Center for Technology & Statecraft

DUV Immersion Lithography: The Foundation of the US AI Advantage

Closing gaps in export controls to China

Nicholas Brown and Saif M. Khan

Executive summary **4**
DUV immersion tools are the strongest chokepoint for limiting Chinese AI 5
Current export controls let China stockpile DUV immersion capable of producing Huawei's AI chips 7
Banning DUV immersion to China would secure America's long-term AI advantage 9
1. Lithography is a key input to chipmaking **12**
2. DUVi is needed for competitive AI chips **15**
Modern AI chips are made using process nodes that require DUVi 15
Less advanced nodes cannot support competitive AI chips 18
3. DUVi is the top chokepoint for China's AI chipmaking **21**
Advanced lithography is the hardest SME to indigenize 21
China is highly unlikely to produce DUVi at a scale that materially expands its AI chipmaking until the mid-2030s 22
4. SME export controls have granted the US a strong lead in AI chipmaking **28**
5. Export controls let China import DUVi capable of mass-producing Huawei AI chips **33**
Problem #1: ASML's NXT:1980i is allowed into China despite enabling production of Ascend AI chips at massive scales 35
Problem #2: Less advanced DUV immersion tools are uncontrolled despite enabling earlier AI chips 37
Problem #3: Dutch servicing rules preserve ASML's ability to support China's installed DUV immersion fleet 38
6. China is building a massive stockpile of DUVi that enable AI chipmaking **39**
China's NXT:1980i imports risk advancing its AI chip production 40
7. A better alternative to exporting DUVi: dry DUV can support China's legacy chips **44**
8. The US government can control DUVi systems with executive branch authorities **47**
9. Banning DUVi exports would secure the US AI chipmaking advantage **52**
Appendix A: AI chips made without DUVi are uncompetitive **56**
Logic scaling 57
DRAM scaling 60
Matching B100 FLOPS and Bandwidth using 28nm logic and HBM1 62
Caveats 63
Appendix B: Technical challenges of developing advanced lithography **64**
Appendix C: DUV frontend logic yields **69**
Caveats 70
Methodology 70
Modeling Yline 71
Modeling Ynonpatterning, residual 72
Modeling Yrandom 73
Modeling Ypatterning 76
Appendix D: DUVi throughputs for Huawei AI chip production **80**
Methodology 81
Caveats 81

Wafer-pass throughput 82
Logic 83
HBM 84
Ascend 910C 86
Appendix E: Assessing China's DUVi installed base 87
China's aggregate installed base 90
Estimate from ASML's financial data 91
Estimate from China's customs data 93
Consolidated estimate 97
Installations at key fabs 98
CXMT 98
SMIC 99
YMTC 100
Multinational fabs 101
Other known installations 102
Allocation to advanced fabs 102
Data sources 104
Appendix F: Lithography-bounded Chinese AI chipmaking 108
US ecosystem production model 111
China ecosystem production capacity model 113
Existing Stockpile 118
Indigenous production 118
New Imports 118
AI Allocation 120
Caveats 122
Authors 123
Acknowledgements 123
Revision history 123
Endnotes 124

DUV Immersion Lithography: The Foundation of the US AI Advantage

Closing gaps in export controls to China

Executive summary

Advanced lithography systems are indispensable for producing competitive AI chips, but China is unlikely to commercialize these systems at scale until the mid-2030s. This gives the United States and its allies, the sole producers of these systems, extraordinary leverage over China's AI ambitions.

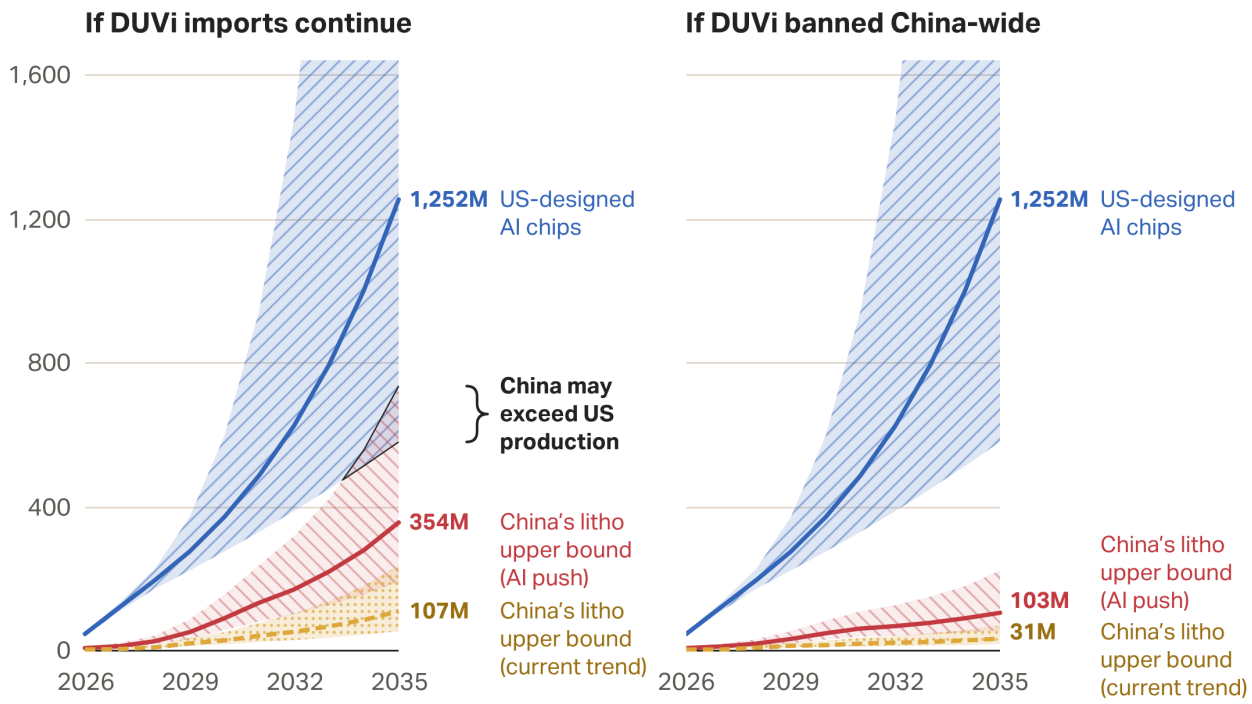
Yet today Washington is giving away its leverage by permitting China to import deep ultraviolet immersion (DUVi) lithography systems capable of producing advanced AI chips. China has aggressively capitalized on this opening, amassing a large, growing stockpile of these systems as of early 2026. If Chinese DUVi imports continue at their 2024–2025 pace or increase moderately over time, China could accumulate enough lithography capacity over the next five to ten years to narrow or even erase the currently sizable US-allied advantage in AI chipmaking.

The United States can prevent this outcome with a simple policy change: replace the current, partial DUVi controls with a China-wide ban on all DUVi exports, with complementary controls on spare parts and servicing. This action would secure the US advanced AI chip production advantage (Figure 1)—and resulting AI capability advantage—for at least a decade and ensure that the next wave of advanced fabs is built in the United States and its allied countries, not China.

Figure 1. Banning DUVi to China secures a US AI chipmaking advantage through 2035

China's existing DUVi fleet cannot match US AI chip production; but continued imports could let it close the gap

Bold lines show median projections (P50); lightly shaded cones span the P10–P90 range



US chip production is estimated for 2026–2027 from TrendForce shipment forecasts and from 2028 onwards by averaging a demand projection and an EUV supply model. H100-equivalents are dense FP8 FLOP/s (INT8 where a chip lacks FP8) divided by the H100's 1,979 TFLOP/s. China's forecasts are for lithography bounded production capacity, assuming it indigenizes all other SME. "Current trend" assumes DUVi are allocated between advanced and legacy fabs in the same proportion as today. "AI push" assumes China mobilizes its DUVi fleet to prioritize AI, obtaining the same allocation forecasted in the US ecosystem. "DUVi banned China-wide" assumes that DUVi imports stop beginning 2027. "DUVi imports continue" assumes a moderate increase in imports over time relative to China's 2024–2025 pace. See Appendix F for assumptions and calculations.

Center for Technology & Statecraft IFP

DUV immersion tools are the strongest chokepoint for limiting Chinese AI

Photolithography systems are the most important type of semiconductor manufacturing equipment (SME) for making advanced AI chips. Photolithography systems use ultraviolet light to print circuit patterns onto a silicon wafer, an essential step in the chipmaking process. The three most

recent generations are dry deep ultraviolet (dry DUV), used since the 1990s; DUVi, introduced in 2006; and EUV, commercialized in 2019. Each successive generation can print finer circuit features, enabling production of increasingly advanced chips. ASML, a Dutch company, is the sole supplier of EUV systems and holds 99 percent of the DUVi market. Nikon, a Japanese company, holds the remaining 1% share of the DUVi market.[1]

There is bipartisan support for limiting China's indigenous chipmaking via controls on SME. The United States and its allies have restricted exports of SME to China since 2018. These controls tightened in 2019, when the Netherlands blocked shipments of ASML's EUV systems, which are essential for producing the most advanced chips. Beginning in 2022, the US and its allies progressively expanded controls to cover DUVi systems and other critical types of SME. These controls have sharply constrained China's AI chip production, allowing the US to accumulate an estimated 50× advantage in the production of AI computing power as of 2026.[2]

Table 1: Lithography system capabilities and export restrictions[3]

<table>
<tr><th>Lithography technology</th><th>ASML tools</th><th>Relevant capabilities</th><th>Example AI chip</th><th>Current export restrictions</th><th>Proposed export restrictions</th></tr>
<tr><td>EUV</td><td>NXE:3100 and more advanced</td><td>Leading-edge logic (3nm and below); HBM4 and better</td><td>NVIDIA Rubin (e.g. R100)</td><td rowspan="2">Banned China-wide</td><td rowspan="4">Banned China-wide</td></tr>
<tr><td rowspan="3">DUVi</td><td>Most advanced: NXT:2000i–2150i</td><td>5nm logic; HBM4[4]</td><td>NVIDIA Hopper (e.g. H100)[5]</td></tr>
<tr><td>Advanced: NXT:1980i, 1970i</td><td>7nm logic; HBM2e/HBM3</td><td>NVIDIA Ampere (e.g. A100), Huawei Ascend 910C, 950DT</td><td>License required China-wide, but many licenses granted</td></tr>
<tr><td>Less advanced: XT:1700i–NXT:1965i</td><td>16/14nm logic; HBM2</td><td>Groq LPU, NVIDIA V100</td><td rowspan="2">No license required</td></tr>
<tr><td>Dry DUV</td><td>NXT:1470 and less advanced</td><td>Legacy production nodes</td><td>Google TPUv1</td><td>No license required</td></tr>
</table>

[1] The dry DUV market contains two types of tools: dry argon-fluoride (ArF) and krypton-fluoride (KrF). ASML and Nikon respectively have 94% and 6% market share in dry ArF.[1] ASML, Nikon, and the Japanese company Canon respectively have 66%, 3%, and 31% market share in KrF.[2]

[2] Data from TrendForce, generally corroborated by 2025 IFP projections.[3]

[3] This table is simplified for brevity. See Section 2 for a comprehensive discussion.

[4] Micron implemented HBM4 using the 1β DRAM node, without EUV.

[5] NVIDIA's H100 was made using TSMC's 4N process, a variant of 5nm, which used EUV. However, it is likely that 5nm can be implemented with NXT:2000i and more advanced systems.

Advanced photolithography (DUVi and EUV) systems are the most effective long-term chokepoint for limiting China's AI chip production, for three reasons:

1. **Advanced lithography is a prerequisite for manufacturing competitive AI chips.** A chip's capability depends on its "node"—a label (e.g., "7nm logic") identifying a generation of chipmaking technologies used together in the manufacturing process. Competitive AI chips can only be made using advanced nodes and these nodes require DUVi or EUV. Chips made using nodes achievable with less advanced lithography (e.g., dry DUV)—28nm and above—are not competitive for AI.
2. **China cannot currently make lithography tools capable of producing AI chips, and it is highly unlikely to do so at the scales needed to materially impact AI chipmaking before the mid-2030s.** Its photolithography system manufacturers remain decades behind ASML and are unlikely to commercialize a DUVi capable of supporting the 7nm logic and advanced HBM used in Huawei's AI chips before the early-to-mid 2030s.[6] Additional years would be required to scale up production enough to materially expand AI chip output. China is even further behind in EUV, which it is unlikely to commercialize before the second half of the 2030s.
3. **Advanced lithography is the hardest SME for China to develop.** Chinese firms are rapidly closing the gap in other SME categories, such as etch and deposition, while making far slower progress in lithography. Advanced lithography will likely be the last chipmaking tool that China indigenizes.

Current export controls let China stockpile DUV immersion capable of producing Huawei's AI chips

US, Dutch, and Japanese export controls allow China to acquire and maintain DUVi capable of producing Huawei's AI chips. Controls ban EUV and ASML's most advanced DUVi systems but leave three major openings for China to exploit:

1. **ASML's advanced DUVi systems, the NXT:1980i and NXT:1970i, are permitted to enter China despite enabling Huawei's AI chips.** The Netherlands requires a license for exporting these systems to any location in China, but public data shows that a large number of shipments are allowed. NXT:1980i supports making the pieces of silicon (called "dies") (7nm logic and a

[6] This estimate is consistent with July 2026 reporting that Shanghai Aishengna has begun small-batch production of a DUVi system capable of 28nm logic and roughly equivalent to ASML's NXT:1950i.[4] A prototype of this system reportedly entered testing in 2025, roughly 14 years after 28nm entered volume production outside of China, consistent with China's historical 13–19 year lag in lithography.[5] Extrapolating a similar 14-year lag beyond TSMC's 2018 introduction of 7nm points to a Chinese 7nm-capable prototype around 2032. ASML's seven-year progression from the NXT:1950i to the 7nm-capable NXT:1980i also implies a 7nm-capable prototype around 2032. Allowing 1–2 years for qualification points to limited commercial use in 2033–2034. See Section 3 for a full discussion, including factors that could speed or slow China's progress.

form of advanced high-bandwidth memory called HBM2e) used in Huawei's AI chips at high throughput.

2. **ASML's NXT:1965i and earlier DUVi are effectively uncontrolled.** These systems can be exported anywhere in China without a license[7] despite supporting 16/14nm logic and the less-advanced HBM2—nodes sufficient for earlier-generation AI chips.
3. **Dutch servicing rules allow ASML to maintain China's installed base of DUVi,** including systems located at fabs making AI chips.[8] Continued servicing and replacement components are necessary to keep these tools online in the long term.

Chinese fabs have aggressively capitalized on this opening, amassing a stockpile of over 330 DUVi systems by the first quarter of 2026, most of which are NXT:1980i systems (Figure 2). These systems can support China's AI chip production through two pathways:

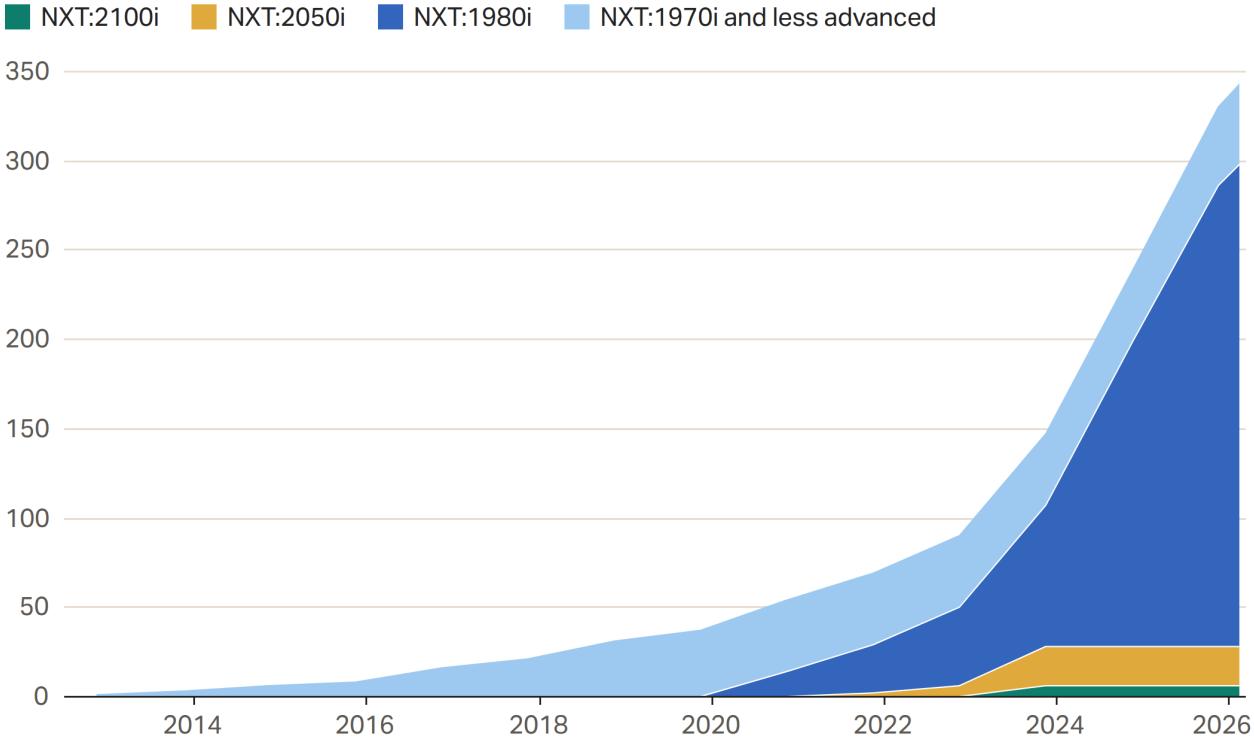
1. **China may be able to directly obtain some NXT:1980i systems at fabs producing or capable of producing AI chips, rather than only at fabs capable of producing legacy chips**. Available evidence suggests that CXMT, China's leading producer of high-bandwidth memory (HBM), a critical input for AI chips, has continued to acquire substantial numbers of DUVi systems in recent years. Hua Hong Semiconductor may likewise have stockpiled these tools for years as it prepared to launch a 7nm advanced logic production line.
2. **Systems initially shipped to legacy fabs may later be repurposed to support AI chip production**. China could divert entire systems or critical subsystems to advanced fabs ("physical diversion"), upgrade the original facility's production node ("diversion in place"), or cycle partially processed wafers between fabs ("wafer diversion").

[7] The United States controls export of these systems (along with all items subject to its jurisdiction) to advanced-node and Entity Listed Chinese fabs. The Netherlands does not, however, and thus ASML exports are effectively uncontrolled.

[8] Dutch servicing rules control transfers of technology and software required to operate, maintain, or repair NXT:1970i and more advanced DUVi systems. However, an exception permits the minimum technical support needed to keep authorized exports operating. Because most controlled DUVi in China were exported under Dutch licenses, ASML can continue providing this support to most of China's installed fleet.

Figure 2. China is building a vast stockpile of DUVi systems

Estimated cumulative DUVi units imported by Chinese-owned fabs since 2012



Estimates exclude imports to multinational fabs located within China. See Appendix E for assumptions and calculations.

Center for Technology & Statecraft IFP

Continued DUVi imports could eventually allow China to erase the US chipmaking advantage. China's current DUVi stockpile is too small to match US chip production. However, if China were to maintain or moderately increase its 2024–2025 pace of NXT:1980i imports over the coming 5–10 years and fully mobilize the resulting fleet of DUVi for AI chip production, it could produce hundreds of millions of H100-equivalents[9] each year—enough to approach or surpass US and allied output under some scenarios (Figure 1).

Banning DUV immersion to China would secure America's long-term AI advantage

The US has all of the authorities needed to close the DUVi loophole. The following actions are necessary:[10]

[9] An H100-equivalent (H100-eq) is a unit of aggregate AI compute: a quantity of chips whose combined dense 8-bit FLOP/s (INT8 where a chip lacks FP8) equals that of one NVIDIA H100 (1,979 TFLOP/s).

1. **Ban all[11] DUVi exports to China, including DUVi produced by ASML and Nikon outside of the United States**. If Japan or the Netherlands is unwilling to do this using their own authorities, the US can do so unilaterally using the foreign direct product rule (FDPR),[12] *de minimis* authority,[13] or by controlling critical US-produced DUVi subsystems when the exporter has knowledge that the "end use" is incorporation into a DUVi system known to be destined to China.[14]
2. **Ban servicing of all DUVi at advanced-node Chinese fabs by foreign persons.** The United States already prevents US toolmakers from performing any form of servicing at advanced fabs. However, the Netherlands permits its citizens to provide the minimum level of servicing needed to maintain China's DUVi fleet. The US government can unilaterally ban servicing of DUVi located in an advanced-node fab using the Export Administration Regulations' (EAR) General Prohibition 10.[15]
3. **Prevent diversion of China's existing DUVi stockpile to AI chip production**. ASML may have the technical capability to neutralize diversion.[16] The US government should work with the Netherlands and ASML to assess this capability and ensure it is robust to countermeasures.
4. **Restrict subcomponents, precursors, and knowledge transfers that enable China's efforts to indigenize advanced DUVi or keep installed DUVi systems online**. Priorities include Zeiss projection optics, rumored to be circulating within China, together with the materials and equipment (e.g., fused silica) needed to make comparable lenses; high-power ArF excimer light sources and their subcomponents and consumables; and precision wafer and reticle stages and their motion sensors. To mitigate impacts on legacy production, licenses could be granted for one-for-one replacements at verified legacy facilities.

These actions would have the following strategic effects:

---

[10] As discussed in Box 1 in Section 8, the MATCH Act, a bill currently under consideration in Congress, calls for restrictions on DUVi exports to and servicing in China.

[11] Extending to ASML's original XT:1700i.

[12] The existing SME FDPR allows the US government to control items produced abroad that contain integrated circuits that are the direct product of US-origin software or technology.[6] All existing DUVi systems meet this criterion.

[13] The *de minimis* authority allows the US government to control items that incorporate more than a specified percentage of US content. The US government can control DUVi extraterritorially by setting a 0% de minimis threshold because all existing DUVi systems incorporate some amount of US content.

[14] All existing DUVi systems incorporate a light source—an irreplaceable technical subsystem used to generate deep ultraviolet light—from either a US company (Cymer) or a Japanese company (Gigaphoton).[7] The US and Japanese governments can jointly prevent DUVi exports to China by controlling domestically produced light sources when the end use is incorporation into a DUVi system known to be destined to China.

[15] General Prohibition 10 prohibits servicing any item when there is "knowledge" that an export control violation has occurred or will occur "in connection with the item."[8] The US government has already declared that fabs known to be producing AI chips are in violation of export controls.[9] It can restrict servicing of any DUVi system at such a fab by informing the public that these DUVi systems are being used in connection with a violation of export controls.

[16] ASML has allegedly assured the Dutch government that it can disable TSMC's EUV systems if China were to invade Taiwan.[10]

1. **A China-wide ban on DUVi exports would ensure the United States an overwhelming AI chipmaking advantage through at least 2035** (Figure 1). China's existing DUVi fleet could not match US-allied output even if every system were redirected to AI chip production, but continued imports could erase this advantage. A China-wide ban would close this pathway and preserve the US advantage at least until China can commercialize domestic 7nm DUVi lithography—unlikely before the mid-2030s.
2. **Tighter servicing controls would slow China's near-term AI chipmaking ramp** by degrading the performance of DUVi installed at China's AI chipmaking fabs. A ban on new exports has a smaller immediate impact[17] but may also degrade near-term production by preventing replacement DUVi systems and subsystems from reaching China's advanced fabs.
3. **New advanced chipmaking capacity would shift away from China and towards the US and its allies.** Demand for advanced chips would not disappear if Chinese fabs can no longer expand; it would move to producers elsewhere. Allied memory manufacturers Micron, Kioxia, Samsung, and SK hynix would face less subsidized competition from China's CXMT and YMTC. Logic foundries TSMC, Intel, and Samsung would absorb demand otherwise filled by SMIC. US and allied chip designers would also benefit. Existing controls already prevent them from manufacturing using Chinese advanced fabs, while Huawei faces no such limitation. Additional constraints on Chinese advanced capacity would therefore disproportionately affect Huawei.
4. **Effects on China's legacy production can be avoided.** China could continue importing dry DUV systems, which are sufficient for all of China's legacy logic and memory nodes. Some DUVi systems are installed in legacy fabs, but policymakers can choose to apply servicing restrictions only to fabs producing advanced chips, leaving legacy fabs untouched.[18]
5. **ASML would bear the primary cost among US and allied toolmakers, but the impact would be modest and mostly recoverable.** ASML's backlog of non-China orders is large enough to keep production lines occupied for months after a cutoff. In the longer term, new fabs built outside of China would create additional demand for ASML equipment. ASML's former CEO, Peter Wennink, stated in 2022 that even if China were "excluded from any growth," new fabs would be built elsewhere and ASML's 2030 outlook would not change much.[11] We estimate that the principal near-term loss is approximately $1–1.5 billion in annual Chinese servicing revenue, equivalent to 3–4 percent of ASML's 2025 sales. Some of this near-term servicing loss could ultimately be recovered if tighter servicing controls degrade China's DUVi fleet, shifting production—and equipment demand—to fabs elsewhere.

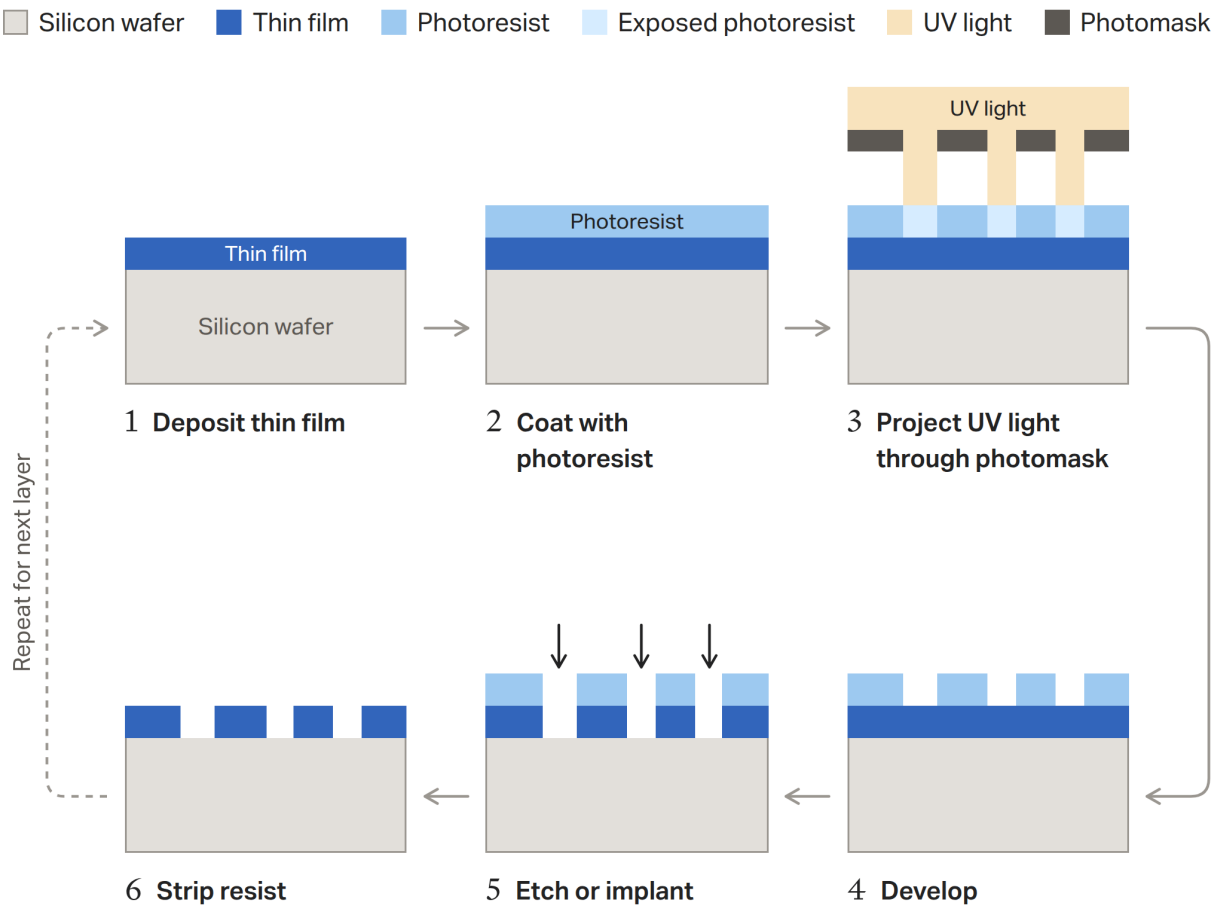
[17] China probably has enough DUVi systems installed at its advanced fabs to support its ramp in the near term. The effects of an export ban would be felt in the mid to long term.

[18] In response, China could attempt to divert untouched DUVi systems in legacy fabs toward advanced production. This would require additional servicing restrictions once such diversion is discovered.

1. Lithography is a key input to chipmaking

Lithography is a critical step in semiconductor manufacturing, in which ultraviolet light is used to print circuit patterns onto a silicon wafer. Modern chips contain many layers of circuitry, built up through repeated manufacturing cycles. In a typical cycle (Figure 3), the wafer is coated with a light-sensitive material called a photoresist; ultraviolet light is projected through a patterned mask onto the resist (the lithography step); the exposed pattern is chemically developed[19]; and that pattern is then transferred, or etched, into the underlying material. The lithography step is performed by a dedicated machine known as a lithography system, which projects ultraviolet light through the mask onto the wafer in ultra-high resolution.

Figure 3. A simplified patterning cycle



Center for Technology & Statecraft IFP

[19] Exposure to ultraviolet light alters the resist's solubility; developing applies a chemical solution that selectively dissolves either the exposed or unexposed regions, leaving the pattern physically formed in the remaining resist.

Lithography technology has advanced through several major paradigms over the past 60 years, each enabling resolution of orders-of-magnitude smaller features. Early systems used mercury vapor lamps to print features below half a micron.[12] In the mid-1990s, the industry introduced systems using deep ultraviolet (DUV) light generated by krypton-fluoride (KrF) excimer lasers.[13] The three subsequent lithography paradigms support virtually all modern advanced chip manufacturing:

- Argon fluoride (ArF) systems, referred to in this report as dry DUV systems (though the label also applies to KrF systems), were introduced in the early 2000s and use 193nm light generated by argon-fluoride excimer lasers to print features down to ~57nm.[14]
- Argon fluoride immersion (ArFi) systems, also known as DUV immersion (DUVi) systems, project ArF light through a thin layer of highly purified water,[20] which enables printing features down to ~38nm.[15]
- Extreme ultraviolet (EUV) systems entered volume production in 2019, following decades of research and development, and use 13.5nm light to print features that would require prohibitively complex patterning with DUVi.[16] These systems support today's most advanced chips.

ASML, a Dutch company founded in 1984, overwhelmingly dominates the supply of advanced lithography systems. ASML is the sole supplier of EUV systems and accounts for approximately 99% of the DUVi market.[17] Its main competitor in DUV lithography—and only remaining commercially viable competitor in DUVi—is Japan's Nikon, while Canon, also of Japan, continues to supply older systems (e.g., i-line lithography tools).[21] (China has reportedly begun small-batch production of domestic DUVi systems, as discussed in Section 3, but none has been shown to be production-worthy for advanced nodes.) Nearly all of ASML's current DUV and EUV systems use its TWINSCAN architecture, a dual-stage design in which one wafer is measured and aligned while another is exposed, improving throughput and overlay. ASML markets two major TWINSCAN DUV platforms: the XT series, used for KrF, dry ArF, and early ArFi (i.e. DUVi), and the NXT series, used for later ArFi and which provides higher overlay precision and throughput. ASML's EUV systems are divided into two series as well: NXE systems use conventional numerical aperture EUV, while the newer EXE series uses high numerical-aperture (high-NA) optics to further improve resolution[22] and is entering volume production in 2026.[19] Several notable ASML systems are enumerated in Table 2:

---

[20] Using ultrapure water extends the capability of ArF systems because water has a higher refractive index (1.44) than air at 193nm, which increases the scanner's numerical aperture—an optical parameter measuring the width of the cone of light that can be focused onto the wafer. A higher numerical aperture enables the scanner to focus a wider cone of light, enabling resolution of finer features.

[21] Several other lithography equipment companies, including GCA, Perkin-Elmer, Ultratech, and Silicon Valley Group (SVG), were active through the 1980s and 1990s but have since exited the market or been acquired.

[22] High-NA systems use advanced optics to increase numerical aperture (NA) from 0.33 to 0.55, enabling resolution of smaller features. Intel has begun using high-NA systems in high-volume manufacturing for select layers of its "Panther Lake" Core Ultra Series 3 processors.[18]

Table 2: ASML systems across lithography paradigms

ASML system	Lithography technology	Production introduction	Notes
PAS 5000	KrF	1991	First ASML KrF system
NXT:870B		2024	Most advanced KrF system
PAS 5500	ArF	1998	First ASML ArF system
NXT:1470		2020	Most advanced ArF system
XT:1700i	ArFi (DUVi)	2006	First volume production ASML ArFi system
NXT:1980i		2015	Most advanced ArFi system widely available in China under current licensing system
NXT:2150i		2024	Most advanced ArFi system
NXE:3400B	EUV	2017	First production EUV system
NXE:3800E		2024	Most advanced EUV system
EXE:5200B	High-NA EUV	2025	First production high-NA EUV system

2. DUVi is needed for competitive AI chips

Modern AI chips incorporate two kinds of silicon: (1) logic dies, used to perform computations, and (2) dynamic random access memory (DRAM) dies, used to store data. Both dies use advanced "nodes"—i.e., successive generations of semiconductor manufacturing technology, labeled in nanometers (e.g., "7nm" logic). Competitive AI chips use logic nodes at 16/14nm and below and DRAM nodes at 20nm or below. These nodes cannot realistically be implemented without DUVi or more advanced (EUV) photolithography, because older (dry DUV) lithography systems lack the resolution and overlay performance needed to reliably pattern their small, tightly spaced features. Falling back to older process nodes sharply reduces chip performance and power efficiency, making such chips uncompetitive for frontier-scale AI. Access to DUVi systems is therefore a prerequisite for competitive AI chipmaking.

Modern AI chips are made using process nodes that require DUVi

Modern AI chips incorporate two kinds of silicon. The first is one or more logic dies, processors specialized for performing computations used for AI. The second is high-bandwidth memory (HBM), a specialized form of vertically stacked DRAM that stores model parameters and other data and is designed to move that data to and from the logic dies at high speeds.

Both types of silicon are manufactured using advanced nodes. A process "node" is a label[23] (e.g., 7nm logic, 1x DRAM) identifying a generation of semiconductor manufacturing technologies. Each successive node combines advances in lithography and complementary manufacturing technologies —including deposition, etch, and metrology—that together enable smaller and more densely packed circuit features, thus improving chip performance and power efficiency. For logic chips, advanced nodes increase transistor density,[24] allowing designers to pack more computational units, wider data paths, and larger caches into each unit of silicon area, while also improving the performance and power-efficiency of individual transistors. This enables processors with better computational performance and power efficiency. For DRAM, advanced nodes increase bit density,[25] enabling more

[23] For logic chips, node names originally tracked the transistor gate length—e.g., the "350nm" node enabled printing chips with a gate length of 350nm. From the mid-1990s onward, however, gate lengths began shrinking faster than node names, and today they have no direct relationship to a measurable transistor dimension. DRAM node names historically corresponded to half of the pitch of the metal bit lines in the memory array. However, once half-pitch entered the sub-20nm range, manufacturers adopted a letter-based scheme—1x, 1y, 1z, then continuing into the Greek alphabet with 1α (1-alpha), 1β (1-beta), 1γ (1-gamma), etc.—to designate successive shrinks within "10nm class" DRAM.

[24] Per Moore's Law, transistor density has increased roughly 150-fold in recent decades: Twenty years ago, the leading-edge 65nm logic node offered a transistor density of approximately 2 million transistors per square millimeter; today, TSMC's N2 (2nm class node) achieves 300 million transistors per square millimeter.[20]

[25] Bit density improved 80-fold over the past two decades but has slowed in recent years, increasing only ~2× over the past decade.[21] A central reason for this slowdown is the difficulty of shrinking the DRAM storage capacitor, which must maintain sufficient capacitance to hold a readable charge even as its lateral dimensions shrink, requiring increasingly tall, high-aspect-ratio structures that are increasingly more difficult to manufacture.[22]

memory to fit on each chip, reducing cost per bit, and improving power efficiency per bit. This enables HBM with better capacity, bandwidth, and power efficiency.

NVIDIA's major AI chips over the past decade reflect this pattern. All have combined logic processors manufactured using 16/14nm or more advanced nodes with HBM manufactured using 20nm or more advanced DRAM (Table 3):

Table 3: NVIDIA's AI chips from the past 10 years have used ≤16/14nm logic and ≤20nm DRAM

Chip	Year	Logic node	HBM generation	DRAM node
Volta (e.g. V100)	2017	TSMC 12FFN (12nm)[23]	HBM2[24]	~20nm[26]
Ampere (e.g. A100)	2020	TSMC N7 (7nm)[26]	HBM2 and HBM2e[27]	~20nm – 1y[28]
Hopper (e.g. H100, H200)	2022	TSMC 4N (custom 5nm)[29]	HBM2e, HBM3, and HBM3e[30]	1y, 1z, and 1β[31]
Blackwell (e.g. B100, B200, B300)	2024	TSMC 4NP (custom 5nm)[32]	HBM3e[33]	1β[34]

Which process nodes a fab can implement depends in part on the capabilities of its available lithography systems. Three capabilities of a lithography system are especially important in determining which process nodes it can implement in a semiconductor fab setting: minimum resolvable feature size (MRF), overlay, and throughput. "Minimum resolvable feature size" (MRF) determines the smallest feature that can be printed in a single exposure—i.e., a dose of light that transfers a pattern onto the wafer. MRF is typically expressed in terms of "half pitch"—i.e., half the distance from the center of one line to the center of the next in a repeating line-and-space pattern, or roughly the width of one line. "Overlay" measures the alignment error between a newly printed pattern and the underlying patterns on the wafer.[29] "Throughput" measures the rate at which the system can expose wafers. A lithography system is viable for a given process node only if it can meet the node's patterning requirements at acceptable yield and cost. It must print the node's smallest features, either in a single exposure or through multi-patterning—i.e., techniques that split a single layer's pattern across multiple exposures or other patterning steps; align successive patterns precisely enough to avoid defects and degraded electrical performance; and do so fast enough to sustain economical high-volume production.

[26] Samsung's 2016 HBM2 used a 20nm process.[25]

[27] A100 40GB and 80GB respectively use HBM2 and HBM2e.[27]

[28] Samsung's HBM2e uses 1y (~16/14nm) DRAM.[28]

[29] Two measures of overlay are relevant: dedicated chuck overlay (DCO), the alignment between layers patterned using a single scanner and wafer chuck, and matched machine overlay (MMO), the alignment between layers patterned by different scanners. Because MMO includes machine-to-machine variation, it is always larger (worse) than DCO.

The process nodes used for modern AI chips require DUVi or more advanced lithography. For logic, all major foundries adopted DUVi at least five years and two nodes before introducing the 16/14nm node.[30] Compared with dry DUV and earlier lithography, DUVi provides two advantages that make advanced nodes viable:

- **Better single-exposure patterning capability.** DUVi systems can resolve patterns with roughly 38nm half-pitch in a single exposure, compared with roughly 57nm for dry DUV. Aggressive multipatterning—particularly self-aligned quadruple patterning—could in principle be used to allow dry DUV to pattern simple features (e.g., dense arrays of straight, regularly spaced lines) at advanced nodes. However, actual logic designs require irregular features (e.g., line ends, cuts, contacts, corners) that cannot be easily split across exposures. Patterning advanced-node chips using dry DUV's resolution would therefore impose extremely restrictive design rules, reducing effective transistor and interconnect density and potentially precluding the layouts needed for high-performance logic. Even if such a design could be developed (no attempt has been made), the multipatterning would greatly increase process complexity, worsen unit economics, and reduce yield because of increasing opportunities for alignment errors.[31]
- **Better alignment between layers.** All existing dry DUV systems, including ASML's most advanced NXT:1470, have far worse overlay performance than DUVi systems. Poorer overlay leaves less margin for aligning patterns between layers or within a single layer (assuming multipatterning is used), increasing alignment errors and lowering yields. With a focused effort, ASML might be able to improve the overlay of dry systems, but even then the disadvantages due to worse resolution would remain.

See Appendix C for further analysis of dry DUV's insufficiency at advanced logic nodes. For DRAM, industry transitioned to DUVi around 2008–2010, several years and multiple shrinks before 20nm-class DRAM entered high-volume manufacturing.[32] Although DRAM's regular arrays are relatively amenable to multipatterning, a complete DRAM process must also pattern irregular features in the periphery. There is no demonstrated dry DUV process for manufacturing 10nm-class DRAM (including the nodes

---

[30] TSMC adopted DUVi in 2007 for its 45nm node and implemented its 16/14nm node in 2015, eight years and three nodes later.[35] Intel adopted DUVi in 2009 for its 32nm node and implemented its 14nm node in 2014, five years and two nodes later.[36] Samsung adopted DUVi in 2008 for its 45nm node and implemented its 14nm node in 2015, seven years and three nodes later.[37]

[31] For example, self-aligned quadruple patterning (SAQP) would be required to print features with 21nm half-pitch using 57nm resolution—but this adds process steps, worsening unit economics, and requires very tight scanner overlay. For the 16/14nm node, multi-patterning reduces the overlay budget for critical layers to just a few nanometers, at or beyond the NXT:1470's ~4nm matched-machine overlay.[38] Even if a dry DUV tool can resolve these features via SAQP, the resulting overlay errors would almost certainly prevent the process from achieving tolerable yields.

[32] Samsung developed its 56nm DRAM node using DUVi in 2007 and implemented its 20nm DRAM node in 2014, seven years and four nodes later.[39] SK hynix adopted DUVi at its 44nm DRAM node in 2009 and implemented its 20nm DRAM node in 2015, six years and four nodes later.[40] Micron adopted DUVi at its 50nm DRAM node by 2009 and implemented its 20nm DRAM node in 2015–2016, six to seven years and four nodes later.[41]

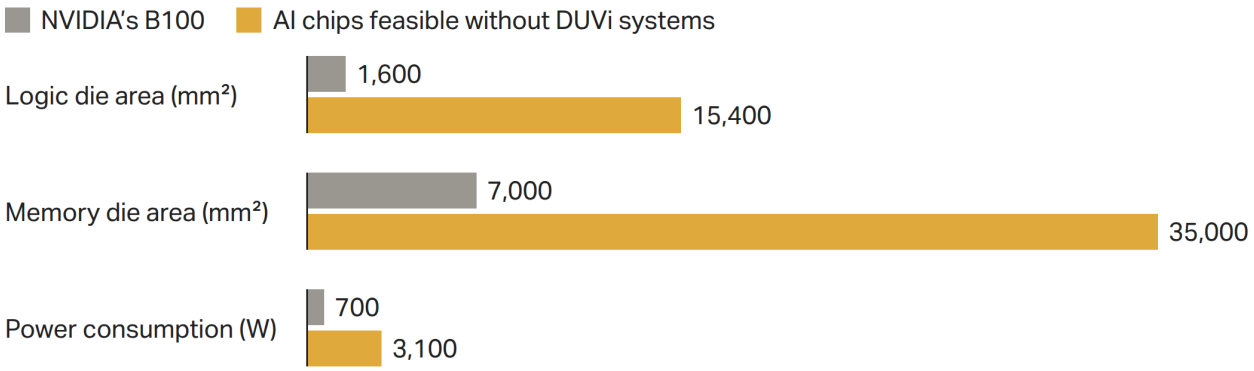
used for HBM2e and more advanced HBM) and the prospect of one being developed (in China or elsewhere) appears extremely unlikely because of these technical and economic constraints.

Less advanced nodes cannot support competitive AI chips

AI chips made on nodes that can be implemented using earlier lithography technologies (e.g., dry DUV) are not competitive with chips made using advanced nodes. Matching the processing performance and memory bandwidth of NVIDIA's B100 AI chip using nodes that can be implemented without DUVi (28nm logic and HBM1) would require ~9.6× as much logic silicon area, 5× as much DRAM silicon area, and ~4.4× as much chip power (Figure 4). See Appendix A for assumptions and calculations.

Figure 4. EUV or DUVi is a hard requirement for making competitive AI chips

Matching B100's FLOP/s and MemBW without DUVi systems consumes far more silicon and power



We estimate requirements to match B100's FP8 FLOP/s and memory bandwidth using 28nm logic and 29nm DRAM/HBM1, process nodes viable without DUVi systems. We scale logic area and power consumption based on transistor densities and iso-performance power consumption across nodes. Memory area and power consumption are derived from HBM1 bandwidth, capacity, bit density, and energy-efficiency characteristics. See Appendix A for assumptions and calculations.

Center for Technology & Statecraft IFP

These technical limitations substantially impede frontier AI training and inference. First, less advanced nodes require far more chips to deliver the same processing power and memory bandwidth. Accommodating these additional chips requires a much larger cluster and datacenter, increasing capex not only for chips but also for packages, boards, networking equipment, power-delivery and cooling equipment, cabling, and other supporting hardware. Second, chips made at less advanced nodes are much less power-efficient, requiring more electricity to enable the same processing throughput. This raises operating costs, increases the amount of power that must be obtained from the

grid or from on-site generation, and necessitates more on-site power distribution and cooling infrastructure. Third, distributing workloads across many more chips substantially reduces FLOP/s utilization (i.e., the share of theoretical processing capacity that is productively used for training or inference), because more time is spent moving data and coordinating work across devices. This lengthens training runs, reduces inference throughput and user-facing responsiveness, and further increases operating costs. Finally, the complexity of managing the cluster increases[33] at least linearly, and potentially superlinearly,[34] with its size, requiring far more engineering investment.

A back-of-the-envelope comparison illustrates the impact of these disadvantages (Figure 5). At full buildout, OpenAI and Oracle's Stargate Abilene supercluster is expected to contain >450,000 NVIDIA GB200 GPUs and draw ~1.2 GW of power across ~4 million square feet of datacenter space.[45] Replicating a cluster of this scale using nodes achievable without DUVi would require ~4.3 million packaged chips, assuming our 9.6× logic silicon scaling factor translates proportionally into package count. It would also require correspondingly more CPUs, boards, networking switches, cabling, cooling systems, power delivery systems, and other supporting hardware. For power, a 4.4× penalty implies ~2.4 GW consumed by logic chips and co-packaged memory. Assuming 100% overhead for CPUs, networking, cooling, and other infrastructure, total datacenter power demand would reach ~4.8 GW—roughly the output of five large nuclear reactors. Assuming datacenter and campus area scales proportionally with chip count, the cluster would require ~38 million square feet of datacenter space. These estimates are illustrative, given that industry has never contemplated building a cluster of this scale using legacy nodes, and an actual design would likely make different engineering tradeoffs.

[33] Large AI training clusters are fragile systems in which single component failures can halt the entire job. When Meta trained Llama 3.1 405B on 16,384 H100 GPUs, it experienced 419 unexpected interruptions over 54 days—roughly one every three hours—with approximately 78% attributed to hardware issues.[42] Meta's research estimates that at 100,000 GPUs, failures occur approximately every 18 minutes and synchronous recovery from each failure can take up to 10 minutes—meaning that under conventional approaches, over half of nominal training time would be consumed by failure recovery.[43] Every chip removed from a cluster reduces the combinatorial surface area for these failures and simplifies the networking, cooling, power delivery, and software orchestration required to keep the system running.

