TSMC

CoWoS and Advanced Packaging Updates

Overweight

2330.TW, 2330 TT
Price (09 Jul 26):NT$2415.0
Price Target (Jun-27):NT$3100.0

- **Raising CoWoS capacity further at TSMC and at a faster pace at OSATs:** We are raising our FY26/27/28 CoWoS capacity estimates further by 1%/9%/11%, with TSMC's CoWoS capacity reaching 115k/190k/225k wfpm by end 26/27/28. TSMC appears to have accelerated its capacity expansion plans in AP7 Phase 2 (now 50-60% allocated to CoWoS, compared to 100% SoIC in prior plans) and also outsourcing interposer wafer manufacturing selectively to Vanguard, to bring on more capacity. TSMC also seems to be again expanding some CoWoS-R capacity, in response to strong CPU and accelerator demand from NVDA and AWS, respectively. At the same time, OSAT capacity expansion for CoWoS-like processes is picking up momentum, and we expect OSATs CoWoS-like capacity to reach 15k/50k/85k wfpm by end 26/27/28. Bulk of OSAT capacity is likely to be deployed for server CPU packaging, followed by networking and small volumes for accelerator chips such as LPU or Trainium 3. Overall, the demand-supply gap has recently expanded to ~20%, given the meaningful upside from agentic AI CPU demand.
- **3DSoIC ramp continues, but more focused on 2028, to intercept mass adoption in accelerators:** We expect SoIC capacity ramp into 2028 to remain quite strong (reaching 65k wfpm by end 2028), but 2027 ramp may be pushed out a bit, as TSMC is trying to accommodate more CoWoS demand in its limited backend capacity at AP7 P2. In addition to AMD (for MI series) and Apple (for M5 Max and Ultra CPUs), we expect AWS Trainium 4, OpenAI, Nexus 2, Meta MTIA 450/500, and some versions of NVDA Feynman to adopt SoIC packaging in 2028. TPUs are likely to adopt SoIC only for v10. While the market is excited about CPO demand for SoIC (for EIC/PIC integration), the actual wafer demand is still quite small, in our view, given small die sizes. 10-15k wfpm can be accommodated with the current pace of capacity ramp.
- **CoPoS is still in the engineering stage, only likely to enter risk production in 2028, with mass production in 2029:** Our research indicates that TSMC's CoPoS is still in the engineering stage and TSMC is still ironing out process issues and finalizing the manufacturing flow. We expect CoPoS to enter risk production in early 2028, with mass production likely in 2029. At this stage, we don't expect NVDA Feynman GPUs to be adopting CoPoS in 2H28, given the very short lead-time from risk production, and expect Feynman Ultra to be the first product adopting CoPoS.
- **NVDA forecasts rising in 2027/28 due to CPUs, LPU and CPO:** We are raising our NVDA CoWoS allocation by 7% for 2027 and expect another 45% growth in 2028, to reflect continued growth in accelerator demand (2-die Rubin Ultra with slightly larger package size in 2027, and likely 4-die Feynman as the high-end SKU in 2028), strong pick-up in CPU demand (Vera CPUs in 2027, Rosa CPUs in 2028), incremental demand from LPUs and Spectrum X CPO packages in 2027-28. Looking into 2028, there could be further upside to CoWoS demand if NVLink Switch CPO for Feynman ramps up in meaningful scale. All accelerators and CPO products for NVDA are on CoWoS-L at TSMC, while CPUs and LPU chips are likely to be on CoWoS-R, shared between TSMC, Amkor and eventually ASE. At this point, we believe Feynman is likely to stick with a CoWoS-L architecture with some SoIC logic

Technology and Telecoms

Gokul Hariharan AC
(852) 2800-8564
gokul.hariharan@jpmorgan.com
J.P. Morgan Securities (Asia Pacific) Limited/ J.P. Morgan Broking (Hong Kong) Limited

Jennifer Hsieh
(886-2) 2725-9868
jennifer.hsieh@jpmorgan.com
J.P. Morgan Securities (Taiwan) Limited

David Chou
(886-2) 2725-9618
david.chou@jpmorgan.com
J.P. Morgan Securities (Taiwan) Limited

Jason Chen
(886-2) 2725-9864
jason.bh.chen@jpmorgan.com
J.P. Morgan Securities (Taiwan) Limited

Subham Singhania
(91-22) 6157-3801
subham.singhania@jpmorgan.com
J.P. Morgan India Private Limited

Quarterly Forecasts (FYE Dec)

Adj. EPS (NT$)

	2025A	2026E	2027E
Q1	13.94	22.08A	31.09
Q2	15.36	24.72	32.67
Q3	17.44	27.39	36.83
Q4	19.50	29.79	37.66
FY	66.24	103.97	138.24

Style Exposure

Quant Factors	Current %Rank	Hist %Rank (1=Top) 6M	1Y	3Y	5Y
Value	43	47	33	44	78
Growth	29	18	11	23	24
Momentum	32	17	29	68	42
Quality	14	14	10	13	7
Low Vol	29	32	37	44	29
ESGQ	27	84	27	16	3

Sources for: Style Exposure – J.P. Morgan Global Markets Strategy; all other tables are company data and J.P. Morgan estimates.

See page 12 for analyst certification and important disclosures, including non-US analyst disclosures.

