THE WEEKLY TAPE

Deep-Dive Edition | Week ending Friday, April 10, 2026

EXECUTIVE SUMMARY

The market has split the AI trade in half — paying all-time highs for the picks-and-shovels (SOXX, NVDA, AVGO) while pricing the diggers (MSFT, META) as if their AI strategy is broken. This deep-dive walks the actual cash-flow math behind both sides. The conclusion is unchanged from last week's brief but the conviction is stronger: buy the hated hyperscalers. Do not chase the chips at these levels.

THE FIVE THINGS THAT MATTER

1. The capex super-cycle is real and accelerating. Combined MSFT + META + GOOGL annual capex went from $53B in 2020 to **$226B in 2025** — a 4.3x increase, with the steepest year-over-year jump (+68%) coming in 2025 itself. This is not a one-off; it is a structural rebuilding of the world's compute infrastructure.

2. Revenue growth at all four names is accelerating, not decelerating. The single most important data point in this report. NVDA went from +56% YoY in Q2 to **+73% in Q4**; META from +16% to +24%; GOOGL from +12% to +18%; MSFT steady at +17%. The bear thesis — that AI spend isn't producing growth — is contradicted by the actual income statements.

3. The FCF squeeze is real but temporary. Operating cash flow at the hyperscalers is growing rapidly; free cash flow has stalled. Microsoft converted 75% of OCF to FCF in 2020; in 2025 that ratio collapsed to 53%. Capex is the entire delta. The key word is *timing*: OCF growth tells you the underlying business is fine; the FCF stall tells you when the spending normalizes, FCF mean-reverts upward violently.

4. NVDA is now 90% data center. Up from 21% five years ago. Every dollar of NVDA's thesis at this point depends on a small handful of hyperscaler and lab customers continuing to spend at current rates. The concentration risk is greater than it has ever been, and it is mispriced because the market keeps treating NVDA as 'diversified semis.' It isn't.

5. The leverage pivot is the canary nobody is watching. META and GOOGL *combined* issued $59B of net new debt in 2025 — both companies that historically ran net cash. They are now funding capex with the bond market because operating cash flow alone can no longer keep up. This is a regime change for capital allocation in mega-cap tech, and it tells you the spending is real.

THE CALL

The market is paying record multiples for the equipment makers selling AI infrastructure and discounting the buyers building it. Both views cannot be right. The same hyperscaler capex that drives NVDA's revenue is the line item the market is punishing MSFT, META, and GOOGL for. This week's Iran-ceasefire relief rally took the existing dispersion and amplified it into a 44-percentage-point spread between SOXX and MSFT year-to-date.

The fundamental data assembled in this report makes a stronger case than the price chart alone. We are not arguing that the chip cycle is broken; we are arguing that the relative value has shifted decisively to the hyperscalers, and that the market's reflexive 'AI capex doesn't pay' narrative is contradicted by every revenue line we can pull.

MARKET BACKDROP — THE 44-POINT SPREAD

A nasty Q1 ended in a tradeable dislocation. The Iran conflict drove broad risk-off in February and March; the hyperscalers were hit harder than the market because the AI capex story compounded with the macro fear. Microsoft fell 23% in Q1 alone — the worst Mag-7 print of the quarter and the steepest single-stock drawdown in mega-cap tech since late 2022. Then on April 8, a U.S.-Iran ceasefire lit up semis: the index ripped 6-11% in a single session as risk assets unwound the war premium. Mega-cap hyperscalers barely participated.

Figure 1 — YTD 2026 Performance: A 44-Point Spread

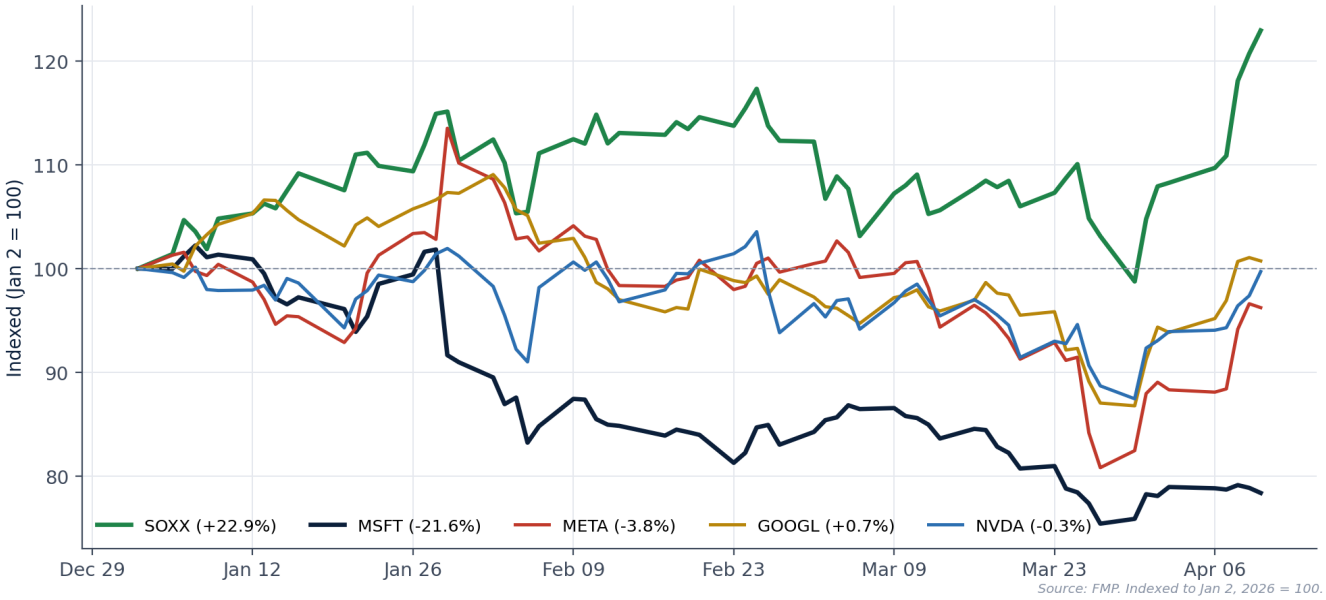

Figure 1. *Year-to-date 2026 normalized to Jan 2 = 100. SOXX +22.9% versus MSFT -21.6% — a 44-point dispersion inside what is supposedly the same trade. The cliff in MSFT in late January is the FY26 Q2 print where capex was guided sharply higher; that single tape-read is the inflection point of the entire year. Source: FMP.*

