KASASV

KOREAN AMERICAN SEMICONDUCTOR ASSOCIATION IN SILICON VALLEY

April 2025

짧지만 생각을 하게하는 글

1. "내가 경애하는 자가 어떤 자인 줄 아느냐? **내 예상을 벗어난 자다.** 매순간 변화하고 성장하는자"

2. "자신이 얼마나 대단한 사람인지 보여주려고 할때, 그때가 바로.... 자기가 쌓은 성벽에 구멍이 뚫리기 시작한다는 것을 직감할 때야. **진짜는 그럴 필요가 없으니까"**

3. "세상에 쓸모없는 것은 없다. **쓸모없다고 여기는 사람만 있을 뿐이다.**

-웹툰에서 나온 대사중.....

INTRO

E-mail 주소는 blind 로 나가기때문에 공유가 되지는 않습니다. 혹시 개인 메일주소를 원하시면 저에게 알려주시면 update 하도록 하겠습니다. 그리고 Newsletter distribution 을 원치않으시면, 알려주시기 바랍니다.

지난 3 월 29 일(토)에는 Q1CY25 KASASV Golf Tournament 가 있었습니다. 참석해주신 모든분들께 감사드립니다.

4 월 16 일(수)에는 한미 AI 혁신센터에서 AI & Machine Learning 을 주제로 강연과 네트워킹자리가 있습니다

4 월 24 일(목)에는 Living trust 와 retirement/investment plan Q&A 세미나가 있을 예정입니다.추가 내용은 저희 웹사이트 announcement 를 참조하시기 바랍니다.

그리고 저희 협회에 KASASV Director 로 도움을 주 Soonho Kim 님, Youngsu Kim 님, 환영하고 앞으로 많은 도움부탁드립니다. Home page 를 참조하시기 바랍니다.

저희는 미국 Internal Revenue Code 501(c)(3)에 의거하여 Federal Income tax exemption 을 받은 비영리단체로 donation/협찬으로 운영이 되고 있습니다. 관심이 있으시면 저희 웹사이트 Partners tab 을 보시면 좀 더 자세히 보실수 있습니다. 올해부터는 off-line event 하고있으니, 스케쥴은 웹사이트(Announcement)를 참조하시기 바랍니다

- 사이먼 리 드림/협회 상근임원

www.kasainsv.com

감사합니다

Intel: New CEO, Same Old Game Plan

March 28, 2025 Mark Hibben

Summary

- Intel appoints Lip-Bu Tan as CEO, continuing Gelsinger's foundry strategy, which I believe is doomed to fail and a red flag for investors.
- Intel's x86 architecture struggles with power efficiency compared to ARM, and the foundry strategy aims to mitigate this disadvantage.
- Intel's 18A process faces delays and won't achieve high volume production until 2026, failing to leap ahead of TSMC's N2 process.
- The combination of Intel Foundry and Products businesses is a prescription for continued losses.

hapabapa

Lip-Bu Tan appointed new Intel CEO, continues Gelsinger strategy (for now)

Intel finally has a new CEO, Lip-Bu Tan, late of Cadence Design, where he engineered an impressive turnaround. And late of the Intel Board, where he recently resigned in protest to the way Gelsinger was running things.

Tan has now been restored to the Board. And apparently he isn't so unhappy with the way Gelsinger was running things to deviate from Gelsinger's key manufacturing strategy. This was to reshape Intel's semiconductor fabrication as a foundry service available to external customers, and that's still the plan, according to the WSJ:

The new CEO will soon have to answer the most burning question about Intel's future: whether he will break apart Intel's design and foundry businesses. Manufacturing chips is an enormous expense that Intel can't currently sustain, say industry leaders and analysts. Former board members have called for a split-up.

But a deal to sell all or part of Intel to competitors seems to be off the table for the immediate future, according to bankers. A variety of early-stage discussions with Broadcom, Qualcomm, GlobalFoundries and TSMC in recent months have failed to go anywhere, and so far seem unlikely to progress.

The company has already hinted at a more likely outcome: bringing in outside financial backers, including customers who want a stake in the manufacturing business.

Tan's statement to Intel employees, as reported by the WSJ, sounds an awful lot like Gelsinger's general approach:

"In areas where we are behind the competition, we need to take calculated risks to disrupt and leapfrog," Tan said in a note to Intel employees on Wednesday. "And in areas where our progress has been slower than expected, we need to find new ways to pick up the pace."

Leapfrogging the competition in process nodes is exactly what Gelsinger tried to do with the "5 nodes in 4 years" strategy/mantra. And we've seen how well that worked out. Intel was left having to turn to TSMC (TSM) to build their flagship mobile (Lunar Lake) and desktop (Arrow Lake) processors.

It begs the question, why is Intel so bound to the foundry strategy? I've pointed out that trying to convert its semiconductor manufacturing to a foundry capable of servicing external customers probably increased manufacturing costs in the near term. Intel had developed a manufacturing process over the years that was optimized for its own products.

I think the key to understanding the underlying motivations for Foundry is Gelsinger's stated goal to not merely achieve parity with competitor TSMC but actually get out in front of them. Why was this essential?

Because it was the only way to save x86. They understood that Intel's x86 architecture was disadvantaged in power efficiency compared to ARM, but they assumed that by having a process advantage, they could mitigate the architecture disadvantage and remain competitive. As I discuss in the next section, even Intel's most advanced chips fall well short of Apple (AAPL) Silicon in power efficiency.

That's what the whole "5 nodes in 4 years" has been about all along, and why Intel can't let go of the strategy even now. Tan has bought into this, because he knows what's at stake. So, the only choice is for Intel to double down and somehow make it work.

I'm absolutely convinced it's the wrong strategy and doomed to failure. But there is an ineffable mandate at Intel that x86 must be preserved at all costs. History is against Intel.

One of the things I've noticed about recent coverage of Intel is the way Intel's failures in smartphones or GPUs are characterized. All too often, articles represent that Intel's management "missed" the smartphone boom or "missed" the AI boom.

Those of us who have followed Intel through the smartphone and AI eras know that it was not for want of trying. Intel's management didn't miss the value of the mobile or AI markets. They simply made bad decisions and offered uncompetitive products for those markets.

Recall that Intel tried to offer processors for mobile devices, including smartphones and tablets. I wrote extensively about Intel's attempt to buy their way into the tablet market. This produced huge losses for Intel and a truly awful wave of cheap Android tablets.

Intel's fundamental mistake was trying to shoehorn x86 processors into mobile devices. It was kind of like Cinderella's ugly step sisters trying to force their feet into the glass slipper. But the x86 mandate was not to be denied.

Over and over again, Intel management insisted that everything was "better with Intel (x86) inside". This meant trying to turn a bunch of small x86 cores into a GPU, dubbed "Larrabee", and then attempting to market it as an AI accelerator when it failed to be competitive as a graphics processor.

Over and over again, we saw Intel management exhausting almost every available option before turning to GPUs as data accelerators for supercomputing and AI. These have included FPGA's through the expensive Altera acquisition, and ASICs through the expensive Habana Labs acquisition. Curiously, Tan was an investor in Habana.

By the time Intel got serious about GPUs, it was already well behind the curve. Intel's great leap forward in GPU accelerators was to be Ponte Vecchio. But the only significant customer for Ponte Vecchio was the Aurora Supercomputer at Argonne National Laboratory.

A successor to Ponte Vecchio was put on hold amidst Gelsinger's belated cost cutting, but according to Reuters, may be on again under Tan:

[Intel] will also restart plans to produce chips that power AI servers and look to areas beyond servers in several areas such as software, robotics and AI foundation models.

Intel management remain behind the curve in the fabless semiconductor industry. But once again, they had to obey the logic of x86 preservation. They couldn't go fabless, because that would level the process node playing field for x86 products, once and for all.

But they knew they couldn't afford to compete with pure play foundries like TSMC (TSM) in process technology. So, they aspired to be a Foundry, hoping that external customers would help finance their process node capital requirements.

It hasn't worked, and it will never work. No major fabless customer in their right minds would trust their designs to a foundry company that is also a direct competitor.

Granted, there have been reports of fabless companies evaluating chips fabricated by Intel, including Nvidia (NVDA), Broadcom and AMD. According to Reuters:

The tests by Nvidia and Broadcom are using Intel's 18A process, a series of technologies and techniques developed over years that is capable of making advanced artificial intelligence processors and other complex chips. The 18A process competes with similar technology from Taiwan's TSMC, which dominates the global chip market.

Nvidia declined to comment. Broadcom did not respond to a request for comment

These tests are not being conducted on complete chip designs but are instead aimed at determining the behavior and capabilities of Intel's 18A process. Chip designers sometimes purchase wafers to test specific components of a chip to work out any kinks before committing to producing a full design at high volume.

Sure, why not? It costs next to nothing to test Intel's chips, and it provides useful intelligence into what 18A is capable of. But will they actually use Intel to fabricate anything?

I'm sure Nvidia won't bother. Nvidia's most advanced GPUs, the Blackwell series for data center and desktop, use TSMC's N4 process. If Nvidia wants a more advanced process, N3 is available. So why risk going to Intel?

And I don't think Apple or AMD would bother either. Both are arch rivals to Intel.

Here's the conundrum Intel faces. The only way Intel Foundry works is if it's truly independent of Intel, which is exactly the one thing Intel can't afford to do.

Intel's most advanced chips still lag Apple in efficiency

When Intel launched Arrow Lake desktop processors last year, Intel claimed "Energy Efficiency Leadership". Certainly Intel did everything in its power to improve the efficiency of the processors. They used TSMC's most advanced N3 process, and they made various architecture changes, including ditching multi-threading, in order to improve efficiency.

But the claims of efficiency leadership were based on some rather transparent view graph sleight of hand:

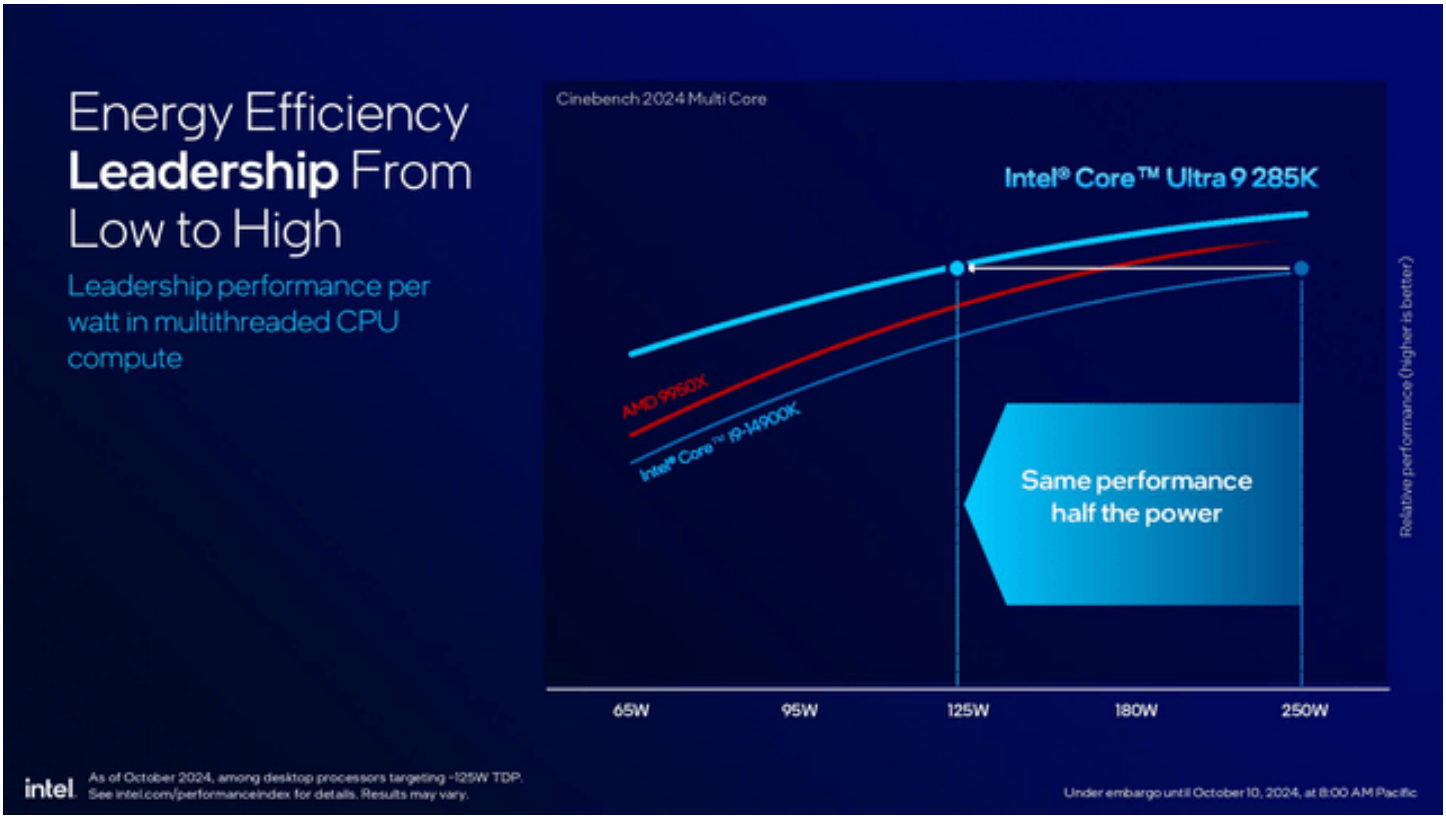

Intel

Anyone who has built a PC or done overclocking knows that there's a law of diminishing returns when increasing the processor clock rate. Power consumption increases dramatically for relatively modest performance gains. That's what the curves in the chart show.

Intel could have just as easily claimed that they had this great new processor that provided about 86% of the performance of the Core i9-14900K for half the power. And it would have been the very same i9-14900K.

It makes more sense to look at the efficiency difference along (vertical) lines of constant power. That's an apples-to-apples comparison. Then the efficiency improvement is just given by the performance improvement. At 250 Watts, the efficiency improvement for Arrow Lake is only about 14%, based on scaling the Intel chart.

