



INTERNATIONAL RESEARCH INSTITUTE — MARKET STUDY

Semiconductors and Critical Supply Chains: Geoeconomics and Market Structure

Semiconductors & supply chains · 2026-03-04 · Licence CC BY 4.0

The industry in brief

Semiconductors have moved from the back office of industrial policy to its centre. For four decades the chip industry organised itself around a single logic — comparative advantage, ruthless specialisation, and the geographic distribution of each production step to wherever it was cheapest and most efficient. That logic produced extraordinary cost declines and an equally extraordinary concentration of capability. The result is a value chain that is globally distributed in aggregate yet locally monopolised at almost every critical node: one company in the Netherlands supplies the lithography tools required for the most advanced chips; one company in Taiwan manufactures the overwhelming majority of leading-edge logic; a small number of firms in Korea and the United States hold the memory market. The system is efficient. It is not, by design, resilient.

This report assesses the industry as both a market and a geoeconomic system. On the market side, global semiconductor revenue reached roughly US\$600–630bn in 2024, and industry bodies and consultancies broadly converge on a trajectory toward US\$1 trillion around 2030 — a projection we treat as a plausible central estimate rather than an established fact. On the geoeconomic side, the period since 2022 has seen the largest coordinated industrial-policy intervention in the sector's history: the United States CHIPS and Science Act, the EU Chips Act, and comparable programmes in Japan, Korea, China, India and elsewhere have together committed well over US\$400bn in public support. Simultaneously, export controls on advanced computing chips and the equipment used to make them have converted a commercial supply chain into a contested strategic terrain.

The sector at a glance

INDICATOR	VALUE
Global device-market revenue, 2024	US\$600–630bn
Widely cited trajectory around 2030 — plausible central estimate, not a forecast	~US\$1tn
Public support announced since 2022	>US\$400bn
Sub-10nm logic fabricated in Taiwan — estimate; stated as a range	80–90%

Our central judgement is that these two forces — reshoring subsidy and strategic denial — are reshaping the industry's map without dissolving its fundamental interdependence. Capacity is being duplicated at the frontier and diversified geographically, but the economics of the sector make full national self-sufficiency uneconomic for every actor, including the largest. The likely outcome to 2030 is not decoupling in the clean sense often implied by political rhetoric, but a partial, expensive and uneven bifurcation: a guarded frontier for the most advanced logic and memory, layered over a still-globalised market for mature nodes, assembly and materials.

For decision-makers the implication is that supply security will be bought in increments, at a premium, and never completely. Governments should calibrate expectations accordingly; investors should price a more persistent capital-intensity and policy-risk environment; and industrial users should treat single-source exposure at specific chokepoints — advanced packaging, high-bandwidth memory, photoresists, lithography — as the binding constraints rather than headline wafer capacity.

What we conclude

- **The market is large but structurally brittle.** A ~US\$600–630bn industry (2024) depends on multiple single-points-of-failure. The most acute is geographic: on our reading of public capacity data, well over 80–90% of sub-10nm logic is fabricated in Taiwan, and effectively all extreme-ultraviolet (EUV) lithography tools originate from one Dutch supplier.
- **Subsidies buy time and capacity, not independence.** Announced public support since 2022 exceeds US\$0.4 trillion. Even so, no economy can economically host the entire chain; the capital and know-how thresholds at the frontier are prohibitive, and duplicated capacity risks future overbuild.
- **Export controls have changed the game's objective.** The 2022–2024 US measures, aligned in part by the Netherlands and Japan, treat advanced computing as a national-security good. This fragments the leading-edge market while leaving the far larger mature-node and materials trade broadly intact.

- **AI demand is the dominant swing factor.** The data-centre build-out has concentrated growth in advanced logic, high-bandwidth memory (HBM) and advanced packaging (notably 2.5D interposer packaging). These are precisely the segments with the least redundancy, tightening the coupling between commercial demand and strategic risk.
- **Mature nodes are the next contested arena.** Rapid capacity additions at 28nm and above raise the prospect of oversupply, price pressure and trade-defence disputes, alongside a persistent dependence of automotive and industrial buyers on legacy chips.
- **Talent, not only capital, is the constraint.** Fab construction can be subsidised faster than the specialised engineering workforce can be trained, making labour and tacit process knowledge a rate-limiter on reshoring ambitions.
- **Full decoupling is improbable by 2030; managed fragmentation is the base case.** The plausible outcome is a guarded frontier over a still-interdependent base — costlier, more redundant in parts, and more politically governed than the pre-2020 system.