[34] For example, Meta found that cluster reliability is inversely proportional to its size.[44]

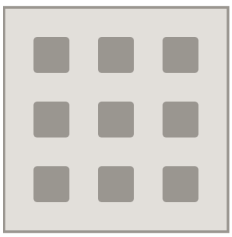
Figure 5. EUV or DUVi is a hard requirement for competing in frontier AI

Replicating OpenAI's Stargate Abilene supercluster without DUVi systems would require far more chips, power, and datacenter space

= 50,000 NVIDIA GB200 (TSMC 4NP, HBM3e)

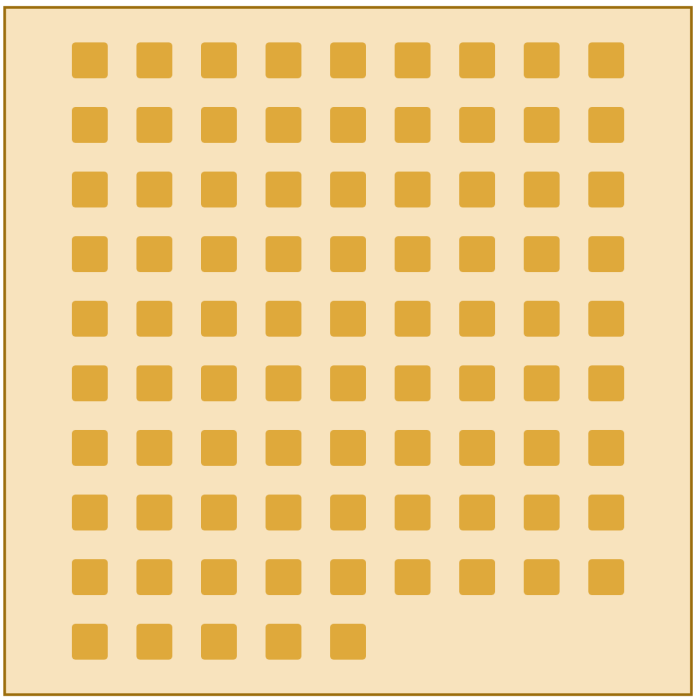
= 50,000 chips made without DUVi systems (TSMC 28HP, HBM1)

Stargate Abilene



450,000 packaged chips
~1.2 GW power capacity
~4M sq ft datacenter space

Cluster X, made without DUVi systems



~4.3M packaged chips
~4.8 GW power capacity
~38M sq ft datacenter space

Stargate Abilene is expected to contain >450,000 NVIDIA GB200 GPUs, with ~1.2 GW of power capacity, ~4 million square feet of datacenter space. Applying a ~4.4× power efficiency penalty and ~9.6× logic silicon area penalty, capturing the disadvantages of legacy nodes (28nm logic, 29nm DRAM, HBM1) compared to the GB200 nodes (TSMC 4NP, 1β DRAM, HBM3e), we estimate a cluster of ~4.3M packages, with ~4.8 GW power capacity, and requiring ~38 million square feet. Estimates are illustrative rather than precise. Campus footprints are drawn to scale by datacenter floor space. Each small square represents 50,000 chips.

Center for Technology & Statecraft IFP

The relative disadvantage of 28nm nodes will only get worse over time. NVIDIA's Rubin AI chips, releasing in the second half of 2026, use TSMC's more advanced 3nm node, and leading chipmakers are now commercializing production of next-generation 2nm nodes.

3. DUVi is the top chokepoint for China's AI chipmaking

DUVi and more advanced lithography systems are the hardest SME to develop. No other category of SME requires comparable R&D investment, development time, or supply chain complexity. As a result, China remains far behind the ASML-led ecosystem and is highly unlikely to produce DUVi systems at a scale that materially expands its AI chip production until the mid-2030s.

Advanced lithography is the hardest SME to indigenize

Advanced lithography is exceptionally difficult to develop, for two interrelated reasons. First, a lithography maker must develop dozens of highly specialized subsystems to an extremely high standard of precision and reliability. Second, these subsystems must be seamlessly integrated to function as a single, high-performance machine, which requires overcoming an additional set of steep engineering challenges. We review these challenges in Appendix B.

The extreme difficulty of developing advanced lithography is reflected in several aspects of the industry's history and present-day structure. **First, new lithography technologies have historically taken decades to mature.** ASML's TWINSCAN platform—the foundation of modern DUVi systems—has been under continuous development for about 25 years, while EUV took over 30 years between initial conception and commercial adoption.[46] As one ASML engineer informed a competitor's CEO, even possessing full technical blueprints would not be enough to reproduce ASML's systems, because they embody "decades, if not centuries" of accumulated tacit knowledge.[47]

Second, even established lithography companies with decades of optical engineering expertise have struggled to execute technological transitions. The shift from dry to immersion DUV effectively pushed Canon out of the market; it never achieved commercially viable production.[48] Nikon was able to commercialize immersion systems but lost substantial market share. Both Nikon and Canon abandoned their EUV programs entirely.[49] The result is an extraordinarily concentrated market: ASML accounts for 99% of DUVi sales, with Nikon accounting for the remaining sliver.[50]

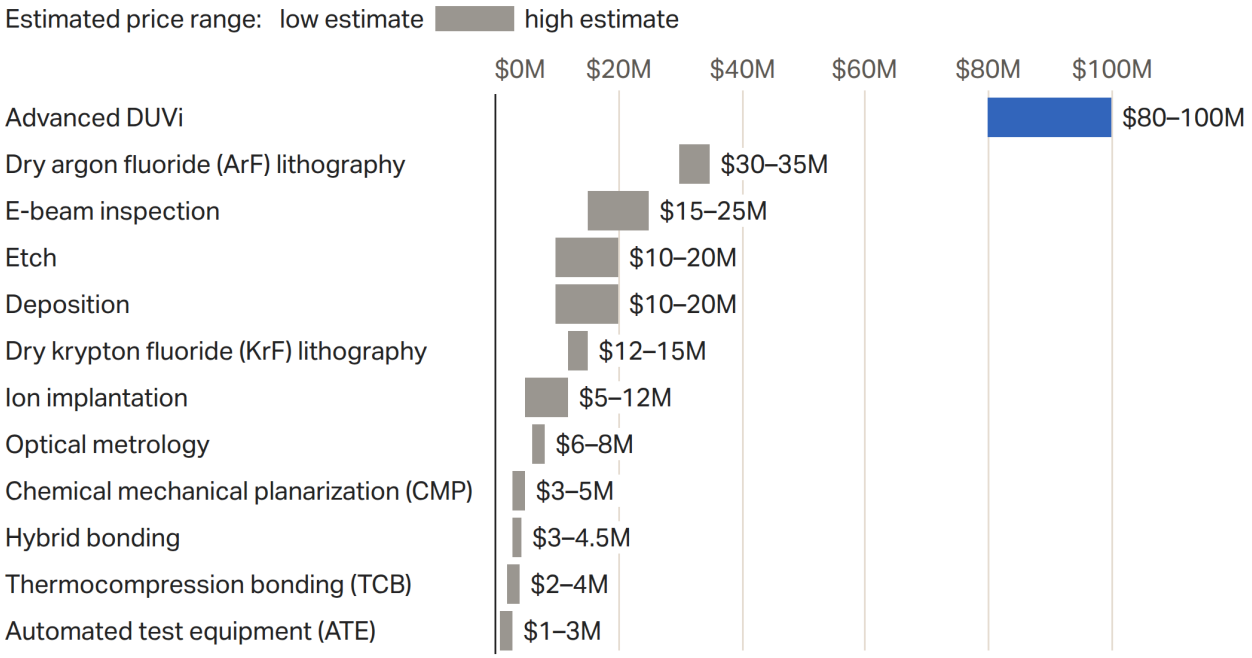
Third, remaining at the technological frontier requires larger R&D investments than in any other SME category. ASML's annual R&D expenditure on lithography exceeds $5.3 billion—more than the total company-wide R&D budgets of diversified equipment makers like Applied Materials (~$3.6 billion) or Lam Research (~$2.1 billion), which are spread across many equipment categories.[35][52] ASML has invested tens of billions of dollars in R&D since its founding in 1984. EUV alone required ~$14–21 billion of total industry investment across ASML, national laboratories, suppliers, and other partners.[53]

[35] Lithography accounts for 97% of ASML's business, with metrology and inspection accounting for the rest, so a small percentage of ASML's R&D spending includes metrology and inspection R&D.[51]

Fourth, lithography tools are far more valuable than—and therefore likely more complex than—other SME. A state-of-the-art DUVi system is worth $80–100 million, far more than the most advanced systems in other categories (Figure 6).[54] By contrast, a typical etch or deposition chamber costs just $1–3 million.[55]

Figure 6. DUVi systems are 3–50× more valuable than other advanced SME

Estimated purchase price of the most advanced systems in each SME category



EUV systems, not shown, are the only semiconductor manufacturing equipment more valuable than DUVi systems, at roughly $200 million each.

Center for Technology & Statecraft IFP

China is highly unlikely to produce DUVi at a scale that materially expands its AI chipmaking until the mid-2030s

Our central estimate is that China will develop commercially viable DUVi capable of 7nm production in the mid-2030s, though commercialization in the early 2030s is plausible. However, commercial viability is only the first milestone. Even if 7nm DUVi were to be commercialized in the early 2030s, scaling up production enough to materially increase AI chip output beyond the levels supported by China's existing DUVi stockpile (Section 6) would require several additional years. China's indigenous DUVi will likely not achieve that milestone before the mid-2030s.

Extrapolations based on ASML's product development timelines and China's historical lag in lithography capabilities both imply commercialization of 7nm-capable lithography in the mid-2030s. China's lithography prototypes have historically reached user-testing milestones 13–19 years after roughly comparable process nodes entered volume production outside China. For example, SMEE's SSB600/10 i-line scanner, which targets 280nm production, was shipped to a customer for testing in 2016,[36] roughly 19 years after 250nm entered high-volume manufacturing.[57] SMEE's SSA600/20 passed expert on-site testing in October 2017,[37] roughly 13 years after TSMC began 90nm volume production.[59] Most recently, the Financial Times reported in September 2025 that SMIC had begun testing a 28nm DUVi system from Shanghai Yuliangsheng, 14 years after TSMC's introduction of 28nm; follow-on reporting in 2026 indicated that the tool had entered low-volume production.[38][61] Extrapolating this 14-year lag beyond TSMC's introduction of 7nm in 2018 implies the arrival of a prototype for 7nm production in 2032.[62]

ASML's product-development history supports a similar timeline. Nearly seven years elapsed between ASML's first shipment of NXT:1950i in early 2009 and first shipment of 7nm-capable NXT:1980i in late 2015.[63] Extrapolating this development interval beyond China's 2025 prototype, which was described as comparable to NXT:1950i, also points to 2032 for a 7nm-capable prototype.[64]

Productizing a 7nm-capable prototype to support high-volume manufacturing (HVM) would likely require at least another 1–2 years of qualification. For example, ASML's NXT:2000i took two years between initial announcement and HVM ramp,[39] while NXT:1980i[40] took about 13–15 months and NXT:1950i at least a year and a half.[41] Together, these factors put our central estimate for the first commercially viable Chinese 7nm-capable DUVi in the mid-2030s (Figure 7). This conclusion also aligns with the historical base rates estimated by the AI Futures Project, which indicate 2032–2038 for commercial-scale 7nm-capable Chinese DUV.[42][68]

---

[36] "June 2016: The company's first—and China's first—front-end scanning lithography machine was delivered to a user/customer."[56]

[37] "In 2017, SMEE's… '90nm lithography machine prototype development' project passed expert on-site testing."[58]

[38] The latest reporting claims that Shanghai Aishengna plans to build approximately five in 2026 and 20 in 2027, for customers including SMIC, Hua Hong, and CXMT. The planned five-system buildout strongly suggests that the program remains in the customer-qualification phase, with volume production contingent on the results of this validation. There is at least some reason for skepticism that Aishengna will achieve volume production by next year. Reports of Chinese 28nm-capable tools have surfaced twice previously, in 2021 and 2023, without subsequent follow-up.[60]

[39] ASML announced early access to NXT:2000i in July 2017 and reported ramping NXT:2000i to HVM in 2019.[65]

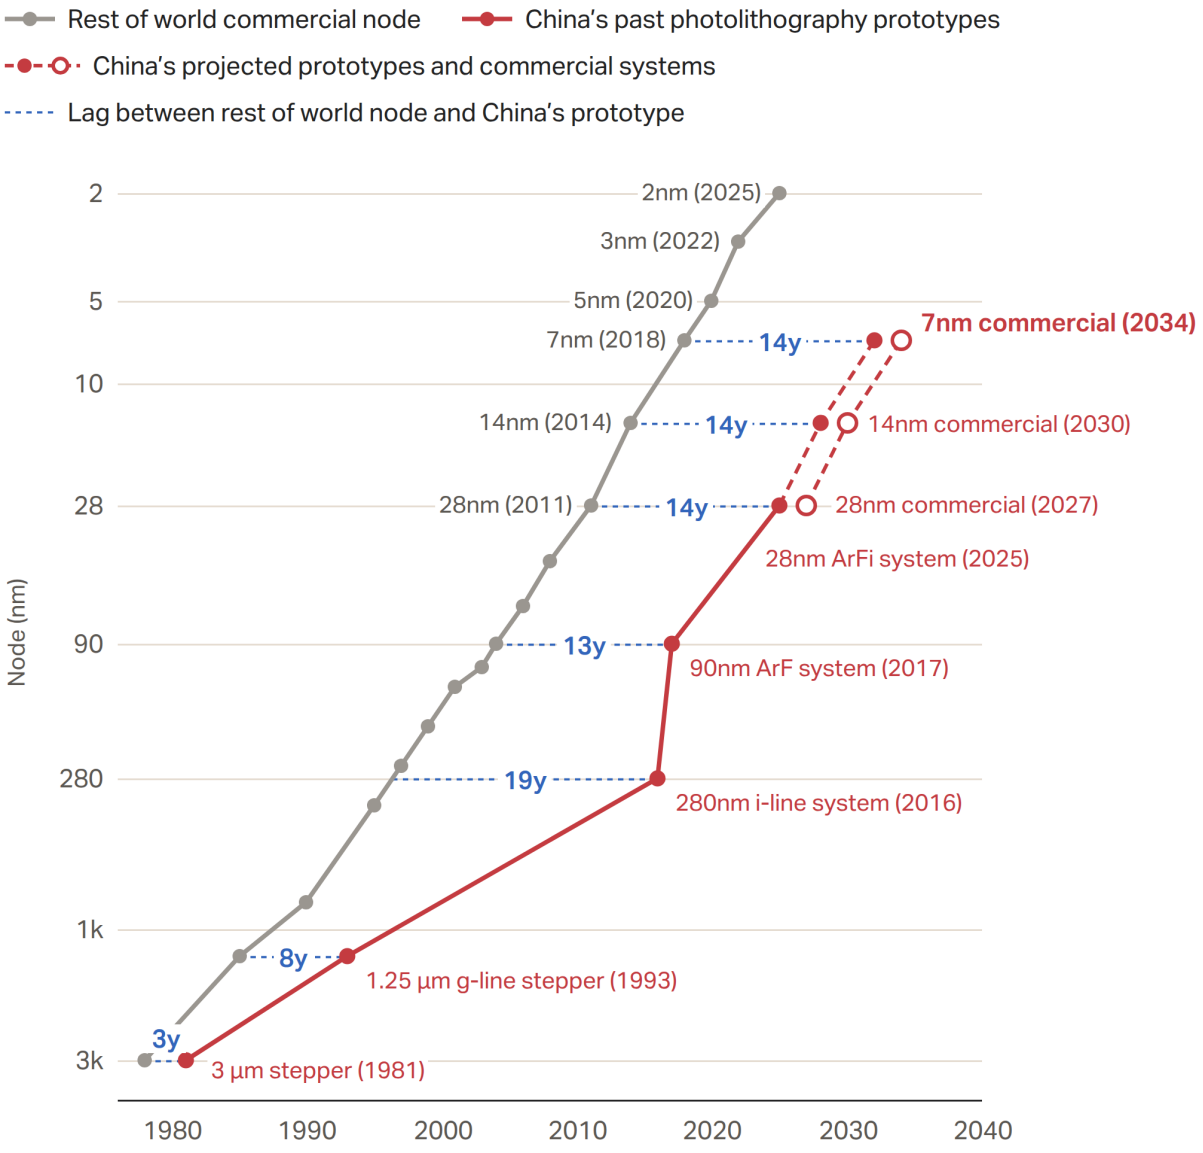
[40] ASML published results of prototyping the package that became 1980i in March 2015 and reported ramping 1980i to HVM for 10nm in 2016Q2.[66]

[41] ASML announced NXT:1950i in December 2008 and reported 1950i achieving volume production in 2010Q2.[67]

[42] This forecast extrapolates the 27–30 years that global industry took to progress from early DUV-development milestones to commercial 7nm production beyond equivalent points in China's DUV indigenization effort.

Figure 7. China's historical progress implies commercial 7nm lithography in the mid-2030s

China's lithography prototypes have lagged rest-of-world nodes by 13–19 years; an additional ~2 years is required to achieve commercial viability



Center for Technology & Statecraft IFP

Several factors might allow China to progress faster than ASML did in the past. Perhaps most importantly, China is not starting from a blank slate. Chinese firms can learn from published research and patents, study the architecture of existing ASML tools, poach former ASML engineers, and draw on China's huge installed base of DUVi systems. For example, Reuters reported in 2025 that China had recruited former ASML engineers for a secret EUV project and was reverse engineering subsystems

from DUV and EUV tools, using a dismantled system.[69] Although physical access cannot transfer all of ASML's accumulated manufacturing and systems-integration knowledge, it can reveal component designs, interfaces, tolerances, and architectural choices, reducing the space of unknowns that must be investigated.

China is also committing large amounts of capital towards acquiring advanced lithography capabilities. The Big Fund III, which began operations at the end of 2024, has $47 billion in registered capital, with equipment makers among its principal targets.[43][72] In parallel, Huawei is building or supporting multiple fabs and investing across the lithography supply chain.[73] Its involvement could help coordinate component suppliers, lithography subsystem developers, and chip manufacturers, accelerating the development of an integrated supply chain.

There is also an outside chance that China can bypass parts of ASML's optical-development pathway. Researchers are exploring alternative paths towards advanced lithography, including free-electron lasers for EUV and entirely different approaches like X-ray lithography.[74] The most significant Chinese effort is a Tsinghua University program that uses steady-state microbunching (SSMB) to generate EUV light. That said, these alternative approaches remain at an early experimental stage and have not been demonstrated even in a prototype.

On the other hand, there are also reasons to expect that China's progress will be slower than ASML's. First, China's lithography makers have an extremely poor track record of commercializing their prototype systems, particularly due to gaps in accumulated implicit process knowledge.[75] China's lithography makers have zero market share in either of the two DUV technology generations (dry ArF and dry KrF) preceding DUVi systems, with no increase from 2019 to 2025.[76] For example, Shanghai Micro Electronics Equipment (SMEE) released tests of a prototype dry ArF 90nm tool (SSA600/10) in 2011; roughly 13 years later, neither that tool nor its successor, the SSA600/20, appeared to have achieved meaningful domestic sales or production-readiness.[77] Even in the ultra-legacy i-line segment, China achieved only a minimal (4%) market share by 2025 and remains unable to satisfy its own domestic demand.[78] For example, Nikon relaunched its legacy NSR-2205iL1 i-line stepper in 2024 specifically to meet unmet Chinese demand.[79] External analysts and members of China's industry consistently express deep skepticism about the production-readiness of domestic systems. In 2023, a Chinese semiconductor executive stated: "Homegrown lithography was examined and verified by academics, not industrial engineers.[80] This equipment is only theoretically usable, and no chip manufacturer has ever dared to activate such a machine in their fabs." This matches views expressed by external analysts, for example: "Nothing has ever come of these reports in terms of equipment in mass production fabs.[81] Instead, SMEE has created demonstration equipment in its laboratory." Commenting on social media in 2024, a semiconductor industry veteran expressed the view that "even if we mobilize the whole nation and everything goes well, it is only in six years—around 2030—that we

[43] However, there is reason to believe that government-led, advanced lithography-related investments have not always been managed well in China. For example, prior Big Fund executives have been investigated for and charged with corruption in the past.[70] Local government-managed funds have also invested in semiconductor projects that have not come to fruition.[71]

might have a chance to produce 28nm chips using fully domestic [immersion lithography] systems." Commenting in 2026, a Chinese analyst stated: "Lithography is the biggest foreign dependency; the domestic share isn't even 5%.[82] SMEE's 28nm DUV is still under development; the lithography link has no near-term solution."

Second, China's supplier ecosystem is less mature than ASML's, both in absolute terms and also relative to the capabilities of its DUVi prototypes. For example, Yuliangsheng's 28nm prototype tool reportedly uses a legacy ArF excimer laser from Keyi Hongyuan.[83] The laser's technical specs (6 kHz / 60W) are comparable to XLR 500i, which Cymer—which supplies light sources for and is now a division of ASML—introduced for the less advanced 45nm node in 2007.[84] For photoresists, Chinese KrF resist suppliers meet only 5 percent of domestic demand and the country remains almost entirely dependent on Japanese suppliers for ArF resists (including ArF immersion and even earlier-generation products for dry ArF).[85] For projection optics, China's key suppliers (Changchun UP Optotech, Nanjing Wavelength, Beijing GuoWang Optics) themselves depend heavily on foreign suppliers for critical subcomponents.[86]

Third, China's funding model may struggle to cultivate the highly specialized supply chains and coordination mechanisms that have enabled the success of ASML's ecosystem. For years, China funded SME development, including lithography, through the 02 Special Project, a state megaproject launched in 2008. This state-backed R&D model succeeded at fostering research breakthroughs and working prototypes but was poorly suited to translating them into commercially viable systems—particularly for lithography, which requires highly specialized supply chains, developed over time through market evolution, and extremely tight iterative feedback loops between integrators, subsystem suppliers, and fabs.

Huawei's recent involvement may partly correct this weakness. But Huawei remains a new entrant in the space, with less experience orchestrating the SME supply chain compared to ASML. The underlying difficulty is that indigenizing commercially viable DUVi is neither the kind of commodity manufacturing at which China's subsidy-based funding model excels, nor a pure research challenge that can be solved through laboratory investment alone. In the words of one Chinese commentator, the fundamental problem in developing a "new-type national system [for directing national resources towards the development of strategic technologies] is coordinating it with market forces."[87] Where China has achieved relatively greater success in its SME industry, such as with AMEC's and Naura's development of etch and deposition tools, it did so by developing a competitive, commercial industry—unlike SMEE's unsuccessful lithography efforts, which relied on state-backed R&D.

Finally, even if China commercializes a 7nm-capable system in the early 2030s, scaling production and adoption to a level that materially expands overall chipmaking capacity would take several additional years. Initial production is likely to be small: for example, Shanghai Aishengna reportedly plans five systems in 2026 and 20 in 2027. Rapid growth thereafter would be difficult to sustain. A CSET review estimated 60% as the highest historical rate of growth across the lithography industry, achieved

first at Nikon and later at ASML.[88] China already has more than 300 imported DUVi systems in-country (Section 6). Although only a subset of these support AI production, an initial batch of domestic scanners would add little to that capacity. To materially expand output, China would need to produce high tens of commercially viable systems annually, which would take several additional years. An impact at that scale before the mid-2030s is unlikely.

4. SME export controls have granted the US a strong lead in AI chipmaking

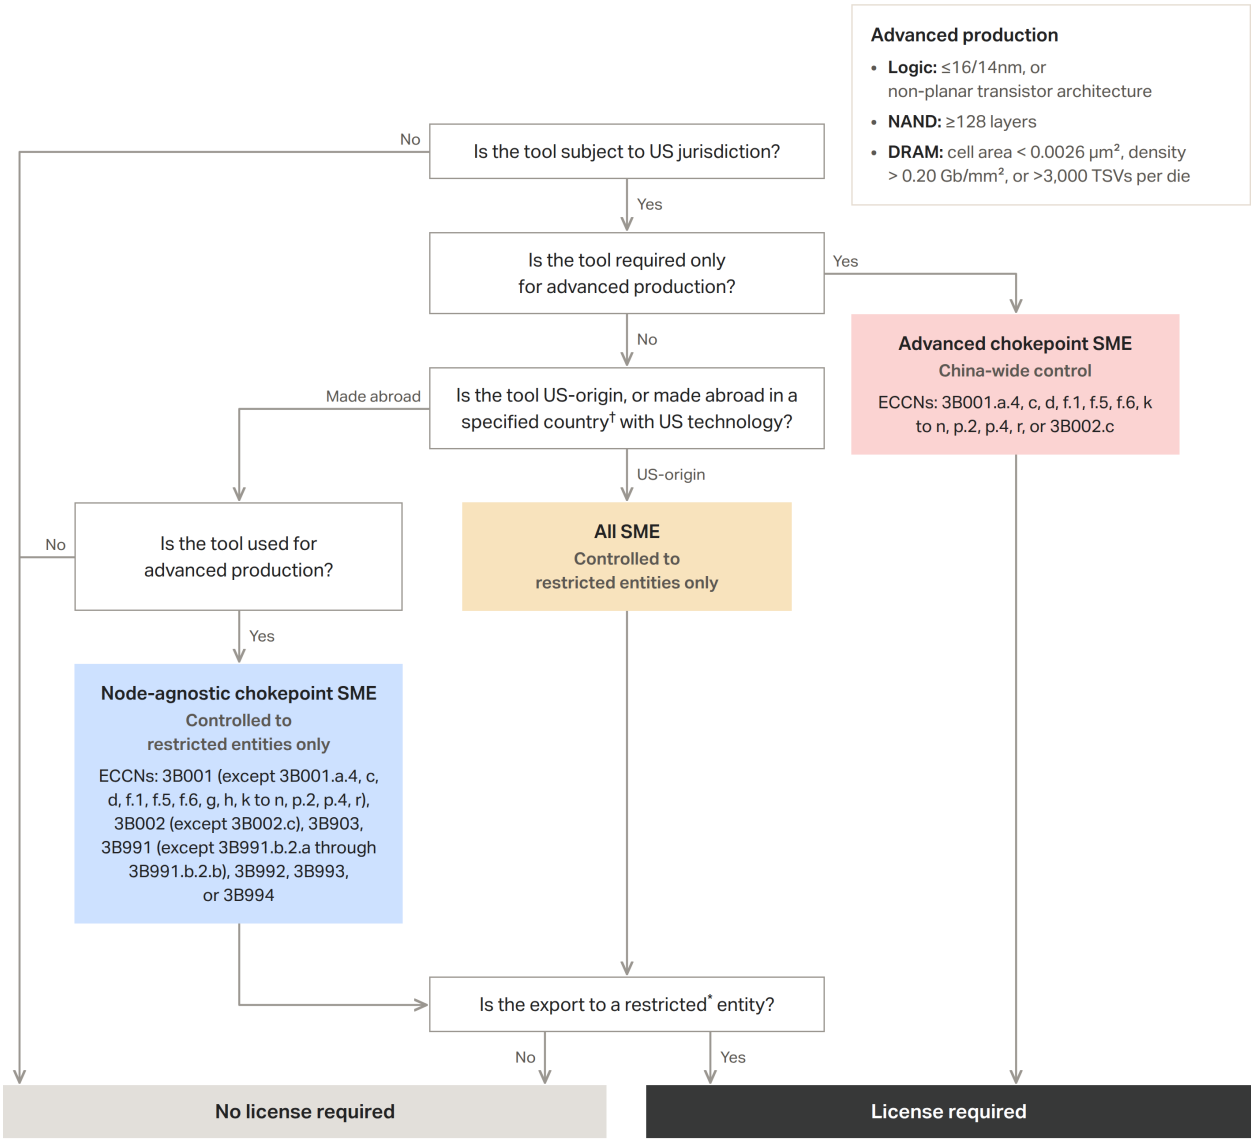
The United States and its allies have restricted SME exports to China since 2018, when the US government put China's then-memory chipmaking champion Fujian Jinhua on its export blacklist, the Entity List, effectively cutting it off from US-origin SME. These controls tightened in 2019, when the Netherlands first withheld authorization for ASML to ship extreme ultraviolet lithography (EUV) systems to China, and again in December 2020, when the United States added China's leading logic foundry, SMIC, to the Entity List. The Biden administration then expanded these measures through a series of rules issued in October 2022, October 2023, December 2024, and January 2025, implemented alongside parallel controls adopted by key allies Japan and the Netherlands.

The current SME control framework is designed to restrain China's production of advanced semiconductors without unduly restricting its ability to produce legacy semiconductors.[44] The US government defines three categories of "advanced-node" chip production: (1) logic chips using a non-planar transistor architecture or produced at or below the 16/14nm node, (2) NAND flash with 128 layers or more, and (3) DRAM meeting specified thresholds for cell area, memory density, or through-silicon-via density.[45] Using these thresholds, the United States implements two types of controls on SME, described below and visualized in Figure 8.

[44] Former US National Security Advisor Jake Sullivan described this framework as "carefully tailored restrictions on the most advanced semiconductor technology exports … not a technology blockade." Official BIS announcements feature similar language, emphasizing the specific intent to control advanced semiconductors. From the October 2022 rule: "The export controls announced in the two rules today restrict China's ability to obtain advanced computing chips, develop and maintain supercomputers, and manufacture advanced semiconductors."[89] From the December 2024 rule: "New controls on semiconductor manufacturing equipment needed to produce advanced-node integrated circuits, including certain etch, deposition, lithography, ion implantation, annealing, metrology and inspection, and cleaning tools."[90]

[45] These are: (i) A memory cell area of less than 0.0026 μm^2; (ii) A memory density greater than 0.20 gigabits per square millimeter; or (iii) More than 3,000 through-silicon vias per die.

Figure 8. A simplified framework for understanding US SME export controls



* Restricted entities: BIS Entity List, facilities where "production" of "advanced-node ICs" occurs.
† Specified countries: all except the 33 allied countries listed in Supplement No. 4 to 15 CFR Part 742. Notably, Japan and the Netherlands are exempted.
Note: This figure simplifies some details for the sake of explaining the controls. Additional license requirements and exceptions apply.

Center for Technology & Statecraft IFP

First, the US imposes China-wide controls on the most advanced tools,[46] including EUV, the most advanced DUVi systems, and advanced etch, deposition, and metrology equipment. These tools are considered to be particularly important chokepoint tools used only for advanced production, justifying a control scope applying throughout China.[91] The US also extends these restrictions to foreign-made equipment using a legal authority called the SME foreign-produced direct product rule (FDPR), but

[46] These tools are identified under ECCNs 3B001.a.4, c, d, f.1, f.5, f.6, k to n, p.2, p.4, r, and 3B002.c. Examples include ASML's EUV, DUVi systems at or more advanced than NXT:2000i, Lam's Kiyo (advanced anisotropic dry etch), and Applied Materials' Endura Volta CVD Co.

exempts qualifying exports and reexports by eligible entities located in 33 allied countries.[92] These include Japan and the Netherlands, both major producers of many of these tools, which have enacted their own parallel, China-wide controls.[93] The combined effect of these overlapping controls is to block exports of any of these tools to China, regardless of their country of origin.

Second, the US additionally controls a broader set of items when they are destined to specific restricted entities within China. For any item subject to US jurisdiction,[47] exports are restricted whenever there is knowledge that the item will be used in "development" or "production" of ICs destined to a facility where "advanced-node" production occurs or to any entity on the Bureau of Industry and Security (BIS) Entity List.[94] Twenty-one Chinese chipmaking entities are on the Entity List.[48] For certain Category 3 items subject to US jurisdiction, restrictions also apply when the exporter knows the destination is an IC fab but does not know whether advanced-node production occurs there. Practically, these restrictions prevent US toolmakers from providing SME, components, software, or technology to any of China's advanced fabs, including SMIC and CXMT.

The US uses the Footnote 5 FDPR to extraterritorially control an enumerated list of foreign-made tools[49] when destined to specific restricted entities—16 of the 21 on the Entity List.[95] This set of tools is considered to be "node-agnostic"—i.e., important for advanced production but also required for legacy production, making a China-wide control detrimental to legacy production—and includes advanced DUVi systems[50] and etch, deposition, and metrology equipment.[96] Covered foreign-made equipment requires a BIS license when it is destined for a Footnote 5 entity on the BIS Entity List or a facility where "advanced-node integrated circuit" production occurs.[97]

---

[47] The legal term for this is "Subject to the Export Administration Regulations," i.e. that the item is subject to US export-control jurisdiction. Items can become subject to US jurisdiction if they are physically located in the US, contain a substantial amount of US-origin content, or are foreign items that are the direct product of US technology or software covered by a foreign direct product rule (FDPR).

[48] The full list is: Chinese Academy of Sciences Institute of Microelectronics; Fujian Jinhua Integrated Circuit Company, Ltd.; Ningbo Semiconductor International Corporation (NSI); Northern Integrated Circuit Technology Innovation Center (Beijing) Co., Ltd.; PXW Semiconductor Manufactory Co., Ltd.; Semiconductor Manufacturing International (Beijing) Corporation; Semiconductor Manufacturing International Corporation (SMIC); Semiconductor Manufacturing International (Shenzhen) Corporation; Semiconductor Manufacturing International (Tianjin) Corporation; Semiconductor Manufacturing South China Corporation; Shanghai Integrated Circuit Equipment & Materials Industry Innovation Center Co., Ltd.; Shanghai Integrated Circuit Research and Development Center; Shenzhen Pengxinxu Technology Co., Ltd.; Si'En Qingdao Co. Ltd.; SMIC Advanced Technology R&D (Shanghai) Corporation; SMIC Northern Integrated Circuit Manufacturing (Beijing) Co., Ltd.; SMIC Semiconductor Manufacturing (Shanghai) Co., Ltd.; SwaySure Technology Co., Ltd.; Wuhan Xinxin Semiconductor Manufacturing Company Limited; Yangtze Memory Technologies Co., Ltd.; and Zhangjiang Laboratory.

[49] These tools are identified under ECCNs 3B001 (except 3B001.a.4, c, d, f.1, f.5, f.6, g, h, k to n, p.2, p.4, r), 3B002 (except 3B002.c), 3B903, 3B991 (except 3B991.b.2.a through 3B991.b.2.b), 3B992, 3B993, and 3B994. Examples include ASML's NXT:1980i and 1970i (advanced DUVi lithography), Lam's Flex (dielectric etch), and KLA's Archer (overlay metrology).

[50] Advanced DUVi systems like ASML's NXT:1980i are the central focus of this report. One of our findings is that these tools are not "node-agnostic", because earlier DUV systems like ASML's NXT:1470 are sufficient for China's legacy production.

Japan and the Netherlands are exempted from the Footnote 5 FDPR but implement their own parallel controls on a similar set of technologies.[51] For example, the Netherlands' 3B801 controls DUVi systems with dedicated chuck overlay (DCO) at or below 2.4nm, matching US Export Control Classification Number (ECCN) 3B993.f.1.b.2.[99] However, there are two key differences. First, unlike the US, Japan and the Netherlands do not employ a mechanism to impose controls on specific entities. Instead, all of their controls on SME are nationwide, with the ability to grant licenses for shipments to specific entities on a case-by-case basis. The license review policy used to determine which entities may receive controlled SME is not public, but there is ample public evidence that both countries approve sufficient licenses to enable a huge volume of shipments into China. Second, Japan and the Netherlands control *only* a specific set of tools (a subset of the tools covered by the Footnote 5 FDPR), leaving less advanced SME —including earlier DUVi systems—uncontrolled (see below).

The US and its allies also restrict servicing of SME. Servicing activities are vital for ensuring continuous operation of advanced SME. For a lithography system, these activities can include on-site installation, calibration and qualification, preventative maintenance, corrective repair, remote troubleshooting, and field upgrades. US rules define covered servicing to include "maintaining, repairing, overhauling, or refurbishing" and require a license whenever a covered US person services certain items involved in "development" or "production" of "advanced-node ICs," with an exclusion for persons working for qualifying US or allied-headquartered companies.[100] Practically, this prevents US persons from performing any servicing in China's advanced fabs.

The Netherlands takes a narrower approach to servicing controls. Unlike the US rules, Dutch controls do not generally prohibit a Dutch person from performing servicing simply because it occurs at an advanced-node facility. Instead, restrictions apply only to transfers of certain software and technology associated with specifically controlled equipment. Dutch rules cover transfers of "technology" or "software" "required" for the "development," "production," or "use" of the tool.[101] Because "use" encompasses activities such as operation, installation, maintenance, and repair, these provisions can in principle prohibit activities typically involved in routine servicing. However, Dutch controls also exempt transfers of the minimum necessary technology to support installation, operation, maintenance, or repair, for legally exported items, such as DUVi systems exported to and installed in China.[52] In practice, these exceptions allow ASML to continue to service DUVi systems exported under a Dutch license, provided the servicing does not transfer technology and software beyond what is minimally necessary.

---

[51] Unlike the US, Japan and the Netherlands do not have a mechanism for controls on specific entities. Instead, SME is controlled nationwide and applications for licenses to specific entities are reviewed on a case-by-case basis. The technical scope of these nationwide controls is public and matches the scope of the US FDPR (e.g., the Netherlands 3B801 control on DUVi captures systems with dedicated chuck overlay (DCO) at or below 2.4nm, matching US 3B993.f.1.b.2).[98] The license review policy is not public, but there is ample public evidence that the Netherlands and Japan grant at least some applications, enabling exports of a substantial volume of SME to specific entities in China, as we discuss below.

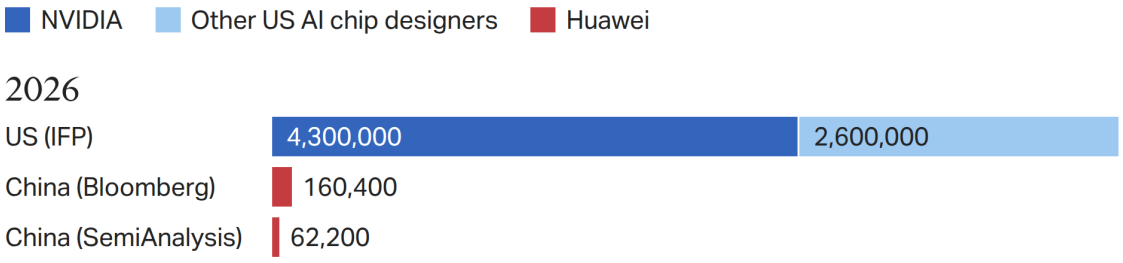
[52] As defined in Annex I of the General Technology Note (GTN) of the EU Dual-Use Regulation, as referenced by the Dutch annex.[102]

Finally, Japan takes arguably an even narrower approach than the Netherlands. For controlled SME only, Japan controls transfers of software programs that support design, manufacturing, or use, where use encompasses operation, installation, maintenance/checking, repair, overhaul, and disassembly repair.[103] By contrast, non-software technology is controlled only when "related to design or manufacturing."[104] Thus, Japanese technicians can provide non-software technical information or know-how, including nearly all ordinary hands-on maintenance, so long as the know-how does not facilitate designing or manufacturing the equipment.

Export controls on SME have provided the US and its allies with a large advantage in AI chip production over China. Controls constrain production of both kinds of silicon used in China's AI chips: advanced logic at 7nm and below, produced primarily at SMIC, and HBM, produced primarily at CXMT. In both cases, controls have limited the volume of advanced equipment available to build and expand advanced production lines and reduced yields on the lines that can be implemented, as fabs are forced to push less capable equipment beyond its normal process envelope while operating without access to servicing support. At SMIC, limited access to advanced deposition, etch, and DUVi systems has constrained expansion of 7nm capacity used for Huawei's Ascend chips and reduced yields, which were reported as "weak" as recently as spring 2026.[105] At CXMT, the same combination of equipment scarcity and low yields, estimated at only 25% for HBM3 as of June 2026, has constrained its HBM2 and HBM3 capacity ramp, likely making HBM the binding constraint on HBM-based AI chip production as of 2026—though Huawei also now markets the non-HBM-based Ascend 950PR which is instead constrained by China's 7nm logic wafer capacity.[106] Together, these constraints have limited China's AI chip production to only a small fraction of the US-led ecosystem's production. In an Institute for Progress report, one of this report's authors (Khan) and colleagues compiled estimates placing Huawei's production between 40,000 and 146,000 B300-equivalent chips (normalized by total processing performance) in 2025, compared with 3.67 million for US firms, and 62,000–160,000 in 2026, compared with 6.89 million for US firms (Figure 9).[107]

Figure 9. US and Chinese AI chip production, 2026

Quantities are normalized to B300-equivalents. Each row is a separate estimate.



5. Export controls let China import DUVi capable of mass-producing Huawei AI chips

Controls on lithography system exports vary depending on the technology generation, national jurisdiction, and facility within China. The most advanced systems—including EUV and the most advanced DUVi—are banned anywhere in China. Advanced DUVi systems—including ASML's NXT:1970i and NXT:1980i—may be shipped into China but are banned to Entity-listed fabs and fabs that engage in "advanced-node production", when US controls apply. Under Dutch or Japanese controls, these systems require a license for shipment anywhere in China, but ample public evidence suggests that many shipments are allowed. The least advanced systems—including ASML's NXT:1965i and earlier DUVi systems, dry ArF and KrF systems, and older lithography technologies—are banned to Entity-listed or "advanced-node" fabs when US controls apply but permitted anywhere in China by Japanese and Dutch controls.

Controls on servicing vary similarly based on technology, jurisdiction, and destination. US toolmakers are banned from performing any servicing activities within "advanced-node" or Entity-listed fabs and on any system exported in violation of the EAR, which *de facto* restricts them from servicing NXT:2000i or more advanced systems, China-wide. Japan's and the Netherlands' servicing controls apply to systems rather than persons. For systems at or more advanced than NXT:1970i, the Netherlands controls transfers of "technology" or "software" for the "use" of the system, but with a general exemption for servicing support required to ensure its continued operation; less advanced systems can be serviced wherever located. Japan restricts transfers of "software" for the "use" of the system but only "technology" transfers when they facilitate design or manufacturing. See Table 4:

Table 4: US and allied controls on lithography system exports and servicing

<table>
<tr><th rowspan="2">Lithography technology</th><th rowspan="2">ASML tools</th><th colspan="2">Controls on tool exports</th><th colspan="3">Controls on servicing</th></tr>
<tr><th>US</th><th>JP/NL</th><th>US Persons</th><th>NL</th><th>JP</th></tr>
<tr><td>EUV</td><td>NXE:3100 and more advanced</td><td colspan="2" rowspan="2">Banned China-wide</td><td rowspan="2">All activities China-wide</td><td rowspan="4">Transfers of “technology” or “software” for the “use”, “design”, or “manufacturing” of the tool
Exemption for minimum support required to ensure continued operation of the tool</td><td rowspan="4">Transfers of “software” for the “use”, “design”, or “manufacturing” of the tool; “technology” for design/manufacturing</td></tr>
<tr><td rowspan="3">DUVi</td><td>NXT:2000i, 2050i, 2100i, 2150i</td></tr>
<tr><td>NXT:1980i, 1970i</td><td rowspan="3">Banned to “advanced-node” and Entity-listed fabs[53]</td><td>License required China-wide but many exports allowed</td><td rowspan="2">All activities at “advanced-node” and Entity-listed fabs</td></tr>
<tr><td>XT:1700i to NXT:1965i</td><td rowspan="2">Uncontrolled</td><td colspan="2" rowspan="2">Uncontrolled</td></tr>
<tr><td>Dry DUV (ArF, KrF); i-line</td><td>NXT:1470 and less advanced</td></tr>
</table>

US controls on the most and least advanced lithography systems reflect a policy objective to restrain China’s advanced production without affecting legacy production. The most advanced DUVi systems, which benefit nodes at 5nm and below[54] but have substitutes at all legacy nodes, are banned China-wide, minimizing the risk of diversion. At the other end of the spectrum, legacy systems (e.g., dry DUV and i-line lithography), which cannot support advanced logic or memory nodes but remain essential to legacy production, are permitted China-wide.

However, controls on advanced DUVi systems (ASML’s NXT:1980i, NXT:1970i, and less capable predecessors) undermine this policy objective, in three ways. First, ASML’s NXT:1980i can support the logic (7nm) and memory (1y) production nodes used for Huawei’s Ascend AI chips, yet is allowed into China.[55] Second, less advanced DUVi systems (e.g. ASML’s NXT:1965i and predecessors) are permitted to enter China’s advanced fabs, despite supporting 14nm logic and 1x memory, nodes used to produce earlier-generation AI chips such as Groq LPUs and the NVIDIA V100. Third, Dutch servicing

[53] This cell is a minor simplification, as the technical scope of US controls differs depending on the origin of the export. For exports originating from the US, all lithography systems (and all items subject to US jurisdiction) are banned to “advanced-node” and Entity-listed fabs. For exports originating outside the US which are made subject to US jurisdiction by the SME or FN5 FDPRs, only systems at or more advanced than NXT:1970i are banned to these fabs.

[54] The most advanced DUVi (NXT:2000i and up) are frequently used alongside EUV systems to pattern less demanding layers at the most advanced nodes (5nm, 3nm, and below). These systems can almost certainly be used to implement 5nm without EUV, and potentially even 3nm with reduced yields. Further, they are fully sufficient for 7nm (Ascend) production, as discussed below.

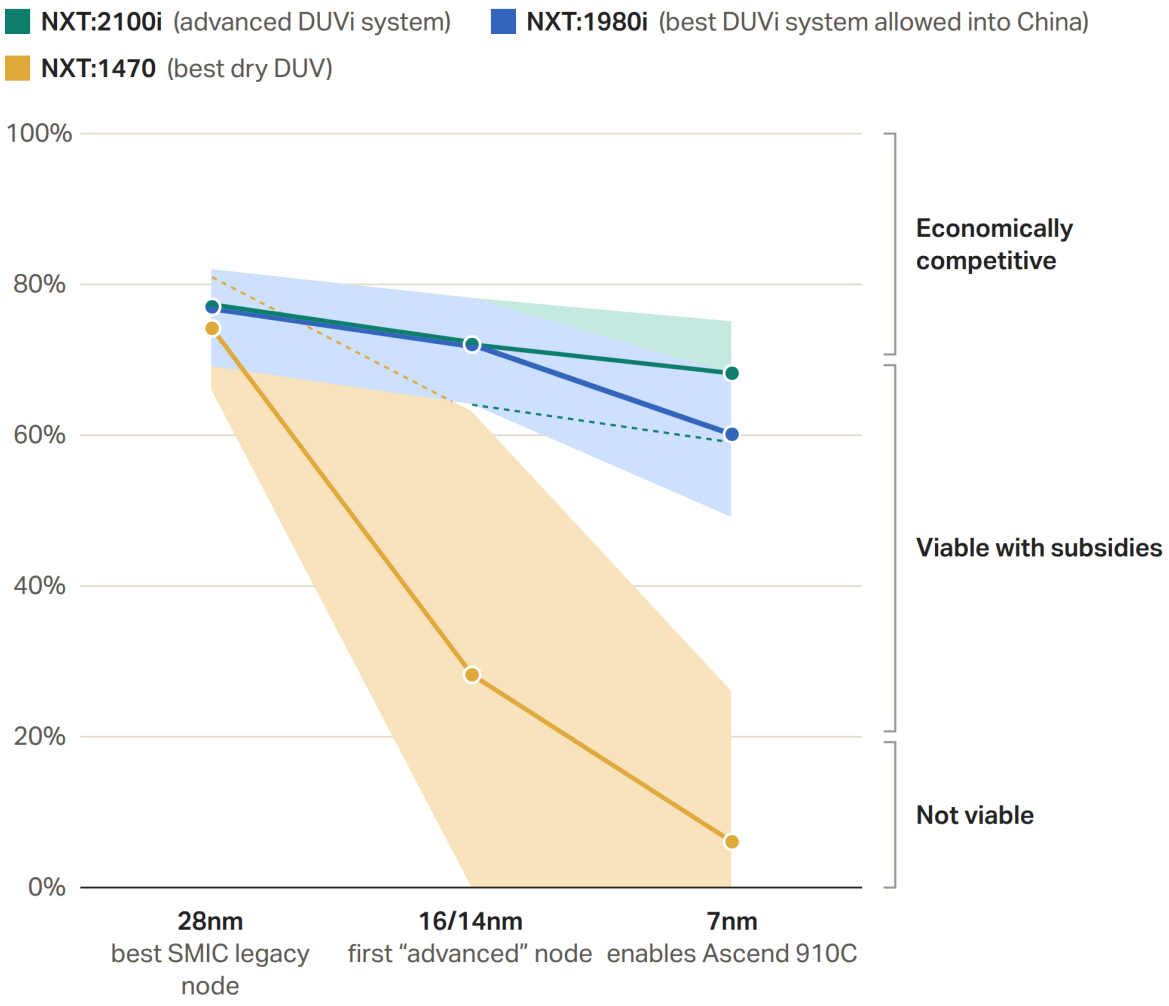
[55] US controls allow NXT:1980i into China so long as it does not enter an “advanced-node” fab. The Netherlands requires a license China-wide, but public data demonstrates that despite these controls, there is a vast influx of these systems into China.

controls permit ASML technicians to support continued operation of DUVi anywhere in China, including at the most advanced fabs, allowing China to sustain its installed base indefinitely.