Three things to note. First, MSFT's January cliff is the late-January earnings print where capex was guided higher — that single tape-read is the inflection of the entire dispersion. Second, semis spent *most* of the year above their starting level even during the broader tech selloff — chips were the resilient sub-sector throughout, then exploded higher this week on the ceasefire. Third, GOOGL is the only mega-cap roughly flat YTD; it has neither the chipmaker rip nor the hyperscaler drawdown, which is exactly why we keep recommending it as the quality compounder of the group.

Rates context: 10-year at 4.29%, 2-year at 3.78%. The curve steepened modestly through Q1 (10-year low ~4.05% in early March, now ~4.30%) — long end repricing on fiscal worries, short end stable as the Fed holds.

Critically: this is *not* the rate environment that causes tech selloffs. If MSFT were down on rates, the bounce-back trade would be obvious. It's down on AI ROI fear, which is a harder problem and a better opportunity.

I. THE CAPEX SUPER-CYCLE IN NUMBERS

Before debating whether this spending will pay off, it is worth grounding ourselves in just how large 'this spending' has become. The numbers are staggering and underappreciated outside of the people who track them weekly.

Figure 2 — Hyperscaler Capex Super-Cycle: $53B→226B in 5 Years

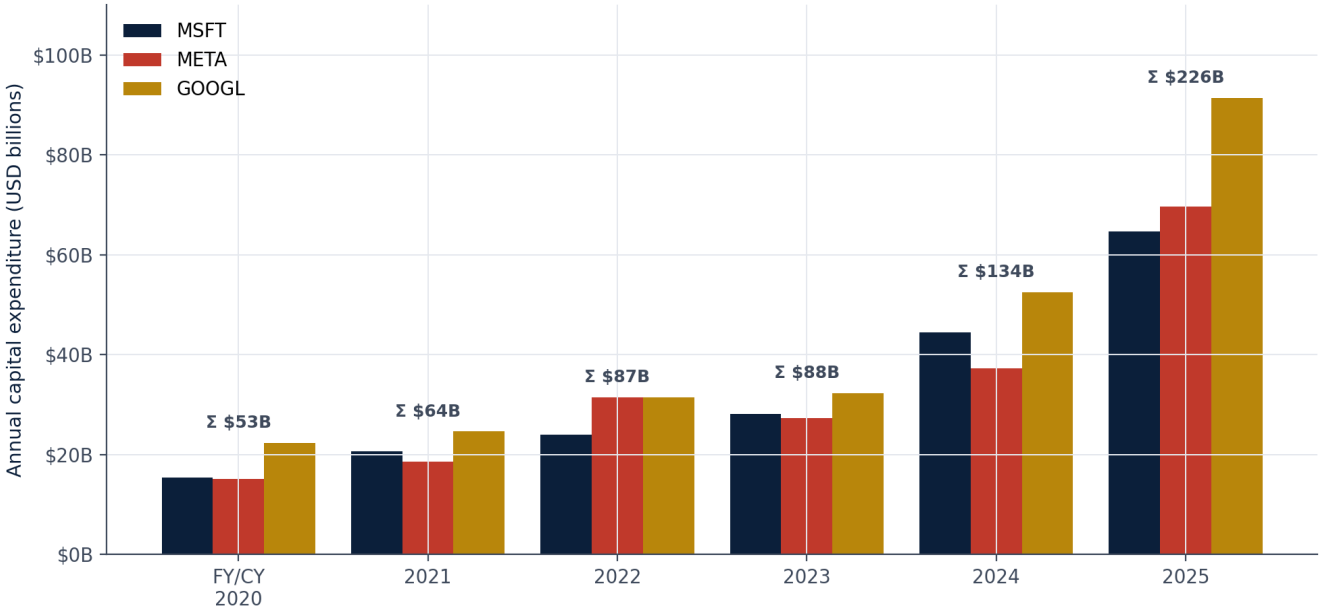

Source: FMP annual cash flow statements. MSFT on FY (June year-end), META & GOOGL on calendar year.

Figure 2. *Annual capital expenditure for the three largest hyperscalers in our coverage. MSFT runs on a June fiscal year; META and GOOGL on calendar year. Combined capex went from $53B in 2020 to $226B in 2025 — a 4.3x increase. The steepest year-over-year jumps came in 2024 and 2025. Source: FMP annual cash flow statements.*

Two observations from the chart make the magnitude visceral. First, the combined 2025 number ($226B) is more than the entire 2020-through-2022 spending of these three companies combined ($203B). They are spending in *one year* what they spent over *three* just five years ago. Second, all three companies are acting in unison. There is no laggard. There is no holdout. The capex pivot is sector-wide and synchronized — and it is structurally different from the asynchronous capex cycles of the prior decade.