1. Context and why it matters

Semiconductors are the foundational input to the digital economy and, increasingly, to the physical one. They sit inside data centres and smartphones, but also inside cars, medical devices, power grids, industrial controllers, and every category of modern weapon system. This ubiquity is why the sector's vulnerabilities are treated as systemic. A shock at a single node does not stay contained: the 2021–2022 automotive chip shortage, triggered by a mismatch between pandemic-era demand swings and mature-node capacity, idled vehicle assembly lines worldwide and cost the automotive industry an estimated hundreds of billions of dollars in lost output. Crucially, the chips in shortest supply were not the most advanced; they were commodity microcontrollers and power components made on older process technologies. The episode was an early, legible demonstration that supply security is a question of specific dependencies, not aggregate volume.

Three structural features make the sector unusually consequential for both markets and statecraft. First, **extreme specialisation**: the modern chip is produced through a globally distributed sequence — design, IP licensing, tool manufacture, wafer fabrication, assembly, test and packaging — in which each step is dominated by a handful of firms. Second, **compounding returns to scale and learning**: leading-edge fabrication now requires capital outlays on the order of US\$20bn or more per facility and decades of accumulated process knowledge, which entrenches incumbents and raises the barrier to entry to near-prohibitive levels. Third, **dual-use character**: the same advanced logic that trains commercial AI models also underpins military simulation, cryptanalysis and autonomous systems, which is why the frontier of the industry has become an object of national-security policy.

The convergence of these features with a deteriorating geopolitical environment — above all the US-China strategic rivalry and the specific exposure of Taiwan — has produced the current moment. What was once a story about consumer electronics pricing is now a story about industrial sovereignty, alliance coordination and the credibility of deterrence. This report holds both lenses simultaneously, because the sector cannot be understood through either alone.

2. Market structure and scale

The semiconductor market can be read along two axes: by **product category** and by **value-chain stage**. Both matter, because concentration risk is distributed unevenly across them.



The device market and the enabling stages that feed it are different businesses with different risk profiles — the strategic leverage sits disproportionately upstream. — IRI

By product, the industry divides into logic (processors and application-specific chips), memory (DRAM and NAND flash), microcomponents, analog, and a residual of discretes, optoelectronics and sensors. By stage, value accrues to design and electronic-design-automation (EDA) software and IP; to the fabrication equipment and materials that enable manufacturing; to the foundries and integrated device manufacturers (IDMs) that make the wafers; and to the assembly, test and packaging (ATP) firms that turn wafers into usable devices.

The table below presents our transparent order-of-magnitude estimates for 2024, anchored to published industry-association revenue totals and cross-checked against major corporate disclosures. Figures are rounded and should be read as **estimates with a plausible error band of roughly $\pm 15\%$** , not audited accounts. They are intended to convey proportion, not precision.

SEGMENT / STAGE	EST. 2024 REVENUE (US\$BN)	APPROX. SHARE	BASIS AND NOTE
Total device market	~630	100%	Anchored to WSTS-style aggregate for 2024; memory-led recovery from 2023
– Logic	~210	~33%	Processors, GPUs, ASICs; AI accelerators the fastest-growing sub-segment
– Memory (DRAM + NAND)	~165	~26%	Highly cyclical; 2024 rebound driven by HBM and price recovery
– Micro (MPU/MCU/DSP)	~95	~15%	Includes automotive/industrial microcontrollers
– Analog	~80	~13%	Fragmented; power management, signal chain
– Discrete, opto, sensors	~80	~13%	Includes power discretes, image sensors
Enabling stages (not additive to device total)			These sit upstream/downstream of the device market
– EDA software + IP licensing	~20	–	Concentrated: three EDA vendors, one dominant CPU-IP licensor
– Wafer-fab equipment (WFE)	~105	–	Five firms hold the large majority; EUV is single-source
– Materials (wafers, gases, photoresist)	~68	–	Japan-heavy in several sub-categories
– Foundry (pure-play)	~130	–	One firm holds ~60%+ and the bulk of leading edge
– ATP / OSAT	~40	–	Assembly/test/packaging; advanced packaging now a bottleneck

The product mix of the ~US\$630bn device market makes the concentration of value in logic and memory visible at a glance.