Problem #1: ASML's NXT:1980i is allowed into China despite enabling production of Ascend AI chips at massive scales

Figure 10. NXT:1980i and more advanced DUVi systems enable Ascend AI chips

Perfect die yields for a 500mm² logic die and a mature manufacturing process



Shaded bands span P10–P90 estimates. See Appendix C for assumptions and calculations.

NXT:1980i is an extremely capable lithography system that enables high-volume production of AI chips like Huawei's Ascend 910C. For logic, the industry record unambiguously confirms that NXT:1980i can support the 7nm node used for the logic dies used in Huawei's AI chips. ASML's original technical paper introduced NXT:1980i as a "scanner for 7nm and 5nm production nodes."[108] TSMC later used the NXT:1980i to develop its first 7nm production line and said that the ramp set company records for "defect density reduction and production volume ramp rate."[109] Our own yield modeling (Appendix C) estimates that NXT:1980i can achieve 50–65% defect-free yield for a large (500mm^2) 7nm die, with overall yield approaching 70–80% after accounting for redundancy (i.e. excluding defects that do not affect chip capability due to redundant circuits) and binning (i.e. chips that achieve lower performance than defect-free versions but are still marketable). These yields are only modestly below those achievable with state-of-the-art (EUV) lithography and would not materially constrain China's logic production. Furthermore, Chinese fabs can also leverage roughly 28 (NXT:2050i/NXT:2100i) systems stockpiled over 2021–2023 (Appendix E) for the one or two most challenging layers in complex designs and NXT:1980i for all other immersion layers, greatly improving yields.

For memory, NXT:1980i can support the 1x and 1y nodes used for HBM2 and HBM2e.[56] ASML's technical writeup introducing the NXT:1980i indicates it was designed with 1x memory as a target application.[111] Public evidence also indicates support for 1y. ASML's 2018 roadmap shows support for 1y a full year before NXT:2000i shipped, implying that early 1y ramps used NXT:1980i.[112] Samsung's initial 1y DRAM mass production in late 2017 was almost certainly implemented primarily using the NXT:1980i, since the successor NXT:2000i had only just begun limited early access testing at that point and did not begin shipping widely until 2018.[113] Samsung claimed a "30 percent productivity gain" for DDR4 based on this 1y node.

NXT:1980i systems produce AI chips at massive scales. We estimate (Appendix D) that a single system operating in a state-of-the-art fab can produce 30,000–47,000 Ascend 910C-equivalents every month—or nearly half a million chips per year (Figure 11).[57] China's fabs achieve substantially lower yields than world-leading fabs, lowering effective throughput. Even so, we estimate a single NXT:1980i deployed in a Chinese fab with sufficient complementary SME could likely produce over 100,000 Ascend 910C-equivalents per year.

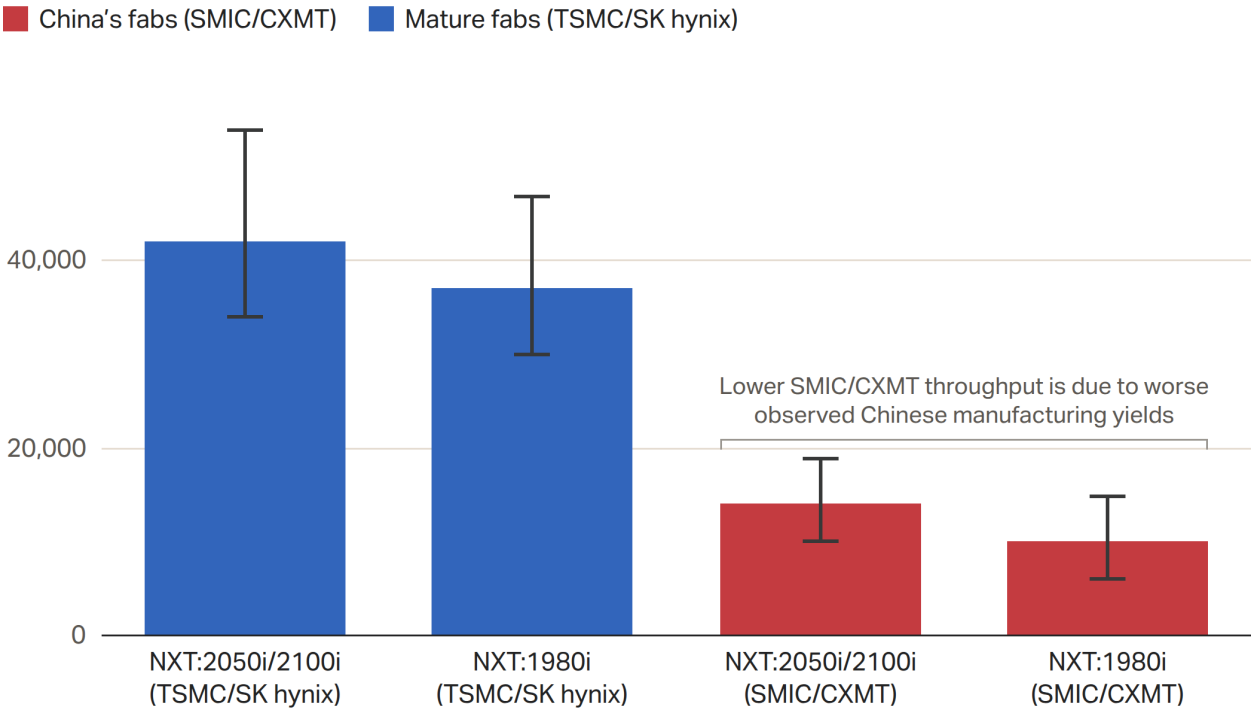
---

[56] Evidence suggests that the NXT:1980i may be viable with reduced yield at the 1z (11-13nm) node as well, potentially enabling HBM3. There is no public evidence that NXT:1980i has been used for critical layers at 1z, but the industry record confirms the broader point that advanced DUVi multi-patterning has been viable for highly advanced DRAM nodes. For example, Micron used DUVi systems for high volume manufacturing all the way to the 1b node—two generations beyond 1z.[110]

[57] As a point of reference, we estimate the annual output from a single NXT:1980i operating in a world-class fab would be sufficient to train Claude Mythos 5 over a 6-month period. Anthropic has not disclosed Claude Mythos 5's training compute, but public estimates cluster around 1×10^{27} FLOP. We speculatively assume Claude Mythos 5's pre-training required ~10^{27} FLOP. An Ascend 910C supports a peak 800 TFLOPS at FP16.[114] Assuming 30–50% utilization during pre-training, 10^{27} FLOP over 6 months would require approximately 160,000–265,000 Ascend 910Cs. Speculatively assuming training time is split evenly between pre-training and reinforcement learning, training Claude Mythos 5 would require approximately 320,000–530,000 Ascend 910Cs. Meanwhile, we estimate a single NXT:1980i scanner can produce 360,000–564,000 Ascend 910Cs.

Figure 11. DUVi throughputs across production settings

Ascend 910C-equivalents per month



Bars reflect P50 estimates for monthly production of finished Ascend 910C chips, assuming throughput is solely constrained by lithography systems. Ranges reflect P10–P90 estimates. See Appendix D for assumptions and calculations.

Center for Technology & Statecraft IFP

Despite this capability, NXT:1980i is allowed into China. The Netherlands requires a China-wide license for NXT:1980i, but public data confirms that many exports are allowed. The Dutch government has not disclosed the criteria it uses to determine which Chinese fabs can receive NXT:1980i. However, China's customs import data (Appendix E) suggests the country is growing its installed base, including at advanced fabs. Even if licenses are consistently denied to certain fabs of concern, there is a serious risk that they may be illicitly diverted or otherwise leveraged to advance China's AI chip production once physically in-country (Section 6).

Problem #2: Less advanced DUV immersion tools are uncontrolled despite enabling earlier AI chips

Older DUVi systems—including ASML's NXT:1965i, 1950i, and XT:1900i—fall outside the scope of the Netherlands' China-wide license requirement. The US restricts transfers of these older systems to

advanced-node fabs but carves out the Netherlands from its FDPR; as a result, these tools are effectively uncontrolled anywhere in China.

The most capable of these systems (NXT:1965i and NXT:1950i) almost certainly enable the advanced logic (16/14nm) and memory (1x) nodes used for earlier AI chips like Groq LPUs.[115] For logic, ASML explicitly markets NXT:1965i as a solution for "sub 20nm nodes," implying viability for 16/14nm production.[116] TSMC almost certainly ramped its 16/14nm node to risk production (i.e. a step before volume commercial production) using 1965i or earlier lithography systems, given that ASML's NXT:1970i first shipped just one month before that milestone.[117] For memory, the NXT:1965i appears to enable 1x production.[58]

ASML does not currently sell NXT:1965i or older systems, but a thriving secondary market is available for these systems,[59] within which China is reported to be a major participant.[119] Recent industry rumors also claimed that ASML was considering restarting production of these systems specifically to support Chinese customers.[120]

Problem #3: Dutch servicing rules preserve ASML's ability to support China's installed DUV immersion fleet

For controlled DUVi—NXT:1970i and more advanced systems—Dutch rules control transfers of "technology" or "software" "required" for the "development," "production," or "use" of the tool.[121] Here, "technology" refers to controlled technical information, like engineering know-how and servicing instructions. "Use" includes operation, installation, maintenance, repair, overhaul, and refurbishing.[122] By default, these provisions would substantially control the support that ASML can provide to China's fabs.

However, Dutch controls also contain an exception that allows transfers of the minimum "technology" needed to install, operate, maintain, or repair[60] an item if either the tool is uncontrolled or its export was authorized.[61] A large fraction of controlled DUVi systems now in China were exported under Dutch licenses. This exception therefore ensures that ASML can continue to provide support needed to maintain most of China's controlled DUVi fleet.

---

[58] Specifically, the NXT:1965i enables adequate overlay for 1x production, provided that the same machine is used for the most critical layer alignments.

[59] ASML claims "~95% of the systems we have sold in the last 30 years are still in use.[118] Many customers are able to generate value by selling systems they no longer require."

[60] Overhaul and refurbishment are notably excluded.

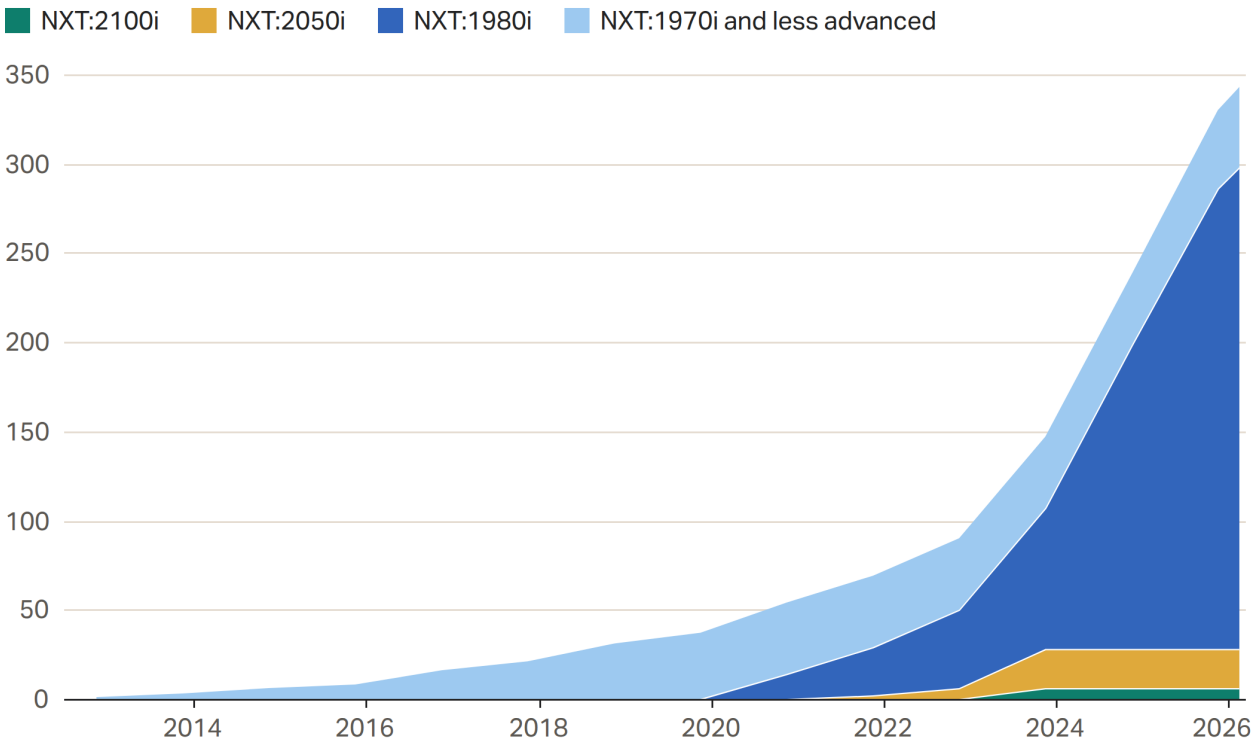
[61] As defined in Annex I of the General Technology Note (GTN) of the EU Dual-Use Regulation, as referenced by the Dutch annex.[123]

6. China is building a massive stockpile of DUVi that enable AI chipmaking

Chinese chipmakers are importing vast numbers of DUVi systems, the large majority of which are NXT:1980i capable of making Ascend AI chips. China first began importing NXT:1980i systems in the late 2010s, after the tool first entered commercial production, but drastically accelerated imports after the Dutch government announced export controls on more advanced DUVi systems in 2023 (Figure 2). As of 2026Q1, we estimate that roughly 270 NXT:1980i systems are installed in Chinese-owned fabs,[62] accounting for ~80% of the DUVi stockpile at these entities. The remainder of China's DUVi stockpile consists of an estimated 22 NXT:2050i systems, imported between 2021 and 2023 before China-wide restrictions took effect in January 2024; 6 NXT:2100i systems, imported over the same period; and ~45 DUVi systems less advanced than the NXT:1980i. See Appendix E for our estimation methodology.

Figure 2. China is building a vast stockpile of DUVi systems

Estimated cumulative DUVi units imported by Chinese-owned fabs since 2012



Estimates exclude imports to multinational fabs located within China. See Appendix E for assumptions and calculations.

Center for Technology & Statecraft IFP

[62] There are also at least 4 multinationally owned fabs with advanced-node production lines operating in China. We model the aggregate DUVi stockpile at these fabs as well, in Appendix E. All top-line findings refer only to stockpiles at fabs owned by Chinese entities, as they are the relevant quantities for understanding China's AI chipmaking capacity.

China's surge of NXT:1980i imports has absorbed a large fraction of ASML's system sales and required a massive investment from China's chipmaking industry. Within the DUVi segment, China's share of ASML shipments surged from 26% in 2022 to 60–70% over 2024Q1–2026Q1. We estimate that China spent over $13 billion on NXT:1980i systems over 2024–2025 alone.

China's NXT:1980i imports risk advancing its AI chip production

China's continued imports of NXT:1980i systems could support AI chip production through two pathways. First, systems delivered to China's advanced fabs can directly advance their capacity to produce AI chips. Although the Dutch government requires licenses for all NXT:1980i exports to China, it does not publicly disclose the criteria used to review individual license applications. It is therefore unclear whether the Netherlands denies licenses to all fabs that produce advanced chips. Second, even if the Netherlands consistently denies licenses to these fabs, systems shipped to other Chinese fabs could later support AI chip production through diversion, cross-fab wafer processing, or unauthorized upgrades.

Public evidence suggests that at least some of China's advanced fabs continue to obtain NXT:1980i systems. First, Anhui customs records and information from third-party contractors (Appendix E) suggest that CXMT, China's flagship DRAM manufacturer and HBM producer, has continued to obtain NXT:1980i systems even after the Netherlands began requiring nationwide licenses in 2024. Second, Hua Hong is almost certainly preparing to launch a 7nm logic fab, consistent with public reporting (Appendix E) that it has previously obtained an NXT:2050i system (an investment that would make little sense absent plans for advanced-node production).[124] Yet neither Hua Hong nor its subsidiary (Huali Microelectronics) has been added to the BIS Entity List, suggesting that US policymakers may have overlooked the company as an advanced-node producer. In April 2026, BIS reportedly issued "is-informed" letters directing US toolmakers to halt certain shipments to Hua Hong's Fab 6 and the planned Fab 8a.[125] It remains unclear whether the Netherlands matched that action through changes to its licensing policy, or whether Dutch authorities had previously granted Hua Hong licenses to import NXT:1980i. Even if Hua Hong is now effectively prohibited from obtaining NXT:1980i, it may have had over two years following China-wide Dutch controls to accumulate an NXT:1980i stockpile.

The same problem may recur over and over as China builds new fabs. China regularly establishes fabs through newly created corporate entities, creating a risk that a facility will obtain NXT:1980i before regulators recognize its plans for advanced-node production. US controls partially address this problem by requiring a license whenever the exporter knows that the covered equipment will be used for "advanced-node" production or if the end-use is unclear. It remains unclear whether the Netherlands applies an equivalent standard when reviewing licenses.

However, even under the most optimistic scenario—in which the Dutch government denies licenses for exports to fabs that support AI chipmaking, matching the scope of the US FN5 FDPR, and reliably identifies all such entities—China has at least four pathways to enlist NXT:1980i systems in support of AI chip production once they are in-country:

First, a fab could obtain NXT:1980i systems while producing legacy chips and subsequently upgrade the production line to an advanced node. Sometimes referred to as "diversion in place," this practice violates export controls but is difficult to prevent. Chinese fabs have previously upgraded existing lines to new nodes, even at advanced nodes,[63] making this scenario plausible.

Second, China could move NXT:1980i systems or their subsystems from an approved fab to a restricted one ("physical diversion"). China regularly diverts controlled technologies to restricted end users and facilities. The very continued existence of Huawei and its fab network (including its own fabs, SMIC, CXMT) involves a massive set of violations and diversion activities, illicit shipments, warehouse stockpiling, and wafer transfers between fabs.[128] These violations have extended to advanced SME. For example, at the start of 2025, SMIC was reportedly able to divert controlled deposition, etch, and inspection/metrology equipment from non-advanced fabs owned by SiEn, Pensun, and Huawei's Dongguan fab to its 7nm SN2 fab.[129]

The risk of diversion is particularly high for NXT:1980i subsystems. Many key subsystems, including Cymer light sources, are smaller, less technically complex, and easier to transport and install than a complete DUVi system.[130] For example, a legacy fab could import a replacement subsystem for an allegedly defective part in an installed NXT:1980i, then divert the replacement subsystem to an advanced fab. The Dutch government may be able to mitigate this pathway by requiring end use verification as a condition of licenses for NXT:1980i subsystems, but enforcing such a condition across China's fleet of approximately 270 NXT:1980i systems would require extensive on-site inspections or verification technology.

Third, China might route partially processed wafers between advanced and non-advanced fabs ("wafer diversion"), potentially dramatically expanding its advanced-node output. For example, an advanced fab at SMIC or CXMT might reserve a small number of more capable NXT:2050i or NXT:2100i systems for the most critical layers, while sending wafers to another fab's NXT:1980i systems to pattern less demanding layers. Moving partially processed wafers between fabs is a recognized industry practice[64] and Huawei has reportedly built infrastructure for transporting wafers between

---

[63] For example, SMIC is reported to have built its first 12nm line as an extension of 14nm.[126] CXMT's 2025 IPO prospectus describes a DRAM mass-production line technical-upgrade project and a DRAM demonstration-line upgrade to a newer-generation process platform at the existing site.[127]

[64] TSMC has patented systems specifically for transferring wafer-filled front-opening unified pods (FOUPs) between fabs; Chartered Semiconductor planned for a "borderless fab" model in which wafers are moved among fabs according to available capacity and process technology; academic research has modeled production routes in which neighboring fabs divide processing operations for a single wafer.[131]

physically adjacent fabs.[132] The principal technical challenge would be qualifying a single production process across multiple fabs. This challenge is substantial but likely surmountable, particularly when the participating fabs are nearby and controlled by a single entity or closely cooperating entities.

Finally, China could upgrade NXT:1980i systems after installation, exacerbating the impact of any of the preceding diversion pathways. ASML designed the NXT platform to allow machines to be upgraded after installation. Upgrades can improve several aspects of scanner performance, including overlay and alignment accuracy, imaging and focus control, and throughput. Dutch controls prevent ASML from improving the NXT:1980i's throughput by more than 1 percent.[65] Nevertheless, public reporting indicates that Chinese fabs have upgraded NXT:1980i systems, thereby improving 7nm production yields, by purchasing upgraded wafer stages, lenses, and alignment sensors through the secondary resale market and relying on Chinese vendors for installation and servicing.[134]

China's NXT:1980i stockpile could also accelerate its development of domestic lithography systems. A large installed fleet of NXT:1980i systems can supply grey-market components for domestic equipment;[66] give Chinese engineers access to reference systems to study and reverse engineer; provide benchmarks for calibrating indigenous prototypes;[67] and build tacit expertise among domestic engineers and maintenance technicians. However, a ban would provide a market opening for Chinese lithography makers; we do not assume either effect dominates.

In the near term, Chinese companies are unlikely to divert large numbers of NXT:1980i systems,[68] though smaller-scale diversion remains plausible. Other SME bottlenecks—particularly limited access to advanced deposition and etch tools restricted by BIS between 2022 and 2025—independently constrain China's advanced chipmaking capacity. These bottlenecks limit the additional production capacity unlocked by diverted NXT:1980i systems, reducing the immediate benefits of diversion.

China's incentive to divert NXT:1980i systems will grow, however, as it develops domestic substitutes for an increasing share of the SME supply chain over time. As discussed in Section 3, lithography is likely to be the last major SME indigenized by China's domestic industry. Once other bottlenecks in deposition, etch, metrology, and related process equipment are alleviated, the opportunity cost of keeping NXT:1980i systems on non-advanced lines will rise. If advanced lithography becomes the

---

[65] See 3B801.f.7.[133]

[66] Reuters claims that the availability of parts on secondary markets was a key enabler of China's domestic EUV prototype.[135]

[67] Reuters claims that China's indigenization efforts include "a team of around 100 recent university graduates… focused on reverse-engineering components from both EUV and DUV lithography machines, according to the people…Each worker's desk is filmed by an individual camera to document their efforts to disassemble and reassemble parts—work… described as key to China's lithography efforts."[136]

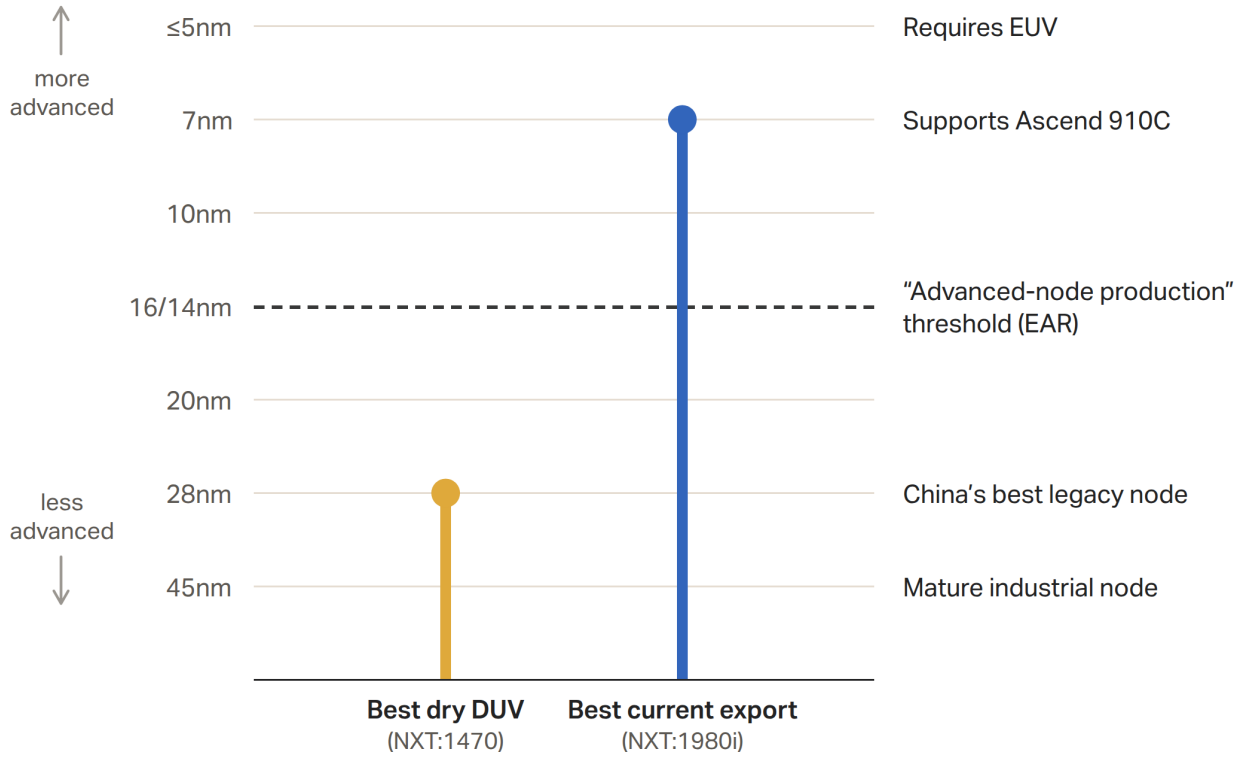
[68] This is because advanced chipmakers may currently be bottlenecked by access to export-controlled SME such as deposition and etch tools, making additional diversion of lithography tools unhelpful for expanding wafer capacity. Additionally, Chinese chipmakers may wish to continue meeting demand for legacy chips made with DUVi systems.

binding constraint on China's AI chip production, China may become increasingly willing to accept the political risks associated with diversion.

7. A better alternative to exporting DUVi: dry DUV can support China's legacy chips

Dry DUV systems—the generation of lithography tools that preceded DUVi—can support all of China's legacy logic and memory nodes *without* enabling the nodes used to produce AI chips and HBM. By restricting DUVi exports while permitting dry DUV exports, the US and its allies can restrain China's AI chip production without damaging its legacy chip production.

Figure 12. Dry DUV enables China's legacy production without supporting AI chip production



Center for Technology & Statecraft IFP

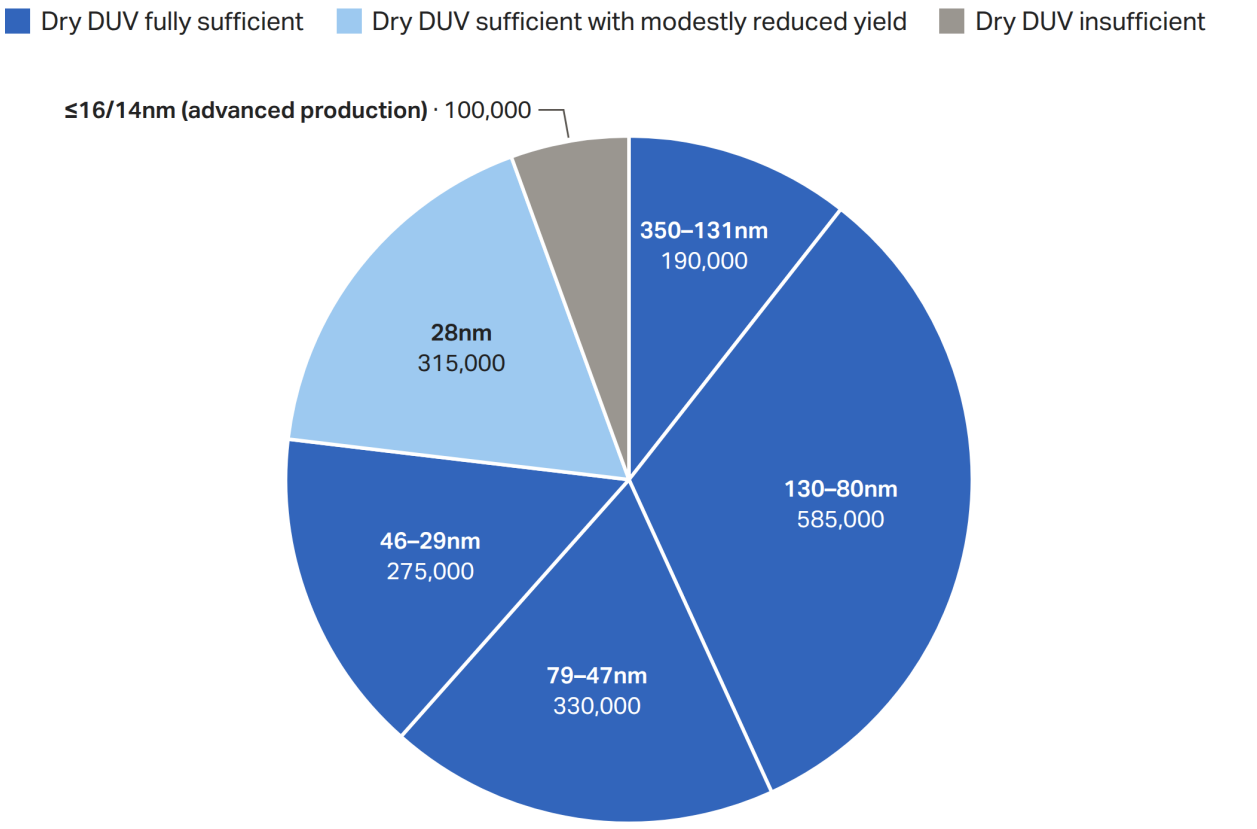
Dry DUV systems can support all of China's legacy logic production. The majority of Chinese logic capacity is at 45nm or older nodes, which enable chips found in cars, household appliances, and industrial equipment. Advanced dry DUV systems are demonstrably sufficient for all of this production[69] (Figure 13). A minority of its capacity extends to 28nm. At this node, China's existing commercial processes incorporate DUVi systems (which were widely available by the time it was implemented), but experimental results and our own modeling find that dry systems can support this

[69] Intel's high-volume 45nm node used dry ArF lithography.[137] This node had gate and M1 pitches of 160nm (80nm half-pitch), corresponding to a k-factor of 0.386, which is well within the viable production regime. The 45nm node's gate critical dimension is 35nm, which Intel achieved using line-and-cut double patterning.

node as well.[138] Our modeling estimates that ASML's best dry DUV system (NXT:1470) can support 28nm with a modest yield penalty (Appendix C), using additional multipatterning. Industry history also supports the plausibility of extending dry DUV to the 28nm node. ASML marketed its XT:1450, an older dry DUV system, as supporting the development of 32nm node processes using double patterning.[139] ASML's NXT:1470 substantially improves on-product overlay compared to XT:1450 (from 7nm to 4.5nm), plausibly enabling it to extend into 20nm-class logic. Researchers have also experimentally demonstrated 45nm and even 32nm half-pitch patterns using dry DUV with double patterning, providing further evidence of its patterning capability.[140]

Figure 13. China's legacy production does not need DUVi

Dry DUV enables 81% of China's legacy production without any yield loss and 19% with minimal loss



Chinese-owned installed capacity as of 2025Q4 in 300mm-equivalent wafers per month. Data from SEMI World Fab Forecast.

Center for Technology & Statecraft IFP

For memory, the evidence is more limited but indicates that dry DUV systems are probably viable for DRAM nodes in the high-20nm range at least, including the 29nm node used for HBM1, with a modest

yield penalty. The most relevant patterning evidence is Nanya's sub-30nm DRAM study, which used a dry DUV tool to produce 30nm line/space patterns using self-aligned double patterning.[141] Samsung's study of 20nm honeycomb DRAM provides weaker, indirect corroboration by demonstrating that self-aligned patterning techniques can reduce a difficult capacitor-patterning sequence from three immersion exposures to one.[142]

However, dry DUV is unlikely to enable the advanced nodes used for AI chips. For logic, we estimate (Appendix C) that NXT:1470 is very unlikely to achieve a viable yield for large 7nm dies. The first problem is alignment. NXT:1470's published on-product overlay does not meet the effective 7nm overlay budget, increasing the risk of misalignment between layers and between multiple exposures within a single layer. Second, it is unclear whether NXT:1470 or any other dry system can implement a 7nm flow with sufficient control of individual features (e.g., cuts, contacts, and other critical features) due to critical-dimension nonuniformity, focus and dose variation, and pattern-dependent imaging errors. It is conceivable that a 7nm flow could be constructed using extremely restrictive design rules—e.g., highly regular, parallel lines with narrowly prescribed locations for cuts and contacts—but no such process has been demonstrated. Even if it proved technically feasible, these additional design restrictions would constrain circuit layout and routing, likely reducing chip density and performance and limiting the kinds of chips that could be produced.

For memory, NXT:1470 is highly unlikely to support 1x and effectively incapable of supporting 1y. At 1x, ASML's illustrative flow specifies a 4nm on-product overlay budget, which is exceeded by NXT:1470's 4.5nm on-product overlay specification, leaving no room for other edge placement errors (e.g., pitch walking). At 1y, both alignment and patterning constraints become tighter. Even with quadruple patterning, each individual exposure would approach NXT:1470's 57nm resolution limit, leaving essentially no margin for error. Repeating these operations over many critical layers probably substantially reduces patterning yield.

8. The US government can control DUVi systems with executive branch authorities

China's large and growing DUVi stockpile is the consequence of a policy choice to restrict only the most advanced lithography systems nationwide, while allowing earlier but still highly capable systems into the country and in some cases directly into advanced fabs.

To ensure that China cannot leverage foreign DUVi for AI chipmaking, the US must stop the flow of these systems into the country and develop enforcement mechanisms to prevent diversion of those systems that have already been stockpiled. The following actions would achieve these objectives.

First, the US government should ban new exports of any DUVi system nationwide to China, extending to ASML's XT:1700i. The US government can implement this change by removing the dedicated chuck overlay (DCO) criterion in ECCN 3B001.f.1, with a policy of denial.[70][144] After the change, China will be unable to import DUVi but still allowed to import dry DUV systems, which are unsuitable for AI chip production but sufficient for China's legacy production.

To be effective, the US government will need to ensure that new controls apply to DUVi produced outside the United States by ASML and Nikon. There are four options for doing so:

- Option 1: The Dutch and Japanese governments apply equivalent controls alongside the US government, with a policy to deny all license applications. The Japanese and Dutch governments currently require a license for DUVi at or more advanced than NXT:1970i,[71] but public evidence (Appendix E) shows that many licenses are granted. Additionally, as for US ECCN 3B001.f.1 and 3B993.f.1, the technical definitions used in Dutch and Japanese controls do not control earlier-generation ASML and Nikon DUVi systems that are nevertheless capable of producing advanced-node chips.[72] Alignment with the US would require that the Dutch and

[70] Ally and partner advanced fabs in China – Samsung Xi'an, SK hynix Wuxi, SK hynix (formerly Intel) Dalian and TSMC Nanjing – could be exempted from these export, parts and servicing restrictions as long as they remain owned and operated by their current parent companies. Because all of their parent companies are CHIPS Act grantees, they are subject to limitations on their ability to expand fab capacity in China.[143]

[71] Japan controls DUVi systems under Appended Table 1, row 7, item (16) to the Export Trade Control Order, with the technical criteria specified in Article 6, item 17, subitem ヲ (wo) of the implementing ministerial ordinance (MITI Ordinance No. 49 of 1991, as amended).[145] The Netherlands requires a license under Article 2 of the Regeling geavanceerde productieapparatuur voor halfgeleiders (Regulation on Advanced Semiconductor Manufacturing Equipment; Staatscourant 2023, 18212, as amended) to export DUVi systems meeting the technical criteria specified in entry 3B801.f.6 of its current Annex, including a DCO value greater than 1.50nm and no greater than 2.40nm.[146] DUVi systems with a maximum DCO value of 1.50nm or less are controlled under entry 3B501.f.1.b of Annex I to Regulation (EU) 2021/821, as amended.[147]

[72] As discussed in Section 5, the Dutch license requirement covers the ASML NXT:1970i and more advanced DUVi systems, but not NXT:1965i and other less advanced DUVi systems. Based on Nikon's published specifications, Japan's list control appears to cover the currently marketed Nikon NSR-S625E, NSR-S635E, and NSR-S636E as well as the discontinued NSR-S621D, NSR-S622D, NSR-S630D, and NSR-S631E. However, an older discontinued Nikon DUVi

Japanese governments remove the DCO parameter from their technical definitions and deny all licenses for DUVi systems covered by the license requirement. In advocating for this outcome, the United States can emphasize (see Section 9) that new controls are likely to shift semiconductor production outside of China, without affecting overall demand for DUVi.

If the Netherlands cannot be convinced to join in controls, the US government can pursue controls on foreign-produced DUVi in collaboration with Japan:

- Option 2A: The Japanese government restricts DUVi made in Japan and the US and Japanese governments control critical DUVi subsystems made in the US and Japan for Dutch DUVi systems otherwise destined to China. First, as in Option 1, the Japanese government can restrict all DUVi made in Japan. Second, all existing DUVi systems incorporate a light source—an irreplaceable technical subsystem used to generate deep ultraviolet light—from either a US company (Cymer) or a Japanese company (Gigaphoton).[148] The US and Japan can jointly prevent DUVi exports to China, without the use of extraterritorial controls or needing cooperation from the Netherlands, by controlling their domestically produced light sources when the exporter has knowledge that the "end use" of that subsystem is incorporation into an ASML DUVi system destined to China.

- Option 2B: The Japanese government restricts DUVi systems made in Japan and the US government controls Dutch DUVi extraterritorially using the *de minimis* authority. First, as in Option 1, the Japanese government can restrict all DUVi made in Japan. Second, the *de minimis* authority allows the US government to control items produced abroad that incorporate more than a specified percentage of controlled US content. All existing ASML DUVi systems, including those without a US light source, incorporate some amount of US content, including wafer stages and US-origin integrated circuits. The US government can therefore control all ASML DUVi extraterritorially, without relying on cooperation from the Netherlands, by setting a 0% *de minimis* threshold for items meeting the revised thresholds in ECCN 3B001.f.1.[73] It is possible that even Nikon DUVi systems made in Japan contain US-origin content that is difficult to remove, in which case Option 2B would not require Japanese collaboration. However, there is limited evidence as to whether this is the case, which is why Option 2B may require Japanese collaboration.

---

system, the NSR-S610C, is not covered by Japanese controls, while it is unclear whether another discontinued Nikon DUVi system, the NSR-S620D, is covered, since its published DCO of ≤3nm is outside the ≤2.4nm control threshold, while demonstrated field results have shown performance at ≤2nm.

[73] In October 2023, the US government previously controlled the NXT:1970i and NXT:1980i using a 0% *de minimis* control but exempted countries which control equivalent commodities.[149] Japan issued equivalent controls in 2023, while the Netherlands issued equivalent controls in 2024. Therefore, the 0% *de minimis* rule no longer applies to exports from these two countries. Option 2B can be enacted by expanding the scope of the de minimis control to apply to all DUVi and removing the carve-out that exempts the Netherlands.

If neither Japan nor the Netherlands can be convinced to join in controls, the US government has the executive authorities required to unilaterally control all foreign-made DUVi systems to China:

- Option 3: The US government controls both Japanese and Dutch DUVi extraterritorially using the Foreign Direct Product Rule (FDPR). The SME FDPR allows the US government to control items produced abroad that are the direct product of US-origin software or technology.[150] All existing DUVi systems meet this criterion because they incorporate integrated circuits that are presumed to be the product of US technology. At present, the US government carves out Japan and the Netherlands from the license requirement for items subject to the SME FDPR, including 3B001.f.1. The US government can enact extraterritorial controls by removing this carveout for 3B001.f.1 items. Ideally, Option 3 should be coupled with the US *de minimis* control in Option 2B to make it more technically and legally robust.

Second, the US government should restrict exports of DUVi subcomponents nationwide, including replacement parts for China's existing DUVi systems. These replacement parts are at high risk of diversion from non-advanced to advanced fabs, as discussed above (see Section 6), making a nationwide license requirement essential.[74] A key challenge is that China currently uses DUVi within many of its legacy fabs in addition to its advanced fabs; restrictions on replacement parts for those DUVi could therefore affect legacy production in addition to advanced production. To mitigate effects on legacy production, the US government could grant individual licenses for exports to legacy fabs while specifying end-use requirements (e.g., one-for-one replacement, regular inspections and reporting by ASML) that deter diversion.

Third, the US government should bar foreign nationals from servicing DUVi at China's advanced fabs. The United States already prevents US toolmakers from performing any form of servicing at advanced fabs. However, Japan and the Netherlands restrict servicing only for specifically controlled tools, and prohibit a smaller set of activities, as discussed in Section 5. At legacy fabs, servicing should continue on condition that there is reporting and verification on the tool's location and usage, one-for-one parts replacement, and periodic inspections.

- One way the US government can achieve this outcome is to convince the Netherlands and Japan to expand their controls to cover the full set of servicing activities and to apply those controls to DUVi systems.
- If either is unwilling, the US government has the authority to unilaterally ban servicing on DUVi located in an advanced fab using the EAR's General Prohibition 10.[152] This prohibition explicitly prohibits servicing of any item subject to US jurisdiction when there is "knowledge" that an export control violation has occurred or will occur "in connection with the item." The US

[74] At a minimum, this should involve China-wide controls on "specially designed" components for any DUVi systems. ASML says its dry DUV and DUVi systems share some subsystems, including wafer stages, reticle stages and reticle handling, so a control limited to DUVi-specific parts could miss components shipped as spares for dry NXT tools.[151] The US government could consider controls on some of these items and license use for dry DUV tools.

government has already declared that Chinese advanced-computing ICs are likely produced in violation of US export controls.[153] After asserting jurisdiction over foreign-produced DUVi (see above), the US government can restrict servicing of DUVi at China's advanced fabs by informing the public that there is a high likelihood that these DUVi are being used in connection with a violation of the US export controls, e.g. production of Huawei Ascend chips.

Fourth, the US government should work jointly with the Netherlands and ASML to detect and deter diversion of China's existing DUVi stockpile. As discussed above (see Section 6), China has already imported a vast number of NXT:1980i DUVi systems to non-advanced fabs, which could massively accelerate AI chip production if diverted from non-advanced to advanced production. ASML probably has the capability to neutralize diversion using technical mechanisms.[75] The US government should work with ASML and the Dutch government to obtain insight into this capability and ensure it is robust to countermeasures. Granting licenses for servicing at legacy fabs may also preserve ASML's visibility into how those tools are used, reducing the risk of diversion.

Fifth, the US government should also work to restrict subcomponents, precursors, and knowledge transfers that enable China's effort to indigenize advanced DUVi. Priorities include Zeiss projection optics, which are rumored to be entering China, together with the materials and equipment (e.g., fused silica) needed to make comparable lenses; high-power ArF excimer light sources and their subcomponents and consumables; and precision wafer and reticle stages and their motion sensors. See Appendix B for further discussion of DUVi subsystems. Although China will eventually indigenize commercially viable DUVi or more advanced lithography, delaying this point even by a few years could constrain its AI chip production even further, by postponing the subsequent ramp-up to reliable, high-volume manufacturing.

[75] ASML has allegedly assured the Dutch Government that it has the capacity to disable EUV systems at TSMC if China were to invade Taiwan.[154]

Box 1: The MATCH Act

The MATCH Act is a bipartisan, bicameral bill included in the manager's amendment to the Senate must-pass annual National Defense Authorization Act (NDAA).[155] As of September 2026, the Senate has not passed its NDAA, and the House-passed NDAA does not include MATCH.

MATCH calls for a countrywide control, with a presumption of denial licensing policy, on exports of DUVi systems to all of China and on servicing them anywhere in China, except to or in fabs that existed at enactment and remain owned and operated by companies headquartered outside China and other countries of concern. It also mandates controls on any item subject to US jurisdiction and its servicing, with a presumption of denial, at certain Chinese-owned fabs of concern (including all Chinese advanced-node fabs) by all persons, including foreign persons.

Within 150 days of enactment, Commerce must impose these controls on US-made SME and other items subject to US jurisdiction (to the extent not already implemented by Commerce), and the executive branch is urged to negotiate with allies to match them. For any ally that has not done so, Commerce must apply US controls on foreign-made SME and servicing by foreign persons. The deadline can be extended once by up to 90 days. The effective date can also be deferred by up to 30 days for an ally that has committed to act.

Congressionally mandated DUVi restrictions would leave little practical flexibility to roll back controls, given the difficulty in passing new legislation revoking MATCH, thus reducing the incentive for Chinese retaliation aimed at convincing the US government to withdraw export controls.[156] However, the potential time for MATCH-directed controls to take effect would enable further Chinese DUVi stockpiling.

