ASIA TECH FEED

Special Report · June 2026

Asia Tech Feed on Semiconductors

The trillion-dollar reset: outlook and predictions for the second half of 2026 through 2027

Memflation	Why memory rewrites the cycle — and who pays for it
The packaging decade	CoWoS, SoIC and the new gatekeepers of AI compute
Silicon sovereignty	China's parallel stack, Japan's 2nm gamble, India's entry
Ten predictions	Falsifiable calls for H2 2026–2027, with confidence levels

EDITOR'S NOTE

An industry that found its second trillion faster than it found its first

Twelve months ago, the question hanging over this industry was whether the AI buildout would carry semiconductors past the long-promised trillion-dollar milestone by the end of the decade. That question is now settled — and almost embarrassingly so. The market crossed $772 billion in 2025, and the forecasting establishment now disagrees only on *how far past* $1 trillion 2026 lands: Gartner sees revenue exceeding $1.3 trillion, while WSTS's spring revision reaches $1.5 trillion, with a path toward $1.9 trillion in 2027.

A spread that wide between the industry's two most-watched forecasters is not noise. It is the defining feature of this moment. Nobody — not the foundries, not the memory makers, not the hyperscalers writing the checks — has a reliable model for an industry in which a single end market (AI infrastructure) now sets the price of memory for every other buyer on Earth, and in which packaging plants, not wafer fabs, decide how many accelerators ship.

This report is our attempt to navigate that uncertainty honestly. We synthesize forecasts from WSTS, Gartner, SEMI, TrendForce, Deloitte and company disclosures; where they disagree, we show the disagreement rather than averaging it away. And in Chapter 7 we do what most outlooks avoid: we make ten specific, falsifiable predictions for the second half of 2026 through 2027, each tagged with an explicit confidence level. We will grade ourselves publicly when the window closes.

Our vantage point is deliberately Asian. More than 75% of the world's chips are fabricated in Asia; the memory supercycle is being decided in Icheon and Pyeongtaek; the packaging bottleneck in Taichung and Chiayi; the sovereignty contest in Shanghai, Chitose and Dholera. Whatever happens to semiconductors in the next eighteen months happens here first.

— **The Editors, Asia Tech Feed**
June 2026

HOW TO READ THIS REPORT

Forecast figures are drawn from named third-party sources and company guidance current as of June 2026. Items marked **F** are third-party forecasts; items marked **E** are Asia Tech Feed estimates. Chapter 7 predictions carry confidence tags — **High** (>75% likelihood in our judgment), **Medium** (50–75%), and **Contrarian** (<50%, flagged where we knowingly depart from consensus). Nothing in this document is investment advice.

CONTENTS

THE NEXT EIGHTEEN MONTHS, AT A GLANCE

$1.3–1.5T
2026 industry revenue
Gartner / WSTS forecast range

~30%
of 2026 chip revenue from AI silicon
Gartner

+125% / +234%
DRAM / NAND price rise, 2026F
Gartner; relief only late 2027

$52–56B
TSMC capex, 2026
70–80% to advanced nodes

+68%
HBM demand growth, 2027F
TrendForce; HBM4E ≈ 40% of mix

$156B
equipment sales, 2027F
SEMI — an all-time record

01 The trillion-dollar reset

The semiconductor industry will clear $1 trillion in 2026 with room to spare. The real story is the $200 billion gap between what the two most credible forecasters think comes next.

After growing 22.5% to $772 billion in 2025, the industry has entered the steepest expansion in its modern history. Gartner's April 2026 forecast puts 2026 revenue above $1.3 trillion — 64% growth, the fastest in two decades. WSTS's spring 2026 revision goes further, to roughly $1.5 trillion, and projects another 27% gain in 2027 to approximately $1.9 trillion. Both agree on the engine: AI infrastructure now accounts for roughly 30% of total semiconductor revenue, and Deloitte estimates the AI chip market alone at about $500 billion in 2026.

Our view is that the truth for 2026 sits in the lower half of that band. The WSTS number leans heavily on memory pricing holding its extraordinary trajectory through the year; Gartner's own analysts caution that "memflation is profound, but it is not perennial." Either way, the structural conclusion is identical: this is no longer a cyclical recovery. It is a re-rating of the industry's baseline scale.

The composition of growth matters more than the headline. Memory is doing the heavy lifting — WSTS sees the segment surging toward $800 billion in 2026, with 32% further growth in 2027 — while logic grows 27–37% on AI accelerators and the 2nm ramp. The legacy economy tells a different story: discretes (+8.0%), optoelectronics (+2.7%) and sensors (+3.0%) are growing at pre-AI rates. The "tale of two markets" that defined 2025 has not closed; it has widened into a canyon.

For Asia, the regional skew is stark. The memory windfall accrues almost entirely to Korea (SK hynix, Samsung) and to Micron's Taiwan and Japan fabs. Taiwan captures the logic and packaging boom. China grows on volume in mature nodes and domestic AI silicon, while Japan's October 2025 sales were still declining year-on-year — a reminder that the AI tide is lifting specific boats, not all of them.