2024 device market by product category

SEGMENT	SHARE
Logic	33%
Memory (DRAM + NAND)	26%
Micro (MPU/MCU/DSP)	15%
Analog	13%
Discrete, opto, sensors	13%

Shares of the ~US\$630bn device market: Logic ~33%, Memory ~26%, Micro ~15%, Analog ~13%, Discrete/opto/sensors ~13%. Estimates with a plausible band of ±15% (§2).

Two observations follow directly from this structure.

First, **the device market and the enabling-stage markets are different businesses with different risk profiles**. The US\$630bn device figure is what most headlines cite, but the strategic chokepoints live disproportionately in the smaller enabling stages — a ~US\$20bn EDA/IP layer and a ~US\$105bn equipment layer whose concentration is far higher than the device market's. Control over these upstream stages confers leverage out of all proportion to their revenue.

Enabling stages by revenue, 2024

PERIOD	VALUE (US\$BN)
EDA + IP	20
ATP / OSAT	40
Materials	68
WFE equipment	105
Foundry (pure-play)	130

Enabling-stage estimates (±15%), not additive to the US\$630bn device total. The smallest layers — EDA/IP (US\$20bn) and WFE (~US\$105bn) — carry the highest supplier concentration (§2).

Second, **the growth trajectory is real but cyclical**. The widely cited path toward a US\$1 trillion market around 2030 implies a compound growth rate in the high single digits from the 2024 base. We regard this as a reasonable central estimate under continued AI and electrification demand, but we flag two caveats: memory revenue is notoriously volatile (it fell sharply in 2023 before rebounding), and a material share of projected growth is concentrated in AI-linked logic and memory, where demand visibility beyond two to three years is genuinely uncertain. A trillion-dollar market is plausible; it is not preordained, and the distribution of that value across segments is more uncertain than the aggregate.

3. The geography of concentration and the real chokepoints

Aggregate manufacturing capacity is more geographically distributed than the headlines suggest — mainland China, Taiwan, Korea, Japan, North America and Europe all host substantial fab capacity. The concentration that matters is not total wafers but **capability at specific points**. Five chokepoints define the system's fragility.

Leading-edge logic fabrication. The manufacture of the most advanced logic — currently the 3nm-class node, with 2nm-class in ramp — is dominated by a single Taiwanese foundry, with a second Korean firm and a US IDM as distant challengers. On our reading of public capacity data, the large majority of sub-10nm output is fabricated on the island of Taiwan. This is the single most consequential concentration in the global economy's supply base, and it coincides with the most acute geopolitical flashpoint. The colloquial "silicon shield" thesis — that this concentration deters conflict — is contested; whatever its deterrent value, it is unambiguously a source of systemic risk.

EUV lithography. Extreme-ultraviolet lithography, indispensable for the most advanced nodes, is supplied by a single Dutch company, itself dependent on a narrow base of sub-suppliers (including a German optics specialist). There is no substitute and no second source. This is the clearest example in the industry of a genuine, near-absolute single point of failure — and it is what makes multilateral export controls on lithography so potent an instrument.

Advanced packaging. As transistor scaling slows, performance gains increasingly come from packaging multiple dies together. The 2.5D interposer packaging used for AI accelerators has become a binding capacity constraint, again concentrated at the leading foundry. Advanced packaging illustrates how a chokepoint can migrate: a segment once treated as low-value "back-end" work is now a gating factor for the highest-value products in the market.

High-bandwidth memory. HBM, the stacked memory that feeds AI accelerators, is supplied by a small number of Korean and US firms, with one Korean firm holding a commanding lead in the most advanced generations. Because each AI accelerator requires HBM in volume, this segment couples memory-market cyclicality directly to the AI capital cycle.

Materials. Several materials sub-categories — notably high-purity photoresists and certain specialty chemicals and gases — are concentrated in Japan, while silicon-wafer supply is dominated by a few Japanese and Taiwanese producers. These dependencies are less visible than lithography but no less real; an interruption to photoresist supply would propagate through fabs worldwide.