Year-over-year capex growth — the acceleration is visible in the deltas:

	2021	2022	2023	2024	2025
MSFT (FY end June)	+34%	+16%	+18%	+58%	+45%
META	+23%	+69%	-13%	+37%	+87%
GOOGL	+10%	+28%	+2%	+63%	+74%
Combined	+21%	+36%	+1%	+53%	+68%

Table 1. *Year-over-year capex growth by hyperscaler. The 2024-2025 acceleration is the structural inflection.*

The 2024-2025 leg is what changed everything. Before 2024, capex grew at high-teens to 30s annually — fast, but normal for a growth business reinvesting. The 2024 step-up to +53% combined growth — followed by another +68% in 2025 — broke the historical pattern. This is no longer 'reinvestment.' This is a strategic land

grab where each company is racing to build out compute capacity faster than the others, on the bet that scale advantage in AI infrastructure will determine the winner of the next decade.

The bear case on these stocks rests on the simple intuition that companies cannot indefinitely accelerate capex faster than they grow revenue without destroying returns. That intuition is correct in general — but it requires the second clause: *faster than they grow revenue*. And as we will show in the next section, revenue at these companies is also accelerating.

It is also worth flagging the historical analogy that keeps coming up in our internal conversations. The 1999-2001 telecom capex cycle saw network operators spend approximately $120B annually at the peak (1999-2000) on equipment from Cisco, Lucent, JDSU, and Nortel. The current AI capex cycle is already running at *nearly twice that pace* in nominal dollars, and approximately 1.4x in inflation-adjusted terms. The scale is bigger; whether the durability is too is the central question.

II. THE BEAR CASE IS WRONG: GROWTH IS ACCELERATING

If you read the financial press for the past sixty days, you would believe that AI monetization is failing — that hyperscalers are spending hundreds of billions on infrastructure that is not producing commensurate revenue growth. This narrative is easy to write because it fits the price action. It is also wrong on the actual numbers.

Figure 3 — Revenue Growth Is Accelerating, Not Decelerating

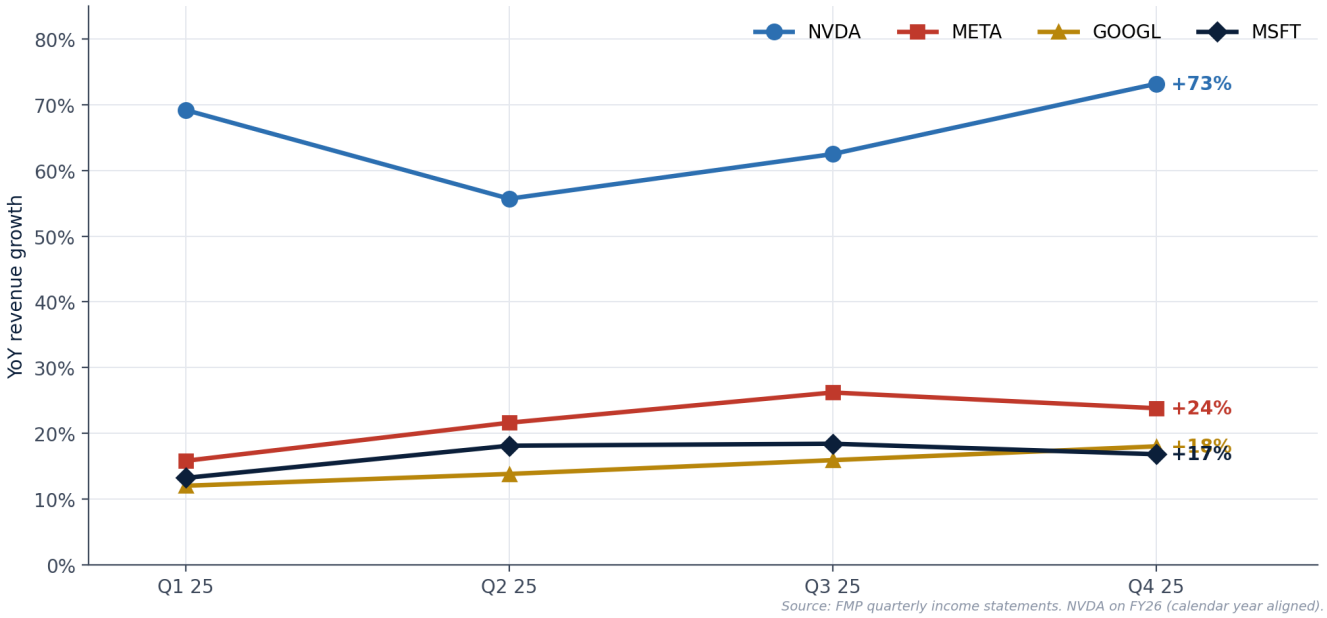

Figure 3. *Year-over-year revenue growth, last four calendar quarters. Every name in the complex is accelerating or holding double-digit growth. NVDA had a brief deceleration in Q2 (the 'air pocket' bears talked about) and is now reaccelerating into Blackwell ramp. The supposedly broken hyperscalers are growing 17-24% YoY at $200B-$400B revenue scale. Source: FMP quarterly income statements.*

Read this chart slowly because it is the spine of our entire thesis. Every line is moving up and to the right. The narrative says NVDA is decelerating; the data says it bottomed at +56% in Q2 and reaccelerated to +73% in Q4 on Blackwell ramp. The narrative says META is broken; the data says revenue accelerated from +16% to +24% over four quarters. The narrative says MSFT can't grow into its capex; the data says revenue compounded at +17% on a $300B+ run-rate while operating margin *expanded* 100bps year-over-year. The narrative says GOOGL is an AI loser; the data says it accelerated from +12% to +18% as Gemini and GCP took off.

The single most important question this chart answers: Are the hyperscalers earning a return on the capex super-cycle? The answer is yes — they are earning it as accelerating top-line growth at scales nobody thought possible at these revenue bases. The cash conversion is delayed (we'll get to that next), but the revenue itself — which is the leading indicator of cash conversion — is fine.

The market is essentially saying: 'I don't believe this growth will continue.' That is a defensible position. But it is a much weaker position than 'this growth isn't happening,' which is what the price action implies. There is a meaningful difference between a stock priced for ~10% terminal growth and a stock priced for 20%+ continued growth, and right now MSFT and META are trading much closer to the former than the latter.

