

The Opportunity: Crowded Theme, Empty Trade

Shay O'Kelly · August 2026

Anyone who has read this far should now raise the obvious objection. If the buildout is the largest in the history of capitalism and the bottlenecks are visible enough that a college student can list them, then surely it is all in the price. And by every measure of positioning, it should be. In Bank of America's August survey of global fund managers, long the Magnificent 7 was once again named the most crowded trade in the world¹. Two months earlier the same survey found 80% of managers crowding into global semiconductors, one of the highest conviction readings in its history². Hedge funds just raised their tilt toward technology by the most in any quarter on record³. These are the most owned, most covered, most discussed companies on earth. I am not bringing anyone a new ticker. So the fair question, the one I would ask any manager pitching this thesis, is simple. Where exactly is the opportunity in the most crowded trade in the world?

Here is my answer, and it is the whole essay. The macro has believers. The micro has none. Money is never static, it flows, and you learn more from where it refuses to go than from where it goes. At the level of the theme, the survey level, the index level, belief is total. Then walk down the ladder toward the actual securities, the actual order books, the actual signatures, and watch the belief drain out at every rung. A survey measures what people say. A price measures what they pay. Ignore the commentary and read the tape. The tape says the market that cannot stop talking about AI is pricing the companies that supply it as if the whole thing ends in about two years.

Start at the top rung, the biggest companies in the world, where the doubt is already visible. Alphabet, which runs the largest custom silicon fleet outside Nvidia's and just guided to as much as \$190 billion of capex, trades near 16 times earnings while the index trades above 20⁴. A discount to the market that contains it. Amazon raised its 2026 capex to \$220 billion, twenty of those billions because memory got more expensive, and its stock has been pinned in a range all year while investors debate whether the spend will earn its cost of capital⁵. And when these companies report, the market grades them on one axis. In July, Meta guided to as much as \$145 billion of capex and fell 7%. Microsoft guided to more, \$175 billion, and rose 7%, because it could point to Azure crossing \$100 billion of annualized revenue, a visible line from spend to sales⁶. Oracle raised capital spending 162%, announced \$40 billion of new debt and equity, and fell 12% in a day. It has been nearly cut in half since June and was downgraded to one notch above junk while sitting on a \$638 billion backlog, eleven years of its current revenue, already signed⁷. One money manager summarized the regime in a sentence. Capex used to be the more the better. Now it is the less the better⁶.

Hold that up against 1999 for a second, because the contrast is the tell. In a bubble, the market pays for the story. Announcing fiber capex made a telecom stock go up, and the checks chased the narrative all the way to the grave. In 2026, announcing AI capex makes your stock go down unless the revenue is already on the

tape. The market is rewarding harvest and punishing planting. That is not a mania. That is a market that owns the theme and does not believe the numbers.

One rung down sits Nvidia, the most important company in the world, trading at roughly 24 times forward earnings, its cheapest multiple in years and less than half its own three-year average of 52⁸. The S&P 500 trades at 20⁹. Four turns above the index. That is the entire premium, while the current quarter is guided to \$91 billion of revenue, a billion dollars a day, with the data center segment growing 92%, and consensus already near \$390 billion for next year¹⁰. Pay four turns over the market multiple, receive growth running roughly ten times the market's. Euphoria was Cisco at over 100 times earnings in March 2000. This is the opposite creature. This is the market renting the most analyzed stock on earth one quarter at a time because it does not trust the earnings past next year.

