KASASV

KOREAN AMERICAN SEMICONDUCTOR ASSOCIATION IN SILICON VALLEY

May 2025

모두를 똑같이 좋아할 수는 없다

리더에게는 각각의 팀원을 위해 옳은 결정을 내리는 일이 가장 어렵다. 팀원중 일부가 리더의 판단과 결정에 반발하거나 화를 낼수도 있다는 것을 예상해야 한다. 그러나 이것은 리더를 외롭게 만드는 많은 일중의 하나이다. 리더는 어려운 최종 결정을 내려야 할 책임을 진다. 반면 다른 대부분은 옆에서 지켜보면서 불평만 늘어 놓기 때문이다. 이처럼 리더는 팀원에게 자신이 그들을 진정으로 걱정하고 사랑한다는 사실을 알게 해야 한다. 모든 조직이 그러하듯이 리더는 팀원과 교육적, 정서적, 정신적으로 깊은 관계를 맺는다.

리더가 팀원에 대한 사랑 없이 조직을 이끌어 가다가는 어느날 문득 뒤를 돌아보고 자신을 따르는 이가 아무도 없음을 알게 될 것이다.

사랑은 경쟁에서, 팀원을 위해, 리더와 팀원이 함께하는 긴 여정에 꼭 필요한 습관이다. – 88 연승의 비밀 , 존 우든

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.

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

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

Linkedin 에 KASASV 를 officially launch 했습니다. 많은 관심 부탁드립니다.

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

Super Micro Computer Results Suggest Nvidia Will Miss Q2 Outlook

April 30, 2025 Tech Stock Pros

Summary

- Nvidia Corporation is going to guide for the July quarter in late May, and we think guidance will miss Wall Street expectations. Then, it'll be time to revisit the stock below $100.

- Super Micro Computer, Inc.'s preliminary results confirm our belief that the July quarter outlook is in trouble, and talks of Trump revamping the AI tier-style restrictions will make things worse for Nvidia stock.
- Huawei is also gaining ground in China market with the Ascend 910C; Jensen's visits to Beijing tell us China is important, too important for Trump to back down.

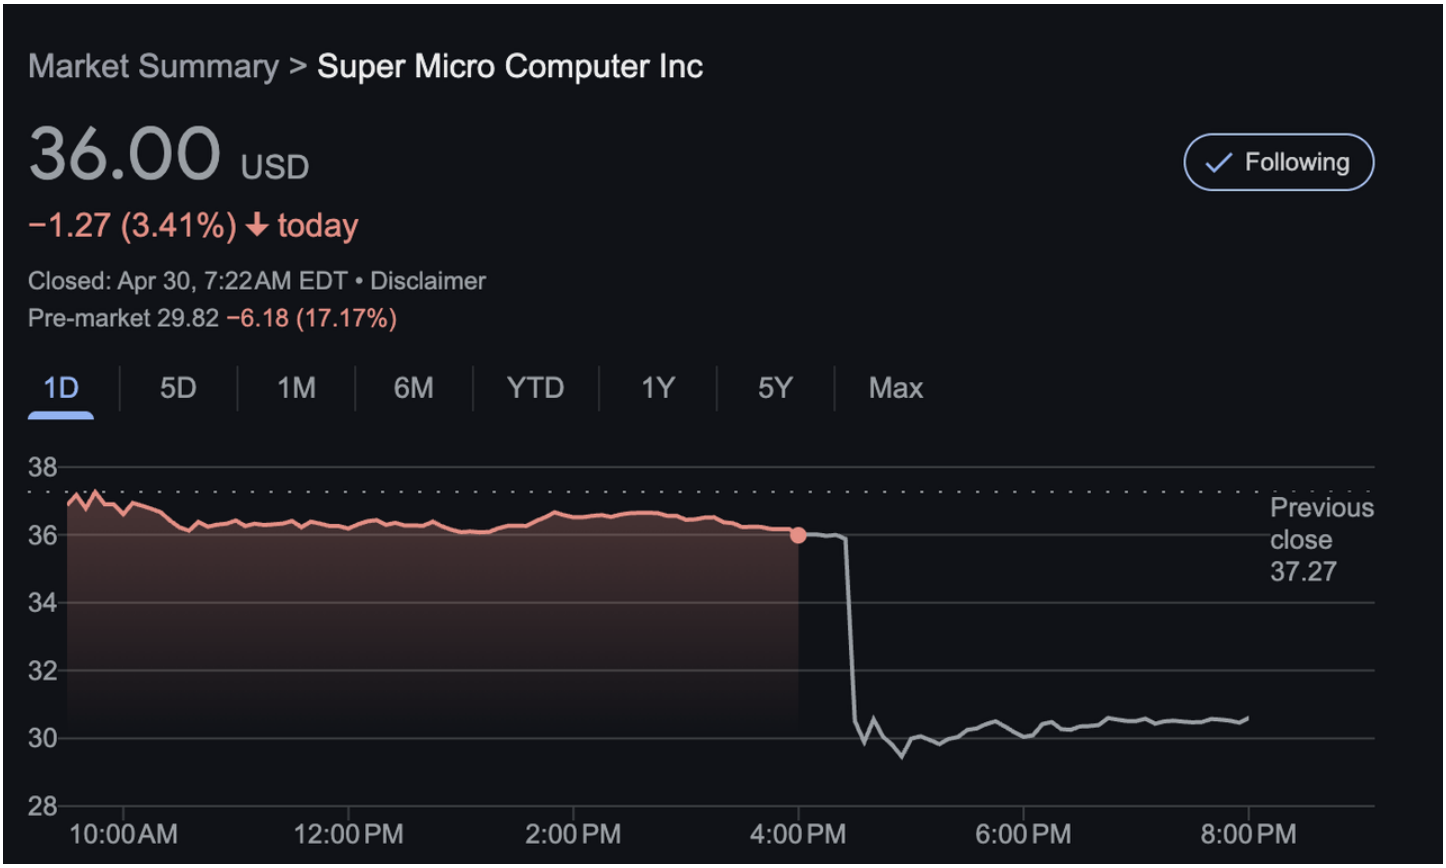

Google

SMCI preliminary results massively trailed Wall Street's expectations, reporting EPS of $0.29 to $0.32, almost half of the estimates for $0.54, and revenue of $4.5B to $4.6B, around $1B lower than the expected $5.5B and lower than their guide in February for revenue in the range of $5B to $6B for the quarter on EPS of $0.46 to $0.62. The interesting part is how management explained it; here's what the statement said:

During Q3, some delayed customer platform decisions moved sales into Q4...higher inventory reserves resulting from older generation products.

This is very telling; the market realizes that what SMCI management is telling us is that demand for Hopper generation is no longer there. In fact, inventory levels are higher on H100 reserves, meaning SMCI can't get H100 off the shelves and is taking a $1B hit because of that, and its gross margins will contract a whopping 220 basis points, too. The cut in sales for the quarter means that, yes, SMCI is going to miss estimates, but it also means that with high inventory levels, they are unlikely to pressure their margins further and place more orders at Nvidia. The fact that Nvidia wrote off $5.5B earlier this month which were supposed to H20 but could also be H100/H200 also confirms that there is something up with demand. Because remember, Nvidia has very high gross margins so when you take that $5.5B number and calculate for it with gross margins then it's a much higher number that they're actually writing off. So, this makes us believe that Nvidia is writing off more than H20s.

This is a signal that 2H demand for Nvidia won't be what the market is expecting it to be. The following chart outlines Wall Street's expectations for Nvidia for the current and July quarter; we expect Nvidia to beat for the current quarter but guide weakly, and that's when we think the negatives will be priced into the stock, and it'll be time to revisit.

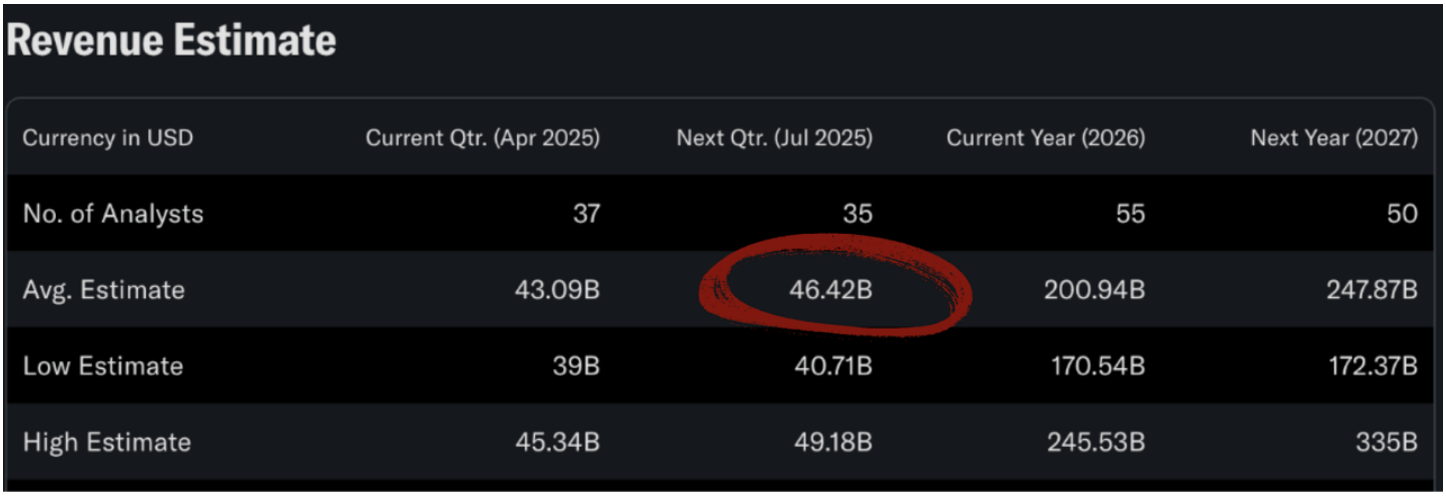

Revenue Estimate

Currency in USD	Current Qtr. (Apr 2025)	Next Qtr. (Jul 2025)	Current Year (2026)	Next Year (2027)
No. of Analysts	37	35	55	50
Avg. Estimate	43.09B	46.42B	200.94B	247.87B
Low Estimate	39B	40.71B	170.54B	172.37B
High Estimate	45.34B	49.18B	245.53B	335B

Yahoo Finance

Let's talk stock price. The stock recovered some of its losses that occurred due to the U.S.'s new licensing requirement for H20, currently trading at $109 versus $96 per share on April 21st. The stock recovered 13% since its low earlier this month, and its RSI followed suit, now at around 50 and even crossing into overbought territory on April 25th. The stock was slightly down earlier this week on the Huawei advancement news this weekend. We cover Nvidia closely, and it's still on our favorite list, but we think the upcoming print has a higher risk profile and recommend investors sell ahead of it to re-enter after the negatives are priced in because the current recovery is on market hope, not fundamentals, and we're here to explain why current hope is badly misplaced.

Trump is turning U.S. Chips into a negotiation card

The newest Reuters update from Trump is that his administration is working on changes to the Biden administration's AI export restrictions; the report stated that the changes include "possibly doing away with its splitting the world into tiers that help determine how many advanced semiconductors a country can obtain." The rule that would be changed is the "Framework for Artificial Intelligence Diffusion," issued by the U.S. Department of Commerce in January, and would be implemented starting mid-May. This rule basically splits the world into three tiers:

- Tier 1: 17 countries plus Taiwan = get unlimited chips
- Tier 2: 120 countries = potential caps on access to AI chips
- Tier 3 (Red tier as we like to call it): China, Russia, Iran, and North Korea = Blocked from chips.

Nothing is certain yet, but it appears that this is a warning that tiers may be removed and replaced with a global licensing regime using "government-to-government agreements," allowing Trump to use U.S. chips as a negotiating card even more effectively. The current way things work is that orders of 1,700 for NVDA H100 don't need to be capped and don't require a license.

Considering that the Trump admin is looking to use chips as more of a negotiation tool, we're not expecting the recently implemented export restrictions to be lifted. Not to mention that Washington has no history of backing down on its AI-related tariffs on China, as seen when it banned the A100, H100 in 2022, or the A800, H800, L40, and L40S in 2023. Nvidia has successfully adapted to U.S. trade rules thrown its way, but we think this time it will be a lot tougher, and the macro backdrop, where AI spending is concerned, won't be there to offset the blow.