9. Banning DUVi exports would secure the US AI chipmaking advantage

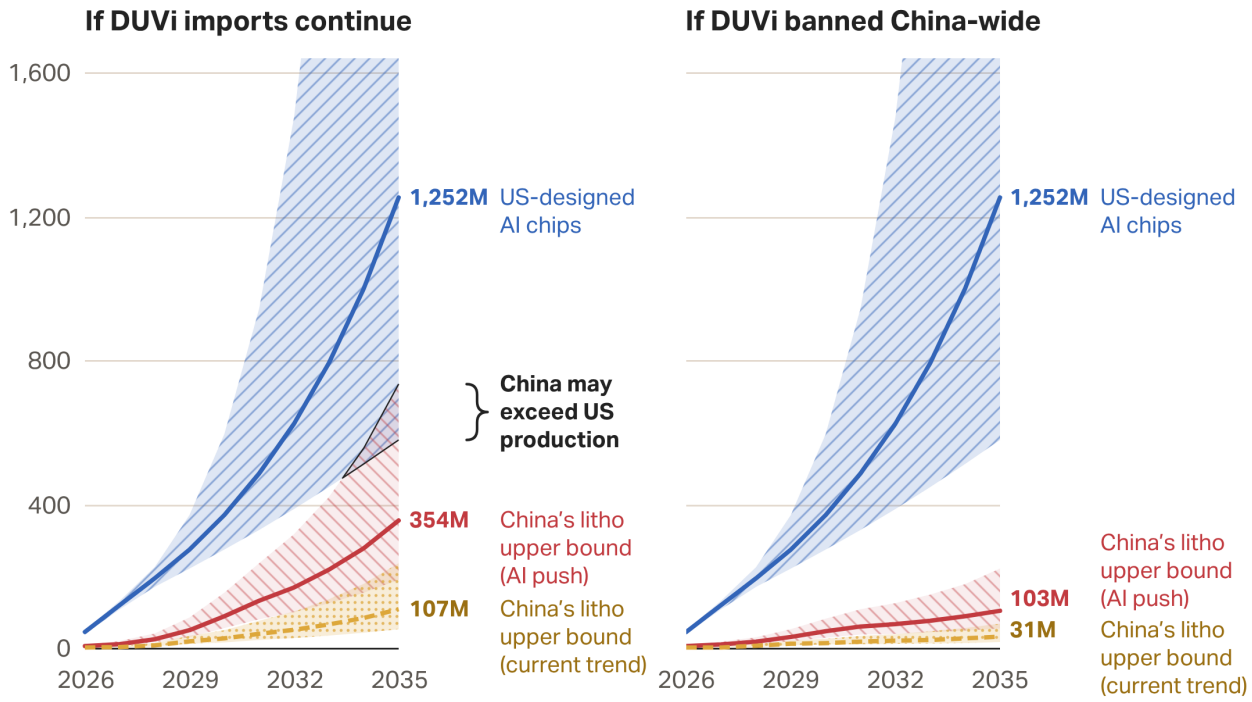
Banning DUVi exports and servicing would reshape US-China competition in AI and semiconductors, with the following effects:

First, it would guarantee the United States and its allies a large AI chip production advantage over China through at least 2035. China's current DUVi stockpile is physically incapable of matching the US-allied ecosystem's AI chip production. Even if China overcomes every other SME bottleneck, diverts every NXT:1980i to AI chip production, and matches the throughput achieved in Western fabs, it would still produce far fewer AI chips than the US ecosystem (Figure 1, Appendix F). Continued DUVi exports would not automatically eliminate this advantage, assuming they are split between advanced and legacy fabs, but they would fundamentally alter the strategic situation. If China continues to import NXT:1980i systems at its 2024–2025 pace or increases its imports moderately over time, and mobilizes this fleet for AI chip production (e.g., after indigenizing other SME bottlenecks), its annual output could reach hundreds of millions of NVIDIA H100-equivalents, enough to erase much of the US advantage and even surpass it under certain scenarios.

Figure 1. Banning DUVi to China secures a US AI chipmaking advantage through 2035

China's existing DUVi fleet cannot match US AI chip production; but continued imports could let it close the gap

Bold lines show median projections (P50); lightly shaded cones span the P10–P90 range



US chip production is estimated for 2026–2027 from TrendForce shipment forecasts and from 2028 onwards by averaging a demand projection and an EUV supply model. H100-equivalents are dense FP8 FLOP/s (INT8 where a chip lacks FP8) divided by the H100's 1,979 TFLOP/s. China's forecasts are for lithography bounded production capacity, assuming it indigenizes all other SME. "Current trend" assumes DUVi are allocated between advanced and legacy fabs in the same proportion as today. "AI push" assumes China mobilizes its DUVi fleet to prioritize AI, obtaining the same allocation forecasted in the US ecosystem. "DUVi banned China-wide" assumes that DUVi imports stop beginning 2027. "DUVi imports continue" assumes a moderate increase in imports over time relative to China's 2024–2025 pace. See Appendix F for assumptions and calculations.

Center for Technology & Statecraft IFP

Second, China's near-term AI production ramp would face additional barriers. First, China-wide DUVi controls would cut off the flow of new systems and subsystems that might otherwise benefit AI chip production, whether through licensed shipments or illicit diversion (Section 6). Second, enhanced restrictions on servicing of DUVi at SMIC, CXMT, and other fabs operating advanced production nodes would degrade the functionality of these tools, reducing overall yields and throughput, though there is some uncertainty as to the magnitude of this effect.

Third, new legacy fab construction in China would be unaffected. China's legacy production nodes can be supported by dry DUV systems, as discussed above (see Section 7). These dry DUV systems would continue to enter China even after a ban on DUVi, enabling expansion of legacy capacity. Restrictions on replacement parts might degrade these DUV systems, but this impact can be mitigated by issuing individual licenses for verified legacy fabs.

Fourth, new advanced logic and memory production would shift away from China and towards the US and its allies. A nationwide ban on DUVi would make it far more difficult for China to expand capacity at advanced logic and memory nodes. With Chinese capacity expansions hobbled, demand for advanced chips would shift to fuel new fab expansions within the allied ecosystem. Recognizing this point, ASML's previous CEO, Peter Wennink, stated in 2022 that even if China were "excluded from any growth," new fabs would be built elsewhere and ASML's 2030 outlook would not change much.[157]

Fifth, US and allied memory foundries, logic foundries, and chip designers would benefit. For memory, US manufacturer Micron, Japanese manufacturer Kioxia, and Korean manufacturers Samsung and SK hynix face serious and potentially existential competitive threats from China's leading fabs, YMTC and CXMT. China-wide DUVi restrictions would slow both firms' near-term ramp at advanced memory nodes, particularly those needed to make HBM, and place a hard limit on their long-term capacity. Allied logic foundries like TSMC, Intel, and Samsung would similarly benefit from limitations on China's leading logic fabs, particularly SMIC, ensuring US and allied leadership at the leading edge. Finally, US and allied chip designers would also benefit relative to the status quo. Existing US rules likely already restrict these designers from manufacturing advanced chips at SMIC, while Huawei faces no such limitation.[76] Further limitations on China's advanced chip manufacturing will disadvantage Huawei in the global chip marketplace, bolstering allied designers.

Sixth, ASML would bear the primary costs among US and allied toolmakers, but these costs are modest and mostly recoverable.[77] Assuming a complete restriction on DUVi sales and servicing to China, ASML could face revenue losses of ~$1–1.5 billion per year (3–4% of its annual revenue) but may be able to recover some of these losses. We break down ASML's revenue losses and recuperation opportunities by tool and servicing revenue, as follows:

- **Near-term tool revenue:** ASML's existing non-EUV backlog for other markets is large enough to cushion a cutoff for several months. At the end of 2025, ASML's Chinese customers accounted for roughly €7.8 billion in backlog orders, while ASML's overall non-EUV backlog

[76] Exporting a US chip design to Chinese fabs on the Entity List, such as SMIC, would require an export license. Additionally, if a Chinese fab acquired SME illicitly, parties involved in producing a chip with that SME could violate US export controls under General Prohibition 10.

[77] Nikon has minimal sales to China; 80% of Nikon's sales have gone to Intel.[158] Other costs would fall on suppliers such as Gigaphoton and Zeiss.

was €13.6 billion, leaving €5.8 billion in non-China non-EUV orders.[78] Even after a complete cutoff of Chinese DUVi sales, it would take months for ASML to run through this backlog.

- **Longer-term tool revenue:** Medium-term demand substitution would further cushion ASML's equipment revenue. As fabs outside China expand to replace lost Chinese capacity, ASML would likely see increased orders from non-Chinese customers. ASML's CEO Christophe Fouquet has indicated that ASML has mostly worked through its pandemic-era Chinese backlog and expects a rebalancing toward non-China demand.[160] As mentioned above, ASML's former CEO, Peter Wennink, stated that even if China were "excluded from any growth," new fabs would be built elsewhere and ASML's 2030 outlook would not change much.
- **Near-term servicing revenue:** ASML would lose up to roughly $1–1.5 billion per year from its Chinese servicing revenue (3–4% of ASML's annual revenue), or less if servicing is restricted only at advanced fabs.[79] Unlike equipment sales, which ASML can redirect to new fabs outside China, servicing revenue for installed-base tools in China would cease following a ban on servicing.[80]
- **Longer-term servicing revenue:** ASML could fully recoup near-term servicing losses if servicing restrictions make China's DUVi installed base non-functional or unusable. In this scenario, surging global chip demand would stimulate non-Chinese fab expansions and ASML would very likely sell equivalent or better DUVi to these new fabs, and then gain servicing revenue from those tools instead. On the other hand, ASML would not recoup near-term servicing revenue losses if China manages to keep DUVi systems online despite ASML withdrawing from servicing.

Even where ASML or other allied toolmakers suffer losses, the revenue at stake was in significant part enabled only by China's subsidy program, which drives Chinese chipmakers to build more capacity than global market demand warrants. The part of the loss that would not shift to fabs elsewhere comes from the elimination of this subsidy-driven overcapacity as a revenue source.

---

[78] EUV tools accounted for 65% of ASML's €38.8 billion year-end backlog in late 2025, placing the non-EUV (including dry DUV and metrology) portion at roughly €13.6 billion. ASML's CEO, Christophe Fouquet, said on the company's January 28, 2026 earnings call that China accounts for around 20 percent of the current system backlog (i.e., €7.8 billion).[159] Since China cannot purchase EUV tools, that entire sum sits within the €13.6 billion non-EUV backlog, leaving roughly €5.8 billion in non-China non-EUV orders.

[79] ASML 2025Q4 results report "Installed Base Management" sales of €8.2 billion for FY2025, approximately 25% of total net sales.[161]

[80] ASML's 2025 Chinese revenue of €9.52 billion, minus estimated equipment revenue of €8.08 billion (33% of €24.47 billion), implies €1.44 billion in Chinese service and upgrade revenue, approximately $1.63 billion at the 2025 average exchange rate.[162] This estimate is approximate because the 33% share is rounded and covers all Chinese service and upgrade revenue, only a subset of which would be affected by a DUVi-only servicing ban. We therefore estimate the loss from a DUVi servicing ban at roughly $1–1.5 billion.

Appendix A: AI chips made without DUVi are uncompetitive

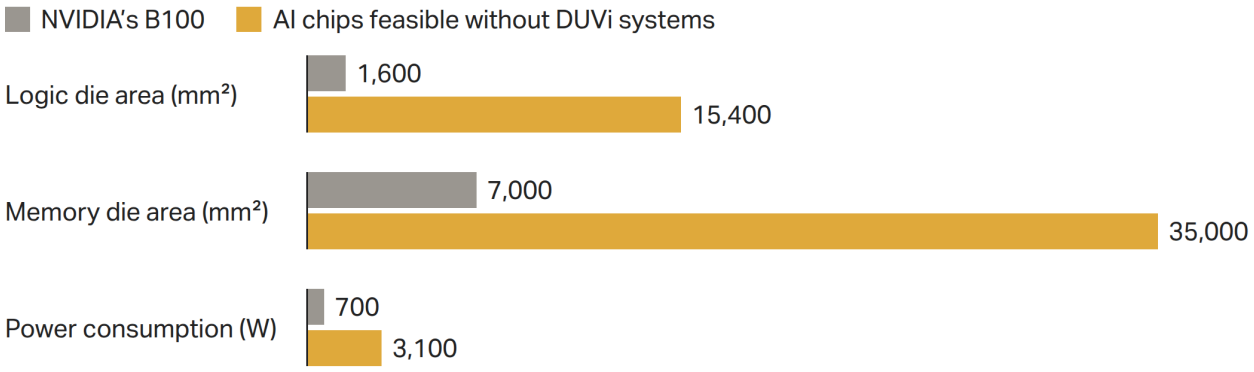
We estimate the silicon area and power required to reproduce the aggregate peak compute throughput and memory bandwidth enabled by modern AI chips, using process nodes that can be implemented without DUVi. We use NVIDIA's announced B100 configuration as a reference, which is manufactured using TSMC's 4NP logic process (a variant of 5nm) and SK hynix's HBM3e, based on the 1β DRAM node. We compare the B100 with hypothetical chips built using 28nm logic and 29nm DRAM, which we separately estimate (Appendix C) to be the most advanced nodes achievable using dry DUV lithography.

Table 5: Comparing leading AI chips with AI chips made without DUV immersion

Metric	NVIDIA's B100 Logic node: TSMC 4NP DRAM node: 1β HBM generation: HBM3e	Equivalent peak FLOPS and bandwidth without DUVi Logic node: 28nm DRAM node: 29nm HBM generation: HBM1
Logic die area (mm^2)	~1,600	~15,400 (9.6×)
Memory die area (mm^2)	~7,000	~35,000 (5×)
Power consumption under a compute-bound workload (W)	700	~3,100 (4.4×)

Figure 4. EUV or DUVi is a hard requirement for making competitive AI chips

Matching B100's FLOP/s and MemBW without DUVi systems consumes far more silicon and power



We estimate requirements to match B100's FP8 FLOP/s and memory bandwidth using 28nm logic and 29nm DRAM/HBM1, process nodes viable without DUVi systems. We scale logic area and power consumption based on transistor densities and iso-performance power consumption across nodes. Memory area and power consumption are derived from HBM1 bandwidth, capacity, bit density, and energy-efficiency characteristics. See Appendix A for assumptions and calculations.

Center for Technology & Statecraft IFP

Logic scaling

Advanced nodes enable logic processors with higher transistor density and improved performance per watt. Transistor density varies between different silicon structures (e.g., high-performance standard cells, SRAM, analog circuitry, and I/O), so the density of a given product depends on its internal composition. We estimate density scaling using two methods:

- Method 1 uses density proxies for high-performance (HP) cell libraries. AI chips devote a large fraction of their area to performance-critical logic implemented with high-drive standard cells. Applying these densities to whole chips is an approximation, because SRAM, analog circuitry, and I/O scale differently from HP logic.
- Method 2 compares densities of NVIDIA GPUs at each node, obtained by dividing published transistor counts by total die areas. This method captures a realistic mixture of logic, SRAM, and I/O, but is potentially confounded by architectural differences between chips.

In addition to transistor density, clock frequency also affects compute throughput per area. Clock frequency depends on a variety of factors, including pipeline design and voltage characteristics, as well as the process node. We assume that a design adapted for 28nm can match the B100's compute clock,

which is supported by high-frequency 28nm arithmetic results.[163] However, this has not been verified and remains a source of uncertainty.

We estimate power scaling using three methods:

- Method 1 backchains TSMC's claimed reductions in iso-performance power efficiency across nodes. This method is internally consistent but reflects scaling for test designs operating at favorable voltages. It also does not fully capture the weaker scaling for SRAM, analog circuitry, and I/O, and assumes that designers use each node to minimize power rather than increase frequency.
- Method 2 estimates the gains realized in product-like implementations. We use data from shipped GPUs and normalize to account for changes in architecture, numerical precision, and voltage/frequency operating point.
- Method 3 compares ports of the same or similar register-transfer level (RTL) across different nodes.

For the chip comparison below, power Methods 2 and 3 provide the more appropriate estimate.

Table 6: Methods for comparing chips

Logic node	Transistor density		Perf / watt vs. TSMC 4NP		
	Method 1: HP Transistor Density (MTr/mm²)	**Method 2: NVIDIA GPU densities**	**Method 1: TSMC's claims[81]**	**Method 2: Realized products[82]**	**Method 3: Vendor-published IP comparisons[83]**
TSMC 28HP	10.6	12.93 (pooled 4 GPU)	0.10	0.17	0.2
TSMC 12FFN	28[84]	25.89 (V100)	0.26	0.48	0.42
TSMC N7	65[171]	65.62 (A100)	0.57	0.75	0.66
TSMC N5	92[85]	—	0.82	Unknown	0.85
TSMC 4NP	98[86]	130 (B100)	1	1	1

[81] We backchain as follows. (1) 4NP/N5: TSMC reported N4P, our proxy for 4NP, had a 22% power efficiency improvement over 5nm.[164] (2) N5/N7: SemiWiki claims "5nm uses 30% less power than 7nm."[165] This corresponds to a 1.43× improvement in power efficiency. (3) N7/12FFN: TSMC N7 is reported to give a 55% power efficiency improvement over 16/14nm; 12FFN is separately reported to have similar power characteristics to 16/14nm.[166] (4) 28HP/12FFN: TSMC claims: "16FF consumes 50% less power at the same speed" than 20SoC, and 20SoC has "30 percent higher speed, 1.9 times the density, or 25 percent less power than its 28nm technology."[167] Overall multiple is 0.75 × 0.5 = 0.375.[168]

[82] Using published specifications for NVIDIA's K40, V100, A100, H100, and B200 chips, we calculate board-level performance per watt and then estimate the portion attributable to the node by adjusting for changes in precision, arithmetic throughput per clock and transistor, and operating point on the frequency/voltage curve (calibrated using different power configurations of the same silicon architecture).

[83] Based on several public estimates: (1) Arm's implementation of Cortex A72 on both 28nm and 16/14nm, (2) Arm's comparison of Neoverse N1 on 7nm with N2 on 5nm, (3) Google's decomposition of TPU v3–v4 efficiency improvement into contributions from design and process technology.[169]

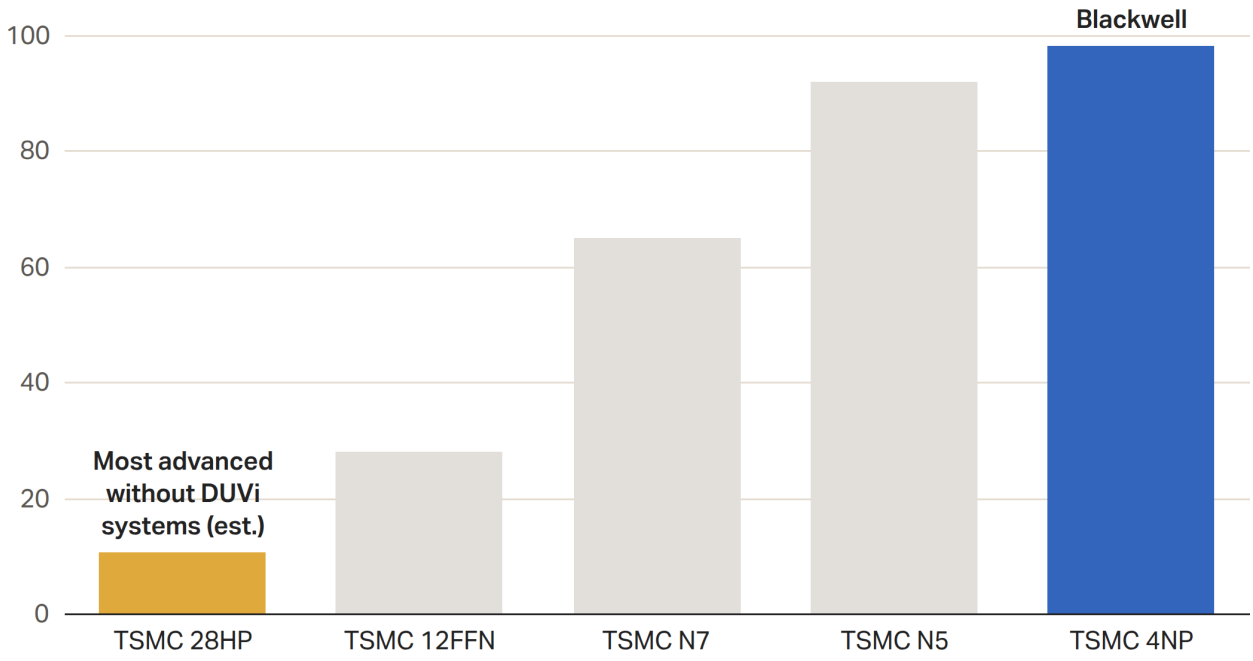
[84] 12nm has 33.8 HD.[170] Scaling by a typical HP/HD ratio, we get ~28.

[85] From Angstronomics: "H280g57 gives a logic density of 92.3Mtr/mm² for 3-fin N5."[172] This is notably much lower than TSMC's marketing claims.

[86] TSMC's N4P, which we use as a proxy for NVIDIA's custom 4NP, offers a 6% density increase over N5.[173]

Figure 14. TSMC's Blackwell node 4NP provides ~10× higher transistor density than 28nm

HP transistor density (MTr/mm²)



Note: Since TSMC does not publish absolute transistor density figures, we derive them from third party estimates and TSMC's reported node over node density scaling.

Center for Technology & Statecraft IFP

DRAM scaling

Advanced DRAM nodes enable HBM with improved capacity, bandwidth, and power efficiency. Higher bit density increases capacity per unit silicon area, which particularly benefits HBM by offsetting the area occupied by TSVs and peripheral circuitry.[87] Lower energy per bit reduces the stack's power consumption and overall thermal load,[88] enabling higher stack heights, increased capacity, and higher sustained bandwidth before the stack reaches thermal limits.[175] Finally, faster transistor switching speeds enable higher per-pin signaling rates and therefore higher bandwidth.

[87] For example, SemiAnalysis found that SK hynix D1z DDR4 achieved 0.296 Gb/mm² versus 0.16 Gb/mm² for the D1z HBM3.[174]

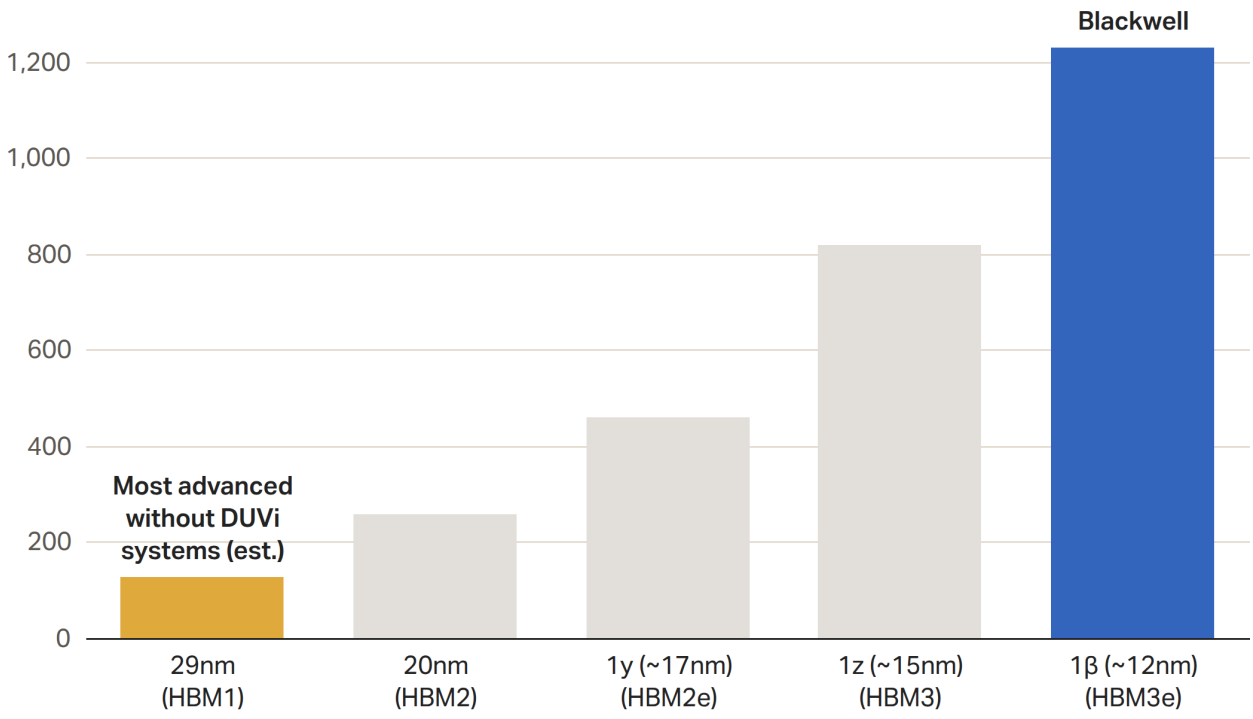
[88] Thermal load is a critical constraint for HBM because its DRAM dies stack directly atop one another, with no airflow between layers, so heat generated by the bottom dies must conduct upward through every die above it to reach the package's heat sink.

Table 7: Comparing DRAM and HBM generations

DRAM node	HBM generation	Bandwidth per stack (GB/s)	Capacity per stack (GB)
29nm	HBM1	128	4[89]
20nm	HBM2	256	8
1y (~17nm)	HBM2e	460	16
1z (~15nm)	HBM3	819	24
1β (~12nm)	HBM3e	1229	24 (B100); up to 36

Figure 15. Blackwell's HBM3e provides ~10× more bandwidth than HBM made on 29nm DRAM

Bandwidth per stack (GB/s)

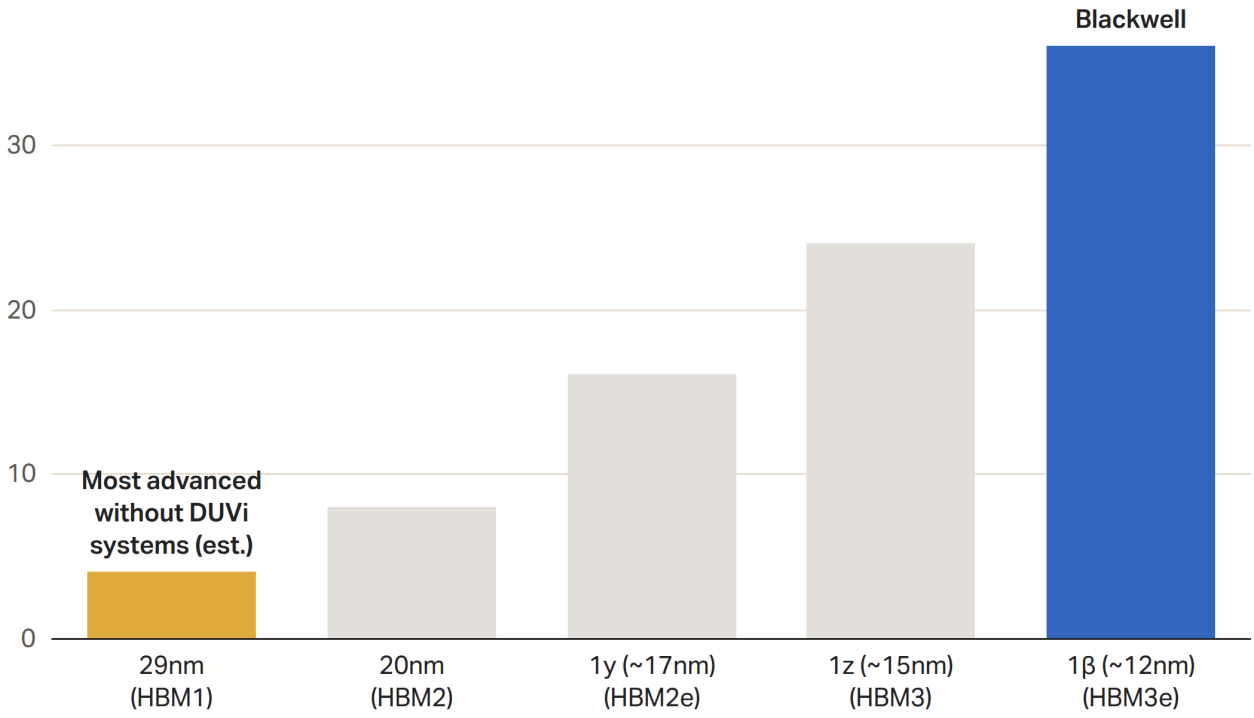


Center for Technology & Statecraft IFP

[89] SK hynix's original HBM1 product shipped with 1GB per stack, but the HBM1 specification supports up to 4GB/stack.

Figure 16. Blackwell's HBM3e provides ~9× more capacity than HBM made on 29nm DRAM

Capacity per stack (GB)



Center for Technology & Statecraft IFP

Matching B100 FLOPS and Bandwidth using 28nm logic and HBM1

We estimate the logic die area, memory die area, and power consumption required to match the B100's peak processing performance (3,500 TFLOPS at dense FP8) and memory bandwidth (8 TB/s).

- **Logic die area**: We estimate a 9.6× density penalty, averaging the 4NP-28nm ratios implied by density Methods 1 and 2. Assuming a similar compute architecture, area composition, and clock speed, matching the B100's 3,500 TFLOPS of dense FP8 compute (~2.2 TFLOPS/mm^2 across an estimated ~1,600 mm^2) therefore requires 15,400 mm^2 of logic silicon, distributed across 20 dies of up to 800 mm^2 each.
- **Memory die area**: Matching the B100's 8 TB/s bandwidth requires 62.5 HBM1 stack-equivalents at 128 GB/s per stack. The per-stack bandwidth already incorporates interface signaling speed, so no additional memory-clock multiplier is applied. We model HBM1 at 4 GB/stack, which is supported by the HBM1 specification, rather than the initial 1 GB/stack

configuration.[90] At an assumed effective HBM1 density of ~0.057 Gb/mm², [91] the resulting 250 GB of memory requires ~35,000 mm² of DRAM die area. This notably corresponds to ~250GB of capacity, 1.3× more than the B100.

- **Power consumption**: We assume a compute-bound workload and decompose the B100's 700 W TDP using measurements from the H200, whose six HBM3e stacks consume ~150 W under a 700 W compute-bound workload.[177] Scaling to the B100's 8 HBM3e stacks implies ~200 W of HBM power and ~500 W for the logic dies. We scale logic power consumption by 5.44×, the midpoint of the power penalties implied by Methods 2 and 3 above, giving 2,720 W. For memory, we assume a 1.73× compute-bound power penalty, using 7 pJ/bit for HBM1 and 4.05 pJ/bit for HBM3e as a proxy. Scaling the same 200 W baseline gives ~346 W of memory power and ~3,066 W in total, rounded to ~3,100 W, or 4.4× the B100's TDP.

Caveats

(1) This comparison calculates the number of legacy AI chips needed to achieve the B100's FLOPS and memory bandwidth, but aggregating legacy chips may not necessarily translate to equivalent system-level performance. This is because a network of less powerful legacy chips would require substantially more networking complexity than a B100, introducing power, latency, packaging, and system availability challenges.

(2) Logic die area is scaled based on transistor density while assuming a similar clock speed. There is some evidence that 28nm chips could be clocked to match the B100, but this is not well established. If clock speed is reduced, it will increase the required die area proportionally. Conversely, Section 7 notes that NXT:1470 could plausibly reach 20nm-class logic. TSMC's stated 1.9× density gain from 28nm to 20nm (footnote 81) would cut the logic-area penalty to roughly 5×.

(3) We assume 29nm DRAM enables HBM1 because HBM1 was historically demonstrated at this DRAM node. But the DRAM node does not uniquely limit HBM bandwidth. Advanced signaling, peripheral circuitry, or packaging could enable 29nm DRAM to achieve greater bandwidth per stack than historical HBM1.

[90] This is a hypothetical configuration: no 4 GB HBM1 stack ever shipped.

[91] SK hynix confirmed that 29nm HBM1 shipped as a 2Gb die measuring 35mm².[176]

Appendix B: Technical challenges of developing advanced lithography

The challenge of developing an advanced lithography system is threefold. First, a large number of specialized subsystems must be developed to an extremely high standard of performance. A few of the most challenging subsystems[92] are as follows and shown in Figure 17:

- The **light source** generates the 193nm DUV light used to print patterns on the wafer and must fire high-power pulses thousands of times per second while keeping wavelength, bandwidth, and delivered dose extremely stable over billions of pulses.[93]
- The **illumination optics subsystem** shapes and homogenizes the DUV light before it reaches the mask in order to achieve the desired balance between resolution, depth of focus, and process window.[94]
- The **projection optics subsystem** projects the patterned light from the mask onto the wafer and must maintain extremely low optical distortion and aberration.[95]

---

[92] There are many additional subsystems and components beyond those listed here, including: mask and reticle handlers, wafer handlers, in-scanner metrology, thermal management, lens manipulators, vibration dampening and structural control, wafer clamping, reticle clamping.

[93] The state-of-the-art technology is an Argon Fluoride (ArF) excimer laser, which generates light by exciting a high-pressure ArF gas mixture using a high-voltage electrical discharge. Challenges in building the light source subsystem include: (i) achieving a spectral bandwidth narrowly clustered around 193nm, which requires compressing the native excimer gain spectrum by over three orders of magnitude, (ii) sustaining energy pulses at a rate of multiple kilohertz while clearing thermal shocks and acoustic resonances from the chamber, (iii) maintaining energy dose stability within a few percent, (iv) containing the reactive fluorine-bearing gas mixture, which naturally tends to degrade the chamber and generates metal-fluoride debris, (v) mitigating DUV-induced color-center formation and compaction in the optics internal to the light source (including chamber windows, line-narrowing-module prisms, and the grating), (vi) synchronizing the master oscillator and power amplifier chambers with sub-nanosecond precision, (vii) extending the lifetime of the discharge chamber against electrode erosion, window damage, and particulate accumulation.

[94] Challenges in building the illumination optics subsystem include: (i) generating freely programmable pupil shapes—including annular, dipole, quadrupole, and computationally optimized freeform sources—at high optical efficiency, (ii) maintaining sub-percent intensity uniformity across the slit-shaped illumination field at the reticle, (iii) controlling polarization purity, including delivering the linear and azimuthal polarization states required for high-NA imaging, (iv) implementing pulse-by-pulse dose feedback to the laser to hold dose stability during scan, (v) suppressing speckle and coherence artifacts inherent to excimer-laser illumination, (vi) sustaining 193nm power without inducing color-center formation, compaction, or thermal distortion in the constituent fused-silica and CaF_2 optics, (vii) maintaining beam pointing and alignment between laser and scanner over years of operation despite mechanical and thermal drift.

[95] Challenges in building the projection optics subsystem include: (i) sourcing raw optical materials, typically synthetic fused silica or single-crystal CaF_2, and growing large CaF_2 boules with low intrinsic birefringence, a slow and low-yield process, (ii) polishing lens surfaces to sub-nanometer RMS error, which can take months, (iii) color-correcting a lens with only two 193nm-transparent materials, which forces an ultra-narrow laser and motivates mirror-lens hybrid designs, (iv) preserving the polarization of the DUV light at steep convergence angles, (v) mitigating DUV-induced lens deformations, including compaction, rarefaction, and color-center formation, which can distort the wavefront over billions of pulses of integrated exposure, (vi) compensating for lens heating, which can distort the wavefront at the sub-nanometer level, in real time, which requires embedded actuators that deform, translate, or tilt individual elements on-the-fly, and (vii) maintaining mechanical, thermal, and contamination stability across the full optical column so that nanometer-scale alignment and 193nm transmission are preserved over the tool's production lifetime.

- **Reticle and wafer stages** move the mask and wafer in precise synchronization during each exposure and must achieve sub-nanometer positioning accuracy while accelerating rapidly enough to enable high throughput.[96]
- **Alignment sensors** ensure that each layer is correctly registered to the layer beneath it, requiring them to correctly measure and correct for tiny positional errors on the wafer in real time.[97]
- **Leveling and focus sensors** ensure the projected image stays in focus by measuring wafer height, topography, and deformation in real time.[98]
- The **immersion hood** maintains the ultra-pure water layer between the final lens and the wafer and must eliminate bubbles, escaped droplets, contamination, thermal non-uniformity, and wafer-edge defects, without sacrificing throughput.[99]

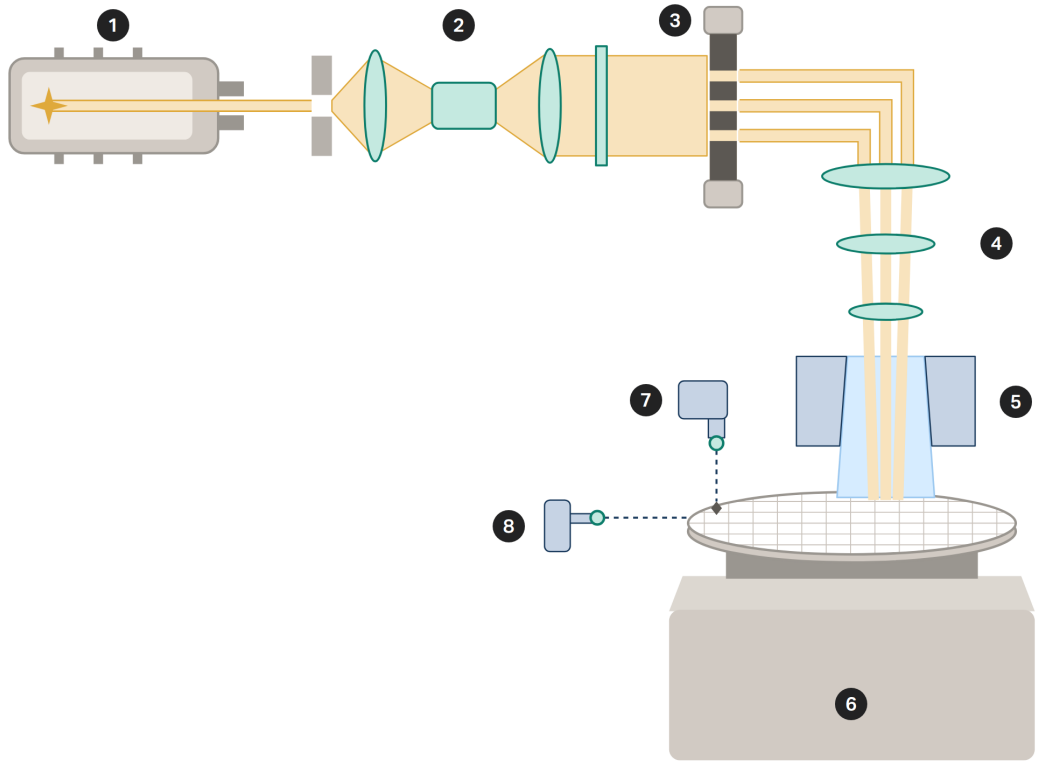
[96] Challenges in building the stage subsystems include: (i) maintaining sub-nanometer positioning accuracy along with high accelerations (16 g for the reticle stage, 7 g for the wafer stage), (ii) synchronizing the trajectories of the reticle and wafer, accounting for the fact that the reticle must move four times faster than the wafer to compensate for the 4× demagnification of the projection optics, (iii) canceling the reaction forces from these accelerations via balance-mass architectures so that no vibration is injected into the exposure point or the fab floor, (iv) measuring stage position to sub-nanometer precision despite air-pressure, temperature, and humidity drift, (v) thermally managing the high-power linear motors so their dissipation does not distort the stages or the surrounding metrology, (vi) clamping the wafer flat to nanometer uniformity without inducing in-plane distortion, (vii) preventing reticle slippage on the reticle stage under 16 g acceleration.

[97] Challenges in building the alignment sensor subsystem include: (i) determining the position of alignment markings that may lie buried beneath several process layers, (ii) handling mark asymmetry and deformation introduced by upstream processing (e.g., CMP, etch), (iii) preserving signal contrast across diverse film stacks, which requires employing multiple illumination wavelengths and diffraction orders, (iv) obtaining enough measurements per wafer to fit a high-order wafer-grid model—capturing translation, rotation, magnification, and higher-order distortion—without reducing throughput, (v) calibrating the geometric relationship between the sensor, the projection lens, and the reticle stage when computing exposure corrections, (vi) maintaining long-term stability of the sensor's optics and electronics, avoiding contamination, drift, and thermal effects.

[98] Challenges in building the level and focus sensor subsystem include: (i) measuring wafer height with nanometer precision across a 300 mm wafer given an immersion depth of focus on the order of 100nm, (ii) generating a complete wafer height map fast enough to support high throughput (>275 wafers per hour in modern DUVi systems), (iii) compensating for wafer warp, bow, and high-topography features (e.g., 3D NAND structures) so the projected image stays within the depth of focus during the scan, (iv) handling process-layer-induced reflectivity variations that distort optical height readings, (v) feeding height data forward to the wafer stage in real time in order to maintain focus across the moving exposure slit, (vi) managing thermal deformation of the wafer during exposure, which can change wafer shape mid-scan.

[99] Challenges in building the immersion-fluid subsystem include: (i) confining the water film against the wafer at scan speeds of hundreds of millimeters per second without leaving droplets, films, or stains, (ii) preventing the formation of bubbles in the film, which scatter 193nm light and cause printing defects, (iii) supplying ultra-pure, degassed, and deionized water at the high flow rates needed to continuously replenish the film and carry away absorbed heat, (iv) maintaining sub-millikelvin thermal uniformity in the film, as even small temperature gradients can change the refractive index, distorting imaging, (v) preventing photoresist components from leaching into the water (which contaminates the lens) and water components from leaching out of the resist, (vi) managing wafer-edge and bevel effects, where the film must transition cleanly between the wafer and surrounding hardware, (vii) ensuring the immersion hood seals reliably during the rapid stage accelerations involved in scanning.

Figure 17. Key subsystems in a DUVi system



1 Light source:
Generates 193nm DUV light pulses

2 Illumination optics:
Shapes and homogenizes the light before the reticle

3 Reticle stage:
Positions the mask during exposure

4 Projection optics:
Projects the patterned light onto the wafer

5 Immersion hood:
Maintains the ultra-pure water layer between the final lens and wafer

6 Wafer stage:
Moves the wafer during scanning exposures

7 Alignment sensor:
Measures wafer marks for layer-to-layer registration

8 Leveling / focus sensor:
Measures wafer height and topography to keep the image in focus

Second, these subsystems must be made to work together as a single machine. ASML does not publicly disclose a parts count for DUVi systems, but its EUV machines, which are built on the same TWINSCAN platform, contain about 100,000 individual parts.[178] Coordinating a machine of this complexity requires an elaborate control system, including dozens of networked microprocessors,

hundreds of actuators, hundreds of thousands of sensors, and millions of lines of code—a codebase ASML says has been under continuous development for 30 years.[179] Complexity of integration scales with the number of components, as changes to one subsystem affect the performance of others. For example, in a DUVi system, the water touches both the lens and the photoresist coating on the wafer, so chemicals from the resist can leach into the water and contaminate the lens.[180] Solving this problem required coordinated industry co-development of the projection optics, immersion hood, water-purification system, and new low-leaching resists and protective coatings.[181]

Third, the integrator must coordinate the enormous network of research, engineering, manufacturing, and transportation required to build machines at production scale. ASML says that roughly 85% of the components in its machines are manufactured by its network of over 5,000 suppliers.[182] At least 700 of these contribute to advanced lithography and 200 contribute components that ASML considers critical and single-source.[183] Some of these partners are effectively irreplaceable. Most notably, Carl Zeiss SMT has been ASML's sole supplier of lenses, mirrors, illuminators, collectors, and other critical optics.[184] US company Cymer[100] and Japanese company Gigaphoton are the dominant suppliers of DUV light sources. Other specialized subcomponents are drawn from specialists scattered across the Netherlands, Germany, the US, the U.K., Switzerland, and Japan. ASML serves as the master integrator. It aligns suppliers' technical roadmaps and production capacity, manages the logistics around component shipments, and coordinates the technical integration itself. ASML treats management of its supplier network as a core capability. Former Chief Strategy Officer Frits van Hout described the company as engineering its supplier relationships "like a machine," and ASML has gone as far as providing capital injections or acquiring suppliers outright[101] when necessary to preserve critical supply capability and production capacity.[186] A non-exhaustive list of key supplier relationships is provided in Table 8:

[100] ASML acquired Cymer in 2013.[185]

[101] Examples include ASML's acquisitions of Cymer in 2013 and Berliner Glas in 2020.

Table 8: Examples of publicly documented producers, co-developers, and component suppliers for key DUVi lithography subsystems

Function	Description	Example suppliers (rest of world ecosystem)	Example suppliers (Chinese ecosystem)
System integration	Assembly, qualification, calibration, and delivery of the complete DUVi system.	ASML, Nikon	Shanghai Aishengna, SMEE, Shanghai Yuliangsheng[102]
Light source	Generates 193nm light pulses with wavelength stability, bandwidth control, and dose uniformity.	Cymer, Gigaphoton Coherent (sensors, internal optics), Linde (ArF gas)	Keyi Hongyuan
Illumination and projection optics	Illumination optics shapes the laser beam into a precisely controlled illumination field at the reticle plane; projection optics projects the reticle pattern onto the wafer.	Zeiss, Nikon Corning (CaF_2 crystals and optical glass blanks), Hellma Materials; Fraunhofer IOF (DUV coatings), Heraeus (synthetic fused silica)	MLOptic, Shanghai Institute of Optics and Fine Mechanics (SIOM), Changchun Institute of Optics (CIOMP), Beijing Guowang Optics
Reticle stage	Carries the reticle and accelerates it in synchronization with the wafer stage during exposure.	ASML, Nikon ASML Berlin (optical modules), Zygo, Renishaw	Beijing U-Precision Tech
Wafer stage	Positions wafers under the projection lens.	ASML/VDL ETG Zygo, Renishaw, Keysight	Beijing U-Precision Tech Chongqing Millison Technologies (bases and vacuum chambers)
Leveling and focus sensors	Maps the wafer surface topography in real time, feeding height (Z) and tilt data to the wafer stage so the resist surface is held precisely at the focal plane.	ASML, TNO	SMEE
Immersion hood	Manages the ultra-pure water meniscus between the final projection lens element and the wafer surface.	ASML, Nikon; TNO, DFS, Exentec, Entegris	Zhejiang Qi'er Electromechanical, SMEE, Yuliangsheng

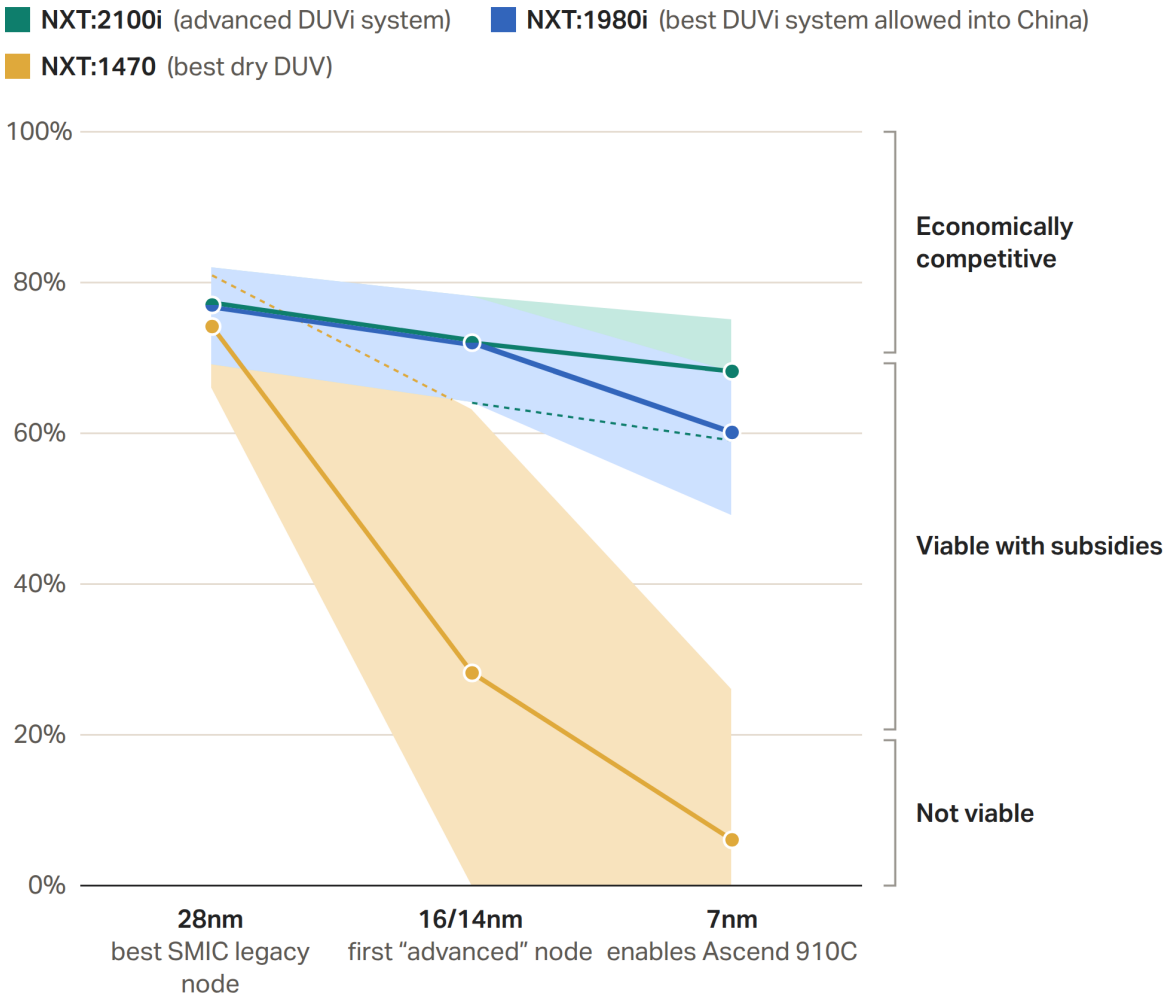
[102] Shanghai Aishengna absorbed teams from Yuliangsheng and SMEE, but the latter two entities still exist.

