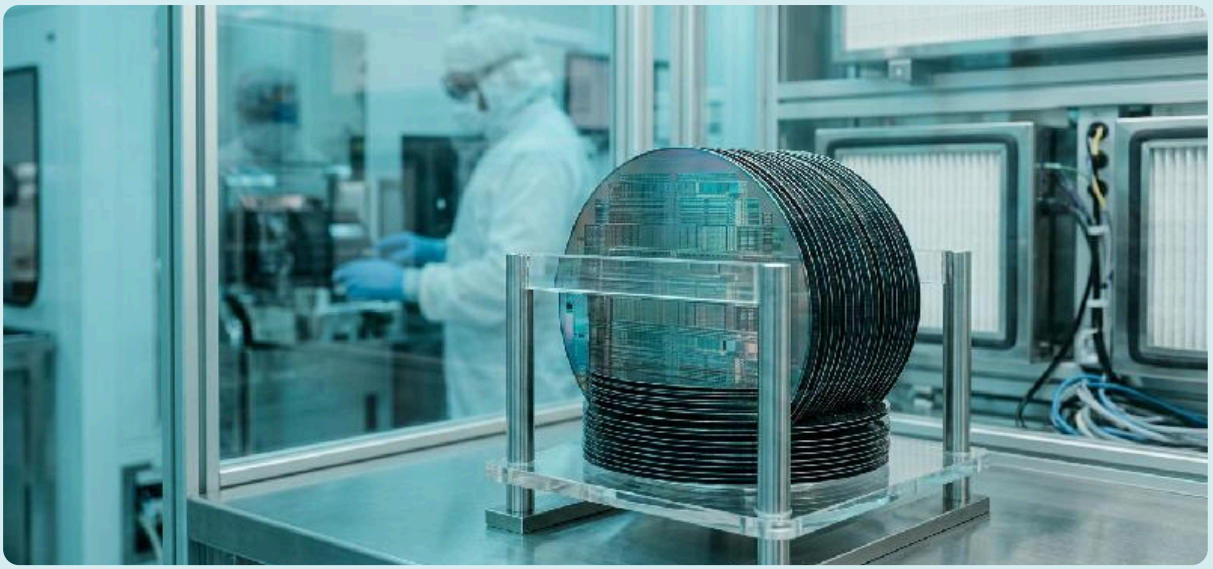
MANUFACTURING DEPENDENCY

As a fabless designer, NVIDIA relies entirely on external foundries for production. The Blackwell architecture is manufactured using TSMC's advanced 4NP process, highlighting the strategic interdependence between NVIDIA's design capabilities and TSMC's manufacturing leadership. This dependency necessitates close collaboration on yield management and capacity allocation, particularly as demand for advanced node wafers exceeds global supply.

Market Position and Competitive Moat

NVIDIA’s competitive advantage extends beyond hardware specifications to include a deeply entrenched software ecosystem. The CUDA parallel computing platform and associated libraries create high switching costs for developers, ensuring that new AI frameworks and models are primarily optimized for NVIDIA hardware. This software-hardware synergy allows NVIDIA to maintain premium pricing despite increasing competition from custom silicon developed by hyperscale cloud providers and alternative GPU manufacturers. The company’s focus on full-stack solutions, including networking via InfiniBand and Ethernet technologies, further consolidates its position as the default infrastructure provider for AI factories.

Looking forward, the deployment of Blackwell-based systems is expected to drive continued revenue growth through fiscal 2026 and beyond. The architecture’s support for larger model parameters and improved inference efficiency addresses the evolving needs of the AI industry, which is shifting from pure training capacity to scalable inference deployment. By integrating confidential computing features and enhanced security protocols directly into the hardware, NVIDIA also addresses growing enterprise concerns regarding data privacy and model protection in multi-tenant cloud environments.[17] This comprehensive approach to infrastructure design ensures that NVIDIA remains at the forefront of the AI supercycle, even as the market matures and diversifies.



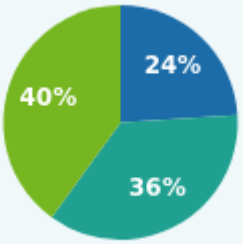
COMPANY PROFILES AND FINANCIAL PERFORMANCE

TSMC: Foundry Leadership and Advanced Node Mix

Taiwan Semiconductor Manufacturing Company (TSMC) operates as the world’s premier pure-play foundry, a business model distinct from Integrated Device Manufacturers (IDMs) such as Samsung Electronics or fabless designers like NVIDIA. By dedicating its entire operational capacity to manufacturing chips designed by external clients, TSMC has established itself as the critical infrastructure provider for the global semiconductor industry. In 2025, the company reported consolidated revenue of $122.4 billion, representing a 35.9% year-over-year increase. This financial performance underscores TSMC’s dominant position in the advanced node market, where it held an estimated 70.2% share of the global pure-play foundry market in the second quarter of 2025.

Revenue Composition and HPC Dominance

The structural shift in semiconductor demand toward artificial intelligence is clearly reflected in TSMC’s revenue mix. High-Performance Computing (HPC), which encompasses AI accelerators, data center processors, and high-speed networking chips, accounted for 58% of total 2025 revenue. This segment has surpassed smartphones as the primary driver of growth, up from 51% in 2024 and 40% in 2022. The concentration of revenue in HPC highlights TSMC’s strategic alignment with the AI supercycle, serving as the exclusive manufacturing partner for leading-edge AI GPUs and custom ASICs. Advanced nodes, defined as manufacturing processes at 7nm and below, made up 74% of wafer revenue in 2025. Within this category, the 3nm and 5nm nodes were the most significant contributors. The 5nm node accounted for 36% of wafer revenue, while the newer 3nm node contributed 24%. The remaining 40% was



generated by other legacy and mature nodes. This distribution demonstrates TSMC’s ability to maintain high utilization rates across multiple generations of technology while transitioning customers to more efficient, smaller geometries.

