Asia Tech Strategy

At a crossroads in the AI Infra trade

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

We are at a crossroads in the AI Infra trade where three questions need decisive answers for the next leg of the rally. (1) Selling Compute is a great business, but is Buying Compute (as an AI model or vertical AI vendor) a great business on a sustainable basis? (2) Do the recent AI safety and regulatory concerns impede the growth of Compute demand?, and (3) What is the sustainable financing model for the next leg of capex growth, now that hyperscalers are moving to negative FCF. We believe that current AI economics indicate that buying compute is still a very healthy business, despite the concerns of open source models, and with AI model advancements accelerating, we believe this could remain sustainable. While AI safety concerns are a near-term hiccup, we believe that frontier AI model training is unlikely to slow down, and the fast pace of AI evolution with RSI etc should create new revenue TAMs beyond coding. Lastly, funding still remains a key area of market worry, but we believe there's sufficient room for strong capex growth in the next two years between accelerating operating cash flow growth, hyperscaler debt and equity issuance. The tech hardware ecosystem continues to remain healthy with steady EPS revisions and compelling valuations. While current sentiment remains poor, we do not see any imminent signals of a capex downturn that is starting to get priced into Tech hardware stocks. Key catalysts would still come from more signs of sustained monetization at the model and AI application layers. Our top picks in Asian Tech are TSMC, ASE, Mediatek, Wiwynn, Aspeed, ASMPT, Tokyo Electron, Advantest, SEC, Chroma and Unimicron.

- **Selling Compute a great business, Buying Compute appears equally good, at least based on current economics:** By now, it has been well established that selling AI Compute is a great business, whether for the semi/hardware ecosystem (record high margins) or for the public cloud/neocloud ecosystem ($ per GPU hour running significantly higher now compared to last year (GPU spot prices) with strong acceleration in public cloud growth (Figure 1) and good IRR on neocloud deals (see detailed research from JPM US analyst, Samik Chatterjee). The concerns have centered around whether Buying Compute (and selling AI tokens or Vertical AI applications) is a good business or not, and whether that can be sustainable due to intense frontier model competition. Based on current economics, we believe that selling tokens remains an extremely attractive business, especially for vendors at the model frontier, with several reports of companies reaching 60-80% inference gross margins. Our back-of-envelope calculation indicates that a hypothetical model vendor could make $20-40bn in annual revenue per GW selling AI tokens, much higher than the $10bn per GW per year that existed in 2025 (based on OpenAI CFO Sarah Friar's comments). We believe that the surge of vertical AI companies that have reportedly reached $100M+ ARR in a very short span of time (80+companies already based on Sapphire Venture report) is also a testament to the value creation downstream of selling Compute.
- **Open Source models should drive increased demand for hardware, as token costs resume their historic cost decline trend:** An added concern on the demand for compute in the last few months has been the rise of Open Source LLMs and the impact on overall token pricing (see Figure 3). We believe that open source models are likely to continue to be important in driving the cost of tokens lower (historical cost curve has been a 80-90% decline in cost per

See page 9 for analyst certification and important disclosures, including non-US analyst disclosures.

token, with 1H26 being an anomaly driven by severe compute shortage and the rise of agentic AI), hence driving higher adoption across various industries. While the market is concerned about monetization at the model layer due to the rise of Open Source, we believe that: (1) In the current compute-constrained environment, both proprietary Frontier models and Open Source models are likely to see strong token consumption growth; (2) Rise of open source is likely to lead to higher AI adoption down the stack; and (3) Increased usage of Open Source should lead to more revenue-generating vertical AI application layers being created down the food-chain, using cheaper Open Source AI tokens to solve domain-specific problems (similar to companies like Fireworks AI etc).