The efficiency improvement for Geekbench multicore scores, turns out to be about the same at 10%:

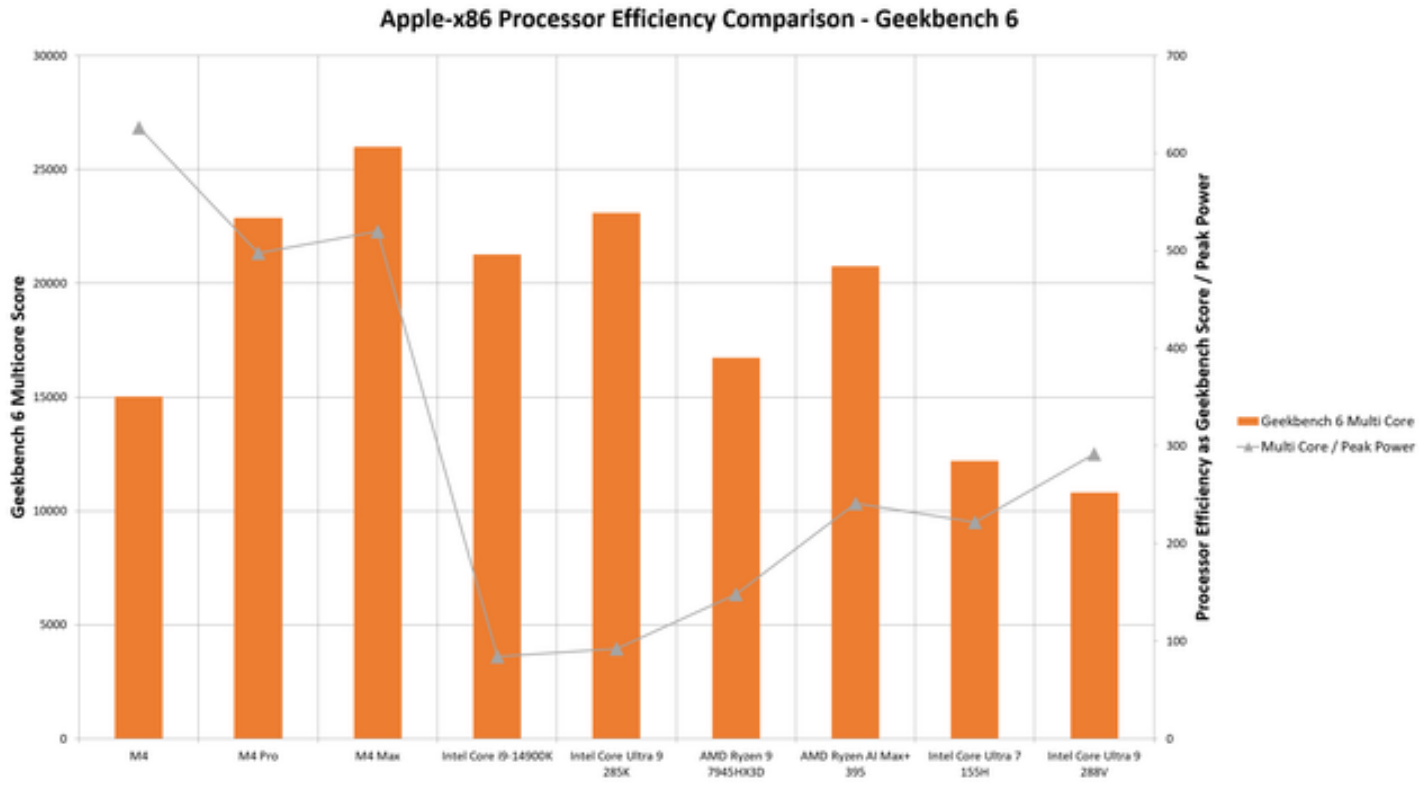

Mark Hibben

The data for the chart was collected from various reviews, as well as my testing of my own M4 Max MacBook Pro. The data is summarized below:

Device	Processor	Chip Manufacturer	Cores	Process Node nm	Clock Rate	Geekbench 6 Multi Core	Peak Power (W)	Multi Core / Peak Power
Apple 14" M4 MacBook Pro	M4	TSMC	4P + 6E 10 threads	TSMC N3	3.94, 2.81 GHz	15022	24	625.92
Apple 16" M4 Pro MacBook Pro	M4 Pro	TSMC	10P + 4E 14 threads	TSMC N3	3.85, 2.59 GHz	22879	46	497.37
Apple 16" M4 Max MacBook Pro	M4 Max	TSMC	12P + 4E 16 threads	TSMC N3	3.85, 2.59 GHz	25999	50	519.98
Intel 14th Gen Desktop	Intel Core i9-14900K	Intel	8P + 16E, 32 threads	Intel 7	5.6, 4.4 GHz	21276	253	84.09
Intel 15th Gen Desktop	Intel Core Ultra 9 285K	TSMC	8P + 16E, 24 threads	TSMC N3	5.6, 4.6 GHz	23095	250	92.38
AMD Ryzen 9 Laptop	AMD Ryzen 9 7945HX3D	TSMC	16P, 32 threads	TSMC N5	2.3-5.4 GHz	16726	112.9	148.15
Asus ROG Flow Z13	AMD Ryzen AI Max+ 395	TSMC	16P, 32 threads	TSMC N4	3.0-5.1 GHz	20754	86	241.33
Acer Swift Go 14	Intel Core Ultra 7 155H	Intel	6P + 8E + 2LPE, 22 threads	Intel 4, Intel 7	4.8, 3.8 GHz	12193	55	221.69
Asus Zenbook S14	Intel Core Ultra 9 288V	Intel	4P + 4E 8 threads	TSMC N3	5.1, 3.7 GHz	10801	37	291.92

Mark Hibben

Performance per Watt is treated as the efficiency metric where CPU performance is given by the Geekbench multicore score. Apple's M4 series performance and efficiency probably represents the best achievable for ARM architecture.

Clearly, Intel's performance doesn't represent the best achievable for x86, given the efficiency of AMD's new Ryzen AI Max+ 395. But even then, the efficiency gap between x86 and Apple Silicon is huge.

Given the fact that both AMD and Intel have spent about a decade trying to equal the efficiency of ARM processors, there is no reasonable expectation that they will ever be able to do so. Even the new Asus ROG Flow Z13 with the AMD Ryzen AI Max+ 395 has to be plugged in to an AC outlet to achieve the scores it does.

On the other hand, Apple's M4 Pro and Max notebooks exceed desktop performance while on battery power alone. In the history of computing, from the earliest mechanical calculators, through vacuum tubes and transistors to the modern integrated circuit, efficiency has always won.

Power efficiency wins because it enables scaling to higher levels of performance. I'm sure Intel's management understands this perfectly well, which explains their quest for process superiority. It's their only hope of closing the efficiency gap and remaining competitive.

And what really is the status of Intel 18A?

The Reuters article about 18A testing by potential customers also claimed that 18A had been delayed to 2026:

The 18A process was already delayed to 2026 for potential contract manufacturing customers. Now, according to supplier documents reviewed by Reuters and two sources familiar with the matter, Intel has pushed back its timeline another six months.

The delay is due to the need to qualify crucial intellectual property for the 18A process, which is taking longer than anticipated. Without the qualified fundamental building blocks of intellectual property that small and mid-size chip designers rely on, a swath of potential customers would be unable to produce chips on 18A until at least mid-2026, according to the two sources and documents.

It is unclear why the intellectual property qualification has been delayed. Qualifying intellectual property includes a guarantee from the supplier that it will work on a given manufacturing process.

Asked about the delay, Intel said, "(We will) begin ramping production in the second half of this year, delivering on the commitments we have made to our customers." The company added that it expects its factories to receive designs from customers this year.

Many chip designers are watching Intel's foundry progress closely in the hope they will be able to use its manufacturing soon, according to industry experts.

Intel's 18A process currently performs at a level between TSMC's most advanced process and its predecessor, Sassine Ghazi, CEO of Synopsys, said in an interview after its financial results. Synopsys supplies some of the crucial intellectual property needed for Intel's foundry.

The last statement about 18A is interesting. TSMC's most advanced current process is N3, but the reference may have been to the forthcoming N2. 18A landing somewhere between N3 and N2 would make sense. I have projected that Intel 3 was essentially equivalent to TSMC's N5, and I thought that 18A would be roughly equivalent to N3.

Following the Reuters article, there were reports that Intel had "refuted" the claim of a delay in 18A, based on statements made by John Pitzer, CVP Investor Relations, at the Morgan Stanley Technology Media & Telecom Conference on March 5. At first glance, it certainly seems that way. Said Pitzer:

Panther Lake is on track to launch in the second half of this year. That launch date has not changed. We feel really good about the progress that we are making. In fact, if you look at where our yields are on Panther Lake today. They're actually slightly ahead at a similar point in time to Meteor Lake. If you look at the development process for Meteor Lake.

It seems like a refutation, until you get to the last paragraph of Pitzer's reply:

Now I'll remind you that we will launch Panther Lake in the second half of this year. In a similar vein, we launched the Meteor Lake in the second half of '23. It wasn't until '24 that it became real volume. We launched Lunar Lake in the second half of last year, and this is the year where it becomes real volume. And as we've said on Panther Lake, we launch in the second half of this year, but it's really not until next year that we get to that volume really starts to help improve the margin profile of the overall company.

So there it is. No high volume production until 2026. And of course, a lot can happen in the interim. What exactly is the Panther Lake launch going to consist of? The tech media have been tossing around the term "paper launch" about the Nvidia RTX 5090, but Panther Lake is looking like a true paper launch.

And it's looking like the purported intellectual property delay for foundry customers is really just an excuse, because Intel won't be able to supply 18A chips to external customers while it's still ramping production for its own chips next year.

In any case, the production ramp for 18A will just happen to coincide with what I expect will be the production ramp for TSMC's N2. So, there won't be any great leap ahead of TSMC for Intel via 18A. And Pitzer also admitted that 18A falls somewhere between N3 and N2:

I think a couple of weeks ago, there was a technical paper out that actually looked at our SRAM density on Intel 18A that compared well with TSMC's N2. Lots of different metrics you can compare technologies on. I think, in general, we think about Intel 18A being an N3 type/N2 sort of comp with the external peers. We're also on track and we've announced this that we're going to tape out our first external design on Intel 18A in the first half of this year. So, we feel good about the progress that we're making there.

Investor takeaways: the handwriting is on the wall

That's nice that Pitzer and company feel good about their progress, but I don't think Intel investors should. The important innovators in semiconductors moved to ARM architecture some time ago. Apple has been using ARM for iPhone since its inception, and custom designing their iPhone SOCs for almost as long. Apple gained experience in making its ARM SOCs more powerful through iPad and had a very successful transition to ARM with its Apple Silicon Macs.

Nvidia has been designing ARM-based SOCs for almost as long as Apple for its Tegra line of mobile SOCs. These have transitioned into powerful edge AI processors for robotics and self-driving vehicles. Nvidia branched out into ARM-based server processors and now relies on them almost exclusively for its Blackwell based "AI factories".

ARM Holdings has announced that it's going to start building its own ARM server processors (as a fabless semiconductor company), and parent company Softbank is going to acquire Ampere Computing, which has its own line of ARM server processors.

Even if Intel suddenly decided to start developing ARM architecture processors, it would have difficulty just getting traction. Virtually all the ARM design expertise that exists in the world has been scooped up by the most important fabless semiconductor companies.

There's still a market for x86 by virtue of the need for legacy software compatibility and PC gaming, which continues to be the most advanced form of computer gaming available. And for those reasons, I still use an x86 workstation for work and gaming.

But I don't expect this to last. The growth of AI in the cloud and at the edge has made energy efficiency even more imperative. Increasingly, x86 is being shut out of the data center in favor of more efficient computing solutions combining ARM CPUs and GPU accelerators.

Personal computing has already shifted in favor of mobile, which has always been dominated by ARM. The last bastion of x86 personal computing is the desktop, and that's eventually going to wither and die.

Intel's current direction of pursuing a foundry business combined with its Intel Products business is unlikely to result in the recovery that Intel's investors need. Large external customers will likely forego using Intel Foundry in favor of pure play foundries such as TSMC.

Lacking the financial support of major external customers will likely mean that Intel never achieves semiconductor process superiority. As a result, Intel will continue to see its share of the server and client CPU markets erode.

The combination of Foundry and Products is a prescription for continued losses in which Product sales decline, and manufacturing continues to be inefficient relative to the pure play foundries. Given the poor prospects for recovery, I think Intel is vastly overvalued.

TSMC's Supercycle Is Just Warming Up

March 31, 2025 Yiannis Zourmpanos

Summary

- TSMC's 3nm and 5nm nodes drove 60% of wafer revenue in Q4-FY24, fueling 38.8% YoY topline growth.
- HPC revenue rose to 53% of total sales in Q4-FY24, up from 43% a year earlier.
- January–February FY25 revenue reached NT$553.3B, marking a 39.2% YoY increase despite seasonal smartphone softness in February.
- Overseas fab expansion may dilute gross margins by 2–3% annually, with Q1-FY25 guidance reflecting a 1-point decline.
- AI accelerator demand is projected to grow at a 40–45% CAGR, driving sustained advanced-node capacity expansion through FY26.

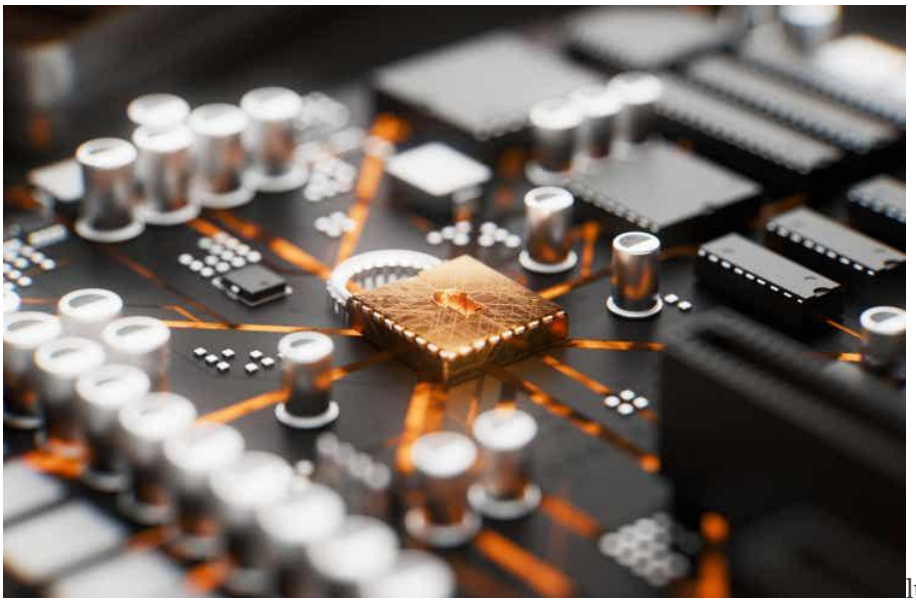
luza studio

TSMC's AI-Fueled Chip Surge: Why 3nm & 5nm Domination Signal a New Growth Supercycle

TSMC's topline growth is based on its advanced manufacturing tech in nodes like 2nm, 3nm, and 5nm. The 3nm-revenue held 26% of total wafer revenue in Q4-FY24 with a solid uplift from the 20% share in Q3-FY24 and 15% in Q4-FY23, whereas its 5nm tech brought in 34% of wafer revenue on a stable demand for high-performance computing (HPC) and

advanced semis applications (like AI). Combined, 7nm and more advanced nodes marked ~74% of total wafer revenue. These advanced tech nodes push TSMC to capture the expanding demand for HPC and energy-efficient chips (for smartphones and data centers).

