INSIGHTS:

Great technology, poor returns

The pattern investors keep repeating. The technology can win. The capital can still lose.

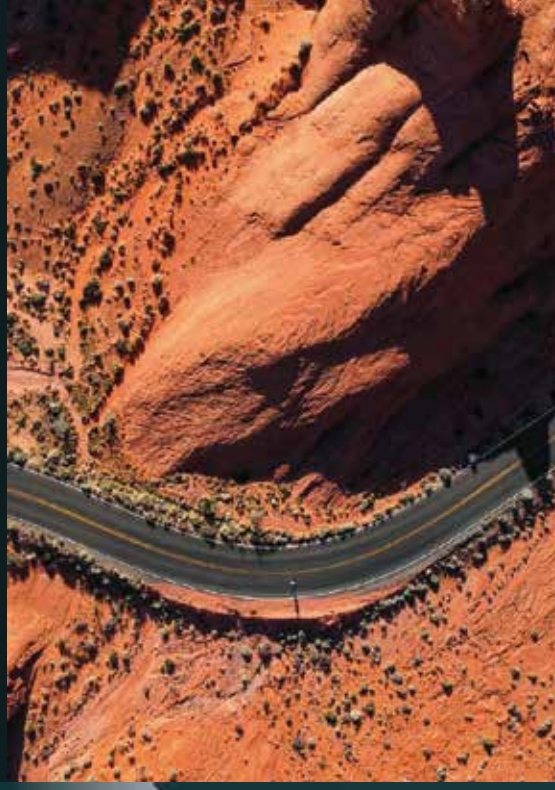

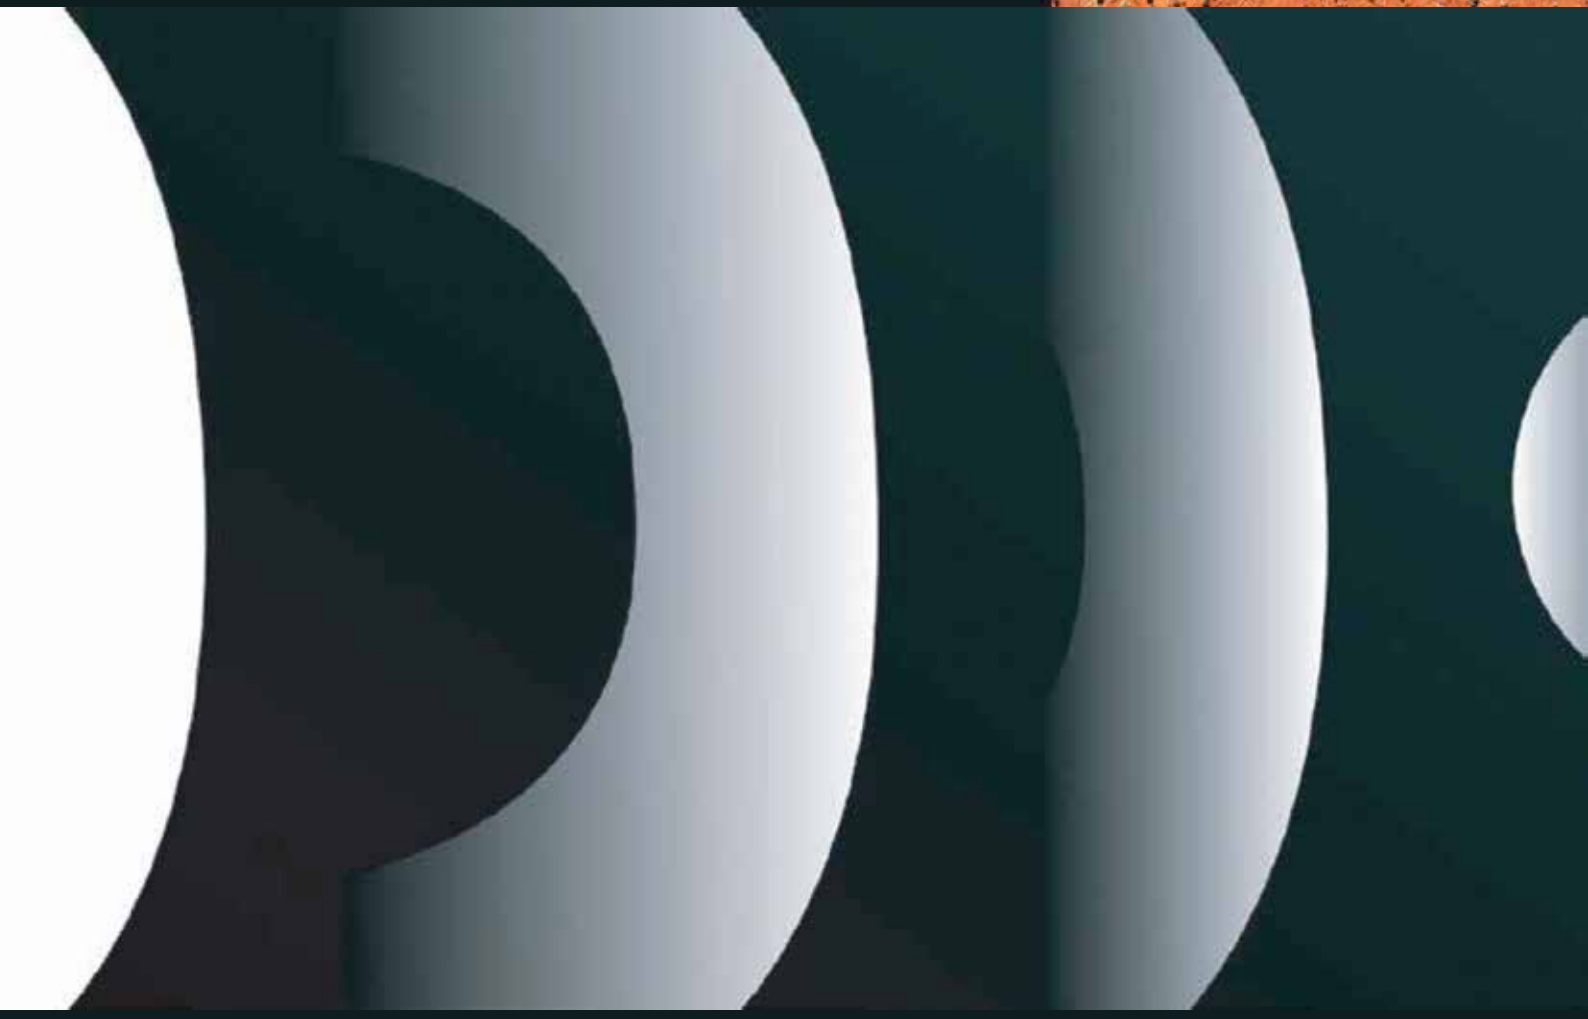

Why the AI build-out may reward its adopters more than its builders - and why most investors are already exposed to the same capital cycle.

INSYNC

THE ARGUMENT IN ONE MINUTE_

AI's importance is settled. Its investment economics are not.

Artificial intelligence is real, useful and likely to be profoundly consequential. The unresolved question is whether the enormous amount of capital being deployed to build the infrastructure will earn an adequate return — and how much of the average portfolio already depends on that answer.

US $725 bn

Approximate 2026 capital expenditure planned by Amazon, Alphabet, Meta and Microsoft, based on announced guidance and calendar-year estimates.

US $650 bn

Estimated annual AI revenue required to earn a 10% return on the build-out, according to J.P. Morgan.

US $2tn

Annual AI revenue Bain estimates may be required by 2030, with an estimated US$800bn shortfall.

~31%

Approximate combined weight of TSMC, Samsung Electronics and SK Hynix in the MSCI Emerging Markets Index at 30 June 2026.

OUR CENTRAL CONCLUSIONS_

First, demand is not the same as return. AI usage can grow rapidlywhile the capital deployed to serve it earns less than investors expect.

Second, the businesses funding the build-out are becoming more capital-intensive. That can dilute return on invested capital even if revenue remains strong.

Third, the suppliers collecting the spending are still exposed to semiconductor cyclicality. Record margins and apparent shortages should not automatically be capitalised as permanent.

Fourth, many portfolios are less diversified than they appear. US, global, emerging-market, quality, growth and thematic allocations increasingly overlap around the same AI capital cycle.

What This Paper Is — And Is Not_

This is not a claim that AI is a bubble, a forecast that the spending cycle will end imminently, or an argument against owning technology companies. It is a capital-allocation argument: investors should distinguish the social and economic importance of a technology from the return earned by the capital that builds it.

THE RECURRING PATTERN_

The technology works. The capital does not always follow.

There is a version of this story that gets told roughly once a generation, and it always opens the same way. A genuinely transformational technology arrives. Demand is real. The companies building it are among the best in the world. Capital rushes in to construct the infrastructure because nobody wants to arrive late.