The strategic point is that **redundancy is thinnest precisely where value and demand growth are highest.** The AI build-out concentrates on advanced logic, advanced packaging and HBM — three of the five chokepoints. Policy that focuses on aggregate "chip capacity" while neglecting these specific dependencies will misallocate resources.

4. Geoeconomic instruments: subsidy and denial

Since 2022 states have deployed two broad instruments in parallel: **positive industrial policy** (subsidy to build domestic capacity) and **negative economic statecraft** (export controls to deny rivals capability). Together they constitute the most significant reordering of the sector's political economy in a generation.

On the subsidy side, the headline programmes are large and broadly comparable in ambition, though structured differently. The table below summarises the principal announced commitments. Amounts reflect headline public figures and mix grants, loans, tax credits and mobilised investment; they are **not directly comparable line-for-line**, and we present them to convey scale and intent rather than as equivalents.

JURISDICTION	PRINCIPAL VEHICLE	ANNOUNCED PUBLIC SUPPORT (APPROX.)	STATED AMBITION
United States	CHIPS and Science Act (2022)	~US\$52.7bn direct + tax credits	Rebuild domestic leading-edge and mature manufacturing
European Union	EU Chips Act (2023)	~€43bn mobilised	~20% of global production value by 2030 (aspirational)
China	National IC "Big Fund" (three phases)	>US\$100bn cumulative, incl. ~US\$47.5bn phase III (2024)	Indigenous capability across the chain; import substitution
Japan	METI subsidy programmes	>US\$25bn committed	Attract leading-edge fabs; revive domestic frontier (2nm venture)
South Korea	Cluster and tax-incentive programmes	Very large; multi-year private-led with public support	Sustain memory leadership; expand foundry
India	Semiconductor Mission	~US\$10bn scheme	Establish assembly/test and initial fabrication

Announced public support by jurisdiction

PERIOD	VALUE (BN)
China (Big Fund)	100
United States	52.7
EU	43
Japan	25
India	10

Headline announced support. Mixes grants, loans, tax credits and mobilised investment; EU shown as €43bn, others in US\$bn; not directly comparable line-for-line (§4). South Korea omitted (no single headline figure). China and Japan values are floors ('>').

Several judgements follow. The programmes are **complementary in some respects and duplicative in others**. The EU's aspiration to reach roughly a fifth of global production by 2030 is, on any realistic reading of the capital and time required, unlikely to be met on the stated timeline; we treat it as a directional signal rather than a forecast. The US programme has already catalysed announced private investment well in excess of the public outlay, but the binding constraints — construction timelines, tool availability and, above all, specialised labour — are not fully soluble with money. China's effort is the most comprehensive in intent because it must build the entire chain, including domestic tooling, under external constraint; its progress at mature nodes has been faster than many expected, while the leading edge remains gated by the equipment it cannot legally import.

On the denial side, the pivotal development was the tightening of US export controls from October 2022, expanded in 2023 and further refined in 2024. These measures restrict the export to specific destinations of advanced AI-capable logic chips and, critically, of the advanced manufacturing equipment needed to make them. Their effectiveness depends on multilateral coordination, because the key tools are made by companies in a handful of allied jurisdictions; the partial alignment of the Netherlands and Japan on equipment controls was therefore the decisive enabling step. In response, the targeted state has accelerated indigenous substitution and deployed its own leverage — notably controls on gallium, germanium and certain other critical minerals in which it holds a dominant processing position.

Geoeconomic trajectory, 2021–2024

WHEN	MILESTONE	DETAIL
2021–22	Automotive chip shortage	Mature-node scarcity idles vehicle assembly lines worldwide; an estimated hundreds of billions of dollars in lost output.
2022	US CHIPS and Science Act	~US\$52.7bn direct funding plus tax credits to rebuild domestic leading-edge and mature manufacturing.
Oct 2022	US export controls tighten	Advanced computing reframed as a national-security good; restrictions extend to the equipment used to make advanced chips.
2023	EU Chips Act; controls expanded	~€43bn mobilised toward ~20% of global production by 2030 (aspirational); Netherlands and Japan partially align on tooling controls.
2024	China 'Big Fund' phase III; controls refined	~US\$47.5bn phase III committed for indigenous capability as US measures are further refined.