MARGIN INTEGRITY CHECK

A critical sub-point: the capex super-cycle is hitting the P&L; through depreciation, but operating margins have not collapsed. MSFT's operating margin in the most recent quarter was 47.1% — versus 45.5% one year earlier. META's quarterly operating margin held above 41%. GOOGL's stayed above 30%. These are world-class margin profiles, intact despite the largest capex acceleration in their corporate histories. The pricing

power is real, the operating leverage is real, and the moat is intact. The market is paying 14-21x EV/EBITDA for these businesses while assuming the margin structure is at risk. It isn't.

III. THE FCF SQUEEZE: TIMING IS NOT BREAKAGE

If revenue and margins are intact, where is the pain coming from? It is coming from one specific line on the cash flow statement: free cash flow. And the FCF compression is real, it is the actual reason these stocks are down, and it is the load-bearing piece of the bear case that we have to address head-on.

Figure 4 — The FCF Squeeze: Cash Generation Up, Free Cash Flow Flat

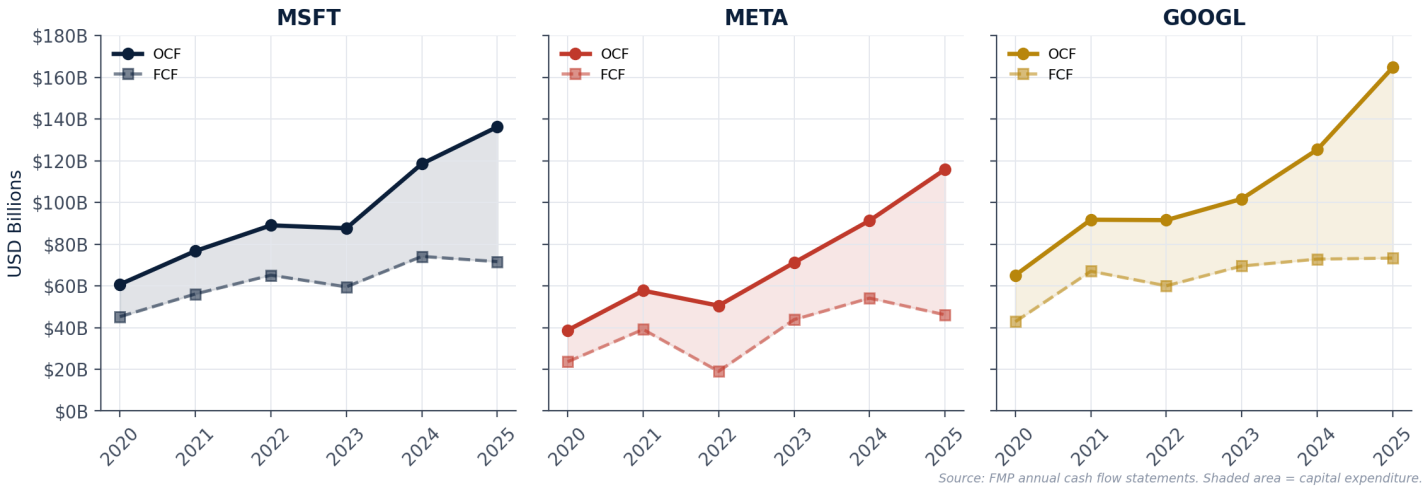

Figure 4. Operating cash flow (solid) versus free cash flow (dashed) for the three spenders, 2020-2025. The shaded area between the two lines is annual capital expenditure — and you can see it widening dramatically from 2023 onward. Source: FMP annual cash flow statements.

Here is the math, in plain English. In 2020, Microsoft generated $61B of operating cash flow and converted $45B of that to free cash flow — a 75% conversion rate. In 2025, it generated $136B of OCF (more than double) but converted only $72B to FCF — a **53% conversion rate**. Operating cash grew 124% over five years; free cash flow grew only 58%. The entire delta is capex.

The story is starker at META. OCF tripled from $39B to $116B over the same period. FCF barely doubled, from $24B to $46B. The conversion rate collapsed from 61% to 40%.

GOOGL is the same picture, slightly delayed. OCF more than doubled to $165B; FCF grew only 70% to $73B. Conversion rate dropped from 66% to 45%.

	2020 OCF	2025 OCF	OCF Δ	2020 FCF	2025 FCF	FCF Δ	FCF/OCF 2020	2025
MSFT	$61B	$136B	+124%	$45B	$72B	+58%	75%	53%
META	$39B	$116B	+199%	$24B	$46B	+95%	61%	40%
GOOGL	$65B	$165B	+153%	$43B	$73B	+71%	66%	45%

Table 2. The FCF compression in numbers. Operating cash flow generation is exceptional; free cash flow conversion is the worst it has been in a decade — and capex is the entire reason.

WHY THIS IS A TIMING ISSUE, NOT A BREAKAGE ISSUE

The natural question: if FCF is compressing, isn't the bear thesis correct? The answer is subtle and matters a lot. There are two ways FCF can compress:

(a) The buy side (capex) grows faster than the sell side (revenue). This is genuinely bad — it means the business is spending money to grow that doesn't translate to sales. This is what happened to telecom carriers in 1999-2001. They spent on fiber that nobody filled.

(b) The buy side (capex) grows faster than the realization curve of the sell side. This is a timing problem, not a value problem. Capex hits the cash flow statement on day one; the revenue from that capex accrues over the asset's useful life — typically 5-7 years for a data center. So in any year where capex accelerates, FCF compresses even if the underlying ROI is great. This is the situation the hyperscalers are in.