EXHIBIT 1

Global semiconductor revenue approaches two trillion dollars

Annual revenue, $ billion; 2026–27 shown as forecast range

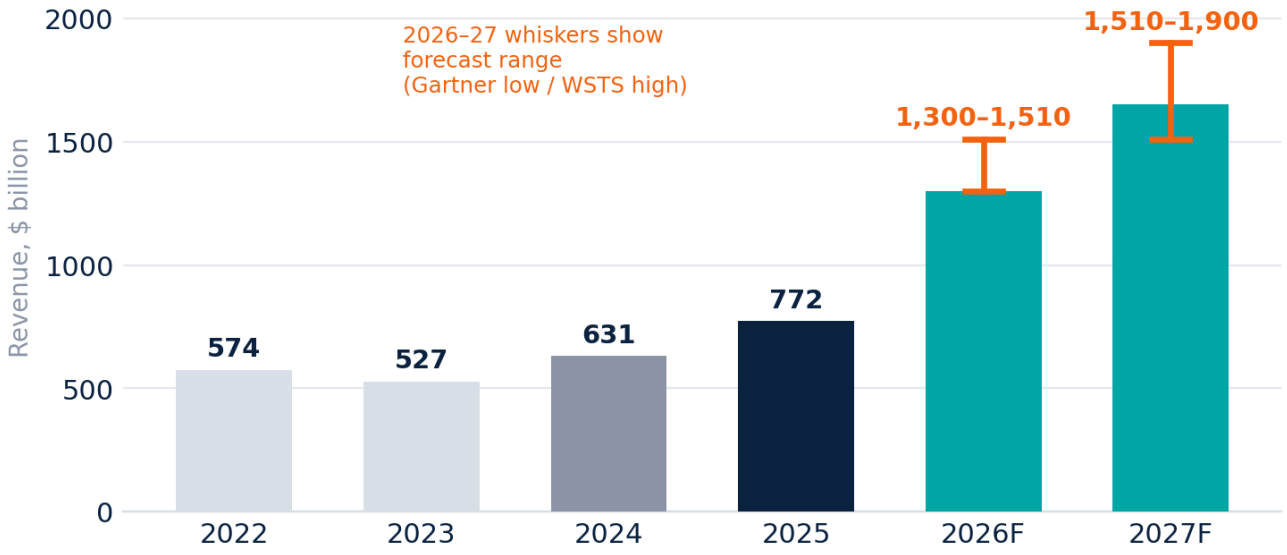

Source: WSTS (historical, spring 2026 forecast); Gartner (April 2026); SIA. ATF illustration of forecast dispersion.

EXHIBIT 2

Growth is violently uneven across product segments

YoY revenue growth by segment, %; note memory's 2026 spike of ~250%

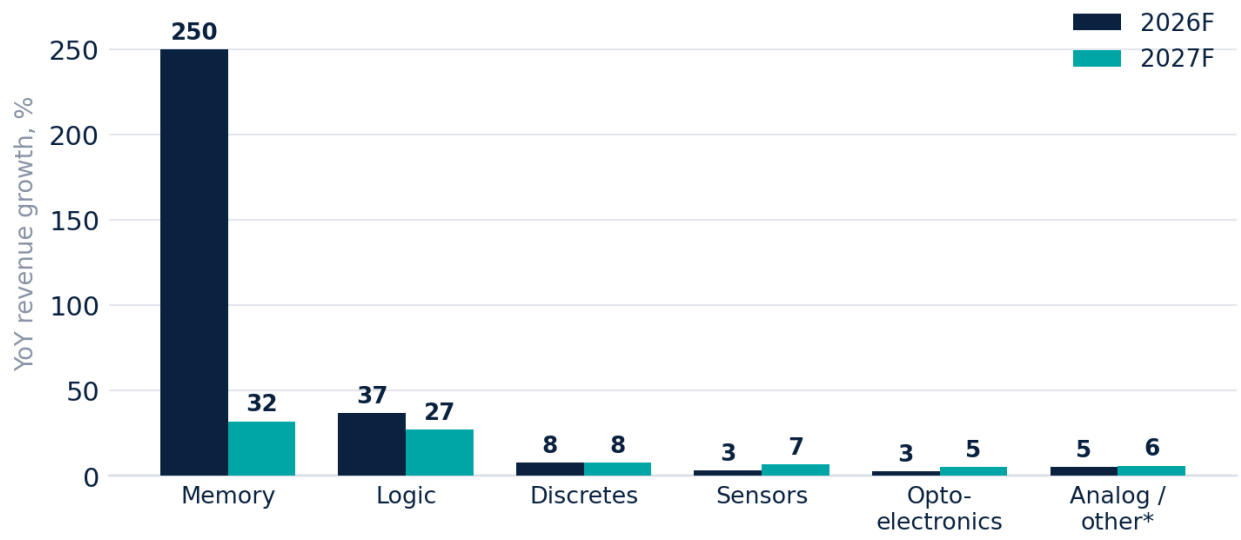

Source: WSTS spring 2026 forecast. *Analog/other: ATF estimate.

Why the forecasters disagree. The $200 billion spread between Gartner and WSTS for 2026 reduces almost entirely to one variable: how long memory contract prices defy gravity. WSTS effectively annualizes the pricing seen in early 2026; Gartner assumes moderation in the second half. Track quarterly DRAM contract resets — they will tell you in real time which house is winning.

What we are watching in H2 2026. Three tells: (1) whether hyperscaler capex guidance for 2027 — already implying cloud capex above $725 billion — survives earnings season intact; (2) whether non-AI demand (autos, industrial, consumer) shows price-driven demand destruction; (3) whether China's import substitution dents multinational unit volumes faster than expected.

THE BOTTOM LINE

Treat $1.3 trillion as the 2026 floor, not the stretch case. But recognize what kind of number it is: roughly half the increment is price, not units. An industry whose growth is this dependent on the scarcity pricing of one component class — memory — has imported a new kind of fragility along with its new scale.

> **“Roughly half the increment is price, not units. The industry has imported a new kind of fragility along with its new scale.”**
>
> Asia Tech Feed analysis, Chapter 1

02 Memflation

Memory has gone from the industry's most cyclical commodity to its scarcest strategic asset. The supercycle will hold through 2027 — and everyone outside AI will pay for it.

The numbers no longer resemble a normal cycle. Gartner expects DRAM average prices to rise 125% and NAND 234% across 2026, with meaningful relief "not expected until late 2027." The cause is structural: HBM production consumes roughly three times the wafer capacity of standard DRAM per bit, and every leading-edge DRAM line converted to HBM removes supply from a conventional market that still has to feed servers, PCs and phones. Add the reported lockup of close to 40% of global DRAM wafer output through 2029 by a single AI buyer for its data-center program, and the conventional market is being rationed.