The essential analytical point is that **controls bite at the frontier and blur below it**. They are effective at slowing access to the most advanced compute, because those depend on chokepoint tools. They are far less effective — and largely not aimed — at the mature-node and materials trade that constitutes the bulk of the market by volume. The result is a market that is being cleaved at the top while remaining deeply interdependent throughout the middle and base.

5. Demand dynamics and the mature-node question

Two demand stories now dominate the sector's outlook, and they pull in different directions.

The first is **artificial intelligence**. The data-centre build-out since 2023 has produced a concentrated surge in demand for advanced accelerators, HBM and advanced packaging. This has been a windfall for the narrow set of firms positioned at those chokepoints and has driven the memory market's recovery. But it also introduces a new form of concentration risk: a large and growing share of the industry's projected growth, and an outsized share of its profit pool, now depends on the durability of AI-related capital expenditure by a small number of hyperscale buyers. Should that capital cycle slow — through model-efficiency gains, disappointing returns on AI deployment, or simple over-provisioning — the correction would fall hardest on precisely the high-value, low-redundancy segments. The AI boom is therefore simultaneously the sector's principal growth engine and a source of concentrated cyclical risk.

The second is the **mature node**. Automotive, industrial, consumer and infrastructure demand runs predominantly on process technologies of 28nm and above. Here the concern is the mirror image of the frontier: not scarcity but potential **oversupply**. Substantial capacity additions at mature nodes — a significant share of it in mainland China, partly as a rational response to being blocked at the frontier — raise the prospect of a coming glut, price erosion and trade-defence disputes. For import-dependent economies this creates a genuine policy dilemma: mature-node chips are strategically important (they are in every car and grid controller), yet building subsidised domestic capacity to make commodity chips risks competing against a wave of low-cost supply. Several jurisdictions have begun to treat legacy-chip dependence as a distinct security concern, separate from the leading-edge contest, and this is likely to become a more prominent axis of trade friction toward 2030.

Underlying both stories is a **workforce constraint** that money addresses only slowly. A leading-edge fab requires thousands of specialised engineers and technicians whose skills take years to develop and whose tacit process knowledge cannot be transferred by subsidy. Multiple reshoring programmes have already encountered labour and know-how bottlenecks that have pushed out timelines. Talent, not capital, may prove the true rate-limiter on the geographic diversification that policy seeks.

6. Risks and uncertainties, ranked

We group the principal risks into four categories and state our confidence in each.

Geopolitical shock (high impact, uncertain probability). A crisis affecting Taiwan — whether blockade, coercion or conflict — is the tail risk that dominates the sector's risk profile. Even a non-kinetic interruption to Taiwanese output or shipping would propagate through the global economy within weeks, given the concentration of leading-edge and advanced-packaging capacity. We make no probability claim here; we note only that the concentration makes the consequence severe irrespective of likelihood, which is itself a reason for redundancy investment.

Policy escalation and fragmentation (high probability, moderate-to-high impact). Further tightening of export controls, retaliatory mineral restrictions, and subsidy-driven overcapacity disputes are, in our assessment, more likely than not to continue. The direction of travel is toward more managed, more politically governed trade in advanced technology, with rising compliance costs and reduced market efficiency.

Cyclical correction (moderate probability, concentrated impact). The AI capital cycle introduces the risk of a sharp demand air-pocket in the highest-value segments. Memory in particular has a long history of boom-bust cycles; a coincidence of memory oversupply with an AI capex slowdown would be painful and is not improbable within the forecast horizon.

Execution and overbuild risk (moderate probability, diffuse impact). The current subsidy wave could produce duplicated capacity that proves uneconomic once demand normalises, particularly at mature nodes. The industry's history includes several episodes of subsidised overbuild followed by consolidation, and publicly funded capacity is more vulnerable to this than market-disciplined investment.

The interaction of these risks matters more than any one in isolation. The scenario that should most concern planners is a **compound event**: a geopolitical shock arriving during a cyclical downturn, when balance sheets are weakest and redundancy investment is being cut. Resilience planning that assumes risks arrive one at a time will understate the required buffers.

Three scenarios for supply security to 2030

We present three named scenarios. They are analytical constructs, not forecasts, and are intended to bound the plausible space rather than to predict a point outcome. Each is stated with its key assumptions.