How can we tell which is which? **We look at revenue growth.** If capex is genuinely wasted, revenue stalls. If capex is just front-loaded vs. its revenue tail, revenue keeps growing. Figure 3 (the previous page) is the answer to this question. Revenue is *accelerating*. The capex is not wasted; it is being recognized on the cash flow statement before the revenue catches up.

This is the mechanical reason we are bullish: when capex normalizes — even slightly — the FCF mean-reverts upward violently. A 10% capex moderation at MSFT in 2026 would translate to an additional $6.5B of FCF (roughly +9%). A flat capex year would translate to closer to +25% FCF growth. The market is currently pricing zero credit for this mean reversion.

IV. THE OTHER SIDE OF THE TRADE: NVDA's CONCENTRATION RISK

Having defended the buyers, we have to be honest about the supplier. NVDA is the most important stock in this complex and it is the one we are most cautious on at current levels. The reason is not its multiple (which is rich but defensible) or its growth (which is spectacular). It is the customer concentration that comes with $200B of data-center revenue from a handful of buyers.

Figure 5 — NVDA's Single-Customer-Cluster Concentration: 90% from Data Center

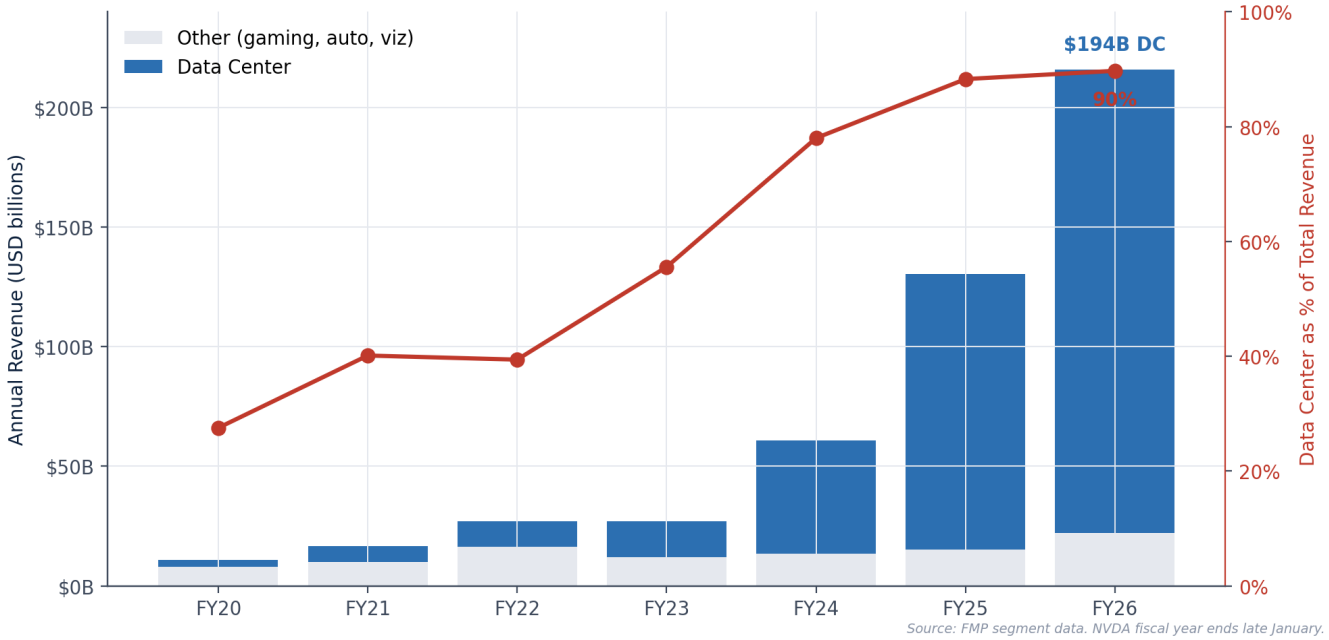

Figure 5. NVDA's revenue mix evolution. Data center went from $3B (21% of revenue) in FY20 to $193.7B (90% of revenue) in FY26. The blue bars show the data center super-cycle in absolute dollars; the red line shows the percentage of total revenue. Source: FMP segment data.

Five years ago, NVDA was a gaming company with a small but exciting data-center business. Today it is a data-center company with a small gaming business. Data center revenue grew from $3B to $194B — a **65x increase**. Gaming, by comparison, grew from $5.5B to $16B in the same period — about 3x. Auto, professional visualization, and other segments are essentially rounding errors at this scale.

The concentration is even more extreme inside the data center segment. NVDA does not disclose exact customer breakdowns, but consensus estimates from sell-side reports suggest the top five customers are responsible for 50-60% of data center revenue, with the top ten accounting for 75-80%. Those top customers are: MSFT (Azure / OpenAI), Meta, Google (GCP / internal), Amazon (AWS / Anthropic), Oracle (Stargate), and a handful of neoclouds (CoreWeave, Crusoe, Nebius). The buyer base is small enough to fit at one conference table.

The PM implication: NVDA's $200B revenue line is not 'diversified semiconductor demand.' It is a concentrated bet on a small group of customers, most of whom are the same companies the market is currently saying have broken AI strategies. If MSFT or META cuts their FY27 capex guidance by even 10% — let alone the 20%+ that would be required to actually discipline the spend — NVDA's order book takes a meaningful hit. The supplier's revenue is, by definition, the buyer's expense, and the market cannot maintain forever the position that the buyers' expense is too high while the supplier's revenue is fairly priced.

We are not bearish NVDA. We think the long-term story is real and the moat (CUDA, Blackwell, ecosystem, software) is genuine. We are simply observing that buying NVDA at the highs, into a customer concentration

that has never been more extreme, into a tape that has never been more euphoric about the chip side of the trade, is the worst possible entry point. We will look at NVDA seriously on any pullback to the 50-day moving average ($182, ~4% lower) or, better, to the 200-day ($180, ~5% lower).