The casualties are predictable. Gartner warns memflation "will destroy, or at least delay, non-AI demand into 2028." PC and smartphone makers face bill-of-materials inflation they cannot fully pass through; automotive and industrial buyers, accustomed to long fixed-price contracts, are discovering those contracts no longer get signed.

The HBM hierarchy hardens — then mutates. TrendForce sees HBM demand growing 77% in 2026 and 68% in 2027. SK hynix enters this window with more than half the market and the deepest NVIDIA relationship; Samsung, first to ship HBM4 commercially in February 2026, is staging the strongest share recovery; Micron expands aggressively from a small base with HBM4 in 2026 and HBM4E in 2027 co-developed with TSMC.

The mutation is customization. HBM4E — expected to reach ~40% of HBM demand in 2027 — introduces customer-specific logic base dies. When Google, Amazon and the GPU vendors can order memory with their own logic embedded, HBM stops being a commodity and becomes a design win. That favors whoever pairs memory with a foundry: advantage Samsung (in-house) and the Micron–TSMC axis, eroding the pure-merchant model that built SK hynix's lead.

EXHIBIT 3

Memflation: the steepest memory pricing cycle on record

Annual average contract price change, %

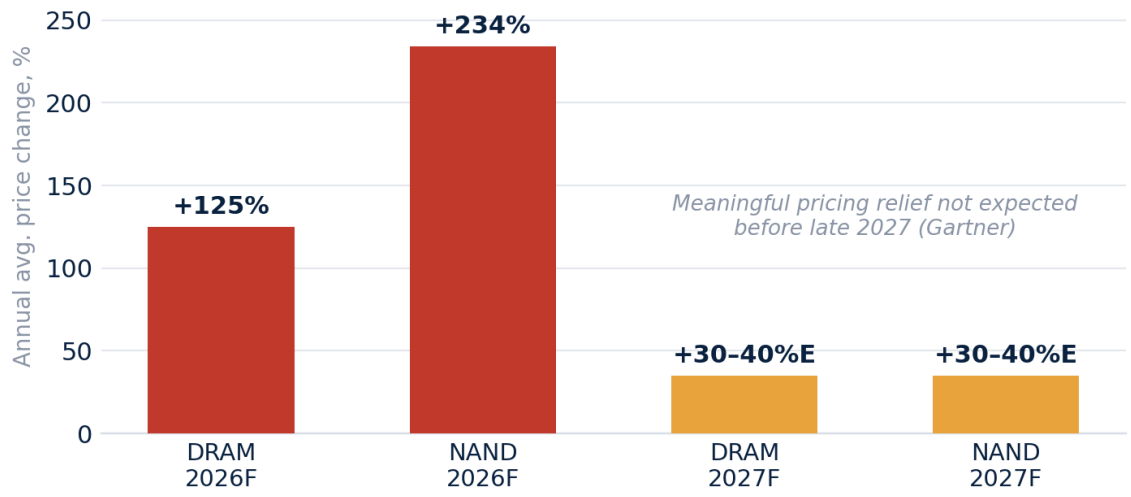

Source: Gartner (April 2026) for 2026; 2027 ATF estimate consistent with Gartner's 'moderating but persistent' guidance.

EXHIBIT 4

HBM demand compounds while the generation mix turns over twice in two years

Left: bit demand index. Right: share of demand by generation

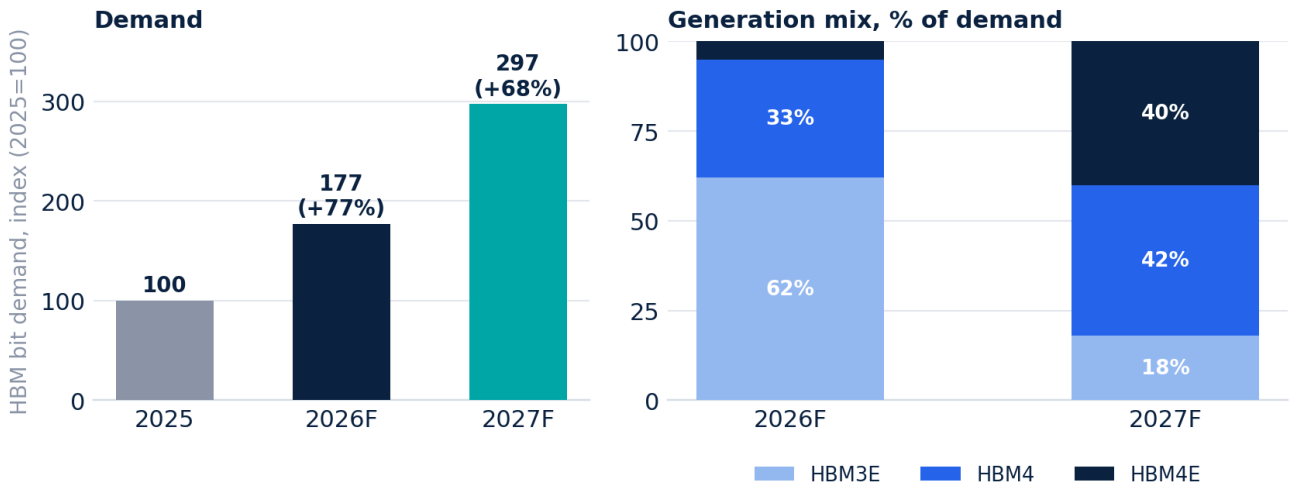

Source: TrendForce (Nov 2025, Feb 2026); Chosun Biz industry estimates; ATF analysis.

Supply arrives — late, and lumpy. The capacity that eventually breaks the cycle is already being poured: Samsung's P4L and SK hynix's M15X reach volume in 2027, and SK hynix is accelerating its Yongin buildout. Our base case: bit supply growth crosses demand growth around mid-to-late 2027, prices plateau before they fall, and 2028 — not 2027 — is when oversupply risk becomes real. Memory makers will exit this cycle with the strongest balance sheets in their history.

