



monsitj/iStock via Getty Images

While the AI bears focus on concentration and circular financing, the last tech overbuild was financed with debt, and this one is being paid for in cash.

Before I discuss why I disagree with the “AI bears,” I want to state that I respect their opinions, have evaluated their concerns, and have simply derived a different set of conclusions. That is an important statement, because this particular group of “AI bears” includes some of the sharpest risk minds in the business, and they have been early to almost every warning that later mattered.

When people this good line up on one side of a trade, you go back and check your own work. That’s what I did, and this article is where I landed. As always, the reason I publish these articles is for accountability later, for you and our clients.


While this group of AI bears may indeed be right about the excess, they could still be potentially wrong about the trade. I care about the latter, and those are two different claims that the market keeps confusing.

The Bear Case Deserves A Hearing

Let's start with the person I admire the most in the AI bear camp: Fred Hickey. Fred has run The High-Tech Strategist since 1987 and has made the cleanest version of the argument. He compares today's datacenter mania to the fiber-optic overbuild that cracked in 2000, only far larger. To wit: he has called it a *"more dire situation than the great fiber-optic capacity overbuilds."*

He is not alone in this view, and that really is the point to address. Michael Burry has been circling the same plumbing, watching Nvidia's (NVDA) credit-default swaps widen as the chipmaker turns into a banker, landlord, and equity partner to its own customers.

But the AI bear roster doesn't stop there. The Bank for International Settlements flagged roughly \$1.65 trillion in off-balance-sheet obligations held by the largest hyperscalers, exceeding the amounts they carry on their books. Then Sequoia's David Cahn put the annual gap between AI infrastructure spending and ecosystem revenue at nearly \$600 billion. Furthermore, Allianz measured the capex-to-revenue divergence at about 46%, well past the 32% that marked the 2001 telecom bust. Then, lastly, in August, an MIT study suggested that most corporate AI pilots had produced no measurable revenue at all.




The gap between what the AI build-out spends and what the ecosystem earns runs near 46% in 2026, wider than the roughly 32% divergence that accompanied the 2001 telecom bust. Source: Allianz Research, 2026.

That is a very serious AI bear group making a very serious case, and you should only ignore it at your peril. When a strategist who has correctly traded five separate Nvidia collapses of 55% or more says a sixth is coming, and a Bank of America survey shows 54% of professional managers are now calling AI a “*bubble*,” you need to factor that into your thinking. As investors, we must work out precisely which parts are right and which parts are borrowed pattern-matching from a different era.

So, let's start with where the AI bears are right.

Where The Bears Are Right

Yes, valuations are stretched, and by the measure that matters most for fragility, concentration is worse now than it was in 2000.



The ten largest S&P 500 stocks now sit near 43% of the index, a record, and well past the roughly 27% peak reached during the dot-com mania. Source: RBC Wealth Management, Guinness Global Investors, InvestmentNews, 2026.

Notice how far the line has traveled in the chart above. The ten largest stocks now make up roughly 43% of the S&P 500, a record, and past the 27% peak the index touched at the height of the dot-com boom. *By that single measure, the market is more top-heavy today than at any point in modern history.* The equal-weight index has already begun to diverge from the headline benchmark, which is exactly the kind of internal crack that tends to show up before the megacaps wobble. Such is the setup the AI bears keep pointing toward, and on that point, they are correct.

Secondly, the circular financing argument is real, too. When Nvidia takes an equity stake in a company that then commits to buying Nvidia chips, part of what gets reported as “*demand*” is the seller funding its own sales. Such is a genuine distortion of the signal, and it deserves the scrutiny that it has been getting. Add the [depreciation math](#), where trailing capex of roughly \$434 billion dwarfs the \$149 billion of depreciation currently running through income statements, and you get a bill that arrives in 2027 through 2029, whether the revenue does or not. The AI bears did not invent any of this; they just read the corporate filings.


Where The Analogy Breaks

So, with all that stated, it seems to be obvious that you should just get out of the AI trade now before the next “*Dot.com*” crash occurs. Here’s the problem with that comparison. The comparison to the fiber-optic “*boom and crash*” is that it turns on the one variable that actually determined the outcome in 2000, and that variable does not read the same today: *who is writing the checks*.

Leading up to the 2000 overbuild, the financing came from companies that had no business borrowing what they borrowed. WorldCom, Global Crossing, and the upstart carriers that were stringing fiber on debt, and the vendor loans that Lucent and Nortel handed to customers who could not pay them back. When revenue failed to arrive on schedule, those balance sheets could not cover the shortfall, and the structure collapsed into bankruptcy court.

But this isn't WorldCom: who's paying for it

Funding of the ~\$725B 2026 hyperscaler AI build-out



Source: FactSet (incremental debt as a share of capex, LTM mid-2026). Figures approximate.