Scenarios for supply security to 2030

A — Managed Fragmentation — Base case — central judgement

Subsidy and control regimes persist and extend incrementally with no kinetic shock. The frontier bifurcates into a guarded advanced-compute ecosystem while mature nodes, materials and packaging stay interdependent; costs rise structurally.

METRIC	VALUE
Aggregate market 2030	approaches ~US\$1tn
Taiwan leading-edge share	majority retained
Mature-node overcapacity	episodic trade disputes

B — Accelerated Decoupling — Downside for efficiency

Friction intensifies short of conflict; controls broaden to more nodes and materials and retaliation escalates. Two increasingly separate technology stacks emerge with duplicated capacity.

METRIC	VALUE
Redundancy	improves
Efficiency	substantial penalty
Innovation	slows at the margin

C — Supply Shock — Tail risk — low probability, severe

A sustained loss of Taiwanese output removes a large share of leading-edge and advanced-packaging capacity; acute shortages of advanced compute cascade through automotive and electronics.

METRIC	VALUE
Probability	low, non-trivial
Impact	global, severe
Recovery	years, not months

Scenario A — Managed Fragmentation (base case; our central judgement). Subsidy and export-control regimes persist and are incrementally extended, but no major kinetic shock occurs. Leading-edge capacity diversifies modestly — new fabs in the US, Japan and Europe reach volume production — while Taiwan retains the majority of the most advanced output through 2030. The market bifurcates at the frontier (a guarded advanced-compute ecosystem) while remaining interdependent for mature nodes, materials and packaging. Costs rise structurally; the aggregate market approaches but may not reach the trillion-dollar mark; mature-node overcapacity triggers episodic trade disputes. *Assumptions: continued but not accelerating geopolitical tension; sustained (if cyclical) AI demand; no Taiwan crisis.*

Scenario B — Accelerated Decoupling (downside for efficiency, upside for redundancy). Geopolitical friction intensifies short of conflict; controls broaden to more nodes and more materials; retaliation escalates. Two increasingly separate technology stacks emerge, with duplicated capacity and higher costs across the board. Redundancy improves but at a substantial efficiency

penalty; some smaller economies are forced to choose alignment. Innovation slows at the margin as the global talent and capital pool splits. *Assumptions: sharp deterioration in relations; aggressive mutual restriction; sustained political will to bear economic cost.*

Scenario C — Supply Shock (tail risk; low probability, severe impact). A sustained loss of Taiwanese output — from conflict, coercion or a major natural event — removes a large share of leading-edge and advanced-packaging capacity for an extended period. The consequences are global and severe: acute shortages of advanced compute, cascading effects through automotive and electronics, and an emergency reallocation of scarce capacity by governments. Recovery is measured in years, not months, given the impossibility of rapidly replicating the affected capability. *Assumptions: a low-probability but non-trivial geopolitical or natural shock; the current concentration persisting at the time of the shock.*

Across all three, one structural feature holds: **full self-sufficiency does not arrive by 2030 for any actor.** The question is not whether interdependence persists, but where it is guarded, how expensively it is duplicated, and how well the system's buffers are pre-positioned.

What this means for decision-makers

For governments and policymakers. Calibrate ambition to economics. Full-chain self-sufficiency is not an achievable or cost-effective goal; targeted resilience at the genuine chokepoints — advanced packaging, HBM, photoresists, tooling — delivers more security per dollar than generic "capacity" subsidy. Prioritise the mature-node dependence problem explicitly, because it is where consumer and industrial exposure is greatest and where oversupply risk and trade defence will collide. Treat allied coordination as the decisive variable: unilateral controls and unilateral subsidies are both markedly less effective than coordinated ones, and the industry's chokepoints reward multilateral action. Finally, invest in the constraint money cannot buy quickly — specialised engineering talent — through immigration, training and university partnerships, and measure programmes on capacity actually operating rather than announced.

For business and industrial users. Map exposure at the level of the specific component and process node, not the aggregate supplier. Single-source dependence at a chokepoint (a particular package, a particular memory generation, a particular materials input) is the binding risk, and it is frequently invisible several tiers up the supply chain. Build buffer inventory selectively for the mature-node components whose loss is cheap to hedge and expensive to suffer. Where feasible, qualify second sources in advance — a slow, costly process that must begin before, not during, a shortage. Treat design flexibility (the ability to substitute nodes or suppliers) as a resilience asset worth paying for.