The Korea question. Memflation is, in effect, the largest transfer of profit pool toward Korean industry since the display era. Watch for second-order effects in H2 2026: aggressive capex (already visible), talent wars spilling into Japan and the US, and renewed political attention — in Washington and Beijing alike — to the fact that the world's AI buildout now has a single-country dependency in memory at least as acute as its Taiwan dependency in logic.

THE BOTTOM LINE

Do not model memory as a cycle that mean-reverts on schedule. The HBM wafer tax, multi-year hyperscaler lockups, and customization all extend the squeeze. We expect tightness through 2027, visible relief only in its final quarter — and the non-AI economy, not the memory makers, absorbing the cost.

03 The packaging decade

Two-nanometer silicon will be abundant before the packages that carry it are. Through 2027, advanced packaging — not lithography — is the gate on the world's AI compute.

TSMC's 2nm ramp is the fastest node introduction in its history: volume production began in Q4 2025 with yields above plan, capacity is on a path from ~40,000 wafers per month to 80,000–90,000 by end-2026, and the company guides a 70% CAGR for 2nm-class capacity through 2028. A16 — N2 with backside power delivery — enters volume in H2 2026 for HPC customers. The node is effectively sold out: both Taiwan 2nm fabs are booked through 2026, with Apple, AMD, and the AI accelerator complex queued behind.

Funding it is a $52–56 billion 2026 capital budget — TSMC's largest ever — with 70–80% directed at advanced nodes and 10–20% at advanced packaging and test. Management's discipline cuts both ways: TSMC historically does not build ahead of committed orders, so the capex number itself is the strongest demand signal in the industry.

The real bottleneck moved downstream. CoWoS capacity doubles again to a projected 130,000–150,000 wafers per month by late 2026, and SoIC grows even faster — TSMC guides 80%+ annual growth for both through 2027. It is still not enough. NVIDIA alone has reportedly secured over 60% of 2026 packaging capacity for the Rubin ramp, and TSMC concedes AI capacity "front-end and back-end, remains very tight."

The physics is forcing a format change. At Rubin's 5.5-reticle package size — now in production at 98% yield — a 300mm wafer yields as few as four to seven units. Hence CoPoS: panel-level packaging, with a pilot line completing in June 2026 and volume targeted for 2028–29. Watch also the challengers: Intel's 18A is in volume with external customers courting, Samsung has reportedly won a marquee 2nm defection from AMD for part of its roadmap, and Rapidus is betting its entire model on fully-automated packaging turnaround. None of them dents TSMC's share before 2028 — but pricing leverage at the margin begins earlier.

EXHIBIT 5

TSMC's largest capital budget in history is a demand signal, not a bet

Capital expenditure, $ billion

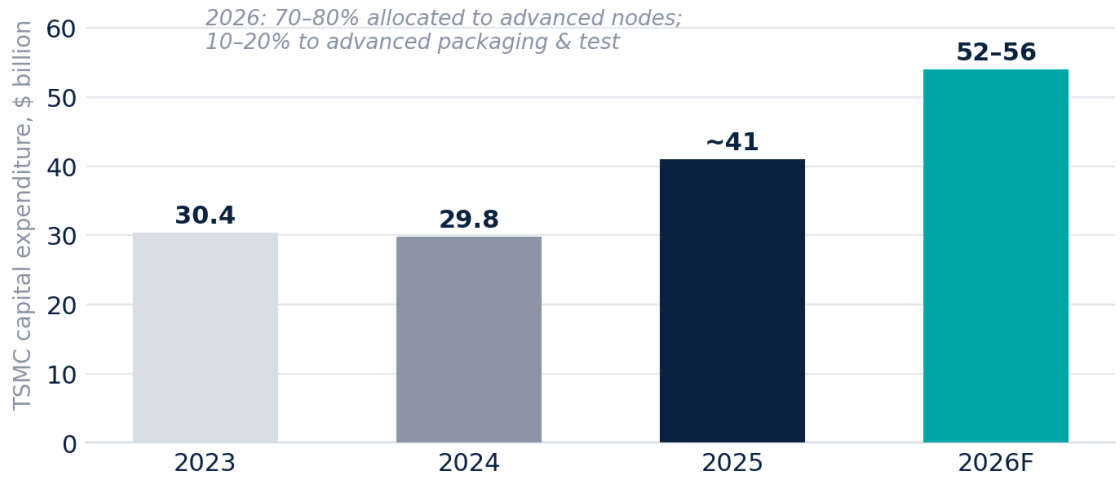

Source: TSMC earnings calls and guidance (Jan 2026); Commercial Times; TrendForce.

EXHIBIT 6

Both gates on AI compute are opening — at different speeds

Capacity, wafers per month (TSMC; 2027 ATF estimates)