Another rung down, the suppliers, and now the doubt turns to open disbelief. Micron has sold out substantially all of its DRAM and high-bandwidth memory output through 2027 and signed sixteen multi-year customer agreements¹¹. The stock is up over 200% this year and trades at 5.7 times forward earnings against a ten-year average of 22. SK Hynix, the HBM leader, trades at 4.9 times after rising roughly ninefold. Samsung sits near 7¹². Super Micro booked over \$60 billion of new orders in a single quarter and trades at 17 times trailing¹³. And the disbelief is not a mood, it is published. Bank of America's base case for 2028 assumes DRAM prices fall 10% and NAND falls 18%, and Micron still earns roughly \$150 a share. Its severe case, a full traditional memory bust, still earns about \$100, eight times the peak of the entire 2018 cycle¹⁴. The stock trades at five and a half times the first number and about eight times the second. Say that in plain English. Micron sells about a fifth of the world's DRAM¹². Buy the stock today and the company earns back its entire share price in about five and a half years in the base case, with memory prices falling the whole way. In a full bust it takes about eight years. The downside case is a slower payback, not a loss. The market is not pricing a downturn. It is pricing the memory of every downturn since the 1990s onto a business whose worst case now out-earns its best pre-AI year by a factor of eight. Meanwhile the buyers are telling you the direction of travel. Amazon just raised its capex guide by \$20 billion and blamed memory prices⁵. A multiple of 5 is not what belief looks like. It is a countdown.

Widen the lens past the names everyone argues about, because the same mispricing runs down every aisle of the store. The equipment makers, Applied Materials, Lam Research, and KLA, the companies that sell the machines every fab on earth is built from, just raised their own 2026 spending forecast by 30%, to as much as \$160 billion, on a consensus path toward roughly \$260 billion by 2028. The stocks sold off 20 to 40% anyway, and the group now trades near 0.7 times its expected growth¹⁵. Storage is the least glamorous aisle in the data center, and it is gone too. Western Digital has sold 100% of its 2026 hard drive production, with purchase agreements reaching into 2028 and 2029, and Seagate's data center capacity is allocated through 2027¹⁶. Sold out for years is now the norm all the way down to the spinning disk. So is the discount.

Go inside the machine itself, because the chips are changing and the market has not repriced the change. Walk the chain. Nvidia alone will sell roughly \$390 billion of AI chips next year¹⁰. Custom chips, the ASICs that Google, Meta, Microsoft, OpenAI, and Anthropic design for their own data centers, are

growing three times faster than GPUs and pass them in units shipped next year. Two companies, Broadcom and Marvell, design 95% of those custom chips¹⁷. Broadcom's slice alone is guided to \$100 billion of AI revenue in 2027, roughly double its entire company today, and the stock costs about 22 times the earnings that revenue implies, the same multiple as an average S&P company¹⁸. A duopoly on the fastest-growing chip category on earth, priced like everything else. The pure inference challengers get the opposite treatment. Cerebras came public at \$48 billion and trades north of 50 times sales¹⁹. The one I watch closest is Qualcomm, because it is not selling a card, it is selling full server racks, shipping this year, built by the company that spent twenty years making phone chips where every watt and every degree of heat mattered, now aimed at data centers that are rationed by exactly those two things²⁰. And here is the detail that ties every one of these machines back to the anchor. Open any of them up. Qualcomm's card carries 768 gigabytes of memory. Nvidia's next flagship carries up to a terabyte per GPU. The compute die is a sliver of the board, and the rest is memory, because serving a model means holding the model and its context right next to the processor²⁰. That is why I do not need to pick the winner of the architecture war. Every one of these designs, whoever builds it, is mostly memory, and the companies selling that memory trade at 5 times earnings.

The bottom rung is CoreWeave, the purest public expression of the buildout, and this is where the gap stops being a discount and becomes an absurdity. Watch the order book move in real time. Signed contracts stood at \$60.7 billion at the end of last year, \$99.4 billion in March, \$104 billion in June, and \$129 billion as of this week, counting the more than \$25 billion of new commitments signed in the first six weeks of the third quarter alone²¹. That last stretch is over \$600 million of new signatures a day. Roughly 40% converts to revenue within two years and about 80% within four, and the business underneath is compounding on schedule, with revenue up 112% last quarter and guided to roughly double for the year²¹. Contracted power reached 4.2 gigawatts, two Hoover Dams of electricity, secured in a country where a new grid connection takes five or more years²². The queue is the moat. A contracted watt is a call option on every GPU generation of the next decade, and the strike was paid when the substation was signed.