Technology Leadership: 3nm and 2nm Transition

TSMC’s technological moat is built on its leadership in extreme ultraviolet (EUV) lithography and transistor architecture. The company’s 3nm process, utilizing FinFET technology, has achieved high yields and serves as the foundation for current-generation flagship mobile and computing processors. Looking forward, TSMC commenced volume production of its 2nm (N2) technology in the fourth quarter of 2025[18]. The N2 node represents a fundamental architectural shift, introducing gate-all-around (GAAFET) nanosheet transistors that replace the traditional fin-based structure[19]. This transition allows for superior current control, reduced leakage, and higher transistor density.

By the end of 2026, TSMC projected its total 2nm capacity to reach 100,000 wafers per month, with initial production centered at the Baoshan Fab 20 and Kaohsiung Fab 22 in Taiwan[19]. Early yield data for 2nm logic test chips stabilized between 70% and 80%, a robust figure for a new architecture[19]. The cost of this technological advancement is significant; a single 2nm wafer is estimated to cost approximately $30,000, a 50% premium over 3nm wafers[19]. Despite the higher price point, demand remains strong, with major clients such as Apple securing substantial portions of initial capacity for future iPhone and Mac silicon[19].

STRATEGIC INSIGHT

TSMC’s Foundry Revenue Efficiency Ratio—annual revenue divided by active customer count—rose to $229.3 million in 2025, up from $94.0 million in 2020. This 143.8% increase indicates that TSMC is generating significantly more value per customer, driven by the complexity and volume of AI-related wafer starts rather than an expansion in the total number of clients.

CoWoS Packaging: The Strategic Bottleneck

Beyond front-end wafer fabrication, TSMC’s advanced packaging technology, Chip-on-Wafer-on-Substrate (CoWoS), has become a critical constraint in the AI supply chain. CoWoS enables the integration of high-bandwidth memory (HBM) with logic GPUs in a single package, a requirement for high-performance AI training and inference workloads. As demand for AI accelerators surged, CoWoS capacity emerged as a bottleneck, limiting the total number of GPUs that could be shipped to data centers. TSMC has aggressively expanded its CoWoS capacity, aiming to double it by the end of 2026 to meet the requirements of NVIDIA and other AI chip designers. This vertical integration of packaging services reinforces TSMC’s role as an indispensable partner in the AI ecosystem, providing a level of service that competitors struggle to match.

Capital Expenditure and Geopolitical Diversification

To sustain its technological lead and expand capacity, TSMC committed to substantial capital expenditure. In 2025, the company spent $40.9 billion on capex, with guidance for 2026 ranging between $52 billion and $56 billion. This level of investment represents the largest single-year semiconductor capital commitment in industry history. A significant portion of this spending is directed toward geographic diversification. While the majority of advanced node production remains in Taiwan, TSMC is constructing fabrication facilities in Arizona (United States), Kumamoto (Japan), and Dresden (Germany). These overseas expansions are intended to mitigate geopolitical risks associated with the concentration of semiconductor manufacturing in Taiwan. However, building fabs outside of Taiwan incurs higher construction and operational costs, which may pressure margins in the near term. Despite these challenges, TSMC’s scale, technological superiority, and deep integration with the global AI supply chain position it as the central pillar of the semiconductor industry’s future growth.

Samsung Electronics: Memory Turnaround & HBM Strategy

Samsung Electronics operates under an Integrated Device Manufacturer (IDM) model, distinguishing its structural position from the fabless-foundry dichotomy that defines NVIDIA and TSMC. As an IDM, Samsung controls both the design and manufacturing of its semiconductor products, a vertical integration that historically provided cost advantages but also exposed the company to the full volatility of the memory cycle. In the third quarter of 2026, this cyclical exposure manifested as a historic turnaround, with the company projecting consolidated operating profit of 107.4 trillion won (approximately $80B)[8]. This guidance represents a dramatic recovery from the losses incurred during the 2023–2024 inventory correction and underscores the critical role of High Bandwidth Memory (HBM) in the current AI-driven supercycle.

> **Q3 2026 GUIDANCE HIGHLIGHTS**
>
> The guidance marks a significant increase from previous quarters, driven by the recovery in memory prices and HBM demand. This marks a significant increase from Q2 2026 operating profit of 89.49 trillion won ($66.68B) and Q3 2025 operating profit of 12.17 trillion won ($9.06B)[9].

The HBM Revenue Mix Shift

The primary driver of Samsung’s profitability recovery is the structural shift in its DRAM product mix toward high-value HBM. As AI training workloads demand greater memory bandwidth, hyperscalers have prioritized HBM procurement over conventional DRAM. Samsung has responded by reallocating wafer capacity to meet this demand, projecting that HBM will account for more than one-third of Samsung's DRAM wafer capacity by 2027[20]. This reallocation is not merely a volume adjustment but a fundamental change in revenue composition.

The transition to next-generation HBM standards is accelerating this shift. HBM4, the sixth generation of high-bandwidth memory, is expected to double its share of Samsung’s HBM revenue from 40% in 2026 to 80% in 2027[20]. Concurrently, the share of conventional DRAM in Samsung’s total DRAM capacity is projected to

shrink from 73% in 2025 to 59% in 2027[20]. This compression of legacy product lines reduces exposure to commoditized pricing pressures while capturing the premium margins associated with advanced AI infrastructure.