WHY WE'RE NOT SHORT NVDA EITHER

A logically consistent bear on hyperscaler capex would short NVDA. We are deliberately not. The reason is the AI labs layer — a demand pillar that sits upstream of the hyperscalers and doesn't necessarily unwind even if MSFT or META cut their internal spend. OpenAI's Stargate commitment, Anthropic's expanded AWS+GCP contracts, xAI's Colossus build, Meta's Llama training infrastructure — these are real, contractually-committed compute demand sources that route through hyperscalers (or directly through neoclouds) into NVDA's order book. *Disclosure: this report was assembled with help from Claude, which is built by Anthropic. Take Anthropic-specific framing accordingly.*

The lab layer makes the demand more durable than a pure 'hyperscaler passthrough' thesis would suggest. It is the reason we won't short SOXX. But the lab layer also moves the question up one level: it converts 'will hyperscaler ROI materialize?' into 'will lab revenue ramps justify the compute commitments funding the lab buildout?' The labs are private. Inference pricing is collapsing. Enterprise adoption is real but slower than the hype implied. The entire chip super-cycle is now resting on a handful of private companies whose ability to monetize $100B+ of compute by 2028 is the central, unanswerable question.

V. THE LEVERAGE PIVOT: A REGIME CHANGE NOBODY IS TALKING ABOUT

Here is the data point that has convinced us this capex cycle is structurally different from anything the hyperscalers have done before. For the entire decade leading up to 2024, Microsoft, Meta, and Google were net *repayers* of debt. They generated more cash than they could deploy and used the excess to retire bonds, buy back stock, and pay dividends. Capital allocation flowed *out* of the bond market and back to equity holders.

That changed in 2024 and accelerated dramatically in 2025.

Figure 6 — The Leverage Pivot: Hyperscalers Now Issuing Debt to Fund Capex

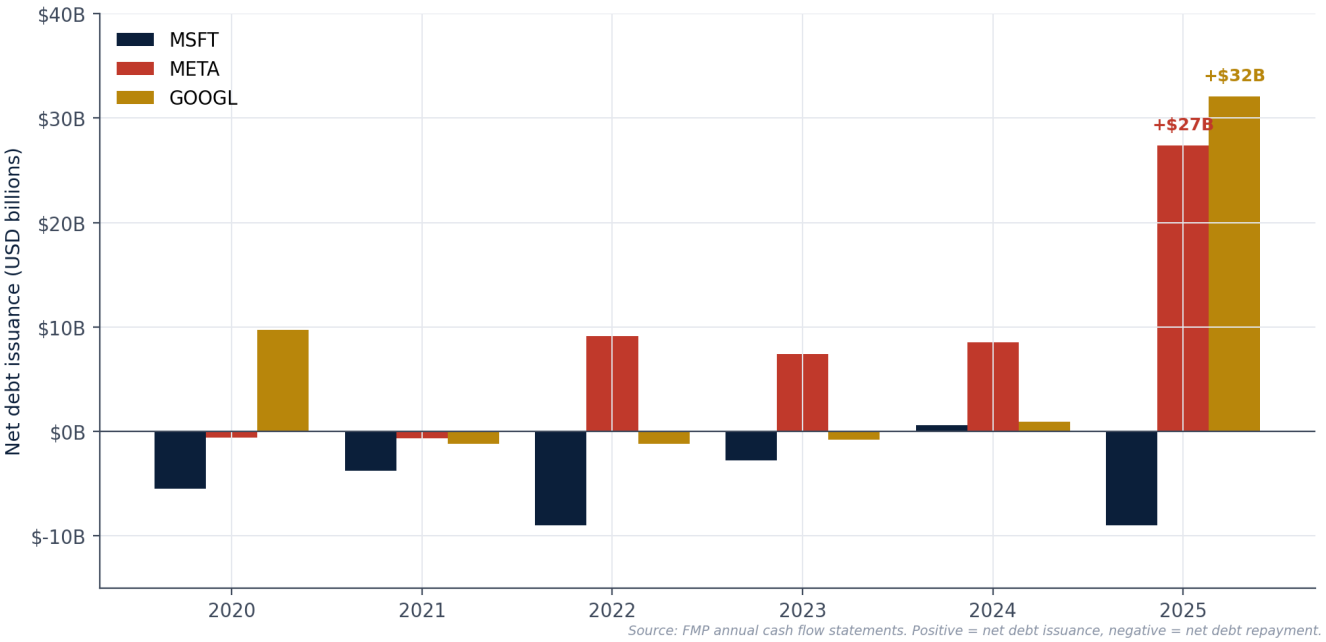

Figure 6. *Net debt issuance by hyperscaler, 2020-2025. Values above zero represent net new debt issued; values below zero represent net repayments. Note the 2025 spike in META and GOOGL — both names that historically ran net cash. Microsoft is still net-repaying. Source: FMP annual cash flow statements.*

In 2025, META issued $27B of net new debt. GOOGL issued $32B. **Combined, $59B of new leverage was added to two companies that historically ran net cash positions.** Microsoft, by contrast, is still net-repaying — but it has the benefit of running a June fiscal year, and we expect its FY26 (June 2026) cash flow statement to show a similar pivot once it lands.

Why does this matter? Three reasons.

First, it confirms the spend is real. Companies do not issue tens of billions of dollars of debt for projects they don't believe in. The capex pivot is not a corporate-comms exercise; it is a balance-sheet commitment. META and GOOGL just stamped their names on a thesis that AI infrastructure spending will produce returns large enough to justify leveraging up. That is meaningfully different from announcing a capex number on an earnings call.

Second, it changes the equity-holder math. When the hyperscalers were net-repaying debt and buying back stock, equity holders received the entire delta as multiple expansion. Now that delta is being absorbed by capex and bondholders. The reason FCF-to-equity went negative for MSFT in TTM ($-22B) is that capex and shareholder returns combined are running above OCF — the gap is being plugged by either cash drawdown or, in META and GOOGL's case, debt issuance. This is a temporary phenomenon, but for as long as it lasts, equity

holders feel the squeeze.