For investors and financial institutions. Price a structurally higher capital-intensity and policy-risk environment. The subsidy wave changes the returns calculus in complex ways: it de-risks some capacity investment while raising the odds of eventual overbuild and margin pressure, particularly at mature nodes. Concentration is a double-edged exposure — the chokepoint incumbents enjoy formidable moats and pricing power, but also carry the sector's most acute geopolitical tail risk. Distinguish the durable structural growth in advanced compute from the cyclical, buyer-concentrated character of the current AI capex surge; the two warrant different discount rates. Policy is now a first-order variable in every semiconductor thesis, not a footnote.

Method, sources and confidence

This report is a synthesis, not a primary-data study. Our approach was to triangulate across four kinds of publicly available sources: (1) industry-association and market-research aggregate data on revenue and segment mix; (2) corporate disclosures from foundries, equipment makers, memory producers and materials suppliers, used to sanity-check segment shares and concentration estimates; (3) official policy documents and legislative texts for the subsidy programmes and export-control measures; and (4) the comparative academic and think-tank literature on semiconductor supply-chain resilience and economic statecraft.

We have been deliberate in distinguishing established facts from estimates. **Established** (reported by primary sources and widely corroborated): the approximate 2024 device-market total; the existence and headline values of the CHIPS Act, EU Chips Act and comparable programmes; the single-source status of EUV lithography; the timing and broad content of the 2022–2024 export controls; and the general fact of leading-edge concentration in Taiwan. **Estimates** (our synthesis, stated with error bands): all segment-level revenue splits in Section 2, which carry a plausible band of roughly $\pm 15\%$; the precise share of sub-10nm capacity located in Taiwan, which we state as a range because firm-level and facility-level capacity data are incomplete; and every forward-looking figure, including the trillion-dollar-by-2030 trajectory, which we treat as a plausible central estimate contingent on stated assumptions rather than a forecast. The subsidy figures in Section 4 mix grants, credits and mobilised investment and are **not** directly comparable line-for-line; they are presented to convey scale and intent.

The scenarios are analytical constructs designed to bound uncertainty, not probabilistic forecasts; where we express a judgement of likelihood we identify it as such. Readers should treat the quantitative content as decision-support of the "right order of magnitude" kind, and consult primary sources for any figure used in a transaction or a formal policy commitment. This report takes no position on the merits of any government's policy and is not funded by any party with a commercial interest in its conclusions.

Sources and references

Industry-association aggregates, legislative texts and control measures below were read directly; corporate capacity and revenue disclosures are used to cross-check segment shares and are referenced by category.

- World Semiconductor Trade Statistics (2024). *WSTS Semiconductor Market Forecast, Autumn 2024*. WSTS.
- Semiconductor Industry Association (2024). *2024 State of the U.S. Semiconductor Industry / Factbook*. SIA, Washington, DC.
- United States Congress (2022). *CHIPS and Science Act of 2022 (Public Law 117-167)*. Washington, DC.
- European Parliament and Council of the European Union (2023). *Regulation (EU) 2023/1781 establishing a framework of measures for strengthening Europe's semiconductor ecosystem (Chips Act)*. Official Journal of the European Union.
- US Bureau of Industry and Security (2022–2024). *Export controls on advanced computing and semiconductor manufacturing items (interim final rules)*. US Department of Commerce, Washington, DC.
- OECD (2023). *Measuring distortions in international markets: the semiconductor value chain*. OECD Trade Policy Papers, OECD Publishing, Paris.
- Ministry of Economy, Trade and Industry (2023). *Semiconductor and Digital Industry Strategy*. METI, Tokyo.
- Annual reports and capacity disclosures of leading foundry, memory, equipment and materials suppliers (2023–2025), used to cross-check segment shares and concentration estimates.

Suggested citation

Vasquez-Moreau, E., Malhotra, R., and Okonkwo-Bauer, N. (2026). *Semiconductors and Critical Supply Chains: Geoeconomics and Market Structure*. Flagship Report. International Research Institute. Licensed under CC BY 4.0.