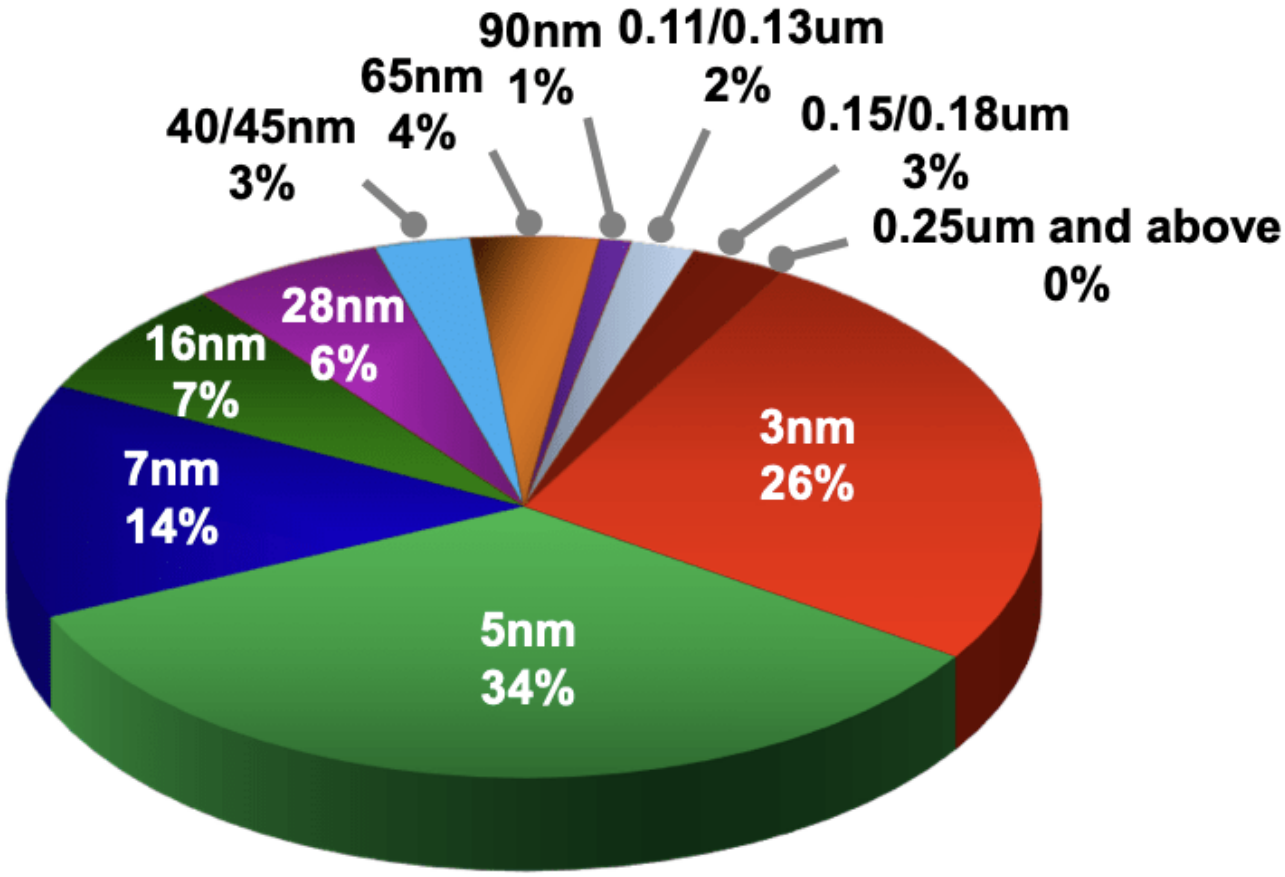

Tsmc presentation

Moreover, the topline from HPC grew to 53% of consolidated topline in Q4-FY24 against 43% in Q4-FY23. Meanwhile, smartphone revenue was 35% of total revenue, with a drop from 43% in Q4 to FY23. This shift in revenue mix points to the increasing demand for AI accelerators, data centers, and HPC chips. As a result, TSMC had a 38.8% YoY increase in consolidated topline in Q4-FY24 (reaching NT$868.46B). These improvements reflect the company's lead in higher pricing for its advanced semiconductor nodes. TSMC's strong pricing power emerged from its scale and process tech lead.

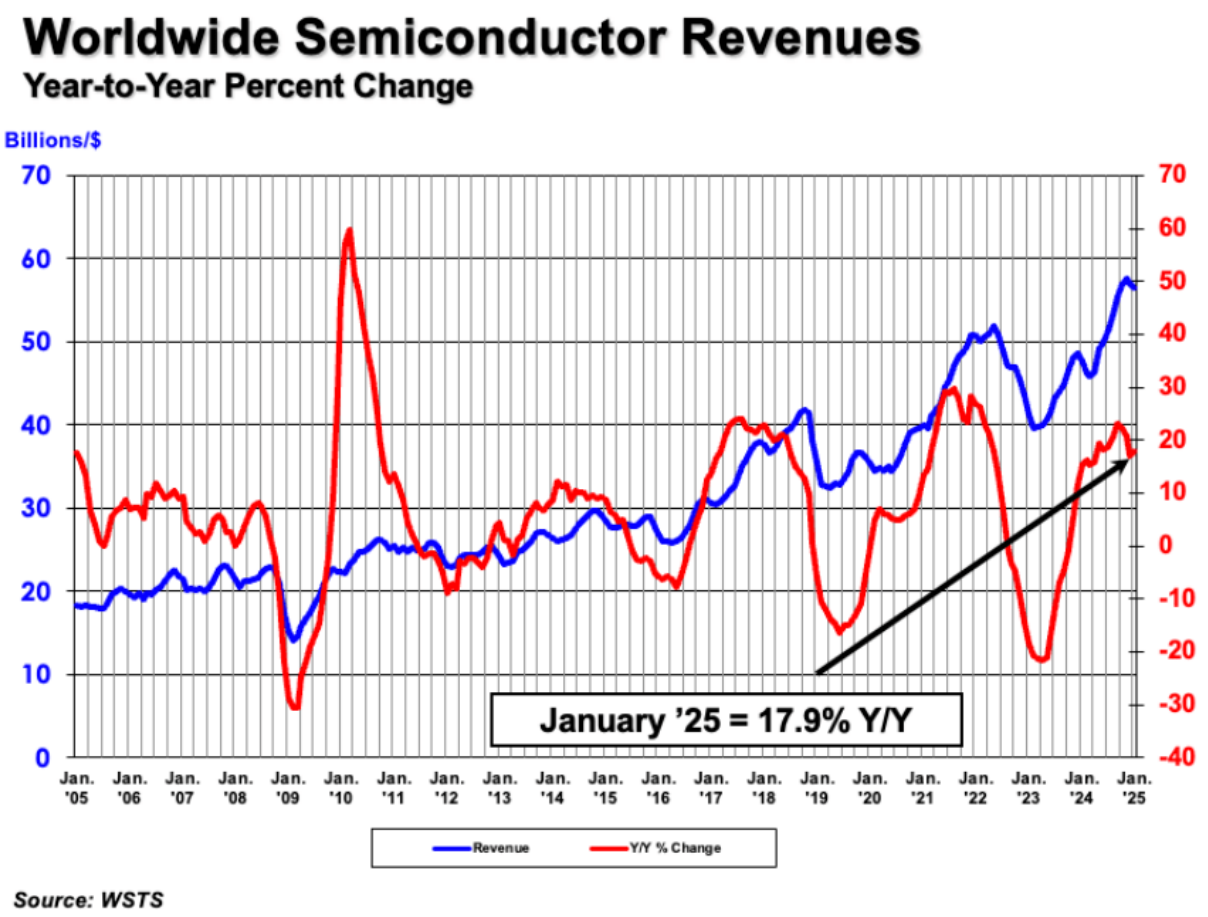

Semiconductors.org

Interestingly, the company's tech roadmap further reinforces its growth potential. TSMC's 2nm (N2) tech may enter volume production in H2 FY25 and may deliver a 10-15% speed improvement at the same power or a 25-30% power reduction at the same speed against its N3E process. Additionally, chip density may increase by over 15%. The subsequent introduction of N2P and A16 nodes set for FY26 also holds performance gains, pushing the demand for advanced semiconductor tech. With that, TSMC's revenue trends through February FY25, despite seasonal fluctuations, show solid growth.

In January FY25, TSMC had a consolidated topline of NT$293.29B representing a 5.4% increase from December FY24 and a 35.9% increase from January FY24. This massive YoY growth points to how the company is capitalizing on strong demand for its leading-edge semiconductor nodes (particularly 3nm and 5nm tech). The January revenue also reflects TSMC's resilience in maintaining high production output despite disruptions caused by a 6.4-magnitude earthquake in Taiwan. Although certain wafers in process were impacted and had to be scrapped, TSMC estimated the financial impact at NT$5.3B (net of insurance claims).

Now, the February FY25 revenue was NT$260.01B, which reflects an 11.3% MoM decline from January FY25 but a massive 43.1% increase from February FY24. The sequential decline is largely attributed to seasonal factors like post-holiday demand softness in the smartphone segment. TSMC typically experiences a slowdown following strong year-end sales.

Despite this short-term dip, the YoY growth reflects TSMC's continued expansion and the increasing adoption of its advanced semiconductor processes. For the cumulative period of January to February FY25, TSMC's topline reached NT$553.30B with a 39.2% YoY increase from NT$397.43B in the same months in FY24. This growth rate is in line with TSMC's broader performance trends.

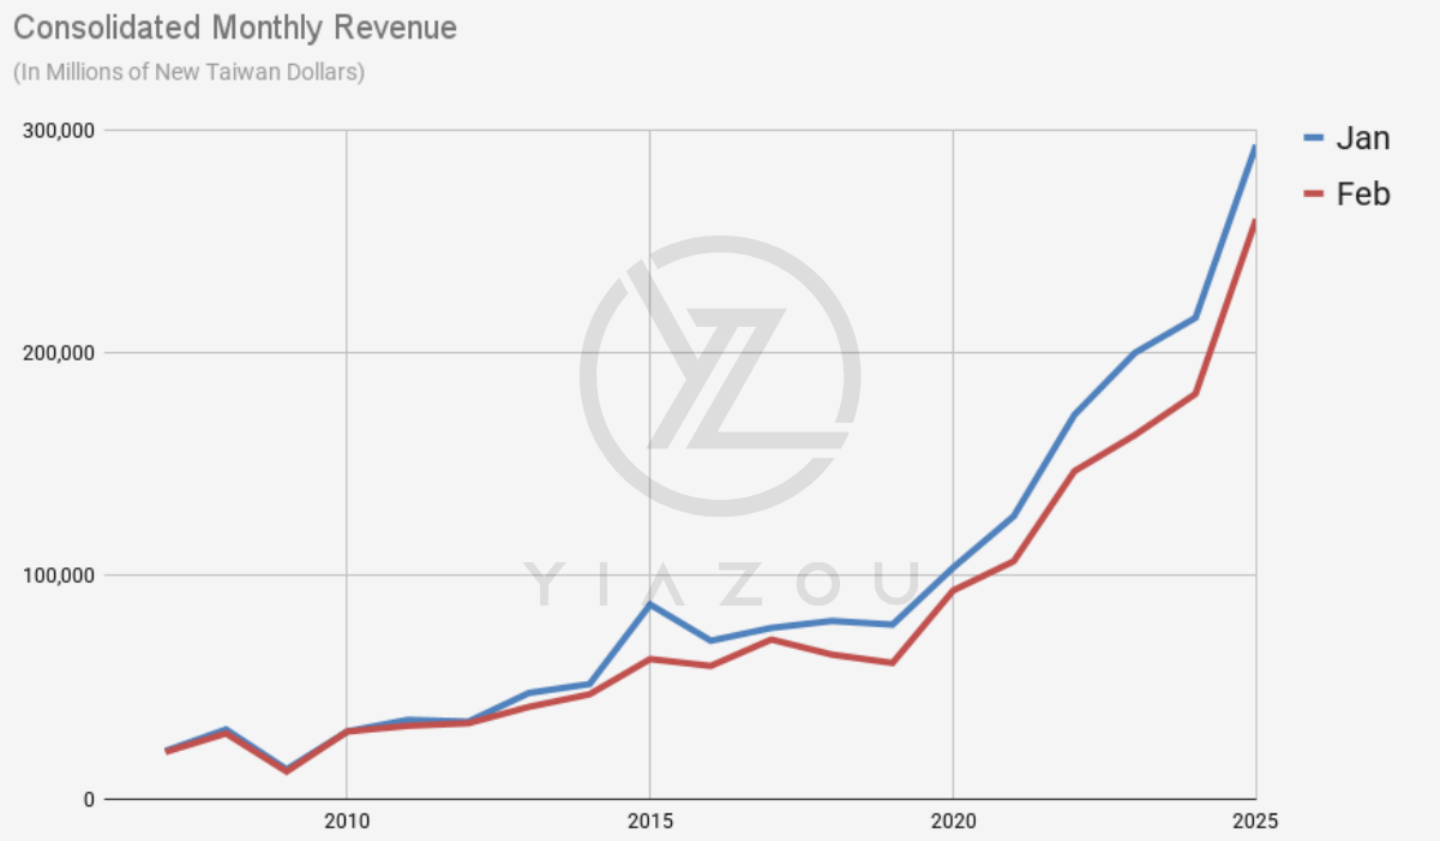

Yiazou

We now forecast the revenue growth from AI accelerators to approach a mid-40% CAGR for the five-year period, starting off the already higher base of 2024.

- C.C. Wei - Vice Chairman & CEO

Looking ahead, TSMC has provided revenue guidance for Q1-FY25 over $25.4B. At the midpoint, this represents a 5.5% sequential decline from Q4-FY24 but a solid 34.7% YoY increase. This projected growth reflects the company's ongoing high-value contracts for advanced chips' fab.

TSMC's Expansion Trade-Off: Global Growth Comes at the Cost of Margins and Flexibility

TSMC's rapid expansion of overseas fabrication facilities in the US, Japan, and Europe led to a financial weakness through margin dilution. Management forecasts a 2% to 3% gross margin dilution impact annually over the next five years due to the ramp-up of overseas fabs, with the effect becoming "more pronounced throughout the year" as facilities in Kumamoto and Arizona scale up.

In Q1-FY25, this dilution is less than 1%, but the midpoint gross margin guidance of 58% (down from 59% in Q4-FY24) reflects an immediate 1% decline due to these overseas costs alongside N2 and CoWoS expansion expenses. The higher cost structure in the US based on smaller scale, elevated supply chain prices, and an immature ecosystem (as explained by Wendell Huang), contrasts with Taiwan's optimized operations.

For instance, the Arizona fab entered high-volume production in Q4-FY24 using N4 tech, attaining yields comparable to Taiwan but incurring higher operational costs. The FY25 CapEx budget of $38B to $42B, with a 27.5% to 41.2% increase from FY24's $29.8B, points to the heavy investment required to support this expansion. However, this escalation in CapEx, coupled with a high single-digit percentage increase in depreciation expenses, pressures free cash flow. In FY24, FCF rose to NT$870.17B from NT$292.15B in FY23.

However, the FY25 outlook suggests a potential squeeze in operating cash flow (NT$1,826.18B in FY24) due to rising fixed costs. This dilution threatens TSMC's long-term gross margin target of "53% and higher," a threshold already challenged by inflationary issues like a 1% margin impact from higher electricity prices in Taiwan.

