

AI Research

SemiBytes: Semis Rally, Fast Inference, Analog, AMAT/LRCX

Equities

Americas

Semiconductors

Timothy Arcuri

Analyst

timothy.arcuri@ubs.com

+1-415-352 5676

Grant Joslin

Analyst

grant.joslin@ubs.com

+1-201-352 0442

Natalia Winkler, CFA

Analyst

natalia.winkler@ubs.com

+1-415-352 4626

Gianmarco Vella

Associate Analyst

gianmarco.vella@ubs.com

+1-415-352 4555

Aaryan Wadhwa

Associate Analyst

aaryan.wadhwa@ubs.com

+1-212-821 6481

Dino Weinstock

Associate Analyst

dino.weinstock@ubs.com

+1-212-713 3172

Earnings previews - Companies reporting the week of 5/4

This week, we are previewing ON, GFS, AMD, SWKS, QRVO, ALAB, SITM, ALGM, MCHP, and INDI. Please refer to our [standalone note](#) for more details.

SOX breaks 10,000 after a 40% April rally – taking a step back, trying to gain some perspective

Semis stocks rose for the 18th straight day on Friday - something that has only happened once in history, during the dotcom period (just before the crash) - with the SOX now up nearly 40% this month. Semis have grown to now >10% of the market cap of all stocks around the world, up from a little more than 3% at the peak of COVID just a few years ago. Semis just surpassed consumer discretionary and are the 2nd largest sector by global capitalization with only financials being larger and now >2x the contribution to global market cap vs. other large tech sectors like software. Yet, on a PEG basis, the group is still as inexpensive as any time in the past 25 years as this rally has been supported by accelerating EPS growth. Within semis, market cap has obviously shifted toward compute, but compute is actually still below its relative contribution in the late 1990s and memory remains notably only ~1/3 of its relative contribution to total sector market cap versus the last major memory supercycle in the mid-2010s. This is particularly surprising to us because AI – and particularly inference – is fundamentally memory enabled and memory bound. This speaks to why we remain bullish on memory overall (and semicaps alongside memory). We show all of this in Figures 1-5.

Speed (and time to first token) is overtaking model scaling as the key bottleneck

As semis have grown in market cap, we have received inbounds from a very wide array of investors that reaches far beyond the traditional tech-focused investors that we have always spoken with and we are getting many questions about LLMs, model scaling, and where we may see the next bottlenecks developing. To that end, we wanted to highlight [a blog post by Amazon Science](#) initially published on Feb 26 which argues that today's scaling of LLMs via parameter count is reaching diminishing returns. Beyond a certain size, models enter a "savant" regime where benchmark gains persist - but true insight actually deteriorates. The authors frame intelligence as a form of time-bounded problem solving, in which the effectiveness of a model (or an agent) is based on its ability to solve tasks efficiently and allocate computational efforts dynamically. Said differently: value is shifting in how efficiently an LLM can solve problems under time constraints such that, when time is limited, systems are forced to learn and prioritize effort. This is an evolution away from slow, brute-force approaches (often correlated to the size of the models). According to the authors, the implications are a shift away from parameter growth and toward architectures/systems that optimize inference efficiency and "temporal adaptability". An example would be the ability to retrieve only portions of memory that matter when reasoning through new solutions - implying that the next leg of AI performance will be driven more by speed and efficiency than sheer size. Along these lines, an [article from The Information](#) out this week points to how OpenAI's Spud (yet to be launched) and Anthropic's models are already reducing their reliance on reasoning while still being capable of providing more intelligent answers. The article cites OpenAI's upcoming Spud model being "noticeably better at answering tough questions without relying on reasoning" and makes the point that Anthropic's Claude Opus 4.7 is more performant than Opus 4.6, but uses fewer tokens - which also suggests less reliance on reasoning. While some of this has to do with much more improved focus on smarter

pretraining, it appears the focus on reasoning to not be growing as fast going forward for two reasons: (1) better pretrained models are cheaper to run given the lower need to output tokens for reasoning, and (2) less time spent on reasoning means models are better for agentic tasks that require fast performance. These are both theoretically interesting and speak to the nature of intelligence - but also, we believe speak to lateral moves from the frontier model developers to attempt to continue scaling AI performance along new dimensions in recognition that leading edge foundry and memory capacity will likely remain very tight for the foreseeable future. This could open up new opportunities for alternative compute architectures (such as LPU with NVDA's Groq acquisition) and maybe ultimately being bullish for CPU.