Now the napkin, in three steps. Step one, what a megawatt costs. Building an AI-ready data center runs \$15 to \$25 million per megawatt before the chips go in, and \$30 to \$45 million with them²³. Step two, what the market charges for CoreWeave's megawatts. The company's enterprise value, meaning the stock plus every dollar of its debt, is about \$84 billion, and its contracted power is 4.2 gigawatts, which is 4,200 megawatts. Divide one by the other and the market's price is about \$20 million per megawatt²². Step three, compare. The market is selling CoreWeave's power pipeline for roughly what the empty buildings cost to construct, which means the \$129 billion of signed contracts, the installed GPU fleet, and the operating business are all priced at zero. The bear case stays on the page. OpenAI is roughly \$22 billion of the commitments, the net loss widened to \$626 million as the capex ramp ate the income statement²¹, and building the contracted capacity means spending about five times trailing revenue this year²⁴. So the position is sized like what it is, a leveraged claim on signatures whose largest counterparty burns cash. But the stock jumped 14% on these numbers and still sits roughly a third below its high, which is a preview of the mechanics. When the signatures print, the repricing runs into this trade, not out of it. And it is not one

company. Nebius carries close to \$50 billion of signed contracts from Microsoft and Meta against roughly half a billion dollars of last year's revenue, and took a \$2 billion equity check from Nvidia²⁵.

Then follow the wire out of the building, because the same mispricing is sitting in the electrons. Power is the one input everyone concedes is scarce, and the honest prices prove it. PJM, the largest grid market in America, just cleared its capacity auction at a record \$329 per megawatt-day, roughly ten times the level of two years ago, pinned against a regulatory cap that kept it from going higher, with data centers driving 40% of the bill²⁶. Microsoft signed twenty years at an estimated \$110 to \$115 per megawatt hour to restart Three Mile Island, and hyperscaler nuclear deals now clear at roughly double wholesale power²⁷. The auction and the contract are the two honest marks for a firm watt, and both have exploded. The equities have not. Vistra, holding more than 5,000 megawatts of hyperscaler agreements including twenty-year Meta nuclear deals, trades at 16 times forward earnings with growth compounding at 37%, a PEG of one half. Talen, whose Susquehanna plant carries a seventeen-year, \$18 billion Amazon contract, trades at about 18 times forward earnings while growing them more than 30% a year²⁸. And the exception proves the rule of this whole essay. GE Vernova, which sells the turbines and books the shortage as revenue today, trades at 63 times earnings after a 243% year²⁹. The market pays the seller of shovels 63 times and the owner of the mine 16. Twenty-year signatures from investment-grade counterparties are the cheapest signatures on the tape.

And underneath the whole ladder, the belief goes negative outright. Short interest against US equities just hit a record \$2.13 trillion. The median S&P 500 stock has not been shorted this heavily since 2011, and total NYSE short interest now exceeds its financial crisis and pandemic peaks³⁰. Nearly half of the managers in that same BofA survey call an AI bubble the market's biggest tail risk, roughly double the share of a month earlier³¹. Put the full tape together. Owned at the index, doubted at the multiple, shorted at the security. The theme is crowded. The belief, walking down rung by rung, thins, sours, and at the bottom flips into an active bet against.

So take the market's side, because it deserves its strongest case. Maybe the earnings really are a peak and a single-digit multiple is fair value on a dying cycle. There is a way to check, and it is the check power traders run every morning. In electricity markets, the spark spread is the gap between the price of power and the cost of the fuel that makes it. As long as the spread is positive, you run the plant. Compute has a spread too, and the best place to read it is the oldest chip still working. The factories essay used the H100 to argue that old chips stay alive. Here is the same machine as a trade. It is four years old now, two full generations behind the frontier. At the peak of the 2023 shortage it rented for \$8 an hour. By last October it had slid to \$1.70, and the depreciation bears looked right on schedule. Since then it has climbed back to \$2.35 on one-year contracts, up almost 40%, because on-demand capacity is sold out across the market³². Now the cost side of the spread. The chip sells for \$30,000 to \$40,000. Spread the middle of that range over four years of round-the-clock work and it costs about a dollar an hour to own, call it a dollar and a half all-in with power and facility. Renting at \$2.35 against a dollar and a half of cost is a positive spread on a machine the accounting says should be nearly worthless, and the spread is widening. And that is only the first spread. The second sits on top, where two dollars of machine time substitutes for fifty dollars of knowledge work. Spreads like that die exactly one way, through a supply glut. So show me the glut. ASML