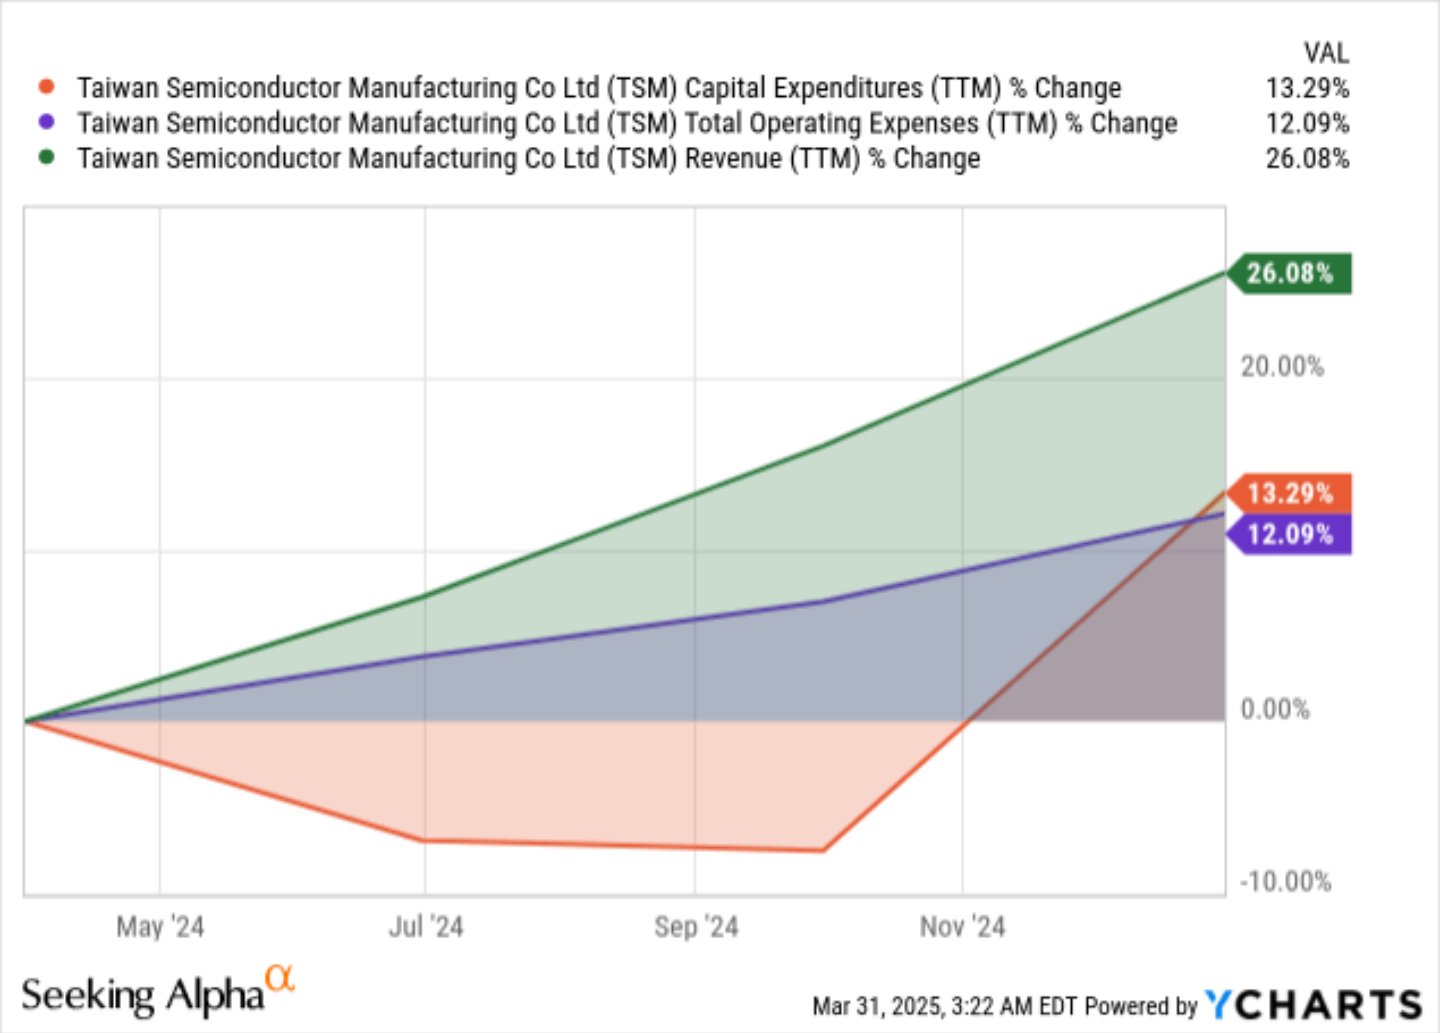

Data by YCharts

Moreover, TSMC's growth is increasingly dependent on concentrated demand segments like AI-related applications within the HPC platform. This introduces a market vulnerability. On a full-year basis, HPC contributed 51% of FY24 revenue, a 58% YoY increase based on AI accelerators (AI GPUs, ASICs, and HBM controllers). This comprised mid-teens percent of total revenue.

While this concentration fuels TSMC's projected mid-20s% revenue growth in FY25 (in $ terms), it also exposes the company to risks of demand fluctuation. Why? Q4-FY24 non-AI segments revenues, such as smartphones (35% of FY24 revenue) and PCs, exhibit only "mild growth," with unit growth in the low single digits. The February FY25 revenue drop from January is due to smartphone seasonality, which points to the limited buffer from other platforms.

Finally, the heavy reliance on CoWoS capacity (predominantly for AI applications) further amplifies this risk. Bruce Lu from Goldman Sachs questioned during the earnings call when non-AI applications might adopt CoWoS, to which C.C. Wei provided no firm timeline beyond "it's coming." This dependency on a single high-growth segment with tight capacity constraints (CoWoS unable to meet current demand) suggests that any slowdown in AI investment (due to potentially macro uncertainty or CSP spending shifts) can hit hard on TSMC's topline growth.

Takeaway

TSMC's advanced 3nm and 5nm tech is powering a 35% YoY revenue surge, driven by booming AI and HPC demand. Despite short-term margin pressure from overseas expansion, its pricing power and leadership in cutting-edge nodes position it for long-term upside. The AI supercycle is just getting started.

KLA Corporation (KLAC) CEO Rick Wallace presents at Cantor Fitzgerald Global Technology Conference (Transcript)

March 11, 2025 SA Transcripts

KLA Corporation (NASDAQ:KLAC) Cantor Fitzgerald Global Technology Conference March 11, 2025 10:40 AM ET

Company Participants

Rick Wallace - CEO

Bren Higgins - CFO

Conference Call Participants

C.J. Muse - Cantor Fitzgerald

C.J. Muse

We better get started. My name is C.J. Muse with Cantor covering semiconductor, semiconductor equipment. Thank you all for coming. I have the honor to introduce to KLA. I always think about KLA Tencor, but KLA team. Rick Wallace, CEO, Bren Higgins, CFO. Welcome, gentlemen.

Rick Wallace

Thanks, C.J., for having us.

Question-and-Answer Session

Q - C.J. Muse

So I have a number of questions, but I figured I would start with the all-important -- where are we in the cycle? And most management teams talked about kind of a mid-single-digit number for 2025. And my sense is everyone thinks they're going to beat that. And therefore, are we not in maybe kind of 5% to 10% kind of world for the year, which considering where expectations were three, six months ago, that's a real kind of positive surprise. So I guess just first question, do you kind of agree with that assessment and what's your early thoughts on 2026?

Rick Wallace

I'll let Bren take that one.

Bren Higgins

Yeah. So, 2024 was a good finish for the company, where WFE levels were about $95 billion company. So it's up somewhere in the mid-single-digit, company was up about 12%. So we had a nice relative performance last year and had a record revenue number, too, for the company, which I think under the circumstances with some of the moving parts we had related to China was a pretty good outcome. You're right. Our view on 2025 is more or less a mid-single-digit type year for the industry, which takes WFE levels up between $100 billion and $105 billion. And look, we'll see how that plays through. Certainly, for KLA, a return to growth at the leading edge, both in logic and foundry, but also in memory is a real driver for our business. We're excited about those opportunities. I'm sure we'll talk more about them. So I think that is the biggest driver, and it's a higher process control opportunity, particularly compared to the last two years, which were more legacy-centric. Easiest way to think about it is that logic and foundry WFE on the front end grows in terms of leading edge growth, offsetting some digestion and export control affected business in China. The legacy parts of the industry are more or less flat. We see a little bit of growth off a very low base in the NAND flash market and then investments in DRAM specifically supporting HBM drives the uptick in the year. Against that backdrop, we feel very good both in terms of our better relative performance both in logic/foundry at the 2-nanometer node versus 3 but also the dynamics around HBM and how that translates into opportunity for KLA. Packaging is also a big opportunity in '25. We talked at earnings about growth of about at least $800 million, up from $500 million in advanced packaging. So another vector there of opportunity supporting what looks to be a pretty strong, not just this year, but as we go forward, high-performance compute market. So just looking at all that, I think it translates into mid-single-digit kind of growth. We feel very good about KLA's relative position. We'll see as we go through the year, whether that view changes. But that's where we sit today.

C.J. Muse

Perfect. So I think I'll kind of pivot now between some of the end markets in '25, '26 trends and key technology inflections for KLA this year and beyond. So maybe to start, optical versus e-beam, obviously, you guys are the best in the world in optical. But we met a few weeks ago and the comment that you made to me was you are all in on e-beam, which seemed like a bit of an inflection. So can you expand on that?

Bren Higgins

Well, the inflection is that it's showing up in the market. We've been spending on e-beam, as I talked to Ahmad, for quite a while. And the question was, when was it going to actually start showing up in wins? And so we've always had a view that e-beam and optical are complementary. And what we needed to have was a fast enough e-beam that was also going to be able to sort through a lot of the noise. It's funny when we talk about AI in our products, the first product to have AI was an e-beam inspector from us back in like 2018. And so we're finding now with the latest generation that we've introduced that we're actually getting wins, and we're seeing growth. And some of our share gains this year, what we anticipate as part of outperform is because of the strength of the e-beam portfolio. And we're seeing -- it's surprised everyone on the operations side that we keep getting increase in the demand for that. Yeah. And it's interesting for us because it's always been something that's been somewhat complementary. The complementary markets to the optical franchise. But now you have other applications, particularly in terms of gate-all-around in characterizing transistors and transistor performance. And so over time, we expect roughly 80% of the dollars in pattern inspection to be spent on optical technologies because those translate to production, e-beam has great sensitivity, but is incredibly slow. But it's more than just pointing our inspectors and increasing the relevancy of optical inspection. There are unique applications that we think we can differentiate that also. That's also a factor in how we see it. Now there's also other parts of the market. There's a metrology market where you're doing more in-die measurements to for overlay metrology as an example. And then there's review. And so when the dynamics is anytime you ship an optical inspector, you have to ship a review system to review and classify those defects. And -- it is something that being able to do that faster is a driver because of all, to Rick's point, all the nuisance that's there. So it's absolutely a factor that we think is a share opportunity. Those are markets we haven't had a strong share over the years, there's growth. And so we'll lay out a deeper analysis on this at our Investor Day in June, but we're certainly encouraged by what we're seeing in the market today.

C.J. Muse

Maybe just to get ahead of that a little bit. So you guys exited CD-SEM, but you just talked about SEM review throughput.

Rick Wallace

In 2007.

C.J. Muse

Yeah, not suggesting near term, but just trying to get the landscape of where you're playing. So it's SEM review and it sounds like throughput, there are some advantages there now. And then the other one, big one is EBI and, I guess for e-beam inspection. Are you attacking all three markets, voltage contrast, physical defects, barrier defects?

Rick Wallace

Yes, and also reticle, I mean we have a system that's being used for advanced reticle inspection too. But the answer is yes. I mean, part of why the strategy took so long for you to hear about was we had debates about this internally. Do you develop one segment and then grow or do you develop for all the segments and bring them out. And we actually want the second path, which is a longer and more expensive. But that's what's showing up now. And so we've got a portfolio of products for our customers, and we're essentially seeing it both in the logic, advanced logic but also advanced memory showing up and getting really strong results. Part of why I wasn't sure I was going to be here as I was in Asia last week, and I was in Taiwan and Korea talking to customers. And it was interesting how their view of our e-beam has changed pretty dramatically from prior visits. So very positive and wondering why they can't get more capacity.

C.J. Muse

And what are the lead times for e-beam inspection?

Rick Wallace

Well, from a new order now, it's...

Bren Higgins

They're pretty consistent with other lead times to somewhere in that 9- to 12-month range. I think we didn't fully anticipate the strength of the demand. And so we're struggling a little bit to ramp, although I think we're in a much better position today than we were before. So, the demand has been strong, and we'll work our way around it. I don't think it's an impediment to the opportunity. It's just one of the usual operational challenges you face.

C.J. Muse

And when you talk about optical and e-beam complementary, obviously, they're very complementary in terms of, as you talked about, what you can see in throughput and balancing that. But is there an opportunity and an advantage for a customer to select you for both where they work well together?

Rick Wallace

Yeah, absolutely, both from -- we're now sharing some of the algorithms, so some of the AI work has been done on both, and we actually have linkage between the system so that to Bren's point, you can point them but also you can validate. So if you could find it on e-beam, can you then turn the optical to run it at higher speed and be able to continually monitor. But that was always the question was, what was the value of the portfolio, was there synergy? And finally, we're seeing realization and customers are pulling it back to us.

C.J. Muse

And I guess, is there a framework for thinking about the driver of adoption from 3 to 2 to I going to get the number wrong A16, A14.

Rick Wallace

I think it's still -- the customer is exploring. I mean there's a waterfall. If you look at 1,000 steps in the fab and there's, let's say, 200 that are being incepted monitored on a routine basis and like right now, for the most part, only two of those in advanced logic are being used by e-beam. I mean mostly everything else is the fastest tool that can do it. So whether it's a dark field tool or whether it's a BBP tool. And so the numbers are expanding as you go forward, provided we can have the coverage and throughput. So the challenge always for any of these technologies was if it's used as an engineering tool, it doesn't drive capacity because you're only debugging a line with it. And so the answer is yes, we're seeing them implement. And there are some layers now, and this is -- was a shock to us that they're inspecting twice the same layer with the optical and e-beam because they're seeing different things.

C.J. Muse

So you also talked about reticle with e-beam. Obviously, very strong 640ES Teron Optical, 8xx multicolumn e-beam but also soon the Teron 710 actinic. So as you look at that portfolio, is there a framework for thinking about your share for overall reticle inspection? And should there be a difference between in-house versus mask houses?

Rick Wallace

Not a huge share. Well, I mean, recall is a different application for sure. And then it's more going to be optically based. But when you say soon, I think soon on 710 is...

C.J. Muse

Sorry, '26 to '28.

Rick Wallace

Okay. It depends what your time frames are for soon. That is -- the answer is yes. And really actinic for reticle, our view has been consistently that when it's really going to be needed is when there's high NA lithography, EUV lithography, which has slipped out. I think there might be a couple of layers at 14, but it's out there away. So I think it will more intersect with HBM for that. So later and right now, the question on 8xx has been, can we continue to provide more capability by adding more columns and getting the throughput out. So that is -- we're getting good feedback on that. And then the optical tools continue to expand based on algorithms and just what we've always done provide more illumination and more algorithms.

Bren Higgins

We're covering most of the market with the optical systems, right, that 90% of the EUV reticles in going through optical systems. Given the nature of the lines and space density that's on the reticles. So we cover most of the production use cases and then you've got what's happening in the fabs where they're requalifying reticles, you all see a wafer inspection tool that doesn't show up in how you calculate reticle inspection, but it's a use case for wafer inspection to print the wafers and use that to validate the fidelity of the reticle. So those are the principal tools in the market, and we've got some advanced capability, obviously, with e-beam and actinic down the road. And we just -- and we talked a little bit about a partnership we have with a key supplier with Carl Zeiss as it relates to the optics for that system. And so the ability for us to make adjustments to the DNA very quickly and there's some synergistic value between the optics and the support that they do the scanner and for the inspector. So I think we're in a pretty good place when the market need materializes.

C.J. Muse

Perfect Maybe moving to advanced packaging, including HBM. You talked about $500 million in '24, $800 million in '25, obviously, great growth. But I believe that's largely CoWoS. And so Curious, as you think about going to HBM 3 and 4, can you kind of walk through how the opportunity opens for KLA? Where are you in terms of having the right price, right gross margin to KLA tool to serve that market? And how do you kind of balance the margin requirements that you need, but also this high-growth market that other players have traditionally supported?