Appendix C: DUV frontend logic yields

We estimate perfect die yields, not accounting for salvaged dies, achievable using a range of DUV systems at various logic and memory nodes, assuming a mid-sized (500 mm²) logic die and a mature production process at an industry-leading fab (e.g., TSMC for logic, SK hynix for DRAM) with state-of-the-art SME (etch, deposition, etc.) outside of lithography. The reported ranges are for the P10, P50, and P90 of the distributions produced by running Monte Carlo simulations sampled from various input distributions, described below. See Figure 10.

Figure 10. NXT:1980i and more advanced DUVi systems enable Ascend AI chips

Perfect die yields for a 500mm² logic die and a mature manufacturing process



Shaded bands span P10–P90 estimates. See Appendix C for assumptions and calculations.

Table 9: ASML DUVi system yield by logic node

NXT tool	Notes	Logic (500mm²)		
		28nm (best SMIC legacy node)	16/14nm (first "advanced" node)	7nm (Ascend 910C)
NXT:2100i	Best DUVi in China	[69, 77, 82]	[64, 72, 78]	[59, 68, 75]
NXT:1980i	Best DUVi allowed into China under current rules	[69, 77, 82]	[64, 72, 78]	[49, 60, 68]
NXT:1470	Best dry DUV	[66, 74, 81]	[0, 28, 63]	[0, 6, 26]

Caveats

Die yield is among the most closely guarded figures in semiconductor manufacturing and depends vitally on proprietary, undisclosed process knowledge. Published figures are scarce and rarely provide context that supports interpretation. Reconstructing yield for a particular node and tool therefore requires assembling evidence from a variety of disparate sources, none of which was intended to enable analysis. The method presented here draws what conclusions the limited public record will support.

We have quantified some uncertainty using Monte Carlo techniques. However, this does not capture uncertainty in the structure of the model itself or in the applicability of the anchoring sources to the processes modelled.

Overall, this analysis should be read as providing limited direct evidence for three claims, supplementing indirect evidence that can be inferred from the industry's historical record:

(1) NXT:2100i, NXT:1980i, and NXT:1470 are all viable at 28nm.
(2) NXT:2100i and NXT:1980i are viable through at least 7nm.
(3) NXT:1470 is almost certainly not viable at 7nm and of limited viability at 16/14nm.

The exact numerical differences between individual cells warrant less weight than these high-level conclusions.

Methodology

We model frontend yield using a Monte Carlo simulation with 100,000 draws. Inputs with uncertainty are specified using quantiles (P10, P50, P90) and fitted using a log-normal distribution for positive quantities and on the logit scale for decimal quantities. Each draw samples every input once and applies the same shared draw across all node and tool combinations.

We model yield as the product of wafer-level line yield and three die-yield factors. Y_{line} and $Y_{nonpatterning,\ residual}$ are not affected by lithography tool quality, while Y_{random} and $Y_{patterning}$ vary with tool quality.

$$Y = \underbrace{Y_{\text{line}}}_{\text{Line yield}} \times \underbrace{Y_{\text{nonpatterning, residual}}}_{\text{Residual non-patterning yield}} \times \underbrace{Y_{\text{random}}}_{\text{Random defect-limited yield}} \times \underbrace{Y_{\text{patterning}}}_{\text{Patterning limited yield}}$$

Line yield
the fraction of wafers that complete wafer fabrication and can be available for die sort

Residual non-patterning yield
systematic and parametric die loss caused by non-patterning process variation or design-process interactions

Random defect-limited yield
die loss due to random, non-systematic killer defects

Patterning limited yield
die loss caused by failure to form patterned features within their allowed geometric tolerances

⌙ Vary with tool quality

Modeling Y_{line}

Y_{line}, line yield, is the fraction of wafers that complete wafer fabrication and can be available for die sort. Line-yield losses include broken or scratched wafers, unrecoverable equipment failures, or gross misprocessing. The general relationship is:

$$\underbrace{Y_{\text{line}}}_{\text{Line yield}} = f_{\text{line}}(\underbrace{PL_{\text{node,litho}}}_{\text{The effective process length required to implement the node using the lithography system}}, \underbrace{F}_{\text{Accounts for the node implementation's reliability, automation, process control, operator proficiency, and excursion containment}})$$

Line yield
the fraction of wafers that complete wafer fabrication

The effective process length required to implement the node using the lithography system

Accounts for the node implementation's reliability, automation, process control, operator proficiency, and excursion containment

We approximate line yield using node-specific constants:

Table 10: Line yield by node

Node	Y_{line}
28nm	[95, 98, 99]
16/14nm	[93, 97, 99]
7nm	[93, 97, 99]

A 2002 survey of 10 high-performing fabs (CSM-52, Figure 3-1) reported roughly 98% line yield per 20 mask layers; a 2010 NAND study independently estimated 98% (though this finding is given slightly less weight as the study was for NAND memory).[187] We estimate modestly lower yields because advanced-node processes have many more layers (e.g., up to 80 for 7nm with immersion), creating more opportunity for losses. On the other hand, improvements since 2023 in automation, reliability, process control, and excursion containment likely counteract this effect. Balancing both considerations and factoring in an estimate from a semiconductor industry expert, we use [95, 98, 99] at 28nm and [93, 97, 99] at the more advanced nodes. Because we do not mechanically scale line yield with every mask equivalent, this approach may slightly overestimate yield for long processes, particularly implementations of advanced nodes using dry DUV tools.

Modeling $Y_{\text{nonpatterning, residual}}$

$Y_{\text{nonpatterning, residual}}$, residual non-patterning yield, accounts for systematic and parametric die loss caused by non-patterning process variation or design-process interactions not already counted in random, patterning, or line yield. Examples include transistor-parameter variation, implant and doping variation, and wafer-edge integration losses not attributable to pattern placement. A mature fab can eliminate most of these failures, but others are unfixable. The general relationship is:

$$Y_{\text{nonpatterning,residual}} = f_{\text{nonpatterning}}(A, n, p, F)$$

- $Y_{\text{nonpatterning,residual}}$: Residual non-patterning yield — systematic and parametric die loss caused by non-patterning process variation or design-process interactions
- A: Die area
- n: The process node
- p: The product design and its sensitivity to process variation
- F: Fab capability and process maturity

We approximate this as a constant with quantiles:

$$\underbrace{Y_{\text{nonpatterning, residual}}}_{\text{Residual non-patterning yield}} \approx \underbrace{K_{\text{nonpatterning, residual}}}_{\text{A constant with quantiles}} = \underbrace{[88, 94, 99]\,\%}_{\text{Quantiles (P10, P50, P90)}}$$

Across the 11 devices in Table 3 of the SEMATECH-sponsored survey, systematic-limited yields are approximately [70, 87, 96]%.[188] These include lithography losses modeled separately here. An industry expert familiar with the program reported that its population also included DRAM, whose more regular, redundant structure may yield better than a large logic die. After separating patterning losses, allowing for improved process control, and biasing moderately upward to reflect the assumption of a world-leading fab, we use [88, 94, 99]%.

Modeling Y_{random}

Y_{random}, random defect-limited yield, accounts for die loss due to random, non-systematic killer defects, such as particles, residues, voids, and localized opens or shorts, whose incidence increases roughly in proportion to the length of the manufacturing process. We model Y_{random} using a negative binomial, reflecting standard modeling approaches from the literature:[189]

$$\underbrace{Y_{\text{random}}}_{\substack{\text{Random defect-limited yield die loss due to random, non-systematic killer defects}}} = f_{\text{random}}(A, D_0(PL_{\text{node,litho}})) = (1 + \underbrace{A}_{\substack{\text{The die's critical area}}} \times \underbrace{D_0}_{\substack{\text{The effective random killer-defect density } (D_{0,\text{random}})}}(\underbrace{PL_{\text{node,litho}}}_{\substack{\text{The effective process length required to implement the node using the lithography system}}}) / \underbrace{\alpha}_{\substack{\text{A clustering factor, which is set to 2, following convention to reflect moderate clustering}}})^{-\alpha}$$

We estimate A as [400, 450, 475] mm², less than the physical 500mm² because some area is allocated to repairable SRAM,[103] which we estimate occupies 5–20% of the overall die.

[103] This includes large shared caches, occupying about 5% of the die for AI processors (e.g., the 60 MB L2 on NVIDIA's GH100) plus most of the memory inside compute blocks, following the long-standing convention that large on-die arrays carry spare rows and columns for repair.[190]

We model $PL_{node,litho}$ as:

$$PL_{\text{node,litho}} = M_{\text{base}} + S_{\text{base}} + \Delta M_{\text{litho}} + \Delta S_{\text{litho}}$$

- $PL_{\text{node,litho}}$: The effective process length required to implement the node using the lithography system
- M_{base}: A physical mask count of a known reference process
- S_{base}: The number of spacer-patterning cycles (e.g., SADP) in the reference process, where we assume each spacer cycle adds one mask-equivalent process content
- ΔM_{litho}: The additional masks required when using a less-capable lithography tool due to additional multipatterning steps
- ΔS_{litho}: The additional spacer-patterning cycles required when using a less-capable lithography tool due to additional multipatterning steps

We assume NXT:2100i and NXT:1980i have the same process length because both have the same 38nm resolution. NXT:1470 has a looser resolution (57nm) and thus requires additional multi-patterning steps above and beyond NXT:2100i/1980i, increasing process length. These additional steps include (i) pitches between 76–114nm, which can be single-patterned with DUVi but require SADP using the NXT:1470, and (ii) pitches between 38-57nm, which can be patterned with SADP using DUVi but which require another pitch multiplication using NXT:1470.

Table 11: Estimated effective process lengths for logic nodes and NXT tools

Node	NXT:2100i/1980i	NXT:1470
28nm	[45, 50, 60][104]	[51, 59, 73][105]
16/14nm	[63, 69, 77][106]	[71, 82, 96][107]
7nm	[82, 88, 98][108]	[98, 108, 122][109]

We anchor estimates of $D_{0,random}$ on 7nm with immersion, for which we estimate $D_0 \approx$ [0.05, 0.06, 0.09]. P90 reflects public reporting that TSMC's N7 process launched in 2018 using NXT:1980i and ramped to $D_0 \approx 0.09$ defects/cm² within three quarters.[199] P50 reflects benchmarking at older nodes, which found the best fabs achieved a D_0 of 0.06, and the likelihood that TSMC's D_0 continued to decline.[200] P10 reflects the lower bound of recent analyst estimates claiming $D_0 \approx$ 0.05–0.15 for a mature node.[201]

We linearly scale D_0 to other nodes and to processes using NXT:1470, deriving the following yields:

[104] We estimate M_{base} at [45, 50, 58]. P50 is anchored to the eBeam Initiative's 50-mask average for the 22–32nm range.[191] P10 reflects the survey's 49-mask average at 32nm. P90 reflects its 57-mask average at 22nm and accommodates the taller eleven- and twelve-level metal stacks available in the 28nm family. We estimate S at [0, 0, 2], since the ~90nm minimum metal pitch sits above the 76nm immersion resolution, with two SADP allowed in P90 for aggressive variants.

[105] We estimate ΔM_{litho} at [5, 6, 8]. The P50 matches ASML's estimate of six critical lithography steps at 28nm, while P10 and P90 add modest room on either side.[192] We estimate ΔS_{litho} at [0, 3, 6], assuming each affected layer adds one spacer cycle if it is line patterned and none if it is hole patterned.

[106] We estimate M_{base} at [60, 65, 72]. The P50 comes from the eBeam Initiative's 65-mask average for ground rules below 16/14nm and at or above 11nm.[193] The P10 reflects the 60-mask average reported at 16/14nm in an earlier survey.[194] The P90 reflects additional metal layers for implementations with taller metal stacks. We estimate S at [3, 4, 5], with P50 reflecting one spacer cycle for the 48nm fin pitch and three for the 64nm 1x metal levels.[195]

[107] We estimate ΔM_{litho} at [4, 6, 9]. The P50 estimate assumes six layers require additional patterning: the 48nm fin pitch, which falls in the 38–57nm window and moves from SADP to SAQP, and the 90nm gate pitch plus four 80–90nm metal and via levels, which fall in the 76–114nm window and move from single patterning to SADP.[196] We estimate ΔS_{litho} at [4, 5, 6]. P50 assumes one of the additional cycles is litho-etch-litho-etch (LELE). The P10 and P90 vary the number of 80–90nm levels in the stack and the split between LELE and self-aligned patterning.

[108] We estimate M_{base} at [75, 80, 88]. P50 is estimated at 80 based on public reporting.[197] P10 is anchored to the eBeam Initiative's 76-mask average for ground rules below 11nm and at or above 7nm (P50 sits above that figure because an all-DUV implementation forgoes the mask reductions available with EUV).[198] P90 reflects taller metal stacks. We estimate S at [7, 8, 10]. P50 is estimated at 8, reflecting two cycles for the ~30nm fin pitch, one for the ~57nm gate pitch, and five for the M0–M4 levels. The P10 estimate assumes that only M0–M3 are at 40nm (S=7), while P90 assumes that M0–M6 are at 40nm (S=10).

[109] We estimate ΔM_{litho} at [8, 10, 12] and ΔS_{litho} at [8, 10, 12]. At P50, we estimate M0–M4 at 40nm fall in the 38–57nm window and move from SADP to SAQP, while M5–M9 at 76nm fall in the 76–114nm window and move from single patterning to SADP, each adding one mask and one spacer cycle. The P10 estimate removes the top two of M5–M9. The P90 adds the 57nm gate pitch and the ~30nm fin pitch.

Table 12: Random defect-limited yield for a 500mm^2 logic die

Node	Y_{random}	
	NXT:2100i/1980i	NXT:1470
28nm	[80, 86, 89]	[77, 84, 87]
16/14nm	[74, 82, 85]	[70, 79, 83]
7nm	[69, 77, 82]	[64, 73, 78]

Modeling $Y_{patterning}$

$Y_{patterning}$, patterning limited yield, accounts for die loss caused by failure to form patterned features within their allowed geometric tolerances. It includes overlay errors between layers and between multiple exposures within a layer, as well as residual non-overlay errors arising from critical-dimension variation, focus and dose variation, and pattern-specific imaging limitations, among others. We model $Y_{patterning}$ as:

$$Y_{\text{patterning}} = Y_{\text{overlay}} \times Y_{\text{residual, non-overlay}}$$

- $Y_{\text{patterning}}$: Patterning limited yield — die loss caused by failure to form patterned features within their allowed geometric tolerances
- Y_{overlay}: Accounts for loss due to overlay
- $Y_{\text{residual, non-overlay}}$: Accounts for loss due to other residual imaging limitations

$Y_{overlay}$ is modeled as the magnitude of two independent Gaussian-distributed registration errors, reflected by a Rayleigh distribution:[202]

$$Y_{\text{overlay}} = [1 - P(O, S)]^{M(A, L)}$$

Accounts for loss due to overlay

A constant determining the shape of the failure probability distribution

$$P(O, S) = e^{[-\kappa(S/O)^2]}$$

The overlay-failure probability for a reference die, node, and lithography system

S is the reference node's maximum permissible overlay error; *O* is the effective on-product overlay achieved by the reference lithography system

$$M(A, L) = [\sqrt{(A/A_0)} \times (L/L_0)]$$

The effective error opportunity multiplier relative to the reference

A is die area, while A_0 is the area of the reference die

L is the number of overlay-critical exposures, while L_0 is the number of overlay-critical exposures required for the reference node

The model assumes that overlay errors scale linearly with the number of critical exposures and with die length, rather than die area, reflecting the fact that overlay errors are typically moderately spatially correlated. Notably, industry experts indicate that actual overlay loss is often dominated by one or two product- and process-specific layers, while other layers with similarly tight specifications contribute little measurable loss. Public information is not sufficient to identify or analyze the key layers in relevant products, so we retain the transparent count-based approach.

We fit the model using published data from ASML at the 20nm node. ASML's 2014 investor-day presentation estimated less than 1% overlay yield loss for a 100mm^2 20nm logic die when the 5nm on-product overlay specification was met, versus ~4% loss when it was exceeded by 20% (6nm on-product overlay, or OPO), assuming 14 immersion passes.[203] Fitting these observations with A_0 = 100 mm^2 and L_0 = 14 yields κ = 4.6.[110]

We apply this model for each combination of NXT tool and node. We assume OPO of [1.5, 1.75, 2.0]nm for NXT:2100i,[111] [3.0, 3.25, 3.5]nm for NXT:1980i, and [4.25, 4.5, 4.75]nm for NXT:1470.[204] We assume

[110] As a backcheck, setting κ to 4.6 implies ~2nm mean overlay, which matches ASML's published numbers.

[111] NXT:2100i has an MMO of 1.3nm; adding additional overlay errors likely takes OPO to 2.0nm.

node OPO specs of [5.75, 6.0, 6.25]nm at 28nm, [3.75, 4, 4.25]nm at 16/14nm,[112] and [2.75, 3.0, 3.25]nm at 7nm.[113][207] We estimate overlay-critical exposure counts by anchoring to ASML's disclosed 14 overlay-critical immersion passes at 20nm and scaling across nodes with two triangulating methods: (i) ASML's published lithography-step counts for a fixed critical-layer stack (6/8/23/34 steps at 28/20/10/7nm) and (ii) a bottom-up decomposition of each node's known patterning schemes, counting every split, mandrel, block, and cut exposure at layers whose overlay budget is within ~1.5× of the node's tightest spec.[208] Counts are identical for all immersion tools. For the NXT:1470, we estimate extra exposures required for layers that newly require advanced patterning.

Table 13: Estimated overlay-critical exposure counts for tool and node combinations

Node	NXT:2100i/1980i	NXT:1470
28nm	[8, 10, 15]	[11, 14, 19]
16/14nm	[16, 18, 22]	[20, 25, 30]
7nm	[32, 38, 45]	[42, 53, 66]

Under these assumptions, we find the following yields:

Table 14: Overlay yield for a 500mm² logic die

NXT tool	$Y_{overlay}$		
	28nm	16/14nm	7nm
NXT:2100i	[100]	[100]	[100]
NXT:1980i	[100]	[99, 100, 100]	[77, 89, 96]
NXT:1470	[99, 100, 100]	[83, 91, 95]	[17, 34, 52]

$Y_{residual,\ non\text{-}overlay}$ captures patterning failures within a single exposure, particularly critical-dimension (CD) variation, focus/dose variation, line-edge roughness, and optical proximity correction (OPC) errors. This factor is probably small when the lithography system can support the node within its process window but expands dramatically when the process window is exceeded.

We estimate [0.97, 0.98, 0.99] for NXT:2100i and NXT:1980i at all three nodes, and for NXT:1470 at 28nm. The UC Berkeley study of mature fab yields, co-sponsored by SEMATECH, found that the best processes achieved 96–98% total systematic yield, with losses split between lithography and other process steps.[209] ASML/Brion's 2019 voltage-contrast study provides complementary evidence that even with overlay minimized, local edge placement error generates a small, residual loss.[210]

[112] ASML estimates OPO requirements below 5nm at 16/14nm.[205]

[113] OPO requirements at 7nm range between 2.5nm and 3.5nm; 3nm is frequently described as a control case.[206]

- For NXT:2100i and NXT:1980i, there is strong evidence that residual loss is minimized. NXT:2100i was used for the most critical layers at 7nm and for critical immersion layers at 3nm and below.[211] ASML explicitly characterized NXT:1980i as an immersion scanner for 7nm and 5nm production nodes, and NXT:1980i provides the same 1.35 NA and 38nm resolution as NXT:2100i;[114] NXT:2100i's capability is thus partial evidence that basic optical resolution is not a limiting problem for NXT:1980i.[212]
- For NXT:1470 at 28nm, evidence also suggests minimal residual yield loss. First, ASML claims that XT:1450, an older dry ArF system with worse process control, is sufficient for high volume manufacturing at the 32nm node.[213] Second, the even older XT:1400E was experimentally found to support extending capability beyond 32nm, "towards the 22nm node," and to enable printing 40nm contact holes with an 80nm pitch (tighter than the 28nm node's ~113nm contact pitch requirement).[214] Third, ASML found that dry ArF systems can achieve 40nm half pitch, tighter than the 45nm half pitch required for 28nm's tightest metal layer (M1).[215]

For NXT:1470 at 16/14nm and 7nm, we estimate [0, 0.5, 0.98], reflecting wide uncertainty as to whether these nodes can be implemented. The P10 reflects a judgment that outright infeasibility cannot be ruled out, particularly at 7nm. On the other hand, it is possible that a workable process can be obtained, particularly for products with highly restrictive design rules, that could minimize this yield factor. This factor is the dominant source of uncertainty for the two affected cells, and the resulting yield ranges should be read as incorporating fundamental uncertainty about the feasibility of implementing these nodes using NXT:1470.

[114] NXT:2100i's advantage over NXT:1980i comes primarily from improved overlay and process control.

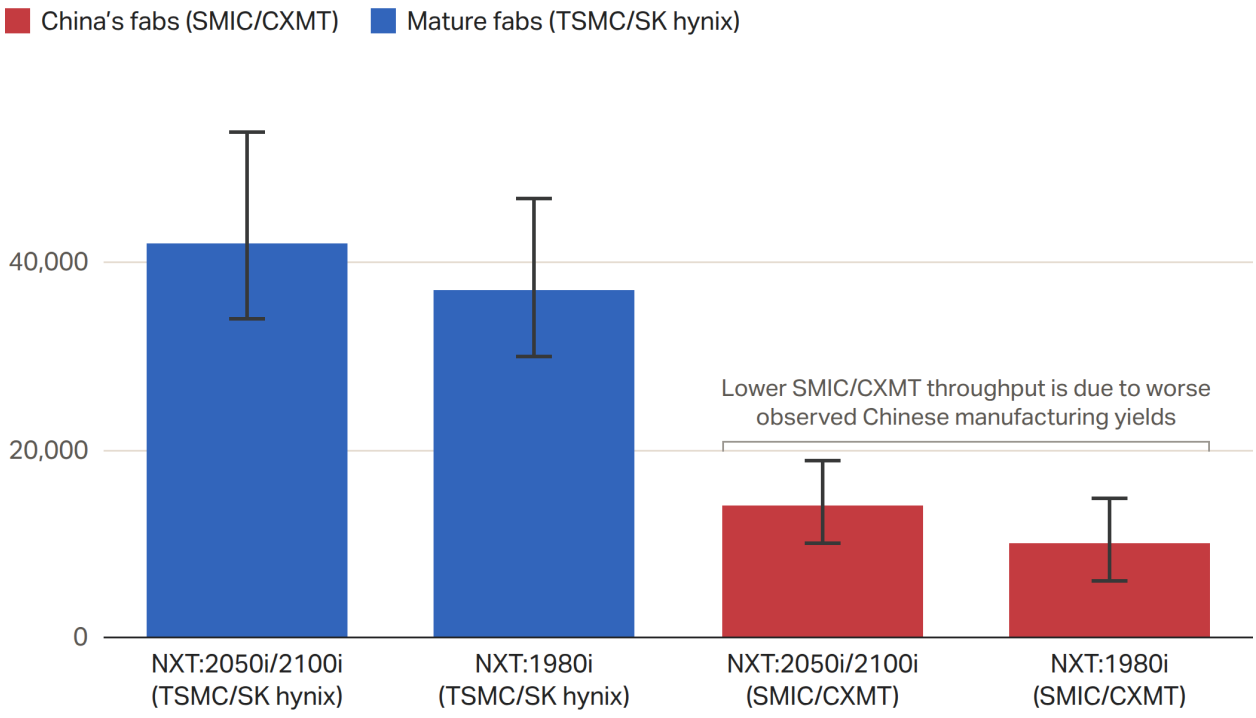
Appendix D: DUVi throughputs for Huawei AI chip production

We estimate the throughputs of the advanced DUVi systems within China's installed base for producing Huawei's Ascend 910C chips, in two scenarios: (1) usage within mature advanced fabs (e.g., TSMC, SK hynix) achieving industry-leading die, stacking, and packaging yields, and (2) usage within China's fabs (e.g., SMIC, CXMT), providing substantially lower yield. The reported ranges are for the P10, P50, and P90 of the distributions produced by running Monte Carlo simulations sampled from various input distributions, as described below. See Figure 11.

The estimates represent the number of finished Ascend 910C-equivalent packages that could be supported by the wafers patterned by one scanner in a single month, conditional on sufficient complementary equipment and packaging capacity.

Figure 11. DUVi throughputs across production settings

Ascend 910C-equivalents per month



Bars reflect P50 estimates for monthly production of finished Ascend 910C chips, assuming throughput is solely constrained by lithography systems. Ranges reflect P10–P90 estimates. See Appendix D for assumptions and calculations.

Table 15: DUVi throughputs across production settings

NXT tool	Ascend 910C-eq / month (leading advanced fab)	Ascend 910C-eq / month (SMIC N+2, CXMT 1y)
NXT:2050i/2100i	[34k, 42k, 54k]	[10k, 14k, 19k]
NXT:1980i class[115]	[30k, 37k, 47k]	[6k, 10k, 15k]

Methodology

We estimate throughputs for each DUVi system when producing the two silicon inputs to the Ascend 910C: 7nm logic dies, and HBM2e stacks using 1y DRAM dies. For each input, we estimate output per scanner as a function of (i) the scanner's wafer-pass throughput, (ii) the number of DUVi wafer passes required per wafer, (iii) the number of dies obtainable per wafer, (iv) die yield, and (v) for HBM2e, backend stacking yield. Dies are converted to Ascend 910C-equivalents assuming two 7nm logic dies and 64 DRAM dies (eight 8-Hi stacks of HBM2e) per chip. This produces an estimate of the number of chips that the scanner could support if solely dedicated to logic or to memory. Last, we estimate an optimal allocation of the tool fleet between logic and memory and apply a CoWoS packaging yield factor, producing the final estimates.

We model production occurring in two scenarios: (1) a mature advanced fab (e.g., TSMC for logic, SK hynix for HBM), with industry-leading die, stacking, and packaging yields, which bounds output from above, and (2) China's advanced fabs (e.g., SMIC for logic, CXMT[116] for HBM), with substantially lower yields. The second scenario is useful for estimating DUVi production in China under current conditions, though only under the critical assumption that sufficient complementary SME exists to reach DUVi throughput limits.

We model input uncertainty by specifying three quantiles (P10/P50/P90). Inputs representing positive quantities are fitted in log space; inputs representing yields are fitted in logit space. We run 100,000 Monte Carlo simulation draws, each of which samples one internally consistent set of parameter values and evaluates the full model, including the tool allocation step.

Caveats

Beyond explicitly modeled uncertainties, this overall analysis has four notable conceptual limitations:

1) We assume DUVi is the sole constraint on Ascend 910C throughput—i.e., that there is a limitless supply of other (non-litho) frontend tools and advanced packaging capacity. In reality,

[115] We model NXT:1980i class scanners as a blend of several generations (Di, Ei, and Fi), approximating throughput of a blended fleet of scanners.

[116] This is a counterfactual, given that CXMT probably skipped HBM2e. We estimate yields by extrapolating from reporting about its HBM3 efforts.

both the US-allied and Chinese semiconductor ecosystems face additional constraints on throughput. In particular, China's advanced node production is constrained by limited availability of etch and deposition tools. Until this bottleneck is overcome, China's production throughput would be unchanged even if many more DUVi systems were diverted to 910C production. These estimates should therefore be understood as assessing the throughput of DUVi *in the absence of these other constraints*.

2) We assume that logic and HBM are independent bottlenecks to production, in which case throughput is maximized by allocating DUVi between logic and HBM in proportion to their contribution to the Ascend 910C. In reality, production is likely to be bottlenecked on either logic or HBM, but not both—e.g., as of spring 2026, HBM, not advanced logic, bottlenecks both the Chinese and the US ecosystem. Each additional DUVi provides a much greater boost to throughput in this scenario, because the fab can allocate its productive capacity towards alleviating one bottleneck rather than splitting it between two.
3) Estimates depend on the Ascend 910C's specific technical architecture, including the specific choice of nodes (7nm logic, 1y DRAM, HBM2e) and the ratios of silicon usage. They would differ for chips using alternate nodes, technologies (e.g., GDDR), and proportions.
4) Each HBM stack also contains a logic base die, which might also require immersion passes if fabricated using advanced nodes. We exclude base dies for simplicity and because public information does not establish that base dies would require advanced immersion. However, this remains a major source of uncertainty.

Wafer-pass throughput

ASML rates the NXT:1980Di at 275 wafers per hour, and the NXT:2050i and NXT:2100i at 295 wafers per hour.[216] These imply theoretical maxima of approximately 198,000 (NXT:1980Di) and 212,000 (NXT:2050i/2100i) wafer passes per month at continuous operation.[117] In practice, throughput is a fraction of the rated maximum, since ratings assume benchmark exposure conditions and continuous availability.

Later NXT:1980i variants have higher rated throughputs (295 wafers/hour for NXT:1980Ei and 330 for NXT:1980Fi). To model a blended mix of NXT:1980i scanners, we apply a flat 5% increase on top of the realized estimated NXT:1980Di throughput.

[117] When ASML reports one "wafer" per hour, it means one exposure. Therefore, for example, a double exposure on a single wafer counts as two wafers for the purposes of calculating throughput.

Table 16: Exposures per month by DUVi system

NXT Tool	Exposures / month
NXT:2050i/2100i	[135k[118], 150k[119], 180k[120]]
NXT:1980i	[131k, 145k[121], 174k]

Logic

The Ascend 910C co-packages two 7nm logic dies, with a combined area of 1,300 mm², or roughly 650 mm² per die. Throughput for manufacturing these dies depends on (i) wafer-pass throughput (see above), (ii) the number of DUVi wafer passes required per 7nm wafer, (iii) the number of dies produced per wafer, and (iv) the fraction of dies produced that are sellable.

Table 17: DUVi system exposures and dies per wafer for logic chips

Factor	Estimate
Exposures per wafer	[32, 38, 45][122]
Dies per wafer	[76, 80, 85][123]

Table 18: DUVi system yield at leading fab versus SMIC

NXT tool	Die yield (leading fab)	Yield (SMIC N+2)
NXT:2050i/2100i	[74, 80, 86][124]	[30, 40, 50][125]
NXT:1980i	[57, 69, 78]	[18, 32, 45][126]

[118] Assumes availability dips from the typical ~97% to ~90%.

[119] Reflects published data from ASML's installed base.[217]

[120] ASML stated that its latest NXT systems "have shown the ability to run in excess of 6,000 wafers per day."[218]

[121] Scales the NXT:2050i/2100i throughput by the ratio of rated throughputs (275 / 295).

[122] See Appendix C.

[123] Dies per wafer is computed as a function of wafer diameter (d), die size (S), and edge losses, as follows: Die per wafer = d * π * (d / (4 * S) - 1 / sqrt(2 * S)), with a 3mm edge exclusion.

[124] We derive sellable die yield for all tools from the yield estimates in Appendix C, assuming that roughly two-thirds of the random defect yield loss can be recovered using redundancy and harvesting.

[125] Reflects a range of recent estimates, including: (1) 30% at the end of 2025, projected to rise to 70% by 2027 (Morgan Stanley), (2) reporting in early 2026 citing low yields as a barrier to SMIC's margins, and (3) 40% in February 2025 (TrendForce).[219]

[126] The P50 estimate applies an additional 20% yield loss factor on top of the P50 estimate for NXT:2100i/2050i. The P10 applies a 40% factor on the P10 for NXT:2100i/2050i. The P90 applies a 10% factor to the P90 of NXT:2100i/2050i. The spread reflects uncertainty as to how lithography capabilities interact with existing yield challenges at SMIC and CXMT.

These factors combine as follows:

$$D = (W_{\text{tool}} / N_{\text{exp}}) \times N_{\text{die}} \times Y$$

Good die throughput per month

Wafer pass throughput per month

Exposures per wafer

Gross dies per wafer

Usable die yield

Table 19: DUVi system logic throughputs

NXT tool	Dies / month (leading fab)	Dies / month (SMIC N+2)
NXT:2050i/2100i	[204k, 258k, 332k]	[90k, 128k, 179k]
NXT:1980i	[160k, 213k, 280k]	[54k, 98k, 152k]

HBM

The Ascend 910C includes 8 stacks of 8-Hi HBM2e, incorporating a total of 64 1y DRAM dies. Throughput for producing HBM depends on the four factors above—dies per wafer, exposures per wafer, exposures per month, die yield—and an additional multiplicative yield factor capturing loss from stacking DRAM dies into HBM. We estimate these as follows:

Table 20: DUVi system exposures and dies per wafer for HBM

Factor	Estimate
Exposures per wafer	[15, 19, 23][127]
Dies per wafer	[550, 575, 610][128]

Table 21: DUVi system HBM yields

	Leading fab		CXMT	
NXT tool	**Frontend yield**	**Backend yield**	**Frontend yield**	**Backend yield**
NXT:2050i/2100i	[90, 95, 99][129]	[90, 95, 99][130]	[35, 45, 55][131]	[70, 80, 90][132]
NXT:1980i	[85, 90, 95][133]		[21, 36, 50][134]	

Combining factors as above, we have:

Table 22: DUVi system HBM throughputs

NXT tool	8-Hi HBM2e stacks / month (leading fab)	8-Hi HBM2e stacks / month (CXMT)
NXT:2050i/2100i	[406k, 527k, 703k]	[144k, 209k, 302k]
NXT:1980i	[374k, 484k, 647k]	[89k, 159k, 254k]

[127] Reflects the span of estimates for Micron, SK hynix, and Samsung in this 2020 presentation from IC Knowledge.[220]

[128] Dies per wafer is computed using the same methodology as for logic, assuming (i) a 100mm² die, slightly below Micron's HBM2e package area, measured at 110mm², and (ii) a 3mm edge exclusion.[221]

[129] Mature DRAM yields, including for HBM-destined dies, exceed 90%, assuming the tool is viable for the node. For example, SK hynix's more advanced 1a DRAM reached 90%, with a Korean report claiming that mature DRAM-die yield can reach 90–100% before HBM's TSV and packaging losses.[222] There is strong evidence that NXT:2100i and NXT:2050i are viable for 1y DRAM—e.g., 1y was implemented originally using NXT:1980i, a less advanced tool.

[130] SK hynix says that MR-MUF, the stacking technology used for HBM2e, enabled stable mass production with high manufacturing yield.[223] Its packaging lead later reported that the validation packaging yield of the (substantially more challenging) 16-high HBM3E was 99% **of** the production packaging yield of 12-high HBM3E.[224]

[131] Moderately adjusted upward from SemiAnalysis's estimate of 35% frontend yield for CXMT's 8-Hi HBM3, which is corroborated by this industry estimate of 25% all-in (frontend and stacking) yield.[225] Note that CXMT did not produce significant HBM2e.

[132] Moderately adjusted upward from SemiAnalysis's 70% stacking yield estimate for HBM3.[226]

[133] As discussed above, there is strong evidence that NXT:1980i is viable for 1y DRAM. For example, ASML's 2018 roadmap claims support for 1y memory a full year before the NXT:2000i, suggesting the 1y ramp initially used the NXT:1980i, and Samsung's initial 1y DRAM mass production in late 2017 was likely achieved primarily using the 1980i, since the successor NXT:2000i had only just begun limited early access testing at that point and did not begin shipping widely until 2018.[227] We assume a modest yield penalty relative to NXT:2100i.

[134] The P50 estimate applies an additional 20% yield loss factor on top of the P50 estimate for NXT:2100i/2050i. The P10 applies a 40% factor on the P10 for NXT:2100i/2050i. The P90 applies a 10% factor to the P90 of NXT:2100i/2050i. The spread reflects uncertainty as to how lithography capabilities interact with existing yield challenges at SMIC and CXMT.

Ascend 910C

We convert the above throughputs to Ascend 910C-eq, assuming two logic dies and 8 HBM2e stacks per chip:

Table 23: DUVi system Ascend 910C-equivalent silicon throughputs

NXT tool	910C-eq logic / month (leading fab)	910C-eq memory / month (leading fab)	910C-eq logic / month (SMIC)	910C-eq memory / month (CXMT)
NXT:2050i/2100i	[102k, 129k, 166k]	[51k, 66k, 88k]	[45k, 64k, 89k]	[18k, 26k, 38k]
NXT:1980i	[80k, 106k, 140k]	[47k, 61k, 81k]	[27k, 49k, 76k]	[11k, 20k, 32k]

Assuming no independent constraints on either HBM or logic production, overall Ascend 910C throughput is balanced by allocating tools in inverse proportion to their Ascend 910C logic and memory throughputs. This allocation is computed separately for each draw, so the ratio varies. For example, at the median of the SMIC/CXMT scenario, approximately 2.4 NXT:1980i systems are allocated to CXMT's HBM production for every tool allocated to SMIC's logic production. With the tool fleet split in these ratios, the blended throughput of 910C-equivalent silicon ($T_{blended}$) is the harmonic combination:

$$T_{\text{blended}} = 1/(1/T_{\text{L}} + 1/T_{\text{M}})$$

- T_{blended}: The blended throughput of 910C-equivalent silicon
- T_{L}: 910C-equivalent logic dies per tool per month
- T_{M}: 910C-equivalent memory dies per tool per month

Finally, we multiply by an advanced packaging (CoWoS or other) yield factor, accounting for loss during advanced packaging, estimated at [97, 98, 99][135] for a mature fab and [65, 75, 85][136] for SMIC/CXMT to get our final estimates for finished 910C chips per DUVi system.

[135] TSMC claims CoWoS yields top 98%, which we take as P50.[228]

[136] CSIS claimed in 2025 that China had CoWoS yields of 75%, which we take as P50.[229]

Appendix E: Assessing China's DUVi installed base

We estimate that China imported 396 DUVi systems between 2012 and 2026Q1 (with a central 80% model range of 361 to 444). Subtracting an estimated 42 / 52 / 67 imports at four multinational fabs yields an estimated 312 / 343 / 383 systems imported for Chinese-owned fabs. Within the Chinese-owned subtotal, we estimate 29 / 45 / 69 NXT:1970i and less advanced systems, 249 / 270 / 294 NXT:1980i systems, 22 NXT:2050i systems, and six NXT:2100i systems. These figures represent cumulative imports rather than operational scanners; retirements and re-exports are not modeled, and imports before 2012 are excluded. We limit our focus to ASML's systems, rather than Nikon, which has minimal market share for advanced systems.[137]

Alongside this aggregate estimate, we estimate installations at (i) SMIC's Beijing and Shanghai fabs, including the Shanghai campus used for 7nm logic, (ii) CXMT's Hefei and Beijing fabs, (iii) fabs owned by non-Chinese multinationals, (iv) all advanced logic and DRAM fabs owned by Chinese companies, (v) YMTC, China's leading NAND producer, and (vi) select other publicly documented fabs. These estimates have varying evidentiary status. The SMIC and CXMT figures are fab-specific estimates; advanced-logic and advanced DRAM totals are broad allocation estimates; "other known installations" are partial observations; and the multinational estimate is an undifferentiated ownership allocation. Public data do not enable a comprehensive fab-by-fab accounting: the entity-level figures below include only systems we can estimate or document, while "unknown" captures the remainder.

Unless otherwise stated, ranges report P10 / P50 / P90. Quantiles are calculated after aggregation and subtraction within each Monte Carlo draw, so they need not add or subtract exactly.

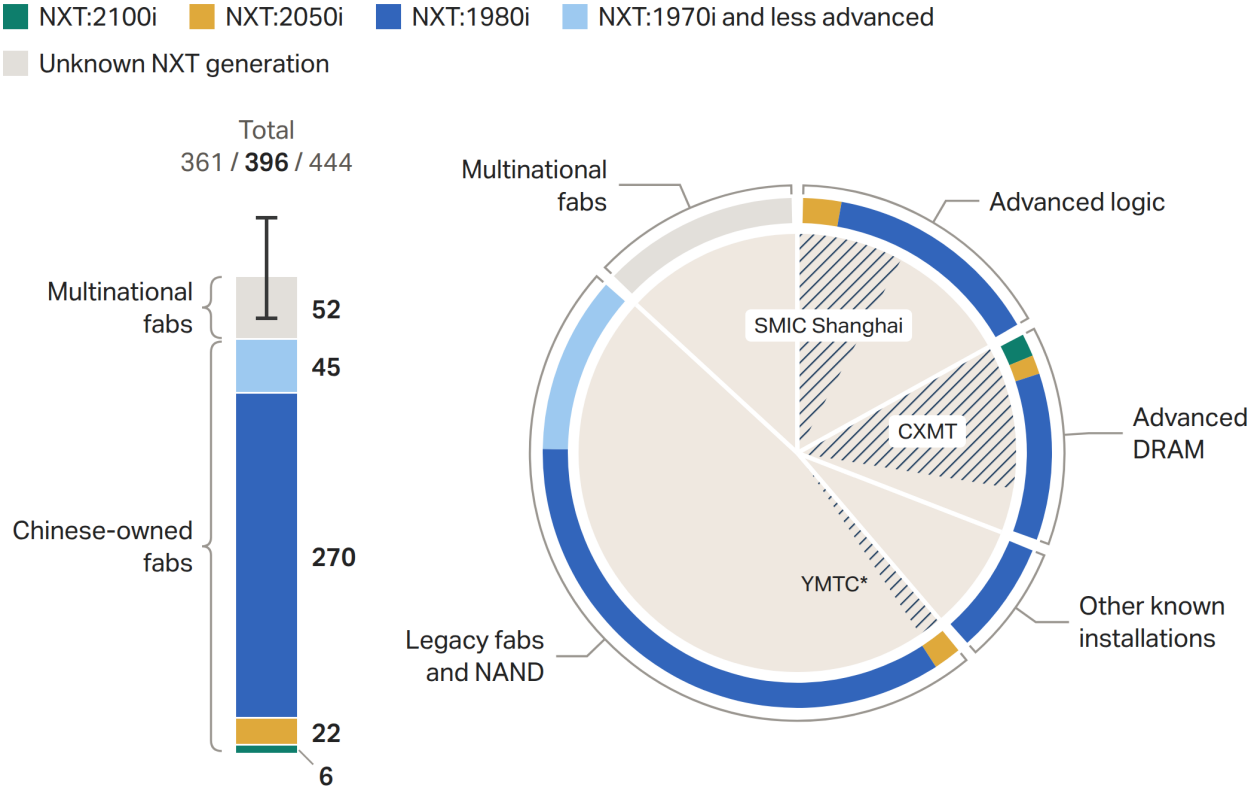
[137] Furthermore, 80% of Nikon's sales have gone to Intel.[230]

Table 24: DUVi systems installed in China

Tools	Total	Advanced logic		Advanced DRAM		YMTC	Other known installations
		SMIC Shanghai	Total	CXMT	Total		
NXT:1970i and less advanced	29 / 45 / 69	Unknown	Unknown	Unknown	Unknown	Unknown	Unknown
NXT:1980i	249 / 270 / 294	19 / 24 / 28	37 / 57 / 89	20 / 33 / 44	29 / 44 / 62	Unknown	31
NXT:2050i	22	6 / 7 / 8	9 / 10 / 11	3	4 / 5 / 6	5 / 7 / 9	0
NXT:2100i	6	0	0	6	6	0	0
Total for Chinese-owned fabs	**312 / 343 / 383**	25 / 31 / 36	47 / 67 / 99	29 / 42 / 53	40 / 55 / 73	Unknown	Unknown
Total for multinational fabs	42 / 52 / 67						
Total for all fabs in China	361 / 396 / 444						

Figure 18. China's 2026 DUVi installed base

Estimated NXT DUVi systems installed in China, by fab and tool type



Wedges are sized by median estimates; hatched areas show the share of a category at the named fab. The outer ring shows each wedge's tool mix; for legacy fabs and NAND it combines YMTC's units with the Chinese-owned units not attributed to another wedge (138 NXT:1980i and 45 NXT:1970i and less advanced). See Appendix E for assumptions and calculations.

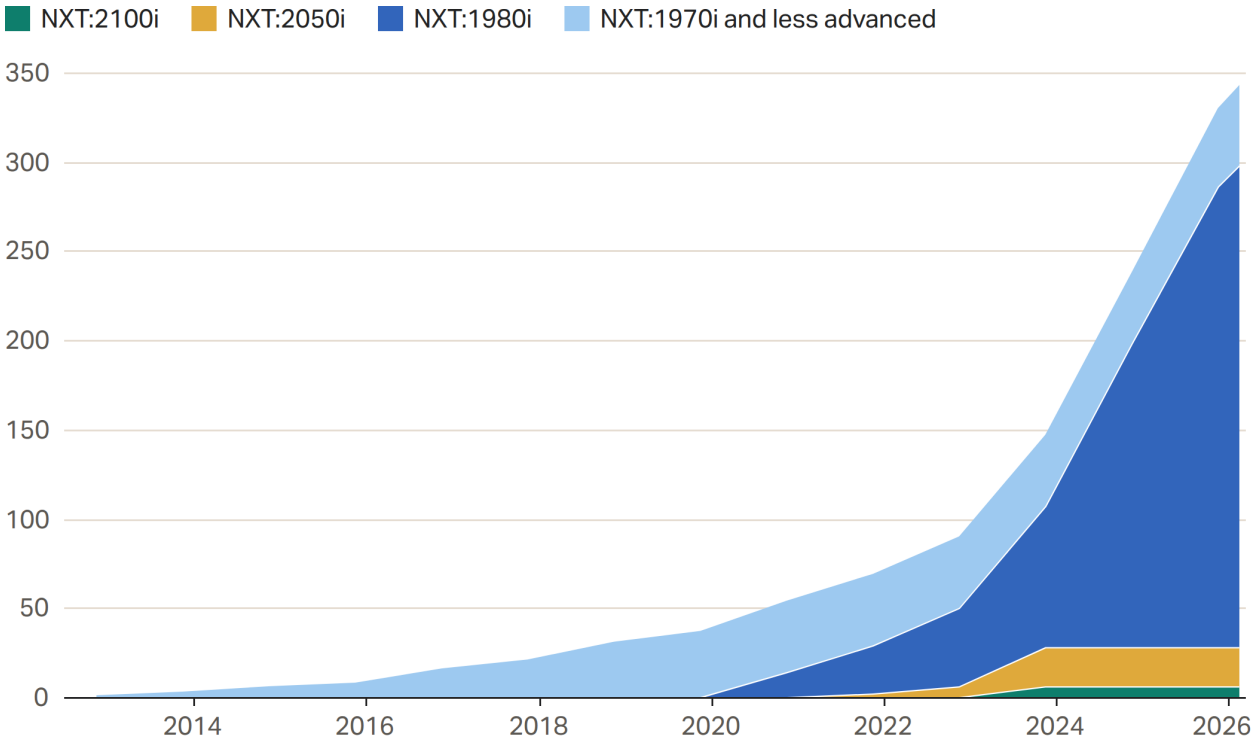
*YMTC share covers only its NXT:2050i systems; its NXT:1980i count is unknown.

Center for Technology & Statecraft IFP

China has probably imported DUVi systems since at least SMIC's first introduction of the 45nm logic node in the late 2000s.[231] From 2012 onward, China imported modest but growing numbers of earlier-generation systems; after NXT:1980i became available in the late 2010s, it began importing those systems as well, reportedly accumulating up to 80 by fall 2023. Imports accelerated sharply in 2023, after the United States and the Netherlands announced tighter controls. During 2023, Chinese fabs imported six NXT:2100i systems, 16 NXT:2050i systems, and 35 NXT:1980i systems. From 2024 to 2026Q1, China more than tripled its stockpile of NXT:1980i systems.