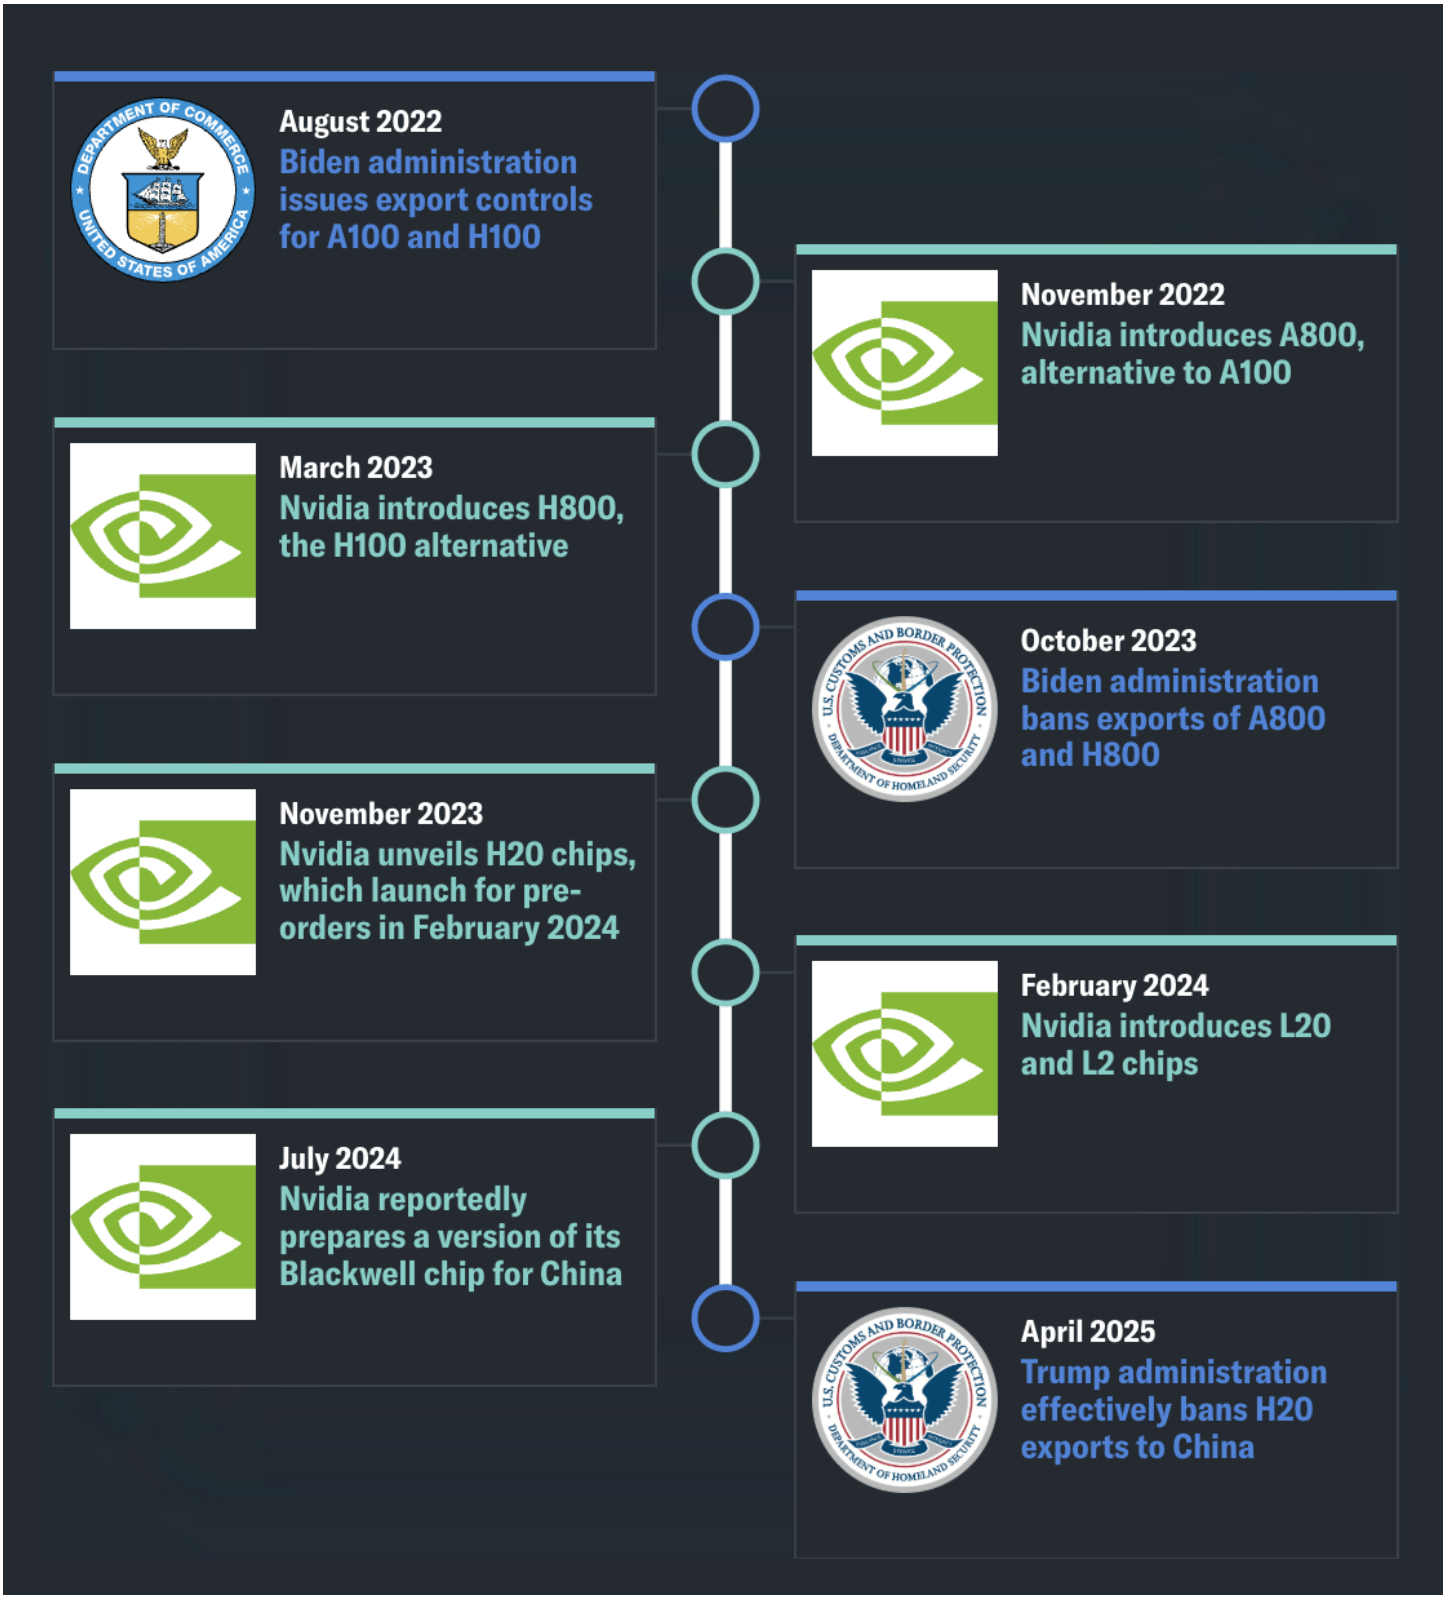

Yahoo Finance

Why is the stock up if the new requirements aren't?

Huawei Tech is reported to be testing its newest AI processors, meant to fill the gap that once belonged to Nvidia's H20, according to a Wall Street Journal report this Sunday, which is the Ascend 910C. Later this year, it's expected to roll out the Ascend 920. Here's the thing: for the most part, Huawei is China's homegrown tech one-stop shop to counter reliance on the U.S. Trump latest efforts to squeeze favorable negotiations out of China by implementing a new licensing

requirement on H20 could backfire for the near-term; Huawei is already reaching out to Chinese tech companies about testing the new Ascend 910C which comes in the line of the company's Ascend 910 series of processors, is supposed to be as powerful as H100 and is going to have ready samples by late May.

That last point is the pain point here for Nvidia; Huawei's plan is to start mass shipments to Chinese customers in May. Even if Ascend 910C doesn't compete meaningfully with Nvidia's chips, which we don't think it will, it could take up the market with pricing power since it is the only option. The Ascend 910C combines two 910B processors into a single package, effectively doubling the compute power and memory capacity of the latter. Reports from Wccftech, Tom's Hardware and TrendForce, cite AGI Hunt (a Chinese media outlet), stating that:

Ascend 910C achieves 60% of NVIDIA H100's performance, delivering strong inference results.

Research from DeepSeek (DEEPSEEK) also tested the chip and confirmed it. Ascend 910C is a fully in house chip that is built with 7nm N+2 process at SMIC (which got about 85% of its total sales from China last year). Part of how Huawei is doing this is through using semi chips made by TSMC (TSM) for a Chinese-based company named Sophgo, which has received 3M custom-designed chips from TSMC over the past couple of years; the U.S. is now tightening regulations around where TSMC's chips go, so this is why Huawei is unlikely to be able to scale this or maintain this in the longer term. Nonetheless, this is significant not because it's the newest high-end tech, as China was already banned from buying H100 since 2022 and has never, to our knowledge, obtained any of the Blackwell series legally. This is a significant news update because it means that even while Nvidia will continue to dominate the global AI processor market, it is losing that market in China.

Jensen told us just how important the Chinese market is to Nvidia by visiting China the week the new restrictions became public knowledge, to "reiterated the company's commitment to the Chinese market." So, while Nvidia's success does not hang on the Chinese market, with the bulk of Nvidia's revenue coming from the tier 1 players in North America, at 47% of total sales, it is not irrelevant.

While we're updating the latest news about the China-Nvidia situation in this article, there is also movement on the second part of our sell thesis that had to do with a potential pullback in AI spend from tier 1 players. When we last wrote on Nvidia, Microsoft (MSFT) was the only company that was showing signs of pulling back on data center leasing, but since then, Amazon (AMZN) has jumped on board with Wells Fargo reporting that AWS was passing some data center discussions.

The bottom line:

Nvidia is the greatest AI company in the world, but even the best go through ups and downs; as far as we're concerned, Nvidia is headed toward a down on its Q2 outlook.

Semiconductor Industry Updates: Tightening Tariffs And CHIPS Act Changes

April 18, 2025 VanEck

Summary

- Tariffs continued under the Section 301 framework now more clearly target semiconductors, amplifying decoupling between U.S. and China.
- Nvidia unveiled its Blackwell Ultra and Vera Rubin platforms, while Micron saw strong demand for high-bandwidth AI memory.
- Intel Investor sentiment improved on news of leadership changes and focus on next-gen 18A process.
- Trump’s March 2025 executive order introduces oversight and potential delays via the U.S. Investment Accelerator, sparking debate over efficiency vs. disruption.

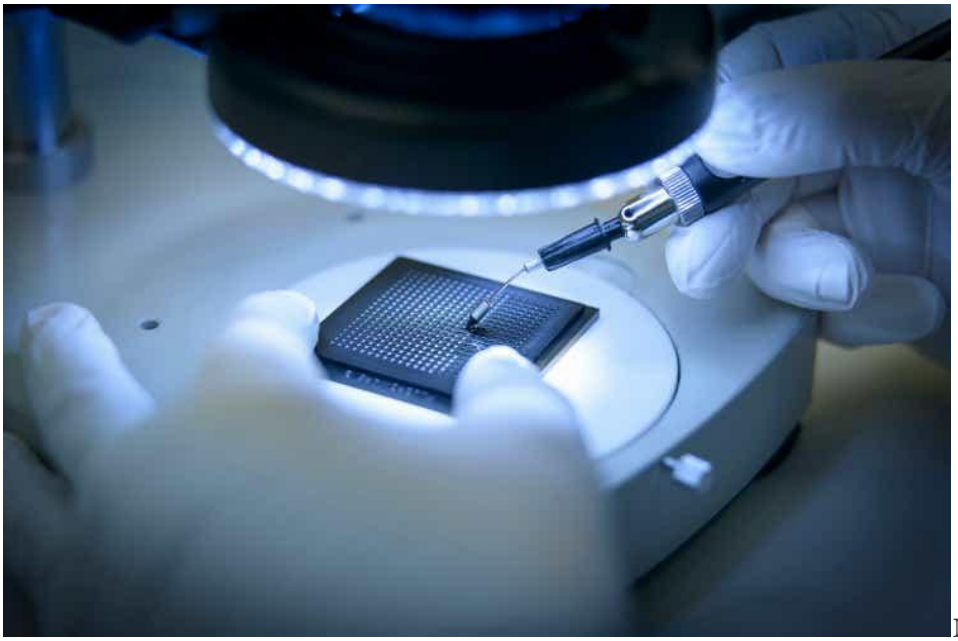
Monty Rakusen

Semiconductor companies continue to drive innovation, while the sector faces policy and supply chain shifts, including tariff expansion and funding uncertainty from Trump's CHIPS Act overhaul.

3 Key Takeaways This Month

- Tariff Pressures Rise: Tariffs continued under the Section 301 framework now more clearly target semiconductors, amplifying decoupling between U.S. and China.
- CHIPS Act Gets a Shake-Up: Trump's March 2025 executive order introduces oversight and potential delays via the U.S. Investment Accelerator, sparking debate over efficiency vs. disruption.
- Global Realignment Accelerates: Friend-shoring trends strengthen across Japan, the Netherlands, and the U.S., with critical projects facing both opportunity and skilled labor constraints.

Semiconductor Industry Snapshot

March Recap Highlights

- Nvidia (NVDA) unveiled its Blackwell Ultra and Vera Rubin platforms, while Micron (MU) saw strong demand for high-bandwidth AI memory.
- Intel (INTC) Investor sentiment improved on news of leadership changes and focus on next-gen 18A process.
- Despite geopolitical tensions, TSMC (TSM) achieved milestones in its U.S. fab expansion.
- Caution from ASML (ASML) reflected weaker demand in non-AI segments like automotive and industrial.
- Despite innovation, investor reactions were sensitive and sometimes muted across the board.

The semiconductor sector in April 2025 remains a case study in duality—booming innovation amid intensifying geopolitical friction. U.S. tariffs on Chinese semiconductors, originally rooted in Biden-era policies, are being reinforced and expanded, further fragmenting global chip ecosystems.

Meanwhile, the CHIPS and Science Act continues to shape domestic momentum, though a major shift occurred in March when President Trump signed an executive order establishing the U.S. Investment Accelerator. This agency is tasked with reviewing CHIPS Act awards for cost-efficiency and performance. While some see this as a streamlining opportunity, others warn of uncertainty and delays—especially at critical sites like TSMC Arizona and Samsung Texas.

Amid these developments, we're also seeing broader diversification beyond AI. Micron's memory investments are increasingly geared toward IoT, and Samsung's logic production targets auto and industrial verticals. These moves reinforce the notion that semiconductor strength depends on more than AI growth alone.

Top Semiconductor Stories

Tariff Expansion Signals Strategic Pressure

Tariffs under the Section 301 framework are being applied more aggressively to semiconductors, reinforcing the divide between U.S. and Chinese chip ecosystems. History shows that companies like Intel and GlobalFoundries tend to weather these policy changes more smoothly than peers with deeper exposure to China. Investors are watching closely for signs of retaliation, including new material export restrictions from China.

CHIPS Act Oversight Reshapes Federal Funding

Following over $30 billion in announced CHIPS Act awards in 2024, the newly formed U.S. Investment Accelerator may renegotiate or reassess key projects. The aim: improve efficiency by loosening prior constraints like DEI and union stipulations. While this could unlock faster execution for some, it also raises questions about delay risks and private sector confidence. Watch for timeline updates on TSMC Arizona, Samsung Texas, and Micron's memory plants in Idaho and New York.

Workforce Capacity Remains a Bottleneck

CHIPS Act-backed projects have created over 115,000 new jobs—but filling them is proving difficult. A majority of roles require four-year degrees, and technical workforce readiness is lagging. Companies like Micron and Intel are expanding training pipelines, but short-term challenges may delay project milestones.

Sector Headwinds & Tailwinds

What's Working:

- Micron: Memory initiatives diversify beyond AI, with tailwinds from IoT and edge computing demand.
- Samsung: Strategic pivot into automotive logic chips aligns with long-term sectoral expansion.
- GlobalFoundries: U.S.-centric operations offer stability amid policy shifts.

What's Challenged:

- TSMC: Oversight from the U.S. Investment Accelerator introduces potential delays to Arizona expansion.
- ON Semiconductor & Qorvo: Tariff-related exposure to China could drive near-term volatility.
- Workforce Readiness: Labor shortages may stall CHIPS Act implementation timelines.

Looking Ahead

Policy remains a powerful force shaping the sector. Upcoming decisions from the U.S. Investment Accelerator, further clarification on China's potential countermeasures (e.g., gallium, antimony), and corporate strategy shifts around foundry partnerships will all influence how the semiconductor industry evolves this quarter.

We'll continue tracking both innovation and policy—as each is now inseparable from the other in semiconductor investing.