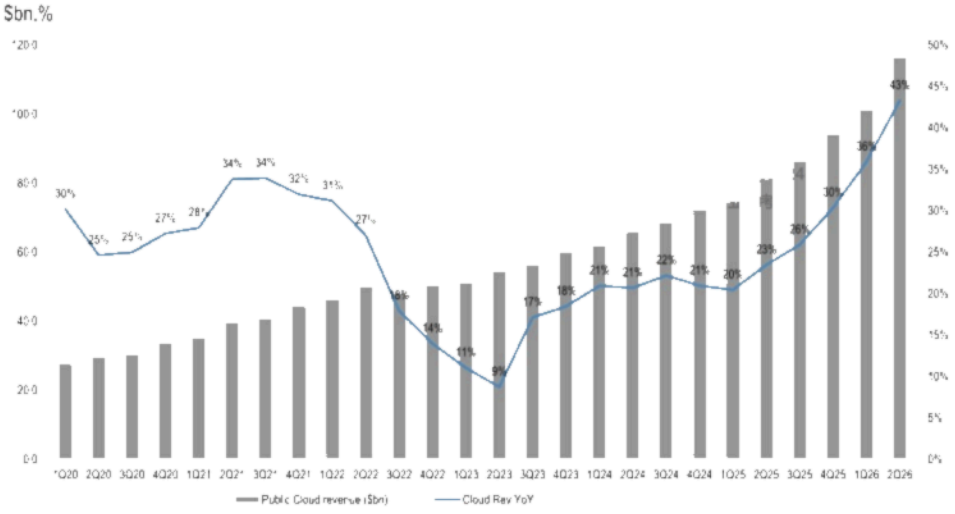
- **Back to first principles - AI safety and regulation triggered by fast AI model advancement rather than a slowdown, promise of stronger future TAM growth:** Over the weekend, the market has gotten increasingly concerned about AI compute demand after reports (see report) citing calls to slowdown the pace of AI deployment with a safety first and regulated approach. Our first-principles-based view is: (1) AI models are evolving faster with impending RSI (recursive self-improvement) which underlines the fact that the pace of technology evolution and AI scaling laws are not slowing – this gives us the confidence that Gen AI could potentially open up new addressable markets in the next 1-2 years, similar to what has happened with coding and agentic workflows in the last 18 months. We do not believe that the pace of AI model training is likely to slow, even if the roll-out of new models could be gated by new alignment restrictions etc. We are also unlikely to see the Open Source models slowing down their advancements, which is likely to put further competitive pressure on frontier model labs to keep accelerating their training efforts to stay ahead.
- **Hyperscaler capex funding: Leverage rising, but feels there is more room to go, OCF growth could accelerate with strong public cloud growth:** Another area of concern for investors has been the room for hyperscaler capex investment growth, given most hyperscalers are now entering negative FCF territory and have to rely on external funding to drive the next leg of capex. We believe that there is still more room to grow capex in the next few years given: (1) Leverage ratios for large hyperscalers are still quite low at 13% net debt to equity; (2) public cloud growth, which could be a key source of operating cash flow growth is accelerating with the strong demand for compute and at higher price points; and (3) equity fund raising from hyperscalers as well as downstream buyers of compute, which could provide an additional layer of funding in the next few years. Still this is a key area to watch, given the environment of rising rates and increased risk perception among some segments of AI infra funding.
- **Among Tech supply chains, SPE looks quite attractive as capex spending accelerates:** Looking into 2027, we believe most of the large semiconductor vendors are likely to highlight equipment as the key bottleneck (so far has been clean room and other areas) as they strive to bridge the demand-supply gap. We expect TSMC to step up investments in N2, N3 and A14 significantly in the next few years, while challengers in Foundry such as Intel, Samsung Foundry and Terafab are also likely to increase capex. Trailing edge Foundries are also starting a new investment cycle for the first time in four years. On the memory front, both DRAM and NAND Flash capex from the top memory vendors should grow rapidly in 2027/28 as clean-room capacity comes online, process migration accelerates and companies build out to cater to their LTA agreements with customers. China memory investment is also likely to accelerate materially in the next two years. SPE vendors are also starting to enjoy some price increases (a rare phenomenon) given tight capacity, increased input costs and expedited tool orders from multiple vendors. Our top picks are Advantest, Tokyo Electron, Chroma, GPTC, ASMPT.
- **Substrates and packaging remain key bottlenecks, with upside from CPUs and EMIB-T:** At the component level, we believe IC substrate and advanced packaging remain the key bottlenecks. Substrate supply is gated by highly concentrated capacity

among tier-1 players (i.e. Unimicron, Ibiden) and long capacity build lead times (up to ~2.5 years), while leading OSATs are additionally constrained by available clean-room space. We expect the supply-demand gap to widen further as strong AI accelerator growth drives higher substrate content per package and greater advanced-packaging intensity, with incremental upside from robust server CPU demand on the back of Agentic AI and rising adoption of substrate-heavy EMIB-T packaging from late 2027. Within our Asia Tech coverage, we therefore prefer tier-1 substrate vendors such as Unimicron, Ibiden and advanced-packaging leaders such as ASE.