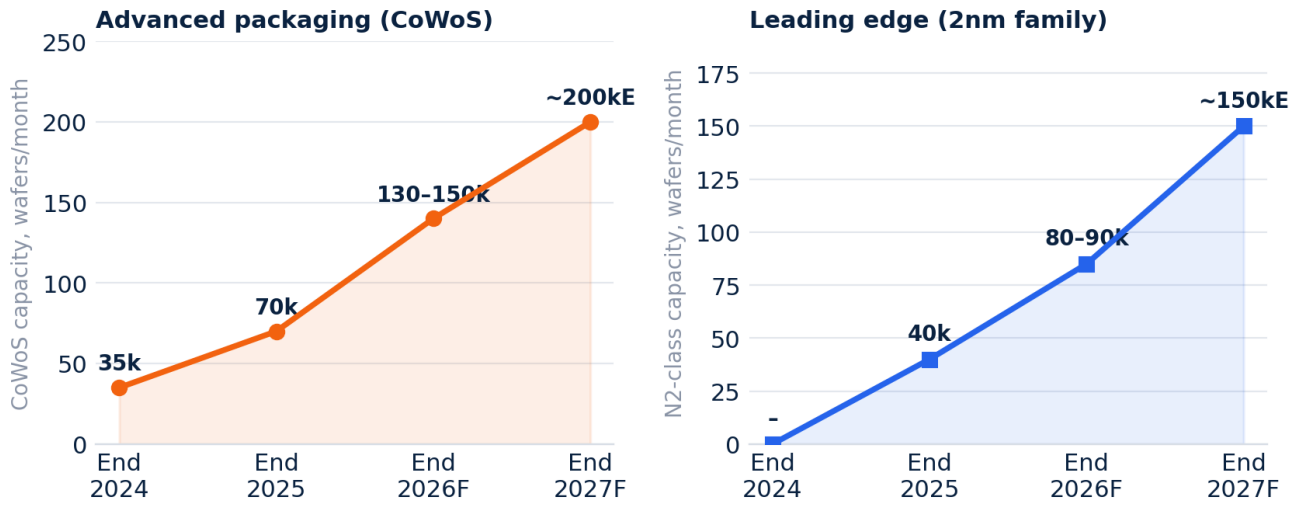

Source: TrendForce; Focus Taiwan / TSMC 2026 Technology Symposium; Tom's Hardware; FinancialContent. 2027E: ATF.

Geography is changing faster than share. Arizona moves from symbol to substance: a fourth fab begins construction in 2026 alongside TSMC's first US advanced-packaging plant, part of a stated path to ~30% of 2nm-and-below capacity in the US over time. Kumamoto and Dresden ramp specialty nodes. For customers, the premium for non-Taiwan wafers becomes a permanent line item — one they are demonstrably willing to pay.

What could break the thesis. Packaging is now single-point-of-failure infrastructure. A tool-delivery slip at AP7/AP8, substrate shortages (ABF remains tight), or a Taiwan power event would propagate to every AI roadmap on Earth within one quarter. Conversely, if CoPoS panelization works early, accelerator unit economics improve abruptly in 2028 — a deflationary surprise almost no one is modeling.

THE BOTTOM LINE

Through 2027, count packages, not wafers. CoWoS/SoIC slots are the true unit of AI compute supply, and TSMC has become their sole meaningful clearinghouse. Every accelerator forecast you read is, implicitly, a forecast of Taichung and Chiayi construction schedules.

> **“Every accelerator forecast you read is, implicitly, a forecast of Taichung and Chiayi construction schedules.”**
>
> Asia Tech Feed analysis, Chapter 3

04 Silicon sovereignty: China's parallel stack

Export controls did not stop China's AI silicon program. They changed its economics — and created a parallel ecosystem that no longer needs to match NVIDIA to matter.

The 2026 policy settlement is best described as managed bifurcation. Washington's January framework permits H200-class exports under a 50% volume cap, third-party security testing and a 25% tariff, while Blackwell-generation parts remain barred; Beijing answers with "parallel purchase" rules pairing every imported Western chip with a domestic deployment. The year of uncertainty preceding the deal did lasting damage to US vendors' China franchise: share ceded to Huawei's Ascend line during the embargo window is not coming back simply because licenses now exist.

On the supply side, China is brute-forcing scale on DUV. SMIC's 7nm-class capacity is on track to reach roughly 60,000 wafers per month in 2026 and 80,000 in 2027; 5nm-class output exists but at yields reported between 20% and 40% — economically irrational by merchant standards, strategically rational under Big Fund III's $47.5 billion umbrella. Hua Hong is readying 7nm at Huali, ending SMIC's domestic monopoly on advanced logic.

Huawei targets 1.6 million Ascend dies in 2026 — including some 600,000 910Cs — roughly eight times its constrained 2025 output, with the MindSpore stack maturing into a credible CUDA substitute for domestic inference. Independent analysis (AEI, April 2026) concludes the installed DUV-immersion fleet can "almost certainly" supply Huawei's domestic accelerator demand this year.

Two caveats keep the gap real. First, compute economics: US-aligned analysts estimate China produces advanced chips at 1–4% of US-allied capacity, implying a 20–50x annual AI-compute disadvantage even in permissive scenarios. Second — and this is the chokepoint we would watch above all others — **HBM**. Ascend's binding constraint is not logic dies but memory: domestic HBM (CXMT-led) remains generations behind, and HBM is squarely inside the export-control perimeter. China's accelerator output in 2027 will be set in memory fabs, not logic fabs.