AI Thrills and Geopolitical Spills

3 Key Takeaways in March

- AI Momentum is Real: Nvidia's new AI processors and Micron’s memory surge highlight continued strong growth in AI-driven segments despite broader volatility.
- Geopolitics Matter: TSMC’s progress in Arizona amid heightened U.S.-China tensions underscores how strategic moves continue shaping semiconductor supply chains and investor sentiment.
- Volatility Remains High: Even industry leaders like NVIDIA and Intel experience sharp headline-driven market swings, reminding investors to balance excitement with cautious optimism.

Semiconductor Industry Snapshot

As we head deeper into 2025, the semiconductor industry continues to ride a strong wave of growth fueled by the AI revolution and expanding data-center demand. Global chip sales jumped broadly in January alone, reflecting just how transformative AI and high-performance computing have become.

But even with this exciting momentum, we've also seen plenty of turbulence. Geopolitical tensions, tariff uncertainties, and uneven recovery across semiconductor markets mean volatility is never far away. While AI-centric players continue to thrive, traditional consumer markets like PCs and smartphones have been slower to bounce back. It's crucial to remember that while growth is compelling, careful navigation is essential in such a dynamic space.

Top Semiconductor Company Stories

Intel's Leadership Refresh Energizes Investors

Intel made headlines with a significant stock rally in March, sparked by fresh leadership rumors and strategic realignment buzz. Intel’s renewed focus—particularly on its promising 18A chipmaking process and enhanced AI offerings—could signal a shift back toward prominence. As Intel’s narrative unfolds, we’re paying close attention to how these strategic moves translate into market confidence and competitive positioning against heavyweights like TSMC and Nvidia.

TSMC’s Arizona Fab Progress Amid Complex Geopolitics

TSMC hit key milestones this month at their Arizona fab, a significant part of their $165 billion global expansion aimed at bolstering supply chain security. The move comes amid heightened geopolitical stakes between the U.S. and China, positioning TSMC at the heart of strategic conversations. For investors, TSMC’s ability to manage geopolitical complexities will offer critical insight into future supply chain stability and market sentiment.

Nvidia Unveils AI Innovations Amid Investor Caution

Nvidia CEO Jensen Huang kicked off the GPU Technology Conference (GTC) on March 17, 2025 by unveiling major advancements in the company’s AI hardware roadmap. Highlights included the introduction of the highly anticipated Blackwell Ultra processors, scheduled for release in late 2025, promising significant performance leaps from previous generations. NVIDIA also showcased the Vera Rubin AI superchip platform, due in the second half of 2026, featuring advanced CPU and networking capabilities with performance double that of the existing Hopper model.

Additionally, Nvidia announced Dynamo, an open-source system designed to accelerate AI model efficiency and reduce inference costs, underscoring the company's strategic focus on making advanced AI accessible and economically viable. A notable new partnership with General Motors was also announced, positioning Nvidia technology at the forefront of GM's autonomous vehicle and robotics ambitions.

Yet, despite these substantial developments, Nvidia's stock saw a slight dip immediately after the keynote. This reaction shows ongoing market caution and investor sensitivity to competitive pressures, reminding us that even leaders face continuous scrutiny in this dynamic industry.

Micron's AI Memory Strength Stands Out

Micron continues to impress with robust demand and impressive profitability for their high-bandwidth memory chips—vital pieces in AI's rapidly expanding puzzle. With broader memory markets facing headwinds, Micron's success underscores how powerfully AI-focused segments are diverging from traditional sectors. We're watching this closely, as Micron's momentum could offer valuable signals about AI demand trends ahead.

ASML Highlights Broader Market Challenges

ASML recently provided a sobering perspective, flagging softer demand outside the red-hot AI sector. Their cautious outlook on automotive and industrial segments reminds us of semiconductor market bifurcation. ASML's commentary is something we're following closely, as it will help signal broader trends in capital expenditures and industry growth dynamics.

Sector Headwinds and Tailwinds

What's Working:

- Intel: Intel's turnaround strategy is resonating as fresh leadership ideas inject optimism into their long-term story.
- Micron: Micron's strategic bet on AI-focused memory chips has paid off significantly, highlighting smart alignment with secular growth trends.
- TSMC: TSMC's strategic global expansion continues to inspire investor confidence, reinforcing their foundational role in advanced chip production.

What's Challenged:

- Marvell (MRVL): Marvell continues to grapple with slower segments like networking and storage, reflecting how critical strategic focus has become in this market.
- Nvidia: Despite unmatched leadership in AI hardware, Nvidia faces increasing competition and fluctuating market expectations—creating near-term uncertainty to watch closely, as evidenced by initial reactions following their GTC keynote.

Looking Ahead

The next few weeks promise plenty to watch. We'll closely follow how the market continues to digest Nvidia's significant announcements, especially the introduction of Blackwell Ultra and Vera Rubin, and whether investor sentiment stabilizes or shifts. Additionally, evolving U.S.-China tariff discussions and key upcoming quarterly earnings reports will provide critical insights into the broader semiconductor landscape.

Longer-term, we're exploring how geopolitical shifts and strategic repositioning by industry giants will impact investor decisions and sector trajectories. Each month, we'll revisit these themes, tracking progress and identifying new developments together.

Why An Intel-TSMC Marriage Makes Sense

April 4, 2025 Kumquat Research

Summary

- Intel and Taiwan Semiconductor have preliminarily agreed on a joint venture to operate some of Intel's foundries.
- Intel would maintain a significant stake, while TSMC would get a 20% stake in exchange for training and sharing expertise.

- While many details remain uncertain, this arrangement has potential benefits for both companies.

During a crazy 24 hours that looked like it couldn't get a crazier, a quite reputable tech news publication called The Information reported that Intel and TSMC had agreed on an early framework for a joint venture that would operate some of Intel's fabs. Intel would likely provide much of the capital expenditures and would retain a significant stake, TSMC would provide training and expertise in exchange for 20%, and other as yet unnamed US semiconductor companies would have stakes as well, likely in exchange for additional capital and other commitments.

Needless to say, this is an unprecedented shift in the chip manufacturing industry. While on the surface, it appears like there are conflicting interests here that would leave such a deal dead in the water, there have been some developments in recent months and years that appear to have forced this outcome. More on this later.

While this is certainly a major event, it isn't exactly a complete surprise. The Trump administration has been pushing for such an outcome for months, and Reuters reported a few weeks ago that TSMC had even approached Nvidia (NVDA), AMD (AMD), Broadcom (AVGO), and Qualcomm (QCOM) in an effort to pitch their investment in such a JV. Well, it seems that these negotiations have progressed fairly quickly!

Okay, so we've mentioned a lot of companies that all have competing interests here, so I think the best way to unravel potential winners and losers would be to evaluate each company's perspective. Let's start with the downtrodden dog everyone wants a piece of: Intel.

For some quick context, after long being an INTC bear during its foundry struggles amid its battles with AMD a few years ago, I switched to the bull side after the company announced its major restructuring. In my last article at the end of January, I argued the stock was a Buy on a turnaround for its potentially lucrative manufacturing business and the enduring prowess of its chip design segment. That piece can be read here.

I think the benefits for Intel here are apparent: the company can avoid bleeding billions from its foundry business, survive even if its 18A node ends up being less successful than hoped, receive chipmaking knowledge and training from the foremost experts in the field, and get some relief on capital investment now that CHIPS Act funding has likely been sidelined. Exactly how much know-how TSMC will be contributing and whether this includes information on leading-edge nodes is unclear, but I think it's more likely that the company will focus more on getting Intel's nodes to sustainable yields rather than sharing its most confidential and valuable trade secrets.

There are a couple of downsides for Intel too, in my opinion. The main downside is that this JV would essentially signal the end of an independent Intel foundry turnaround and the benefits that would come along with such an accomplishment. Though the odds weren't exactly stellar at this point, it does remove one of the high-ceiling bull scenarios. The second main downside, to me, is that letting TSMC peek under the hood will give it insight into Intel's processes that Intel will not receive into TSMC's. This likely has more short-term implications rather than long-term from a competitive advantage standpoint, but it is still a significant disclosure.

I think TSMC's position here is more nuanced. At first glance, it would seem to be a bad deal for the company -- it has to share knowledge and expertise with its main rival in order to manufacture chips that will likely steal business away from its own factories. In addition, TSMC has already committed $165 billion in US investment so why is it now concentrating even more investment in the country?

TSMC Fab in Phoenix, Arizona (TSMC)

I think the answer is likely more geopolitical than anything. The Trump administration has been pushing TSMC to assist Intel for a while, gutted the CHIPS Act and instead got concessions from the company to invest an additional $100 billion in the US, and has threatened Taiwan with semiconductor tariffs. The US is also Taiwan's major military benefactor when it comes to keeping China from trying to seize the island by force. These factors all exert pressure for TSMC to assist Intel even when from a purely competitive or profit-driven perspective, this JV might not have been pursued.

With that said, I do see some benefits here. First, by owning 20% of this JV, TSMC is hedging its competitive risk -- even if Intel succeeds with 18A and creates a competitive node, its success will be partially owned by the company it would be competing against. Second, the JV allows TSMC to hedge political risk as well since this outcome would likely insulate it from the Trump administration's ire and would go a long way towards securing a tariff exemption. Third, as I mentioned earlier, in exchange for bailing out its spiraling rival, TSMC will get a valuable look into Intel's technologies and facilities that could prove useful in its own designs.

The last perspective here is that of the chip design companies who, in a vacuum, would likely not be too fond of the prospect of TSMC becoming intertwined with its primary rival. Not that Intel was lighting up the industry with design wins for its leading-edge nodes, but more competition is good for buyers whereas consolidation is good for the sellers. However, offering these companies a stake in the JV to operate these foundries mitigates some of those concerns; by investing in management of manufacturing facilities, these companies can likely secure capacity at lower costs with the caveat that there is much work to be done to get production up and running at scale.

Further, in exchange for their stakes in the JV, it's likely that these companies will have to commit to manufacturing at least some of their chips at JV-run plants for a certain number of years in order to get them off the ground and bootstrap the foundry business (à la AMD spinning off GlobalFoundries). While TSMC will own a stake, I still think this move would help to break some of the company's current pricing power leverage over the industry, and companies like Nvidia and Broadcom would love to see that happen.

So could this be the rare win-win-win for all involved? I think so, but perhaps not at equal magnitudes. I think the biggest winner here is Intel, followed by the chip design companies, and then TSMC ekes out a benefit, though mostly geopolitical.

Takeaway

I'll preface this section with the warning that the current market environment is and will likely remain extremely volatile for the foreseeable future. Tariffs, counter-tariffs, and an impending economic slowdown are all weighing on the market and individual stocks will get caught in the crossfire even if there are bullish developments. The following recommendations are for long-term investors only. I don't claim to know what will happen in the short-term so invest at your own risk.

For long-term investors, I think this development adds to the bull thesis for INTC. I have not been a fan of a spin-off since I think Intel's foundry business is one of its main core advantages, so I think this JV idea is a good compromise and could end up being the beginning of the turnaround for which shareholders have been waiting. With TSMC assisting Intel fabs to run efficiently and scale up complicated nodes, the odds of stemming the cash bleed and returning to consistent profitability have increased significantly, in my opinion.

While this agreement is still preliminary, and I think it enhances the company's prospects, I think Intel is well-positioned regardless of the outcome. New CEO Lip-Bu Tan has espoused the foundry as the backbone of Intel, and I think this remains the case with or without the JV.

Tariffs Aren't As Scary For Tesla

April 4, 2025 The Techie

Summary

- Tesla, Inc. is best positioned in the current tariff wars between the EV/auto group, and it has two cars coming this summer: Model Y and FSD.
- Tesla's vertical integration and U.S.-based production should lessen the impact of new tariffs compared to competitors.
- The stock is down around 29% since the start of the year, and I think the current valuation reflects reset expectations, making it a good buying opportunity.
- I'm watching for Tesla's earnings and developments in Canada and the new Lithium Refinery Plant in Texas for future growth.

What I’m trying to say is that the hate train won’t last, and that was apparent with Tesla’s performance post deliveries missing estimates by a margin. However, the stock still shot up over 5% after Trump announced Musk’s departure from the DOGE soon.

I said this once, and I’ll say it again: I think this remains a good chance to keep adding in bunches before the incoming upside due to full self-driving, or FSD, in June and the new Model Y. I think negatives of the plunging Europe sales are now well priced into the stock, so if this month tells the same story, I believe it won’t make a dent on the stock price.

Stock price and valuation:

The stock is now down 29.5% year to date and is trading at $267, down over 44% from its mid-Jan highs of $479. I’m maintaining my buy on Tesla, and while I think there might be some more downside ahead due to the newly and freshly announced tariffs that sent the market on a nosedive Thursday, Tesla is still better positioned to be an outperformer in

FY25 with expectations, so low at current levels. I'll be discussing the tariff effect on Tesla in a second, but dare I say I appreciated Musk's bluntness regarding tariffs late last month.

Wall Street is currently on the fence regarding the stock, with over 13% of Street analysts giving the stock a strong buy, 35.2% a buy, 29.6% a hold, and the rest a sell and a strong sell. Price target, or PT, mean and median have been on a downward trend over the past 3 months as well, with PT median down from $283 in January to $320 currently and PT mean also down from $284 to $318 currently. I would've personally flagged this if the number had stayed the same or maintained an upward trend. I think expectations are now reset.