HBM4 REVENUE SHARE PROJECTION

80%

Expected share of Samsung's HBM revenue in 2027, up from 40% in 2026[20].

CONVENTIONAL DRAM CAPACITY

59%

Projected share of total DRAM capacity in 2027, down from 73% in 2025[20].

Pricing Power and Supply Constraints

The tightness in the HBM supply chain has granted manufacturers significant pricing power. Market analysts estimate that HBM prices will rise more than 100% year-over-year in 2027, reflecting persistent supply shortages and the gradual ramp-up of HBM4 production in the second half of the year[20]. TrendForce, a market tracker, projects an even steeper increase, with blended average selling prices rising by 121%[20]. This pricing environment allows Samsung to narrow the gross margin gap between HBM and conventional DRAM, improving overall profitability despite the higher complexity of advanced packaging.

However, the reallocation of capacity to HBM has created secondary effects in the broader memory market. As HBM consumes an increasing share of wafer output, the supply of conventional DRAM has tightened, leading to potential shortages in standard memory products. Industry sources note that it would be "strange if conventional DRAM was not in short supply" given the rapid expansion of HBM capacity[20]. This dynamic supports higher prices across the entire DRAM spectrum, further bolstering Samsung’s bottom line.

Cyclicality vs. Structural Stability

Despite the record-breaking guidance, Samsung’s financial profile remains inherently cyclical, contrasting sharply with the structural stability observed in TSMC’s foundry business. While TSMC benefits from long-term capacity reservations and multi-year customer agreements, Samsung’s memory division is

subject to sharper fluctuations in spot prices and inventory levels. The company's Mobile eXperience (MX) division, for instance, has reportedly begun cutting smartphone production for Q4 2026 by up to 30% due to rising component costs[9]. This internal friction highlights the trade-offs of the IDM model: while semiconductor profits surge, downstream device margins may compress. Furthermore, Samsung's foundry business continues to trail TSMC in advanced node market share, limiting its ability to fully capitalize on the logic chip boom driven by AI accelerators. Unlike TSMC, which captures value from both NVIDIA's designs and AMD's innovations through its CoWoS packaging technology, Samsung's foundry revenue remains a smaller fraction of its total semiconductor sales. Consequently, Samsung's performance is disproportionately tied to the memory cycle, making it a higher-beta play on AI infrastructure demand compared to its peers.

METRIC	Q3 2025	Q2 2026	Q3 2026 (GUIDANCE)
Consolidated Sales	86.06 trillion won ($64.13B)	171.5 trillion won ($127.78B)	195 trillion won ($145.33B)
Operating Profit	12.17 trillion won ($9.06B)	89.49 trillion won ($66.68B)	107.4 trillion won ($80B)

Table 1: Samsung Electronics Consolidated Financial Guidance. Source: Samsung Electronics[9].

Strategic Implications for Investors

For investors, Samsung Electronics offers a distinct risk-return profile characterized by high cyclicality and significant upside during memory upcycles. The company's P/E ratio of 12.3 reflects this cyclical discount compared to the semiconductor industry average of 23.6. While NVIDIA and TSMC command premium valuations due to their dominant positions in AI logic and advanced packaging, Samsung's valuation is anchored by its legacy memory businesses and lower-margin foundry operations.

Nevertheless, the ongoing AI supercycle provides a tailwind for Samsung's HBM strategy. As the company successfully qualifies its HBM3E and HBM4 products with major hyperscalers, it stands to capture a larger share of the growing AI memory market. The projected doubling of HBM4 revenue share by 2027 suggests

that Samsung is effectively executing its technology roadmap, albeit with a lag behind competitors in some areas. Investors should monitor the company's ability to maintain pricing power in HBM and manage the supply constraints in conventional DRAM, as these factors will determine the sustainability of its recent profit surge.

> **RISK ALERT: INTERNAL DIVISION FRICTION**
>
> Rising semiconductor prices are impacting Samsung's own device divisions. The Mobile eXperience (MX) division has cut smartphone production for Q4 2026 by up to 30%, as rising memory costs erode device margins[9]. This internal cannibalization is a unique risk of the IDM model.

FUTURE HORIZONS: EDGE AI & AUTONOMOUS DRIVING

Beyond the Data Center: The Expansion into Edge Inference and Autonomous Systems

While the semiconductor supercycle of 2023–2026 was defined by centralized data center capital expenditure, the next phase of industry growth is projected to shift toward distributed computing architectures. Edge Artificial Intelligence (Edge AI) represents the deployment of machine learning algorithms on local hardware devices rather than relying exclusively on cloud-based servers. This structural transition addresses critical limitations of cloud-centric models, specifically latency, bandwidth constraints, and data privacy concerns. By processing data at the point of generation, Edge AI enables real-time decision-making capabilities essential for autonomous systems, industrial automation, and consumer electronics. The market trajectory indicates a robust expansion, with the global Edge AI market size projected to grow from $46.96 billion in 2026 to $445.75 billion by 2034, exhibiting a compound annual growth rate (CAGR) of 32.5% during this period.[14]