- **Optics a key element of solving the interconnect challenge; NPO favored near-term:** As AI infra providers and model labs put greater emphasis on improving AI cluster efficiency, we expect interconnect to be the key bottleneck. The data transfer between GPUs, GPU-to-CPU, and memory-to-storage is increasingly gating the performance. We believe this challenge drives higher adoption of optics, favoring NPO (Near-Packaged Optics) in the near term, given it is a simpler solution. We believe that the two large ASIC vendors, Google and AWS, as well as NVDA will all embrace NPO-based solutions in the next 2 years. CPO will be the long-term future, but could take more time for the supply chain to mature, especially for scale-up solutions.
- **Given ongoing despeccing and algorithmic efficiency concerns, climbing the wall of worry in memory requires more faith:** From a top-down perspective, the memory sub-sector remains quite noisy due to the high degree of investor focus on HBM despeccing, increased efforts to drive higher algorithmic efficiency (with the advent of looped transformers etc.) to reduce memory consumption per function, and offloading more of the KV cache storage to lower tiers of memory (such as DDR or NAND). Industry fundamentals continue to remain healthy with demand supply balance not likely before 2028 and pricing should continue to rise into 2027, which should drive some stock price appreciation if AI sentiment comes back in 4Q26. Still, investor sponsorship may remain lukewarm in the near-term, based on our feedback, compared to other sub-sectors in tech. This sentiment may change if the next iteration of AI model evolution drives a significant increase in memory demand per function, similar to what happened with the original reasoning breakthrough in late 2024.
- **Where could price increases be stronger in 2027 vs 2026?** The sub-sectors where we see stronger price increases (in % terms) in 2027 vs. 2026 are Foundries (especially 8" mature Foundries), OSAT, substrate, SPE and high-end CCL, while memory is one area where the pace of price increase is likely to slow down, along with PCB materials. For Foundries, we see TSMC hiking prices across all process nodes in 2027 (compared to only leading edge in 2026), while mature foundry price increases could be higher (10-15% for 8" vs. 8-10% in 2026). OSAT price increases are accelerating and spreading to wirebonding and mature flip chip packaging given the lack of clean room availability due to the AI crowding out effect. Substrate price increases are starting, and the introduction of EMIB-T in late 2027 is likely to create a tighter supply-demand environment for substrates compared to 2025-26, as advanced substrates tend to be very capacity consuming (on top of larger area size and higher layer count). High-end CCL (M7 and above) might also see stronger price hikes as demand (both AI and general server spec migration) growth continues to outpace supply growth.

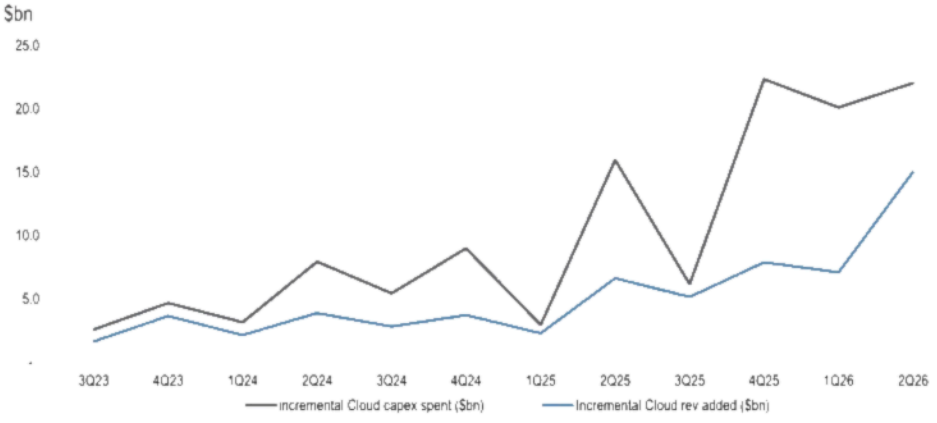
Gokul Hariharan AC
(852) 2800-8564
gokul.hariharan@jpmorgan.com

Figure 1: Public Cloud Quarterly Revenue Growth (YoY)



Companies included are Google, Microsoft, Amazon, and Oracle
Source: Bloomberg Finance L.P.

