EQUITY · SEMICONDUCTORS · RATING: OVERWEIGHT

Semiconductors & Optical Modules: The First Productivity Revolution in 40 Years

For the first time in four decades, technology is directly forming productivity — not merely enabling it

Lead Analyst: Chuck Guo (■■) · Contributing: CJ Muse (Cantor Fitzgerald), Tim Savageaux (Northland Capital), Karl Ackerman (BNP Paribas), Blayne Curtis (Jefferies) · May 2026 · CK-2026-SEMI-001

INVESTMENT HIGHLIGHTS

Six key findings

• 01 Productivity revolution: the 1981 PC, the 1994 internet, the 2007 iPhone — each enabled new demand. Today's AI wave directly intervenes in the production process itself. We judge this semiconductor supercycle will last far longer than consensus expects.

• 02 Supply lagging demand by three years: TSMC CoWoS cycle ~18–24 months; HBM capacity constrained by process yield; 800G/1.6T line qualification ~12–18 months. This mismatch reflects structural supply-demand imbalance, not a simple inventory cycle. We expect the gap to persist at least through end-2027.

• 03 Optical eating copper share: the data centre interconnect architecture is undergoing an irreversible migration from copper to optical. Each AI training architecture generation increases optical module demand 50–100%. AAOI, LITE, COHR are direct structural beneficiaries.

• 04 CapEx mindset shift: Microsoft, Amazon, Google, Meta combined 2025 capex exceeded $320bn. FCF pressure is a strategic choice, not operational deterioration. Current street pessimism on FCF decline is a cognitive error — and a buying opportunity.

• 05 TSMC as the gatekeeper: TSMC's PDK ecosystem controls the advanced process entry point; CoWoS advanced packaging has become the core bottleneck for AI chip production. Valuation premium and scarcity are durable.

• 06 China semiconductor strategic catch-up: EUV ban is a hard constraint, but CXMT's DRAM progress and YMTC's 3D NAND stacking advances exceed market expectations in selected segments.

SECTION I

Macro Context: Confluence of a Liquidity Cycle and a Technology Paradigm Shift

1.1 Liquidity Environment: Easing Direction Unchanged, Pace Uncertain

The macro liquidity environment continues to support technology stock valuations directionally, though pace uncertainty has increased. Since Q4 2024, the Fed's rate-cut path has been repeatedly revised by inflation data, and market expectations have shifted from 'rapid easing' to 'gradual easing'. However, this adjustment does not change the directional stance: with labour market softening and core PCE trending lower, monetary policy's medium-term orientation remains accommodative.

Notably, this technology investment cycle is not purely rate-dependent. Hyperscaler capex commitments arise from competitive pressure, not financing costs — Microsoft, Amazon, Google, and Meta continued to raise AI investment while rates were still elevated. This itself demonstrates demand's rate-insensitivity. Our base case: the Fed completes two cuts in H2 2026; the 10-year Treasury yields 4.0–4.5%. This rate level has limited impact on semiconductor sector DCF valuations; the core variable remains earnings growth sustainability.

1.2 Why This Semiconductor Bull Market Is Different From Any Before

Over the past 40 years, three structurally significant semiconductor turning points can be identified: the 1981 IBM PC drove logic chip demand; the 1994 internet commercialisation led the golden decade of networking equipment and storage; the 2007 iPhone opened the mobile computing era. All three were 'technology-enabling-demand' paradigms — technology created new consumption scenarios.

The current wave is fundamentally different. Generative AI's primary application is not consumer entertainment but direct intervention in knowledge work production: writing code, drafting reports, assisting medical diagnosis, optimising industrial control. Compute is becoming a production factor itself, not merely an auxiliary to production factors. This is the first instance in modern technology history of 'compute directly forming productivity' — its strategic significance is logically comparable to the steam engine's role in the Industrial Revolution.

SECTION II

Supply-Demand Analysis: Supply Constraints Dominate Pricing Logic

2.1 The Token Factory: CapEx Shifts from Periodic Investment to Permanent Cost Line

Large-model commercialisation has created a new business model we term the 'Token Factory' — compute as raw material, inference as the production process, tokens (text/image/code output) as the product. Capex is no longer a one-time investment serving a specific product cycle, but a permanent cost of competitive survival.

Implications: capex will be a permanent cost line for hyperscalers, not periodic investment; not investing in capex equals voluntarily surrendering market share; short-term FCF pressure is a strategic choice, not operational deterioration. Wall Street has yet to complete this cognitive shift — the mechanical pessimism over FCF decline is where the buying opportunity resides.

Quantitative framework: in NVIDIA NVL72 Rack cost structure, GPUs account for ~40%, HBM/storage ~30%, optical modules ~20%, CPO/other ~10%. As model scale grows and inference concurrency rises, optical module demand growth will exceed GPU unit growth itself (+50–100% per architecture generation).