ships EUV machines in the dozens per year. Advanced packaging is sold out into 2027. Memory is sold out through 2027. TSMC, the one company that sees every real order book on earth, just raised its 2026 growth outlook above 40% and lifted its own capex to \$64 billion³³. The fuel is rationed, the power price is rising, and the market is pricing the plants for the end of demand. And the spread steepens with every generation. CoreWeave's published benchmark of Nvidia's next platform runs ten times the tokens per megawatt of the current one, so the same watt earns more rent with every refresh³⁴. There is an old version of this setup. In the 1860s the money in oil was never in the wells, it was in refining, and the market of that era needed a decade to notice. The refiners of intelligence are sitting in plain sight at five times earnings. The market owns them. It just refuses to believe the spread.

Now the shape of the bet, because the asymmetry is the point of the whole structure. The downside is a multiple that already assumes the earnings die, on order books that run two years past the assumed funeral. The upside is what happens when cycle multiples meet regime earnings, and it is not capped by anything except the size of the buildout. Known downside, uncapped upside. Crowded trades are dangerous when the crowd has priced perfection and has to run for one exit at once. This crowd has priced the funeral. There is nothing to give back, and the repricing, when it comes, runs into the trade rather than out of it. You do not have to imagine what that repricing looks like, because it has already happened in exactly one aisle. NAND prices are forecast to rise 234% this year with the shortage running into 2028, and SanDisk, the purest way to own that corner, rose 726% in the first half of the year, the best stock in the S&P 500³⁵. One aisle repriced, and it repriced violently. And even the repricing was disbelief in disguise. SanDisk's earnings grew twenty-two fold while the stock rose sevenfold, which means the best performer in the index got cheaper as it climbed. The rest of the store is still marked for clearance. My edge is not information. Every number in this essay is public and most of them sit on free websites. The edge is structure and nerve. Underwrite the magnitude consensus refuses to, in the layers where the signatures already prove it, and hold without leverage while the market argues with its own tape.

Strong beliefs deserve tripwires, so here is what would prove the market right and me wrong. HBM contract prices cracking while capacity is still being added. CXMT, China's memory champion, reaching competitive yield at scale years ahead of schedule. Rental rates rolling over instead of rising. Backlogs and remaining performance obligations stalling at the neoclouds, or an anchor counterparty failing to pay. Hyperscaler power contracts and capacity auctions repricing lower at renewal. Token growth flattening while ASML's shipment count ramps. Those are the signatures of a real glut forming, and the day the order books stop growing before the multiples re-rate, the cycle case wins and I will say so in writing. Nothing on the tape says that today.

That is the collection. The machine is real, the constraint is physical, the value pools at the bottlenecks, and the market, for all its crowding, still prices the constraint like a cycle. The gap closes one way or the other. Either the bust arrives on schedule and these essays are wrong, or the earnings keep arriving and the multiple has nowhere left to go but up. I know which side the signatures are on. The only question left is not about the machine at all. It is about who is holding the brush.