MARKET PROJECTION

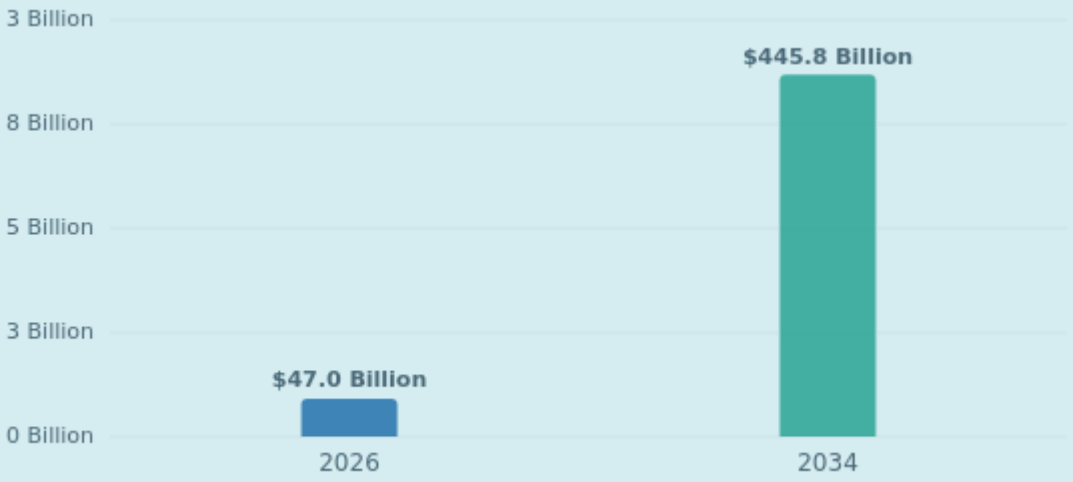
The Edge AI market is forecast to expand from **$46.96 billion (2026)** to **$445.75 billion (2034)**, driven by the integration of AI capabilities into IoT devices and the demand for low-latency processing.[14]

The AI Expansion Pipeline

The evolution of AI infrastructure follows a distinct pipeline from centralized training to decentralized inference. Large language models and complex neural networks are initially trained in data centers using high-performance GPUs, such as those based on NVIDIA's Blackwell architecture. Once trained, these models are optimized and deployed to edge devices for inference—the process of executing the model to generate predictions or decisions. This pipeline extends further into autonomous driving systems, where vehicles require ultra-low latency processing for object detection, path planning, and vehicle control. The interdependence of these stages creates a sustained demand for semiconductors across both the cloud and edge segments.

Market Size Projections: 2026–2034

The financial magnitude of the Edge AI sector underscores its strategic importance to the semiconductor value chain. As IoT devices proliferate and 5G networks enable higher data throughput, the volume of data generated at the edge necessitates local processing capabilities. The projected growth from $46.96 billion in 2026 to $445.75 billion by 2034 reflects not only an increase in device shipments but also a rise in the computational complexity handled by each device.[14] This expansion is supported by advancements in hardware efficiency, including the development of Neural Processing Units (NPUs) and AI accelerator chips that optimize power consumption for battery-operated devices.



Concurrently, the specialized semiconductor segment supporting these applications is experiencing parallel growth. The Edge AI semiconductor market, which includes the physical chips required for edge processing, is projected to grow from USD 29.85 billion in 2026 to USD 107.86 billion by 2034, exhibiting a CAGR of 17.4%.[21] This subset of the broader Edge AI market highlights the direct revenue opportunity for fabless design firms and foundries. Companies such as NVIDIA, Qualcomm, and Intel are increasingly focusing on system-on-chip (SoC) designs that integrate CPU, GPU, and NPU cores to meet the specific power and performance requirements of edge deployments.

Autonomous Driving: The High-Performance Edge

Autonomous driving represents the most computationally intensive application of Edge AI, requiring real-time processing of vast amounts of sensor data from cameras, LiDAR, and radar systems. The global autonomous driving chips market is estimated at USD 24.22 billion in 2024 and is expected to grow from USD 29.73 billion in 2025 to USD 191.07 billion in 2034, at a CAGR of 23%.[22] This growth is driven by the rising adoption of SAE Level 2+ and Level 3 vehicles, as well as increased investments by automotive original equipment manufacturers (OEMs) and Tier-1 suppliers.

The competitive landscape for autonomous driving chips is characterized by a mix of established automotive semiconductor providers and new entrants from the consumer electronics and AI sectors. NXP Semiconductors led the market with over 13% share in 2024, while the top five players—including NXP Semiconductors, Infineon Technologies, Texas Instruments, Renesas Electronics,

and STMicroelectronics—collectively held a 47% market share.[22] However, the increasing complexity of autonomous systems is driving demand for high-performance AI accelerators capable of handling deep learning workloads directly within the vehicle. This shift favors companies with expertise in both high-performance computing and automotive-grade reliability standards, such as ISO 26262 functional safety certification.

LATENCY & PRIVACY

Edge AI Advantages

- Real-time processing within milliseconds
- Reduced dependency on internet connectivity
- Enhanced data privacy via local processing
- Lower bandwidth costs for cloud transmission

AUTONOMOUS CHIPS

Market Leaders (2024)

- NXP Semiconductors (>13% share)
- Infineon Technologies
- Texas Instruments
- Renesas Electronics
- STMicroelectronics

The convergence of Edge AI and autonomous driving technologies creates a synergistic effect on semiconductor demand. As vehicles become more autonomous, they effectively transform into mobile data centers, requiring significant computational power that must be delivered within strict power and thermal constraints. This requirement drives innovation in chiplet architectures and advanced packaging technologies, such as TSMC’s CoWoS, which allow for the integration of heterogeneous compute elements. For investors, the expansion into Edge AI and autonomous driving suggests that the semiconductor supercycle will not end with the saturation of data center build-outs but will instead evolve into a broader, more diversified phase of embedded intelligence across multiple industries.