EXHIBIT 7

China's parallel stack scales on volume, not parity

Left: ≤7nm-class wafer capacity. Right: Huawei Ascend die output

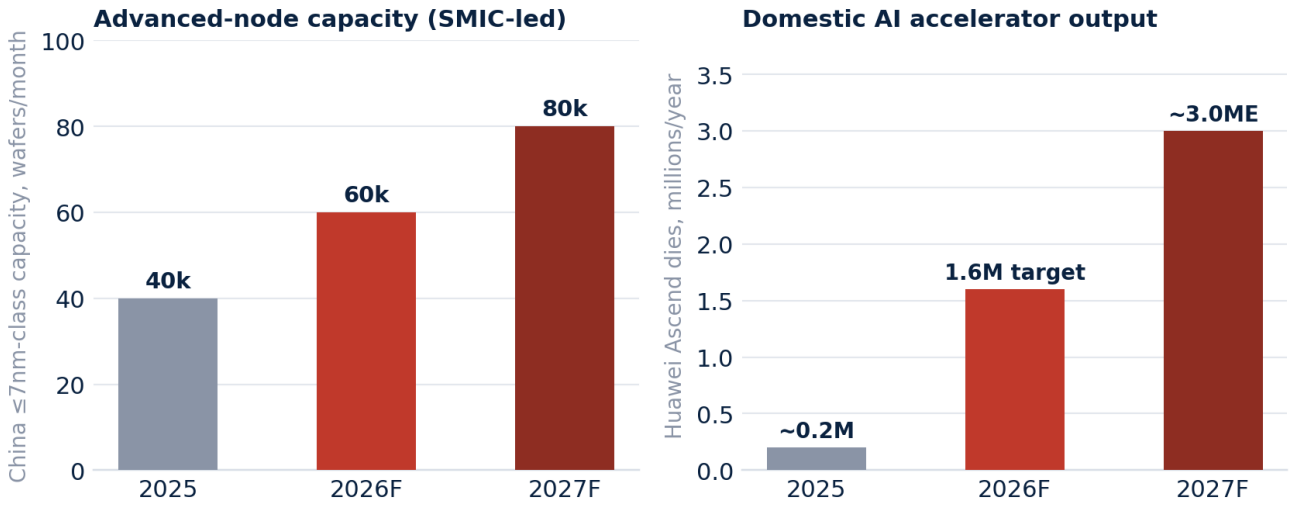

Source: Industry estimates via Oplexa, AEI (Apr 2026), Reuters; US Dept. of Commerce testimony (2025). 2027E: ATF.

THREE SCENARIOS FOR THE BIFURCATED MARKET, H2 2026–2027

BASE CASE — MANAGED BIFURCATION

55%
ATF subjective probability

The January framework holds with periodic friction. H200-class flows resume at capped volumes; Chinese hyperscalers split procurement between NVIDIA and Ascend under parallel-purchase rules. SMIC reaches ~80k wpm advanced capacity in 2027; HBM scarcity caps Ascend output near 3M dies. Multinationals retain a smaller, tariffed, but real China business.

ESCALATION — THE PERIMETER WIDENS

30%
ATF subjective probability

A security incident or political cycle triggers revocation of H200 licenses and extension of controls to subsystems (HBM brokering, EDA cloud access, tool parts). China retaliates via rare earths, legacy-chip dumping and procurement bans. Non-AI chip prices fall as Chinese mature-node capacity floods export markets; AI silicon bifurcates completely.

DÉTENTE — TRANSACTIONAL OPENING

15%
ATF subjective probability

Broader trade settlement loosens caps in exchange for purchases and tariff revenue. NVIDIA's China share partially recovers, slowing — but not stopping — Ascend adoption: state-directed demand now has institutional momentum independent of price or performance.

Second-order effects worth pricing now. (1) Equipment: China remains the largest single buyer of wafer-fab equipment as it stockpiles against future restrictions — a tailwind for Tokyo Electron and the US toolmakers that policy hawks keep trying to close. (2) Legacy-node deflation: Chinese mature-node overcapacity is the one segment of this industry where prices fall through 2027, pressuring Taiwanese and Japanese specialty foundries. (3) Talent: Big Fund III's horizon runs to 2039; the recruiting wave across Korean and Taiwanese memory and packaging engineers is just beginning.

05 The toolmakers' supercycle

Equipment spending sets records in both 2026 and 2027 — and for the first time, the back end grows as a strategic asset rather than an afterthought.

SEMI's projections compound across every category: total equipment sales of $145 billion in 2026 and a record $156 billion in 2027; wafer-fab equipment reaching $135 billion; 300mm fab equipment spending up 18% in 2026 and another 14% in 2027, crossing $150 billion for the first time and continuing toward $172 billion by 2029. Three drivers stack: the 2nm gate-all-around transition (lithography- and deposition-intensive), the memory buildout (HBM tooling, hybrid bonding), and sovereignty programs from Arizona to Chitose to Dholera that buy tools for strategic rather than purely economic reasons.

The structural novelty is the back end. Test equipment grew 48% in 2025 and keeps compounding as AI devices demand longer, costlier test flows; assembly and packaging tools grow near double digits on heterogeneous integration. The decades-old hierarchy in which back-end suppliers traded at a discount to front-end peers is dissolving.

Where the money lands. ASML's EUV and High-NA backlog stretches years out, and TSMC's accelerating 2nm buildout flows straight to Veldhoven, Applied Materials, Lam, KLA and Tokyo Electron. Korea's memory expansion (P4L, M15X, Yongin) is the swing factor for 2027 — if it accelerates, SEMI's record forecast will prove conservative. China remains the largest single regional buyer, concentrated in mature nodes and increasingly in domestic toolmakers (Naura, AMEC) for everything outside the chokepoints.

The fragility. Equipment is the industry's leading indicator in both directions. If hyperscaler capex guidance cracks in late 2026, tool orders are the first casualty — and the 2027 record evaporates. Watch ASML bookings and TEL's order commentary as the earliest honest signals of any AI digestion phase.

EXHIBIT 8

Two consecutive record years for semiconductor equipment

Sales, $ billion

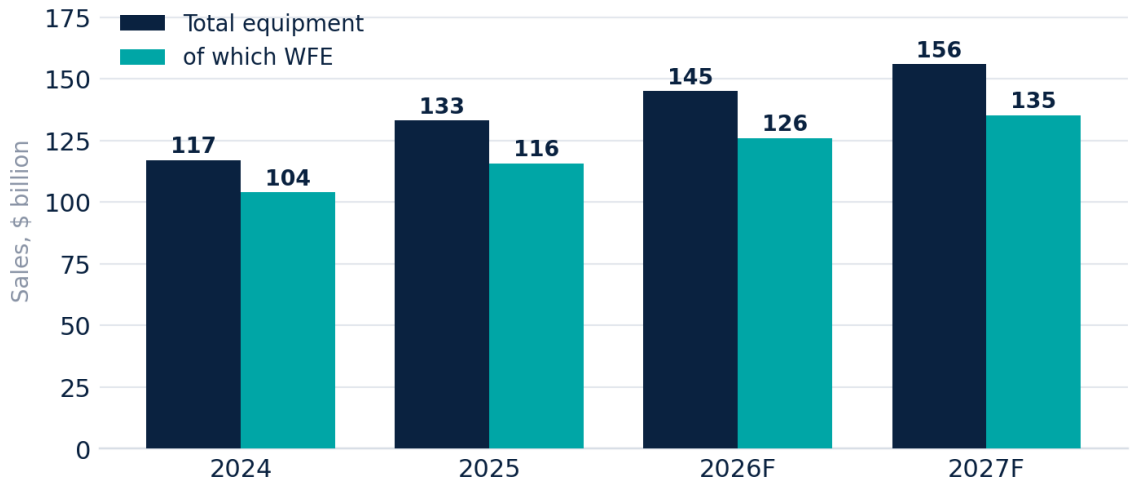

Source: SEMI year-end 2025 Total Equipment Forecast and 300mm Fab Outlook (Apr 2026); 2026 WFE: ATF interpolation of SEMI growth rates.