Then comes the part that tends to be forgotten: the technology works. The capital does not.

British railways in the 1840s. American railroads in the 1880s. Fibre-optic networks around 2000. US shale through the 2010s. Each changed the economy. Each also consumed enormous amounts of capital, and in each case many investors who assumed that "mattering" and "earning" were the same thing were disappointed.

The internet did change everything. It simply created more durable value for many of the companies that used the fibre than for the companies that laid it. US shale made America the world's largest oil producer while the industry, in aggregate, consumed more cash than it generated for much of the following decade.

We think artificial intelligence has reached the point where investors must make the same distinction. This is not a view about whether AI is real or useful. It plainly is both. It is a view about the price being paid to build it, the durability of the earnings being created along the supply chain and the return that the deployed capital is likely to earn.

THE SCALE IS NOT IN DISPUTE_

Amazon, Alphabet, Meta and Microsoft plan to spend around US$725 billion between them on capital expenditure in 2026, based on announced guidance and calendar-year estimates. That compares with roughly US$410 billion in 2025 and US$226 billion in 2024. The number is moving rather than fixed: guidance has generally been revised upwards as the year has progressed. [1]

The broader build-out is larger again. Goldman Sachs' baseline model estimated approximately US$765 billion of annual AI capital expenditure in 2026 and around US$7.6 trillion cumulatively between 2026 and 2031 once chips, data centres, power, cooling and supporting infrastructure are included. [2]

One feature genuinely is different this time, and not necessarily in a reassuring direction. Railroads and fibre networks were often financed by equity and speculative debt issued by companies that could fail. The current build-out began with some of the strongest balance sheets in corporate history. Yet even these companies are increasingly supplementing internal cash flow with bonds, leases, project finance and private capital.

THE ECONOMIC HURDLE_

The revenue gap is the whole argument

AI revenue is growing exceptionally quickly. The issue is not whether monetisation exists; it is whether monetisation can grow fast enough, at adequate prices and utilisation, to earn an attractive return on the capital already committed.

Anthropic's annualised revenue run rate rose from approximately US$9 billion at the end of 2025 to around US$47 billion by May 2026. OpenAI was reported to be running near US$25 billion. Between them, two businesses that were barely commercial a few years ago are producing revenue at a scale few companies ever reach. [3]

That is genuine progress and should not be dismissed. But the comparison with the build-out remains demanding. Together, the two companies' annualised revenue represents only about one-tenth of one year's capital expenditure by the four largest hyperscalers. Attribution outside the model providers is also difficult: part of cloud growth reflects AI, part reflects the migration of existing workloads, and part may represent revenue shifted from other products rather than entirely new economic value.

J.P. Morgan estimates that the AI build-out may need to generate approximately US$650 billion of annual revenue to earn a 10% return on the capital deployed. Bain estimates that AI companies could require around US$2 trillion of annual revenue by 2030 to fund the computing capacity implied by projected demand, while still facing an estimated shortfall of roughly US$800 billion. [4] [5]

> **The gap may be narrowing as a ratio while still widening in dollars. Dollars are what must be financed.**

A second pressure works against the revenue hurdle: the unit price of AI processing continues to fall. One analysis of 2.4 billion enterprise API calls found that the effective blended cost per million tokens fell 67% between the first quarter of 2025 and the first quarter of 2026, from US$18.40 to US$6.07. Falling prices stimulate usage, but they also require substantially more volume to support the same revenue pool. [6]

This creates a difficult combination. The industry must grow into an enormous revenue requirement while the price of the output is deflating and the hardware used to produce it is improving rapidly. Cheaper inference may ultimately accelerate adoption through a Jevons-style effect, but higher usage does not automatically guarantee that the owners of the infrastructure capture attractive economics.

THE MECHANISM_

Why this becomes a return-on-capital problem

For a quality-focused investor, this is the mechanism that matters most.

The companies at the centre of the build-out have historically been among the finest businesses in the world: asset-light, high-margin and capable of producing exceptional returns on invested capital because they needed comparatively little capital to grow. Software scaled without factories.

AI infrastructure is different. It is physical, power-hungry, capital-intensive and depreciating. Data centres require chips, land, buildings, electricity, cooling, networking and redundancy. Accounting lives for servers and network equipment often extend over five or six years, while economic obsolescence can arrive much earlier as newer processors deliver sharply better performance per dollar.

THE ARITHMETIC_

A company earning 30% on its existing capital that begins deploying very large sums at 5% or 10% does not remain a 30% return business. **It becomes a blended one.** That dilution can occur without a single quarter of disappointing revenue.