Rick Wallace

Maybe I'll intro and Bren can fill on the details. So I remember needing probably 18 months ago, one of our major suppliers where they told us that we want your WFE your front-end stuff in the back end. And we said, well, I don't think you mean that because it's too expensive, right? You're going to say that, and then we're going to put it in there and you're going to say that -- and they said, no, no, we're dead serious because what we're seeing in packaging is the requirements are looking a lot more like the front end. And so you take what was a low end in the front end and it becomes a high end in packaging. And a lot of our growth and success in addition to some of the process stuff, but a lot of our growth has been that conversion of some of those capabilities and just the adoption of optical inspection, the adoption of some overlay, the adoption of flatness. And those are -- what I miss, films? Yeah. And so those are the reasons. And so what they're telling us is the cost is so high. And C.J. I always look at what is the cost of what they're making and what percent can they add on an inception. Right? So if you think about reticle, you can understand why they can spend as much as they do. The cost of the reticle is so high. The issue was the cost of the packaging wasn't high enough to spend a lot of extra money on the inspection. That's not true anymore. And the cost of failure on a package. And so they're throwing a lot more process control intensity at it. And what they're telling us is that's not slowing down, right? So they're trying to streamline the HBM process. But even as they do, they need more. So the answer is we say $800 million, but we don't think it stops there in terms of growth. And the beauty is we've got all the roadmap. We have all the products. Right? Because all you have to do is continue to adapt the stuff that we have in the leading edge. The challenge for the guys that have come at this market from the low end was that they're already stretched to the top of their capability. So when we talk about packaging as our low end is their high end and they're kind of at a limit, whereas we've got a huge road map plus the opportunity that we have because we have the processing SPTS division. So those are kind of -- the road map we see from the packaging guys is fully integrated to the conversations we're having with our leading customers on the front end. And I think that market keeps growing for as far as we can tell.

C.J. Muse

So if we isolate to HPM. Can you kind of walk through, I guess, bumps are getting smaller, wafer thinning, what are the tools that you see as an opportunity for KLA?

Bren Higgins

So those are the things that are happening, right? I mean, generally, is that you'll run into height restrictions. And so that will -- you see will have thinning or smaller bumps, thinning wafers and then you also run into eventually to hybrid bonding techniques. So right now, historically, bump inspection has been more of a legacy packaging market, but as you start to increase the capability there, it creates opportunities for us to address with higher-value systems. So they're the same systems in CoWoS but it's inflecting more from a technology point of view. So most of our exposure today has been, to your point earlier, has been logic and but we're increasingly optimistic about the momentum in memory as we go forward for some of the things that you mentioned.

C.J. Muse

And as you think about going to 5 or 3-nanometer on the logic die for HBM, freeing up space there for adding incremental logic. Is that an opportunity that you see as well?

Rick Wallace

Yeah. So HBM -- you mean in memory or you were still talking package.

C.J. Muse

The logic die within the HBM4 stack.

Rick Wallace

Yeah, of course. And so I think that they're back to the every conversation we have, customers are frustrated that they're having to add process control as a percent of their spend, right? So that's a natural thing for them, but they're also saying they have to do it. Right? So we know that there is a driver and opportunity. They want cost-effective solutions to do it, and they want us to reengineer some of our tools for those applications. I don't know what the...

Bren Higgins

Yeah. I mean one of the things it's roughly 12-nanometer today, and to your point, going to something more aggressive. And if you look at what the biggest opportunity in HBM is how it changes DRAM intensity on the front end, Forget the packaging part of larger die because you have to drill the TSVs, you got microcontrollers, you have less redundancy. You got to make -- do you have to do this logic-based eye, then you have reliability checks to make sure that each of the chips functions well in the stack. And so all that's driving the bigger opportunity for KLA in terms of just overall intensity as it relates to DRAM, then you take that HBM do some of the connection stuff that we talked about earlier and integrated the interposer with the logic chip. So that's the evolving opportunity, but we're really encouraged by what we're seeing just in terms of the DRAM die and the DRAM processing relative to advanced DRAM in terms of the intensity.

C.J. Muse

So we hit on advanced packaging, but really focused on HBM. Are there other technology inflections that we should be thinking about for other parts of advanced packaging?

Bren Higgins

I think the biggest issue is that you start to shrink the interconnect density and the interposer, which drives you to need more sensitive solutions. To Rick's earlier point, we're addressing this part of the market with very high sampling rates. Your 100% inspection with our lowest end capability. But we have the portfolio all the way up to e-beam, not saying e-beams part of this. But we would -- as you start to shrink the density, then you're going to need more advanced systems. And typically, when you need more sensitivity, there's a throughput trade-off. And so it's why I'm less concerned in the long run about the margin -- gross margin implications of packaging today, which is a negative from a put and take point of view to the gross margin because it is our lowest in solutions. But over time, we would expect that we're going to ship more advanced solutions, which will carry a stronger margin profile that, I think, turns it into more kind of normal KLA over time. Now we talk about these deltas. I mean it doesn't change our longer-term view of the margin profile over time. But certainly, in the near term, has been a bit of a headwind as those tools carry margins that are different than some of our advanced front-end tools that make up a bigger part of our mix.

C.J. Muse

Got you. Maybe focusing on leading-edge foundry, which you alluded to a little bit earlier. And I think there's a couple of areas where you're very excited. So first one, going to N2, I think you talked about process control intensity higher by 100 bps? And then secondly, as you move to HPC very quickly on 5 and 3 you start talking larger die and requiring more process control insertions because you don't want to have the yields on those larger die be a disaster. So as you think about those two trends, how are you thinking about overall kind of foundry for KLA process control intensity into the '25, '26, '27 time frame?

Rick Wallace

Well you can ask the number.

Bren Higgins

Yes. So for N2, and that's KLA share of market, right? We have talked about our versus N3 that our share of market expectation is about 100 basis points higher and a piece of that is share, we talked about e-beam earlier, but also there's intensity changes. The nature of a gate-all-around device is complex. You still have -- while you don't have an increase in the number of litho layers, you still have very advanced litho, EUV layers or increasing EUV layers. So they're scaling defects down in the interconnects of the transistor design that are still a challenge, but then you have new defect mechanisms related to residual defectivity, buried defects and so on. So all that, we think, is contributing opportunity. The last time we had an architecture change in the industry, KLA's share a market approach to 8%. I think when you look at the numbers we talked about earlier, I think we're more or less approaching that again overall across the market. So if you go forward, you look at the ramp of N2, particularly with the number of designs, some of the issues you talked about large die, higher value die, all those things we think are conducive, less reuse because of the adoption of the node, the strength of N3 and so we're pretty optimistic about that, but also the road map going forward, both in terms of shrink, but also then you started talking about power distribution, other opportunities here over the next couple of years. So I'd like to say that you have all -- over the last five years or so, you've had a lot of things that have driven KLA share market up, logic/foundry mix relative to memory, well-adopted nodes, limited reuse. I think all those things continue. But now the mix is shifting. The mix is shifting to high-value die, big die and so that's creating, I think, opportunities for us to see it continue to grow.

Rick Wallace

Well, and the one continually surprising fact is the number of leading-edge designs keep going up, right? And so you think about the conventional wisdom years ago was there'd be fewer because the costs are so high. But yeah, what we're seeing is -- and this is a unique benefit for KLA. If the number of designs go up, it means it's a higher mix fab, which means process control becomes more important. It's good for the reticle market because it means there's more reticles being produced. And the cost of these are so much back to my point about the need for inspection and measurement is also supported by the cost structure because you have to make sure that these are right. We also know that there's an undersupply relative to demand of 2-nanometer right now in all the forecasts. We know that there's more stated demand and when you translate all the hyperscaler CapEx they could come down a lot and you'd still be short of demand. And so that's a great driver for us. But the other thing is that it's playing to our technology strength. So we're gaining share because the performance differentiation becomes more important. So all the work we did on AI in our systems is showing up with our customers. So that's why we're pretty bullish on that, but also on the memory side of everything that's happening in high-performance compute and AI, those memory chips are getting much harder to inspect. I mean it's literally memory is more important, the logic is harder and the packaging. Right? We have all three at the same time. Memory, I was thinking about this, there was a time where memory was our big driver. And then years ago, people went to 100% redundancy. And the reason they did that and said, if you had a problem with your computer, you just reboot it because it was a memory leak, right? That's not going to be okay these high-performance compute, right? So suddenly, the intensity around memory manufacturing for inspection and measurement has gone up again because they're terrified of being the weak part of that very expensive stack. So we see process control intensity going up there. So it's kind of high-performance compute is if we had to design a technology to drive our markets, we couldn't have come up with a better one, and everything is happening with AI and high-performance compute.

C.J. Muse

Maybe a question on Intel. And I know you can't be too specific, but I guess following on your comment around any which way you slice 2-nanometer, there's shortages. My question is geared around they either have to bring wafers in-house or they have to recommit to more wafers at TSMC. And so this is more of a high-level kind of foundry/logic question in terms of -- do you think the industry is ready, already capacity-wise to support that kind of binary decision or whenever it happens, we're going to need to spend more money?

Rick Wallace

Somebody is going to need to spend more money. Right? I mean, it's not -- because the demand is already above the forecast supply, nobody is allocated for that when we talk to customers, and there's only a couple you could talk to and so if there's success on those leading nodes, they're going to need some level, whether it's 40,000, what we think 50,000. They're going to need some level of capacity to support it. And the 2-nanometer stuff is already spoken for. Now you always plan for less than what customers are telling you want. But yeah, there's going to -- if there's success there, there is going to need to have more capacity added. So that -- and that's upside to the model we talk about for 2025, if it were to happen in '25.

Bren Higgins

Yes, we contemplate different demand scenarios in terms of how we plan for our production, given our lead times, the nature of the lead times, the intrinsic lead time to some of the components. So implicit in our model is carrying the flexibility to be able to support surprises and demand that are inside of lead time. So I think we're in a pretty good place from a capacity and supply chain point of view to support the different scenarios that are potentially out there. And so we'll see how that plays out. But to your point, and to Rick's point, it seems like that there's opportunities for upside out there, and we'll have more to say about it when I think we materializes.

C.J. Muse

And just to clarify, Rick, do you think that surprise could actually happen in '25? My sense was there wasn't enough clean room space.

Rick Wallace

No, no, I'm saying they could recognize they need it within '25. And so it would be -- you couldn't bring it on in '24.

C.J. Muse

Got you. Makes sense. How about China? You previously talked about a $500 million headwind and now you found other business kind of to offset it. Do you feel like you've completely derisked any sort of headwinds there? And what do you think kind of normalized China is as a percentage of your business over time?

Rick Wallace

So I talked about it in earnings that we would come -- our business percent would come down from about a little over 40%, I think 41% last year to somewhere around 29%, plus or minus this year. So a pretty significant reduction. There's capacity digestion that's happening. There's obviously the export issue you mentioned. So the last two years have been fairly unique in that you had a drop-off in investment from your non-China customers. And your China customers with the greenfield facilities and a lot of public funding we're able to basically fill that void. So it drove the percent of business to elevated levels, but it was really unique, I think, to the time frame and the constraints we had in '21 and '22 in terms of being able to satisfy overall demand. So I think over time, it drops down somewhere to 20% to 25%. There's self-sufficiency directives in China over time. And if you look at an 80% target and where they're at sort of 20-ish percent today, it does mean that you're going to see a level of investment going forward, $25 billion to $30 billion of WFE. So I think as you model that out towards the end of the decade, we see it somewhere in that range. So yes, we had the impact the overall estimates didn't really change. We did see some marginal improvement at the leading edge, some additional strength in China from where we were when we sort of set those expectations. And we'll see what happens in terms of that impact over time. Is there a licensing opportunity or not less clear. So you can reallocate certain systems that are strongly demanded but that's just a pull forward of business. So from a gross impact level, I think we're comfortable with the estimate we gave plus or minus $100 million. And then we'll see if there are opportunities to mitigate that, and that changes moving forward, but at least that's where we see things today.

C.J. Muse

Makes sense. So if I were to look at my bottoms out model, I look at domestic China, Samsung, less so TSM or Hynix due to lack of clean room space as sources of upside for 2025. Would you generally agree with that view?

Rick Wallace

Yeah. Sure.

C.J. Muse

Okay. How about gross margins? So you've guided 62% for the full year. You've contemplated decline, I assume, of higher-margin China embedded in that. At the same time, you've exited or divested flat panel business. I guess when you see an uptick in volumes, what kind of incremental gross margins should we be thinking about?

Bren Higgins

So, our margins by region are generally the same. What drives our margin variability quarter-to-quarter tends to be product mix. Customers are buying different types of tools and carry different margins outside of volume-driven deltas, that's really the only difference across different customers. And in fact, if you look at customers in China, given the nature of some of our older systems, it may actually carry lower margins. So I don't think on a go-forward basis that that's necessarily a headwind. We talked about the packaging dynamic earlier. So we feel like we're operating today about 62%. As you mentioned, that's the guidance we gave for the year at roughly $12 billion business levels. So our target for our '26 model was 63% plus at $14 billion business levels. So given the mix we have, I feel pretty comfortable about our ability to continue to drive gross margins between 60% and 65%, which translates into a number in line with our target model. The exit of FPD will have a marginally better impact probably 20 basis points on the overall. So it's a factor, growth in our software business is also offsetting some of the mix dynamics I talked about earlier. So it's, I think, very consistent with what you've come to expect from KLA. The biggest dilution factor we have is our service business. Service is growing at a faster rate. And so we have a lot of focus in our service business about driving the incremental margins of service, but that's contemplated in that 60% to 65% long-term incremental gross margin target. One of the challenges in service is supporting a much broader footprint. So you have to make investments more broadly, but you'll ultimately get leverage on those investments. So we feel pretty good about the ability to manage whatever that diluted for us, I mean, we love the business. It's a great business for us. But -- and again, as I said, I think it's contemplated that overall target.

C.J. Muse

Perfect, I think we've run out of time. So, Rick, Bren, really appreciate it. Thank you very much.

Rick Wallace

Thank you.

Bren Higgins

Thank you.

Lam Research: Time To Buy Despite Tariff Concerns

March 7, 2025 FutureRich Investing

Summary

- I rate Lam Research as a Buy due to its leading position in semiconductor processing equipment, crucial for the AI and technology boom.
- Despite geopolitical risks and potential tariffs, Lam's strong market position and significant revenue growth make it a resilient investment.
- Lam's financial health is robust, with impressive 5-year growth in revenue, net income, and free cash flow, supporting long-term profitability.
- Intrinsic value analysis shows Lam is undervalued by 28.50%, with a buy price of $91.42, making it a compelling investment opportunity.