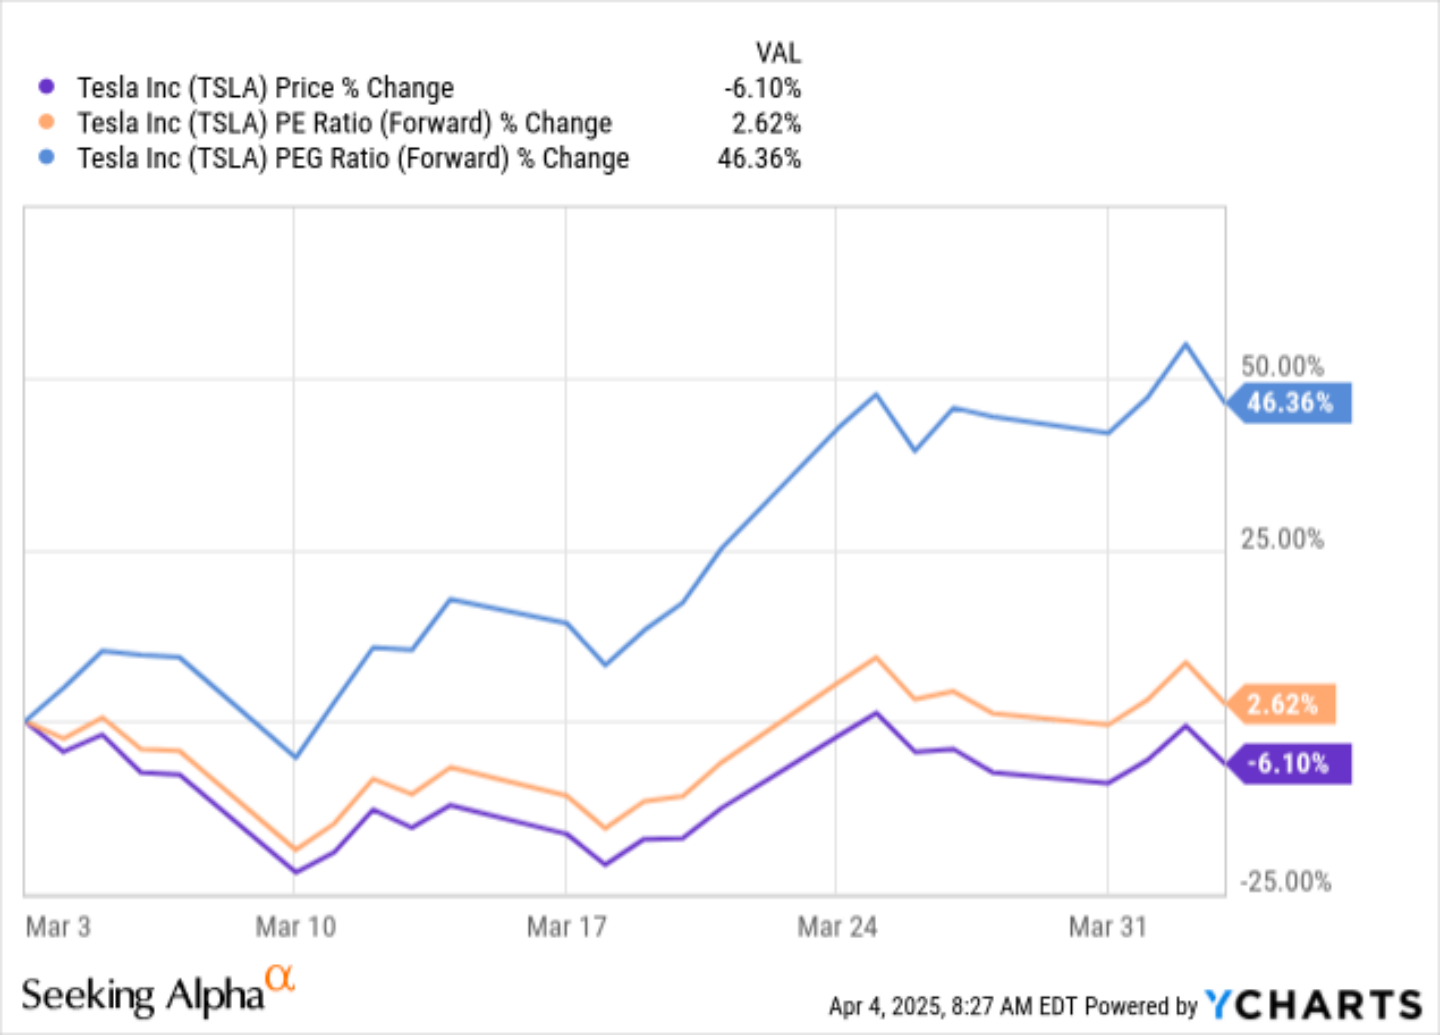

YCbarts

The stock is down over 6% on the one-month chart, while forward P/E was up 2.6% and forward PEG ratio was up 46.3%. I think while the stock performance doesn't reflect Tesla's potential as an AI and robotics company, the PEG and P/E ratio do.

Tariffs:

A creator on X revealed in an infographic of "where cars sold in America are made" that Tesla was the only one that had "100% made in the U.S.," as seen below. Musk responded to this saying: "Important to note that Tesla is NOT unscathed here. The tariff impact on Tesla is still significant."

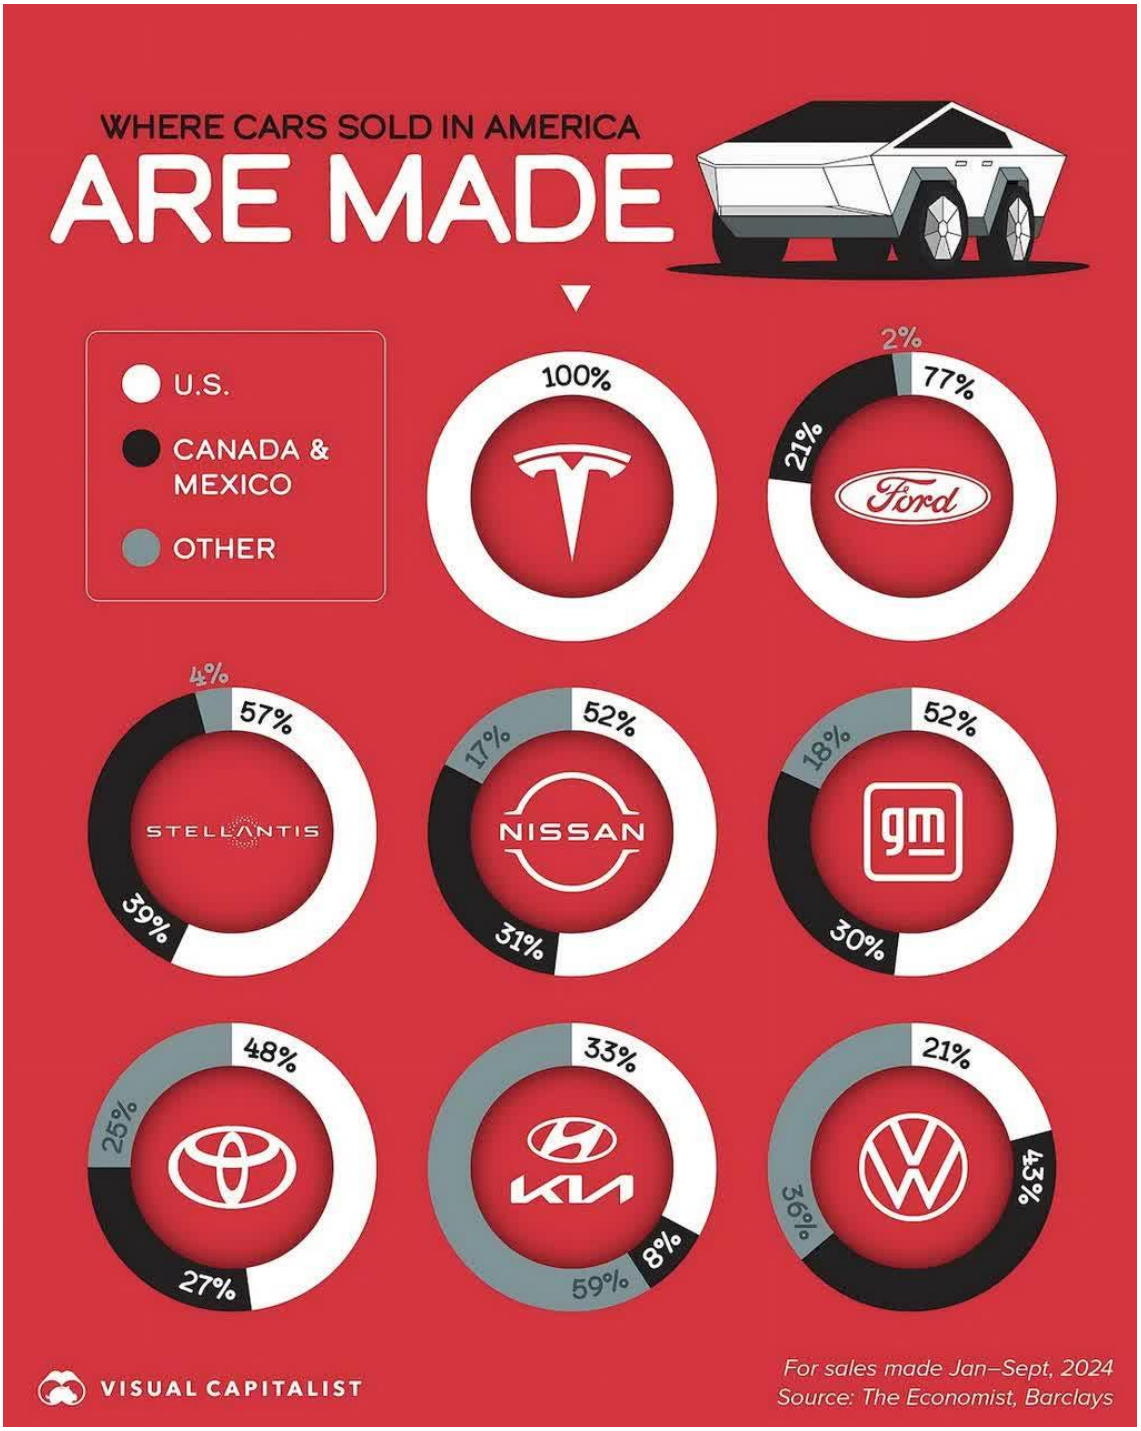

The Economist

So yes, while Tesla might perhaps lose a finger, others are losing limbs. Tesla produces all its North American vehicles in the U.S. at factories in Texas and California, which could lessen the impact of tariffs in comparison to others dependent on manufacturing outside the U.S., like Ford (F), General Motors (GM), Nissan (OTCPK:NSANY), Toyota (TM), Honda (HMC), Hyundai (OTCPK:HYMTF), Rivian (RIVN), and Stellantis (STLA). According to Cox Automotive's director of industry insights, Tesla is better positioned against peers due to its vertical integration, where the company produces the vehicles, the batteries, and so on, putting them in control of the supply chain.

Basically, Tesla has a head start as of now since others still need to move production to the U.S., which not only costs a lot but is also a lengthy process. The fact is, EVs are already more expensive than gasoline-powered vehicles, making the odds of consumers picking them already slim. And now, the tariffs will definitely force companies to raise prices, causing

demand, in my opinion, to be significantly hit. Tesla isn't exempt, but the company has Model Y and Model 3 as the top 2 best-selling EV models in the U.S. last year, all while being significantly more expensive than others, snatching over 28.6% and 14.6% of the market share, respectively, as seen below.

	EV model	2024 sales*	2024 market share
1.	Tesla Model Y	372,613	28.6%
2.	Tesla Model 3	189,903	14.6%
3.	Ford Mustang Mach-E	51,745	4%
4.	Hyundai IONIQ 5	44,400	3.4%
5.	Tesla Cybertruck	38,965	3%
6.	Ford F-150 Lightning	33,510	2.6%
7.	Honda Prologue	33,017	2.5%
8.	Chevrolet Equinox EV	28,874	2.2%
9.	Cadillac Lyriq	28,402	2.2%
10.	Rivian R1S	26,934	2.1%

Electrek

Sure, plans for a cheaper model are now challenged, but I think the tariffs are now a priority, and maneuvering them will be an easier job for Tesla than the others. According to a document from the National Highway Traffic Safety Administration, around 20% to 25% of Tesla's parts are imported from other countries, while the rest of the 60% to 75% are produced in the U.S. and Canada, as seen below.

Vehicle	Pct in US/Canada
Model 3 LR AWD/RWD	75%
Model 3 Performance	70%
Model Y (All Variants)	70%
Cybertruck	65%
Model S	65%
Model X	60%

Not a Tesla App

Back in June, the Tesla Model 3 was on top of the list produced by Cars.com. According to the latest ranking from the Kogod School of Business, the 25 cars built in the U.S. that have the most American parts are the Tesla Model 3 at 87.5%, the Tesla Model Y Long Range, and the base Model Y tied for second place at 85%. Below is a list of the rest.

Rank	Model	Total Domestic Content
1	Tesla Model 3 Performance	87.5%
2 (tie)	Tesla Model Y Long Range	85%
2 (tie)	Tesla Model Y	85%
3	Tesla Cybertruck	82.5%
4 (tie)	Tesla Model S	80%
4 (tie)	Tesla Model X	80%
4 (tie)	Ford Mustang GT automatic transmission	80%
4 (tie)	Ford Mustang GT 5.0-liter	80%
4 (tie)	Ford Mustang GT Coupe Premium	80%
5 (tie)	Honda Passport AWD (all-wheel drive)	76.5%
5 (tie)	Honda Passport Trailsport	76.5%
6 (tie)	Jeep Wrangler Rubicon	76%
6 (tie)	Jeep Wrangler Sahara	76%
7 (tie)	Volkswagen ID.4 AWD 82-kWh	75.5%

Kodod School of Business

Tesla has been an apparent leader in the EV space, and a strong contender in the face of competition in the U.S. and China with players like BYD (OTCPK:BYDDF), Li Auto (LI), XPeng (XPEV), and Nio (NIO) who offer a pricing advantage. I don't want to go off track, as I'll discuss Tesla in China shortly. What I'm trying to say is that though the company is a recognized leader in EVs, it hasn't had any new EV model aside from the Cybertruck since 2020. For a few earnings calls now, Musk has been talking about Tesla's transformation to an AI robotics company, with robotaxis taking the spotlight. I think investors are starting to fully grasp that, and the deliveries miss, as I previously mentioned, yielded no substantial negative reaction from the market because eyes are now somewhere else. Call it the start of something new.

Tesla in China:

Tesla reported deliveries of 336,681, the lowest quarterly performance in over two years, down 13% year over year and missing estimates that have been revised down over the past few months from around 470,000 to 370,000.