Gokul Hariharan [AC]
(852) 2800-8564
gokul.hariharan@jpmorgan.com

J.P.Morgan

Price Performance



	YTD	1m	3m	12m
Abs	55.8%	4.8%	23.5%	121.6%
Rel	-0.8%	3.3%	-6.6%	20.2%

Company Data

Shares O/S (mn)	25,930
52-week range (NT$)	2535.0-1065.0
Market cap ($ mn)	1,954,246
Exchange rate	32.04
Free float (%)	92.1%
3M ADV (mn)	41.90
3M ADV ($ mn)	2,967.8
Volatility (90 Day)	36
Index	TAIEX
BBG ANR (Buy | Hold | Sell)	39|1|0

Key Metrics (FYE Dec)

NT$ in millions	FY25A	FY26E	FY27E	FY28E
Financial Estimates				
Revenue	3,809,054	5,342,260	7,192,208	8,813,415
Adj. EBIT	1,936,092	3,175,201	4,275,861	5,262,336
Adj. EBITDA	2,624,188	3,968,123	5,280,152	6,544,105
Adj. net income	1,717,883	2,696,350	3,584,875	4,411,972
Adj. EPS	66.24	103.97	138.24	170.13
BBG EPS	64.73	99.62	128.33	159.03
Cashflow from operations	2,437,716	3,413,412	4,421,425	5,578,226
FCFF	1,292,759	1,860,700	1,949,175	3,105,976
Margins and Growth				
Revenue Growth Y/Y (%)	31.6%	40.3%	34.6%	22.5%
EBIT margin	50.8%	59.4%	59.5%	59.7%
EBIT Growth Y/Y (%)	46.4%	64.0%	34.7%	23.1%
EBITDA margin	68.9%	74.3%	73.4%	74.3%
EBITDA Growth Y/Y (%)	32.2%	51.2%	33.1%	23.9%
Net margin	45.1%	50.5%	49.8%	50.1%
Adj. EPS growth	46.4%	57.0%	33.0%	23.1%
Ratios				
Adj. tax rate	16.0%	17.7%	17.7%	17.7%
Interest cover	NM	NM	NM	NM
Net debt/Equity	NM	NM	NM	NM
Net debt/EBITDA	NM	NM	NM	NM
ROE	35.4%	41.4%	40.1%	36.9%
Valuation				
FCFF yield	2.1%	3.0%	3.1%	5.0%
Dividend yield	0.9%	1.2%	1.4%	1.7%
EV/Revenue	16.2	11.3	8.3	6.5
EV/EBITDA	23.5	15.2	11.3	8.8
Adj. P/E	36.5	23.2	17.5	14.2

Summary Investment Thesis and Valuation

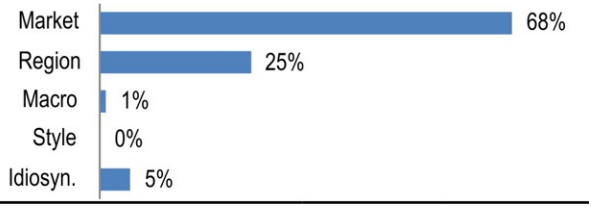
Investment Thesis

We expect TSMC's structural growth drivers to remain very strong as leading-edge supply (N4, N3, and N2) stays tight well into 2027 or early 2028, given accelerating AI demand, continued growth in N3, N2 and advanced back-end revenues, and long lead time for capacity build despite accelerated capex. Near-term, we believe GM is also likely to remain strong, nudging 70% levels, given N3 improvement, high UTR and decent pricing, despite N2 and overseas dilution. We expect broader price hikes in 2027, given the widening supply-demand gap and a stronger Datacenter AI CAGR (we now estimate 69% CAGR from 2024-29E, driven by both higher accelerator estimates and a surge in AI CPU numbers).

Valuation

Our Jun-27 PT of NT$3,100 is based on ~20x 12-month forward P/E and reflects stronger near-term profitability, better visibility on AI demand, stronger capacity build and firmer pricing into 2028. Our target multiple is higher than TSMC's five-year average historical multiple.

Performance Drivers



Factors	6M Corr	1Y Corr
Market: MSCI Asia Pac ex JP	0.88	0.83
Region: Taiwan	0.83	0.89
Macro:		
JPM EM Currency(EMCI) Fixing	-0.24	-0.22
JP Morgan GBI-EM Global Div	-0.17	-0.17
Emerging Central Bank Rate	-0.07	-0.16
Quant Styles:		
Momentum	0.18	0.20
Quality	0.14	0.15
DivYld	-0.18	-0.10

Source: J.P. Morgan Global Markets Strategy for Performance Drivers; company data, Bloomberg Finance L.P. and J.P. Morgan estimates for all other tables. Note: Price history may not be complete or exact.

stacking for certain variants and could approach 9-10x reticle size on CoWoS-L for the high-end versions (we expect 2-die and 4-die package variants for this product). We have not seen evidence of CoPoS based products for NVDA in 2028 so far.