Figure 2: Incremental Public Cloud Revenue Addition (QoQ) vs. Capex Growth (QoQ)



Companies included are Google, Microsoft, Amazon, and Oracle.
Note: Capex may not be all for Public Cloud since the CSPs have internal workloads which may not be captured in Cloud revenues
Source: Bloomberg Finance L.P.

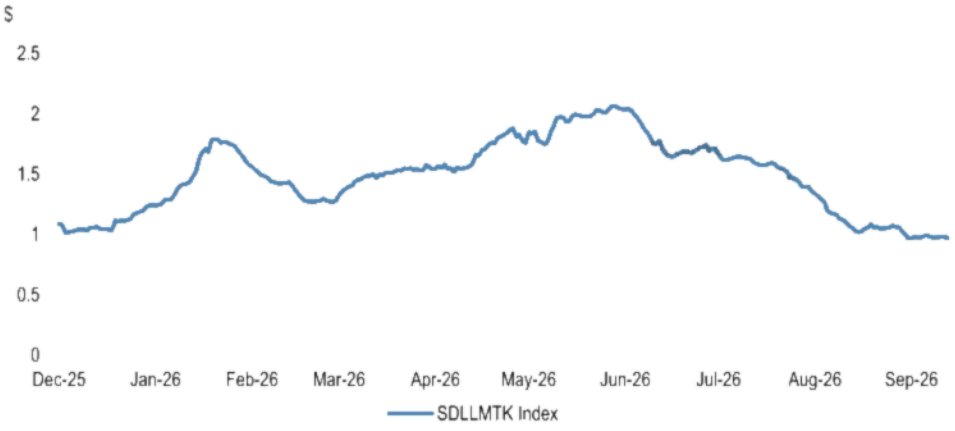
Table 1: Public Cloud Revenue Backlog

US$bn	3Q25	4Q25	1Q26	2Q26
Google (Cloud Backlog)	155	240	460+	514
Amazon (AWS Backlog)	200	244	364	496
MSFT (Commercial remaining performance obligation)	392	625	627	678

Source: Company reports.

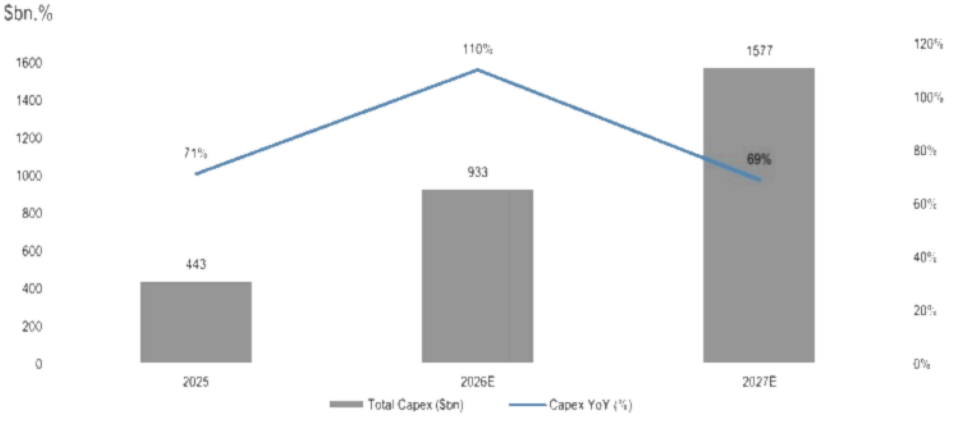
Gokul Hariharan AC
(852) 2800-8564
gokul.hariharan@jpmorgan.com

Figure 3: Price per million tokens - Silicon Data Index



Source: Bloomberg Finance L.P.