We remain bullish analog even after this move; MCU not even back to trendline, so MCHP may be next

Following TXN's strong guidance last week, we received a flood of investor calls - many of which carried an underlying concern around the sustainability of this analog upcycle and whether possibly fueled by pull-ins. Put simply, we think this exactly what happens during upturns - purchasing managers see capacity tightness increasing and prices potentially going higher and respond by building some inventory. This is what creates the upturn and can often be sustained for 2yrs+. We have updated our long-term cyclical framework based on our bottom-up rollup of 10 top analog companies ([Figure 6](#)); the industry has a long term 5% total revenue CAGR which is comprised of ~8-9% for industrial and autos and more of a flattish trend for telecom and personal electronics. What this work shows is that relative to trendline, analog is only back to long-term trend and has yet to even go above trend. MCU on the other hand, remains far below trendline and we think maybe sets up as the next big vertical where we could see big positive surprises (MCHP in our coverage). For the sector as a whole, we remain above Street through 2028 and our estimates contemplate industrial only returning to its long term trendline growth with no premium reflecting secular drivers like advanced robotics, renewables/smart grid, or data center (which is ~10% of revenue and growing at >50% CAGR). We would also note that profitability across the sector remains depressed ([Figure 8](#)). In the US we continue to favor TXN and see MCHP as maybe the next to report and guide to much higher revenue. Internationally, UBS continues to prefer Renesas and STM.


TSMC's push out of High NA - positive for AMAT/LRCX

In conjunction with [TSMC's US Tech Symposium this week](#) (covered by Sunny Lin), we received a host of investor questions around the significance of the company's decision to [delay High NA EUV adoption to nodes beyond A13](#) (ASML covered by Francois-Xavier Bouvignies). We view the implications for deposition and etch equipment suppliers as structurally positive. Specifically, we believe that TSMC's ability to shrink metal pitches to <25nm via the use of 0.33 NA should result in a meaningful increase of litho-etch-dep loops across the process flow. While we would argue that TSMC would have not pursued this path if it would have not been able to maintain a ~4 mask limit at key M0-M2 layers, we think this approach could require an incremental 5-7 deposition steps, and 12-5 additional etch steps in aggregate compared to the case of a single 0.55 NA exposure, one hard-mask transfer etch, and one direct metal etch. In quantifying the incremental steps needed, we assume a M0 pitch of 20-21nm via SAQP, M1 pitch of 22-23nm via SADP, and a M2 pitch of 23-24nm via SADP. We also note that TSMC's transition to backside power distribution (TSMC coined this as Super Power Rail or SPR) at A14 likely alleviates routing congestion in lower metal layers, again diminishing the case for High NA adoption. Net, we see the postponement of High NA adoption at TSMC as a clear tailwind for deposition and etch suppliers (AMAT and LRCX), both from an intensity standpoint (higher pattern complexity, more alternative pattern "healing" techniques necessary, etc) and potential upside in capacity demand (due to longer cycle times). We also believe this could increase the likelihood of adoption of alternative patterning solutions like dry resist, growth in ALD patterning, hard mask optimization, and line edge roughness mitigation techniques - all of which reinforce our view for a very strong setup for AMAT and LRCX over the next several years.




SEMIS RALLY IN CONTEXT

Figure 1: SOX M/M return history




Source: UBS research, FactSet. Weekly data based on trailing 4 weeks' performance.

Figure 2: SOX & S&P 500 PEG - this rally has been more than paid for in growth...




Source: UBS research, FactSet.

Figure 3: SOX & S&P 500 LTG - ...with EPS LT growth estimates continuing to expand on what was already a record high




Source: UBS research, FactSet.

Figure 4: Semis now surpass 10% of global market cap



Source: UBS research, FactSet.


Figure 5: Key semis stocks' market cap share by subsector



Source: UBS research, FactSet. Includes 47 key semiconductor stocks by current market cap

ANALOG CYCLE UPDATE

Figure 6: Analog thru-cycle long term view and Industrial/Auto and Analog/MCU decomposition




Source: Company reports, UBS research, FactSet, Visible Alpha.