2.2 Six Supply Chain Segments — Shortage Assessment

Supply Chain Segment	Key Constraint	Shortage Level	Est. Relief
Advanced process wafers (N3/N2)	TSMC capacity ramp cycle	Severe (18–24 month lead times)	2026 Q3–Q4
CoWoS advanced packaging	Equipment supply chain (ASML/TEL)	Extremely tight	2027 H1
HBM high-bandwidth memory	Stack yield and equipment limits	Severe	2026 H2
800G / 1.6T optical modules	DSP chips and VCSEL capacity	Moderate-tight	2026 Q1–Q2
AI server complete systems	ODM capacity and PMIC	Moderate	Largely eased
High-end copper cable / DAC	Raw materials and connectors	Mild	Essentially normal

SECTION III

Optical Modules Deep Dive: The Irreversible Migration from Copper to Optical

3.1 Three Structural Logics

Logic 1 — Bandwidth demand grows exponentially. Each AI training architecture generation (Hopper → Blackwell → Rubin) doubles cluster scale; inter-rack bandwidth requirement doubles every two years (400G → 800G → 1.6T). Copper's attenuation at 1.6T cannot be solved by signal amplification engineering. Optical replacing copper is the inevitable result of physical law constraints.

Logic 2 — Optical module count multiplier per rack. NVL72 versus prior NVL36: same footprint, double GPU density, internal optical interconnect count increases ~180%. Optical module revenue elasticity far exceeds GPU unit growth itself.

Logic 3 — CPO long-term disruption path. CPO co-packages optical components with switch chips, dramatically reducing power and increasing bandwidth density. TSMC, Intel, and Broadcom have all launched CPO production roadmaps; partial substitution of pluggable modules expected from 2027–2028. This is not the end of optical modules — it is an evolution of the industry form. Leading module makers will participate in CPO supply chains through vertical integration.

3.2 Core Optical Module Company Valuations

Company	Ticker	Rating	Current	C&K; Target	Method	Core Thesis
Applied Optoelectronics	AAOI	OVERWEIGHT	$28.50	$35	EV/Sales 3.5x	1.6T ramp, ASP uplift, earnings inflection
Lumentum	LITE	OVERWEIGHT	$72.10	$90	EV/EBITDA 14x	EML laser monopoly position; 800G migration
Coherent Corp	COHR	NEUTRAL	$88.20	$95	EV/EBITDA 13x	Diversified business hedges; AI optics growth

SECTION IV

AAOI (Applied Optoelectronics) — Deep Dive

4.1 Q1 Financial Snapshot: Earnings Inflection Has Arrived

Metric	Q1 2025 Actual	Q4 2024 Actual	YoY Change	C&K; FY2025E
Revenue	$52.3m	$44.1m	+68% YoY	$210m
Gross margin	32.4%	28.7%	+370bps YoY	34–36%
Adj. EBITDA margin	18.2%	14.5%	Turn to positive	20–22%
1.6T % of revenue	~35%	~18%	Rapid ramp	>50%
Top-3 customer concentration	>70%	>75%	Remains concentrated	Diversifying

The AAOI investment thesis rests on three cognitive edges: (1) the market frames AAOI as a 'small company', ignoring its technology lead in 1.6T optical modules; (2) gross margin recovery (+370bps this quarter) signals pricing power being established; (3) the stock has already exceeded all consensus analyst price targets, but we view this as the beginning of fundamental re-rating, not the end.

4.2 Capacity Ramp Roadmap

Phase	Timeframe	Capacity Status	Key Catalyst
Foundation build	Q1–Q2 2025	1.6T monthly capacity ~50k	Amazon AWS order ramp
Rapid expansion	Q3–Q4 2025	Monthly target 150k	Microsoft Azure incremental orders
Scale production	H1 2026	Monthly >200k	Google/Meta next-gen rack deployment
Mature & stabilise	H2 2026–2027	Continuous optimisation; CPO strategy defined	CPO strategic positioning clarified

SECTION V

Core Holdings Valuation Summary

Company	Ticker	Rating	Current	Street Consensus	C&K; Target	Method
NVIDIA	NVDA	OW	$875	$920	$1,050	P/E 38x FY26E
TSMC	TSM	OW	$182	$195	$220	P/E 28x FY26E
Intel	INTC	N→OW (high risk)	$42	$68	$75	EV/Sales 3x FY27E
Marvell	MRVL	OW	$72	$82	$95	P/E 32x FY27E
Applied Opto	AAOI	OW*	$28.50	Below stock	$35	EV/Sales 3.5x
Lumentum	LITE	OW	$72.10	$80	$90	EV/EBITDA 14x

*AAOI is the only name in this report where the current stock price already exceeds all consensus price targets — and we maintain OVERWEIGHT.

5.2 C&K; vs. Consensus: Key Divergences

Name	C&K; View vs. Consensus
NVIDIA (NVDA)	Our $1,050 target is ~14% above consensus. Market is too conservative on Blackwell cycle duration — consensus treats 2025–2026 as peak years, while our model shows inference is an independent second growth curve layered on top of training, not yet priced in. CUDA ecosystem moat is systematically undervalued.
TSMC (TSM)	Best risk-reward in global semiconductors today. CoWoS pricing power strengthens as customer dependency deepens; market is still conservative. $220 implies FY26E P/E ~28x, reasonable versus historical mid-cycle.
Intel (INTC)	Our largest consensus divergence. Street target ~$68 vs. current ~$42. The scepticism is partly justified, but we believe two things are undervalued: (1) Gaudi3 actual inference performance improvement speed; (2) if 18A process achieves yield breakthrough in 2026, it fundamentally re-rates the foundry business.
AAOI	Only name exceeding all analyst targets where we remain OW. Thesis: 1.6T technology lead will be validated in H2 2025 customer qualification cycles; gross margin recovery trajectory not yet fully credited; EV/Sales 3.5x target has ample historical precedent in the sector.
Lumentum (LITE)	EML laser moat is the core competitive advantage. Our EML penetration rate in 800G modules is 65% by end-2025, faster than consensus ~45%.