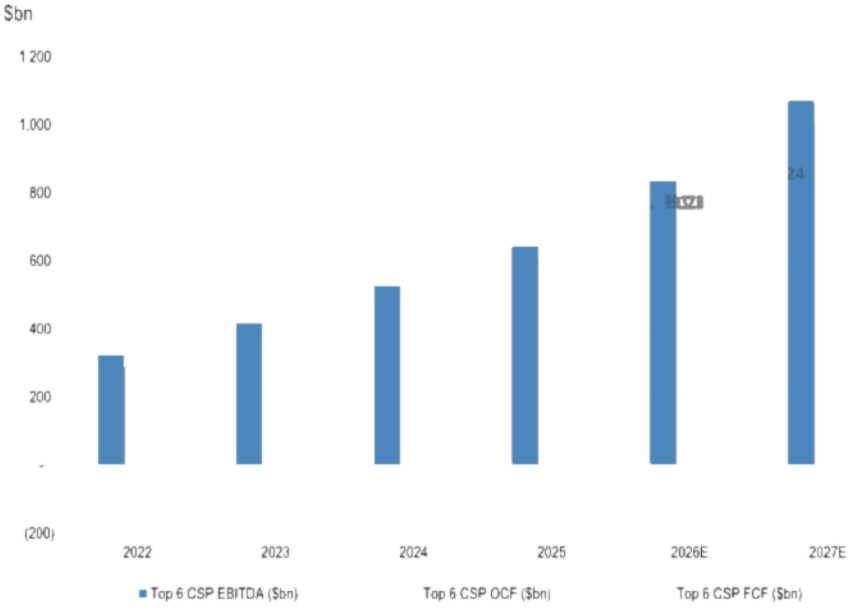
Figure 4: Top 7 Hyperscaler capex forecast into 2027 - maintaining a heady pace of growth



Companies included are Amazon, MSFT, Google, Meta, Oracle, Coreweave, and SpaceX. Data is adjusted to CY.
Source: Bloomberg Finance L.P.J.P. Morgan estimates.

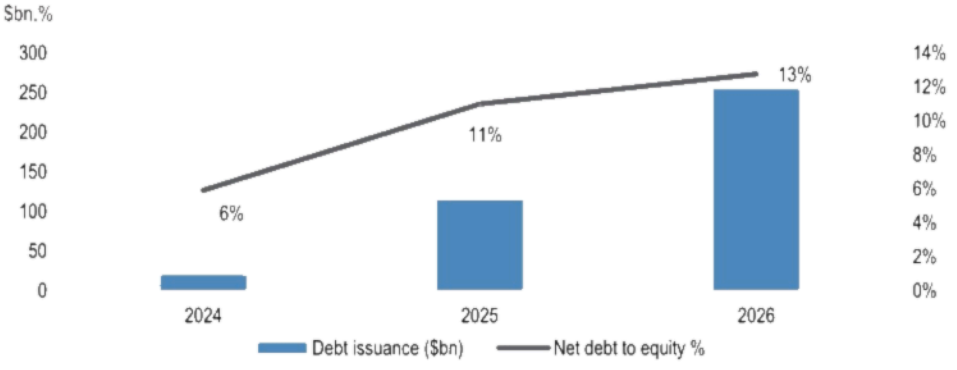
人在草木间

Figure 5: Operating cash flow driven funding reaching limits



Source: J.P. Morgan estimates, Bloomberg Finance L.P.

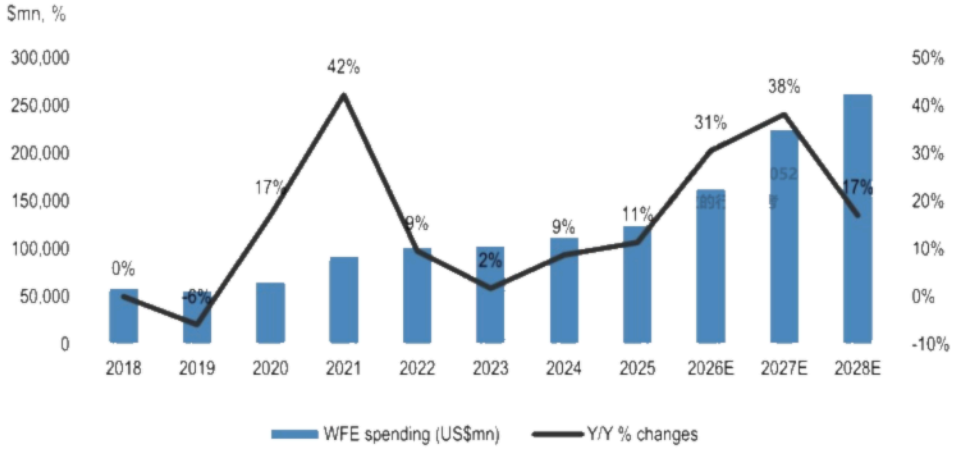
Figure 6: Hyperscaler debt issuance and net debt to equity



Note: Top hyperscalers include: Google, Meta, Microsoft, Amazon, ORCL, Coreweave, SpaceX
Source: Bloomberg Finance L.P.