- **Significant upside for AMD from server CPUs, MI450 seeing slight upside:** Our AMD CoWoS estimates are rising significantly in 2027-28, to reflect a very strong demand uptick for the Venice family of server CPUs (using CoWoS-L). We believe that AMD could produce ~3M/4.5M units of Venice CPUs in 2027/28, largely utilizing CoWoS-L like process from ASE. MI450 also appears to be ramping up smoothly, reaching 1.5M units in 2027, while MI500 with ultra-large reticle size (14-15x on CoWoS-L) should enter production in 2028, based on our checks. AMD is also using variants of CoWoS and 2.5D Wafer Level Fan Out (called Elevated Fan Out Bridge) for several of its high-end PC CPUs, gaming GPUs and Xbox Custom silicon from 2027, and hence the aggregate demand for wafer-level Fanout is likely to reach 50-60k wfpm, with fulfillment ratios only running at 50-60% currently. In our exercise, we are limiting our forecasts to AMD's AI accelerators and server CPUs only.
- **TPU volumes nudging higher, Mediatek seeing bigger upside in 2027:** We raise 2027 TPU shipment forecast from 8M to 8.9M units, with the upside mainly from MediaTek's TPU v8t. We now expect v8t to contribute 3M units (vs. 2.5M previously) in 2027, alongside 4.2M for v8i Sunfish (vs. 3.8M previously) and the remainder from v7 Ironwood. For 2028, we expect MediaTek's TPU v9 (Humufish) to enter mass production (shipping 3M units in 2028) alongside Sunfish Ultra (four Sunfish compute dies vs. two in Sunfish), driving total TPU volumes to grow 30% to 11.6M units (using a 2-die equivalent for Sunfish). Mediatek's CoWoS allocation is likely to drop in 2028, given Humufish v9 is likely to be on Intel EMIB-T packaging. Our supply chain checks do not indicate Pumafish v9 ramping in 2028 at this point in time. We currently expect Mediatek's TPU share to be at 30–40% in 2027/28 and believe risk remains to the upside, depending on the ability to secure more substrate and wafer capacity.
- **Trainium3 volumes continue to move up; Trn4 likely a meaningful step up in design complexity:** Our checks indicate steady uptick in Trainium 3 demand and forecasts; we are now modeling 4.8M units through the lifecycle, with demand remaining well ahead of supply. We believe that AWS/Alchip will account for the vast majority of the CoWoS procurement, while some volume is being outsourced to Marvell, given the very tight capacity. We expect Trainium 4 to ramp up in 2028, with a meaningful step-up in design complexity, using 3D-SoIC for logic-stacking and multiple interconnect options. We expect AWS/Alchip to remain the primary suppliers (for Neuronlink and UALink versions) while Marvell could have some share for variants that employ NVLink fusion based interconnect. We expect Trainium 4 volumes to be relatively flattish in 2028, compared to Trainium 3 in 2027, but the wafer consumption and chip ASPs are likely to be significantly higher.
- **Other ASICs still small; 2027 key ramp year with focus on MTIA, MAIA and OpenAI, with Space X as a wildcard in 2028:** TPU and Trainium remain the only two large-scale ASIC programs in 2027, in our view. Several additional programs bear watching—including Microsoft Maia, Meta MTIA, OpenAI Jalapeño, and Tesla/SpaceX AI ASIC—but we believe they will remain at a much smaller scale in 2027 or even 2028. Microsoft is ramping Maia 200 in 2H26, though volumes remain modest (0.2–0.3M lifetime units, in our view); larger scale likely requires Maia 300 (likely ramp in 2H27) or Maia 400 thereafter. Meta has outlined 3-4 ASIC projects in its 2026-28 roadmap focused mainly on inference; but supply chain indicators still suggest a modest ramp in volumes in 2027. OpenAI introduced its first ASIC, Jalapeño, in 2H26, but we expect volumes to remain small in 2027, with its second chip, Nexus 2, likely to have higher volumes in 2028/29. We believe SpaceX/Tesla is the largest wildcard for 2028, with its AI ASIC project currently in RFQ stage with a likely ramp in late 2028.
- **CPUs remain a key upside catalyst in 2027, more CPUs could embrace CoWoS in**

2028: Server CPUs have emerged as the primary upside catalyst for CoWoS estimates into 2027, with NVDA Vera (~5M units), AMD Venice (3M+ units) all adopting CoWoS packaging for the first time. Looking into 2028, we believe that some of the ARM-based CPUs from hyperscalers could also follow this path and embrace 2.5D packaging, which could result in further demand upside. Much of the CPU CoWoS packaging demand is likely to benefit OSATs (ASE for Venice, Amkor and ASE for Vera CPUs), as TSMC is focusing its resources on higher reticle size CoWoS for accelerators and ramping up 3DIC packaging.

- **Intel EMIB seeing traction with Mediatek TPUv9, follow-on demand is still uncertain:** Given the acute shortage in advanced packaging capacity, Intel's EMIB-T process is also seeing incremental demand. The marquee project for Intel is Mediatek's TPUv9 Humufish, and our checks indicate that the progress seems satisfactory. We believe that substrate level yields are around 60%, but should be acceptable 1-1.5 years away from ramp. Intel's EMIB process is positioned as a cheaper option, with costs at ~50% of CoWoS-L, and enabling larger reticle size packages. However, substrate capacity is likely to be a key limiting factor for capacity ramp, in our view, since the substrate build-up and embedding of the silicon bridge steps are still at low yields. Beyond Google TPUv9, we've seen some interest from Amazon (for small volumes of Trainium 3/4) and certain CPU customers. A key watchpoint for EMIB progress would be Google's TPU v10 design choice, given the likely adoption of 3DSoIC logic die stacking, could push Mediatek to revert back to TSMC CoWoS-L for better end-to-end process simplicity.

Note: Alphabet (GOOGL, OW), Amazon.com (AMZN, OW) and Meta Platforms (META, OW) are covered by Doug Anmuth. Advanced Micro Devices (AMD, N), Broadcom (AVGO, OW), Marvell Technology (MRVL, OW), Intel (INTC, UW) and NVIDIA (NVDA, OW) are covered by Harlan Sur. Amkor Technology (AMKR, OW) is covered by Peter K. Peng. Apple (AAPL, OW) is covered by Samik Chatterjee, CFA. Microsoft (MSFT, OW) is covered by Mark R. Murphy. Tesla (TSLA, UW) is covered by Ryan Brinkman. Open AI is covered by Brenda Duverce.