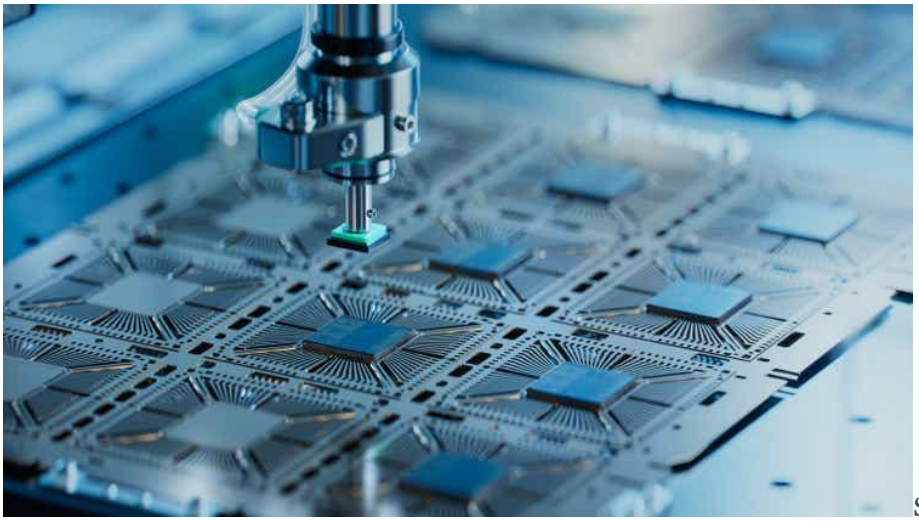
SweetBunFactory

Today I want to explain why I rate Lam Research (LRCX) as a BUY, based on my belief that they will be one of the main players in the coming technological boom, thanks to their leading position as a semiconductor processing equipment manufacturer, in two of the six critical wafer fabrication steps: Etching and Deposition.

In this analysis, we will go through the company's latest Q2 '25 earnings report, as well as get to understand their market position, their competitors and geopolitical threats. We will further analyze the company through our 9 Pillar analysis and finally calculate its intrinsic value using our 4 valuation models.

Introduction

Lam is one of the three top companies worldwide engaging in semiconductor processing equipment manufacturing, alongside with ASML Holding (ASML) and Applied Materials (AMAT). First, let's understand what these three companies actually do.

To create a semiconductor, a wafer (the blank base of a semiconductor often made of silicon (Si)) has to undergo six processes: Photolithography, Etching, Doping, Deposition, Chemical Planarization (CMP), testing and packaging. Lam is the world leader in the creation of machines that carry out two of these processes: Etching and Deposition, with a 35.00% and 29.00% market share, respectively. This makes it quite obvious that they play a key role in the AI and technology boom of the early 21st century.

Analysis

Market Overview

As I already mentioned, the company plays a key role in wafer manufacturing worldwide, but let's dig in deeper and uncover some potential issues. Firstly, I want to point out that Lam operates in 14 primary locations, of which only three are in the USA. You probably already know where I'm going with this, yes, danger from tariffs.

Primary Locations (Lamresearch.com)

With Trump proposing to implement higher import tariffs on South Korea and Taiwan, the two countries that hold more than a third of the company's primary locations, proposing up to 100.00% tariffs on semiconductors made in Taiwan, and suggesting increasing the tariff on South Korea, Lam's profitability might greatly suffer.

Furthermore, 17.00% of the company's revenue comes from sales in Taiwan, mainly TSM, with 25.00% coming from South Korea and 31.00% coming from China. This accounts for a combined percentage of 73.00%, and 91.00% when we consider Japan, which is more than 3/4 of the company's total revenue coming from Asia, which, as we all know, is the US's competitor, and therefore the current Trump administration has hinted towards and proposed the banning of a portion of semiconductor exports into that region. This would greatly affect Lam's ability to generate revenue.

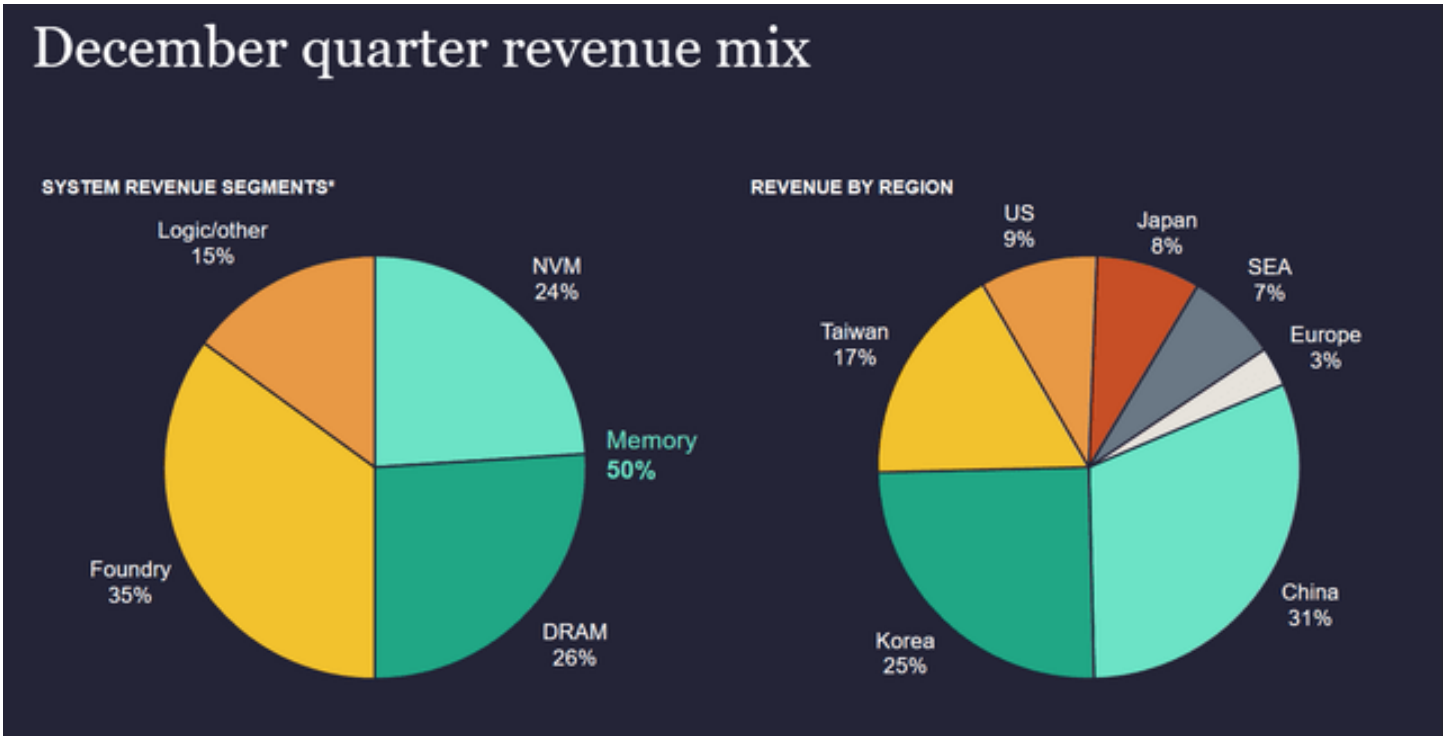

Revenue By Region (Lamresearch.com)

The Tariffs Won't Go Through

But let's not get all negative. I personally believe that the Trump administration will not go through with most of their tariff threats, at least not to the full extent as they promise. The reason for my belief is that if the tariffs were to go through fully, despite Trump's belief that they would strengthen the US economy, it would cripple it. The country would suddenly generate less income due to cut demand. The cost of labor would also increase, which would further decrease companies profitability, forcing them to increase product prices, further decreasing the profitability of their customers, throwing the US into a strong cycle of weakening economy, while other countries would bloom. This is something I believe Trump is aware of, and is rather using the tariff threats as a show of the US's strength.

Market Predictions

With all the negativity out of the way, let's look at the positive future of the company. From 2025, to 2034, the AI semiconductor market is expected to grow at a CAGR of 15.23%, rising from $56.42B to $232.85B.

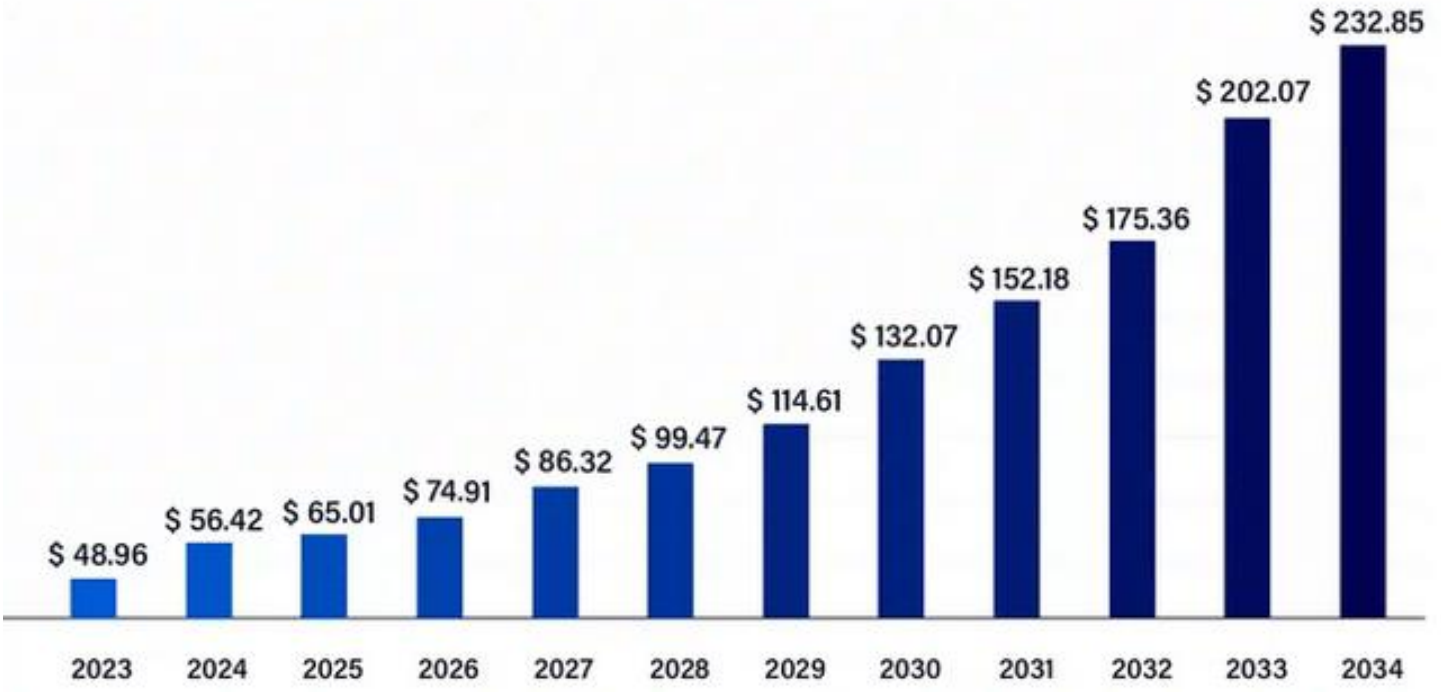

Semiconductors in AI Projected Growth (precedenceresearch.com)

As I mentioned before, Lam is the leading company in Etch, Deposition and Cleaning technologies, which play a fundamental role in the creation of HBM (High Bandwidth Memory) chips, which play a curtail role in any advanced AI. Essentially, the smaller and more powerful the chip, the better, but to create chips like this, more steps are required, and that directly benefits Lam as it is an area they focus on.

The company's HBM revenue is growing at astronomical rates, growing by more than 50.00% from Q1 '25, and expected, by the management, to grow twice the corporate average at least until 2026. My further bullish sentiment is supported by the chip size always getting smaller, As I already mentioned, the creation of smaller chips requires more steps, allowing Lam to charge more for their product and potentially raise their margins.

Q2 '25 Earnings

Before looking through the long-term fundamentals of the company, let's take a look at their most recent earnings. The company beat their non-GAAP EPS by $0.03, from the analysts' expectations of $0.88. Their revenue was also beaten by $70.00M, expected at $4.31B but reported at $4.38B.

YOY the quarterly revenue has increased from $3.76B to $4.38B, accounting for a 16.45% increase. The YOY net income grew by 24.86% from $954.30M to $1.19B. This is some nice, high growth.

Furthermore, the management predicted higher than expected revenue for Q3, at $4.65B +/- $300.00M, beating the analysts' consensus estimate of $4.33B. EPS is also expected higher, at $1.00 +/- $0.10, than the estimates of $0.88.

So overall, Lam is looking forward to some decent growth, but as I am a value long-term investor, I look for undervalued companies that I could hold for the next 10 years+, so instead of the short-term quarterly reports, let's look at the historical fundamentals using our 9 Pillar Analysis.

9 Pillar Analysis

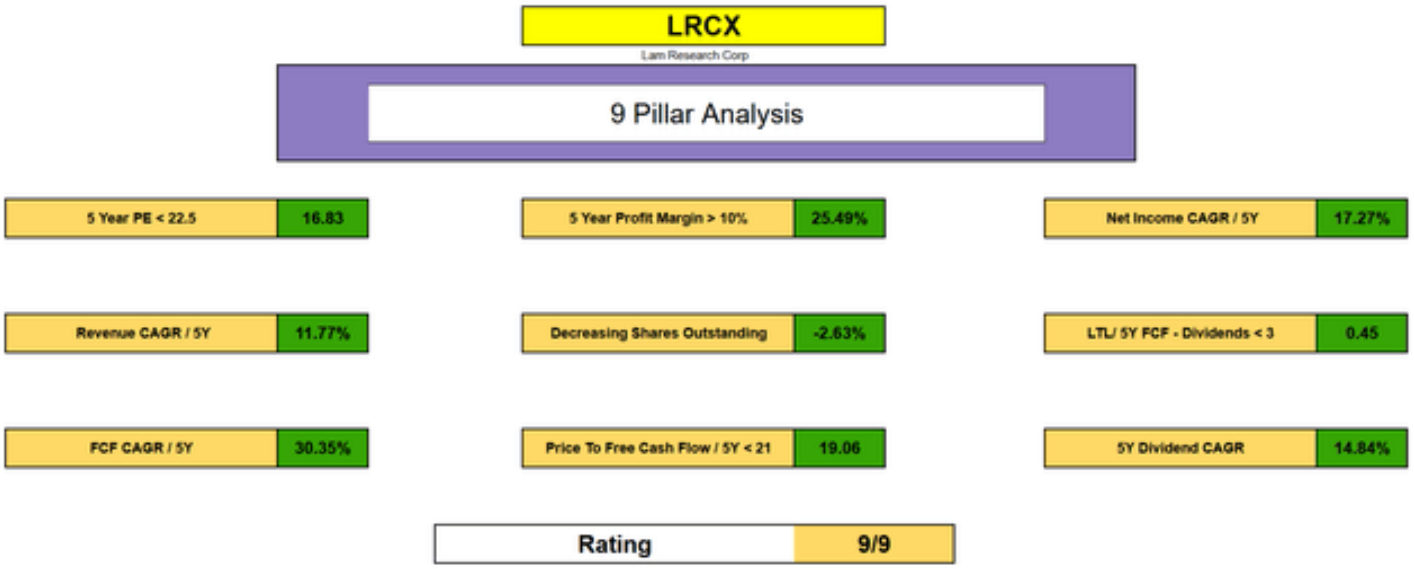

9 Pillar Analysis (FutureRich Valuation Sheets)

For those of you who aren't familiar with the 9 Pillars, let me quickly sum them up. Pillars 1 and 2 represent the 5-year Price-to-Earnings (P/E) and Price-to-Free Cash Flow (P/FCF) ratios, respectively. Pillars 3, 4, and 5 illustrate the growth in revenue, net income, and free cash flow over the past 5 years. Pillar 6 focuses on the 5-year trailing profitability, while Pillar 7 highlights the change in shares outstanding. Pillar 8 tracks dividend growth, and Pillar 9 calculates the ratio of long-term liabilities to 5-year free cash flow minus dividends, which assesses a company's capability to pay down debt within a reasonable period while still paying dividends.