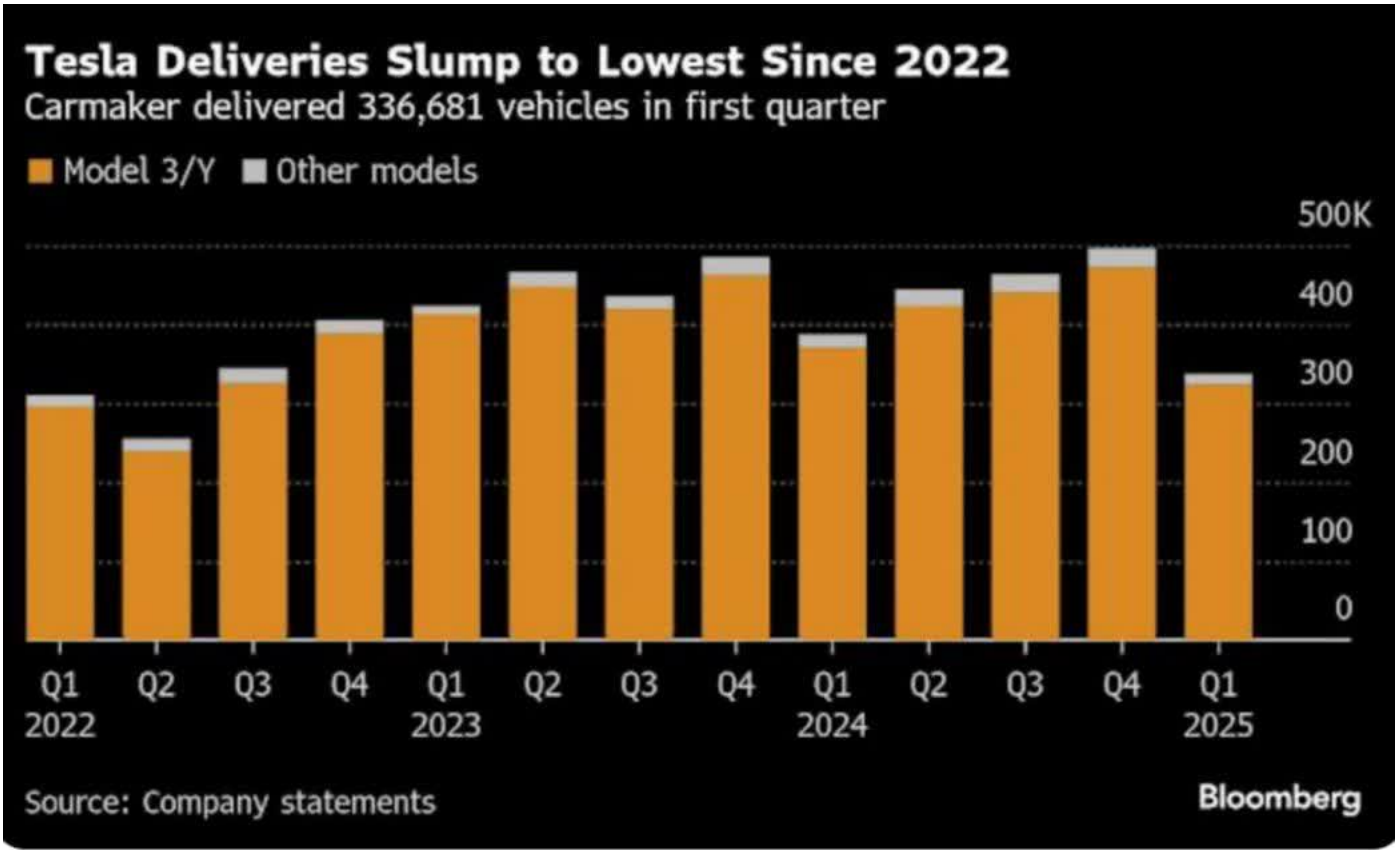

Bloomberg

Tesla dealt with partial shutdowns in some of its factories, and these pauses enabled the company to begin production of a revamped version of the Model Y, which has been very well received in both the U.S. and China (which the company started delivering in late February). I think this should help with better deliveries next quarter.

Tesla's China-insured units have been soaring over the past few weeks, as seen below, which tells me demand there is healthy, despite strong competition.

week	2022	YTD	2023	YTD	2024	YTD	2025	YTD
1	3.680	3.680	2.110	2.110	3.200	3.200	5.500	5.500
2	7.189	10.869	12.654	14.764	7.400	10.600	7.800	13.300
3	3.598	14.467	7.496	22.260	11.700	22.300	10.000	23.300
4	5.074	19.541	3.356	25.616	12.700	35.000	11.600	34.900
5	379	19.920	8.643	34.259	10.600	45.600	1.800	36.700
6	3.601	23.521	6.963	41.222	5.700	51.300	6.200	42.900
7	5.000	28.521	5.913	47.135	2.600	53.900	7.500	50.400
8	13.655	42.176	10.705	57.840	10.800	64.700	6.900	57.300
9	11.317	53.493	13.266	71.106	12.500	77.200	12.400	69.700
10	14.583	68.076	17.032	88.138	13.200	90.400	13.800	83.500
11	15.414	83.490	18.712	106.850	12.300	102.700	15.300	98.800
12	15.860	99.350	15.886	122.736	13.200	115.900	17.400	116.200
13	9.650	109.000	14.275	137.011	17.300	133.200	21.000	137.200
14	128	109.128	6.973	143.984	1.907	135.107		
15	134	109.262	12.500	156.484	6.212	141.319		
16	108	109.370	10.300	166.784	5.168	146.487		
17	1.330	110.700	11.500	178.284	14.800	161.287		
18	3.200	113.900	5.900	184.184	11.000	172.287		
19	1.609	115.509	9.990	194.174	9.800	182.087		
20	2.135	117.644	10.200	204.374	13.800	195.887		
21	1.510	119.154	12.800	217.174	13.200	209.087		
22	9.852	129.006	14.500	231.674	15.000	224.087		
23	18.211	147.217	16.400	248.074	12.000	236.087		
24	17.949	165.166	14.500	262.574	11.700	247.787		
25	18.186	183.352	16.700	279.274	17.500	265.287		
26	14.939	198.291	17.400	296.674	14.100	279.387		
27	162	198.453	3.200	299.874	6.500	285.887		
28	424	198.877	10.000	309.874	11.400	297.287		
29	3.034	201.911	7.700	317.574	10.500	307.787		
30	5.117	207.028	10.600	328.174	13.500	321.287		
31	6.700	213.728	12.800	340.974	12.500	333.787		
32	7.179	220.907	14.000	354.974	15.500	349.287		
33	3.620	224.527	13.900	368.874	14.300	363.587		
34	12.107	236.634	17.000	385.874	14.200	377.787		
35	15.264	251.898	11.800	397.674	14.400	392.187		
36	18.215	270.113	10.700	408.374	16.200	408.387		
37	15.856	285.969	8.500	416.874	15.600	423.987		
38	23.109	309.078	13.500	430.374	13.800	437.787		
39	10.651	319.729	7.200	437.574	22.600	460.387		
40	2.218	321.947	1.000	438.574	1.800	462.187		
41	1.724	323.671	7.500	446.074	8.500	470.687		
42	5.268	328.939	8.100	454.174	13.200	483.887		
43	8.079	337.018	10.800	464.974	10.500	494.387		
44	11.195	348.213	14.000	478.974	15.700	510.087		
45	13.939	362.152	12.700	491.674	17.300	527.387		
46	14.366	376.518	16.300	507.974	17.100	544.487		
47	16.121	392.639	16.700	524.674	16.700	561.187		
48	11.670	404.309	17.400	542.074	18.600	579.787		
49	12.977	417.286	15.400	557.474	21.900	601.687		
50	10.254	427.540	18.300	575.774	18.452	620.139		
51	8.915	436.455	18.500	594.274	17.600	637.739		
52	4.338	440.793	15.800	610.074	18.600	656.339		
sum	440,793		610,074		656,339		137,200	

Tesla China insured units [@Tslachan on X]

2025 Jan	(30)-5	5,500	Jul	(30)-6	
	6-12	7,800		7-13	
	13-19	10,000		14-20	
	20-26	11,600		21-27	
	27-(2)	1,800		(28)-3	
Feb	3-9	6,200	Aug	4-10	
	10-16	7,500		11-17	
	17-23	6,900		18-24	
	24-(2)	12,400		25-31	
Mar	3-9	13,800	Sep	1-7	
	10-16	15,300		8-14	
	17-23	17,400		15-21	
	24-30	21,000		22-28	
Apr	(31)-6		Oct	(29)-5	
	7-13			6-12	
	14-20			13-19	
	21-27			20-26	
	28-(4)			27-(2)	
May	5-11		Nov	3-9	
	12-18			10-16	
	19-25			17-23	
	26-(1)			24-30	
Jun	2-8		Dec	1-7	
	9-15			8-14	
	16-22			15-21	
	23-29			22-28	

• The data on the left is a table to compare Tesla China's domestic insurance numbers (2022 vs 2023 vs 2024 vs 2025)
- I expressed it as 52 weeks.
- National holidays, days of the week, etc. vary from base year to year and are data *for simple reference or comparison.
- Don't take it as an absolute criterion.

Tsla Chan

What I'm watching:

1. The Tesla-Canada situation. Back in late March, Canada froze rebate payments to Tesla, and two days ago, rumors and claims of Tesla being banned in Canada surfaced. Later, these were denied by Canadian government officials. Still, Tesla is now barred from future incentive programs, and I don't think it will end there due to high tensions between Canada and the U.S. over tariffs.
2. On the batteries side, which I discussed last time as Tesla's new bread and butter, I'm optimistic about the upcoming Lithium Refinery Plant in Robstown, Texas, and think the company is best positioned due to the massive head start it has in that regard.

3. I'm also watching Musk's departure from his DOGE position, which caused the backlash against him. He's expected to step back from his role once he finishes the 130-day period as a temporary advisor to Trump.

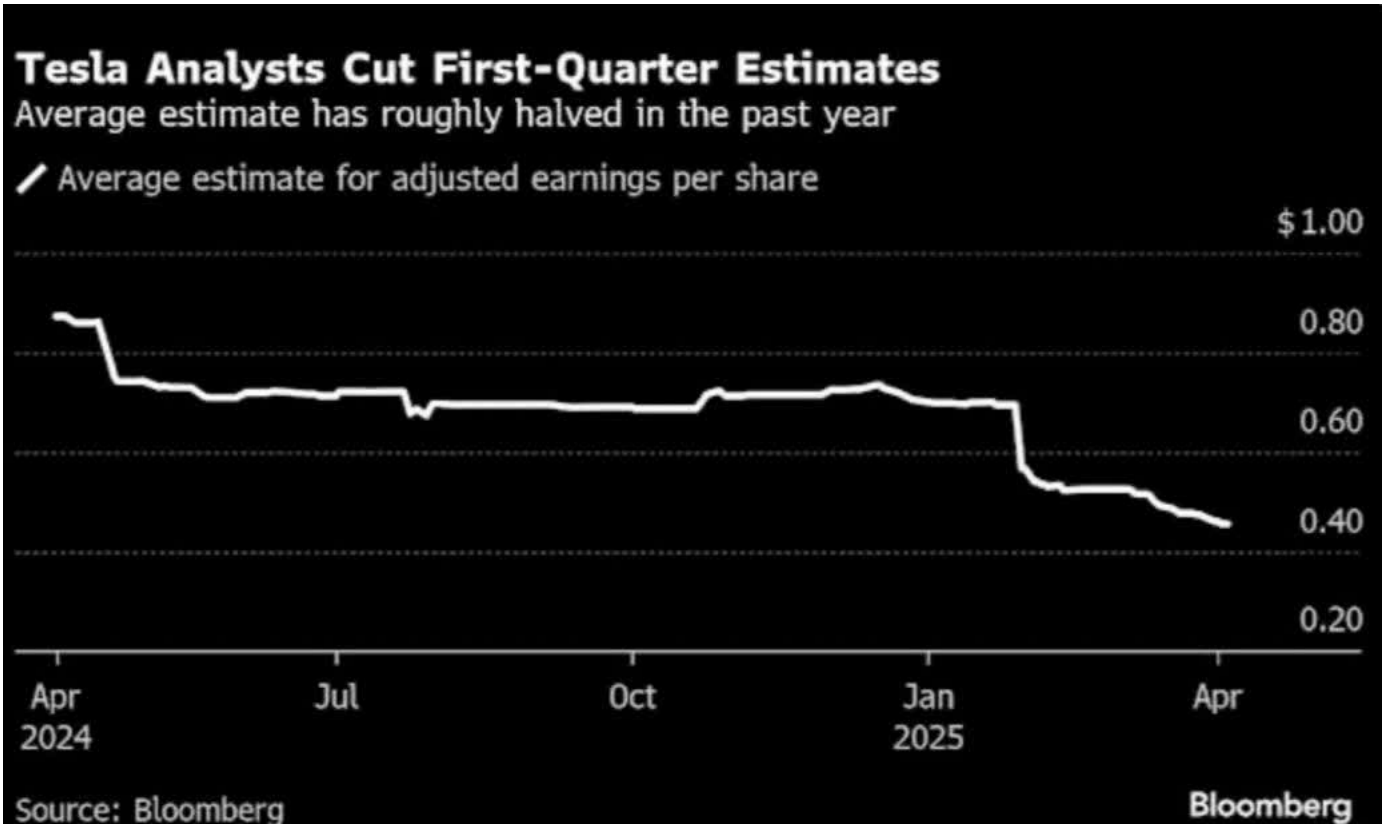

Bloomberg

All eyes are now on the coming earnings call. JPMorgan reduced its estimates of the company's earnings further and now expects Tesla's first-quarter earnings to plunge to 36 cents per share, less than the previous projection of 40 cents per share, and analysts' estimates of 46 cents per share.

ASML: How Secure Is Its EUV Monopoly?

April 2, 2025 Shareholders Unite

Summary

- ASML is the sole producer of cutting-edge EUV lithography machines for advanced 3nm chip manufacturing.
- Huawei and SMEE are developing an LDP-based EUV system, aiming for mass production by 2026, potentially reducing costs and energy use compared to ASML's LPP method.
- Huawei's SAQP technique can achieve 5nm-equivalent transistor density using DUV tools, though it's costly and suffers from low yields.
- Nanoimprint Lithography, another EUV alternative, has been successfully produced by Canon but faces precision limitations for high-volume manufacturing.
- ASML's Zeiss optics remain a significant hurdle for competitors in the EUV lithography market.

Sundry Photography

ASML (NASDAQ:ASML) is the 800-pound gorilla in the lithography market due to being the only one able to produce cutting-edge EUV (Extreme Ultraviolet) lithography machines able to produce the most advanced chips with the most miniature circuitry (3nm).

The company looks set to be a main beneficiary of the AI revolution, while training huge LLMs is running into diminishing returns, the market is shifting towards inference and reasoning models.

The inference market is likely to be multiples in size and reasoning models can take up to 100x the processing power compared to their predecessors, so the demand for AI chips is likely to remain very strong in the coming years as this is still in the very early innings (Q4 conference call, our emphasis):

Looking to 2025, we see full year revenue between €30 billion and €35 billion and gross margin between 51% and 53%. Consistent with our view from last quarter, artificial intelligence has become the key driver for growth in our industry at this moment. As we have witnessed in 2024, AI has created a shift in the market dynamics that is not benefiting all customers equally in the short term. If AI demand continues to be strong and customers are successful in bringing on additional capacity online to support that demand, there is potential opportunity towards the upper end of our range. On the other hand, there are also risks related to customers and geopolitics that could drive results towards the lower end of the range.

Other segments also see an upturn (Q4CC):

Looking at market segments, we currently expect Logic to be up versus 2024 with the ramp of leading-edge nodes while we expect Memory to remain strong, similar to 2024. With respect to our Installed Base business, we expect revenue to grow versus 2024 driven by both service and upgrades as part of a growing install base, to which EUV's contribution is continuing to grow.