Gokul Hariharan AC
(852) 2800-8564
gokul.hariharan@jpmorgan.com

J.P.Morgan

Table 1: CoWoS allocation breakdown

							% Revision vs prior ests		
CoWoS Consumption by customer (k wafer per year)	**2023**	**2024**	**2025**	**2026E**	**2027E**	**2028E**	**2026E**	**2027E**	**JPM Comments**
NVIDIA	70	192	400	725	1196	1735	*1%*	*7%*	Mild uptick in 2026 and meaningful uptick in 2027 from Vera CPU. 2028 includes upside from Feynman (4-die), LPX, Rosa CPU and Spectrum X ramp. Likely to see upside from NVLink Switch CPO. No CoPoS in 2028.
At TSMC	*70*	*127*	*385*	*685*	*1076*	*1385*	*1%*	*8%*	2026 CoWoS-L at TSMC, others at OSAT, 2027 Rubin Ultra assuming 2-die-on-interposer solution.
CoWoS-S / R	*70*	*97*	*30*	*10*	*40*	*105*	*100%*		TSMC adding some CoWoS-R to support Vera and future LPX projects now.
CoWoS-L	*0*	*30*	*355*	*675*	*1036*	*1280*	*0%*	*4%*	Majority of CoWoS to migrate to CoWoS-L by 2H25. Rubin Ultra likely to remain on CoWoS-L. Feynman likely to be still on CoWoS, given slower progress for CoPoS - likely to need 9-10x reticle size package.
Outside TSMC (CoWoS-S/R)	*0*	*65*	*15*	*40*	*120*	*350*	*0%*	*-4%*	Outsourcing ratio picking up in 2026 again as Vera outsourcing to Amkor starts to pick up more momentum in 2027, followed by ASE/SPIL. LPU40 on CoWoS-R could also be outsourced to OSATs in 2028.
AMD	10	45	52	95	355	540	*19%*	*97%*	MI450 ramp in late 2026, moving to CoWoS-L for MI450 onwards, forecasts incl. Xilinx. Meaningful upside forecasts to 2027 due to MI400 and Venice CPU demand. Venice in ASE from 2027, continued upside in 2028.
CoWoS-S	10	45	52	25	0	0	*0%*	*-100%*	MI350 and Xilinx - main CoWoS-S customer. AMD also has non-AI projects migrating to 2.5D packaging in 2H26, likely to be mostly at ASE but not captured in these forecasts.
CoWoS-L	0	0	0	55	160	230	*0%*	*-6%*	CoWoS-L move likely only for MI400 family and beyond, Venice family of CPUs also to use CoWoS-L.
Outside TSMC (CoWoS-L/ FoCoS / EFB)	0	0	0	15	195	310			Mostly Venice CPU, but several PC, Gaming and Xbox chips also using Wafer-Level Fan-Out from 2027 - we are modeling only the server CPU parts.
Broadcom	40	70	96	250	425	580	*0%*	*0%*	TPU v7 (Ironwood) and v8i most of the volume in 2026 and 2027, v8 family likely to extend to 2028 as well. MTIA CoWoS ramp seems slower, OpenAI in 2027. Tomahawk 6 also using CoWoS-S.
CoWoS-S at TSMC	40	70	91	220	260	130	*0%*	*0%*	CoWoS-S is still mainstream in 2026, but more adoption of CoWoS-L from 2027 for Meta, OpenAI and future versions of TPU starting late 2027.
CoWoS-S outside TSMC	0	0	5	10	70	130	*0%*	*0%*	Outsourcing starting to pick up - trialling ASE and Amkor for Networking projects in 2027. OSAT qualification slower than expected.
CoWoS-L	0	0	0	20	95	320	*0%*	*0%*	Starting to use CoWoS-L for some projects including advanced TPU models; also likely to incorporate SoIC for some advanced ASICs for TPUv10, MTIA 450/500 and OpenAI Titan2 in 2028.
Alchip + AWS	7	14	5	65	175	250	*0%*	*6%*	Trn2 upside in 1H26, Trn3 ramp in mid 2026, upside to lifetime volumes for Trn3. Most Trn3 CoWoS allocation through AWS or Alchip. Small overflow to other vendors like Marvell due to excess demand.
Marvell (incl AWS for Trn2)	2	13	75	35	40	80	*0%*	*33%*	Near-term upside to Trn2. Minor Trn3 CoWoS or OSAT allocation. MSFT MAIA 3 to ramp in 2027. Not seeing a big ramp for other CoWoS projects in 2027. In 2028, expect some allocation for Trn4 (for NVLink fusion versions).
MediaTek	0	0	0	30	150	75	*0%*	*20%*	2026 TPUv8t volume strengthening in 2026/27. TPUv9 Humufish now firmly on Intel EMIB-T - hence lower 2028 allocation.
Others	5	16	24	40	40	60	*14%*	*-20%*	Includes Cisco, Microsoft, GUC, Socionext, other ASIC design houses and smaller startup customers etc.; some non-AI chips also using CoWoS (Networking chips for various vendors, some small volume server CPUs etc.).
Overall Total CoWoS Consumption	**134**	**350**	**652**	**1,240**	**2,381**	**3,320**	***2%***	***14%***	

Source: Company data, J.P. Morgan estimates. Note: Capacity allocation is before adjusting for yield losses, which may vary by accelerator and package type.

Table 2: WMCM forecasts

							% Revision vs prior ests		
WMCM Consumption (K wafer per year)	**2023**	**2024**	**2025**	**2026E**	**2027E**	**2028E**	**2026E**	**2027E**	**JPM Comments**
Apple	0	0	0	417	720	1,044	*0%*	*0%*	All high-end iPhone 18 models in 2H26E using WMCM.
Number of iPhones (M units)				88	151	219	*0%*	*0%*	Likely to adopt for A20 Pro based models on N2. Encompassing iPhone foldable, Pro and Pro max.
Utilization				100%	100%	100%	*0%*	*0%*	
Gross Die per wafer				210	210	210	*0%*	*0%*	Expect Package size to be 50% larger than InFO with ~200 dies per Wafer due to side-by-side stacking of LPDDR5.
% Yield				98%	98%	98%	*0%*	*0%*	
TSMC WMCM Capacity (k wafer per year)				**426**	**735**	**1,065**	***0%***	***0%***	

Source: Company data, J.P. Morgan estimates.