INVESTMENT RECOMMENDATIONS & CONCLUSION

Investment Recommendations & Valuation

The structural divergence in semiconductor performance between October 2023 and October 2026 mandates a portfolio construction strategy that prioritizes risk-adjusted efficiency over raw cumulative returns. While NVIDIA delivered exceptional absolute gains of +268.4%, its volatility profile and maximum drawdown of -36.88% necessitate careful position sizing. TSMC emerged as the superior vehicle for risk-adjusted capital appreciation, achieving a Sharpe Ratio of 1.5366—the highest among the three major equities analyzed—while maintaining a moderate correlation (0.6510) with NVIDIA. Samsung Electronics, despite its pivotal role in the High Bandwidth Memory (HBM) supply chain, exhibited the lowest risk-adjusted performance with a Sharpe Ratio of 0.9854 and the deepest maximum drawdown of -42.85%. Consequently, investment recommendations are stratified by investor risk tolerance, leveraging the distinct financial architectures of the Fabless, Foundry, and Integrated Device Manufacturer (IDM) models.

AGGRESSIVE GROWTH

NVIDIA Core

Target investors seeking maximum exposure to AI infrastructure scaling. Capitalize on the Blackwell architecture cycle and FY2025 revenue surge to $130.5 billion[7]. Accept higher volatility for superior absolute return potential.

BALANCED ALLOCATION

TSMC Anchor

Optimal for investors prioritizing capital preservation alongside growth. Leverage TSMC’s 1.5366 Sharpe Ratio and pure-play foundry dominance in 3nm/2nm nodes to mitigate single-company execution risk.

DIVERSIFIED INCOME

Samsung Satellite

Suitable for value-oriented portfolios targeting memory cycle turnarounds. Utilize Samsung’s low correlation (0.1343 with NVDA) to hedge against logic-chip concentration risks while capturing HBM demand recovery.

Strategic Portfolio Allocations

Portfolio construction must account for the asymmetric correlations observed in daily returns. The moderate positive correlation between NVIDIA and TSMC (0.6510) reflects their symbiotic relationship in the AI GPU supply chain, where TSMC's CoWoS packaging capacity directly constrains NVIDIA's shipment volumes. Conversely, Samsung's near-zero correlation with NVIDIA (0.1343) provides significant diversification benefits, insulating the portfolio from logic-semiconductor-specific shocks. The following table outlines recommended weightings based on the risk-return profiles established in prior analysis.

PORTFOLIO TYPE	PRIMARY HOLDING	SECONDARY HOLDING	DIVERSIFIER
Aggressive Growth	NVDA (60%)	TSM (30%)	Cash / Short-term Treasuries (10%)
Balanced Allocation	TSM (50%)	NVDA (30%)	005930.KS (20%)
Diversified Income	TSM (30%)		NVDA (20%)

For the Aggressive Growth portfolio, a 60% allocation to NVIDIA captures the upside of the AI supercycle, while a 30% TSMC position hedges against potential fabless design bottlenecks. The Balanced Allocation strategy inverts this ratio, placing TSMC at the core due to its superior Sharpe Ratio of 1.5366, thereby reducing overall portfolio volatility without sacrificing exposure to advanced node demand. The Diversified Income profile utilizes TSMC as a stable yield and growth engine, complemented by smaller positions in NVIDIA and Samsung to capture cyclical upswings in both logic and memory sectors. Notably, Samsung's inclusion in diversified portfolios is justified not by its standalone risk-adjusted returns, but by its unique correlation structure which dampens overall portfolio variance during logic-chip downturns.

Valuation Context & Market Multiples

Valuation metrics as of October 2026 indicate divergent market expectations for each entity. Samsung Electronics trades at a relatively conservative 1.9x revenue multiple, reflecting its conglomerate structure and the cyclical nature of its

memory business. With last twelve months (LTM) revenue reaching $487 billion and EBITDA of $275 billion, Samsung presents a value-oriented entry point for investors willing to endure the inherent volatility of the IDM model. Its Price-to-Earnings (P/E) ratio of 12.3 places it fifth among top semiconductor companies by P/E ratio, below the average P/E ratio of 23.6 for the top 169 semiconductor companies.

In contrast, NVIDIA and TSMC command premium valuations commensurate with their dominant positions in the AI infrastructure stack. NVIDIA's FY2025 revenue growth of 114% year-over-year to $130.5 billion justifies its elevated multiples, as the market prices in sustained demand for Blackwell-based systems[7]. TSMC's revenue growth underscores its pricing power in advanced nodes, where it faces negligible competition in 3nm and 2nm manufacturing. Investors should monitor the convergence of these multiples; a compression in NVIDIA's valuation relative to TSMC would signal a shift in market preference from design innovation to manufacturing scarcity.

ANALYST CONSENSUS

Goldman Sachs maintains buy ratings on key technology enablers, citing Microsoft's AI platform breadth and unit economics in Azure as primary drivers for continued infrastructure spend[29]. This institutional confidence supports the thesis that semiconductor demand will remain robust through 2027, favoring companies with direct exposure to cloud AI capex cycles such as NVIDIA and TSMC.