2024 - Business summary

ASML:

- Total net sales of €28.3 billion at 51.3% gross margin
- Net income at €7.6 billion resulted in an EPS of €19.25
- Returned €3.0 billion to shareholders through dividends and share buybacks

EUV lithography:

- EUV system sales reduced 9% to €8.3 billion, recognized revenue on 44 EUV systems (NXE & EXE)
- Shipped and recognized revenue from multiple EXE:5000 systems and NXE:3800 systems

DUV lithography:

- DUV system sales increased 4% to €12.8 billion, recognized revenue on 374 systems of which 34% immersion systems
- Shipped and recognized revenue from the first NXT:2150i immersion system as well as the first NXT:870B KrF system

Applications:

- Metrology & Inspection systems sales increased 20% to €646 million due to higher YieldStar and MBI systems sales
- First revenue generation from eScan1100 multibeam inspection (MBI) systems at multiple customers

Installed Base:

- Installed Base Management[1] sales increased 16% to €6.5 billion due to higher service and upgrade business

ASML Q4/24 earnings deck

New products

From the Q4 conference call (Q4CC):

In 2024 we have also successfully achieved a number of technology milestones, including the release of a number of new products critical for our customer's technology roadmaps. On our Low NA EUV technology, the NXE:3800E, we demonstrated the full system specification in our factory with 220 wafers per hour throughput at a new record overlay. We are on track to deliver a new systems at full specification and start upgrades for the system already at our customers during the first half of 2025. We continue to work with our customers to drive the maturation of the system to support their ramp to high volume manufacturing.

First NXT:2150i immersion system shipped

Provides productivity and overlay improvement

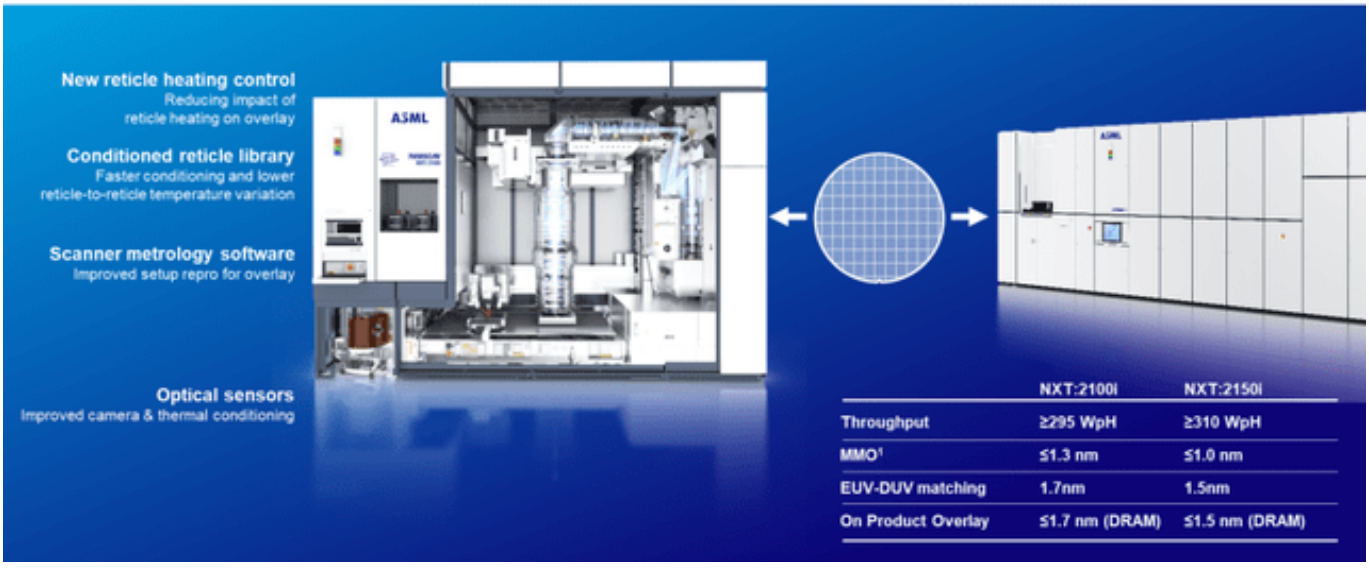

ASML Q4/24 earnings deck

Finances

After a pretty rough year, the company is back on track:

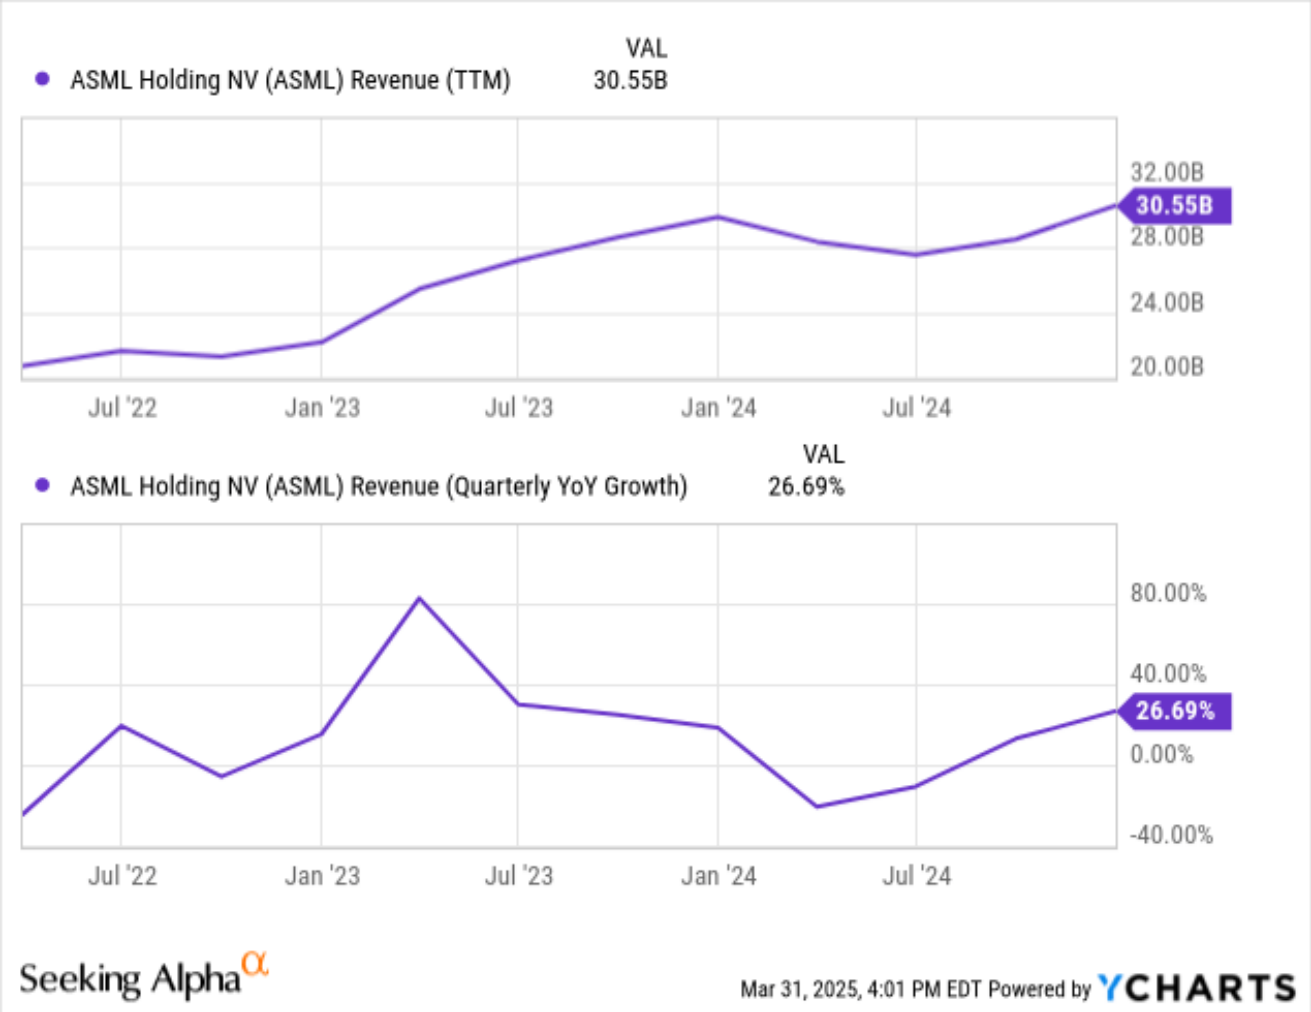

Data by YCharts

This was somewhat surprising as the company's China business actually boomed during that period (Q4CC):

I think to be honest our view on China hasn't really changed in comparison to when we last spoke. Right. So when we last spoke we said the key dynamic on China is that and why China was so very strong in 2023 and 2024 now because for a number of years we have been building up a very significant backlog because all the fill rate for China was so low. That is the reason why the China sales in 2023 and 2024 was so high.

Revenue (+22.5% to E9.3B) came in above the high end of guidance mainly because of installed base revenue and the recognition of the sales of two High-NA systems.

Net bookings came in very strong, up 169% sequentially to E7.09B and way above analyst expectations (E3.99B). Management is discontinuing quarterly bookings in FY25 as these are too lumpy (Q4CC):

Q4 net system bookings came in at €7.1 billion which is made up of €3 billion of EUV and €4.1 billion of non EUV. Net system bookings in the quarter were driven by Logic with 61% of the bookings with Memory accounting for the remaining 39%.

Backlog is also substantial at E36B.

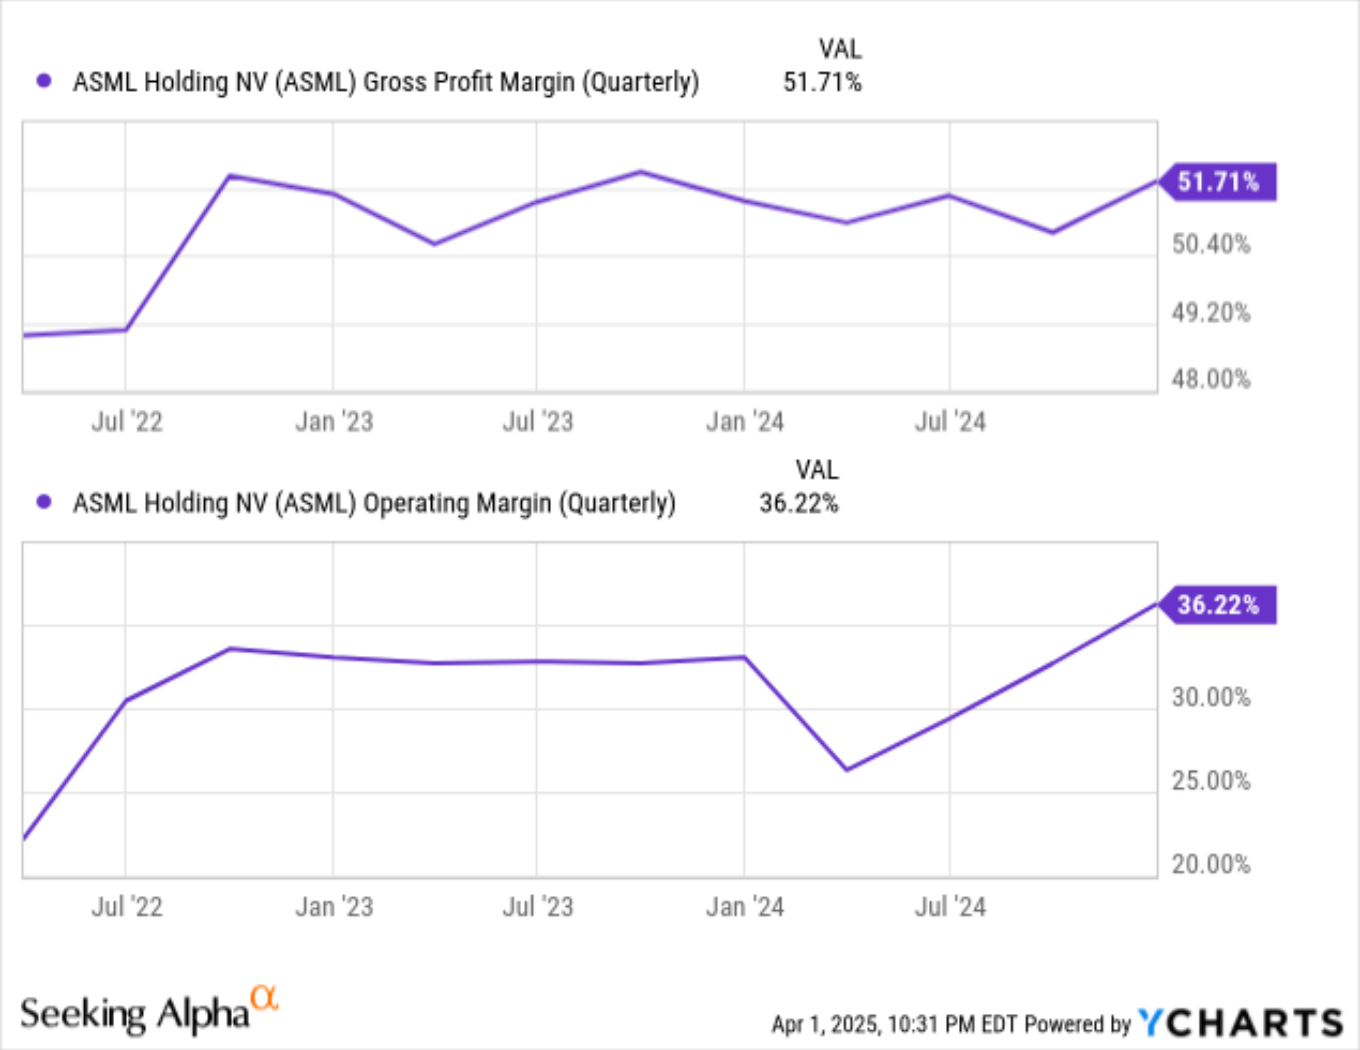

Data by YCharts

The operating margin has recovered nicely and so did cash flow, with a bumper Q4 with free cash flow at E8.84B, almost all the free cash flow for the year:

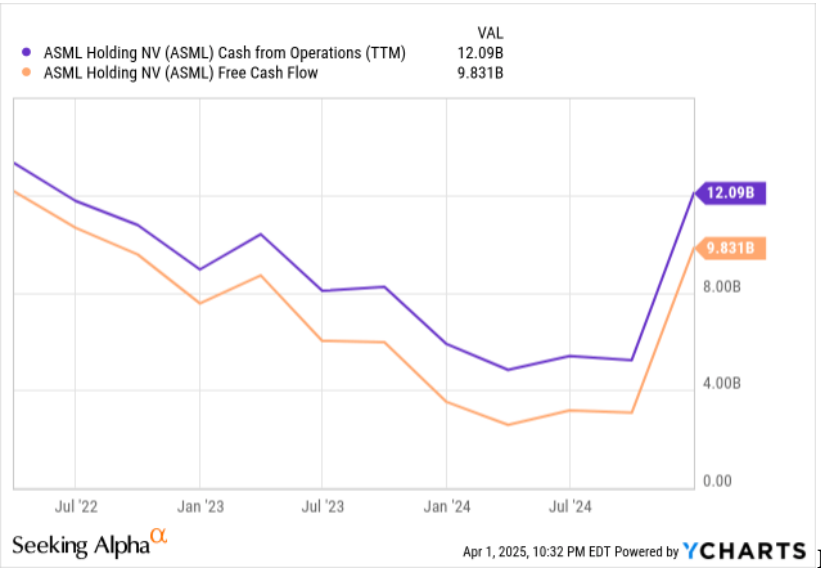

Data by YCharts

There hasn't been a lot of operating leverage in the model but if growth keeps at present levels that will produce some operating leverage.