This is why we treat rising capital-expenditure guidance as information about future returns, not simply as evidence of future demand.

The key variables are utilisation, pricing, useful life, maintenance capital and the extent to which revenue is genuinely incremental rather than cannibalised from existing services.

The funding structure is changing accordingly. Alphabet, Amazon, Meta, Microsoft and Oracle had issued approximately US$236 billion of debt through May 2026, more than four times the amount issued over the same period a year earlier. Private credit, leases, project finance and supplier-backed structures are also playing a larger role, making the ultimate distribution of risk less transparent than headline company balance sheets imply. [7]

WHY IT IS DIFFICULT TO STOP_

If the prospective returns are questionable, why does somebody not simply spend less? Because the build-out has become load-bearing at three levels at once.

For the economy. Computing infrastructure is now an important contributor to US business investment. A material slowdown would therefore be a macroeconomic event, not merely an earnings disappointment.

For the companies. No participant can comfortably stop unilaterally. Each is spending partly from conviction and partly from the fear of lacking capacity if demand does materialise. That can be rational for each company and poor for the industry as a whole.

For the financing. Once projects are supported by bonds, long-dated leases, purchase commitments and private capital, they acquire contractual momentum. Committed capital tends to be spent even when the return calculation becomes less attractive.

The market has begun to react to this tension. Alphabet raised its 2026 capital-expenditure outlook to US$195–205 billion in July. Its shares fell despite Google Cloud revenue growth of more than 80%, illustrating that investors are beginning to distinguish strong demand from the price paid to serve it. [8]

The practical conclusion is that the cycle may run longer and become larger than a pure return-on-capital analysis would justify. That is what over-investment cycles do. It is also why we are not attempting to call the precise top.

THE SUPPLIER SIDE_

The most cyclical industry in the world, priced as though it is not

There is another under-appreciated element: who is collecting the money today.

A substantial portion of hyperscaler capital expenditure becomes current-period revenue for semiconductor and adjacent infrastructure suppliers. What a hyperscaler records as a long-lived asset, a chipmaker, networking company or equipment supplier records as a sale. The same transaction is viewed from opposite ends of the supply chain.

That matters because semiconductors have been among the most reliably cyclical industries in modern markets. The product changes; the mechanism does not.

Demand accelerates and lead times lengthen. Customers order ahead of need, often from multiple suppliers, to secure allocation. Reported orders begin to exceed genuine consumption. Capacity is then committed on multi-year lead times into what appears to be a permanent shortage.

The turn does not require demand to collapse. End demand merely has to flatten, utilisation expectations have to disappoint or customers have to discover that they ordered too much. Inventory then unwinds through the supply chain.

Revenue can fall much faster than end consumption, and because fabrication and manufacturing costs are largely fixed, margins often fall faster than revenue.

The recent history remains instructive. In the 2023 downturn, global semiconductor sales fell 9.4%, while memory revenue declined by roughly 31%. Sector drawdowns of 60% to 70% followed both the dot-com collapse and the global financial crisis, while less severe corrections in 2018 and 2020 still exceeded 30%. [9]

LATE-CYCLE SIGNALS ARE VISIBLE_

Margins are at or near records in several parts of the memory supply chain. Customers are signing multi-year agreements to secure high-bandwidth memory. Suppliers describe themselves as unable to meet a large proportion of potential demand. At the same time, capacity commitments are expanding dramatically.

South Korea announced an investment programme involving Samsung Electronics, SK Group and suppliers that was valued at more than US$500 billion at prevailing exchange rates. Micron subsequently lifted its planned US manufacturing and technology investment to more than US$250 billion through 2035. These are long-duration programmes rather than one-year spending forecasts, but they demonstrate how quickly scarcity can attract future supply. [10]

The industry may remain undersupplied for longer than sceptics expect. But the investment consequence is straightforward: the stronger the price signal and the longer it persists, the more capacity ultimately arrives. The best cure for a shortage is often the shortage itself.

THE VALUATION TRAP_

Peak-cycle earnings produce the cheapest multiple of the cycle

> **A cyclical business at the top of its cycle always looks inexpensive — precisely when the risk can be greatest.**

This is an old valuation trap. Record earnings produce a low price-to-earnings multiple, so the shares appear to offer a margin of safety at exactly the point where current earnings may provide the least useful guide to normal profitability.