06 The rest of Asia rises

Beyond the Taiwan–Korea–China triangle, three storylines mature in this window: Japan's 2nm gamble meets its deadline, India breaks ground at scale, and ASEAN quietly captures the packaging overflow.

Japan: Rapidus meets its moment. The boldest industrial-policy experiment in modern Japan hits its decisive year. Rapidus — 2nm GAA transistors demonstrated in July 2025, pilot line running in Chitose, ¥267.6 billion in fresh funding and the government now its largest shareholder with a golden share — targets mass production in 2027. Its wedge is deliberate: small batches, the shortest turnaround in the industry, and fully automated advanced packaging co-located with the fab from 2027. It will not threaten TSMC's volume. It does not need to: a credible second source for low-volume leading-edge silicon — defense, edge AI (the Tenstorrent partnership), automotive — is worth more to Tokyo than market share. The honest risks are customers and yield, in that order. Kumamoto (TSMC), Micron's DRAM expansion and a reshored tool ecosystem give Japan its strongest semiconductor hand in thirty years regardless of how the Rapidus bet itself resolves.

India: from approvals to concrete. The India Semiconductor Mission has now approved ten fab and OSAT projects across two waves, anchored by Tata's Dholera fab and assembly plants from Micron and others.

The 2026–27 window is when India's program becomes falsifiable: first commercial OSAT output ships, Dholera's tool move-in begins, and design-linked incentives either produce indigenous silicon tape-outs (C2i's AI power chip is an early proof point) or don't. Our view: India wins meaningful share in packaging, test and power/analog this decade — not logic fabrication — and that is a perfectly good prize.

ASEAN: the overflow economy. Malaysia's Penang corridor, Vietnam and Singapore are the quiet beneficiaries of both packaging scarcity and China de-risking. NVIDIA's server-manufacturing hiring in Vietnam alongside expanding Taiwanese ODM capacity signals where AI hardware assembly diversifies next; Singapore adds fab capacity in specialty nodes; Malaysian OSATs run at capacity on the test-time inflation that AI devices created. None of this shows up as a headline number. All of it compounds.

The power constraint is regional, too. AI data centers need an estimated 92 GW of additional power by 2027 globally. In Asia outside China, grid interconnection queues — Japan, Malaysia, India — are now a harder gate on data-center growth than chip supply. Land, power and cooling, not GPUs, set the 2027 ceiling for several ASEAN markets.

THE BOTTOM LINE

The semiconductor map of Asia is fragmenting productively. Japan re-enters the leading edge as a boutique, India industrializes the back end, ASEAN absorbs the overflow — and every one of these stories is gated by electricity before it is gated by silicon.

07

Ten predictions for H2 2026–2027

Specific, falsifiable, and time-boxed. Each call carries our confidence level; we will publish a public scorecard when the window closes at end-2027.

01

2026 closes between $1.30 and $1.45 trillion

HIGH

Below WSTS's $1.51T spring call, at-or-above Gartner's $1.3T. Memory pricing moderates in H2 as hyperscalers resist annualized triple-digit increases; unit growth stays intact.

02

Memory peaks as a share of industry revenue in 2027 — above 40%

HIGH

Memory exceeds $800B in 2026 and grows ~30% further in 2027, becoming the largest product category for the first time since 2018. DRAM contract prices do not decline year-on-year until Q4 2027 at the earliest.

03

HBM4E customization triggers the first major share shift since HBM began

MEDIUM

By end-2027, SK hynix's HBM revenue share falls below 50% as custom base dies advantage Samsung's foundry pairing and the Micron–TSMC axis. SK hynix remains #1 — but the moat narrows visibly.

04

Packaging stays the binding constraint for all of 2026 — and most of 2027

HIGH

CoWoS/SoIC slots, substrates and HBM stacks, not 2nm wafers, cap accelerator shipments. TSMC's advanced packaging revenue share roughly doubles from ~10% toward 15–20% of total revenue by late 2027.

05

TSMC's 2027 capex guides above $60 billion

MEDIUM

The 2nm/A16 demand wave, US packaging buildout and CoPoS industrialization push the January 2027 guide past $60B. Consensus today clusters in the mid-$50s; the surprise is upward.

06 **Rapidus ships first commercial 2nm wafers in 2027 — with fewer than five named customers** MEDIUM

Chitose makes its date in volume terms that are strategically meaningful and commercially small. The stock of announced customers, not yield rumors, is the metric to watch through 2027.

07 **Huawei clears 2.5M+ Ascend dies in 2027 — but HBM caps real deployments** HIGH

Logic capacity supports the target; domestic HBM does not. Effective deployed Ascend compute undershoots die output materially, and HBM smuggling/brokering becomes the export-control story of 2027.

08 **At least one $10B+ hyperscaler data-center program slips publicly on power, not chips** HIGH

The 92GW global power gap bites. A named campus delay — grid interconnect, turbines, or transformers — becomes the canonical evidence that electricity replaced silicon as the AI constraint.