SOURCES

1. BofA Global Fund Manager Survey, August 2026 ("Long Magnificent 7" most crowded trade, cited by 45% of 169 managers, \$413B AUM): Reuters; Investing.com.
2. BofA Global Fund Manager Survey, June 2026 (80% of managers long global semiconductors, third-highest conviction reading on record): Benzinga.
3. Goldman Sachs Q2 2026 hedge fund positioning data (largest quarterly increase in net Information Technology tilt on record): Goldman Sachs Prime Services via press reports.
4. Alphabet at ~16x earnings vs. S&P 500 above 20; 2026 capex guidance of \$180-190B: Yahoo Finance; 24/7 Wall St.
5. Amazon 2026 capex raised from \$200B to \$220B on higher memory costs; forward P/E ~29; EV/EBITDA ~14: 24/7 Wall St; App Economy Insights.
6. Meta \$130-145B capex guidance and -7% reaction, ~29% off its highs; Microsoft \$175B guidance and +7% on Azure crossing \$100B of annualized revenue; "less the better" quote (Jason Lemire, Bold Wealth Partners): Fortune; 24/7 Wall St; FinanceFeeds.
7. Oracle capex up 162% to \$55.7B, \$40B of new debt and equity, one-day -12%, ~47% decline since June, S&P downgrade to BBB-, \$638B backlog: Reuters; The Motley Fool; EBC Financial Group.
8. Nvidia forward P/E ~24.5 vs. its three-year average of 52.4 and five-year average of 60.5: GuruFocus; The Motley Fool.
9. S&P 500 forward P/E ~20 (August 2026): FactSet Earnings Insight.
10. Nvidia Q2 FY2027 guidance (\$91B revenue, data center +92%); consensus FY2027 revenue ~\$391B: S&P Global Market Intelligence; Simply Wall St.
11. Micron DRAM and HBM output sold out through 2027; sixteen multi-year customer agreements: company commentary via TS2 and MarketWise.
12. Micron at 5.7x forward earnings (ten-year average 22, up 214% YTD); SK Hynix at 4.9x forward after a ~9x move; Samsung near 7x forward; global DRAM revenue share Q1 2026 (Samsung 38%, SK Hynix 29%, Micron 22%): The Motley Fool; GuruFocus; Investing.com; Korea Times.
13. Super Micro \$60B+ of new orders in fiscal Q4, record backlog, ~17x trailing earnings: 24/7 Wall St; Yahoo Finance.
14. Bank of America 2028 memory scenarios (base case of DRAM -10% and NAND -18% with ~\$150 Micron EPS; severe downturn ~\$100 EPS, roughly 8x the 2018 cycle peak): BofA Global Research via Benzinga and 24/7 Wall St.
15. Semicap 2026 WFE guidance raised ~30% to \$150-160B; consensus \$156B/\$213B/\$262B for 2026-28; sector off 20-40% at ~0.7x three-year PEG: BofA via Investing.com; Morningstar; Morgan Stanley.
16. Western Digital 100% of 2026 HDD production sold out with agreements into 2028-2029; Seagate nearline capacity allocated through calendar 2027: Tom's Hardware; Trefis; Yahoo Finance.
17. Custom ASIC shipments tripling 2024-2027, surpassing GPU shipments by 2027; Broadcom and Marvell at ~95% of co-design: TrendForce via TechTimes; Tom's Hardware; Bloomberg Intelligence.
18. Broadcom AI revenue +143% to \$10.8B in Q2 FY26; \$100B AI revenue target for 2027; ~22x on FY27 estimates: Broadcom via SEC; The Motley Fool; TIKR.
19. Cerebras IPO at ~\$48B (May 2026, \$5.5B raised, +68% day one); subsequent ~28% decline on margin guidance; FY26 core revenue guide ~\$860M: TechCrunch; CNBC; Morningstar.
20. Nvidia licensing Groq technology (December 2025); Qualcomm AI200/AI250 rack-scale inference servers (768GB LPDDR per card, mobile-derived power efficiency, AI200 shipping 2026); Nvidia Rubin Ultra HBM capacity up to ~1TB per GPU: Data Center Knowledge; NAND Research; Spheron; Tom's Hardware.
21. CoreWeave Q2 2026 results (revenue \$2.58B, +112%; revenue backlog \$60.7B at Dec 31, \$99.4B at Mar 31, \$104.2B at Jun 30, +246% YoY, \$129.2B as of Aug 11 including \$25B+ of Q3 commitments; net loss \$626M; FY26 guide \$12.4-13.2B revenue; shares

+14% after hours; ~40%/~80% backlog conversion within 2/4 years): CoreWeave Q2 2026 earnings release (SEC); CNBC, Aug 11, 2026; Investing.com; TradingView.