Roughly two-thirds of 2026 hyperscaler AI capital spending is funded from operating cash flow and equity, with new debt the minority slice and investment-grade. Source: FactSet, mid-2026.

Today's buildout is a different animal on this exact axis. Roughly two-thirds of the 2026 capex is funded directly from the operating cash flow and equity of Microsoft ([MSFT](#)), Alphabet ([GOOGL](#)), Amazon ([AMZN](#)), and Meta ([META](#)), four of the most profitable enterprises ever assembled. The existing borrowing is investment-grade and still a minority of spending. *The balance sheets carrying this cycle are not WorldCom's, and that difference is close to the whole ballgame.*

Revenue Is Real

Second, "no revenue" is not the same thing as revenue that simply hasn't caught up to the spending yet. Inference now clears roughly 70% gross margins. Microsoft's AI business is past a \$37 billion run rate, Amazon's AI revenue is growing in the triple digits, and Anthropic went from about \$9 billion to a reported \$47 billion run rate in a single year.

More notably, even Nvidia, the bears' favorite "whipping boy," has seen forward earnings climb so rapidly that its multiple has actually compressed as fundamentals caught up to what was believed to be overly exuberant expectations. That is the mirror image of Cisco in 2000, which peaked at nearly 30 times sales on earnings that then evaporated. *The revenue-trailing capex is a timing issue, not the zero-payback story the headlines imply.*

Unlike 2000, the revenue is real

The dot-com leaders were priced on revenue that never came. Today's earn it.

2000

The dot-com leaders were priced near 30x sales.

The earnings never justified it, and many had no revenue at all.

2026

\$81.6B

Nvidia revenue last quarter, +85% YoY

~70%

Inference gross margin

~\$175B

AI ecosystem revenue run-rate

Sources: NVIDIA (Q1 FY27); SemiAnalysis; industry estimates. 2026, approximate.

The dot-com leaders were priced on revenue that never arrived. Today's AI leaders post large, fast-growing revenue: Nvidia alone booked \$81.6B last quarter, up 85% year over year. Sources: NVIDIA, SemiAnalysis, industry estimates, 2026.

Third, the AI bears predict a glut, yet the binding constraint right now is the opposite of a glut. Microsoft is sitting on something like \$80 billion of Azure orders it cannot fill for lack of electricity, with GPUs idle in inventory waiting on power.

Today, more than 60% of the data center capacity planned for 2027 is not yet under construction. If or when datacenter demand is rationed by the power grid rather than by customers walking away, you do not have a capacity glut; you have a shortage.

However, a fair objection at this point, and it is the strongest one the bears have: build two or three years' worth of power and transmission, and today's shortage becomes tomorrow's oversupply. That is true concern, and it is a timeline risk worth watching closely, but it is also a 2028 question, not a 2026 one.

	FIBER & DOT-COM, 2000	AI BUILD-OUT, 2026
WHO FUNDS IT	Debt-laden carriers, vendor loans	~68% operating cash flow and equity
BUILDERS' PROFITABILITY	Many cash-flow negative, pre-earnings	Most profitable firms in history
PEAK VALUATION	Cisco near 30x sales	Leaders below that on price-to-sales
EARNINGS UNDER THE LEADER	Collapsed after the peak	NVDA earnings up sharply, multiple compressed
BINDING CONSTRAINT	Too much capacity, too little demand	Too little power, demand rationed

What The AI Bears Debate Means For Investors

Let me be clear. The AI bears have a real case but no timing. This is the same problem we noted in “*Debt Trap: A Crisis Without A Calendar.*” I am definitely not arguing that investors should be buying the AI complex with both hands and closing their eyes. The question is *not* whether there is excess, because there plainly is. The real question is what a disciplined investor does with a genuine, extreme, but cash-funded overbuild.

Start with position sizing, because it is the one tactic that survives contact with a drawdown. Nvidia has fallen by 55% or more on five separate occasions since 2000, and it has recovered to new highs after each. Investors who were sized to hold through the pain benefited tremendously. They did even better if they managed their exposure risk during those drawdowns. *Own your AI exposure at a weight where a 50 percent drawdown is uncomfortable rather than fatal. Sizing comes first.*

Secondly, the rules are simple.

- Favor the self-funders over the borrowers, and*
- Spread your exposure across the layers of the trade, the chips and the clouds, and the power underneath them, rather than staking the whole thesis on a single chip name.*
- Always insist that the price you pay is backed by existing earnings and not by a total addressable market slide.*

Lastly, keep some dry powder (ie, cash), because the volatility in this complex is a feature rather than a defect, and a real correction turns into a gift the moment you have cash and a shopping list ready.

Where you take the risk matters as much as how much you take. Not all AI exposure carries the same danger, and the map below is how I would sort it.

Where to take the AI risk

Same trade, very different risk. Own the self-funders. Rent the rest with care.