Figure 2. China is building a vast stockpile of DUVi systems

Estimated cumulative DUVi units imported by Chinese-owned fabs since 2012



Estimates exclude imports to multinational fabs located within China. See Appendix E for assumptions and calculations.

Center for Technology & Statecraft IFP

China's aggregate installed base

We estimate cumulative DUVi imports attributable to Chinese-owned fabs over 2012–2026Q1, in two stages. First, we reconstruct gross annual imports into China, using ASML's public financial disclosures, China's customs statistics, shipment figures published by the Select Committee on the CCP, and public estimates of China's installed base as of 2023. Second, we estimate imports to four multinational fabs (Intel Dalian in Liaoning, Samsung Xi'an in Shaanxi, and SK hynix Wuxi and TSMC Nanjing in Jiangsu) and subtract these imports from the gross estimate.

We report ranges for the P10, P50, and P90 of distributions produced by running 100,000 Monte Carlo simulations sampled from uncertain input distributions specified by P10, P50, and P90 quantiles. Positive-valued inputs are sampled from two-piece lognormal distributions fitted separately to their upper and lower quantiles. Pricing and ASML-finance assumptions are drawn once per simulation and applied across years, because errors in these assumptions affect all such years. By contrast,

assumptions about customs data are drawn separately for each province-year observation, reflecting observation-level uncertainty in customs classification, coverage, and valuation.

Table 25: DUVi imports attributable to Chinese-owned fabs

Year	NXT:2100i	NXT:2050i	NXT:1980i	NXT:1970i and less advanced	Total
2026Q1	0	0	9/11/13	0/0/2	9/11/14
2025	0	0	72/89/111	0/0/14	78/94/114
2024	0	0	90	0/0/0	90
2023	6	16	35	0/0/0	57
2022	0	4	17	0/0/0	21
2021	0	2	14	0/0/0	16
2020	0	0	6/12/16	0/2/11	12/15/19
2019	0	0	0/0/7	1/7/10	6/8/11
2018	0	0	0/0/0	7/10/14	7/10/14
2017	0	0	0/0/0	4/5/7	4/5/7
2016	0	0	0/0/0	6/8/10	6/8/10
2015	0	0	0/0/0	1/2/4	1/2/4
2014	0	0	0/0/0	1/3/6	1/3/6
2013	0	0	0/0/0	1/3/5	1/3/5
2012	0	0	0/0/0	0/1/1	0/1/1
Total for Chinese-owned fabs	**6**	**22**	**249/270/294**	**29/45/69**	**312/343/383**
Multinational fabs					42/52/67
Total including multinational fabs					361/396/444

Estimate from ASML's financial data

ASML's financial disclosures report its ArFi revenue share, worldwide ArFi shipments, EUV revenue share, and share of revenue attributable to customers in China. All revenues use net system sales, excluding servicing and field options. We estimate ASML's China ArFi revenue as:[138]

[138] There is some uncertainty here. On the one hand, the lack of access to EUV might mean the China market skews heavier into DUVi (the next best substitute) than the rest of the world. On the other hand, export controls restrict China from accessing the best DUVi systems, potentially reducing its share of DUVi.

$$\underbrace{R_{\text{China, DUVi}}}_{\text{ASML's DUVi revenue from China}} = \underbrace{R_{\text{China}}}_{\text{ASML's China revenue}} \times \underbrace{R^{\text{frac}}_{\text{Worldwide, DUVi}}}_{\text{ASML's DUVi share of worldwide system revenue}} / (1 - \underbrace{R^{\text{frac}}_{\text{Worldwide, EUV}}}_{\text{ASML's EUV share of worldwide system revenue}})$$

$$\underbrace{S_{\text{China, DUVi}}}_{\text{ASML's China DUVi shipments}} = \underbrace{S_{\text{Worldwide, DUVi}}}_{\text{ASML's worldwide DUVi shipments}} \times R_{\text{China, DUVi}} / \underbrace{R_{\text{Worldwide, DUVi}}}_{\text{ASML's worldwide DUVi revenue}}$$

This method has two important limitations.

- First, it assumes that the composition of ASML's China system revenue resembles the composition of ASML's non-EUV system revenue outside of China. The direction of the resulting bias is unclear. On the one hand, non-Chinese customers plausibly spend less on dry DUV because they typically invest less in legacy nodes. On the other hand, access to EUV may reduce their spend on DUVi as well. We account for this uncertainty using a multiplier on ASML's China ArFi revenue, with quantiles [0.8, 1, 1.2].
- Second, it assumes that Chinese customers purchase the same mix of DUVi models as customers outside of China. This assumption is almost certainly incorrect, because controls from 2023 onwards restrict NXT:2000i and more advanced systems, concentrating later purchases in the less expensive NXT:1980i. Comparing ASML's worldwide blended DUVi revenue to China's DUVi revenue may understate the number of systems China purchased. We account for this uncertainty using a multiplier on China's ArFi price, with quantiles [0.8, 1, 1.1].

Table 26: ASML's ArFi revenue and shipments, estimated from public financial disclosures

	Cumulative	2026Q1	2025	2024	2023	2022	2021	2020
ArFi China revenue ($M)	29,380	945	7,375	6,853	4,865	1,433	1,723	1,418
ArFi China shipments	[306, 395, 514]	[8, 10, 13]	[66, 85, 111]	[68, 87, 114]	[50, 64, 83]	[17, 21, 28]	[19, 25, 32]	[17, 22, 29]

	2019	2018	2017	2016	2015	2014	2013	2012
ArFi China revenue ($M)	918	1,397	602	629	365	448	337	71
ArFi China shipments	[11, 15, 19]	[17, 22, 28]	[8, 10, 13]	[9, 12, 15]	[5, 7, 9]	[6, 8, 10]	[5, 6, 8]	[1, 1, 2]

Estimate from China's customs data

China's public customs statistics portal reports the quantity and declared import value of lithography equipment under HS codes 84862031 ("Step and repeat aligners (steppers) for manufacturing semiconductor devices or integrated circuits") and 84862039 ("Other lithography equipment"), for each Chinese province.[232] The limitation of this dataset is that it does not identify scanner models. NXT:1980i systems are aggregated with earlier DUVi, dry DUV, and i-line, and potentially other equipment classified under the same codes.

For each province and year, we treat the reported quantity and value as deriving from a three-part mixture of:

- NXT:2050i and NXT:2100i systems, for which shipments are independently known;
- residual immersion-priced systems, labeled "DUVi-priced units" in 2015–2020 and used as an NXT:1980i proxy from 2021 onward; and
- dry DUV and less advanced (e.g., i-line) systems.

$$\text{Value}_{\text{total}} = \underbrace{\text{Shipments}_{\text{immersion}}}_{\text{Residual immersion-priced systems}} \times \underbrace{\text{ASP}_{\text{immersion}}}_{\text{Estimated for 2015–2020 using ASML's worldwide blended ArFi ASP; from 2021 onward, a modeled NXT:1980i}} +$$

$$+ \underbrace{\text{Shipments}_{\geq 2050i}}_{\text{NXT:2050i and NXT:2100i systems, for which shipments are independently known; estimated in Table 27}} \times \underbrace{\text{ASP}_{\geq 2050i}}_{\text{Estimated at \$90M (Table 35) with multiplier [0.8, 1, 1.2]}} +$$

$$+ (\text{Shipments}_{\text{total}} - \text{Shipments}_{\text{immersion}} - \text{Shipments}_{\geq 2050i}) \times \underbrace{\text{ASP}_{\text{dry}}}_{\text{Estimated at \$[9, 12.5, 15] million}}$$

$$\underbrace{\text{Shipment}_{\text{immersion}}}_{\text{Residual immersion-priced systems}} =$$

$$= \text{clamp}\left[\begin{array}{l} (\text{Value}_{\text{total}} - \text{Shipments}_{\geq 2050i} \times \\ \quad \times (\text{ASP}_{\geq 2050i} - \text{ASP}_{\text{dry}}) - \\ \quad - \text{Shipment}_{\text{total}} \times \text{ASP}_{\text{dry}}) \, / \\ \quad / (\text{ASP}_{\text{immersion}} - \text{ASP}_{\text{dry}}), \\ 0, \text{Shipments}_{\text{total}} - \text{Shipments}_{\geq 2050i} \end{array}\right]$$

where:

- ASP_{dry} is estimated at \$[9, 12.5, 15] million, calibrated based on (i) ASML financial data (Table 35) and (ii) customs import values and quantities over 2021-2024, in which the Select Committee data confirms import quantities for NXT:1980i, 2050i, and 2100i. The customs method estimates ASP of \$12.6 million, while the ASML method estimates \$12.5 million.

- Shipments$_{>=2050i}$ are estimated in Table 27:

Table 27: Shipments of NXT:2050i and NXT:2100i, by province

System	Year	Quantity	Province	Recipient
NXT:2100i	2023	5	Anhui	CXMT
	2023	1	Beijing	CXMT
NXT:2050i	2021	2	Shanghai	SMIC
	2021	1	Shanghai	Huahong
	2023	6	Shanghai	SMIC
	2022-2023	3	Anhui	CXMT
	2022-2023	7	Hubei	YMTC

- $ASP_{>=2050i}$ is estimated at $90M (Table 35) with multiplier [0.8, 1, 1.2]
- $ASP_{immersion}$ is estimated for 2015–2020 using ASML's worldwide blended ArFi ASP: worldwide ArFi net system sales divided by worldwide ArFi units. From 2021 onward, residual immersion units are priced as NXT:1980i, supported by data from the Select Committee on China, and modeled as a linear interpolation between $63M in 2018 (Table 35) and $76M in 2024 (estimated using ASML financial disclosures and Select Committee shipment numbers as $6,853M / 90 = $76M). For 2025-2026Q1, it is scaled in proportion to the blended DUVi ASPs (Table 35). We apply a [0.8, 1, 1.2] price multiplier.

We can now generate the following estimates for residual immersion-priced units.

Table 28: Estimated immersion-priced imports by province, excluding NXT:2000i and above

Province	Total	2026Q1	2025	2024	2023	2022	2021	2020	2019	2018	2017	2016	2015
Shanghai	77/97/127	6/7/9	34/43/57	15/20/27	3/6/10	2/3/4	0/1/3	9/11/15	0/1/1	1/2/3	1/1/2	0/0/0	0/0/0
Jiangsu	25/31/42	0/0/0	0/0/1	1/2/3	2/3/4	0/0/0	0/0/0	6/7/10	5/6/8	6/8/11	3/4/5	0/0/0	1/1/2
Shaanxi	6/8/11	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	2/2/3	0/1/1	1/2/2	1/2/2	0/0/0	0/0/0	1/2/3
Liaoning	6/8/11	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	3/4/5	1/2/3	2/3/4	0/0/0
Beijing	50/63/81	0/1/1	13/16/20	17/22/29	9/11/15	3/3/5	3/4/6	0/0/0	0/0/0	0/0/0	1/1/2	3/4/5	0/0/0
Hubei	32/41/54	0/0/0	9/12/16	7/9/12	4/7/11	2/3/4	2/3/4	1/2/3	2/3/4	1/2/3	0/0/0	0/0/0	0/0/0
Guangdong	32/41/54	2/2/3	16/21/27	5/7/10	8/11/14	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0	0/0/0
Anhui	26/33/42	1/1/1	3/3/4	10/13/17	0/1/4	4/6/7	2/3/4	0/0/1	2/2/3	2/3/4	0/0/0	0/0/0	0/0/0
Other provinces	34/43/57	1/2/2	4/5/6	17/22/29	2/4/5	0/0/0	3/4/5	0/0/0	1/1/2	2/2/3	1/2/2	2/3/4	0/0/0
Total	292/364/476	10/13/16	81/101/131	75/94/124	33/43/59	11/15/20	13/17/22	17/22/30	12/15/20	18/22/29	8/10/13	7/9/13	2/3/5

The customs method has several limitations. First, the HS codes might include a variety of equipment beyond new ASML scanners, including non-ASML equipment, used equipment, or ancillary items classified as lithography machinery. Second, declared value may also differ from ASML pricing because of freight, insurance, configuration, transfer pricing, or the treatment of system modules. Third, the reported province is normally the registration location of the importer or consignee, which may differ from the system's ultimate installation location. Finally, known shipments cover only 19 of 22 reported NXT:2050i systems shipped to Chinese fabs, leaving three unassigned. Those systems may remain embedded in the residual estimates above.

We account for uncertainties in customs accounting using multipliers with quantiles [0.9, 1.0, 1.05] but do not attempt to account for incomplete identification of advanced systems or errors in reported provinces.

Consolidated estimate

The consolidated model selects the strongest available source for each period and triangulates between sources when no direct shipment is available. For 2012–2014, before customs data is available, we use the ASML financial estimate alone. For 2021–2024, we use the Select Committee's shipment numbers (Table 38). These are the most direct public estimates available and distinguish NXT:1980i, NXT:2050i, and NXT:2100i, and are not modeled with uncertainty. For 2015–2020 and 2025–2026Q1, we combine the ASML financial estimate and the customs estimate. Over their common 2015–2026Q1 span, the ASML and customs methods broadly align, at 294 / 379 / 494 and 317 / 389 / 501, respectively. We combine them as follows:

$$\text{Shipment}_{\text{combined}} = \underbrace{w_{\text{customs}}}_{\substack{\text{The weight on the customs estimate, modeled with uncertainty, with quantiles [0.25, 0.5, 0.75]}}} \times \underbrace{\text{Shipment}_{\text{customs}}}_{\text{The customs estimate}} + [1 - w_{\text{customs}}] \times \underbrace{\text{Shipment}_{\text{ASML}}}_{\text{The ASML financial estimate}}$$

We use a public estimate from November 2023, which claims that China had no more than 80 NXT:1980i systems at that time, to estimate installations prior to 2021 (the first year for which we have data from the Select Committee), with modeled uncertainty [0.25, 0.5, 0.75] as to the fraction of known 2023 shipments incorporated into the estimate.[233] All remaining pre-2021 systems are classified as NXT:1970i and older.

For 2021–2024, the Select Committee estimates represent Chinese-owned fabs. For other years, we subtract estimated imports to multinational fabs (described below) from the gross estimate, year by year within each Monte Carlo draw. For 2012–2017, we deduct multinational imports from the NXT:1970i and less advanced bucket, and from the NXT:1980i bucket beginning in 2018. This allocation is an accounting convention that reflects our modeling of China's purchases during these periods rather than an independent estimate of the scanner mix at multinational fabs.

Installations at key fabs

We estimate installations at key fabs involved in AI chip production, using public shipment records and reporting. We examine three firms in depth: SMIC, China's flagship logic foundry; CXMT, China's flagship HBM maker; and YMTC, China's flagship NAND maker. YMTC is relevant because it is a major purchaser of advanced lithography systems; understanding its fleet helps bound how many of the most advanced NXT:2050i/2100i systems remain available to logic and DRAM producers.

CXMT

Table 29: DUVi systems installed at CXMT

Site	NXT:1980i	NXT:2050i	NXT:2100i	Total
Hefei	12 / 25 / 33	1/2/3	4/5/6	19 / 32 / 42
Beijing	6 / 8 / 11	0/1/2	0/1/2	8 / 10 / 15
Total	20 / 33 / 44	3	6	29 / 42 / 53

We estimate NXT:2100i and NXT:2050i installations by cross-referencing shipments reported by the Select Committee (Table 38) with a 2025 presentation from Jiufu Technology Group (Table 40), a contractor that supplies lithography scanner foundations.[234] Jiufu reports delivering 2 foundations classified as NXT:2100i to CXMT Beijing and 7 NXT:2100i foundations to CXMT Hefei. However, the Select Committee reports that only six NXT:2100i systems have entered China in total. The most likely explanation is that the foundations reported by Jiufu can be used for either NXT:2050i or NXT:2100i. We treat the nine foundations as supporting six NXT:2100i and three NXT:2050i, with uncertainty over allocation between Hefei and Beijing.

For NXT:1980i, Jiufu reports 12 delivered foundations at Hefei, 8 delivered foundations at Beijing, and 13 outstanding Hefei orders (without specifying the model). We set:

- Hefei P50 at 25, equal to 12 delivered orders plus the 13 outstanding orders, assuming all of these became NXT:1980i deliveries. This estimate is supported by Anhui imports and the capex cross-check below. P10 is set at 12, allowing for the possibility that none of the orders were NXT:1980i. P90 is set at 33, consistent with the upper end of the customs-based estimate.

- Beijing P50 at 8, equal to Jiufu's reported foundation deliveries. We set P90 at 11 to allow for additional deliveries since the 2025 presentation and P10 at six in case some foundations are used for other NXT systems.

As a cross-check, we independently estimate NXT:1980i based on CXMT's $32.2 billion capex on gross machinery and equipment as of 2025 (Table 39). First, we assume that ~90% of CXMT's machinery is front-end wafer fabrication equipment (WFE).[139] Second, we estimate that lithography accounts for approximately 19% of China's WFE spending.[140] Multiplying the factors, we estimate that lithography accounts for ~17% of CXMT's machinery capex. Next, we estimate that immersion accounts for 75% of this capital, which matches ArFi's share of ASML's DUV scanner revenue. Applying this share to CXMT's estimated capex implies $4.1 billion of capital on DUVi systems.

From the DUVi capital stock, we subtract nine reported NXT:2050i/2100i systems at $90 million each, allocate 5% for installation and capitalized upgrades, and assume 95% of the remainder corresponds to NXT:1980i rather than predecessors. Dividing by NXT:1980i's $76 million ASP, we estimate 39 NXT:1980i systems, a close match to the estimate from the Jiufu presentation.

One overarching uncertainty is whether CXMT continues to receive NXT:1980i systems to this day. The Dutch government requires a license for these systems nationwide, but the Select Committee's data confirm that they continue to enter the country.[235] The Anhui customs data provides some evidence that CXMT has been able to obtain NXT:1980i systems in recent years.[141]

SMIC

Table 30: DUVi systems installed at SMIC

Site	NXT:1980i	NXT:2050i	Total
Shanghai	19 / 24 / 28	6/7/8	25 / 31 / 36
Beijing	6 / 9 / 12	0	6 / 9 / 12
Total	25 / 33 / 40	6/7/8	31 / 40 / 48

We estimate that SMIC has 6 / 7 / 8 NXT:2050i systems at SMIC South in Shanghai, based on the following.

[139] SEMI reports that WFE accounted for $104.3 billion, or 89% of the industry's $117.1 billion in total equipment sales.

[140] ASML reported €10.2 billion of total Chinese revenue in 2024, but this includes service, field options, and metrology equipment. Applying ASML's global revenue composition (77% system sales, of which 97% was lithography) implies €7.6 billion. Dividing this by estimated Chinese WFE spending of approximately $44 billion produces an approximately 19% scanner share.

[141] Some of these systems could have gone to NexChip, but the volume of imports, together with Jiufu's 13 outstanding CXMT Hefei orders, strongly suggests that a substantial share went to CXMT.

1. The Select Committee on the CCP reports that ASML sold 22 NXT:2050i systems to China during 2021–2023 and none in 2024.[236] Public reporting indicates that two of these were delivered to SMIC South in 2021.[237]
2. A September 2023 equipment ledger reproduced on the Sino Defence Forum reported four NXT:2050i systems at SMIC South—two delivered in 2021 and two in June 2023—and the accompanying discussion identified another three systems scheduled for delivery by year-end.[238]
3. A later Sigmaintell report stated that SMIC acquired three NXT:2050i systems in 2024.[239] Because the Dutch government reportedly began denying licenses for NXT:2050i beginning January 2024, we interpret these as systems imported in late 2023 but installed in 2024.

We estimate SMIC's NXT:1980i fleet drawing from a September 2023 SemiAnalysis report claiming that SMIC already possessed "well over 30" advanced ASML ArFi systems and had ordered additional NXT:1980i systems.[240] SemiAnalysis also estimated that SMIC's first 7nm fab would contain 15 ArFi systems when fully equipped by the end of the second quarter of 2024. We translate this into approximately 30 / 36 / 42 total ArFi systems in September 2023, rising to 31 / 40 / 48 after allowing for tools already ordered or undergoing installation before January 2024, and subtract the estimated NXT:2050i systems above.

The main uncertainty is whether SMIC continued to receive NXT:1980i systems after 2023. SMIC's principal Beijing and Shanghai entities—including SMIC Beijing, SMIC North, SMIC Shanghai and SMIC South—are designated with Footnote 5 on the US Entity List and barred from receiving NXT:1980i by US rules.[241] However, although the Netherlands has required a license for NXT:1980i, the Dutch government does not publicly disclose its licensing decisions for individual customers or fabs. Our estimate assumes that SMIC's Beijing and Shanghai fabs have not obtained additional NXT:1980i systems from 2024 onward. If the Netherlands has granted further licenses, the installed base will be higher.

The geographic allocation is less certain. The same Sino Defence Forum equipment ledger identified six NXT:1980i systems at SMIC Beijing and six at SMIC South, alongside five purported NXT:2000i systems at SMIC South.[242] We use six Beijing systems as the lower estimate and assume the rest of the remaining fleet is allocated to Shanghai, where SMIC's advanced production is concentrated.

YMTC

Table 31: DUVi systems installed at YMTC

Tool	Number
NXT:2050i	5 / 7 / 9

We estimate that YMTC has 5/7/9 NXT:2050i systems, based on five considerations:

1. Supportive (though limited) direct evidence. A February 2022 industry-forum post claims that YMTC had obtained NXT:2050i systems, and contemporaneous Chinese reporting identifies ASML as a major supplier.[243]
2. The national allocation of NXT:2050i systems. We estimate that China imported 22 NXT:2050i in total, with 6 / 7 / 8 estimated for SMIC, three for CXMT, and one for Hua Hong. This leaves 11 unaccounted for.
3. YMTC's status as a major purchaser of WFE, at a scale comparable to SMIC and CXMT. YMTC's gross machinery capital at year-end 2023, excluding Wuhan XMC, was RMB120 billion, 25% greater than CXMT's and 60% as large as SMIC's.
4. The timing of YMTC's capacity expansion. Fab 1 ramped during 2021 and Fab 2 ramped during 2022–2023, closely matching the period during which China imported NXT:2050i.
5. YMTC's manufacturing scale. Fab 1 reached 100,000 wafer starts per month (WSPM), close to the 120,000 WSPM scale of CXMT's Phase I project, for which we estimate CXMT acquired 9 NXT:2100i and NXT:2050i systems.

In sum, YMTC's scale and the timing of its expansion make it the leading candidate for bulk purchases of NXT:2050i in the relevant period. We estimate it acquired the majority of the 11 remaining systems, matching the scale of SMIC and CXMT, with P10 and P90 bounds to reflect uncertainty as to the specifics of its purchases.

Multinational fabs

We estimate imports at four known multinational fabs operating within China: Samsung Xi'an (Shaanxi province), SK hynix (formerly Intel) Dalian (Liaoning province), and SK hynix Wuxi and TSMC Nanjing (both Jiangsu province). For 2015–2026Q1, we use the customs model above to isolate imports to these provinces, which are a plausible upper bound on imports to the relevant fabs. We do not attempt to identify system models.

We assign all inferred Shaanxi and Liaoning immersion imports to the multinational fabs, because of an absence of evidence of other plausible DUVi purchasers in those provinces. Jiangsu has several fabs plausibly capable of receiving DUVi systems, including Hua Hong's Fab 7 and Rong Semiconductor's Huai'an line. We set Jiangsu's multinational fab share at 70/85/100%, which is drawn once per simulation and shared across years.

For 2012–2014, before the customs series begins, we estimate 4 / 6 / 8 systems at SK hynix Wuxi and 1 / 3 / 6 systems at Samsung Xi'an. Together, these represent roughly half of the period's estimated 12 / 15 / 20 national imports. We allocate more to Wuxi because it was already producing 29nm DRAM, for which DUVi is used, while Samsung Xi'an began production only in 2014. Both estimates are capped at the national estimate, based on ASML financial data, within each draw. We did not find evidence that the other two fabs ramped during this time period.

We obtain an estimate of 42 / 52 / 67 systems across fabs.

Other known installations

Public sources identify several additional installations beyond SMIC, CXMT, and YMTC. These are not estimates of the entities' complete installed bases and should not be summed across sources. Notably, some of the sites listed by Jiufu are anonymous and might therefore overlap with previously discussed SMIC/CXMT fabs or the separate Hua Hong observation. The allocation in the summary table assumes these observations are distinct; any overlaps would increase the unallocated NXT:1980i residual.

Table 32: DUVi systems installed at other known installations

Entity	Installation
Hua Hong Fab 6	1 NXT:1980i, 4 additional unknown DUVi 1 NXT:2050i (Sino Defence Forum)[244]
JHICC	At least 1 DUVi of unknown model (China BaoGao)[245]
Other Jiufu-listed fabs	30 NXT:1980i foundations: Shanghai 12" #1 (5), PST Shenzhen (10), Chuxing Wuhan (4), UNT Shaoxing (1), NAURA "N8" line Beijing (1), XMC Wuhan (1), Jihai Hangzhou (1), Guangzhou 12" #2 (1), Chengdu 12" (1), and Shanghai 12" #2 (5)

Allocation to advanced fabs

Building on the estimates above, we estimate the fraction of China's DUVi stockpile that is installed within fabs producing advanced logic or DRAM suitable for AI chips. These estimates are highly uncertain, especially for NXT:1980i, because of uncertainties regarding (i) the total number and locations of China's advanced fabs and (ii) the implementation of Dutch licensing policy on exports to these fabs.

Table 33: DUVi system allocation to Chinese advanced fabs

NXT Tool	Logic	DRAM	Total
NXT:2100i	0	6	6
NXT:2050i	9 / 10 / 11	4 / 5 / 6	13 / 15 / 17
NXT:1980i	37 / 57 / 89	29 / 44 / 62	74 / 101 / 142

For NXT:2050i and NXT:2100i, we use the allocations estimated above: 5 / 7 / 9 NXT:2050i systems at YMTC (assumed to be used for NAND production and excluded from the totals above), six NXT:2100i and three NXT:2050i at CXMT, 6 / 7 / 8 NXT:2050i at SMIC, and one NXT:2050i at Hua Hong Fab 6. This leaves an estimated 2 / 4 / 6 NXT:2050i unaccounted for. We assume that these systems are located at advanced fabs, rather than legacy fabs, because they (i) are unnecessary for legacy production, (ii) cost more than earlier systems that enable legacy production, and (iii) were primarily imported in 2023, when

Chinese buyers anticipated impending countrywide restrictions. We assume an even split between logic and DRAM.

For NXT:1980i, no comparable fab-level allocation exists. We estimate the advanced-node fleet by running a Monte Carlo simulation over five components:

- **SMIC's Shanghai fabs (19 / 24 / 28) and CXMT (20 / 33 / 44)**, based on the estimates above. We exclude SMIC Beijing's 6 / 9 / 12 systems, because these fabs are believed to produce only at legacy nodes.
- **Guangdong advanced fabs (3 / 6 / 22)**. Several Huawei-linked fabs are located in Guangdong, including PXW, a rumored 7nm facility in Shenzhen, and SwaySure's DRAM lines. Guangdong's estimated NXT:1980i-priced imports (32 / 41 / 54) provide an upper bound on the installed base at these facilities; subtracting the 10 NXT:1980i foundations that Jiufu reported as delivered to PST Shenzhen—a legacy logic (40nm / 28nm) fab—leaves 22 / 31 / 44. There is further uncertainty as to how many systems were obtained by SMIC's 28nm Shenzhen fab, how many of PST Shenzhen's remaining 17 orders were NXT:1980i, and what other legacy facilities might have received orders.
- **Shanghai's advanced logic fabs other than SMIC (5 / 19 / 49)**. Shanghai's estimated NXT:1980i-priced imports total 77 / 97 / 127. We attribute 19 / 24 / 28 to SMIC Shanghai. Subtracting the central estimate of 24 leaves 53 / 73 / 103, including advanced facilities like Hua Hong's Fab 6 (reported to hold at least one NXT:1980i, one NXT:2050i, and four unidentified DUVi) and Fab 8a, which is reported to have a 7nm line. This total also includes imports to legacy fabs such as SMIC Lingang and Hua Hong's 300mm Fab 5. We use a wide P90 to reflect the uncertainty in the division of these tools between advanced and legacy fabs.
- **Beijing's advanced fabs other than SMIC and CXMT (1 / 2 / 22)**. Beijing's estimated imports total 50 / 63 / 81. Subtracting the SMIC Beijing and CXMT Beijing allocations within draws leaves 32 / 46 / 64. We are not aware of other Beijing fabs with advanced logic or DRAM production, and it is possible that these systems all went to legacy fabs (e.g., SMIC Jingcheng) but there is still substantial uncertainty. We estimate a low median but leave significant upside uncertainty, reflected in the P90.
- **Advanced fabs in other provinces (1 / 3 / 30)**. The only documented candidate is JHICC in Fujian, a DRAM fab reportedly participating in Huawei's HBM consortium. Collectively, provinces not listed in Table 28 imported an estimated 34 / 43 / 57 systems, including a 2024 spike of 17 / 22 / 29. Our P90 allows for the possibility that most of this spike went to advanced fabs not known to the authors.

Data sources

Table 34: ASML's financial data since 2012

	2026Q1	2025	2024	2023	2022	2021	2020	2019	2018	2017	2016	2015	2014	2013	2012
EUV % worldwide sales	66	48	38	42	46	46	43	31	23	17	7	2	7	2	0
ArFi % worldwide sales	23	42	44	41	34	36	38	53	58	63	77	76	80	78	71
ArFi total shipments	17	131	129	125	81	81	68	82	86	76	70	67	76	77	72
China % worldwide sales	19	33	41	29	14	16	18	12	19	11	15	10	9	8	2

Table 35: Lithography system average sales price by year

Year	DUVi ASP ($M)			Blended dry ArF + KrF ASP ($M) (from ASML's financial statements)
	Blended (from ASML's financial statements)	Modeled NXT:1980i	NXT:2050i/2100i	
2015	54			
2016	56			
2017	60			
2018	66	63[142]	–	12
2019	64	65	–	11
2020	66	67	–	11
2021	72	70	–	12
2022	68	72	–	11
2023	78	74	90[143]	13
2024	80	76	–	14
2025	89	85	–	13
2026Q1	99	95	–	14

[142] In 2016 and 2017, ASML recognized €3.540 billion and €4.029 billion of ArFi revenue on 68 and 76 systems, respectively, including 46 and 66 NXT:1980i.[246] Solving the two revenue equations while assuming stable ASP through 2018 yields an NXT:1980i ASP of €53.7 million, or $63.4 million.

[143] Contemporaneous reporting claims that NXT:2050/2100i systems sold to China in 2023 cost €80 million to €90 million (~$90M at 2023 exchange rates).

Table 36: Lithography system (HS 84862031 and 84862039) import quantities, from China's customs data

Province	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026Q1	Total since 2015
Shanghai	11	5	2	18	16	21	34	31	51	58	79	9	335
Jiangsu	13	1	10	19	17	26	47	49	11	38	22	5	258
Shaanxi	17	5	0	2	8	12	12	0	1	2	2	0	61
Liaoning	1	8	6	16	0	1	1	0	0	0	0	0	33
Beijing	2	9	6	1	5	15	19	12	35	46	21	8	179
Hubei	0	1	1	9	8	13	18	9	34	23	19	3	138
Guangdong	0	0	1	2	5	5	8	14	28	41	37	6	147
Anhui	0	0	0	8	6	10	8	13	23	18	8	1	95
Other	1	6	7	12	2	18	21	18	42	49	19	9	204
China total	45	35	33	87	67	121	168	146	225	275	207	41	1,450

Table 37: Lithography system (HS 84862031 and 84862039) import value (USD $M), from China's customs data

Province	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026Q1	Total since 2015
Shanghai	125.3	29.9	92.2	340.9	243.7	876.2	730.9	561.0	1,507.0	2,004.7	4,105.2	704.7	11,321.7
Jiangsu	206.9	16.6	308.3	669.4	527.4	728.5	511.2	579.9	298.7	594.3	303.7	59.5	4,804.5
Shaanxi	303.4	54.2	0.0	108.2	179.2	194.2	278.4	0.0	13.2	22.3	16.5	0.0	1,169.6
Liaoning	1.8	223.0	166.2	399.0	0.0	4.3	10.5	0.0	0.0	0.0	0.0	0.0	804.9
Beijing	18.4	287.6	130.9	0.1	23.4	191.6	473.1	352.6	1,225.1	1,965.0	1,442.1	142.7	6,252.6
Hubei	0.0	0.7	3.1	215.9	263.3	261.0	398.3	281.5	1,424.4	861.1	1,096.4	40.6	4,846.3
Guangdong	0.0	0.0	19.6	2.4	45.7	53.5	74.8	114.5	1,019.9	969.1	1,948.4	260.7	4,508.6
Anhui	0.0	0.0	0.0	242.1	199.2	152.1	253.9	493.0	985.1	1,083.2	332.9	108.7	3,850.2
Other	4.6	186.8	163.3	283.4	83.3	128.9	469.4	165.0	756.9	2,015.2	576.1	250.1	5,083.1
China total	660.6	798.8	883.7	2,261.4	1,565.2	2,590.3	3,200.5	2,547.5	7,230.2	9,514.9	9,821.3	1,566.8	42,641.1

Table 38: ASML's DUVi sales (2021–2024) to Chinese-owned fabs,[144] reported by the Select Committee on the CCP[247]

DUVi system	2021	2022	2023	2024	Total
NXT:2100i	0	0	6	0	6
NXT:2050i	2	4	16	0	22
NXT:1980i	14	17	35	90	156

Table 39: CXMT's equipment capital stock and expenditures (RMB billions), from its IPO prospectus and audit[248]

Year	Fixed-asset machinery additions	Gross machinery and equipment at year-end
2022	38	73
2023	24	96
2024	82	174
2025	Not separately presented	226

Table 40: Jiufu Technology Group's reported orders and deliveries of lithography scanner foundations, by fab[249]

Fab	Ordered	Delivered	NXT2100	NXT1980	NXT1470	NXT870	XT1460	XT1060	XT860	XT400
CXMT Hefei	61	48	7	12	8	16	–	–	–	5
CXMT Beijing	24	24	2	8	2	8	–	–	–	4
Shanghai 12" #1	35	23	–	5	–	–	1	–	1	2
PST Shenzhen	33	16	–	10	–	5	–	–	–	1
Chuxing Wuhan	35	27	–	4	–	8	–	–	–	4
Tianjin 12"	7	7	–	–	–	–	1	–	3	1
UNT Shaoxing	4	4	–	1	–	–	–	–	1	–
NAURA "N8" line Beijing	1	1	–	1	–	–	–	–	–	–
XMC Wuhan	8	8	–	1	–	–	–	–	4	–
Jihai Hangzhou	5	5	–	1	–	–	1	1	1	1
Guangzhou 12" #1	10	7	–	–	–	–	–	–	–	–

[144] The data appears to exclude DUVi sales to multinational fabs in China given the focus on Chinese chipmakers.

Fab	Ordered	Delivered	NXT2100	NXT1980	NXT1470	NXT870	XT1460	XT1060	XT860	XT400
Guangzhou 12" #2	6	6	–	1	–	–	1	–	1	–
Hefei 12"	2	2	–	–	–	–	–	–	–	–
Chengdu 12"	10	9	–	1	–	1	–	–	–	–
Shanghai 12" #2	10	9	–	5	–	2	–	1	–	1
Wuxi 12"	1	1	–	–	–	–	–	–	1	–
Total	252	197	9	50	10	40	4	2	12	19

Appendix F: Lithography-bounded Chinese AI chipmaking

This appendix compares projected AI chip production in the US and allied ecosystem through 2035 with bounds on Chinese production imposed by limitations on access to DUVi lithography systems. For China, we model production as bounded by its existing ASML scanner stockpile, potential future ASML imports, and indigenously produced scanners. We evaluate five import environments: imports banned beginning in 2027, imports continuing at the 2024–2025 pace, and modeled "Diminished", "Robust", and "Acceleration" trajectories. We also model three allocations of DUVi systems within China: a "Baseline" (also called "current trend") that maintains the current split of China's DUVi system fleet between advanced and legacy fabs; a "Mobilization" (also called "AI push") in which China achieves the allocation found in the US and allied ecosystem; and a "Physical Ceiling" that assigns every usable scanner to AI production.

The China estimate assumes that every complementary chipmaking bottleneck—including etch, deposition, metrology, HBM stacking, and advanced packaging—is eliminated, leaving lithography as the sole constraint. The resulting estimates are therefore lithographic capacity bounds rather than forecasts of realized production. These bounds are unrealistic in the near term, given China faces numerous additional constraints, but they reveal the production capacity enabled by varying policy choices. Additionally, they may become informative in the mid term if China indigenizes complementary SME bottlenecks.

China estimates are reported as P10 / P50 / P90 quantiles from 100,000 Monte Carlo draws, conditional on each import and allocation scenario. US and allied estimates are low, central, and high forecasts. All production figures are reported in millions of H100-equivalents (H100e): a chip's dense FP8 FLOP/s (INT8 where a chip lacks FP8) divided by the H100's 1,979 TFLOP/s. For chips with FP8, one B300 equivalent equals about 2.53 H100e.

We find:

1. **Banning new DUVi system exports ensures a large US and allied advantage in AI chipmaking through 2035.** Even under extreme assumptions, in which China: (1) mobilizes 100% of its in-country DUVi systems for AI chip production, violating existing export controls; (2) reaches world-class yields and throughputs by 2031; (3) indigenizes every complementary tool required for AI chipmaking; and (4) indigenizes DUVi systems capable of matching the NXT:1980i in 2030 (much earlier than our median estimate) and scales up indigenous DUVi system production at an unprecedented rate of 100% year over year, we estimate China's installed DUVi systems will support a maximum of 104 / 182 / 382 million H100-equivalents annually by 2035, of which 79 / 133 / 267 million comes from the ASML stockpile as of 2026Q1, 15 / 25 / 50 million from systems estimated to be imported during 2026Q2–Q4, and 10 / 24 / 66

million from indigenously produced DUVi systems. In the same year, US and allied production is estimated at 578 / 1,252 / 5,404 million. China's P50 estimate is 15% of the US central case.

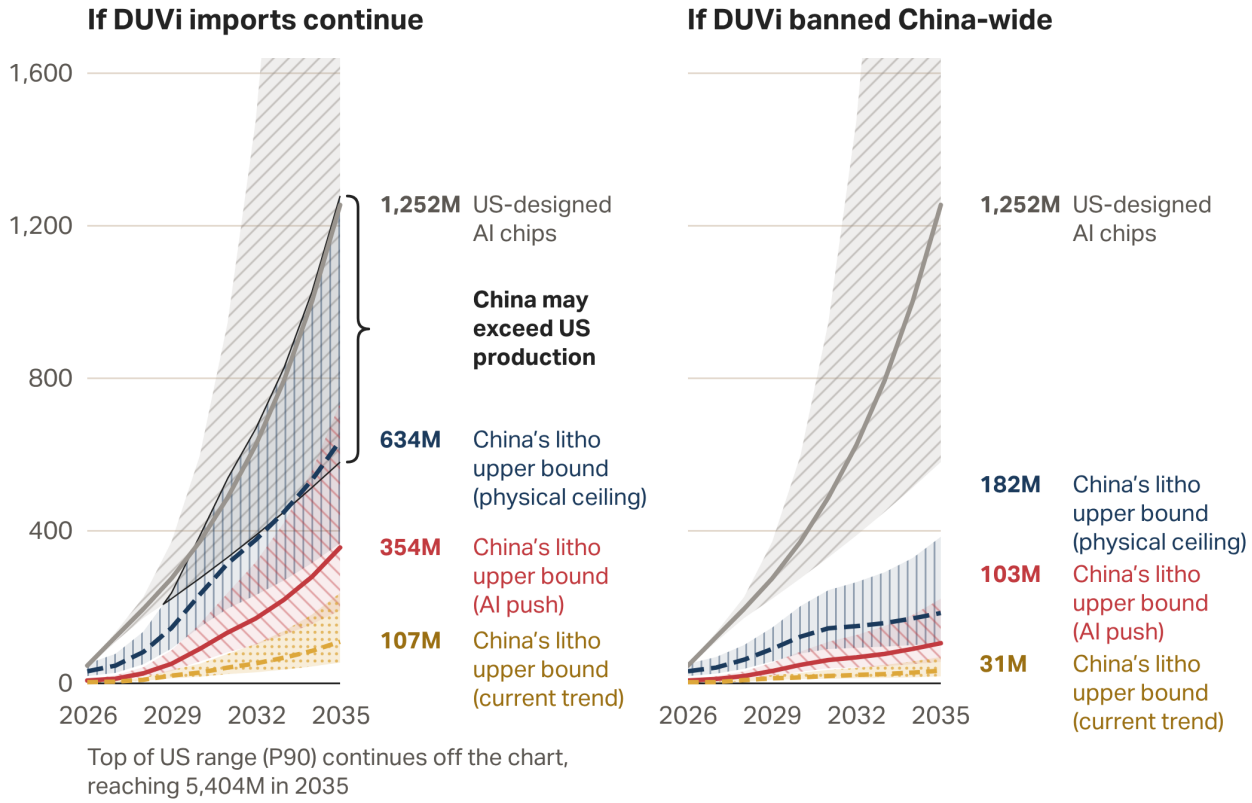
2. **Continued NXT:1980i imports could erode—and under some assumptions eliminate—the US chipmaking advantage.** If imports continue at the 2024–2025 pace, China could reach 154 / 290 / 600 million H100-eq if mobilized for AI production and 297 / 517 / 1,045 million at the Physical Ceiling, by 2035. Assuming a moderate increase in imports over time relative to the 2024–2025 pace (the "Robust" scenario shown in Figures 1 and 19), these rise to 189 / 354 / 733 million under Mobilization and 365 / 634 / 1,276 million at the physical ceiling. These medians would still trail the US median estimate (1,252 million) but could surpass the low estimate (578 million).

Figure 19. Banning DUVi to China secures a US AI chipmaking advantage through 2035

China's existing DUVi fleet cannot match US AI chip production; but continued imports could let it close the gap

Bold lines show median projections (P50); lightly shaded cones span the P10–P90 range

Annual AI chip production (Million H100-equivalents)



Top of US range (P90) continues off the chart, reaching 5,404M in 2035

US chip production is estimated for 2026–2027 from TrendForce shipment forecasts and from 2028 onwards by averaging a demand projection and an EUV supply model. H100-equivalents are dense FP8 FLOP/s (INT8 where a chip lacks FP8) divided by the H100's 1,979 TFLOP/s. China's forecasts are for lithography-bounded production capacity, assuming it indigenizes all other SME. "Current trend" assumes DUVi are allocated between advanced and legacy fabs in the same proportion as today. "AI push" assumes China mobilizes its DUVi fleet to prioritize AI, obtaining the same allocation forecasted in the US ecosystem. "Physical ceiling" assumes every DUVi in-country is used for AI chip production. "DUVi banned China-wide" assumes that DUVi imports stop beginning 2027. "DUVi imports continue" assumes a moderate increase in imports over time relative to China's 2024–2025 pace.

Center for Technology & Statecraft IFP

US ecosystem production model

We estimate annual AI chip production available to the US and allied ecosystem, in millions of H100-equivalents, using market forecasts for 2026 and 2027 and an average of a demand projection and a lithography capacity model for 2028 to 2035.

Table 41: Annual US AI chip production projections

Year	P10	P50	P90
2026	39	44	48
2027	107	119	131
2028	169	193	224
2029	220	273	369
2030	273	369	594
2031	328	485	943
2032	386	623	1,475
2033	447	792	2,289
2034	511	999	3,527
2035	578	1,252	5,404

For 2026 and 2027, we use estimates from TrendForce of 17.2 and 47.1 million GB300-equivalents for US-designed chips, corresponding to ~44 million and ~119 million H100-equivalents using dense FP8 throughput, with 10% uncertainty for P10 and P90. For 2028, the demand projection described below grows 37.5% / 50% / 62.5% over 2027, combining 10% / 20% / 30% chip revenue growth with AI 2040's forecast of 25% annual growth in H100-eq per dollar over 2025–2028 rather than the slower rates used from 2029.[250]

For 2028–2035, we average a demand forecast with an estimate of AI production supported by available EUV capacity. We estimate demand growth at [16 / 32 / 56] percent, as the product of chip revenue growth and growth in H100-eq per revenue dollar.

Table 42: 2029–2035 H100-equivalent demand growth

Scenario	Revenue growth	H100-eq per dollar	H100-eq demand growth
P10	10%	5%	16%
P50	20%	10%	32%
P90	30%	20%	56%

The P50 is anchored to hyperscaler capex growth over 2018–24, reflecting the assumption that chip spending growth returns to the long-run trend for cloud and AI infrastructure after a fast growth period in the 2020s. Capital expenditure at Alphabet, Microsoft, and Meta increased from $50.7 billion to $136.2 billion, implying a 17.9% CAGR. Including Amazon increases this to 22.9%, but includes substantial fulfillment infrastructure. We use a rounded 20%. For yearly improvements to H100-equivalents per dollar, we take AI 2040's forecast that by 2028, compute cost-efficiency growth will have slowed to 10% compared to 25% across 2025–2028 and 40% before that.[251] The P10 assumes 10% revenue growth, close to ASML's 9% semiconductor-market outlook. The P90 assumes 30%, representing continued exceptional expansion. Its 20% yearly improvement in H100-eq per dollar is deliberately optimistic, to capture advances in architecture and performance. The resulting unconstrained demand estimates are:

Table 43: Demand-based annual AI chip production projections

Year	P10	P50	P90
2028	147	179	213
2029	171	236	332
2030	198	312	518
2031	229	412	809
2032	266	543	1,261
2033	309	717	1,968
2034	358	947	3,070
2035	415	1,250	4,789

The EUV model separately estimates gross EUV supply and the residual capacity available for AI after non-AI production, and depends on fleet size, scanner productivity, H100-eq per exposure, and competing non-AI demand.

- EUV fleet size. We estimate approximately 310 installed low-NA EUV systems at the start of 2026.[252] ASML's guidance implies a production capacity of 65 systems in 2026 and a 30% capacity increase to 85 in 2027; it is investigating a further capacity 30% increase, to roughly 110, in 2028. Bank of America projects 85 shipments in 2027 and 87 in 2028.[253] We assume shipments match this capacity in 2026 and 2027 (65 and 85) and are 85 / 95 / 110 in 2028 (P10 / P50 / P90), so that full capacity buildout and utilization in 2028 is the P90 case. After 2028, P10 holds additions at 85 systems annually; P50 and P90 respectively increase by 5% and 30% year over year after 2028. Scanners delivered during a year receive half weight, approximating an evenly distributed delivery schedule.
- EUV exposure throughput. Following SemiAnalysis, we assume that an EUV averages 75 wafer-exposures per hour at 90% uptime, corresponding to ~590,000 exposures per year.[254] We

assume that scanner productivity increases by 0% / 1.5% / 3% annually. The central assumption is conservative relative to marketed improvements to reflect the fact that fleet-average performances improve more slowly as older scanners remain in use. We assume H100-eq per exposure improves at 1% / 3% / 8% annually after 2027, capturing net improvements in performance per unit of lithographic capacity.