Figure 7: Analog Semis UBS estimates vs Visible Alpha Consensus

Calendar Year, USD \$MM	CY26e			CY27e			CY28e			CY29e		
	Street	Δ	UBS	Street	Δ	UBS	Street	Δ	UBS	Street	Δ	UBS
Total Revenue												
ALGM	998	+31	1,029	1,173	+110	1,284	1,384	+147	1,531	1,651	+33	1,684
ADI	14,302	+407	14,709	15,733	+836	16,569	17,218	+865	18,083	18,409	+619	19,028
ON	6,302	-112	6,189	6,951	+60	7,010	7,637	+406	8,043	9,026	-167	8,859
MCHP	5,468	+189	5,657	6,427	+581	7,008	7,214	+956	8,170	8,740	+248	8,988
TXN	21,016	+583	21,599	23,176	+1,615	24,791	25,379	+1,639	27,017	29,001	-393	28,607
IFX	19,102	+103	19,205	21,657	-277	21,380	24,598	-1,165	23,433	27,675	-2,381	25,294
NXPI	13,557	+48	13,605	14,953	+361	15,314	16,068	+733	16,802	18,226		18,226
STM	14,109	+387	14,497	16,248	+534	16,783	17,879	+423	18,302	19,798	-71	19,727
Renesas	10,020	-428	9,593	11,091	-667	10,424	12,194	-1,194	11,000	11,639		11,639
Rohm	3,188	-31	3,157	3,333	-4	3,329	3,543	-63	3,480	3,671	-32	3,639
Total	108,063	+1,176	109,239	120,742	+3,149	123,891	133,114	+2,747	135,860	147,834	-2,143	145,691
% Y/Y	15.5%		16.8%	11.7%		13.4%	10.2%		9.7%	11.1%		7.2%
Gross Margin (Non-GAAP)												
ALGM	51.0%	+114 bp	52.2%	53.0%	+191 bp	54.9%	55.4%	+91 bp	56.3%	56.9%	+15 bp	57.0%
ADI	72.0%	+24 bp	72.3%	72.5%	+146 bp	73.9%	73.1%	+193 bp	75.0%	74.4%	+72 bp	75.2%
ON	39.6%	-10 bp	39.5%	43.2%	-22 bp	42.9%	45.6%	+44 bp	46.0%	48.3%	+20 bp	48.5%
MCHP	61.1%	+221 bp	63.3%	63.0%	+278 bp	65.8%	65.4%	+143 bp	66.8%	66.0%	+98 bp	66.9%
TXN	59.4%	+141 bp	60.8%	60.7%	+312 bp	63.8%	61.4%	+426 bp	65.7%	61.7%	+512 bp	66.8%
IFX	43.7%	-70 bp	43.0%	46.0%	-117 bp	44.8%	47.2%	-218 bp	45.0%	45.8%	-223 bp	43.6%
NXPI	57.9%	-21 bp	57.7%	59.1%	+1 bp	59.1%	60.2%	+18 bp	60.4%	61.1%	-0 bp	61.1%
STM	35.9%	+67 bp	36.6%	39.2%	+303 bp	42.2%	42.0%	+356 bp	45.5%	45.0%	+100 bp	46.0%
Renesas	58.9%	-66 bp	58.3%	60.2%	-68 bp	59.5%	61.3%	-70 bp	60.6%	61.7%		61.7%
Rohm	24.4%	+495 bp	29.3%	27.6%	+420 bp	31.8%	28.3%	+333 bp	31.7%	27.5%	+428 bp	31.8%
Total	52.8%	+51 bp	53.3%	54.6%	+142 bp	56.0%	55.9%	+161 bp	57.5%	56.5%	+130 bp	57.8%
Operating Income (Non-GAAP)												
ALGM	188	+27	216	288	+91	379	427	+104	531	602	+24	626
ADI	6,827	+375	7,202	7,688	+924	8,612	8,645	+1,010	9,654	9,581	+585	10,166
ON	1,326	-61	1,265	1,786	-1	1,786	2,195	+262	2,457	3,070	-32	3,038
MCHP	1,736	+223	1,959	2,362	+530	2,892	2,963	+684	3,647	3,974	+133	4,107
TXN	8,517	+622	9,139	9,862	+1,796	11,658	10,932	+2,493	13,425	12,601	+2,020	14,621
IFX	3,863	+36	3,900	5,149	+30	5,178	6,493	-599	5,894	7,811	-1,495	6,316
NXPI	4,735	+116	4,851	5,536	+487	6,023	6,229	+880	7,109	8,105		8,105
STM	1,320	+188	1,508	2,529	+647	3,175	3,358	+995	4,353	4,611	+318	4,929
Renesas	2,166	-222	1,944	2,662	-232	2,430	3,201	-368	2,833	3,281		3,281
Rohm	184	+56	240	323	+22	344	438	-58	381	492	-82	411
Total	30,862	+1,361	32,223	38,183	+4,293	42,477	44,881	+5,404	50,285	54,130	+1,471	55,601
% Y/Y	31.5%		37.3%	23.7%		31.8%	17.5%		18.4%	20.6%		10.6%


Source: Company reports, UBS research, FactSet, VisibleAlpha

Figure 8: Analog Cycle Table




Source: Company reports, UBS research, FactSet.

Figure 9: Analog historical margins and multiples



Source: Company reports, UBS research, FactSet.

Figure 10: Analog historical margin/multiple trajectory - the margin continues to work higher while the multiple range remains similar in successive cycles



Source: Company reports, UBS research, FactSet.

Valuation Method and Risk Statement

Investing in the technology sector entails above-average risk given low sales visibility, rapid pace of innovation and technological change, intense competition, frequent M&A, and low barriers to entry in many markets.