Gokul Hariharan [AC]
(852) 2800-8564
gokul.hariharan@jpmorgan.com

Table 3: GPU and AI accelerators chip-level shipment forecasts

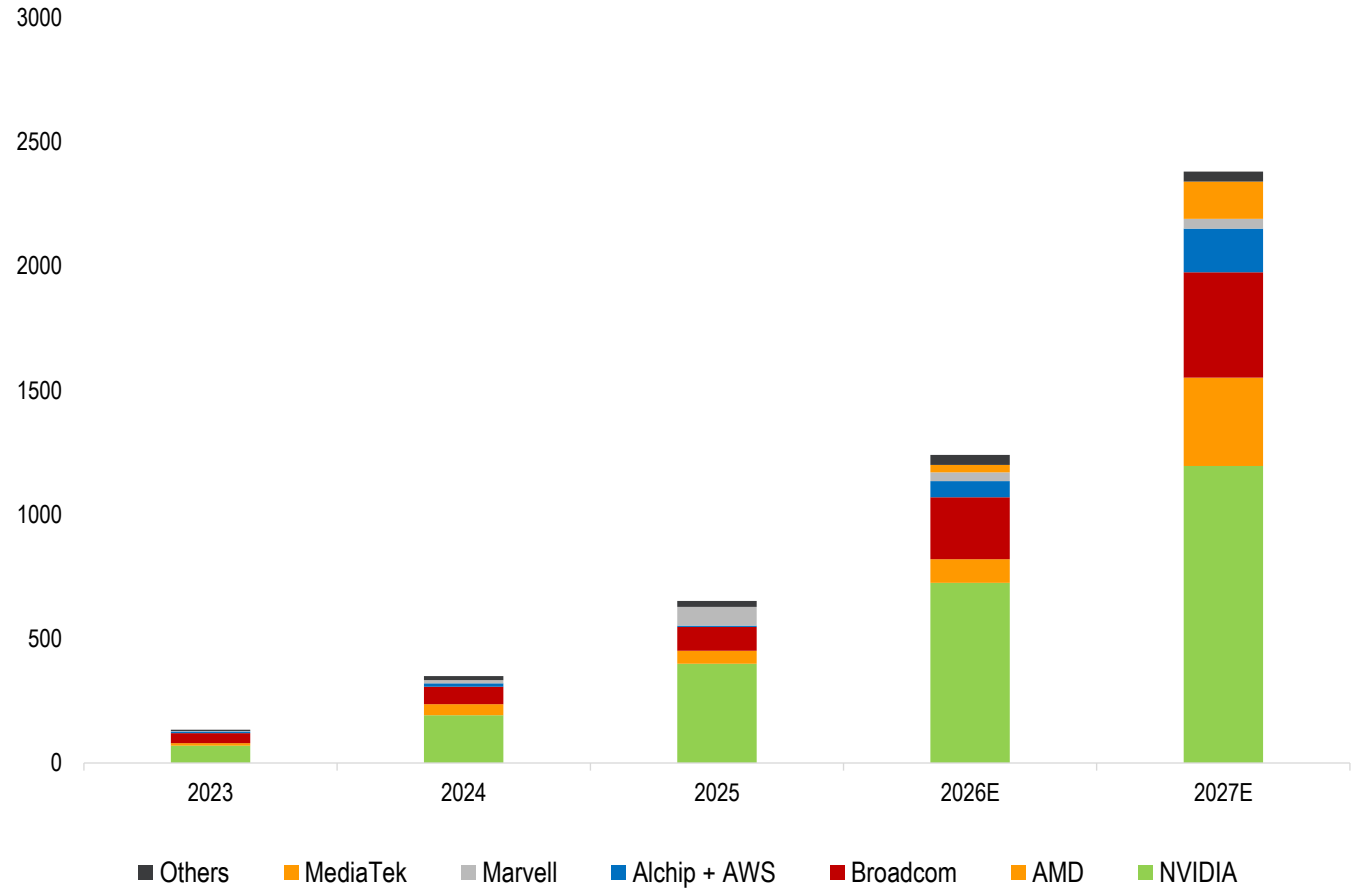
GPU and Accelerator Units (M)	2023	2024	2025	2026E	2027E	2028E
NVDA						
Hopper Series	1.7	4.3	1.0	0.4	-	-
Blackwell Series	-	0.2	5.2	6.4	1.0	-
Rubin Series	-	-	-	2.1	7.3	1.5
Rubin Ultra Series (2 die equivalent)	-	-	-	-	1.7	8.5
Feynman Series	-	-	-	-	-	1.6
Total NVDA GPU Units	**1.7**	**4.5**	**6.2**	**8.9**	**10.0**	**11.6**
NVDA eq GPU die units	*1.7*	*4.7*	*11.4*	*17.4*	*20.0*	*26.6*
Vera CPU	-	-	-	0.6	5.0	3.0
Rosa CPU	-	-	-	-	-	4.0
Others Including Spectrum X	-	-	-	0.2	0.4	0.9
LPX (LP40)						1.0
AMD						
MI300/325	0.1	0.5	0.4	-	-	-
MI350	-	-	0.3	0.2	-	-
MI450	-	-	-	0.5	1.5	2.0
MI500				-	-	0.3
Total AMD GPU Units	**0.1**	**0.5**	**0.7**	**0.7**	**1.5**	**2.3**
Venice CPU			-	0.4	3.2	4.5
Google TPUs						
v4	0.7	0.4	-	-	-	-
v5	0.5	1.3	-	-	-	-
v6 /v7 (Trillium/Ironwood)	-	-	1.7	2.7	1.7	-
v8t - Zebrafish - Mediatek	-	-	-	0.5	3.0	1.5
v8i - Sunfish /Sunfish Ultra (in 2 die equivalent)	-	-	0.1	1.2	4.2	7.1
v9 - Humufish - Mediatek				-	-	3.0
Total TPU Units	**1.2**	**1.7**	**1.8**	**4.4**	**8.9**	**11.6**
AWS						
Inferentia 2	0.3	0.5	0.1	-	-	-
Trainium 2	-	0.3	1.5	0.6	-	-
Trainium 3	-	-	-	1.3	3.5	-
Trainium 4				-	-	3.5
Total Inferentia / Trainium Units	**0.3**	**0.8**	**1.5**	**1.9**	**3.5**	**3.5**
Meta						
MTIA 300 and beyond				0.3	0.5	1.1
Total MTIA units (with CoWoS)				**0.3**	**0.5**	**1.1**
Microsoft						
MAIA 200 - GUC				0.1	0.1	0.0
MAIA 300 - Marvell				0.1	0.7	0.8
Total MAIA units				**0.2**	**0.8**	**0.8**

Source: Company data, J.P. Morgan estimates.