Emerging competitive threats?

ASML does have legacy competition from Japan, but these haven't been able to keep up. Nikon and Canon offer complementary equipment, they primarily focus on legacy technologies like KrF (Krypton Fluoride) and I-line lithography systems that cannot compete at the cutting edge where EUV reigns.

However, Canon has spent 15 years developing a potential EUV alternative; Nanoimprint Lithography, which:

In contrast to conventional photolithography equipment, which transfers a circuit pattern by projecting it onto the resist coated wafer, the new product does it by pressing a mask imprinted with the circuit pattern on the resist on the wafer like a stamp. Because its circuit pattern transfer process does not go through an optical mechanism, fine circuit patterns on the mask can be faithfully reproduced on the wafer. Thus, complex two- or three-dimensional circuit patterns can be formed in a single imprint1, which may reduce the cost of ownership (CoO). Canon's NIL technology enables patterning with a minimum linewidth of 14 nm2, equivalent to the 5-nm-node3 required to produce most advanced logic semiconductors which are currently available. Furthermore, with further improvement of mask technology, NIL is expected to enable circuit patterning with a minimum linewidth of 10 nm, which corresponds to 2-nm-node.

The company already sold a machine to the Texas Institute for Electronics in 2024, a semiconductor consortium that includes AMD, Applied Materials, Intel, and Micron, but so far, demand isn't exploding despite the scarcity of cutting-edge lithography machines.

The sale was for a price that CEO Fujio Mitarai said would be "one digit less than ASML's EUVs" (as well as using 90% less energy compared to ASML's EUV systems) which would enable smaller chip manufacturers to use advanced machines producing chips with smaller circuitry.

There are problems though:

because NIL is incompatible with both DUV and EUV, its insertion into current design flows is complicated, to say the least. Meanwhile, it is unclear whether it is possible to develop a fabrication process that can rely solely on nanoimprint lithography.

For now, this seems more of a niche solution, but there are improvements possible that could give it a wider appeal.

Chinese efforts

The US-imposed export restrictions which apply to ASML as well are not without risks:

- The obvious risk is lost sales for Western companies in China, which is a large market for many of the cutting-edge chip equipment technology producers, ASML isn't an exception (they can still sell the older DUV machines).
- It greatly incentivizes the Chinese to develop alternatives, and they are putting a lot of resources into this.

Sales to China are being affected as the backlog is tapering off, from the Q4CC:

Our China business in 2023 and 2024 was relatively high because of our ability to execute on a backlog that was created after low order fill rates in previous years. For 2025 and beyond, we expect our China business to go back to a more normalized percentage of our sales.

This demand won't go away as China (or more precisely SMIC) can produce chips with 7nm (and possibly stretch that to 5nm) circuitry and there are enough segments that don't need cutting-edge 2nm or 3nm based chips, but it's clear that selling EUV machines to China would be a large revenue generator for ASML if it was allowed.

The second risk, China developing home-grown alternatives to EUV begins to come into focus, here are some emerging solutions:

- LDP (Huawei and SMEE)
- Discharge plasma extreme ultraviolet lithography light source (Harbin Institute of Technology)
- SAQP or Self-Aligned Quadruple Patterning (Huawei)
- SiCarrier

To the surprise of many, China has already managed to produce chips with 7nm circuitry, like Huawei's Kerin 9010 produced by SMIC. These chips were produced by extending the capabilities of DUV machines, the type of lithography machines ASML can still sell to China.

SMIC has found ways to extend the capabilities of these machines through multiple pattering, which can take circuitry to 7nm, and Huawei's patented Self-Aligned Quadruple Patterning, which combines etching and lithography and might take it even to 5nm while using last-generation DUV machines.

But the following should make instantly clear that this is stretching things:

With EUV, the wavelengths of the lithography printers are 14 times smaller than DUV, at only 13.5 nm, compared to 193 nm of ArF immersion DUV systems... As an example, Intel tried to use SAQP in its first 10 nm nodes to reduce reliance on EUV, which resulted in a series of delays and complications, eventually pushing Intel into EUV.

So it's not surprising this approach has considerable disadvantages in terms of increased complexity, leading to lower yields, slower throughput, and higher production costs.

The last two solutions are the most revolutionary as they propose a different, and possibly better and cheaper way to generate EUV light.

ASML's technology is based on LPP (laser-produced plasma) and relies on high-energy lasers to bombard liquid tin droplets, creating plasma that emits EUV light. This process requires complex FPGA-based real-time control electronics to achieve precise timing and coordination.

While it has proven to be highly effective, LPP systems require sophisticated components and precise engineering, these ASML EUV systems are notoriously complex (they cost between $120M-$150M and can exceed $200M when long-term service and maintenance are factored in). Its latest high-NA EUV tool costs around $380M.

Huawei and SMEE are testing a prototype EUV system LDP (Laser-Induced Discharge Plasma) Technology as an alternative to ASML's laser-produced plasma (LPP) method. Trials are scheduled for Q3 2025, targeting mass production by 2026:

The LDP approach employed in the Chinese system generates 13.5 nm EUV radiation by vaporizing tin between electrodes and converting it to plasma via high-voltage discharge, where electron-ion collisions produce the required wavelength. This methodology offers several technical advantages over ASML's laser-produced plasma (LPP) technique, including simplified architecture, reduced footprint, improved energy efficiency, and potentially lower production costs.

It's not all plain sailing though:

this Huawei system must still answer questions about resolution capabilities, throughput stability, and integration with existing semiconductor manufacturing flows.

One should also keep in mind that LDP isn't a new solution. Early LDP systems used rotating tin supply discs as electrodes. A laser would create a pre-plasma on the tin film before a large electrical discharge generated the Z pinch and EUV light.

LDP's wheel setup offered efficiency benefits by ensuring the laser always hit tin, unlike the droplet targeting in LPP. It also provided some debris mitigation.

ASML used an LDP light source in its first Alpha demo tool but ultimately adopted LPP. The exact reasons for this shift are unclear, but a primary theory is that LDP could not scale output power as rapidly as LPP.

Extreme Technologies (which, like ASML, also comes from Philips) achieved 500W EUV power at the plasma source with LDP in 2008 but the power at intermediate focus (reaching the wafer) was significantly lower. Scaling LDP might have been hindered by the plasma becoming larger rather than brighter at a higher power.

It's not clear whether, or to what extent the Chinese efforts have overcome these power and scaling problems, if they have, they have kept that to themselves.

Asianometry concludes in a video on the subject that a domestically produced Chinese EUV machine is unlikely to be economically competitive with ASML's systems initially. It might suffer from lower wafer throughput or resolution.

That might not matter all that much to the Chinese, given the preeminent strategic importance of the technology and their ability to produce at sustained losses in order to capture learning economies.

Additionally, their LDP solution could incorporate innovative solutions, such as a different optic system to compensate for lower light source brightness.

They exhort ASML to closely monitor any developments, no doubt they are already doing just that, insofar as the Chinese produce papers about the progress they're making.

Conclusion

Several positive things are happening:

- ASML is a premier beneficiary of the AI revolution that is still in the early innings.
- Growth is back on track with booking booming.
- Cash flow was particularly strong in Q4.
- The share price has become more reasonable and is close to or at 3-year lows on multiple valuation metrics.
- While the company can't sell its most advanced machines to China, it still generates considerable business from China selling older DUV machines.
- However much trade relations escalate, China is unlikely to put any restrictions on ASML.
- Its stranglehold on the top end of the market is intact, Canon's NIL machine hasn't gotten much traction and looks to be a niche solution at best, at least for now. Huawei might spin a 'DeepSeek' surprise on all of us with its LDP-based technology, but if they do, they have kept a lot of information to themselves as LDP isn't a new solution. While it has undeniable advantages, there are also important issues to be solved. And based on the arguably limited information that's out, there are no indications that Huawei has solved all of these. For the Chinese market, they don't need to, as they can sell at a subsidized loss, since ASML can't sell its EUV machines to China anyway. This isn't a disaster. Although it provides Huawei with a path to producing for learning economies, so this could eventually become a serious threat.
- A more immediate threat is the state of the world economy (we're writing on the eve of Trump's tariff announcement). While we don't think a significant downturn in the world economy would have a big effect on AI

spending, it can very well affect other chip segments relevant for ASML, so that is a risk, even though management did give a positive outlook for memory and logic, things on the ground can change quickly.

Micron: If You're Questioning Gross Margins, I Have Your Answer

April 2, 2025. Joe Albano

Summary

- Investors and analysts are scratching their heads about Micron's gross margins and how there isn't sequential growth as high bandwidth memory (HBM) continues to grow wildly.
- HBM and high-capacity DRAM for servers are performing well, but are not enough to offset the overall lower margins of consumer DRAM and the like.
- Moreover, Micron's gross margins are being affected by the downcycle in NAND which has more of an effect on margins than you might expect, and my numbers prove the point.
- Meanwhile, the stock chart is mirroring this uneasiness in the fundamentals, as it continues to move through a long consolidation period.

As we headed into Micron's (NASDAQ:MU) FQ3 '25 earnings report the other week, I said there shouldn't be any surprises. And that was mainly the case. High bandwidth memory (HBM) is up with yields a bit better than expected. NAND continues to be in a downcycle, and gross margins are hovering around the mid-30s due to mix and different internal levers to navigate certain end market situations. There really wasn't too much there.

However, what has mounted now, after a guide for gross margins to be down incrementally quarter-over-quarter (again), is confusion among investors and observers as they see the lack of gross margin growth as HBM continues to dominate the headlines and grow massively.

But it's the last part people are taking for granted.

Data from company financials, chart author's

It's correct to believe if it weren't for HBM, DRAM would join its memory brethren NAND in the downcycle camp. After all, everything but the Data Center has been lackluster this past quarter and into this one, and even then only strong in the HBM and LPDRAM products for servers, specifically.

People are starting to question management on why this is the case.

"If HBM is doing so well, why isn't the company doing better?"

There are a few reasons, including the relative size of HBM for the company and how non-HBM products are faring in terms of demand (and supply with NAND).

To put this into perspective, HBM accounted for "over $1B" in revenue during the quarter. Let's call it $1.2B to be conservative and provide the benefit of the doubt. Likewise, high-capacity DRAM and LPDRAM for servers also reached $1B in revenue during the quarter. These are both considered high margin but not as high as HBM. Let's call this $1.1B.

The combination of our revenue from high-capacity DRAM modules and our industry-leading LPDRAM for the data center also exceeded the $1 billion milestone for the quarter.

- Sanajy, Mehrotra, CEO, Micron's FQ2 '25 Earnings Call

Therefore, only $2.3B of the quarter's $8.1B, or 28% were these high-margin, high-demand products. This means it made up almost 38% of DRAM revenue in the quarter. As you can see, 62% of DRAM revenue and 72% of overall revenue do not see this level of demand and margins.

To me, this isn't confusing as to why gross margins are only moderating in the 30s.

But let's dig deeper to understand how DRAM and NAND are doing otherwise and compare it to Micron's performance overall.

Digging Into The Two Memory Types

As we know, NAND is in a downcycle as the company continues to underutilize WFE and has begun to structurally reduce output by mimicking the same strategy it used in DRAM last fiscal year to slow bit growth but still move to leading-edge nodes to get the cost downs. However, DRAM is also in a modest (but not extreme) downcycle in mobile, consumer, automotive, and industrial. These sectors were down quarter-over-quarter and, in most cases, year-over-year.

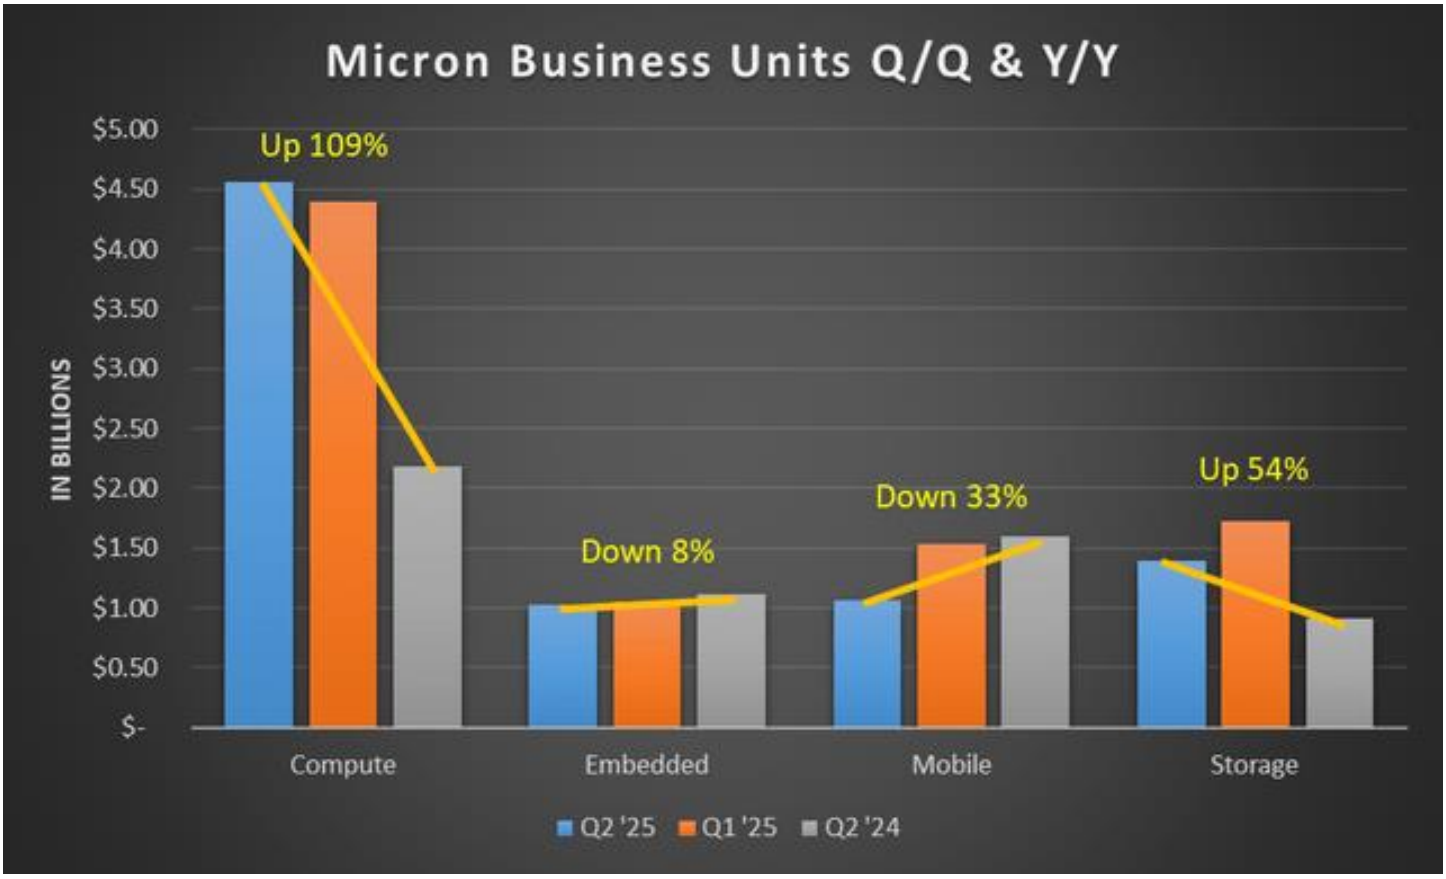

Data from company financials, chart author's

Now, you have to remember, NAND is mixed into these business units except for Compute in varying degrees. This is why the Compute (and Networking) unit is not seeing much of a hit, though it includes PC and graphics DRAM, which has slowed a bit, but not to the extent NAND, particularly in the Data Center, has, which was down 20% quarter-over-quarter, though up 54% year-over-year.