09 **China's mature-node exports trigger formal trade action by the US or EU in the window** MEDIUM

Legacy-chip overcapacity meets industrial-policy politics: anti-dumping duties or tariff escalation on ≥28nm Chinese chips, with Taiwanese and Japanese specialty foundries as collateral damage.

10 **No AI winter — but an 'equipment air pocket' in late 2027** CONTRARIAN

Contrarian to the record-forecast consensus: as memory capacity lands and one hyperscaler digests, tool orders flatten for two to three quarters from late 2027. The secular story resumes in 2028; the straight line does not.

HOW TO HOLD US ACCOUNTABLE

Each prediction resolves against named public sources — WSTS/Gartner totals, TrendForce HBM share, TSMC disclosures, SEMI equipment data, company announcements. The scorecard publishes at asiatechfeed.com in January 2028. We expect to be wrong at least twice; if we are not, we were not specific enough.

08 Risk map and watchlist

The bull case is consensus. These are the variables that would change our numbers — ranked by probability and impact over the forecast window.

EXHIBIT 9

What we worry about, and how much

ATF subjective assessment, H2 2026–2027 horizon

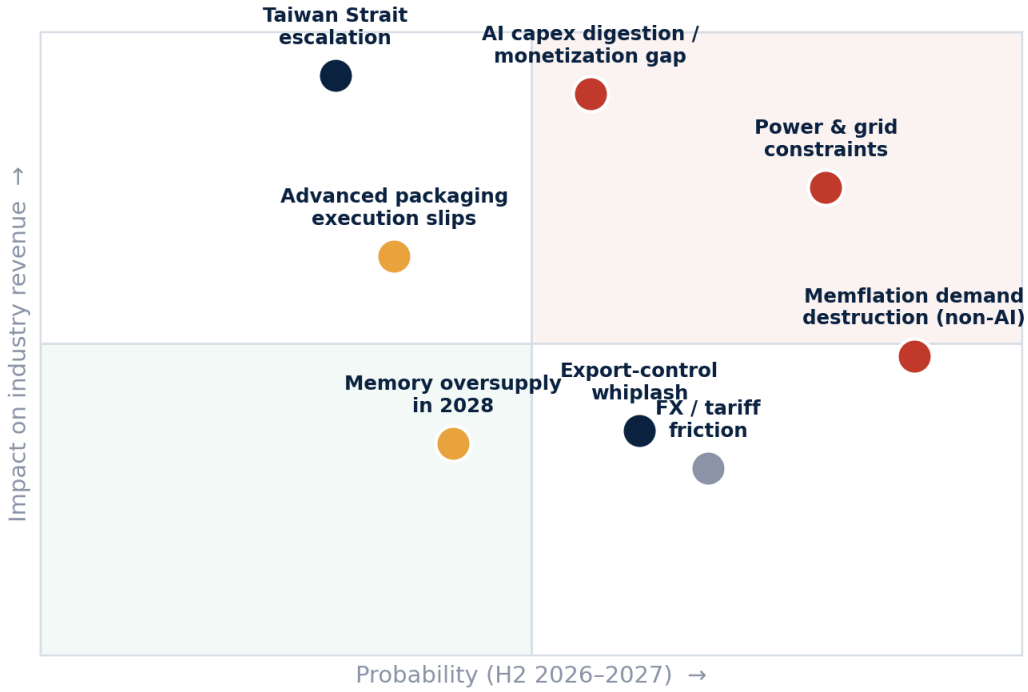

Source: Asia Tech Feed analysis.

The quarterly watchlist. (1) DRAM/NAND contract resets — the single best real-time test of the memflation thesis. (2) Hyperscaler capex guidance and AI revenue disclosure — the demand side's honesty check. (3) ASML and Tokyo Electron bookings — the earliest leading indicator of any digestion phase. (4) CoWoS/SoIC tool move-in schedules at AP7/AP8 — the supply side of every accelerator forecast. (5) US license and rule-making actions on H200-class exports and HBM — the bifurcation thermostat. (6) Grid-interconnection announcements in Japan, Malaysia, India — the power ceiling. (7) Rapidus customer announcements. (8) Chinese mature-node pricing and any EU/US trade filings. When three or more of these move against the bull case in a single quarter, re-underwrite everything.

Methodology, sources and disclosures

Scope. This report covers the global semiconductor industry and adjacent infrastructure (memory, foundry, advanced packaging, equipment, and AI data-center demand) over the period July 2026 through December 2027, with an Asian editorial lens.

Sources. Market sizing draws on WSTS (autumn 2025 and spring 2026 forecasts), Gartner (April 2026), SIA monthly data, and Deloitte's 2026 industry outlook. Memory and HBM analysis draws on TrendForce, Counterpoint, company disclosures from SK hynix, Samsung Electronics and Micron, and Korean trade press via Chosun Biz. Foundry and packaging analysis draws on TSMC earnings calls and 2026 Technology Symposium disclosures, TrendForce, Focus Taiwan/CNA, Commercial Times, and Tom's Hardware reporting. China analysis draws on AEI (April 2026), CSIS, Reuters, US Department of Commerce testimony, and trade-press estimates. Equipment data is from SEMI's year-end 2025 forecast and April 2026 300mm Fab Outlook. Regional coverage draws on Rapidus corporate disclosures, The Register, the India Semiconductor Mission, and DigiTimes.

Estimates and judgment. Figures marked F are third-party forecasts; figures marked E are Asia Tech Feed estimates derived from the growth rates and capacity disclosures cited. Scenario probabilities and prediction confidence levels are the editors' subjective judgments. Where credible sources conflict — most notably the 2026 market total — we present the range rather than a false point estimate.

Independence. Asia Tech Feed has no commercial relationship with any company discussed. This document is editorial analysis, not investment research, and nothing herein constitutes investment, legal or procurement advice. All trademarks belong to their owners. McKinsey & Company is not affiliated with, and has not reviewed or endorsed, this publication.

Corrections. Errors will be corrected at asiatechfeed.com/corrections. The prediction scorecard publishes January 2028.