Gokul Hariharan AC
(852) 2800-8564
gokul.hariharan@jpmorgan.com

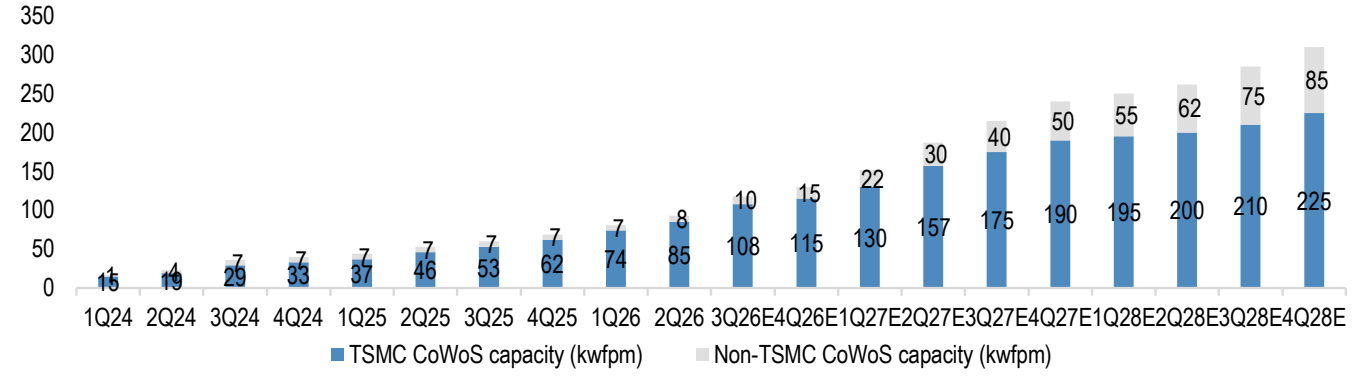
Figure 1: Customer mix for CoWoS

k wafers per year



Source: Company data, J.P. Morgan estimates.

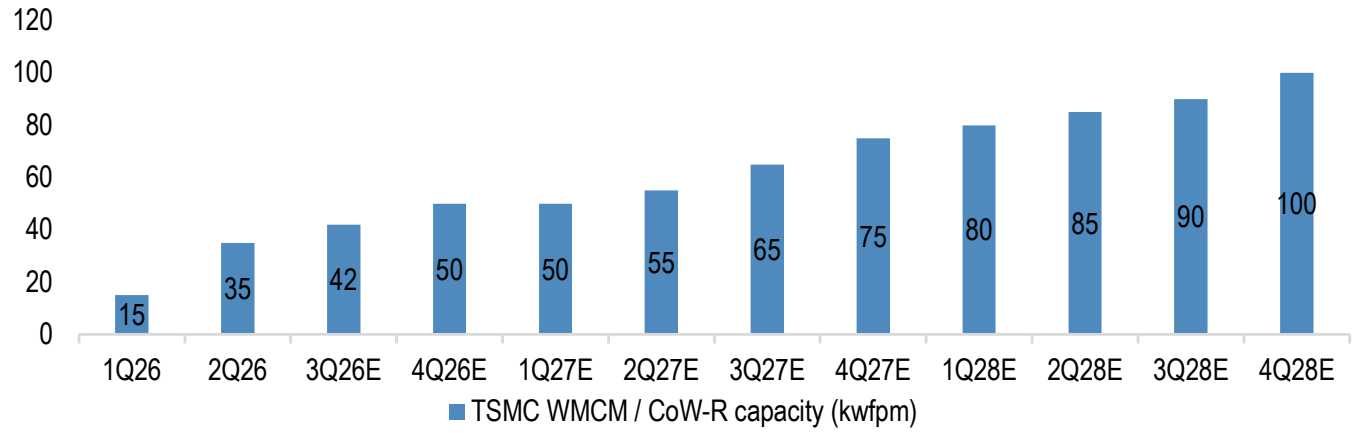
Figure 2: CoWoS capacity – TSMC/non-TSMC



Source: Company data, J.P. Morgan estimates.

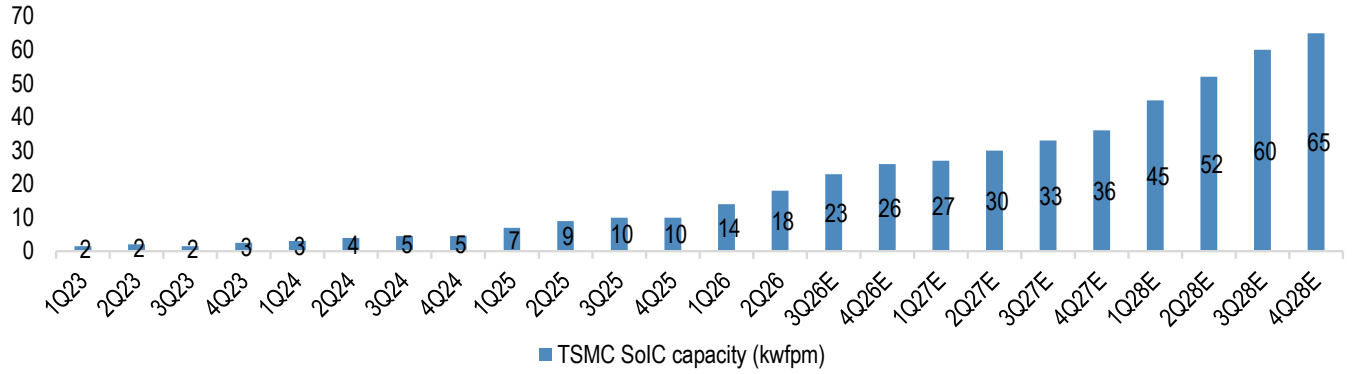
Gokul Hariharan AC
(852) 2800-8564
gokul.hariharan@jpmorgan.com

Figure 3: WMCM/CoW-R capacity – TSMC



Source: J.P. Morgan estimates.

Figure 4: SoIC capacity – TSMC



Source: Company data, J.P. Morgan estimates.