22. CoreWeave contracted power expanded from 3.7 to 4.2 GW (Q2 2026 call); market capitalization ~\$50B and enterprise value ~\$84B post-report: CoreWeave earnings call via 24/7 Wall St; GuruFocus; StockAnalysis.
23. AI data center build costs (AI-optimized shell and infrastructure \$15-25M/MW excluding GPUs; \$30-45M/MW fully loaded): Axis Intelligence; JLL benchmarks.
24. CoreWeave 2026 capex guidance of \$31-35B against ~\$6.2B of trailing revenue; stock roughly a third below its 52-week high post-report: The Motley Fool.
25. Nebius: Microsoft contract up to \$19.4B, Meta up to \$27B, ~\$50B total contracted backlog vs. \$530M 2025 revenue; \$2B Nvidia equity investment: CNBC; The Motley Fool; TradingKey.
26. PJM 2026/27 capacity auction record \$329.17/MW-day, ~10x in two years, price-capped; data centers ~40% of the \$16.4B cost: Utility Dive; IEEFA; Citizens Utility Board.
27. Hyperscaler nuclear PPAs at \$100-140/MWh on 20-year tenors; Microsoft-Constellation Three Mile Island restart at ~\$110-115/MWh (Jefferies estimate), 20-year, ~\$16B: Jefferies via AOL; industry reports.
28. Vistra (5,000+ MW hyperscaler PPAs incl. ~2,600 MW 20-year Meta nuclear deals; ~37% forecast earnings growth; ~16x forward P/E, ~0.5 PEG); Talen (17-year, ~1.9GW, ~\$18B AWS agreement; ~18.6x forward P/E with ~32% forecast annual earnings growth): TIKR; The Motley Fool; Simply Wall St; GuruFocus; Utility Dive.
29. GE Vernova at ~63x next-twelve-month earnings after a +243% one-year run; record 116 GW gas backlog: TIKR; The Motley Fool.
30. Record \$2.13T of short interest against US equities; S&P 500 short interest ~3.79% of float, highest since 2010; NYSE ~9%, above financial crisis and pandemic peaks; median S&P 500 stock most shorted since 2011: Bloomberg; S&P Global Market Intelligence.
31. Roughly half of BofA survey managers naming an AI bubble the biggest tail risk, up from 28% the prior month: BofA Global Fund Manager Survey via Bloomberg.
32. H100 rental history (\$8/hr peak 2023; \$1.70/hr October 2025; \$2.35/hr one-year contract March 2026, up ~40%; on-demand sold out); purchase price \$30-40K: GPUSmith; IntuitionLabs; CloudZero.
33. TSMC 2026 revenue growth outlook raised above 40%, capex lifted to \$64B, HPC/AI at 61% of revenue: TSMC earnings commentary via Yahoo Finance, Investing.com, CNBC.
34. CoreWeave-published Vera Rubin benchmark at 10x tokens per second per megawatt vs. Grace Blackwell: NVIDIA; Introl.
35. Gartner forecast of NAND prices rising 234% in 2026 with shortage into 2028; SanDisk up 726% in H1 2026 (best S&P 500 performer) on ~22x EPS growth: The Motley Fool; TheStreet; Forbes.

No Offer or Solicitation. The only purpose of this site is to provide general background information on RAXIA. This is not an offer to sell or the solicitation of any offer to buy any securities. Any such offer would be made only to qualified investors after distribution and return of the Private Placement Memorandum and governing agreement for the applicable Fund (the "Offering Materials"). Nothing on this site is investment advice or a recommendation of any security. The author may hold positions in companies discussed. Past performance is not indicative of future results.