Memory manufacturers have historically achieved very strong gross margins at the top of a cycle and materially lower margins at the bottom. During the 2018–19 downturn, one leading producer saw its gross margin fall from close to 60% to below 30%, revenue decline by around 30% and its share price more than halve while the company was still reporting near-peak profits.

The discipline required is to value the businesses on normalised, through-cycle earnings rather than on the current year. A low multiple on a record margin is not automatically value. It may simply be the market's traditional way of discounting the next downcycle.

This assumption is even more demanding when the customer base is financing purchases with more debt and when the return earned by that customer base remains unresolved. A supplier can report excellent current revenue even while the economics of its customers are weakening.

THE COUNTER-CASE DESERVES TO BE TAKEN SERIOUSLY_

There is a strong argument on the other side. AI demand is different from a personal-computer replacement cycle. Training requirements continue to rise, inference is broadening into consumer and enterprise workflows, and agentic systems may consume substantially more computing capacity than conventional chat applications.

Long-term supply agreements could also reduce the traditional inventory cycle by improving visibility and shifting some risk towards customers. Custom silicon, higher utilisation and software optimisation may improve the economics of the installed base. Falling token prices could expand the addressable market so dramatically that aggregate revenue continues to rise despite unit-price deflation.

The cycle may therefore extend considerably. We are not relying on an imminent collapse. The investment question is whether current valuations already assume that exceptional demand, margins and capital intensity persist together — and whether investors are being adequately compensated if one of those assumptions proves wrong.

WHAT WOULD CHALLENGE OUR THESIS?_

The evidence would become more constructive if AI-attributable revenue approached the return hurdle faster than capital expenditure; hyperscaler incremental returns and free cash flow remained resilient; hardware utilisation stayed high across successive processor generations; and semiconductor margins remained structurally above history after a full capacity cycle.

PORTFOLIO IMPLICATIONS_

Our response is positioning, not prediction.

We are not attempting to forecast the month in which the investment cycle turns. We are trying to own the parts of the value chain where durable economics are most likely to accrue, while avoiding assumptions that require peak conditions to become permanent.

1

We are not trying to time it.

Over-investment cycles end when funding and confidence stop, not when the arithmetic first becomes uncomfortable. The gap between those moments can be years, and attempting to trade it precisely often costs more than it earns.

2

We prefer adopters to builders.

The largest economic gains from a general-purpose technology have often accrued to businesses that use cheaper infrastructure to improve products, productivity or customer economics — not necessarily to the companies that constructed the infrastructure.

3

We insist on return on invested capital, not simply growth.

Growth financed at inadequate returns can destroy value while appearing impressive on the revenue line. For every AI-related holding, we want to understand the incremental capital required and the return expected on that capital.

4

We normalise cyclical earnings.

A low multiple applied to record margins is not a margin of safety. We prefer businesses whose earnings we can broadly recognise five years from now to businesses whose current profitability depends on scarcity and allocation.

5

We watch the balance sheet and the risk transfer.

The move from internally funded expansion towards bonds, leases, private credit and purchase guarantees changes the risk profile of companies and projects. That change should be reflected in the price investors are willing to pay.

THE PRACTICAL DISTINCTION_

We are not against owning the infrastructure layer. Some builders will possess genuine advantages in architecture, software, networking, distribution or cost. The discipline is to identify which companies control scarce economics and which are merely enjoying scarce supply.

The most attractive opportunities may ultimately emerge among businesses that can use abundant, cheaper intelligence to raise revenue, improve margins, reduce working capital or deepen customer relationships without having to fund the underlying infrastructure themselves.

THE PORTFOLIO MOST INVESTORS ALREADY OWN_

The diversification illusion

The AI capital cycle does not sit only inside specialist technology portfolios. It has become a major driver of broad market, global and emerging-market benchmarks — often without investors consciously choosing that exposure.

Semiconductors have reached approximately 18% to 19% of the S&P 500, a record share. One capital cycle therefore influences a very large portion of US index returns, earnings growth and valuation sensitivity. [11]

Geographic diversification does not necessarily escape it. At 30 June 2026, TSMC, Samsung Electronics and SK Hynix represented approximately 31% of the MSCI Emerging Markets Index. An investor moving from the US into emerging markets in search of diversification can therefore be buying more exposure to the same AI infrastructure cycle in a different currency. [12]

A global tracker is more concentrated than it appears. With the US representing roughly 70% of a developed-market global index and the largest US technology companies accounting for a substantial share of the S&P 500, a plain global allocation can hold close to a quarter of its equity exposure in a small group of companies tied to the same investment cycle.