Gokul Hariharan AC
(852) 2800-8564
gokul.hariharan@jpmorgan.com

J.P.Morgan

Table 4: Glossary of important terms

HBM	High Bandwidth Memory	HBM is a high-speed computer memory interface for 3D-stacked synchronous dynamic random-access memory (SDRAM) from Samsung, Micron and SK Hynix.
CoWoS	Chips-on-Wafer-on-Substrate	TSMC's 2.5D interposer-based packaging technology for high-performance computing applications.
CoWoS - S	Chips-on-Wafer-on-Substrate with Silicon Interposer	A CoWoS technique using silicon interposer as the bridge between chip and substrate.
CoWoS - R	Chips-on-Wafer-on-Substrate with Redistribution Layer Interposer	A CoWoS technique using redistribution layer interposer (RDL) to replace silicon interposer, enabling higher packaging size flexibility and lower cost.
CoWoS - L	Chips-on-Wafer-on-Substrate with Local Silicon Interconnect	A CoWoS technique using local silicon interconnect (LSI) to connect dies at high speed rate, while the other area on the chip will use RDL.
COW - R	Chip-on-Wafer with Redistribution Layer Interposer	The first part of the CoWoS process, where the logic chip and HBMs are attached to the RDL interposer.
RDL	Redistribution Layer	A layer of metal interconnects that redistributes I/O access across the chip, facilitating flexible and efficient chip-to-chip connections, particularly in advanced packaging techniques like 2.5D and 3D integration.
WMCM	Wafer-level Multichip Module	An approach that integrates multiple chips (CPU, GPU, memory, etc.) at the wafer stage before slicing them into individual modules. Unlike traditional InFO (Integrated Fan-Out) packaging, focusing on single-die integration, WMCM enables more flexibility in combining different components while maintaining a small footprint.
XPU	X processing unit	The "X" in "XPU" stands for any compute architecture (CPU, GPU, ASIC etc.) that is suitable for different applications.
TPU	Tensor processing unit	Google Cloud TPUs are cusomized AI accelerators, optimized for training and inference of AI models.
LPDDR	Low Power Double Data Rate	Also known as LPDRAM, LPDDR is used in devices requiring speedy, high-bandwidth data rates without compromising power efficiency.
CPO	Co-Packaged Optics	An integration of optical componnents and ASICs on a single packaged substrate, enabling high-speed data transmission and less power loss.
ASIC	Application-Specific Integrated Circuit	ASICs (application-specific integrated circuits) are customized chips for certain special purposes, and right-of-use is limited to the chip designer only.
CSP	Cloud Service Provider	A third-party company that provides a cloud-based platform for businesses to access on demand over a network, including compute, storage, and infrastructure services.
FOCoS	Fan Out Chip on Substrate (similar to TSMC's CoWoS)	A fan-out package flip-chip placed on a high pin count ball grid array (BGA) substrate. The fan-out package has an RDL allowing the construction of shorter die-to-die (D2D) interconnections between chips.
Trn	Trainium	Trainium are chips developed by AWS specifically for AI training and inference, which deliver high performance while saving costs.
InFO	Integrated Fan-Out	A wafer-level integration technology features high-density RDL and Through InFO Via (TIV) for enhanced interconnect performance. It supports various 2D and 3D package schemes optimized for mobile and high-performance computing applications.
EIC-PIC	Electronic/Photonic Integrated Circuit	EICs utilize electronic components (transistors, capacitors) for electrical signal processing, while PICs use optical components (waveguides, lasers) for light manipulation in data transmission. EICs handle electrons, and PICs handle photons.
CoPoS	Chip-on-Panel-on-Substrate	TSMC's CoPoS is a rectangular substrate evolution of CoWoS, offering more usable space than traditional round wafers, enhancing output efficiency and reducing costs.
EMIB	Embedded Multi-die Interconnect Bridge	Intel's EMIB is an advanced packaging technology that embeds silicon bridges into an organic substrate. It enables multiple chiplets to communicate side-by-side with high bandwidth and low latency, without the need for full-size silicon interposers.

Source: J.P. Morgan.

Gokul Hariharan [AC]
(852) 2800-8564
gokul.hariharan@jpmorgan.com

J.P.Morgan

Investment Thesis, Valuation and Risks

TSMC *(Overweight; Price Target: NT$3100.0)*

Investment Thesis
We expect TSMC's structural growth drivers to remain very strong as leading-edge supply (N4, N3, and N2) stays tight well into 2027 or early 2028, given accelerating AI demand, continued growth in N3, N2 and advanced back-end revenues, and long lead time for capacity build despite accelerated capex. Near-term, we believe GMs are also likely to remain strong, nudging 70% levels, given N3 improvement, high UTR and decent pricing, despite N2 and overseas dilution. We expect broader price hikes in 2027, given the widening supply-demand gap and a stronger Datacenter AI CAGR (we now estimate 69% CAGR from 2024-29E, driven by both higher accelerator estimates and a surge in AI CPU numbers).

Valuation
Our Jun-27 PT of NT$3,100 is based on ~20x 12-month forward P/E and reflects stronger near-term profitability, better visibility on AI demand, stronger capacity build and firmer pricing into 2028. Our target multiple is higher than TSMC's five-year average historical multiple.

Risks to Rating and Price Target
Key downside risks to our rating and price target include: (1) the debate on the duration of the AI capex growth cycle and (2) any negative impact from weak PC/Smartphone demand in 2H26.

Gokul Hariharan [AC]
(852) 2800-8564
gokul.hariharan@jpmorgan.com

J.P.Morgan

TSMC: Summary of Financials

Income Statement - Annual	FY24A	FY25A	FY26E	FY27E
Revenue	**2,894,308**	**3,809,054**	**5,342,260**	**7,192,208**
COGS	(1,269,954)	(1,527,760)	(1,725,452)	(2,307,962)
Gross profit	**1,624,354**	**2,281,294**	**3,616,807**	**4,884,246**
SG&A	(98,119)	(98,775)	(121,360)	(171,005)
Adj. EBITDA	**1,984,850**	**2,624,188**	**3,968,123**	**5,280,152**
D&A	(662,797)	(688,096)	(792,922)	(1,004,290)
Adj. EBIT	**1,322,053**	**1,936,092**	**3,175,201**	**4,275,861**
Net Interest	76,718	93,369	95,785	78,819
Adj. PBT	**1,405,839**	**2,041,663**	**3,276,816**	**4,358,963**
Tax	(233,407)	(326,266)	(579,180)	(772,802)
Minority Interest	836	2,486	(1,286)	(1,286)
Adj. Net Income	**1,173,268**	**1,717,883**	**2,696,350**	**3,584,875**
Reported EPS	45.24	66.24	103.97	138.24
Adj. EPS	**45.24**	**66.24**	**103.97**	**138.24**
DPS	**16.50**	**22.00**	**29.00**	**34.50**
Payout ratio	51.0%	48.6%	43.8%	33.2%
Shares outstanding	25,933	25,933	25,933	25,933