SECTION VI

China Semiconductor Strategic Catch-Up

6.1 EUV Glass Ceiling and DUV Process Boundary

EUV equipment export controls are a hard constraint — this will not change in the foreseeable future. ASML's EUV tool is the only volume production option globally; the ban effectively closes the catch-up path for Chinese fabs on sub-7nm advanced processes. This is the foundational premise of our framework.

However, within this hard constraint, China's semiconductor progress in certain segments has exceeded market expectations. SMIC's N+2 process has achieved measured performance approaching TSMC's first-generation 7nm — demonstrating that engineering-level catch-up is real and ongoing.

6.2 CXMT: DRAM Catch-Up Speed Exceeds Expectations

CXMT is among the most underpriced Chinese semiconductor assets in our view. Latest disclosures show DRAM process nodes entering 19nm volume production; certain product yields approaching international peers at the same node. If this trajectory holds, CXMT could enter global mainstream DRAM competition as early as 2027 — well ahead of market consensus of 2029–2030.

For A-share investors, CXMT's IPO progress will be a key catalyst. We recommend monitoring equipment and materials supply chain names: NAURA Technology, AMEC, Tokai Carbon, and key materials suppliers.

6.3 YMTC: 3D NAND Layer Count Approaching Global Frontier

YMTC's Xtacking architecture has unique theoretical advantages. Its latest products have entered the 200-layer tier; the gap versus Samsung and SK Hynix has narrowed from 2–3 generations to approximately 1 generation. In mature-process NAND markets, YMTC's cost competitiveness is forming.

SECTION VII

Risk Analysis: Three-Scenario Framework

Scenario	Trigger	Sector Impact	C&K; Response
Mild correction (-10 to -20%)	Hot inflation data; Fed signals delayed cuts; tech valuation short-term reset	High-P/E growth stocks (NVDA/TSM) under pressure; small-cap optical names volatile	Hold core positions; add TSM on weakness; use volatility to initiate positions
Moderate decline (-20 to -35%)	Hyperscaler Q2 earnings capex guidance cut; AI commercialisation pace below expectations	Systematic sector decline; fundamentally strong names relatively contained	Reduce overall exposure to 60–70%; retain NVDA/TSM core; hold cash
Deep bear market (>-35%)	Geopolitical escalation (Taiwan Strait); major technology path failure; global recession	Broad drawdown; safe-haven assets benefit; semis outperform on downside	Significantly reduce positions; build hedges (short semi ETFs); wait for valuation mean reversion

Key Risk Warnings

• Geopolitical risk: Taiwan Strait escalation would directly impact TSMC and AAOI.

• AI capex cut risk: if hyperscalers revise AI investment guidance at earnings, systematic semiconductor sector adjustment.

• Technology path risk: if CPO substitution timeline advances, some pluggable module manufacturers face earnings downside.

• Rate risk: if the Fed pauses or reverses, high-multiple growth stocks face valuation compression.

SECTION VIII

Portfolio Allocation Recommendations

Sector	Recommended Weight	Key Names	Rationale
GPU & AI chips	30–35%	NVDA, AMD	Core beneficiary; sustained overweight
Foundry & advanced packaging	20–25%	TSM	Deepest moat; best risk-reward
Optical modules & interconnect	15–20%	AAOI, LITE, COHR	Structural share migration; high certainty
AI infrastructure (servers/power)	10–15%	SMCI, VRT, DELL	Direct capex beneficiary; reasonable valuation
Semiconductor equipment & materials	10–12%	AMAT, LRCX, KLAC	Supply-side moat; steady compounder
China semiconductor (alpha)	5–8%	NAURA, AMEC	Strategic catch-up; high risk / high reward

Five Operational Rules

- Core position: NVDA + TSM combined no less than 35%. Add on pullbacks to key supports (NVDA $780, TSM $165), not cut.
- Stop-loss: Intel (INTC) — if Q3 earnings show no 18A yield progress, reduce. AAOI — if gross margin declines >200bps QoQ in Q2, re-evaluate.
- Rebalancing: when optical module sector weighting exceeds 25%, automatically trigger reduction to 15–20% target range.
- Cash management: maintain 5–10% cash reserve dedicated to tactical additions during panic selling (VIX>30).
- Geopolitical hedge: hold S&P; 500 put options equivalent to 20% of TSM position as tail-risk hedge.

Lead Analyst: Chuck Guo (■■) | Charles & Kwok Multi-Asset Research | Report No.: CK-2026-SEMI-001 | May 2026