Above you can see that Lam is a 9 Pillar company, as everything is sitting in the green, but let's take it all apart.

Financial growth (Pillars 3, 4, 5)

Over the past 5 years, the company's revenue grew from $10.05B in 2020 to $14.90B in 2024, accounting for a CAGR of 11.77% and an overall increase of 48.21%

Year:	2020	2021	2022	2023	2024
Revenue:	$10,045	$14,626	$17,227	$17,429	$14,905

5Y Revenue Growth(M) (FutureRich Valuation Sheets)

The 5Y net income CAGR is sitting at 17.27%, increasing from $2.25B in 2020 to $3.83B in 2024, again accounting for a very high increase of 70.01%, higher than the revenue increase, which I love as it hints that the company is becoming more profitable. The EPS also increased from $1.51 to $2.90 from 2020 to 2024.

Year:	2020	2021	2022	2023	2024
Net Income:	2,252	3,908	4,605	4,511	3,828

5Y Net Income Growth (FutureRich Valuation Sheets)

Finally, the FCF, my favorite metric, increased by a whopping 121.30% over the past 5Y, accounting for a CAGR of 30.35%. The reason why I love the FCF more than anything, is because it is essentially the blood of the company. Used for

innovation, paying out dividends and buying back shares, which we will get to a little further in the analysis. The FCF/S increased from $2.23 to $3.11 within the same time frame.

Year:	2020	2021	2022	2023	2024
Free Cash Flow:	1,923	3,239	2,554	4,677	4,256

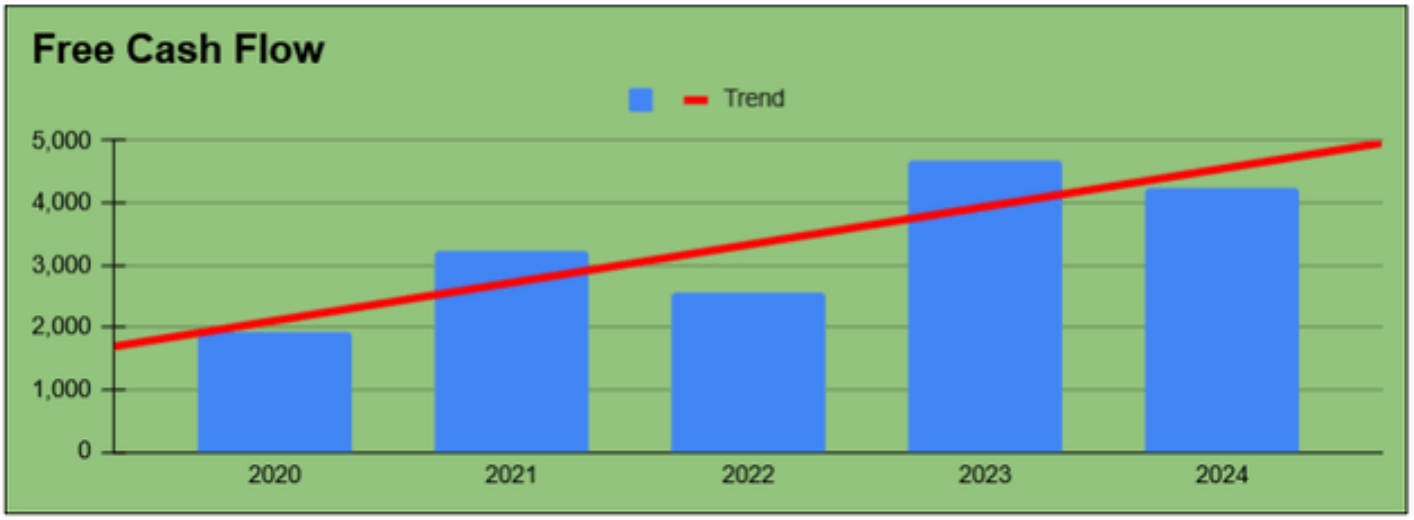

5Y FCF Growth (FutureRich Valuation Sheets)

Overall, there is no doubt that over the past 5Y the company has done really, really well regarding financial growth, and if they grow like this into the future, I wouldn't be surprised if their stock price increased by close to 100.00%, within the next 5 Years.

Valuation (Pillars 1, 2)

Unlike many tech and semiconductor companies, Lam seems to be valued quite fairly. Their 5Y PE is sitting at 16.83 compared to the current of 24.03, which is quite high, but that is because YOY their net income dropped from $4.51B to $3.82B, which, I believe, isn't a long-term trend, and therefore the sudden increase in their cost doesn't bother me that much.

Regarding their P/FCF, as the FCF decreased as well, the current ratio is sitting at 24.55, higher than the 5Y average of 19.06. So overall, through the long term, the company does seem undervalued, even though in 2024 they might have become more expensive for that year particularly.

Dividends (Pillar 8)

It might not be obvious at first glance, but Lam is a great company for long-term dividend investors. They have been increasing their dividends consecutively over the past 10 years, much longer than what is the sector average of 1.6. Their 5Y FCF payout ratio is also sitting nice and low at 24.24% leaving enough FCF for innovation, dividend growth, share buybacks, and paying down debt. Making the company's dividends very safe, and with a 5Y dividend CAGR of 14.84%, it truly is a gem for long-term dividend investors.

Year:	2020	2021	2022	2023	2024
Dividend:	0.46	0.52	0.60	0.69	0.80

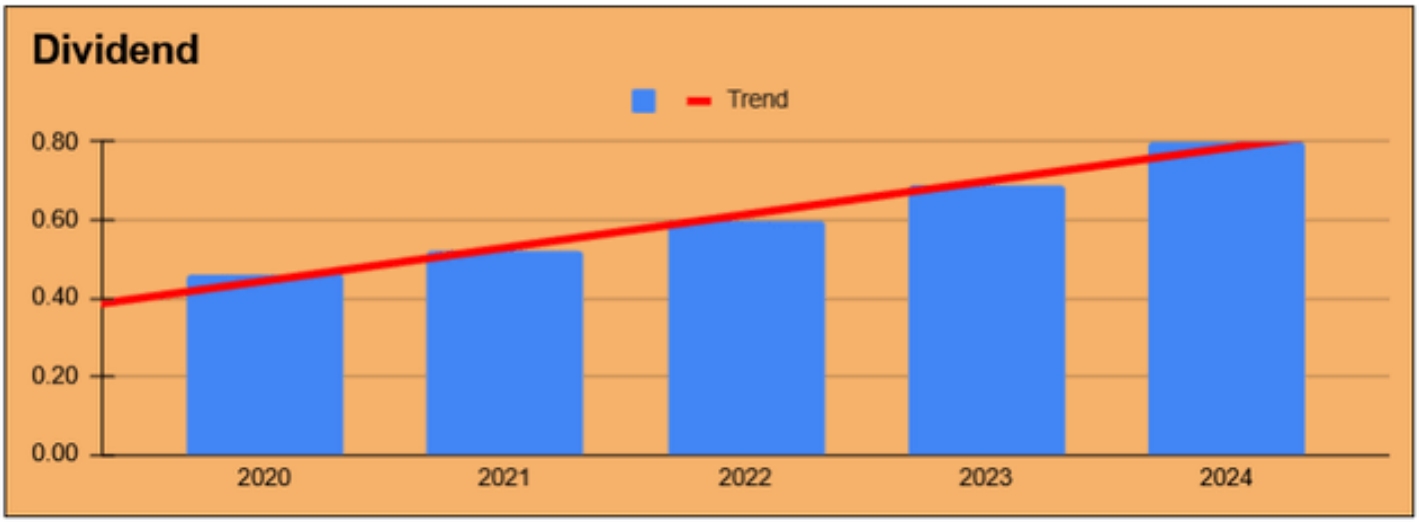

5Y Dividend Growth (FutureRich Valuation Sheets)

Further Pillars

Something I further want to mention is that the company's profit margin has been growing steadily at a 5Y CAGR of 3.68%, increasing from 22.42% in 2020 to 25.68% in 2024. I also love the fact that the company has been repurchasing a lot of shares at an average rate of 2.63% YOY over the past 5Y, allowing for the company's value to grow.

Year:	2020	2021	2022	2023	2024
Shares Outstanding:	1491	1453	1406	1358	1320

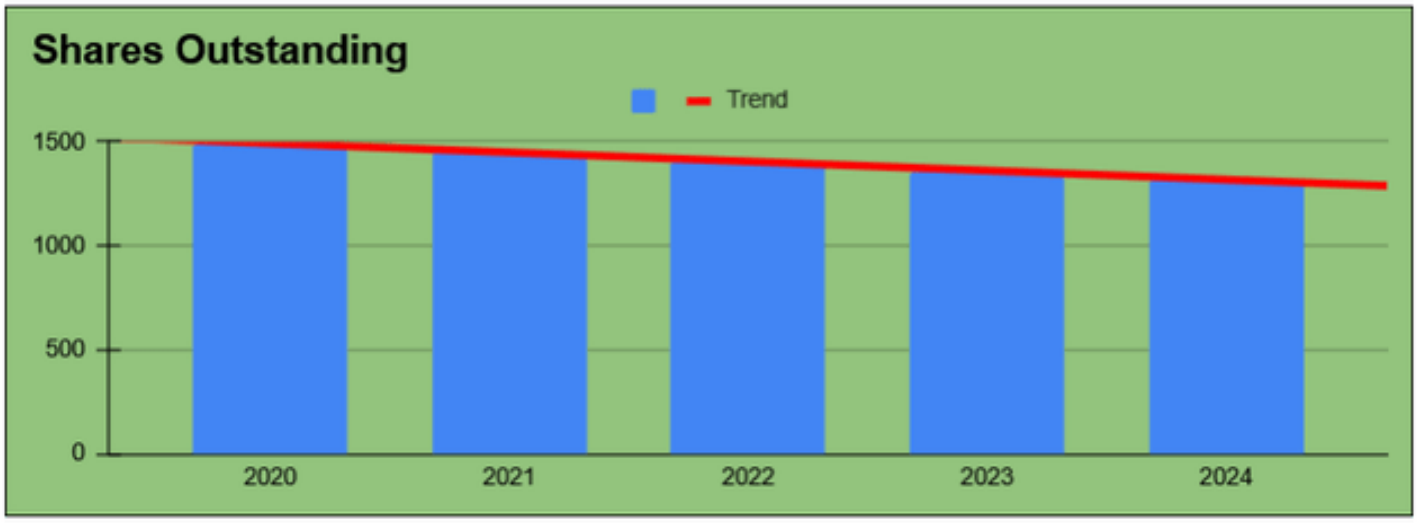

5Y Shares Outstanding Growth (FutureRich Valuation Sheets)

Finally, before calculating the intrinsic value, let's take a look at the last pillar I deem very important, and that is Pillar 9, which shows the debt levels of the company. The LTL/5Y FCF - Dividends is sitting at only 0.45, with the current ratio

sitting nice and high at a 5Y average of 1.67. These two metrics tell me that Lam generates enough FCF to cover all their LTL while paying out dividends in about 2.20 years and has enough short-term cash to pay off its short-term debt.

Summing up the 9 Pillars, it is quite clear that the company has been doing extraordinarily well, and with the growing demand for its product, which we went through at the start of the analysis, in my eyes, they are set to grow at their current rate or even faster into the next years.

Intrinsic Value

In every one of my analyses, when possible, I like to calculate the intrinsic value of a company to figure out at which price I should start buying. Today, we will have Lam undergo all of our 4 Intrinsic value models.

Graham's Formula

Using this formula, we judge the company's value based on their 5Y EPS and their future 5Y EPS growth. After plugging in the EPS YOY growth estimate at 19.67% and combining it with the 5Y average EPS of $3.02, I got an intrinsic value of $103.69, 31.16% higher than the current price of $79.05

Discounted Cash Flow Model (DCF)

In this model, we predict the long-term YOY FCF growth of the company for the next 10 years +. I predicted a YOY growth of 11.00% and combined that with a discount rate of 8.33% and average perpetual growth of 2.50% to get an intrinsic value of $96.45.

Multiples Valuation Model

Here I compare the current PE of Lam to its competitors, so AMAT and ASML. Taking these two companies averaged out PE of 30.11, and multiplying by the EPS of Lam sitting at $3.29, I get an intrinsic value of $99.06.

Dividend Discount Model

Finally, for the alternative DDM I will predict the YOY dividend growth at 14.00%, and combined with a discount rate of 8.33% and the analysis price estimate for next year of $97.30, get an intrinsic value of $107.11.

Ultimate Valuation

Averaging the 4 values in our Ultimate Valuation, we will get an intrinsic value of $101.58, 28.50% higher than the current price of $79.05. This means that with a 10.00% margin of safety, we will get an acceptable buy price at $91.42, well above the current trading price, making the company a BUY based on its intrinsic value.

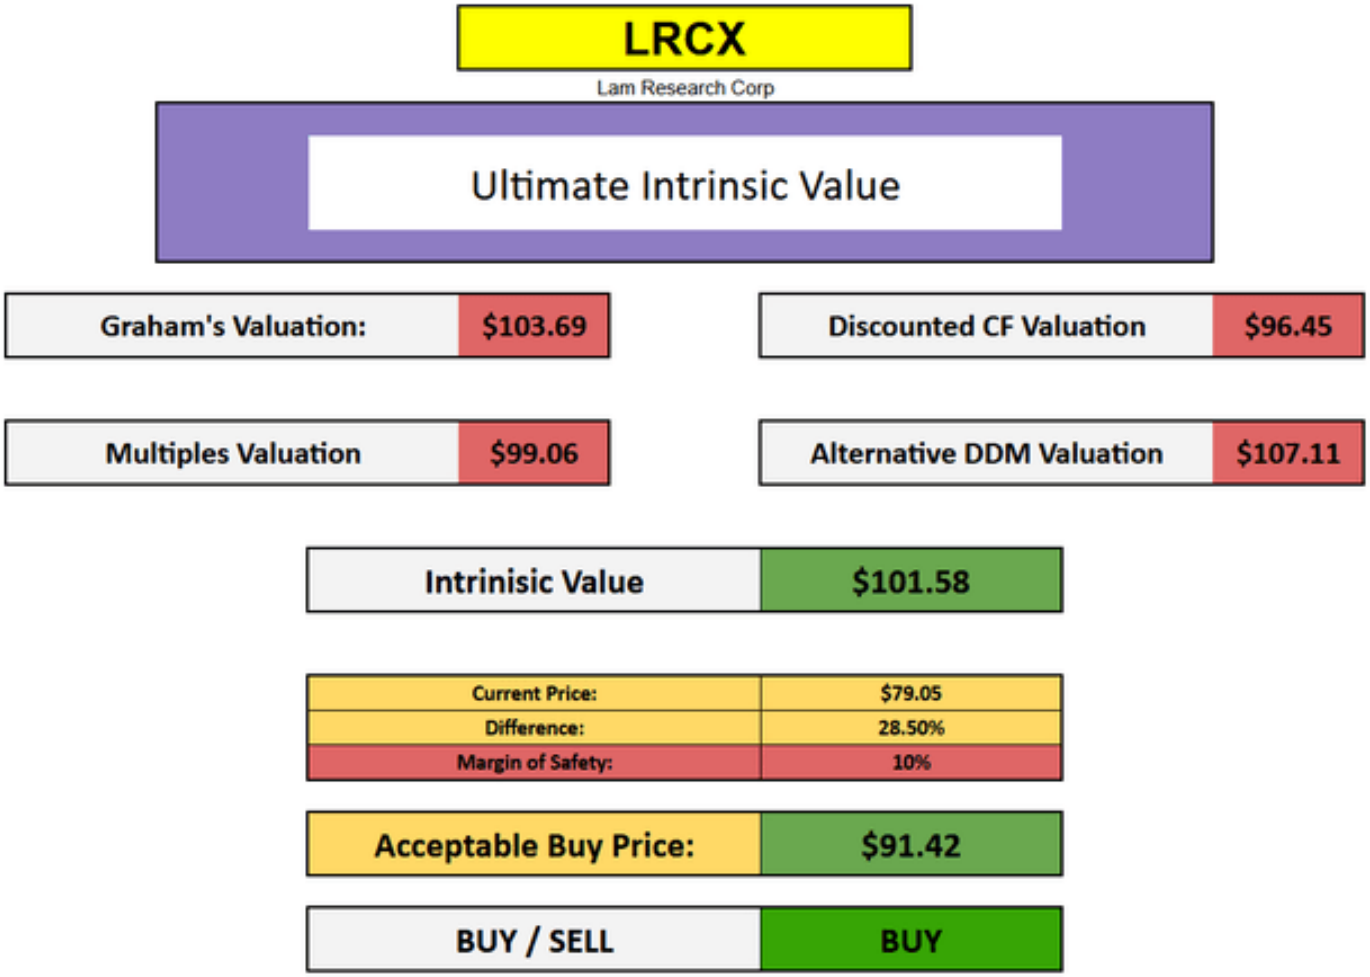

Ultimate Valuation (FutureRich Valuation Sheets)

Conclusion

In conclusion, I have highlighted both the risks and positives that come with investing in the company. Despite the risks from tariff wars, I believe that Lam Research's future potential comes mainly from the predicted high growth in the AI and technology sector, their undervaluation based on our models and great growth track history make them a compelling BUY.