Ultimately, the choice between these equities depends on the investor's view of the AI lifecycle stage. Early-stage investors may prefer NVIDIA's explosive growth trajectory, while late-cycle allocators should favor TSMC's entrenched manufacturing moat and superior risk-adjusted returns. Samsung serves as a tactical hedge, offering exposure to the memory intensity of AI workloads without the direct execution risk of GPU architecture transitions. A disciplined rebalancing strategy, guided by the Sharpe Ratio hierarchy and correlation dynamics, remains the most effective method for navigating the semiconductor supercycle.

INVESTMENT RECOMMENDATIONS & CONCLUSION

Conclusion & Key Risks

The semiconductor industry’s transition into an AI-driven supercycle has fundamentally altered the risk-return profiles of its leading participants. While the structural demand for compute remains robust, the divergence in financial performance and operational resilience among NVIDIA, TSMC, and Samsung Electronics necessitates a differentiated investment approach. The analysis confirms that NVIDIA and TSMC represent core holdings for exposure to the AI infrastructure buildout, whereas Samsung Electronics functions primarily as a cyclical diversifier with limited direct correlation to the GPU-led growth narrative. However, this bullish thesis is contingent upon the industry’s ability to navigate a complex matrix of geopolitical, physical, and regulatory constraints that have intensified in 2026.

> **STRATEGIC IMPERATIVE**
> Portfolio construction must prioritize risk-adjusted efficiency over raw cumulative returns. TSMC’s superior risk-adjusted returns and moderate correlation with NVIDIA provide optimal diversification benefits, while Samsung’s low correlation with NVIDIA offers a hedge against concentrated AI hardware volatility.

Geopolitical Fragility and Supply Chain Resilience

The most immediate threat to the semiconductor supply chain is no longer purely technological but geopolitical. The concentration of advanced logic manufacturing in Taiwan and memory production in South Korea exposes the global AI infrastructure to regional instability. Recent disruptions in the energy sector, such as the March 2026 missile and drone strikes on Qatar’s Ras Laffan hub, which removed 20% of global LNG supply, have demonstrated how quickly physical constraints can throttle chip production[24]. For TSMC and Samsung, energy security is now a critical component of geopolitical resilience. The integration of ESG principles, particularly the transition to renewable energy sources, is no longer merely a compliance issue but a strategic imperative to mitigate exposure to volatile fossil fuel markets and ensure uninterrupted operations[3]. Furthermore, the United States has increasingly leveraged its control over high-bandwidth memory and advanced chip exports as a tool of foreign policy. As chip shortages persist into 2026, every unit exported to competing nations represents a potential shortfall for domestic AI development[25]. This dynamic creates a

precarious environment for companies like Samsung and TSMC, which must balance their commercial interests in China with strict export controls imposed by Washington. The risk of secondary sanctions or forced divestment remains a latent but significant drag on valuation multiples for firms with substantial exposure to the Chinese market.

The Physical Bottleneck: Energy and Infrastructure

Beyond geopolitics, the AI supercycle faces a hard physical limit: electricity. The five largest hyperscalers committed over $660 billion in capital expenditures in 2026, yet they are encountering a grid capacity wall[24]. AI-optimized data centers now require 100-500 megawatts each, equivalent to powering entire cities, and existing infrastructure cannot scale at the required pace. This "energy chokehold" threatens to delay the deployment of next-generation GPUs, such as NVIDIA’s Blackwell architecture, regardless of manufacturing yield. If power availability does not keep pace with chip production, the industry risks a scenario where supply exceeds deployable capacity, leading to inventory backlogs and downward pressure on pricing power.

RISK ALERT: CAPEX SUSTAINABILITY

- Hyperscaler spending growth may decelerate if ROI on AI models fails to materialize.
- Physical infrastructure delays (power, cooling) could push out revenue recognition for chipmakers.
- Regulatory scrutiny on energy consumption may impose new operational costs.

Regulatory and Antitrust Headwinds

As NVIDIA’s market capitalization approaches the $4 trillion threshold, it has become a primary target for global antitrust regulators. Concerns about the sustainability of its stock rally have grown amid investigations into its dominant position in the AI chip market[13]. Regulatory actions could force changes in business practices, such as unbundling software from hardware or limiting exclusive partnerships with cloud providers. Similarly, Samsung Electronics has faced intense scrutiny from the Korea Fair Trade Commission (KFTC) regarding

abuse of superior bargaining positions, as seen in the Broadcom consent decree case[27][28]. These regulatory pressures introduce legal uncertainty and potential fines that could erode margins and distract management from core innovation efforts.

Final Verdict: A Tiered Allocation Strategy

In conclusion, the semiconductor landscape of 2026 demands a nuanced portfolio strategy. NVIDIA remains the undisputed leader in AI acceleration, but its valuation reflects near-perfect execution expectations. TSMC offers the most compelling risk-adjusted profile, backed by its irreplaceable role in the global supply chain and superior Sharpe Ratio. Samsung Electronics, while benefiting from the HBM memory boom, remains a cyclical play with lower correlation to the core AI narrative, making it suitable only as a tactical diversifier. Investors must remain vigilant against the converging risks of geopolitical fragmentation, energy scarcity, and regulatory intervention, which collectively define the upper bound of the current AI supercycle.

STRATEGIC ACTIONS

- Maintain core positions in NVDA and TSM for long-term AI exposure.
- Utilize Samsung (005930.KS) as a low-correlation hedge against memory cycles.
- Monitor hyperscaler capex guidance for early signs of demand saturation.
- Diversify geographic exposure to mitigate regional geopolitical shocks.