SAFER TO OWN

Hyperscalers funding their own capex out of operating cash flow.
Power and the grid, the real bottleneck: utilities, power gear, cooling.

OWN, BUT SIZE IT

The chipmakers with real, growing earnings and pricing power.
Watch price and concentration. Trim the winners back to target weight.

HIGHER RISK: TRIM OR AVOID

Unprofitable model labs carrying commitments they can't fund.
Neoclouds and single-tenant SPVs levered to depreciating chips.

RIA Advisors. Segments shown for illustration, not individual investment advice.

Where the AI trade looks least like 2000 (self-funders and the power bottleneck) versus where it looks most like it (levered, single-tenant, pre-profit names). For illustration, not individual advice.

The 5 Signals

The self-funders and the power bottleneck are part of this trade that looks least like 2000. Conversely, the levered edges are the part that looks most like it. That levered part is where a revenue disappointment does the real damage, and those are the first positions to shed when the story starts to wobble. *The profitable compounders funding their own buildout sit in a different bucket, and selling them because a bear called a top is how investors miss years of compounding while waiting on a crash that shows up late, or never.*

Which raises the harder question. *How do you know when the story is actually wobbling?*

This is crucial and the trap that most investors fall into. *You do not need to call the top.* What you need is a short list of signals that fire before the top is obvious to everyone and the discipline to act on the list rather than argue with it.

Five signals to start getting out

You don't need to call the top. You need a checklist that fires before it's obvious.

- 1 Hyperscaler capex guidance gets cut**
A 15% cut can take roughly 22% off Nvidia's forward earnings.
- 2 AI revenue growth stalls**
The ecosystem ramp flattens, or a headline target gets missed.
- 3 Credit stress hits the plumbing**
AI credit-default swaps widen, or a neocloud or SPV cracks.
- 4 The depreciation wall bites**
2027 to 2029 charges outrun revenue and squeeze margins.
- 5 Breadth breaks or liquidity is pulled**
Equal-weight rolls over, or the Fed drains the punchbowl.

Sources: Morgan Stanley; company disclosures; RIA Advisors. 2026.

A short de-risking checklist. The single most important trigger is a cut to hyperscaler capex guidance, which resets the revenue math for the whole chain. Sources: Morgan Stanley, company disclosures, 2026.

Which brings me to the question I get most often: *“Why not skip the stock-picking and just own the index?”*

Here is my opinion. The index has quietly become the *“bet.”* With the ten largest names accounting for nearly 43% of the S&P 500, buying the market today is a concentrated wager on those same few companies, made passively, without anyone ever deciding it was a good idea. *Owning the index is not a way to sidestep the AI trade because it is the AI trade, whether you meant it that way or not.*

[Bob Farrell's Rule #9](#) is always worth repeating here:

“When all the experts and forecasts agree, something else usually happens.”

Conclusion

With more than half of managers now calling AI a “bubble” and “long the Magnificent 7” ranked the most crowded trade on the Street for nearly two years, the consensus has already tilted bearish. That does not make the AI bear case wrong, but it does suggest the obvious crash may refuse to arrive on the obvious schedule.

One of my favorite quotes from Howard Marks is that “being *too far ahead of your time is indistinguishable from being wrong.*” When it comes to investing, timing is critical. *Most importantly, notice that Hickey himself holds his AI-bear book at roughly 1% of his portfolio in puts, suggesting he treats it as a hedge rather than a conviction short. That is the posture worth borrowing.* Own the compounders, hedge the tail, and let the revenue prove or disprove itself on the tape.

The AI bears will eventually be right about a drawdown, because everyone is eventually right about a drawdown. Whether they are right about the trade depends on a question their favorite analogy cannot answer:

“What happens when the richest companies on earth overbuild with their own money rather than borrowed money?”

Such is the question actually on the table, and that is the question you must answer before you sell.

Sources & Notes

1. *Fred Hickey commentary: CNBC (Feb 2026); The Market Ear via Benzinga (five NVDA collapses, sixth expected); Kitco News (Jul 2026).*
2. *Circular financing and off-balance-sheet obligations: Bloomberg; 24/7 Wall St. citing the Bank for International Settlements 2026 Annual Report and Nikkei (~\$1.65T off-balance-sheet).*
3. *Capex-to-revenue gap: Sequoia Capital (David Cahn, ~\$600B); Allianz Research (46% divergence vs 32% in the 2001 telecom cycle).*
4. *Funding mix and depreciation: FactSet (incremental debt ~32% of capex, LTM mid-2026); SiliconAnalysts / SEC filings (trailing capex ~\$434B vs ~\$149B depreciation).*
5. *Inference margins and AI revenue: SemiAnalysis via Forbes (~70% inference gross margins); company disclosures (Azure AI >\$37B run rate, Anthropic ~\$9B to ~\$47B).*