As you can see, there's no direct data on where NAND is taking a hit vs. DRAM relative to the specific sector. The only pure-play business unit is Compute and Networking, and it only tells us four DRAM markets (data center, PC, graphics, and networking). This doesn't include mobile, auto, industrial, and consumer embedded DRAM, which is in Mobile and Embedded.

But, it's enough to understand everything except for Compute and Networking was down quarter over quarter, with Mobile and Storage the biggest losers. So, with the idea demand isn't chipper across the board, let's get back to the gross margin point.

Digging Into Their Respective Gross Margins

While Compute was up, it doesn't mean data center was all of it. Going forward, consumer PC and graphics DRAM growth is a headwind to the massive margins HBM contributes (some estimate 70% GMs) since consumer DRAM (and NAND) are of lower margins. This is called mix. I'll get to why this is important in a moment.

To come back to a statement I made earlier in this article, if we back out the HBM and high-margin products, DRAM would also look like a moderate cycle correction right now. This is important to keep in mind while considering the bigger picture.

But here's what's happening. NAND is hurting, and outside of HBM and high-margin LPDRAM for AI servers, DRAM is treading water in terms of bits shipped with indications of it slowly returning to growth in average selling prices (ASP). The cycles in the two memory types are somewhat misaligned in 2025, where generally they are in synchronous declining sine waves of sorts (it's not usually gradual up and gradual down, but you get the point).

So if NAND prices and bits shipped aren't up when DRAM prices and bits shipped are up, there's not a nice tracking of metrics, including margins, to become a clean sequential upcycle like we're used to seeing. Then, you add in HBM which has thrown a wrench into this as it "makes up for" other weaker DRAM products in the consumer, mobile, and auto businesses, and margins aren't so easy to read.

But there's more to it then just demand.

The next question is, have DRAM and NAND prices actually recovered from prior peaks?

Obviously, there isn't an expectation for the same dollar prices since technology costs come down over time, related to Wright's Law, and separately but similarly, Moore's Law. Moreover, memory isn't so cut and dry. New technologies aren't like transistors in a CPU, which get smaller and smaller but function largely the same way. DRAM technologies get more complicated. DDR4 to DDR5 is an example, as it has two new logic components that DDR4 doesn't have. This adds to the costs of DDR5; thus, the gross margin profile won't be the same as the later generations of DDR4. However, the price of DDR5 is higher than DDR4. In the end, you can see there's this give-and-take on both sides of the equation.

But, above all, we haven't seen the surge in DRAM pricing like in years past. It's been more subdued in terms of the strength of pricing. To illustrate this, I took a starting point of $8.00 in ASPs (the dollar amount is really for presentation purposes) and followed Micron's earnings reports using their percentage of quarter-over-quarter changes each quarter. Similarly, I took Micron's FQ2 '18 reported DRAM gross margin of 66% (yes, it was still providing DRAM versus NAND gross margins then), giving me its average cost of DRAM, and then followed its guidance (and confirmations) for cost-downs over the years. We wind up in FQ2 '25 with a gross margin for DRAM of about 54% (excluding HBM). This is what it looks like:

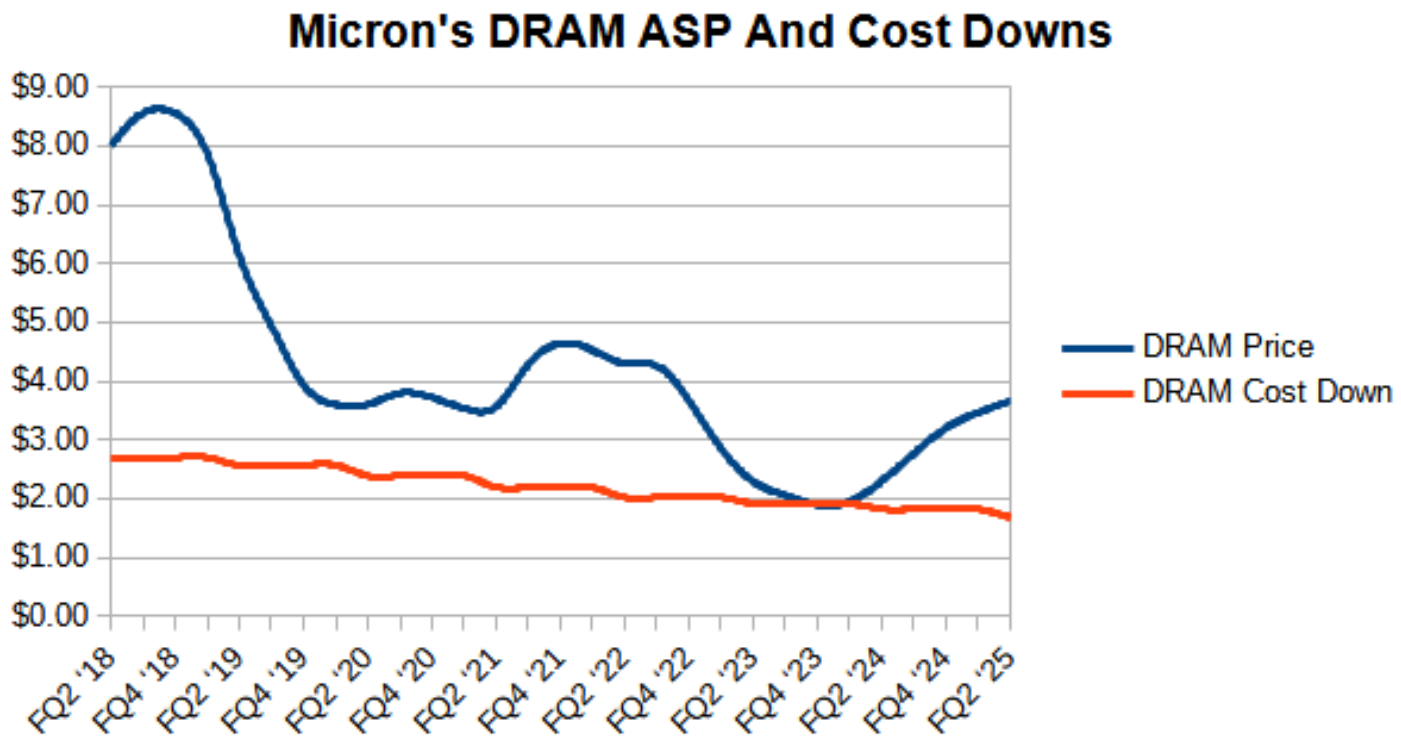

Company filings, transcrips, and Grok3, chart author's

Now, obviously, that's not the complete picture, so I ran the same exercise for NAND. Things are much tighter in NAND. In 2018, the company's gross margins were 47% for NAND. However, NAND very quickly came down in price from 2018 to 2020, steeper than DRAM. Over the last several years, NAND has had negative gross margins several times, while DRAM barely had it occur once.

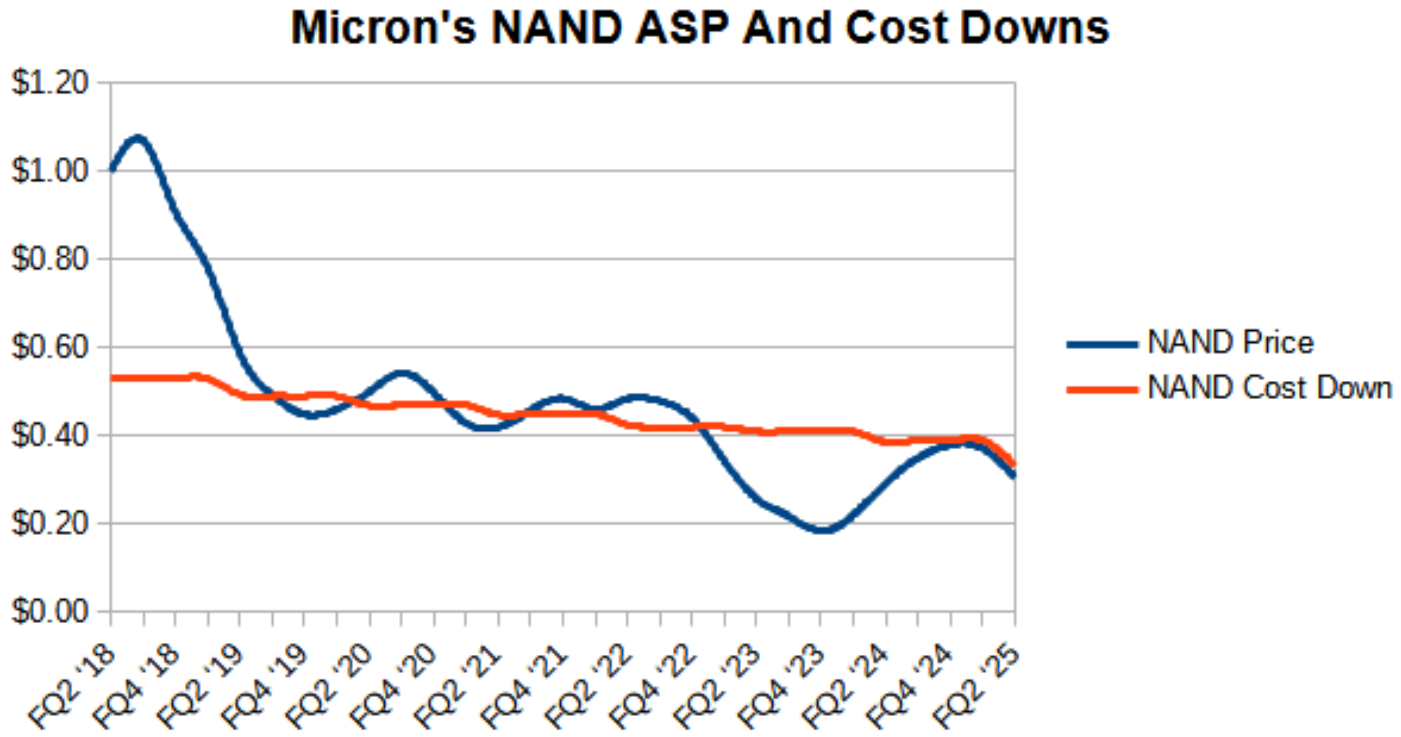

Company filings, transcrips, and Grok3, chart author's

With this in mind, we can build a gross margin profile to explain why, even with HBM and the spending for all things AI, it doesn't matter if the upturn in DRAM is happening or HBM is growing wild percentages, the downturn in NAND with its typical razor-thin margins is keeping overall company margins compressed.

Using the assumption HBM revenue of $1.2B tracks at 70% gross margins, I'll place the high-capacity server and LPDRAM at 61%. And through my pricing exercise, the remaining DRAM comes to about 54% (though for DRAM and NAND in 2025, I factored in the full-year cost downs already, so this calculation will back out half of that), giving us an average gross margin for DRAM of 59% in FQ2.

For NAND, it's currently looking at a -20% gross margin, and on its $1.9B in revenue, it lost $380M in FQ2. With that in mind, the total company gross margin calculates to 39%, just slightly higher than the company's reported 38%. That's not bad for a back-of-the-envelope tracking of pricing and cost-downs from 2018 to now! It means I'm pretty close. I may be off on individual categories of products, and I'm willing to fine-tune them, but overall, the bigger picture is clear.

So, this exercise says the bottom line is NAND has more sway on the company than you might have first thought, while it's currently more of a drag on the company than first expected. But not all is to blame on NAND. Even with HBM, DRAM margins are still down seven percentage points from those 2018 highs. Pricing matters, and right now, the price of DRAM is not at the level we've seen in other upcycles compared to the costs. Thus, combined with NAND in a downcycle, this upcycle seems (and is) muted.

All Is Not Lost Moving Forward

But I'd be remiss if I didn't give you the picture before us.

If NAND is in the dumps now, but the cost downs for this year are some of the highest the company has seen historically, then when NAND demand recovers, and the producers get supply under control, Micron will have a more structurally sound cost basis for when ASPs find an upward trend.

We expect fiscal 2025 NAND front-end cost reductions to be in the low-teens percentage range.

- Sanajy, Mehrotra, CEO, Micron's FQ1 '25 Earnings Call

In other words, when NAND recovers, and DRAM simply continues on this slow recovery path, margins could quickly get into the 40s. For example, if NAND margins were breakeven this past quarter, the same revenue would have provided a 44.5% gross margin. A 10% NAND gross margin would push the same quarter to 46.7% overall.

The question now is, what does the rest of the fiscal and calendar year look like in terms of NAND and DRAM pricing, cost downs, and demand (bits shipped)? I've already tired you out in this article, so let's save it for the next one. After all, this understanding was needed first to know how the outlook truly affects the financials.

Chart Analysis

In addition to this confusion over gross margins, the stock price appears to reflect the uncertainty surrounding this upcycle. As I've discussed for the last several months, Micron is in a very messy consolidation, and after the last few weeks, it may be more complex than I initially had in mind. Nevertheless, I now understand better what it's trying to do.

Chart autho's

Micron has been working through sideways consolidation for the better part of eight months. More recently, it has been working in a triangle pattern, but it's nearing the end of that pattern and should break out to the upside once complete, at least in the short term. It has a shot to head to the (B) wave target shown above between $115 and $120. Once it reaches the (B) wave target, wave (C) should start to the downside, targeting the $70s. All in all, there should be one more move higher (the wave [C] of (B)) to exit long positions before we move into the final stage of this correction (or at least this leg of it).

Overall, Micron is in a messy situation, both in the memory cycle and the chart. For now, I'm not going to make any serious moves with my long position, though, I may trim again on a move to $115-plus. For the time being, both the memory market and the chart need to work through some baggage before we can see Micron settle into some type of advancement.