RISK AVOIDANCE

- Avoid overconcentration in single-supplier dependencies without hedging.
- Do not ignore energy infrastructure constraints in deployment forecasts.
- Refrain from speculative trading based on short-term geopolitical headlines.
- Do not assume regulatory environments will remain static for dominant players.

Data Basis and Limitations

The analytical framework, financial metrics, and strategic conclusions presented in this report are grounded in publicly available data verified as of October 28, 2026. This section delineates the methodological boundaries, temporal constraints, and structural limitations inherent in the underlying datasets. Investors and stakeholders must interpret the preceding performance analyses—specifically regarding NVIDIA, TSMC, and Samsung Electronics—within these defined parameters to avoid misapplication of historical trends to future projections.

TEMPORAL BOUNDARY

All figures, financial statements, and market capitalizations reflect the state of public information available up to October 28, 2026. Subsequent corporate earnings releases, macroeconomic policy shifts, or geopolitical developments occurring after this date are not incorporated into the quantitative models or risk assessments presented herein.

1. Temporal Scope and Data Latency

The three-year cumulative return analysis, spanning from October 2023 to October 2026, captures a period of extreme volatility and structural transition in the semiconductor industry. The maximum drawdown periods identified for each entity are specific to the observation window. These historical stress tests do not guarantee similar behavior in future market corrections, particularly as the industry matures from an AI-infrastructure build-out phase to an application-deployment phase.

Furthermore, financial data derived from fiscal year reports may exhibit timing mismatches. For instance, NVIDIA’s reported revenue reflects a fiscal calendar that differs from the standard calendar year used by some competitors. While adjustments have been made to align comparative risk metrics, residual latency in the reporting of supply chain constraints or order backlogs may introduce a lag between real-time operational reality and reported financial outcomes.

2. Methodological Derivation of Risk Metrics

The risk-adjusted performance indicators, including the Sharpe Ratios and correlation matrices, are computed directly from daily price returns over the specified three-year horizon. The calculated Sharpe Ratios represent a historical average of excess returns per unit of volatility. These values are sensitive to the chosen risk-free rate benchmark and the specific start and end dates of the observation period. A shift in the observation window by even a few weeks can materially alter these ratios, particularly for assets with high short-term variance such as NVIDIA.

The calculated correlation coefficients are statistical measures of co-movement, not causal relationships. They indicate how prices have moved together in the past but do not predict future dependency structures. In times of systemic market stress, correlations across asset classes often converge toward unity, potentially diminishing the diversification benefits suggested by the historical heatmap. Therefore, the portfolio construction recommendations based on these historical correlations should be viewed as a baseline rather than a static hedge strategy.

3. Heterogeneity in Source Definitions

Market size projections and industry growth rates cited from third-party research firms—including Gartner, IDC, Deloitte, and McKinsey—utilize varying methodological scopes. Discrepancies exist in geographic coverage (e.g., global versus ex-China), product categorization (e.g., standalone logic chips versus integrated modules), and temporal definitions (e.g., calendar year versus fiscal year). For example, projected market sizes for Edge AI semiconductors and autonomous driving chips are based on specific definitions of "edge" and "autonomous" that may exclude certain embedded systems or hybrid cloud-edge architectures included by other analysts.

Consequently, figures from different sources are not directly interchangeable. Reported compound annual growth rates for specific sub-segments cannot be linearly extrapolated to other areas without adjusting for the specific base year and scope definition used by the originating firm. Users of this report should treat these projections as directional indicators of market momentum rather than precise deterministic forecasts.

4. Structural Limitations of Forward-Looking Statements

The valuation multiples and revenue forecasts for future years rely on assumptions of continued technological scaling and demand elasticity. These models assume no catastrophic disruptions in the global supply chain, such as prolonged geopolitical conflicts affecting advanced node manufacturing in Taiwan or South Korea. Additionally, the regulatory environment surrounding artificial intelligence and semiconductor exports is dynamic. Changes in export controls, antitrust enforcement, or subsidy regimes could fundamentally alter the competitive landscape for fabless designers like NVIDIA and foundries like TSMC, rendering current valuation multiples obsolete.

Finally, the financial health of Samsung Electronics, particularly its memory division turnaround and HBM strategy, is assessed based on recent guidance and operating profit trends. The cyclical nature of the memory market means that current profitability improvements may be subject to rapid reversal if inventory levels rise or if demand for consumer electronics stagnates. The report’s positive outlook on Samsung’s recovery is contingent on the successful execution of its advanced packaging roadmap and the sustained pricing power of HBM products, both of which face intense competition from emerging rivals.

ANALYTICAL DISCLAIMER

Interpretation Guidance

- Historical performance metrics (Sharpe Ratio, MDD) are backward-looking and do not guarantee future results.
- Correlation matrices describe past co-movement and may break down during systemic crises.
- Market size projections vary by source definition; compare only within consistent methodological frameworks.
- All data is current as of October 28, 2026; subsequent events are not reflected in this analysis.