Balance Sheet & Cash Flow Statement	FY24A	FY25A	FY26E	FY27E
Cash and cash equivalents	2,422,020	3,068,595	4,392,302	5,423,224
Accounts receivable	272,088	282,059	419,840	533,954
Inventories	287,869	288,109	419,840	533,954
Other current assets	106,376	178,367	136,215	173,239
Current assets	**3,088,352**	**3,817,131**	**5,368,196**	**6,664,371**
PP&E	3,234,980	3,691,841	4,451,630	5,919,590
LT investments	149,040	172,370	169,709	173,992
Other non current assets	219,565	251,682	274,192	274,192
Total assets	**6,691,938**	**7,933,024**	**10,263,728**	**13,032,145**
Short term borrowings	59,858	136,926	156,242	156,242
Payables	74,227	84,330	109,158	123,877
Other short term liabilities	1,130,441	1,236,763	1,362,147	1,443,654
Current liabilities	**1,264,525**	**1,458,019**	**1,627,548**	**1,723,773**
Long-term debt	958,429	896,062	901,018	901,018
Other long term liabilities	145,408	118,147	97,132	79,114
Total liabilities	**2,368,362**	**2,472,229**	**2,625,697**	**2,703,906**
Shareholders' equity	4,288,545	5,419,596	7,596,602	10,286,810
Minority interests	35,031	41,199	41,429	41,429
Total liabilities & equity	**6,691,938**	**7,933,024**	**10,263,728**	**13,032,145**
BVPS	**165.4**	**209.0**	**292.9**	**396.7**
y/y Growth	*24.0%*	*26.4%*	*40.2%*	*35.4%*
Net debt/(cash)	(1,403,733)	(2,035,607)	(3,335,042)	(4,365,964)
Cash flow from operating activities	**1,975,661**	**2,437,716**	**3,413,412**	**4,421,425**
o/w Depreciation & amortization	662,797	688,096	792,922	1,004,290
o/w Changes in working capital	(90,016)	(108)	(244,683)	(213,510)
Cash flow from investing activities	**(928,044)**	**(1,200,404)**	**(1,572,560)**	**(2,476,533)**
o/w Capital expenditure	(958,198)	(1,272,219)	(1,842,153)	(2,472,250)
as % of sales	*33.1%*	*33.4%*	*34.5%*	*34.4%*
Cash flow from financing activities	**(313,242)**	**(590,737)**	**(517,144)**	**(913,970)**
o/w Dividends paid	(428,731)	(573,004)	(750,753)	(893,381)
o/w Net debt issued/(repaid)	90,711	14,701	24,272	0
Net change in cash	**734,375**	**646,575**	**1,323,707**	**1,030,922**
Adj. Free cash flow to firm	**1,142,360**	**1,292,759**	**1,860,700**	**1,949,175**
y/y Growth	*164.7%*	*13.2%*	*43.9%*	*4.8%*

Income Statement - Quarterly	1Q26A	2Q26E	3Q26E	4Q26E
Revenue	**1,134,103**	**1,271,833**	**1,403,908**	**1,532,416**
COGS	(382,808)	(387,611)	(454,944)	(500,089)
Gross profit	**751,295**	**884,222**	**948,964**	**1,032,326**
SG&A	(24,573)	(29,252)	(32,290)	(35,246)
Adj. EBITDA	**824,417**	**957,295**	**1,054,322**	**1,132,089**
D&A	(165,450)	(178,636)	(221,883)	(226,953)
Adj. EBIT	**658,966**	**778,659**	**832,439**	**905,136**
Net Interest	26,145	21,941	22,799	24,900
Adj. PBT	**687,800**	**801,641**	**856,286**	**931,090**
Tax	(114,998)	(160,328)	(145,569)	(158,285)
Minority Interest	(322)	(322)	(322)	(322)
Adj. Net Income	**572,480**	**640,992**	**710,396**	**772,483**
Reported EPS	22.08	24.72	27.39	29.79
Adj. EPS	**22.08**	**24.72**	**27.39**	**29.79**
Shares outstanding	25,932	25,932	25,932	25,932

Ratio Analysis	FY24A	FY25A	FY26E	FY27E
Gross margin	56.1%	59.9%	67.7%	67.9%
EBITDA margin	68.6%	68.9%	74.3%	73.4%
EBIT margin	45.7%	50.8%	59.4%	59.5%
Net profit margin	40.5%	45.1%	50.5%	49.8%
ROE	30.3%	35.4%	41.4%	40.1%
ROA	19.2%	23.5%	29.6%	30.8%
ROCE	22.7%	27.7%	34.6%	35.2%
SG&A/Sales	3.4%	2.6%	2.3%	2.4%
Net debt/equity	NM	NM	NM	NM
P/E (x)	53.4	36.5	23.2	17.5
P/BV (x)	14.6	11.6	8.2	6.1
EV/EBITDA (x)	31.4	23.5	15.2	11.3
Dividend Yield	0.7%	0.9%	1.2%	1.4%
Sales/Assets (x)	47.4%	52.1%	58.7%	61.7%
Interest cover (x)	NM	NM	NM	NM
Operating leverage	128.3%	147.0%	159.0%	100.1%
Revenue y/y Growth	33.9%	31.6%	40.3%	34.6%
EBITDA y/y Growth	36.5%	32.2%	51.2%	33.1%
Tax rate	16.6%	16.0%	17.7%	17.7%
Adj. Net Income y/y Growth	39.9%	46.4%	57.0%	33.0%
EPS y/y Growth	39.9%	46.4%	57.0%	33.0%
DPS y/y Growth	26.9%	33.3%	31.8%	19.0%

Source: Company reports and J.P. Morgan estimates.
Note: NT$ in millions (except per-share data).Fiscal year ends Dec. o/w - out of which