Third, it creates a hard ceiling on how long this can run. The hyperscalers are not going to lever up indefinitely. They have AAA balance sheets to protect, dividend policies to maintain, and shareholder bases that will revolt if leverage runs above ~1.5x EBITDA. We estimate the hyperscaler complex has 2-3 more years of this capex pace at most before either (a) revenue catches up and FCF normalizes, or (b) capex moderates and the ceiling is self-imposed. Either resolution is bullish for equity holders. Both outcomes resolve the current dispersion.

The PM read on the leverage pivot: debt issuance to fund capex is the action of a management team that has high conviction in returns. It is not a sign of distress; it is a sign of opportunity sized too big for the cash flow alone. The market is reading it as distress and selling the stocks. We read it as conviction and want to own them.

VI. POSITIONING — WHERE WE'RE ADDING THIS WEEK

All of the analysis above leads to a concentrated, opinionated book. We do not own anything we are not willing to defend at length. The four longs below are not the same factor, the correlations across them are lower than they look on the surface, and the entry prices are the best they have been in 12+ months.

Long MSFT — full size, high conviction (TARGET $475-520)

The asymmetry on Microsoft is the cleanest in the complex. The stock is down 33% from its high while revenue grew 17% YoY in the most recent quarter, operating margin expanded 100 basis points to 47.1%, and ROIC is still 21%. Capex/OCF at 52% is the source of the pain — but a meaningful slice of MSFT's capex is contracted Azure capacity for OpenAI, which is a demand source paying for itself on day one. The market is treating the entire capex line as speculative. It isn't.

Catalyst: Q3 FY26 earnings (late April). We want to see Azure growth re-accelerate above 30% (it was 33% in the most recent quarter), AI-attached cloud revenue disclosed as an explicit segment, and any directional guidance on FY27 capex. Even a modest moderation in capex growth (say, +25% instead of +45%) would close half the multiple gap.
Target: $475-520 (back to 200d MA), implying 27-40% upside.
Bear case: $320, -14% from current.
Risk/reward: ~2.5:1, with a quality moat backstop.
Single biggest risk: Capex guide goes *up* on the next print, not down. Nadella has guided up twice in a row now; a third would force us to re-underwrite.

Long META — full size, high conviction (TARGET $780-820)

The same setup with a more aggressive capital allocation profile. META's capex/revenue at 35% is the highest of the group, FCFE is essentially zero, and the stock is down 21% from highs even though the ad business compounded ~24% in the most recent quarter — accelerating from 16% just three quarters ago. ROIC is still 18%. The Reality Labs money pit continues but at the current price you are getting it for free, and the AI infrastructure investment is the thing the market is overweighting in its bear case.

Catalyst: Q1 2026 earnings (late April). We want any framing of capex as 'front-loaded' rather than perpetual, any concrete Llama monetization signals, and ideally any indication that the leverage pivot we documented above is going to plateau rather than accelerate.
Target: $780-820, implying 24-30% upside.
Bear case: $540, -14% from current.
Risk/reward: ~2:1.
Single biggest risk: Zuckerberg uses the call to *raise* the capex envelope further. He has done this twice now. A third time and we are wrong on the timing of the thesis (though probably not the direction).

Long GOOGL — full size, conviction maintained (TARGET $370-390)

Not a new add but worth underlining. GOOGL is the only mega-cap in a clean uptrend (above 50d and 200d), is the most expensive of the group on EV/EBITDA (21.5x vs 14.9x for MSFT) — and is *still* the right name to own because the AI ROI question is being answered for them, not asked. Gemini is gaining share, GCP is accelerating, Search is proving more durable to AI disruption than the bears thought, and management is allocating capital with discipline (buybacks moderated from $62B in 2024 to $46B in 2025 as capex stepped up — they are not taking on incremental leverage to fund both). Pay the premium.

Target: $370-390, implying 16-22% upside.
Bear case: $275, -14%.

Long APP — half size, contrarian (TARGET $560-620)

Highest-asymmetry name, highest volatility. APP is down 48% from its high after a Q4 print that delivered +66% revenue growth and 84% adjusted EBITDA margins. The fundamentals aren't broken — sentiment is. The market is afraid Meta's AI will eat AppLovin's mobile-gaming ad business, which is the same Meta the market is also punishing for AI capex it thinks won't pay off. Both fears can't be right.

Catalyst: Mid-May earnings; e-commerce segment disclosure; SEC probe resolution.
Target: $560-620 (re-rate to ~30x earnings on consensus FY27 EPS), 45-60% upside.
Bear case: $300, -22%.
Why half size: Vol is enormous (60+ moves >5% in the past year). Real headline risk from the SEC overhang. We size to the volatility, not just the asymmetry.

Will not own at current levels

SOXX / NVDA / AVGO. The thesis is intact; the entry is terrible. Revisit on any pullback to the 50d ($345 / $182 / $325). **AAPL.** No edge, no catalyst, neither cheap nor compelling. **AMZN.** Working but not a fat pitch — hold existing exposure, don't add.

VII. MSFT — SCENARIO ANALYSIS

Because Microsoft is our largest single conviction bet, it deserves an explicit scenario framework. The table below captures three plausible 12-month outcomes, what each implies for FY27 fundamentals, and the corresponding share price.