- Non-AI EUV demand. We estimate that EUV scanners will perform ~250 million exposures in 2027 (based on an estimated fleet size of (375+460)/2 ≈ 418 scanners and ~599,000 exposures per scanner per year, after 1.5% productivity growth). AI chips use ~60 million (24%) of these: TrendForce's ~119 million H100-eq at ~0.50 exposures per H100-eq (0.49–0.57), which we derive from a chip-by-chip estimate of EUV use in each US-designed AI chip expected to ship in 2027 per TrendForce data.[145] That leaves ~190 million exposures for non-AI demand in 2027, which we assume grows at 6% / 9% / 12% annually. The P50 is anchored on the ASML 9% outlook, while P10 and P90 add uncertainty.

Table 44: EUV-based annual AI chip production projections

Year	P10	P50	P90
2028	190	207	235
2029	269	310	405
2030	348	426	670
2031	427	557	1,076
2032	506	703	1,689
2033	585	867	2,610
2034	663	1,050	3,983
2035	741	1,253	6,019

China ecosystem production capacity model

We model China's lithography-bounded AI chip production capacity under five import environments and three allocation scenarios:

- **Import environments.** All five environments include an estimated 69 systems imported during 2026 but add only the 56 / 58 / 60 systems estimated after the existing stockpile measurement at the end of Q1. "Imports banned" sets imports to zero beginning in 2027. "2024–2025 Pace" fixes imports at 90 systems annually, matching the 2024–2025 pace of imports. "Diminished"

[145] To generate the ~0.50 ratio, in addition to counting AI chips, we also count the host CPUs, NVLink switch chips, and NICs shipped with them. These add about 11%. We also count HBM, which accounts for about half the total; EUV use depends on DRAM nodes, e.g., no EUV layers for Micron's 1β and seven for Samsung's 1c. Our China estimates count only AI chip packages (including logic dies and HBM). We assume China can import uncontrolled CPUs; if it is forced to produce them domestically, then the lithography capacity China has left for AI chips would be lower than we model in this appendix.

assumes that China's imports pull back from the recent pace. "Robust" assumes a moderate increase in imports over time relative to the 2024–2025 pace: 90 systems in each of 2027 and 2028, then 10% annual growth. "Acceleration" assumes an extreme expansion in global lithography capacity, extending ASML's announced 30% capacity ramp indefinitely, while maintaining China's 2024–2025 share of imports.[255] These scenarios are described at greater length, below.

- **Allocation scenarios.** "Baseline" broadly continues the present allocation, as shaped by commercial incentives, Dutch licensing policy, and export control effectiveness. "Mobilization" assumes China reorganizes lithography capacity towards AI production, including potentially by circumventing export controls. "Physical Ceiling" assigns every usable scanner to AI chip production, estimating the upper bound enabled by China's in-country DUVi systems.

The following table reports P10 / P50 / P90 annual output for every combination of scenarios. Contributions from indigenous production and the existing stockpile are Shapley-attributed within the "Robust" import environment; they can vary across import environments because additional scanners may relieve a logic or DRAM bottleneck and thereby change the marginal value of other scanners. The final column sums production from all three sources under the "Robust" import scenario.

Table 45: China's AI chip production (million H100-eq)

Year	Scenario	Production from existing stockpile	Production from indigenous DUVi systems	Production from newly imported DUVi (2026 onwards)					Total (assumes "Robust" imports)
				Imports banned	2024–2025 Pace	Diminished	Robust	Acceleration	
2026	Baseline	0 / 1 / 2	0 / 0 / 0	0 / 0 / 0	0 / 0 / 0	0 / 0 / 0	0 / 0 / 0	0 / 0 / 0	0 / 1 / 2
2026	Mobilization	2 / 5 / 8	0 / 0 / 0	0 / 0 / 0	0 / 0 / 0	0 / 0 / 0	0 / 0 / 0	0 / 0 / 0	2 / 5 / 8
2026	Physical ceiling	15 / 27 / 45	0 / 0 / 0	1 / 1 / 4	1 / 1 / 4	1 / 1 / 4	1 / 1 / 4	1 / 1 / 4	17 / 29 / 49
2027	Baseline	1 / 1 / 2	0 / 0 / 0	0 / 0 / 0	0 / 1 / 1	0 / 1 / 1	0 / 1 / 1	0 / 1 / 1	1 / 2 / 4
2027	Mobilization	5 / 8 / 14	0 / 0 / 0	1 / 2 / 3	1 / 3 / 5	1 / 2 / 4	1 / 3 / 5	2 / 3 / 6	6 / 10 / 19
2027	Physical ceiling	19 / 33 / 56	0 / 0 / 0	4 / 6 / 11	6 / 11 / 19	5 / 10 / 15	6 / 11 / 19	7 / 12 / 21	25 / 44 / 75
2028	Baseline	2 / 5 / 8	0 / 0 / 0	0 / 1 / 1	1 / 2 / 5	1 / 2 / 4	1 / 2 / 5	1 / 4 / 7	4 / 7 / 13
2028	Mobilization	9 / 15 / 24	0 / 0 / 0	2 / 2 / 5	6 / 9 / 16	4 / 6 / 11	6 / 9 / 16	8 / 12 / 20	15 / 24 / 40
2028	Physical ceiling	30 / 50 / 82	0 / 0 / 0	6 / 10 / 15	18 / 30 / 51	13 / 21 / 37	18 / 30 / 51	24 / 39 / 67	48 / 80 / 133
2029	Baseline	6 / 10 / 17	0 / 0 / 0	1 / 1 / 4	5 / 8 / 15	4 / 6 / 11	5 / 8 / 17	7 / 13 / 24	11 / 18 / 33
2029	Mobilization	15 / 25 / 43	0 / 0 / 0	3 / 5 / 8	15 / 23 / 40	9 / 16 / 27	15 / 23 / 41	22 / 35 / 61	30 / 49 / 84
2029	Physical ceiling	46 / 74 / 122	0 / 0 / 0	8 / 14 / 23	42 / 68 / 111	27 / 45 / 74	43 / 69 / 113	63 / 101 / 167	89 / 143 / 235
2030	Baseline	7 / 11 / 20	0 / 0 / 0	2 / 3 / 5	8 / 14 / 27	5 / 9 / 17	8 / 15 / 28	13 / 24 / 47	15 / 26 / 49
2030	Mobilization	24 / 39 / 67	0 / 0 / 1	4 / 7 / 13	29 / 47 / 82	19 / 30 / 52	31 / 49 / 85	50 / 82 / 142	55 / 88 / 153

Year	Scenario	Production from existing stockpile	Production from indigenous DUVi systems	Production from newly imported DUVi (2026 onwards)					Total (assumes "Robust" imports)
				Imports banned	2024–2025 Pace	Diminished	Robust	Acceleration	
2030	Physical ceiling	64 / 101 / 168	0 / 0 / 1	12 / 19 / 31	77 / 123 / 204	49 / 79 / 130	81 / 129 / 214	135 / 213 / 356	145 / 230 / 384
2031	Baseline	8 / 14 / 27	0 / 0 / 0	2 / 3 / 6	13 / 23 / 45	8 / 14 / 27	13 / 24 / 48	25 / 45 / 90	22 / 39 / 75
2031	Mobilization	30 / 50 / 88	1 / 1 / 2	6 / 9 / 17	45 / 74 / 133	28 / 46 / 83	49 / 80 / 145	91 / 150 / 271	79 / 131 / 235
2031	Physical ceiling	73 / 118 / 201	1 / 1 / 2	14 / 23 / 38	110 / 176 / 303	68 / 110 / 188	120 / 193 / 331	223 / 359 / 617	194 / 312 / 535
2032	Baseline	9 / 16 / 31	0 / 0 / 1	2 / 4 / 7	15 / 30 / 61	10 / 18 / 37	18 / 34 / 69	37 / 71 / 145	27 / 50 / 101
2032	Mobilization	32 / 55 / 102	1 / 2 / 4	6 / 10 / 19	58 / 98 / 184	35 / 60 / 112	66 / 112 / 210	138 / 235 / 439	99 / 168 / 316
2032	Physical ceiling	75 / 122 / 216	1 / 4 / 7	14 / 23 / 41	133 / 218 / 388	81 / 133 / 236	152 / 249 / 443	317 / 521 / 926	229 / 374 / 666
2033	Baseline	10 / 17 / 36	0 / 1 / 3	2 / 4 / 7	20 / 39 / 81	12 / 24 / 49	24 / 47 / 96	56 / 108 / 227	34 / 65 / 135
2033	Mobilization	35 / 61 / 117	2 / 4 / 8	6 / 11 / 22	72 / 127 / 245	44 / 76 / 148	87 / 153 / 294	205 / 358 / 692	124 / 217 / 419
2033	Physical ceiling	76 / 125 / 231	4 / 7 / 15	14 / 24 / 43	157 / 263 / 484	95 / 158 / 291	188 / 314 / 579	442 / 739 / 1,365	268 / 447 / 825
2034	Baseline	10 / 19 / 41	1 / 3 / 6	2 / 4 / 8	25 / 49 / 105	15 / 30 / 63	31 / 61 / 131	81 / 163 / 350	42 / 83 / 178
2034	Mobilization	38 / 68 / 135	3 / 7 / 17	7 / 13 / 25	89 / 161 / 320	54 / 96 / 191	113 / 202 / 404	298 / 537 / 1,072	153 / 277 / 555
2034	Physical ceiling	77 / 130 / 248	6 / 13 / 31	14 / 25 / 48	182 / 310 / 588	110 / 185 / 353	230 / 388 / 746	608 / 1,034 / 1,979	313 / 531 / 1,024
2035	Baseline	11 / 22 / 46	2 / 4 / 12	3 / 5 / 9	30 / 61 / 134	18 / 37 / 79	39 / 81 / 177	118 / 243 / 534	52 / 107 / 235
2035	Mobilization	41 / 75 / 153	6 / 14 / 37	7 / 14 / 29	107 / 201 / 409	65 / 120 / 244	143 / 265 / 542	427 / 799 / 1,638	189 / 354 / 733
2035	Physical ceiling	79 / 133 / 267	10 / 24 / 66	15 / 25 / 50	208 / 360 / 716	125 / 214 / 423	276 / 476 / 943	827 / 1,429 / 2,851	365 / 634 / 1,276

China's annual AI chip production for a given year is modeled as:

$$\underbrace{\text{H100-equivalent production}}_{\text{China's annual AI chip production for a given year, in H100-equivalents}} = (\underbrace{\text{Existing Stockpile}}_{\text{China's stockpile of DUVi systems as of 2026Q1}} + \underbrace{\text{Indigenous Production}}_{\text{The sum of indigenously produced scanners equivalent to the NXT:1980i}} + \underbrace{\text{New Imports}}_{\text{The sum of NXT:1980i-equivalent systems imported from ASML after 2026Q1 through the given year}}) \times \underbrace{\text{AI allocation}}_{\text{The share of scanners devoted to advanced logic and HBM for AI chips}} \times \underbrace{\text{Physical chip throughput}}_{\text{The rate at which each scanner enables production of complete chip packages}} \times \underbrace{\text{Performance per chip}}_{\text{Computing power per chip package, in H100-equivalents}}$$

where:

- Existing Stockpile is China's stockpile of DUVi systems as of 2026Q1.
- Indigenous production is the sum of indigenously produced scanners equivalent to the NXT:1980i.
- New Imports sums NXT:1980i-equivalent systems imported from ASML after 2026Q1 through the given year.
- AI allocation is the share of scanners devoted to advanced logic and HBM for AI chips.
- Physical chip throughput is the rate at which each scanner enables production of complete chip packages, incorporating scanner exposure throughput, lithography passes per wafer, dies per wafer, frontend die yield, HBM stacking yield, and packaging yield. It is calibrated using the Ascend 910C, as a stand-in for a reference chip with a comparable fabrication footprint. We sample 2026 throughputs of NXT:1980i and NXT:2050i/2100i systems from the SMIC / CXMT distributions in Appendix D, then increase them smoothly to the mature levels (TSMC / SK hynix) by 2031. Where the fleet is inefficiently allocated between logic and memory, output is limited by the smaller capacity.
- Performance per chip captures computing power, modeled in H100e, provided by the reference chip package: 0.76 H100e for the Ascend 910C (1,504 dense INT8 TOPS ÷ 1,979). It captures performance improvements that are independent of lithographic capability (e.g., better chip architecture). We model it as a 1% / 3% / 8% annual increase, matching the US production model's factor.

The implementation separately models NXT:1980i and NXT:2050i/2100i scanners, and separately tracks logic and HBM. The multiplicative expression above is a conceptual simplification.

Existing Stockpile

China's existing DUVi system stockpile is estimated in Appendix E as 249 / 270 / 294 NXT:1980i systems, 22 NXT:2050i systems, and six NXT:2100i systems, imported by Chinese-owned fabs through 2026Q1. We exclude NXT:1970i and less advanced systems, which we have not established as capable of making advanced logic and HBM. This assumption does not meaningfully affect our top-line conclusions, as these systems are a minority of China's installed base and are less capable than the scanners included in the model.

Indigenous production

Indigenous production is modeled using two independently sampled factors.

- Commercial-equivalence year: the first year in which China produces a commercially viable scanner capable of matching NXT:1980i's performance for advanced AI chip production, modeled at 2030 / 2033 / 2036. The P50 reflects the extrapolations from ASML's history and China's past progress described in Section 3, while leaning to the earlier end of the plausible range. P10 allows for faster progress, potentially enabled by massive state investment. P90 allows for some technical challenges to remain intractable for substantially longer.

- Production ramp: production begins at five systems, doubles to 10 and 20 in the next two years (reflecting rapid scale-up from a very low initial base), and then grows at 30% / 50% / 100% annually. The 30% case resembles ASML's announced near-term capacity expansions; the 50% case is an exceptionally fast but historically grounded ramp (a CSET review of the industry's longer history separately estimated that 60% annual growth had been achieved only twice in the lithography industry's history: first at Nikon and later at ASML); the 100% case assumes unprecedented mobilization.[256]

New Imports

We model five import scenarios. All scenarios assume China imports 69 systems during 2026, the midpoint between an extrapolation of its 2026Q1 pace and the 2024–2025 pace.

Table 46: Scenarios for China's DUVi system imports

Scenario	Interpretation	China's share of worldwide shipments	Worldwide shipment growth after 2028
Imports banned	Imports fall to zero.	0%	Not applicable
2024–2025 Pace	China's 2024–2025 import pace.	Not modeled	Not modeled
Diminished	Imports pull back from the 2024–2025 surge.	25%	0%
Robust	Imports increase moderately over time relative to the 2024–2025 pace.	40%	10%
Acceleration	ASML's announced near-term capacity ramp persists through 2035 while China retains its recent shipment share.	70%	30%

Table 47: China's annual NXT:1980i-equivalent imports

Year	Imports banned	2024–2025 Pace	Diminished	Robust	Acceleration
2027	0	90	55	90	120
2028	0	90	50	90	150
2029	0	90	50	99	195
2030	0	90	50	109	254
2031	0	90	50	120	330
2032	0	90	50	132	428
2033	0	90	50	145	557
2034	0	90	50	159	724
2035	0	90	50	175	941

We anchor near-term estimates to (i) China's roughly 90 NXT:1980i imports over 2024–2025 and (ii) ASML's planned immersion capacity expansions: 130 systems in 2026, a 30% expansion in 2027, and a potential 30% expansion in 2028, to roughly 220 systems.[257] The 2028 estimates correspond to China receiving 25% / 40% / 70% of that output. The "Diminished" scenario settles slightly above China's pre-2023 share, "Robust" is near the midpoint between its pre- and post-2023 shares, and "Acceleration" preserves its approximately 70% share over 2024–2025.

After 2028, we model imports as a function of two factors:

- **Growth in worldwide DUVi system shipments** (0% / 10% / 30%), matching EUV growth estimated in the US production model. "Robust" is anchored to 10% growth, consistent with ASML's projection that DUV wafer exposures will rise from ~525 million in 2024 to 900 million in 2030. "Diminished" assumes the 2026-28 expansion is followed by a flat cycle, with growing

exposure demand absorbed by the enlarged installed base. "Acceleration" mechanically extends ASML's announced capacity ramp through 2035. The resulting scale is extreme, but useful for contemplating the consequences if ASML's near-term expansion becomes a sustained trend.

- **China's share of worldwide DUVi system shipments** (25% / 40% / 70%). "Diminished" settles slightly above its pre-2023 share. "Acceleration" assumes China maintains the 2024–2025 share. "Robust" assumes the share settles near the midpoint between the pre- and post-2023 imports.

Table 48: China's historical shares of global DUVi system output

Year	ASML global DUVi system shipments	Estimated Chinese DUVi system imports	China share of global DUVi system shipments
2018	86	14	16%
2019	82	15	18%
2020	68	22	32%
2021	81	16	20%
2022	81	27	33%
2023	125	51	41%
2024	129	90	70%
2025	131	92	70%

Imports receive half weight in their delivery year, reflecting evenly distributed arrivals, and full weight thereafter.

AI Allocation

We model three scenarios for allocation of DUVi systems to advanced-node AI production.

Table 49: Scenarios for allocation of China's DUVi systems

Scenario	Interpretation	Modeled allocation
Baseline	Continuity with the present allocation, reflecting commercial incentives, Dutch licensing policy, enforcement efforts, and circumvention efforts.	Existing systems remain where currently installed. New foreign and indigenous systems enter advanced fabs at the current 27% / 38% / 53% share. Within advanced fabs, AI use rises from 5% / 10% / 20% in 2026 to the outside-China allocation in 2030, then tracks it through 2035.
Mobilization	China mobilizes to prioritize AI and circumvents restrictions on NXT:1980i locations within China	A single factor models allocation to AI production, which rises to the outside-China share in 2030 (34% / 38% / 44%) then tracks it through 2035 (42% / 57% / 70%)
Physical ceiling	Upper physical limit on Chinese AI output	100% of existing, imported, and indigenous systems are devoted to AI production

Current allocations are estimated as follows:

- **Allocation to advanced fabs.** We adopt the estimates from Appendix E. Advanced logic is estimated at 37 / 57 / 89 NXT:1980i systems and 9 / 10 / 11 NXT:2050i systems. Advanced DRAM is estimated at 29 / 44 / 62 NXT:1980i, 4 / 5 / 6 NXT:2050i, and 6 NXT:2100i. Accounting for joint uncertainty, 27% / 38% / 53% of the Chinese-owned NXT:1980i fleet is located at advanced logic or DRAM fabs.
- **Allocation of advanced-fab scanners to AI.** We estimate that 5% / 10% / 20% of advanced lithography capacity is devoted to AI chips. For logic, Epoch estimates that Huawei will produce around 1 million Ascend 950 packages in 2026, each containing two compute dies of approximately 400 mm^2.[258] A 300mm wafer holds roughly 140 dies of that size. Assuming 30–50% die yield, this would require 2,500–4,200 wafers per month, roughly 4–7% of SMIC's advanced-node capacity, estimated by SemiAnalysis at 60,000 wafer starts per month.[259] Epoch also projects 430,000 Ascend 910Cs, but some may use foreign dies fabricated at TSMC. Accounting for production by other Chinese AI chip designers, potential domestic 910C production, and uncertainty in output and yields supports a wider 5–20% range. For memory, Epoch estimates CXMT's 2026 HBM allocation at 30,000 wafer starts per month by the end of 2026.[260] This comes out of roughly 350,000 total DRAM wafer starts per month. Assuming advanced nodes constitute between one-half and all of CXMT's capacity, the implied advanced-node share for HBM is 9–17%.

For later years, the allocation varies by scenario:

- **Baseline** assumes continuity with the present allocation, as influenced by commercial and policy factors. As discussed in Section 5, the Dutch government requires a license for exports of NXT:1970i and NXT:1980i within China. There is substantial uncertainty regarding which licenses are currently being granted and, more broadly, how effectively the broader control structure can prevent diversion, diversion in place, and other tactics that China might use to obtain NXT:1980i at fabs making AI chips. As discussed in Appendix E, there is evidence that NXT:1980i systems have reached advanced fabs, such as CXMT and Hua Hong. Rather than making assumptions about Dutch licensing policy or the effectiveness of enforcement, Baseline assumes that approximately the present mixture of policy impact and commercial incentive persists into the future: existing tools remain at their estimated fab types (Appendix E), while imported and indigenous systems enter advanced fabs in the same proportion (27% / 38% / 53%) as the current fleet. Within those advanced fabs, AI allocation rises smoothly from 5% / 10% / 20% to the outside-China allocation (see US production model, above) by 2030, then tracks it through 2035.
- **Mobilization** assumes that Chinese fabs strongly prioritize AI production and find ways to circumvent any future policies that suppress the current allocation at advanced fabs—whether

through overt violation, illicit diversion, misleading end-use claims, diversion in place, or other sophisticated evasion. We use the single allocation factor estimated in the US production model above, tracking the share of the fleet devoted to AI production.

- **Physical ceiling** allocates the entire existing stockpile and every imported or produced system to advanced fabs.

In the "Mobilization" and "Physical Ceiling" scenarios, we allow DUVi systems to be divided between logic and DRAM to maximize Ascend 910C output, reflecting strong prioritization of AI output. The "Baseline" scenario assumes the existing logic/DRAM split of currently installed DUVi systems.

Caveats

- The China estimates are AI chip production upper bounds based on aggregate DUVi system capacity, not forecasts of actual AI chip production. Full utilization of China's DUVi installed base could be bottlenecked by lack of etch, deposition, metrology, advanced packaging, and HBM fabrication and stacking tools, as well as masks, servicing, skilled labor, power, and demand.
- The model assumes no scanner retirements, failures, re-exports, or servicing losses. This may overstate usable capacity later in the forecast and means that a ban on servicing would probably constrain China more than the modeled import ban alone.
- The model assumes that ASML continues producing China-eligible NXT:1980i systems rather than retiring it in favor of the most advanced DUVi systems (i.e. the NXT:2000i and more advanced), which cannot be sold to China.
- The model assumes that Chinese fabs reach mature TSMC and SK hynix AI chip yield and throughput by 2031. This requires China to rapidly improve process control, HBM stacking, and advanced packaging.
- The "Mobilization" and "Physical Ceiling" scenarios assume widespread diversion in violation of US export controls and an ideal balancing of DUVi systems across logic and HBM fabs. They do not model fab construction time, process qualification time, downtime, or the opportunity cost of reducing legacy and non-AI chip production, all of which could reduce Chinese AI chip output.
- The DUVi system indigenization and scale-up estimates are highly speculative, relying on uncertain estimates of when China achieves commercially viable DUVi systems and how quickly Chinese DUVi providers can scale commercial production.
- The model's US production forecast depends on an average of a demand model and a model of EUV supply for AI chipmaking. It does not separately constrain HBM, advanced packaging, substrates, power, data-center construction, DUVi machines for certain layers, or High-NA EUV.

Authors

Nicholas Brown is a Senior Fellow at the Center for Technology & Statecraft. Saif M. Khan is Director of the Center for Technology & Statecraft.

Acknowledgements

For helpful discussions, comments, and input, we thank Cynthia Light Brown, Jack Freed, Joseph Fridman, Jimmy Goodrich, Lennart Heim, Jeremy Lansford, Hamish Low, Mary Clare McMahon, Amelia Michael, Konstantin Pilz, and James Sanders. The authors are solely responsible for all mistakes.

Revision history

We first published this report in September 2026. In November 2026, we improved the AI chip production model in Appendix F and redrew Figures 1 and 19. The report's conclusions are unchanged.

- **H100-equivalents** now use dense 8-bit performance instead of dense 16-bit performance.
- **US projections (Tables 41 and 43)** use TrendForce's October 2026 forecast for 2026–2027 and average the demand- and EUV-based estimates from 2028 instead of 2029. US production in 2035 is 578 / 1,252 / 5,404 million H100e (previously 422 / 974 / 3,566).
- **The EUV-based estimate (Table 44)** starts from a chip-by-chip estimate of AI chips' 2027 EUV use and assumes 85 / 95 / 110 EUV scanner shipments in 2028 (previously 110), growing 5% a year at the median (previously 10%).
- **China projections (Tables 45 and 49)** count the Ascend 910C at 0.76 H100e (previously 0.64). With a DUVi ban, China's 2035 maximum is 104 / 182 / 382 million H100e (previously 87 / 153 / 321).

Endnotes

[1] Emerging Technology Observatory, "ArF Dry (DUV) Lithography Tools," Supply Chain Explorer: Advanced Chips, accessed September 26, 2026, https://chipexplorer.eto.tech/?parentNode=N25&selectedNode=N21.

[2] Emerging Technology Observatory, "KrF (DUV) Lithography Tools," Supply Chain Explorer: Advanced Chips, accessed September 26, 2026, https://chipexplorer.eto.tech/?parentNode=N25&selectedNode=N105.

[3] Georgia Adamson et al., "Should the US Sell Blackwell Chips to China?" Institute for Progress, October 25, 2025, https://ifp.org/the-b30a-decision/.

[4] Fanny Potkin, "China Starts Production of Home-Grown Immersion DUV Chipmaking Tools—Source," *Reuters*, July 28, 2026, https://www.reuters.com/world/china/china-starts-production-home-grown-immersion-duv-chipmaking-tools-source-2026-07-28/.

[5] "China Trials Its First Advanced Tools for AI Chipmaking," *Financial Times*, September 2025, https://www.ft.com/content/8fd79522-e34f-4633-bc87-ef0aae2d9159.

[6] Export Administration Regulations, 15 C.F.R. § 734.9(k) (2026), https://www.ecfr.gov/current/title-15/part-734/section-734.9#p-734.9(k).

[7] "Light Source Technology," Cymer, accessed September 26, 2026, https://www.cymer.com/light-source-technology/; Gigaphoton Inc., home page, accessed September 26, 2026, https://www.gigaphoton.com.

[8] Export Administration Regulations, 15 C.F.R. § 736.2(b)(10) (2026), https://www.ecfr.gov/current/title-15/part-736/section-736.2#p-736.2(b)(10).

[9] Bureau of Industry and Security, "Guidance on Application of General Prohibition 10 (GP10) to People's Republic of China (PRC) Advanced-Computing Integrated Circuits (ICs)," US Department of Commerce, May 13, 2025, https://www.bis.gov/media/documents/general-prohibition-10-guidance-may-13-2025.pdf.

[10] Diederik Baazil, Cagan Koc, and Jordan Robertson, "ASML and TSMC Can Disable Chip Machines If China Invades Taiwan," *Bloomberg*, May 21, 2024, https://www.bloomberg.com/news/articles/2024-05-21/asml-tsmc-can-disable-chip-machines-if-china-invades-taiwan.

[11] Toby Sterling, "ASML CEO: Would Not Change 2030 Growth Targets If China Excluded," *Reuters*, November 11, 2022, https://www.reuters.com/technology/asml-ceo-would-not-change-2030-growth-targets-if-china-excluded-2022-11-11/.

[12] "Light and Lasers," ASML, accessed September 26, 2026, https://www.asml.com/en/technology/lithography-principles/light-and-lasers.

[13] K. Jain, C. G. Willson, and B. J. Lin, "Ultrafast Deep UV Lithography with Excimer Lasers," *IEEE Electron Device Letters* 3, no. 3 (1982): 53–55, https://research.ibm.com/publications/ultrafast-deep-uv-lithography-with-excimer-lasers.

[14] M. Rothschild et al., "Photolithography at 193 nm," *Journal of Vacuum Science & Technology B* 10, no. 6 (1992): 2989–96, https://www.researchgate.net/publication/224462378_Photolithography_at_193_nm.

[15] Yayi Wei and David Back, "193nm Immersion Lithography: Status and Challenges," SPIE, March 22, 2007, https://www.spie.org/news/immersionlitho-intro; Sander Hofman, "How Immersion Lithography Saved Moore's Law," ASML, August 2, 2023, https://www.asml.com/en/news/stories/2023/how-immersion-lithography-saved-moores-law.

[16] Patrick Naulleau, "EUV Lithography," *Synchrotron Radiation News* 32, no. 4 (2019): 2, https://www.tandfonline.com/doi/full/10.1080/08940886.2019.1634429.

[17] Brian Potter, "How ASML Got EUV," *Construction Physics* (blog), November 20, 2025, https://www.construction-physics.com/p/how-asml-got-euv; Emerging Technology Observatory, "ArF Immersion (DUV) Lithography Tools," Supply Chain Explorer: Advanced Chips, accessed September 26, 2026, https://chipexplorer.eto.tech/?parentNode=N25&selectedNode=N22.

[18] ASML, "High NA EUV Reaches New Readiness Milestone with First High-Volume Logic Product," press release, July 15, 2026, https://www.asml.com/en/news/press-releases/2026/high-na-euv-reaches-new-readiness-milestone.

[19] Danilo De Simone et al., "High-NA EUV Lithography: The Next Major Step Forward," imec, October 4, 2021, https://www.imec-int.com/en/articles/high-na-euvl-next-major-step-lithography.

[20] "65 nm Process," Wikipedia, last modified May 2, 2026, https://en.wikipedia.org/wiki/65_nm_process; Anton Shilov, "Intel's 18A and TSMC's N2 Process Nodes Compared: Intel Is Faster, but TSMC Is Denser," *Tom's Hardware*, February 13, 2025, https://www.tomshardware.com/tech-industry/intels-18a-and-tsmcs-n2-process-nodes-compared-intel-is-faster-but-tsmc-is-denser.

[21] Dylan Patel et al., "The Memory Wall: Past, Present, and Future of DRAM," *SemiAnalysis* (blog), September 3, 2024, https://newsletter.semianalysis.com/p/the-memory-wall.

[22] Sony Varghese, "DRAM Scaling Requires New Materials Engineering Solutions," Applied Materials (blog), April 27, 2021, https://www.appliedmaterials.com/us/en/blog/blog-posts/dram-scaling-requires-new-materials-engineering-solutions.html.

[23] *NVIDIA Tesla V100 GPU Architecture: The World's Most Advanced Data Center GPU*, white paper WP-08608-001_v1.1 (NVIDIA, 2017), https://images.nvidia.com/content/volta-architecture/pdf/volta-architecture-whitepaper.pdf.

[24] *NVIDIA Tesla V100 GPU Architecture*.

[25] Nasir Rabik, "Samsung Electronics Start Mass Production of HBM2 DRAM," *Lowyat.NET*, January 20, 2016, https://www.lowyat.net/2016/92625/samsung-electronics-start-mass-production-of-hbm2-dram/.

[26] Ronny Krashinsky et al., "NVIDIA Ampere Architecture In-Depth," *NVIDIA Technical Blog*, May 14, 2020, https://developer.nvidia.com/blog/nvidia-ampere-architecture-in-depth/.

[27] "NVIDIA A100 Tensor Core GPU," NVIDIA, accessed September 26, 2026, https://www.nvidia.com/en-us/data-center/a100/.

[28] Samsung Electronics, "Samsung to Advance High Performance Computing Systems with Launch of Industry's First 3rd-Generation (16GB) HBM2E," press release, February 4, 2020, https://news.samsung.com/global/samsung-to-advance-high-performance-computing-systems-with-launch-of-industrys-first-3rd-generation-16gb-hbm2e.

[29] Michael Andersch et al., "NVIDIA Hopper Architecture In-Depth," *NVIDIA Technical Blog*, March 22, 2022, https://developer.nvidia.com/blog/nvidia-hopper-architecture-in-depth/.

[30] SK hynix, "SK hynix to Supply Industry's First HBM3 DRAM to NVIDIA," press release, June 8, 2022, https://news.skhynix.com/sk-hynix-to-supply-industrys-first-hbm3-dram-to-nvidia/.

[31] Moore Morris and Ray Wang, "Deep Dive on HBM," *Nomad Semi* (blog), June 3, 2025, https://www.nomadsemi.com/p/deep-dive-on-hbm; WhyCry, "SK hynix HBM3E (HBM3 Extended) 5th Generation High-Bandwidth-Memory to Enter Mass Production in 2024," *VideoCardz*, May 30, 2023, https://videocardz.com/newz/sk-hynix-hbm3e-hbm3-extended-5th-generation-high-bandwidth-memory-to-enter-mass-production-in-2024.

[32] "NVIDIA Blackwell Architecture," NVIDIA, accessed September 26, 2026, https://www.nvidia.com/en-us/data-center/technologies/blackwell-architecture/.

[33] SK hynix, "SK hynix Begins Volume Production of Industry's First HBM3E," press release, March 19, 2024, https://news.skhynix.com/sk-hynix-begins-volume-production-of-industry-first-hbm3e/.

[34] WhyCry, "SK hynix HBM3E."

[35] Dick James, "A Trip down TSMC Memory Lane – Part 2," *TechInsights* (blog), accessed September 26, 2026, https://www.techinsights.com/blog/trip-down-tsmc-memory-lane-part-2.

[36] Intel Corporation, "Intel Completes Next Generation, 32nm Process Development Phase," press release, December 9, 2008, https://www.intc.com/news-events/press-releases/detail/1330/intel-completes-next-generation-32nm-process-development.

[37] "Xilinx Enters into Partnership with Samsung," *Tech Monitor*, February 21, 2010, https://www.techmonitor.ai/technology/xilinx_enters_into_partnership_with_samsung_100222.

[38] Mark LaPedus, "Getting Over Overlay," *Semiconductor Engineering*, June 18, 2015, https://semiengineering.com/getting-over-overlay/.

[39] Y. K. Park et al., "Fully Integrated 56 nm DRAM Technology for 1 Gb DRAM," in *2007 IEEE Symposium on VLSI Technology* (IEEE, 2007), 190–91, https://doi.org/10.1109/VLSIT.2007.4339688; "So… About Samsung Mass Producing the Most Advanced 20nm DDR3 DRAM," Samsung Global Newsroom, March 18, 2014, https://news.samsung.com/global/about-samsung-mass-producing-the-most-advanced-20nm-ddr3-dram.

[40] Tae-Seung Eom et al., "Comparative Study of Binary Intensity Mask and Attenuated Phase Shift Mask Using Hyper-NA Immersion Lithography for Sub-45nm Era," in *Optical Microlithography XXI*, Proceedings of SPIE 6924 (2008), 69240H, https://www.researchgate.net/publication/252282185_Comparative_study_of_binary_intensity_mask_and_attenuated_phase_shift_mask_using_hyper-NA_immersion_lithography_for_sub-45nm_era_-_art_no_69240H; "10,000 Times Thinner than Human Hair? How SK hynix's Miniaturization Innovation Secured Its DRAM Leadership," SK hynix Newsroom, June 27, 2025, https://news.skhynix.com/one-team-spirit-miniaturization-innovation-dram-leadership/.

[41] Carl Wintgens, "The 50-nm DRAM Battle Rages On: An Overview of Micron's Technology," *EE Times*, March 23, 2009, https://www.eetimes.com/the-50-nm-dram-battle-rages-on-an-overview-of-microns-technology/; Micron Technology, Inc., *Annual Report on Form 10-K for the Fiscal Year Ended September 3, 2015* (Boise, ID: Micron Technology, 2015), https://www.sec.gov/Archives/edgar/data/0000723125/000072312515000112/a2015q4.htm.

[42] Charlotte Trueman, "Meta Report Details Hundreds of GPU and HBM3 Related Interruptions to Llama 3 Training Run," *DatacenterDynamics*, July 29, 2024, https://www.datacenterdynamics.com/en/news/meta-report-details-hundreds-of-gpu-and-hbm3-related-interruptions-to-llama-3-training-run/; Bugra Akyildiz, "Scaling and Reliability Challenges of LLama3," *MLOps Newsletter* (blog), September 8, 2024, https://mlops.substack.com/p/scaling-and-reliability-challenges.

[43] Omkar Salpekar et al., "Training LLMs with Fault Tolerant HSDP on 100,000 GPUs," preprint, arXiv, January 2026, https://arxiv.org/pdf/2602.00277.

[44] Apostolos Kokolis et al., "Revisiting Reliability in Large-Scale Machine Learning Research Clusters," preprint, arXiv, October 2024, https://arxiv.org/html/2410.21680v2.

[45] *AI Changes Everything: Business Agility, Real Results*, e-book, Latin America 2025 (Oracle, 2025), 10, https://www.oracle.com/latam/a/ocom/docs/gated/artificial-intelligence-ai-changes-ebook-en.pdf; Crusoe, "Crusoe Expands AI Data Center Campus in Abilene to 1.2 Gigawatts," press release, March 18, 2025, https://crusoe.ai/newsroom/crusoe-expands-ai-data-center-campus-in-abilene-to-1-2-gigawatts.

[46] Jessica Timings, "TWINSCAN: 20 Years of Lithography Innovation," ASML, August 18, 2021, https://www.asml.com/en/news/stories/2021/twinscan-20-years-innovation; Peter Westberg, "ASML: Architecting Earth's Most Complex Machines," Quartr, September 27, 2024, https://quartr.com/insights/edge/asml-architecting-earths-most-complex-machines.

[47] Neil Hacker, "The World's Most Complex Machine," *Works in Progress*, April 23, 2026, https://worksinprogress.co/issue/the-worlds-most-complex-machine/.

[48] Jon Y, "A Deep Dive into Immersion Lithography Technology," *The Asianometry Newsletter* (blog), October 26, 2022, https://www.asianometry.com/p/a-deep-dive-into-immersion-lithography.

[49] Hacker, "World's Most Complex Machine."

[50] Emerging Technology Observatory, "ArF Immersion (DUV) Lithography Tools"; "Nikon Aims to Close the Gap on ASML with New ArF Lithography System in FY28," *TrendForce*, February 21, 2025, https://www.trendforce.com/news/2025/02/21/news-nikon-aims-to-close-the-gap-on-asml-with-new-arf-lithography-system-in-fy28/.

[51] ASML Holding N.V., "ASML 2025 Fourth-Quarter and Full-Year Results" (investor presentation, Veldhoven, Netherlands, January 28, 2026), exhibit 99.2 to Form 6-K, January 28, 2026, https://www.sec.gov/Archives/edgar/data/937966/000162828026003701/a2026_01x28xpresentation.htm.

[52] "ASML Holding Research and Development Expenses," Macrotrends, accessed September 26, 2026, https://www.macrotrends.net/stocks/charts/ASML/asml-holding/research-development-expenses; "Applied Materials Research and Development Expenses," Macrotrends, accessed September 26, 2026, https://www.macrotrends.net/stocks/charts/AMAT/applied-materials/research-development-expenses; "Lam Research Research and Development Expenses," Macrotrends, accessed September 26, 2026, https://www.macrotrends.net/stocks/charts/LRCX/lam-research/research-development-expenses.

[53] Mark LaPedus, "Billions and Billions Invested," *Semiconductor Engineering*, April 17, 2014, https://semiengineering.com/billions-and-billions-invested/.

[54] Katie Tarasov, "Exclusive Look at the Creation of High NA, ASML's New $400 Million Chipmaking Colossus," *CNBC*, May 22, 2025, https://www.cnbc.com/2025/05/22/exclusive-look-at-high-na-asmls-new-400-million-chipmaking-colossus.html.

[55] "Semiconductor Manufacturing Equipment Price: Understanding the Cost Factors," Hillman Curtis, October 13, 2023, https://hillmancurtis.com/semiconductor-manufacturing-equipment-price/.

[56] Ni Zhengyang 倪正洋, *半导体行业深度报告：复盘 ASML 发展历程，探寻本土光刻产业链投资机会* [In-Depth Semiconductor Industry Report: Reviewing ASML's History and Exploring Investment Opportunities in China's Domestic Lithography Supply Chain] (西南证券 [Southwest Securities], June 17, 2020), https://pdf.dfcfw.com/pdf/H3_AP202006181385713139_1.pdf?1592474684000.pdf.

[57] Intel Corporation, *Annual Report 1997* (Santa Clara, CA: Intel Corporation, 1998), https://www.chiphistory.org/chc_upload/content/pdf/19/1506018723/1506018723.pdf.

[58] "全球光刻机发展概况以及半导体装备国产化" [An Overview of Global Lithography Machine Development and the Localization of Semiconductor Equipment], *国际电子商情* [ESM China], May 23, 2018, https://www.esmchina.com/news/201805230941.html.

[59] TSMC, "TSMC Announces Nexsys 90 Nanometer Volume Production," press release, December 29, 2004, https://pr.tsmc.com/english/news/1319.

[60] Anton Shilov, "China's 28nm-Capable Chip Fabbing Tool on Track amid Trade War," *Tom's Hardware*, December 6, 2020, https://www.tomshardware.com/news/chinas-28nm-capable-chip-fabbing-tool-on-track-amid-trade-war; Vishakha Saxena, "China Expecting Key Native Lithography Machine This Year: SCMP," *Asia Financial*, August 4, 2023, https://www.asiafinancial.com/china-expecting-key-native-lithography-machine-this-year-scmp.

[61] "China Trials Its First Advanced Tools"; TSMC, "TSMC 28nm Technology in Volume Production," press release, October 24, 2011, https://pr.tsmc.com/english/news/1689; Potkin, "China Starts Production."

[62] Lipen Yuan, "Celebrating One Billion 7nm Chips: Why Scale Matters," *TSMC* (blog), August 20, 2020, https://www.tsmc.com/english/news-events/blog-article-20200801.

[63] ASML, "ASML Launches TWINSCAN NXT:1950i Immersion Lithography System," press release, December 3, 2008, https://www.asml.com/en/news/press-releases/2008/asml-launches-the-twinscan-nxt1950i-immersion-lithography-system; ASML, "ASML Ships New TWINSCAN NXT Immersion Lithography Platform," press release, September 29, 2015, https://www.asml.com/en/news/press-releases/2015/asml-ships-new-twinscan-nxt-immersion-lithography-platform.

[64] "China Trials Its First Advanced Tools."

[65] ASML, "ASML Announces Holistic Lithography Product Suite for the 7/5 Nanometer Node at SEMICON West," press release, July 10, 2017, https://www.asml.com/en/news/press-releases/2017/asml-announces-holistic-lithography-product-suite-for-the-75-nanometer-node-at-semicon-west; ASML Holding N.V., *Annual Report on Form 20-F for the Fiscal Year Ended December 31, 2019* (Veldhoven, Netherlands: ASML Holding N.V., 2020), https://www.sec.gov/Archives/edgar/data/937966/000093796620000011/a2019integratedreportb.htm.

[66] Stefan Weichselbaum et al., "Immersion and Dry Scanner Extensions for Sub-10nm Production Nodes," in *Optical Microlithography XXVIII*, Proceedings of SPIE 9426 (2015), 942616, https://www.spiedigitallibrary.org/conference-proceedings-of-spie/9426/942616/Immersion-and-dry-scanner-extensions-for-sub-10nm-production-nodes/10.1117/12.2087112.short; ASML, "ASML Reports First-Quarter Sales and Gross Margin in Line with Guidance," press release, April 20, 2016, https://www.asml.com/en/news/press-releases/2016/asml-reports-first-quarter-sales-and-gross-margin-in-line-with-guidance.

[67] ASML, "ASML Launches TWINSCAN NXT:1950i"; ASML, "ASML Announces 2010 Second Quarter Results," press release, July 14, 2010, https://www.asml.com/en/news/press-releases/2010/asml-announces-2010-second-quarter-results-sales-set-for-new-high.

[68] Hamish Low, Joshua Turner, and Romeo Dean, "A Forecast of Chinese DUV and EUV Photolithography Progress," *AI Futures Project* (blog), June 19, 2026, https://blog.aifutures.org/p/a-forecast-of-chinese-duv-and-euv.

[69] Fanny Potkin, "How China Built Its 'Manhattan Project' to Rival the West in AI Chips," *Reuters*, December 17, 2025, https://www.reuters.com/world/china/how-china-built-its-manhattan-project-rival-west-ai-chips-2025-12-17/.

[70] Qin Min, Qu Yunxu, and Han Wei, "'Big Fund' Graft Probe Ensnares Three More Individuals," *Caixin Global*, August 10, 2022, https://www.caixinglobal.com/2022-08-10/big-fund-graf-probe-ensnares-three-more-individuals-101924232.html; Qin Min and Han Wei, "China Piles $47.5 Billion into 'Big Fund III' to Boost Chip Development," *Caixin Global*, May 28, 2024, https://www.caixinglobal.com/2024-05-28/china-piles-475-billion-into-big-fund-iii-to-boost-chip-development-102200633.html; Ann Cao, "Former China Chip Fund Executive Expelled from Communist Party for Bribery, Graft," *South China Morning Post*, January 20, 2023,

https://www.scmp.com/tech/policy/article/3207561/former-china-chip-fund-executive-expelled-communist-party-bribery-graft.

[71] Stewart Randall, "Why Is China Investigating the State-Backed Semiconductor 'Big Fund'?" *TechNode*, August 12, 2022, https://technode.com/2022/08/12/silicon-why-is-china-investigating-the-state-backed-semiconductor-big-fund/; Qiu Xiaofen and Su Jianxun, "Billion Dollar Heist: How Scammers Rode China's Chip Boom to Riches," trans. Alexander Boyd, *ChinaTalk* (blog), March 30, 2021, https://www.chinatalk.media/p/billion-dollar-heist-how-scammers.

[72] Anton Shilov, "China Starts Big Fund III Spending: $47 Billion for Ecosystem and Fab Tools," *Tom's Hardware*, January 8, 2025, https://www.tomshardware.com/tech-industry/china-starts-big-fund-iii-spending-usd47-billion-for-ecosystem-and-fab-tools.

[73] Antonia Hmaidi, *Huawei Is Quietly Dominating China's Semiconductor Supply Chain*, MERICS Report (Mercator Institute for China Studies, 2024), https://merics.org/en/report/huawei-quietly-dominating-chinas-semiconductor-supply-chain.

[74] Low, Turner, and Dean, "Forecast of Chinese DUV and EUV."

[75] Will Hunt, Saif M. Khan, and Dahlia Peterson, *China's Progress in Semiconductor Manufacturing Equipment: Accelerants and Policy Implications*, CSET Policy Brief (Washington, DC: Center for Security and Emerging Technology, March 2021), https://cset.georgetown.edu/wp-content/uploads/CSET-Chinas-Progress-in-Semiconductor-Manufacturing-Equipment.pdf.

[76] Jacob Feldgoise and Hanna Dohmen, "Inside Beijing's Chipmaking Offensive: Where Is China Gaining Ground?" Center for Security and Emerging Technology, July 14, 2025, https://cset.georgetown.edu/article/inside-beijings-chipmaking-offensive/.

[77] Douglas B. Fuller, "Tech War or Phony War? America's Flimsy Controls on Semiconductor Fabrication Equipment and China's Response," preprint, SSRN, December 2023, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4704203; Jan-Peter Kleinhans and John Lee, "Lithography: Is the EU's Semiconductor Manufacturing Equipment a Strategic Chokepoint?" in *Reverse Dependency: Making Europe's Digital Technological Strengths Indispensable to China*, ed. Tim Rühlig (Digital Power China, 2024), 30–43, https://www.jpkleinhans.de/home/DPC2024_Lithography_Chapter.pdf.

[78] Feldgoise and Dohmen, "Inside Beijing's Chipmaking Offensive."

[79] Bloomberg, "Strong Demand in China for Chip Machines: Nikon," *Taipei Times*, July 18, 2024, https://www.taipeitimes.com/News/biz/archives/2024/07/18/2003820941.

[80] "China Gives Chipmakers New Powers to Guide Industry Recovery," *Financial Times*, March 2023, https://www.ft.com/content/d97ca301-f766-48c0-a542-e1d522c7724e.

[81] Douglas Fuller, "Tech War or Phony War? China's Response to America's Controls on Semiconductor Fabrication Equipment," *China Leadership Monitor*, no. 78 (December 2023), https://www.prcleader.org/post/tech-war-or-phony-war-america-s-porous-controls-on-semiconductor-equipment-and-china-s-response.

[82] Macro_Lin (@LinQingV), "长鑫存储扩产的设备供应链全景…" [Panorama of the equipment supply chain for CXMT's capacity expansion…], X, May 18, 2026, https://x.com/LinQingV/status/2056320187138453905.

[83] Paul Triolo, "Innovation under Pressure: China's Semiconductor Industry at a Crossroads," *American Affairs*, February 20, 2026, https://americanaffairsjournal.org/2026/02/innovation-under-pressure-chinas-semiconductor-industry-at-a-crossroads/.