References

1 Nvidia

Wikipedia · Accessed: 2026-10-28

https://en.wikipedia.org/wiki/Nvidia

2 TSMC

Wikipedia · Accessed: 2026-10-28

https://en.wikipedia.org/wiki/TSMC

3 Global Value Chain Resilience: Can ESG Practices Mitigate Geopolitical Risk? The Perspective from the Asian Semiconductor Industry

Asuming, Daniel · 2024 · *OPUS HWR Berlin* · DOI: 10.4393/opushwr-4336 · Accessed: 2026-10-28

https://doi.org/10.4393/opushwr-4336

4 Inverse DEA for Portfolio Volatility Targeting: Industry Evidence from Taiwan Stock Exchange

Temitope Olubanjo Kehinde, Sai-Ho Chung, Oludolapo Akanni Olanrewaju · 2025 · *International Journal of Financial Studies* · DOI: 10.3390/ijfs13040192 · Citations: 7 · Accessed: 2026-10-28

https://doi.org/10.3390/ijfs13040192

5 nvda-20260125

sec.gov · Accessed: 2026-10-28

https://www.sec.gov/Archives/edgar/data/1045810/000104581026000021/nvda-20260125

6 nvda-20251026

sec.gov · Accessed: 2026-10-28

https://www.sec.gov/Archives/edgar/data/1045810/000104581025000230/nvda-20251026

7 NVIDIA Announces Financial Results for Fourth Quarter ...

sec.gov · Accessed: 2026-10-28

https://www.sec.gov/Archives/edgar/data/1045810/000104581025000021/q4fy25pr.htr

8 Samsung flags $80 billion profit on AI boom, highest ...

reuters.com · Accessed: 2026-10-28

https://www.reuters.com/business/samsung-q3-profit-jumps-783-ai-memory-boom-lift

9 Samsung Electronics announces its earnings guidance for ...

gsmarena.com · Accessed: 2026-10-28

https://www.gsmarena.com/samsung_q3_2026_earnings_guidance-news-74956.php

10 Gartner Forecasts Worldwide Semiconductor Revenue to ...

gartner.com · Accessed: 2026-10-28
https://www.gartner.com/en/newsroom/press-releases/2026-04-08-gartner-forecasts-

11 Gartner Forecasts Worldwide Semiconductor Revenue to ...
gartner.com · Accessed: 2026-10-28
https://www.gartner.com/en/newsroom/press-releases/2026-08-24-gartner-forecasts-

12 IDC Semiconductor Market Outlook
semi.org · Accessed: 2026-10-28
https://www.semi.org/sites/semi.org/files/2025-09/1%20%EA%B9%80%EC%88%98%EA%B2%B

13 Nvidia's $4 Trillion Stock Rally Faces More Threats Than ...
bloomberg.com · Accessed: 2026-10-28
https://www.bloomberg.com/news/articles/2026-01-06/nvidia-s-4-trillion-stock-ral

14 Edge AI Market Size, Share, Trends | Growth Forecast [2034
fortunebusinessinsights.com · Accessed: 2026-10-28
https://www.fortunebusinessinsights.com/edge-ai-market-107023

15 Semiconductor market forecast 2030: AI-driven growth
mckinsey.com · Accessed: 2026-10-28
https://www.mckinsey.com/industries/semiconductors/our-insights/the-2-point-3-tr

16 Microbenchmarking NVIDIA's Blackwell Architecture
arxiv.org · Accessed: 2026-10-28
https://arxiv.org/html/2512.02189v3

17 NVIDIA Blackwell Architecture
nvidia.com · Accessed: 2026-10-28
https://www.nvidia.com/en-us/data-center/technologies/blackwell-architecture/

18 2nm Technology - Taiwan Semiconductor Manufacturing ...
tsmc.com · Accessed: 2026-10-28
https://www.tsmc.com/english/dedicatedFoundry/technology/logic/l_2nm

19 TSMC Commences Mass Production of 2nm Chips ...
investor.wedbush.com · Accessed: 2026-10-28
https://investor.wedbush.com/wedbush/article/marketminute-2026-1-22-the-silicon-

20 Samsung heads for 100 trillion won profit milestone in Q3, ...
koreajoongangdaily.com · Accessed: 2026-10-28
https://www.koreajoongangdaily.com/business/samsung-heads-for-100-trillion-won-p

21 Edge AI Semiconductor Market Size, Share | Forecast [2034
fortunebusinessinsights.com · Accessed: 2026-10-28
https://www.fortunebusinessinsights.com/edge-ai-semiconductor-market-117383

22 Autonomous Driving Chips Market Size & Share 2025 - 2034
gminsights.com · Accessed: 2026-10-28
https://www.gminsights.com/industry-analysis/autonomous-driving-chips-market

23 Global Semiconductor Industry Outlook - GSA

gsaglobal.org · Accessed: 2026-10-28

https://www.gsaglobal.org/global-semiconductor-industry-outlook/

24 The great data center delay: Why your AI chips are stuck in ...

manufacturingdive.com · Accessed: 2026-10-28

https://www.manufacturingdive.com/news/opinion-omdia-ai-semiconductor-chip-scarc

25 American AI Companies Can't Get Enough Chips

cnas.org · Accessed: 2026-10-28

https://www.cnas.org/publications/reports/american-ai-companies-cant-get-enough-

26 Samsung Electronics Seen Deeply Undervalued, Target Kept ...

en.sedaily.com · Accessed: 2026-10-28

https://en.sedaily.com/finance/2026/10/03/samsung-electronics-deeply-undervalued

27 Implications of the Recent KFTC Decision to Accept ...

kimchang.com · Accessed: 2026-10-28

https://www.kimchang.com/en/insights/detail.kc?sch_section=4&idx=25758

28 Korea fines Broadcom €13.4 million over Samsung abuse

globalcompetitionreview.com · Accessed: 2026-10-28

https://globalcompetitionreview.com/article/korea-fines-broadcom-eu134-million-o

29 Top tech stocks to buy, according to Goldman Sachs

cnbc.com · Accessed: 2026-10-28

https://www.cnbc.com/2026/10/24/top-tech-stocks-to-buy-according-to-goldman-sach