AI-Driven Semiconductor Growth

March 5,2025 VanEck

Summary

- AI model specialization and monetization are accelerating, reshaping semiconductor demand.
- Major players like AWS, Meta, and Microsoft are rapidly optimizing their AI compute strategies, influencing semiconductor investment and CAPEX allocation.
- Nvidia remains dominant, but the evolving AI landscape is creating new demand across GPUs, ASICs, and custom silicon, benefiting a wider range of semiconductor players.

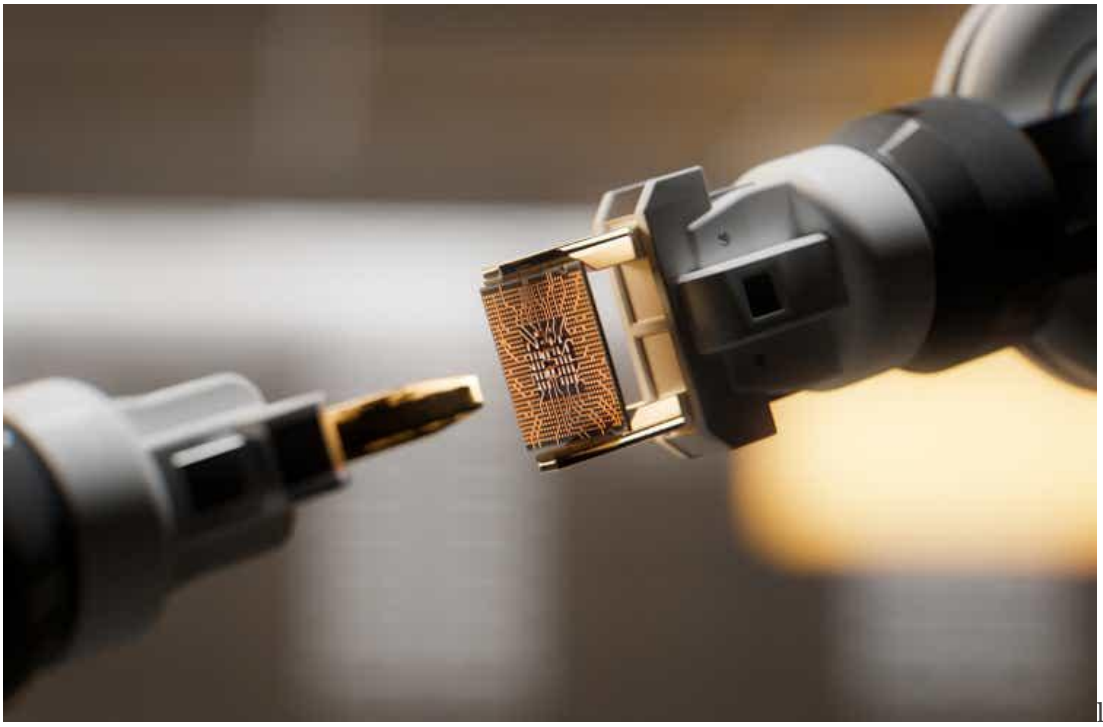
luza studio

In our latest semiconductor outlook discussion, we broke down the latest shifts in AI-driven semiconductor demand, Nvidia's evolving role, and what these trends mean for investors. Here are our top three takeaways:

AI model development is moving faster than expected: AI model specialization and monetization are accelerating, reshaping semiconductor demand.

Hyperscalers are restructuring AI infrastructure: Major players like AWS, Meta, and Microsoft are rapidly optimizing their AI compute strategies, influencing semiconductor investment and CAPEX allocation.

The semiconductor opportunity continues to broaden: Nvidia remains dominant, but the evolving AI landscape is creating new demand across GPUs, ASICs, and custom silicon, benefiting a wider range of semiconductor players.

1. AI Models Are Scaling and Optimizing at the Same Time

The conversation started with the recent AI breakthroughs from DeepSeek and X AI/Grok. Initially, DeepSeek's efficiency gains sparked concerns that AI compute scaling might slow. However, Nvidia's latest results reaffirm that scaling laws are still intact, particularly for US-based operators. The key takeaway? AI still requires significant semiconductor hardware, whether through scaling (GPUS) or optimization (custom silicon and ASICs).

2. Hyperscalers Are Monetizing AI Faster Than Expected

AI monetization has moved beyond expectations, with Meta exceeding forecasts and OpenAI, Microsoft (MSFT), and AWS carving out distinct AI business models. Rather than competing for the same customers, each is building its own lane, reinforcing demand for different semiconductor solutions. Faster monetization likely translates to higher overall CAPEX spending, benefiting the semiconductor supply chain.

3. Reevaluating Nvidia's Position in a Broader Semi Landscape

Nvidia remains the dominant AI chip provider, but the discussion highlighted how AI infrastructure is being unbundled. Hyperscalers are ramping up custom silicon efforts (Amazon's Tranium, Google's TPU, Meta's in-house chips), and market share definitions are shifting as AI compute needs fragment. While Nvidia remains a major player, the broader

semiconductor ecosystem is expanding, creating more opportunities across different chip architectures. The VanEck Semiconductor ETF (SMH) and the VanEck Fabless Semiconductor ETF (SMHX) both offer a diversified ETF approach to capture opportunities across the entire industry.

4. Case Study: ARM and the Acceleration of AI at the Edge

One of the biggest surprises of the quarter was how quickly AI is being integrated into mobile devices. ARM's licensing revenue has surged as companies like Apple, Qualcomm (QCOM), and MediaTek (OTCPK:MDTTF) adopt its ARMv9 architecture to power AI workloads at the edge. This rapid adoption suggests AI compute is moving beyond just data centers, creating a new growth driver for semiconductors outside traditional GPU markets.

5. The Geopolitical Factor: Intel, TSMC, and US-China Dynamics

The discussion also touched on Intel's struggles, TSMC's strategic position, and the long-term implications of US semiconductor policy. With China accelerating its domestic semiconductor efforts and the US relying heavily on Taiwan-based manufacturing, the semiconductor supply chain remains a key area of geopolitical focus.

Final Thoughts from the Past Quarter

The pace of AI innovation continues to exceed expectations, reinforcing demand across various segments of the semiconductor industry. While Nvidia remains a leader, the broader opportunity in AI-driven semiconductors is expanding beyond just GPUs, creating new investment considerations.

As always, our focus remains on mapping these industry shifts as they happen—understanding not just where the market is today but how it is reflexively evolving.

Stargate Project to require 64,000 Nvidia GPUs by 2026: report

March 07, 2025 Chris Ciaccia

- The artificial intelligence-focused Stargate Project announced by President Trump in January is likely to require around 64,000 Nvidia (NASDAQ:NVDA) graphics processing units by 2026, Bloomberg reported.
- The initiative, which is being led by OpenAI, Oracle (NYSE:ORCL) and SoftBank (OTCPK:SFTBY), will need 64,000 of Nvidia's GB200 GPUs by the end of next year, the news outlet added, citing a person familiar with the matter. There will be an initial tranche of 16,000 GB200 GPUs sent to the project by this summer.
- The aforementioned GPU figures are just for one data center for a single customer, the news outlet explained further.
- Nvidia, OpenAI, SoftBank and Oracle did not immediately respond to a request for comment from Seeking Alpha.
- Following the January announcement, OpenAI, which is largely backed by Microsoft (MSFT), said the project would deploy $100B immediately, with a total of $500B spent over four years.
- Arm (ARM), along with Microsoft and Nvidia, are also key partners of the project.
- The build out of Stargate is currently underway, starting in Texas, and there may be as many as 10 data center sites around the country for its needs

TSMC's $100B U.S. Investment Leaves Intel On The Sidelines

March 4, 2025 Oliver Rodzianko

Summary

- TSMC's $100B U.S. expansion secures semiconductor supply chains, strengthens U.S. tech leadership, and mitigates geopolitical risk, establishing Arizona as the new hub for advanced chip manufacturing.
- With projected 30% EPS growth in FY 2025 and 20% in FY 2026, TSMC trades near its five-year average valuation, offering a compelling risk-reward profile. My 2026 PT implies a 27.78% upside.
- This investment de-risks TSMC's geographic exposure, enhances AI and robotics development, and cements its dominance as the world's most advanced semiconductor foundry.

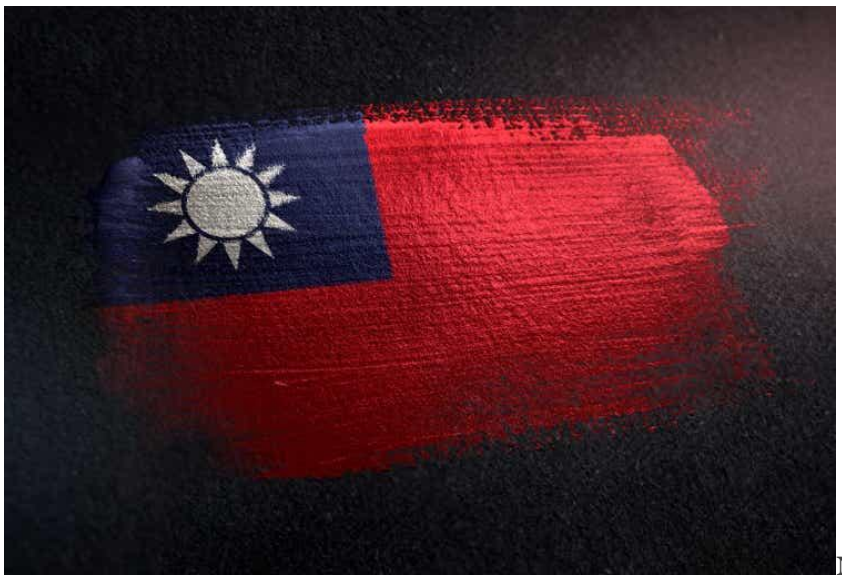
NatanaelGinting/iStock via Getty Images

TSMC (NYSE:TSM) today announced its intention to expand its investment in advanced semiconductor manufacturing in the United States by an additional $100 billion. Building on the company's ongoing $65 billion investment in its advanced semiconductor manufacturing operations in Phoenix, Arizona, TSMC's total investment in the U.S. is expected to reach US$165 billion. The expansion includes plans for three new fabrication plants, two advanced packaging facilities and a major R&D team center, solidifying this project as the largest single foreign direct investment in U.S. history. — TSMC's $100B U.S. Investment News Announcement

Of all the workings of the Trump administration to date, this is by far the most consequential for the world's long-term geopolitical security. There is no better company than Taiwan Semiconductor Manufacturing Company Limited aka TSMC to enable America to diversify the advanced semiconductor supply chain away from Taiwan and the economic and military threat from China. I covered this threat in depth in my last analysis of TSMC.

While international relations between China and the U.S. are likely to develop diplomatically, it is still of critical geopolitical importance for supply chains to be immunized from ideological tensions between Western democratic values and Eastern communism. As such, the $100 billion U.S. investment from TSMC is bullish for global economic stability, which in turn is bullish for Big Tech and the wider stock market in general. Arguably the greatest benefactor from the newly announced deal is TSMC itself, which will benefit from closer relationships with market-leading American technology companies and more revenue security from much greater diversification into the West.

Deal Analysis

The enormous TSMC deal contains plans for three new state-of-the-art fabs, two advanced packaging facilities, and one major R&D center in the United States. To illustrate why the packaging facilities are so crucial to bring to the United States, consider that Nvidia (NVDA) is projected to take 70% of TSMC's 2025 Chip-on-Wafer-on-Substrate-L

packaging capacity. In addition, by establishing a U.S. R&D team, TSMC can collaborate more closely with American Big Tech companies and U.S. academics on cutting-edge research.

This infrastructure investment from TSMC opens the path for faster AI and robotics growth as leading American tech firms require the proximity of fabs and packaging for faster production and iteration cycles. Therefore, the initiative is not just a geopolitical play, it is customer-driven.

This deal also ushers in the “Silicon Valley of Chips” in Arizona, with a complete ecosystem for advanced chip manufacturing in the United States, not just front-end wafer fabrication. The U.S.'s greatest victory here is in advanced semiconductor packaging arriving on U.S. soil. This eliminates a choke point where chips would previously have needed to be sent to Asia for packaging. As of March 2025, TSMC does not have operational packaging facilities in the U.S.

By domesticating U.S. advanced semiconductor chip manufacturing, TSMC will also be avoiding the Trump-era tariffs that are affecting other Asian manufacturers, such as those operating in China. This will put TSMC even further ahead of its Chinese competitors, and likely solidifies TSMC as the most powerful semiconductor company in the world for the indefinite future. The company's moat is remarkably wide, not only technologically with the most advanced manufacturing capabilities, but soon also geographically. With its prior geographic concentration risk now solved through greater U.S. investment, the bull case for TSMC stock contains much less risk and increased reward.

Intel Implications

There have been rumors recently of Intel (INTC) Foundry being acquired by TSMC. However, after the recent announcement with Trump and TSMC Chairman Wei at the White House on Monday, March 3, Intel appears essentially sidelined. This is unfortunate for Intel, but current dynamics dictate an actual solution to geopolitical insecurity. Intel may pose more of an operational burden than an efficient solution at this time.

While Intel has been intent on delivering advanced semiconductor manufacturing to the United States for many years, it has failed to achieve a competitive position comparable to TSMC. Intel is also highly unlikely to achieve this in the future, considering it is now operating with interim management after the ousting of fallen turnaround star Pat Gelsinger. Intel appears to have deep structural issues and misalignments with TSMC's advanced semiconductor manufacturing processes, but the company still has potential to carve a moderately good semiconductor manufacturing position in the United States. Ultimately, however, TSMC Arizona will continue to dominate and become America's flagship advanced semiconductor manufacturer.