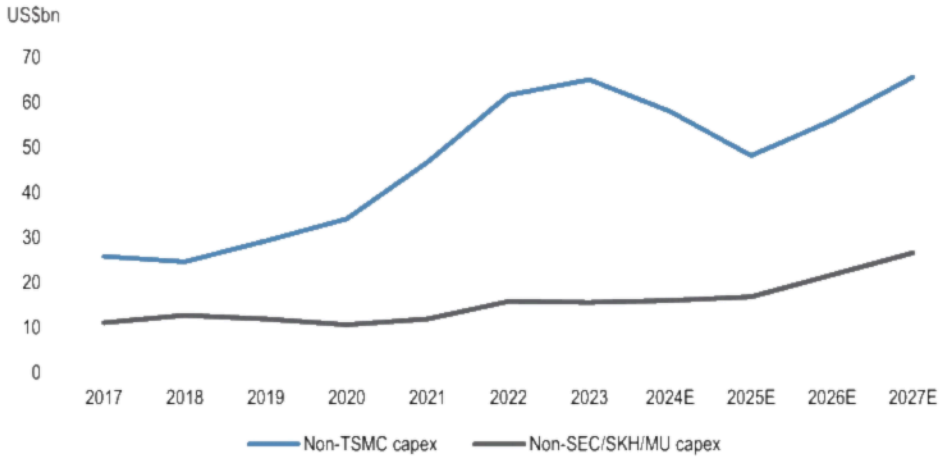
Gokul Hariharan AC
(852) 2800-8564
gokul.hariharan@jpmorgan.com

Figure 7: WFE Spending



Source: Company data, Gartner® and J.P. Morgan estimates.

Figure 8: Semi capex increase from non-TSMC and non-SEC/SKH/MU



Source: Company data, Gartner® and J.P. Morgan estimates

Figure 9: Asia tech EPS revisions and revision breadth

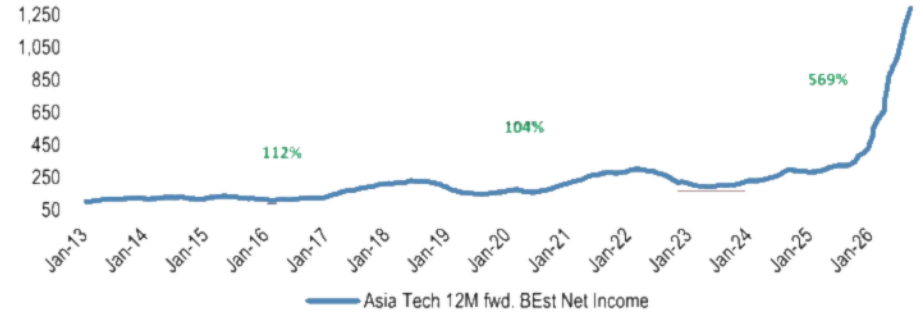


Source: J.P. Morgan estimates, Bloomberg Finance L.P.

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

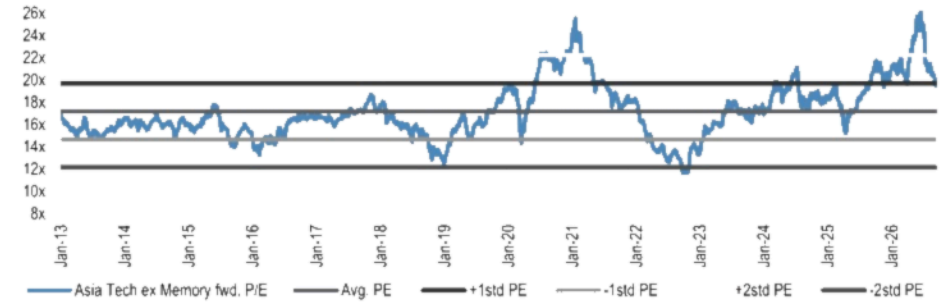
Figure 10: Asia tech 12M fwd BEst Net Income

Rebased from Jan-11 as 100



Source: Bloomberg Finance L.P.

Figure 11: Asia tech ex-memory (fwd) P/E



Source: Bloomberg Finance L.P.

Figure 12: Asia Tech (fwd) P/B



Source: Bloomberg Finance L.P