Sector and thematic funds can amplify the exposure further. Semiconductor ETFs are often presented as satellites around a diversified core, yet their largest holdings frequently overlap with companies already dominant in US, global and emerging-market benchmarks.

Style diversification is also less reliable than it once was. The same companies appear across growth, quality, momentum and increasingly core mandates. Allocating across labels does not necessarily allocate across economic outcomes.

> **Cheaper is not the same as diversified.**

A de-rating in the largest technology names may reduce valuation risk, but it does not remove common-factor exposure. The relevant question is not only whether each holding appears attractively valued in isolation. It is how many holdings depend on the same assumptions about AI spending, semiconductor scarcity, funding availability and terminal returns.

A USEFUL PORTFOLIO EXCERCISE_

Look through the labels.
Aggregate exposure across direct AI beneficiaries, hyperscalers, semiconductor suppliers, data-centre equipment, power infrastructure, global indices, emerging-market indices and thematic funds. Then ask how much of the portfolio would be affected if capital-expenditure growth merely slowed — even if AI adoption continued to rise.

A BETTER INVESTOR CONVERSATION_

The questions that matter now

Clients often ask whether AI is real. It is. That question is largely settled. The more valuable questions concern capital, price, durability and portfolio dependence.

- What portion of reported growth is genuinely incremental AI revenue rather than cloud migration, bundling or cannibalisation?
- What utilisation rate, token price and asset life are required for the investment to earn its cost of capital?
- Are current semiconductor earnings being valued on normalised margins or on shortage economics?
- Which companies can adopt AI to improve their economics without having to fund the infrastructure themselves?
- What incremental return on invested capital is implied by the next dollar of infrastructure spending?
- Who bears the risk if demand arrives later than expected — the hyperscaler, the supplier, the project financier or the customer?
- How much portfolio exposure is duplicated across US, global, emerging-market, quality, growth and thematic allocations?

THE DISTINCTION THAT MATTERS_

We are not sceptics about artificial intelligence. We expect it to be as consequential as many of its advocates claim, and we expect the productivity benefits to be widely shared across the businesses that adopt it successfully.

Our caution is narrower and more durable. Transformational technologies have a long history of enriching their users while disappointing many of their financiers — and of paying suppliers handsomely until the moment scarcity becomes capacity.

The capital that builds the infrastructure has repeatedly earned less than the capital that patiently waited to use it. The companies that ultimately create the most value may not be those spending the most today, but those able to convert abundant intelligence into better products, lower costs, stronger customer relationships and higher returns on capital tomorrow.

> **The technology can win. The capital can still lose.**

Keeping portfolios on the right side of that distinction is, in our view, one of the most valuable things an investor can do in this cycle.

RESEARCH NOTES_

Selected sources and data references

The paper combines company guidance, public filings, index data and third-party estimates. Figures are rounded and may differ because of reporting periods, exchange rates, definitions of AI capital expenditure and subsequent guidance changes. Sources accessed or current to 31 July 2026 unless otherwise stated.

1. Company filings and guidance; Financial Times compilation, "Meta, Microsoft, Amazon and Alphabet are about to spend a shocking amount of money to dominate the AI era", 3 June 2026.
2. Goldman Sachs, "The assumptions shaping the scale of the AI build-out", 1 May 2026.
3. Anthropic company revenue updates compiled in May 2026; public reporting on OpenAI annualised revenue.
4. J.P. Morgan Asset Management, "How is AI being monetized?" — estimate of annual revenue required for a 10% return.
5. Bain & Company, "US$2 trillion in new revenue needed to fund AI's scaling trend", Global Technology Report press release, 23 September 2025.
6. AI API Infrastructure Report / CloudEagle analysis of 2.4 billion enterprise API calls, 2026.
7. Barron's, "As AI anxiety spreads, hyperscaler default bets are spiking", 29 July 2026.
8. Reuters, "Google increases capex forecast again after cloud-driven quarterly beat", 22 July 2026.
9. Semiconductor Industry Association / WSTS data for the 2023 semiconductor downturn; company reports and market data for historical cycle examples.
10. Reuters and company announcements on South Korean semiconductor investment plans and Micron's US$250 billion programme through 2035, June–July 2026.
11. S&P Dow Jones Indices and constituent market-cap data; semiconductor weightings are time-sensitive and should be checked at the date of use.
12. MSCI Emerging Markets Index factsheet, 30 June 2026; constituent weights are time-sensitive and may change materially.

IMPORTANT INFORMATION_