[84] Rostislav Rokitski et al., "Enabling High Volume Manufacturing of Double Patterning Immersion Lithography with the XLR 600ix ArF Light Source," in *Optical Microlithography XXII*, Proceedings of SPIE 7274 (2009), 72743O, https://www.spiedigitallibrary.org/conference-proceedings-of-spie/7274/1/Enabling-high-volume-manufacturing-of-double-patterning-immersion-lithography-with/10.1117/12.816049.short.

[85] "Japan Rumored to Curb Photoresist Exports as China Targets 40% Self-Sufficiency by 2026," *TrendForce*, December 3, 2025, https://www.trendforce.com/news/2025/12/03/news-japan-rumored-to-curb-photoresist-exports-as-china-targets-40-self-sufficiency-by-2026/.

[86] Eduardo Jaramillo, "China's Semiconductor Industry Can't Quit German Optics," *The China Project*, May 1, 2023, https://thechinaproject.com/2023/05/01/chinas-semiconductor-industry-cant-quit-german-optics/.

[87] Kleinhans and Lee, "Lithography."

[88] Will Hunt, Saif M. Khan, and Dahlia Peterson, *China's Progress in Semiconductor Manufacturing Equipment: Accelerants and Policy Implications*, CSET Policy Brief (Washington, DC: Center for Security and Emerging Technology, March 2021), https://cset.georgetown.edu/publication/chinas-progress-in-semiconductor-manufacturing-equipment/.

[89] Bureau of Industry and Security, "Commerce Implements New Export Controls on Advanced Computing and Semiconductor Manufacturing Items to the People's Republic of China (PRC)," press release, October 7, 2022, https://www.bis.gov/press-release/commerce-implements-new-export-controls-advanced-computing-semiconductor-manufacturing-items-peoples.

[90] Bureau of Industry and Security, "Commerce Strengthens Export Controls to Restrict China's Capability to Produce Advanced Semiconductors for Military Applications," press release, December 2, 2024, https://www.bis.gov/press-release/commerce-strengthens-export-controls-restrict-chinas-capability-produce-advanced-semiconductors-military.

[91] House Select Committee on the Strategic Competition between the United States and the Chinese Communist Party, *Selling the Forges of the Future: How the CCP Gets Its Key Equipment for Making Semiconductors from U.S., Dutch, and Japanese Companies* (Washington, DC: US House of Representatives, October 7, 2025), https://chinaselectcommittee.house.gov/sites/evo-subsites/selectcommitteeontheccp.house.gov/files/evo-media-document/selling-the-forges-of-the-future.pdf.

[92] 15 C.F.R. § 734.9(k).

[93] Gregory C. Allen, Emily Benson, and Margot Putnam, "Japan and the Netherlands Announce Plans for New Export Controls on Semiconductor Equipment," Center for Strategic and International Studies, April 10, 2023, https://www.csis.org/analysis/japan-and-netherlands-announce-plans-new-export-controls-semiconductor-equipment.

[94] Export Administration Regulations, 15 C.F.R. § 744.23 (2026), https://www.ecfr.gov/current/title-15/subtitle-B/chapter-VII/subchapter-C/part-744/section-744.23.

[95] Export Administration Regulations, 15 C.F.R. § 734.9(e)(3) (2026), https://www.ecfr.gov/current/title-15/part-734/section-734.9#p-734.9(e)(3).

[96] House Select Committee on the CCP, *Selling the Forges of the Future*.

[97] Export Administration Regulations, Entity List, 15 C.F.R. pt. 744, supp. no. 4 (2026), https://www.ecfr.gov/current/title-15/subtitle-B/chapter-VII/subchapter-C/part-744/appendix-Supplement%20No.%204%20to%20Part%20744; Export Administration Regulations, 15 C.F.R. § 734.9(e)(3)(ii) (2026), https://www.ecfr.gov/current/title-15/part-734/section-734.9#p-734.9(e)(3)(ii).

[98] Netherlands Ministry of Foreign Affairs, "Regeling van de Staatssecretaris van Buitenlandse Zaken van 20 november 2025, nr. BZ2520952, houdende intrekking van de Regeling aanvullende controlemaatregelen op de Verordening producten voor tweeërlei gebruik en de wijziging van de Regeling geavanceerde productieapparatuur voor halfgeleiders" [Regulation of the State Secretary for Foreign Affairs of November 20, 2025, No. BZ2520952, Repealing the Regulation on Additional Control Measures to the Dual-Use Items Regulation and Amending the Regulation on Advanced Semiconductor Manufacturing Equipment], *Staatscourant* 2025, no. 38653, November 24, 2025, https://zoek.officielebekendmakingen.nl/stcrt-2025-38653.html.

[99] Netherlands Ministry of Foreign Affairs, "Regeling van 20 november 2025"; Export Administration Regulations, Commerce Control List, 15 C.F.R. pt. 774, supp. no. 1 (2026), https://www.ecfr.gov/current/title-15/subtitle-B/chapter-VII/subchapter-C/part-774.

[100] Export Administration Regulations, 15 C.F.R. § 744.6 (2026), https://www.ecfr.gov/current/title-15/subtitle-B/chapter-VII/subchapter-C/part-744/section-744.6.

[101] Netherlands Ministry of Foreign Affairs, "Regeling van 20 november 2025."

[102] Regulation (EU) 2021/821 of the European Parliament and of the Council of 20 May 2021 Setting Up a Union Regime for the Control of Exports, Brokering, Technical Assistance, Transit and Transfer of Dual-Use Items (Recast), *Official Journal of the European Union* L 206 (June 11, 2021): 1, consolidated text as of November 15, 2025, https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?qid=1782099200814&uri=CELEX%3A02021R0821-20251115; Netherlands Ministry of Foreign Affairs, "Regeling van 20 november 2025."

[103] "技術関連" [Technology-Related Questions], 安全保障貿易管理 Q&A [Security Export Control Q&A], Ministry of Economy, Trade and Industry, last updated April 23, 2026, https://www.meti.go.jp/policy/anpo/qanda25.html.

[104] "輸出貿易管理令別表第一及び外国為替令別表の規定に基づき貨物又は技術を定める省令" [Ministerial Order Specifying Goods and Technologies Pursuant to the Provisions of Appended Table 1 of the Cabinet Order on Export Trade Control and the Appended Table of the Foreign Exchange Order], MITI Ordinance No. 49 of 1991, Japanese–English parallel text, translation current through METI Ordinance No. 5 of 2020, Japanese Law Translation Database System, Ministry of Justice, accessed September 26, 2026, https://www.japaneselawtranslation.go.jp/ja/laws/view/4401/tb.

[105] Fanny Potkin, "China's No. 2 Chipmaker Readies 7 nm Production as Beijing Ramps Up Self-Sufficiency Drive," *Reuters*, March 16, 2026, https://www.reuters.com/world/asia-pacific/chinas-no-2-chipmaker-readies-7-nm-production-beijing-ramps-up-self-suffiency-2026-03-16/.

[106] Ray Wang et al., "China's CXMT Is Set to Challenge DRAM Incumbents," *SemiAnalysis* (blog), June 23, 2026, https://newsletter.semianalysis.com/p/chinas-cxmt-is-set-to-challenge-dram; Venkat Somala, "Will Huawei Catch Up to Nvidia by 2030?" Epoch AI, September 4, 2026, updated September 24, 2026, https://epoch.ai/publications/huaweis-roadmap-to-2031.

[107] Adamson et al., "Should the US Sell Blackwell Chips."

[108] Roelof de Graaf et al., "NXT:1980Di Immersion Scanner for 7nm and 5nm Production Nodes," in *Optical Microlithography XXIX*, Proceedings of SPIE 9780 (2016), 978011, https://www.spiedigitallibrary.org/conference-proceedings-of-spie/9780/978011/NXT1980Di-immersion-scanner-for-7nm-and-5nm-production-nodes/10.1117/12.2220589.short.

[109] Anton Shilov, "Update: ASML Denies Report It Was Working on China-Specific Lithography Tools," *Tom's Hardware*, July 5, 2023, https://www.tomshardware.com/news/asml-may-release-china-specific-versions-of-lithography-tools-report.

[110] Mark LaPedus, "DRAM Scaling Challenges Grow," *Semiconductor Engineering*, November 21, 2019, https://semiengineering.com/dram-scaling-challenges-grow/.

[111] De Graaf et al., "NXT:1980Di Immersion Scanner."

[112] Martin van den Brink, "Industry Roadmap and Technology Strategy" (presentation, ASML Investor Day, Veldhoven, Netherlands, November 8, 2018), exhibit 99.3 to ASML Holding N.V., Form 6-K, November 9, 2018, https://www.sec.gov/Archives/edgar/data/937966/000119312518323057/d650002dex993.htm.

[113] Samsung Electronics, "Samsung Now Mass Producing Industry's First 2nd-Generation, 10-Nanometer Class DRAM," press release, December 20, 2017, https://news.samsung.com/global/samsung-now-mass-producing-industrys-first-2nd-generation-10-nanometer-class-dram; ASML, "Q3 Sales Strong across the Full Product Portfolio," press release, October 18, 2017, https://www.asml.com/en/news/press-releases/2017/q3-sales-strong-across-the-full-product-portfolio; ASML, "Earnings Growth Continues, Driven by Strong Sales across Full Product Portfolio," press release, July 18, 2018, https://www.asml.com/en/news/press-releases/2018/earnings-growth-continues-driven-by-strong-sales-across-full-product-portfolio.

[114] Rui Ma (@ruima), "The most important China Tech news today is Huawei's chip roadmap announcement and associated breakthroughs…," X, September 18, 2025, https://x.com/ruima/status/1968592824380878931.

[115] Dylan Patel and Daniel Nishball, "Groq Inference Tokenomics: Speed, but at What Cost?" *SemiAnalysis* (blog), February 21, 2024, https://newsletter.semianalysis.com/p/groq-inference-tokenomics-speed-but.

[116] "TWINSCAN NXT:1965Ci," ASML, accessed September 26, 2026, https://www.asml.com/en/products/duv-lithography-systems/twinscan-nxt-1965ci.

[117] TSMC, *TSMC 2013 Business Overview* (Hsinchu, Taiwan: TSMC, 2014), https://investor.tsmc.com/sites/ir/annual-report/2013/2013_Business_Overview_E_1.pdf; ASML, "ASML Reports 2013 Results as Guided," press release, January 22, 2014, https://www.asml.com/en/news/press-releases/2014/asml-reports-2013-results-as-guided.

[118] ASML, *Annual Report 2025*, strategic report section (Veldhoven, Netherlands: ASML, 2026), https://ourbrand.asml.com/m/8ab959d4926657b/original/asml-2025-annual-report-strategic-report-section.pdf.

[119] Potkin, "How China Built Its 'Manhattan Project'."

[120] "ASML Poised for Revenue Surge as Zeiss Expands Capacity for Chip Tools—Kuo," *Investing.com*, January 27, 2026, https://www.investing.com/news/stock-market-news/asml-poised-for-revenue-surge-as-zeiss-expands-capacity-for-chip-tools--kuo-93CH-4468159.

[121] Netherlands Ministry of Foreign Affairs, "Regeling van 20 november 2025."

[122] Regulation (EU) 2021/821.

[123] Regulation (EU) 2021/821; Netherlands Ministry of Foreign Affairs, "Regeling van 20 november 2025."

[124] Potkin, "China's No. 2 Chipmaker Readies 7 nm Production."

[125] Karen Freifeld, "US Orders Chip Equipment Companies to Halt Some Shipments to Hua Hong, China's Second-Largest Chipmaker," *Reuters*, April 28, 2026, https://www.reuters.com/world/china/us-orders-chip-equipment-companies-halt-some-shipments-hua-hong-chinas-second-2026-04-28/.

[126] Keji Meixue 科技美学, "中芯国际 12nm 工艺已开始试产：功耗降低 20％、性能提升 10％" [SMIC's 12nm Process Has Entered Trial Production: 20% Lower Power, 10% Higher Performance], *新浪财经* [Sina Finance], May 18, 2020, https://finance.sina.cn/stock/relnews/hk/2020-05-18/detail-iirczymk2316159.d.html?vt=4.

[127] CXMT Corporation, *首次公开发行股票并在科创板上市招股说明书（申报稿）* [Prospectus for Initial Public Offering and Listing on the STAR Market (Draft for Filing)] (Hefei: CXMT Corporation, 2025), https://pdf.dfcfw.com/pdf/H2_AN202512301811513791_1.pdf?t=1767447526446.

[128] "Huawei's Shenzhen Facilities Reveal New Push into Advanced Chipmaking," *Financial Times*, May 4, 2025, https://www.ft.com/content/afd618f8-12c9-4297-b2a9-49f7dc548da4.

[129] Gregory C. Allen, *DeepSeek, Huawei, Export Controls, and the Future of the U.S.-China AI Race* (Washington, DC: Center for Strategic and International Studies, 2025), https://csis-website-prod.s3.amazonaws.com/s3fs-public/2025-03/250307_Allen_AI_Race.pdf?VersionId=8Qpp5OWwGj9OKB25qzQX4tQ_DzlZ9NS_.

[130] Cymer, "Light Source Technology."

[131] Chieh Hsu et al., "Wafer Transportation," US Patent Application Publication 2023/0395415 A1, filed June 4, 2022, and published December 7, 2023, https://patents.google.com/patent/US20230395415A1/en; Chartered Semiconductor Manufacturing Ltd., *Prospectus Supplement*, Form 424B5 (Singapore: Chartered Semiconductor Manufacturing, 2005), https://www.sec.gov/Archives/edgar/data/1095270/000114554905001413/u92561b5e424b5.htm; Muh-Cherng Wu, Chen-Fu Chen, and Chang-Fu Shih, "Route Planning for Two Wafer Fabs with Capacity-Sharing Mechanisms," *International Journal of Production Research* 47, no. 20 (2009): 5843–5856, https://scholar.nycu.edu.tw/en/publications/route-planning-for-two-wafer-fabs-with-capacity-sharing-mechanism/.

[132] "Huawei's Shenzhen Facilities."

[133] Netherlands Ministry of Foreign Affairs, "Regeling van 20 november 2025."

[134] Eleanor Olcott, "China Boosts AI Chip Output by Upgrading Older ASML Machines," *Financial Times*, December 19, 2025, https://www.ft.com/content/d10398db-b8b4-40f3-8c6d-b340470f5f3c?syn-25a6b1a6=1.

[135] Potkin, "How China Built Its 'Manhattan Project'."

[136] Potkin, "How China Built Its 'Manhattan Project'."

[137] K. Mistry et al., "A 45nm Logic Technology with High-k+Metal Gate Transistors, Strained Silicon, 9 Cu Interconnect Layers, 193nm Dry Patterning, and 100% Pb-Free Packaging," in *2007 IEEE International Electron Devices Meeting* (2007), 247–250, https://ieeexplore.ieee.org/document/4418914.

[138] "SMIC," AnySilicon, accessed September 26, 2026, https://anysilicon.com/vendors/smic/.

[139] ASML, "ASML Extends ArF Leadership," press release, December 6, 2006, https://www.asml.com/en/news/press-releases/2006/asml-extends-arf-leadership.

[140] Huixiong Dai et al., "45nm and 32nm Half-Pitch Patterning with 193nm Dry Lithography and Double Patterning," in *Optical Microlithography XXI*, Proceedings of SPIE 6924 (2008), 692421, https://doi.org/10.1117/12.772260.

[141] Weicheng Shiu et al., "Advanced Self-Aligned Double Patterning Development for Sub-30-nm DRAM Manufacturing," in *Optical Microlithography XXII*, Proceedings of SPIE 7274 (2009), 72740E, https://doi.org/10.1117/12.813986.

[142] J. M. Park et al., "20nm DRAM: A New Beginning of Another Revolution," in *2015 IEEE International Electron Devices Meeting (IEDM)* (2015), 26.5.1–26.5.4, https://doi.org/10.1109/IEDM.2015.7409774.

[143] National Institute of Standards and Technology, "Preventing the Improper Use of CHIPS Act Funding," 88 Fed. Reg. 65600 (September 25, 2023), https://www.federalregister.gov/documents/2023/09/25/2023-20471/preventing-the-improper-use-of-chips-act-funding.

[144] 15 C.F.R. pt. 774, supp. no. 1.

[145] "輸出貿易管理令" [Export Trade Control Order], Cabinet Order No. 378 of 1949, as amended through Cabinet Order No. 194 of 2026, e-Gov 法令検索 [e-Gov Law Search], accessed September 26, 2026, https://laws.e-gov.go.jp/law/324CO0000000378/; "輸出貿易管理令別表第一及び外国為替令別表の規定に基づき貨物又は技術を定める省令" [Ministerial Order Specifying Goods and Technologies Pursuant to the Provisions of Appended Table 1 of the Cabinet Order on Export Trade Control and the Appended Table of the Foreign Exchange Order], MITI Ordinance No. 49 of 1991, as amended through METI Ordinance No. 72 of 2025, e-Gov 法令検索 [e-Gov Law Search], accessed September 26, 2026, https://laws.e-gov.go.jp/law/403M50000400049/.

[146] Netherlands Ministry of Foreign Affairs, "Regeling geavanceerde productieapparatuur voor halfgeleiders" [Regulation on Advanced Semiconductor Manufacturing Equipment], *Staatscourant* 2023, no. 18212, June 30, 2023, https://zoek.officielebekendmakingen.nl/stcrt-2023-18212.html; Netherlands Ministry of Foreign Affairs, "Regeling van 20 november 2025."

[147] Regulation (EU) 2021/821 of the European Parliament and of the Council of 20 May 2021 Setting Up a Union Regime for the Control of Exports, Brokering, Technical Assistance, Transit and Transfer of Dual-Use Items (Recast), *Official Journal of the European Union* L 206 (June 11, 2021): 1, consolidated text as of November 15, 2025, https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:02021R0821-20251115.

[148] Cymer, "Light Source Technology"; Gigaphoton Inc., home page.

[149] Bureau of Industry and Security, "Export Controls on Semiconductor Manufacturing Items," 88 Fed. Reg. 73424, 73446 (October 25, 2023), https://www.federalregister.gov/d/2023-23049/p-amd-2.

[150] Export Administration Regulations, 15 C.F.R. § 734.9(k)(1)(ii)(B) (2026), https://www.ecfr.gov/current/title-15/subtitle-B/chapter-VII/subchapter-C/part-734/section-734.9#p-734.9%28k%29%281%29%28ii%29%28B%29.

[151] Herman Boom, "DUV Products and Business Opportunity" (presentation, ASML Investor Day, Veldhoven, Netherlands, November 14, 2024), exhibit 99.5 to ASML Holding N.V., Form 6-K, November 18, 2024, https://www.sec.gov/Archives/edgar/data/937966/000093796624000026/exhibit995.htm.

[152] Bureau of Industry and Security, "Guidance on Application of General Prohibition 10."

[153] Bureau of Industry and Security, "Guidance on Application of General Prohibition 10."

[154] Baazil, Koc, and Robertson, "ASML and TSMC Can Disable Chip Machines."

[155] Multilateral Alignment of Technology Controls on Hardware (MATCH) Act, S. 4281, 119th Cong. (2026), https://www.congress.gov/bill/119th-congress/senate-bill/4281/titles; "Multilateral Alignment of Controls on Semiconductor Manufacturing Equipment," sec. 6040 of Senate Amendment 6691 (Mr. Wicker) to S. 4784, 119th Cong., *Congressional Record* 172, no. 114 (July 14, 2026): S3971, https://www.govinfo.gov/content/pkg/CREC-2026-07-14/pdf/CREC-2026-07-14-pt1-PgS3947.pdf#page=25.

[156] Arnab Datta et al., "How to Implement an Operation Warp Speed for Rare Earths," Institute for Progress, October 21, 2025, https://ifp.org/how-to-implement-an-operation-warp-speed-for-rare-earths/#managing-export-control-and-trade-policy-and-chinese-retaliation.

[157] Sterling, "Would Not Change 2030 Growth Targets."

[158] René Raaijmakers and Paul van Gerven, "Japan's Export Restrictions Threaten to Marginalize Nikon in Top Litho Segment," *Bits&Chips*, April 5, 2023, https://bits-chips.com/article/japans-export-restrictions-threaten-to-marginalize-nikon-in-top-litho-segment/.

[159] ASML, "Investor Call ASML CEO Christophe Fouquet and CFO Roger Dassen: Q4/FY 2025 Results," transcript, January 28, 2026, https://ourbrand.asml.com/m/285dcba22c2bf203/original/2026_01_28-ASML-Transcript-investor-call-Q4-2025.pdf.

[160] Ryan Browne, "ASML Shares Jump as Surge in Orders Defies Fears of DeepSeek Hitting AI Chip Demand," *NBC10 Philadelphia*, January 29, 2025, https://www.nbcphiladelphia.com/news/business/money-report/critical-chip-firm-asml-posts-fourth-quarter-sales-and-profit-beat/4091773/.

[161] ASML, "ASML Reports €32.7 Billion Total Net Sales and €9.6 Billion Net Income in 2025," press release, January 28, 2026, https://www.asml.com/en/news/press-releases/2026/q4-2025-financial-results.

[162] ASML Holding N.V., *2025 Annual Report Based on IFRS* (Veldhoven, Netherlands: ASML Holding N.V., 2026), exhibit 99.1 to Form 6-K, February 25, 2026, https://www.sec.gov/Archives/edgar/data/937966/000162828026011377/asml-2025xannualxreportx.htm; ASML, "ASML 2025 Fourth-Quarter and Full-Year Results," investor relations presentation, January 28, 2026, https://ourbrand.asml.com/m/3136300aa4999bc1/original/2026_01_28_Presentation-Investor-Relations-Q4-

2025.pdf; "Foreign Exchange Rates—G.5A Annual," statistical release, Board of Governors of the Federal Reserve System, January 5, 2026, https://www.federalreserve.gov/releases/g5a/20260105/.

[163] Boyu Li et al., "A Reconfigurable Processing Element for Multiple-Precision Floating/Fixed-Point HPC," *IEEE Transactions on Circuits and Systems II: Express Briefs* 71, no. 3 (2024): 1401–5, https://www.scribd.com/document/895832459/A-Reconfigurable-Processing-Element-for-Multiple-Precision-Floating-Fixed-Point-HPC.

[164] TSMC, "TSMC Expands Advanced Technology Leadership with N4P Process," press release, October 26, 2021, https://pr.tsmc.com/english/news/2874#:~:text=Compared%20to%20N5%2C%20N4P%20will%20also%20deliver,investment%20in%20continuous%20improvement%20of%20our%20process.

[165] Daniel Nenni, "TSMC N5 Process Technology Wiki," SemiWiki, July 13, 2025, https://semiwiki.com/wikis/industry-wikis/tsmc-n5-process-technology-5nm-wiki/.

[166] "DCE Platform – N7/N6 Technology," TSMC, accessed September 26, 2026, https://www.tsmc.com/english/dedicatedFoundry/technology/platform_DCE_N7_N6.

[167] TSMC, "TSMC 16FinFET Plus Process Achieves Risk Production Milestone," press release, November 12, 2014, https://pr.tsmc.com/english/news/1835.

[168] "TSMC 20nm Technology," StatNano Nanotechnology Products Database, entry registered September 5, 2016, https://product.statnano.com/product/6776/tsmc-20nm-technology.

[169] Brian Jeff, "A Walk Through of the Microarchitectural Improvements in Cortex-A72," *Architectures and Processors Blog*, Arm Community, May 4, 2015, https://developer.arm.com/community/arm-community-blogs/b/architectures-and-processors-blog/posts/a-walk-through-of-the-microarchitectural-improvements-in-cortex-a72; Norman P. Jouppi et al., "TPU v4: An Optically Reconfigurable Supercomputer for Machine Learning with Hardware Support for Embeddings," preprint, arXiv, April 2023, https://arxiv.org/pdf/2304.01433.

[170] Kashish Kumawat, "7nm vs 10nm vs 14nm: Fabrication Process," Tech Centurion, November 26, 2019, https://www.techcenturion.com/7nm-10nm-14nm-fabrication.

[171] Illumi Huang, "Apple, Huawei Use TSMC, but Their 7nm SoCs Are Different," *EE Times*, January 22, 2020, https://www.eetimes.com/apple-huawei-use-tsmc-but-their-7nm-socs-are-different/.

[172] SkyJuice, "The Truth of TSMC 5nm," *Angstronomics* (blog), June 25, 2022, https://www.angstronomics.com/p/the-truth-of-tsmc-5nm.

[173] David Schor, "TSMC Extends Its 5nm Family with a New Enhanced-Performance N4P Node," *WikiChip Fuse*, October 26, 2021, archived February 15, 2026, at the Wayback Machine, https://web.archive.org/web/20260215150257/https://fuse.wikichip.org/news/6439/tsmc-extends-its-5nm-family-with-a-new-enhanced-performance-n4p-node/.

[174] Dylan Patel et al., "Scaling the Memory Wall: The Rise and Roadmap of HBM," *SemiAnalysis* (blog), August 12, 2025, https://newsletter.semianalysis.com/p/scaling-the-memory-wall-the-rise-and-roadmap-of-hbm.

[175] "HBM4 Thermal Path Design and Hot-Spot Mitigation in Advanced Stacks," Patsnap Eureka, September 12, 2025, https://eureka.patsnap.com/report-hbm4-thermal-path-design-and-hot-spot-mitigation-in-advanced-stacks.

[176] Hiroshige Goto 後藤弘茂, "8G-bit 品の投入で大容量化を図る HBM ロードマップ" [HBM Roadmap Aims for Larger Capacities with the Introduction of 8Gb Parts], *PC Watch*, May 15, 2014, https://pc.watch.impress.co.jp/docs/column/kaigai/648474.html.

[177] Aditya Ujeniya et al., "Architectural Trade-Offs in the Energy-Efficient Era: A Comparative Study of Power-Capping NVIDIA H100 and H200," preprint, arXiv, April 2026, https://arxiv.org/abs/2604.11391.

[178] Glenn Busch, "Powering the Future: ASML's Indispensable Role in Chip Manufacturing," American Money Management, March 8, 2024, https://www.amminvest.com/asml-indispensable-role-in-chip-manufacturing/.

[179] Youbao Zhang and Huijie Huang, "Photolithography Control System: A Case Study for Cyber-Physical System," preprint, arXiv, February 2024, https://arxiv.org/pdf/2402.15693; "Software," ASML, accessed September 26, 2026, https://www.asml.com/en/careers/working-at-asml/united-states/technical-field-software.

[180] S. Gaugiran et al., "Leaching Mechanisms in Immersion Lithography with or without Top Coat," *Microelectronic Engineering* 84, no. 5–8 (2007): 1054–57, https://www.sciencedirect.com/science/article/abs/pii/S0167931707000949.

[181] Wei and Back, "193nm Immersion Lithography."

[182] "Sourcing & Supply Chain," ASML, accessed September 26, 2026, https://www.asml.com/en/careers/teams/sourcing-and-supply-chain; ASML, *2024 Annual Report* (Veldhoven, Netherlands: ASML, 2025), https://www.asml.com/en/investors/annual-report/2024.

[183] "ASML's Supply Chain, Bill of Materials, and the Devastating Effects of Potential Tariffs on US Fabs," *Entropy Capital* (blog), April 23, 2025, https://entropycapital.substack.com/p/asmls-supply-chain-bill-of-materials; ASML, "Sourcing & Supply Chain."

[184] "Lenses & Mirrors," ASML, accessed September 26, 2026, https://www.asml.com/en/technology/lithography-principles/lenses-and-mirrors.

[185] "Cymer," ASML, accessed September 26, 2026, https://www.asml.com/en/company/about-asml/cymer.

[186] Westberg, "Architecting Earth's Most Complex Machines."

[187] Robert C. Leachman, *Competitive Semiconductor Manufacturing: Final Report on Findings from Benchmarking Eight-Inch, Sub-350nm Wafer Fabrication Lines*, CSM-52 (Berkeley: Competitive Semiconductor Manufacturing Program, University of California at Berkeley, 2002), https://ieor.berkeley.edu/wp-content/uploads/2019/10/CSM-52.pdf; Sarah Boyd, Arpad Horvath, and David Dornfeld, "Life-Cycle Assessment of NAND Flash Memory," *IEEE Transactions on Semiconductor Manufacturing* 24, no. 1 (2011): 117–124, https://escholarship.org/content/qt2wr9b3t1/qt2wr9b3t1.pdf?v=lg.

[188] Robert C. Leachman, "Yield Modeling and Analysis" (course notes, IEOR 130: Methods of Manufacturing Improvement, University of California, Berkeley, Spring 2014), https://fog.misty.com/perry/cod/references/yield_models.pdf.

[189] International Technology Roadmap for Semiconductors (ITRS), "Yield Enhancement," in *International Technology Roadmap for Semiconductors: 2005 Edition* (2005), https://www.semiconductors.org/wp-content/uploads/2018/08/2005Yield.pdf.

[190] NVIDIA, *NVIDIA H100 Tensor Core GPU Architecture*, white paper, version 1.02 (NVIDIA Corporation, 2022), https://www.hpctech.co.jp/assets/images/info/catalog/pdf/gtc22-whitepaper-hopper_v1.02.pdf; James S. Blomgren, "Microprocessor with a Large Cache Shared by Redundant CPUs for Increasing Manufacturing Yield," US Patent 5,828,578, filed November 29, 1995, and issued October 27, 1998, https://image-ppubs.uspto.gov/dirsearch-public/print/downloadPdf/5828578.

[191] Jan Willis and Aki Fujimura, *2019–2020 Mask Maker Survey Results* (eBeam Initiative, 2020), https://www.ebeam.org/docs/eBeam-Mask-Maker-Survey-2020.pdf.

[192] Mark LaPedus, "Litho Options Sparse after 10nm," *Semiconductor Engineering*, October 16, 2014, https://semiengineering.com/litho-options-after-10nm/.

[193] Mark LaPedus, "Survey: Optimism Grows for EUV," *Semiconductor Engineering*, September 11, 2017, https://semiengineering.com/survey-optimism-grows-for-euv/.

[194] Mark LaPedus, "Survey: Mask Complexity to Increase," *Semiconductor Engineering*, September 29, 2015, https://semiengineering.com/survey-mask-complexity-to-increase/.

[195] Mark LaPedus, "IBM, Intel and TSMC Roll Out FinFETs," *Semiconductor Engineering*, December 14, 2014, https://semiengineering.com/ibm-intel-and-tsmc-roll-out-finfets/.

[196] LaPedus, "IBM, Intel and TSMC Roll Out FinFETs."

[197] Mark LaPedus, "7nm Fab Challenges," *Semiconductor Engineering*, April 21, 2016, https://semiengineering.com/7nm-fab-challenges/.

[198] LaPedus, "Optimism Grows for EUV."

[199] Snowdog, "TSMC Actual 7nm Defect Rate and Therefore Yield Revealed," HardForum (forum thread), December 9, 2019, https://hardforum.com/threads/tsmc-actual-7nm-defect-rate-and-therefore-yield-revealed.1989911/.

[200] Leachman, "Yield Modeling and Analysis."

[201] Vikram Sekar, "How Foundries Calculate Die Yield," *Vik's Newsletter* (blog), September 15, 2024, https://www.viksnewsletter.com/p/how-foundries-calculate-die-yield; "How Many Chips Fit on a Wafer? GDPW Calculator & Die Yield Guide," Silicon Analysts, accessed September 26, 2026, https://siliconanalysts.com/guide/chips-per-wafer.

[202] Rani S. Ghaida, Mukul Gupta, and Puneet Gupta, "A Framework for Exploring the Interaction between Design Rules and Overlay Control," in *Metrology, Inspection, and Process Control for Microlithography XXVII*, Proceedings of SPIE 8681 (2013), 86810C, https://nanocad.ee.ucla.edu/wp-content/papercite-data/pdf/c69.pdf.

[203] Peter Jenkins, "Market Update" (presentation, ASML Investor Day, London, November 24, 2014), exhibit 99.7 to ASML Holding N.V., Form 6-K, November 24, 2014, https://www.sec.gov/Archives/edgar/data/937966/000119312514423554/d826396dex997.htm.

[204] Martin van den Brink, "Many Ways to Shrink: The Right Moves to 10 Nanometer and Beyond" (presentation, ASML Investor Day, London, November 24, 2014), exhibit 99.3 to ASML Holding N.V., Form 6-K, November 24, 2014, https://www.sec.gov/Archives/edgar/data/937966/000119312514423554/d826396dex993.htm; "TWINSCAN NXT:1470," ASML, accessed September 26, 2026, https://www.asml.com/products/duv-lithography-systems/twinscan-nxt1470.

[205] Emil Schmitt-Weaver et al., "Overlay Improvements Using a Real Time Machine Learning Algorithm," in *Metrology, Inspection, and Process Control for Microlithography XXVIII*, Proceedings of SPIE 9050 (2014), 90501S, https://www.spiedigitallibrary.org/conference-proceedings-of-spie/9050/90501S/Overlay-improvements-using-a-real-time-machine-learning-algorithm/10.1117/12.2046914.short.

[206] Van den Brink, "Many Ways to Shrink"; Jan Mulkens et al., "Overlay and Edge Placement Control Strategies for the 7nm Node Using EUV and ArF Lithography," in *Extreme Ultraviolet (EUV) Lithography VI*, Proceedings of SPIE 9422 (2015), 94221Q, https://proceedings.spiedigitallibrary.org/conference-proceedings-of-spie/9422/94221Q/Overlay-and-edge-placement-control-strategies-for-the-7nm-node/10.1117/12.2085761.full.

[207] Yuntao Jiang et al., "Discussion on Overlay Control for 2X nm Technology Node and Beyond," in *2015 China Semiconductor Technology International Conference* (2015), 1–3, https://ieeexplore.ieee.org/document/7153357/.

[208] LaPedus, "Litho Options Sparse after 10nm."

[209] Leachman, *Competitive Semiconductor Manufacturing*.

[210] Cyrus E. Tabery et al., "Voltage Contrast Edge Placement Estimation for Overlay, CD, and Local Uniformity Metrology" (conference presentation), in *Metrology, Inspection, and Process Control for Microlithography XXXIII*, Proceedings of SPIE 10959 (2019), 109591U, https://www.spiedigitallibrary.org/conference-proceedings-of-spie/10959/109591U/Voltage-contrast-edge-placement-estimation-for-overlay-CD-and-local/10.1117/12.2516613.full.

[211] ASML Holding N.V., *Annual Report on Form 20-F for the Fiscal Year Ended December 31, 2022* (Veldhoven, Netherlands: ASML Holding N.V., 2023), https://www.sec.gov/Archives/edgar/data/937966/000093796623000014/asml-20221231.htm.

[212] De Graaf et al., "NXT:1980Di Immersion Scanner."

[213] ASML, "ASML Extends ArF Leadership."

[214] Xiangqun Miao et al., "Double Patterning Combined with Shrink Technique to Extend ArF Lithography for Contact Holes to 22nm Node and Beyond," in *Optical Microlithography XXI*, Proceedings of SPIE 6924 (2008), 69240A, https://www.spiedigitallibrary.org/conference-proceedings-of-spie/6924/1/Double-patterning-combined-with-shrink-technique-to-extend-ArF-lithography/10.1117/12.772024.full.

[215] William Arnold, Mircea Dusa, and Jo Finders, "Manufacturing Challenges in Double Patterning Lithography," in *2006 IEEE International Symposium on Semiconductor Manufacturing* (2006), 283–286, https://ieeexplore.ieee.org/document/4493084.

[216] ASML, "ASML Ships New TWINSCAN NXT"; "TWINSCAN NXT:2050i," ASML, accessed September 26, 2026, https://www.asml.com/en/products/duv-lithography-systems/twinscan-nxt2050i; Boom, "DUV Products and Business Opportunity."

[217] Boom, "DUV Products and Business Opportunity."

[218] "DUV Lithography Systems," ASML, accessed September 26, 2026, https://www.asml.com/en/products/duv-lithography-systems.

[219] Ramish Zafar, "Morgan Stanley Guts SMIC's Huawei AI GPU Revenue by More Than 50% Due to Abysmal Yields," *Wccftech*, September 6, 2025, https://wccftech.com/morgan-stanley-guts-smics-huawei-ai-gpu-revenue-by-more-than-50-due-to-abysmal-yields/; "SMIC Posts Record $9.3B in 2025 Sales; 7nm Yields Reportedly Weigh on Margins," *TrendForce*, February 11, 2026, https://www.trendforce.com/news/2026/02/11/news-smic-posts-record-9-3b-in-2025-sales-7nm-yields-reportedly-weigh-on-margins/.

[220] Scotten W. Jones, "Economics in the 3D Era" (presentation prepared for Nikon LithoVision 2020, IC Knowledge LLC), posted on SemiWiki, March 4, 2020, https://semiwiki.com/wp-content/uploads/2020/03/Lithovision-2020.pdf.

[221] *Integrating and Operating HBM2E Memory: Achieve Maximum Performance from the Highest-Bandwidth DRAM in Micron's Ultra-Bandwidth Solutions*, technical brief (Micron Technology, 2021), https://assets.micron.com/adobe/assets/urn:aaid:aem:275edf31-79e3-4b6c-8bbd-a233babe9281/original/as/micron-hbm2e-memory-wp.pdf.

[222] Hwang Min-gyu 황민규, "SK 하이닉스, DDR5 시장서 독주… '14 나노 D 램 수율 90% 수준'" [SK hynix Dominates the DDR5 Market: "14nm DRAM Yield at the 90% Level"], *조선비즈* [Chosun Biz], June 13, 2023, via Daum, https://v.daum.net/v/20230613134559399; Kim Do-hyeon 김도현, "[삼성·SK 메모리 레이스] '전영현 오자마자 견제구' 하이닉스, 영업기밀 오픈" [Samsung–SK Memory Race: Hynix Fires a Warning Shot, Revealing Trade Secrets as Soon as Jun Young-hyun Arrives], *더벨* [The Bell], May 23, 2024, https://www.thebell.co.kr/front/newsview.asp?code=0402&key=2024052310234282801020 54.

[223] "SK hynix's Kangwook Lee Becomes Industry's First Packaging Executive to Win Dawon Kahng Award," SK hynix Newsroom, February 19, 2025, https://news.skhynix.com/en/sk-hynix-kangwook-lee-becomes-industry-first-packaging-executive-to-win-dawon-kahng-award/.

[224] Park Jin-hyung 박진형, "SK 하이닉스 'HBM3E 16 단 패키징 수율 문제 없다'" [SK hynix: "No Problems with HBM3E 16-High Packaging Yield"], *전자신문* [Electronic Times], November 5, 2024, https://www.etnews.com/20241105000299.

[225] Wang et al., "China's CXMT"; Hwang Min-gyu 황민규, "'중국발 HBM' 공포 과장됐나… 中 CXMT 상장 서류서 사라진 HBM 투자 계획" [Was the "Chinese HBM" Scare Overblown? HBM Investment Plans Vanish from CXMT's Listing Documents], *조선비즈* [Chosun Biz], July 8, 2026, https://cbiz.chosun.com/svc/bulletin/bulletin_art.html?contid=2026070802326.

[226] Wang et al., "China's CXMT."

[227] Van den Brink, "Industry Roadmap and Technology Strategy"; Samsung Electronics, "2nd-Generation, 10-Nanometer Class DRAM"; ASML, "Q3 Sales Strong"; ASML, "Earnings Growth Continues."

[228] Markus Kasanmascheff, "TSMC Says CoWoS Yield Tops 98 Percent as Capacity Expands," *WinBuzzer*, May 18, 2026, https://winbuzzer.com/2026/05/18/tsmc-cowos-yields-top-98-as-capacity-expands-xcxwbn/.

[229] Gregory C. Allen, *DeepSeek, Huawei, Export Controls, and the Future of the U.S.-China AI Race* (Washington, DC: Center for Strategic and International Studies, 2025), https://www.csis.org/analysis/deepseek-huawei-export-controls-and-future-us-china-ai-race.

[230] Raaijmakers and van Gerven, "Japan's Export Restrictions."

[231] Semiconductor Manufacturing International Corporation (SMIC), "SMIC Achieves First 45-Nanometer Silicon Success," press release (Exhibit 99.1 to Form 6-K), December 8, 2008, https://www.sec.gov/Archives/edgar/data/1267482/000130901408000703/exhibit1.htm.

[232] General Administration of Customs of the People's Republic of China, "海关统计数据在线查询平台" [Online Customs Statistics Query Platform], accessed September 26, 2026, http://stats.customs.gov.cn/queryData/queryDataByWhere.

[233] Gu Lingyu 顾翎羽 and Wu Junyu 吴俊宇, "美国禁运，ASML 光刻机中国销量却在飙升，为啥？" [Despite the US Embargo, ASML's Lithography Machine Sales in China Are Soaring—Why?], *虎嗅网* [Huxiu], November 9, 2023, https://m.huxiu.com/article/2278507.html.

[234] House Select Committee on the Strategic Competition between the United States and the Chinese Communist Party, *Selling the Forges of the Future: How the CCP Gets Its Key Equipment for Making Semiconductors from U.S., Dutch, and Japanese Companies* (Washington, DC: US House of Representatives, October 7, 2025), https://chinaselectcommittee.house.gov/media/reports/selling-the-forges-of-the-future; Jiufu Technology Group, "ASML 光刻机基座简介及案例-20250730" [Introduction to ASML Lithography Machine Foundations and Case Studies, 20250730], presentation, Scribd, https://www.scribd.com/document/963670109/ASML 光刻机基座简介及案例-20250730?doc_id=963670109.

[235] House Select Committee on the CCP, *Selling the Forges of the Future*.

[236] House Select Committee on the CCP, *Selling the Forges of the Future*.

[237] "调研国产半导体制造——中芯国际" [Researching Domestic Semiconductor Manufacturing: SMIC], Zhihu, accessed September 26, 2026, https://zhuanlan.zhihu.com/p/603619800.

[238] olalavn, "The number of advanced lithography machines currently owned in mainland China…," post in "Chinese Semiconductor Industry" thread, Sino Defence Forum, September 6, 2023, https://www.sinodefenceforum.com/t/chinese-semiconductor-industry.8479/page-2790#post-1019003.

[239] Sigmaintell, "中芯国际传 2025 年完成 5nm 芯片开发 但成本可能比 TSMC 版本高出 50%" [SMIC Reportedly to Complete 5nm Chip Development in 2025, but Costs May Be 50% Higher Than TSMC's], Sigmamirror, April 10, 2025, https://www.sigmamirror.com/index/article/detail/cid/6/id/1667.html.

[240] Dylan Patel, Afzal Ahmad, and Myron Xie, "China AI & Semiconductors Rise: US Sanctions Have Failed," *SemiAnalysis* (blog), September 12, 2023, https://newsletter.semianalysis.com/p/china-ai-and-semiconductors-rise.

[241] 15 C.F.R. pt. 744, supp. no. 4.

[242] olalavn, "The number of advanced lithography machines."

[243] Hyper, "Probably ymtc has 2050i…," post in "Chinese Semiconductor Industry" thread, Sino Defence Forum, February 8, 2022, https://www.sinodefenceforum.com/t/chinese-semiconductor-industry.8479/page-1091; Shen Yiran 沈怡然, "荷兰限制部分型号光刻机出口 中国市场影响几何" [Netherlands Restricts Exports of Certain Lithography Machine Models: How Much Will the Chinese Market Be Affected?], *经济观察报* [Economic Observer], via *澎湃新闻* [The Paper], March 11, 2023, https://www.thepaper.cn/newsDetail_forward_22262232.

[244] foofy, "From zhihu. Huahong Grace got one. SMIC got one…," post in "Chinese Semiconductor Industry" thread, Sino Defence Forum, April 5, 2021, https://www.sinodefenceforum.com/t/chinese-semiconductor-industry.8479/page-402#post-684397.

[245] "2018 年中国半导体行业现状：销量对外依存度依然强烈（图）" [State of China's Semiconductor Industry in 2018: Dependence on Foreign Supply Remains Strong (with Charts)], 观研报告网 [Guanyan Report Network], July 16, 2018, https://www.chinabaogao.com/detail/349332.html.

[246] ASML Holding N.V., *Annual Report on Form 20-F for the Fiscal Year Ended December 31, 2016* (Veldhoven: ASML Holding N.V., 2017), https://www.sec.gov/Archives/edgar/data/937966/000093796617000007/a20-fasml2016.htm; ASML Holding N.V., *Annual Report on Form 20-F for the Fiscal Year Ended December 31, 2017* (Veldhoven, Netherlands: ASML Holding N.V., 2018), https://www.sec.gov/Archives/edgar/data/937966/000093796618000007/a2017integratedreportbased.htm.

[247] House Select Committee on the CCP, *Selling the Forges of the Future*.

[248] CXMT Corporation, *首次公开发行股票并在科创板上市招股说明书* [Prospectus for Initial Public Offering and Listing on the STAR Market] (Hefei: CXMT Corporation, 2026), https://static.sse.com.cn/disclosure/listedinfo/announcement/c/new/2026-07-22/688825_20260722_Q1T7.pdf; CXMT Corporation, *财务报表及审计报告：2025 年 1 月 1 日至 6 月 30 日止期间、2024 年度、2023 年度及 2022 年度* [Financial Statements and Audit Report for the Period from January 1 to June 30, 2025, and for 2024, 2023, and 2022], audited by Deloitte Touche Tohmatsu Certified Public Accountants LLP, 新浪财经 [Sina Finance], December 30, 2025, https://vip.stock.finance.sina.com.cn/corp/view/vCB_AllBulletinDetail.php?id=11891212&stockid=688825.

[249] Jiufu Technology Group, "ASML 光刻机基座简介及案例."

[250] Romeo Dean, "Compute in Plan A," AI 2040, accessed September 26, 2026, https://ai-2040.com/supplements/compute-supplement.

[251] Dean, "Compute in Plan A."

[252] Dylan Patel, interview by Dwarkesh Patel, "Dylan Patel—Deep Dive on the 3 Big Bottlenecks to Scaling AI Compute," *Dwarkesh Podcast*, March 13, 2026, https://www.dwarkesh.com/p/dylan-patel.

[253] ASML Holding N.V., "ASML Reports €9.3 Billion Total Net Sales and €2.9 Billion Net Income in Q2 2026," press release (Exhibit 99.1 to Form 6-K), July 15, 2026, https://www.sec.gov/Archives/edgar/data/0000937966/000162828026048235/pressreleasefinancialresul.htm; "ASML Plans 30% Increase in Low-NA EUV Production Capacity by 2027," *Crypto Briefing*, July 15, 2026, https://cryptobriefing.com/asml-low-na-euv-production-capacity-2027/.

[254] Dwarkesh Patel (@dwarkesh_sp), ".@dylan522p lays out how we know the hard upper bound on how much compute can be produced annually by 2030: around 200 GW/year…," X, March 13, 2026, https://x.com/dwarkesh_sp/status/2032514120461988204.

[255] ASML Holding N.V., "ASML Reports €9.3 Billion Total Net Sales and €2.9 Billion Net Income in Q2 2026," press release, July 15, 2026, https://www.sec.gov/Archives/edgar/data/937966/000162828026048235/pressreleasefinancialresul.htm.

[256] Hunt, Khan, and Peterson, *China's Progress in Semiconductor Manufacturing Equipment*.

[257] ASML Holding N.V., "ASML Reports €9.3 Billion Total Net Sales."

[258] Somala, "Will Huawei Catch Up."

[259] Dylan Patel, AJ, and Myron Xie, "Huawei Ascend Production Ramp: Die Banks, TSMC Continued Production, HBM Is the Bottleneck," *SemiAnalysis* (blog), September 8, 2025, https://newsletter.semianalysis.com/p/huawei-ascend-production-ramp.

[260] Somala, "Will Huawei Catch Up."