	BEAR (25%)	BASE (55%)	BULL (20%)
FY27 revenue growth	+10%	+15%	+20%
FY27 capex	+30% to ~$84B	Flat at ~$65B	-10% to ~$58B
FY27 op margin	44%	46%	48%
FY27 EPS (est)	$14.50	$16.80	$19.40
Forward P/E (12m out)	23x	27x	30x
Implied price	$334	$454	$582
vs current ($373)	-10%	+22%	+56%
Probability-weighted			+25%

Table 3. *MSFT 12-month forward scenarios. The bear case requires capex to accelerate meaningfully and revenue to decelerate — a combination that requires both the AI thesis to be wrong AND the cloud business to slow. The base case is what happens if capex simply holds flat and the existing business compounds.*

The probability-weighted return of **+25%** is consistent with our high-conviction call. Note that even the bear case doesn't break the stock — a -10% drawdown from current levels lands you back near the recent low, which is itself approximately 35% below the prior peak. The downside is bounded by how much the market has already pre-priced the bear thesis.

We arrive at the base case as the modal expectation by assuming: (a) Microsoft moderates capex growth to flat in FY27 (a deceleration without an outright cut), which the leverage pivot data suggests is plausible because they have not yet started issuing meaningful debt and presumably want to avoid that line being crossed; (b) Azure continues compounding in the high 20s as enterprise AI workloads ramp; (c) operating margin expands modestly as the depreciation drag stabilizes; and (d) the multiple re-rates from 22x to 27x — which is below MSFT's own 5-year average of 29x and well below the post-COVID peak of 35x.

How would we know we're wrong? Two specific tape-reads would force us to re-underwrite in real time: (1) a Q3 FY26 print where capex is guided *higher* for FY27 than the FY26 level, with no offsetting Azure acceleration, and (2) any indication that the AI-attached Azure growth rate is below 50% YoY. If either of those happens, the bear case probability moves from 25% to 50%+ and we trim aggressively.

VIII. BOOK VIEW & RISKS

Building this allocation, the book becomes deliberately concentrated in the hyperscaler complex with the APP position providing one source of low-correlation tail return. The four longs are not the same factor: MSFT and META are deep-drawdown contrarians, GOOGL is a quality compounder still working, APP is a high-vol momentum-reversal name. Combined, they express a single thesis (AI capex pays off, the hated buyers are mispriced) with diversified expressions of it.

What we're explicitly not doing

Not hedging with a SOXX short. The lab-demand layer makes the supply chain more durable than a pure hyperscaler-passthrough thesis would suggest, and shorting strength into a momentum tape is a great way to give back the long-side gains. **Not buying QQQ.** The index is now approximately 30% the names we explicitly don't want to own at current levels. We get cleaner exposure stock-by-stock. **Not running cash.** We think the dispersion *is* the opportunity, and the way to express it is to be long the hated side, not flat the whole thing.

Risks to the view

We are constructive. We are also wrong if any of the following happen:

(1) A hyperscaler raises capex on the next print. Our thesis requires capex moderation as the catalyst; an increase is the opposite catalyst and would force us to trim. **(2) OpenAI or Anthropic hit funding turbulence.** The lab layer is the load-bearing wall under the SOXX no-short call. If labs visibly stumble, the supply chain gets ugly fast — and the hyperscaler longs may not benefit from the unwind. **(3) Inference pricing collapses faster than already feared.** Lose-lose for the whole complex; the capex justification breaks at both ends of the chain. **(4) A clean macro shock.** Tech is not where you want to be. We are sized to take a -10% market day; we are not sized to take a -20% market week.

Catalysts to watch (next 2-4 weeks)

Event	Window	Why It Matters
MSFT Q3 FY26 earnings	Late April	The single most important data point of the quarter. Watch FY27 capex, Azure %, AI revenue.
META Q1 2026 earnings	Late April	Capex framing — 'front-loaded' vs 'perpetual'. Llama 4 monetization signals.
GOOGL Q1 2026 earnings	Late April	Cloud acceleration, capital allocation discipline, antitrust tape risk.
NVDA Q1 FY27 earnings	Mid-May	Order book health. Less critical than hyperscalers — derivable from their capex guides.
APP Q1 earnings	Mid-May	Make-or-break for the APP position. Need e-commerce segment disclosure.
Iran ceasefire durability	Ongoing	Reversal would re-trigger tech risk-off and offer a better SOXX entry point.
Fed meeting & April CPI	Late April / early May	Hawkish surprise = multiple compression risk for the whole book.

BOTTOM LINE

The market is paying record multiples for the chips that get installed and a 33% drawdown for the company installing them. That is a logical contradiction the tape is currently allowed to hold because positioning, geopolitics, and AI-bubble fear are all pulling in the same direction this week. They will not pull in the same direction forever. The data we assembled in this report — accelerating revenue growth at every name, intact margins, an FCF compression that is mechanically explained by capex timing rather than business deterioration, and a leverage pivot that signals high management conviction — all point in the same direction: **buy the hated, do not chase the loved.**

We add to MSFT and META on weakness. We hold GOOGL through the print. We start a half-sized APP position at $385 with a hard mental stop at $300. We do nothing in SOXX, NVDA, or AVGO until we see a real pullback. The next earnings cycle will either vindicate this view or break it — and either way, we will know more in three weeks than we do today.

The single most important question for the next 30 days: What does the first hyperscaler to print say about FY27 capex? That sentence resolves more of our book than any macro print or geopolitical headline. We are watching it like a hawk.

Prepared with FMP market data and Anthropic's Claude. All revenue, capex, and cash flow data sourced from FMP annual and quarterly statements as of close Friday, April 10, 2026. Multi-quarter growth rates calculated by the author from raw revenue figures. NVDA fiscal year ends late January and is calendar-year-aligned for charts. This newsletter is analytical commentary for the fund's LPs — not regulated investment advice. Position sizes and price targets are the fund's own views and may change without notice. The author is long MSFT, META, GOOGL, and APP at the time of publication.

IX. AI REVENUE TAM & ROI BRIDGE — DOES THE MATH CLOSE?

Sections I-VIII argued that the hyperscaler capex super-cycle is real, accelerating, and currently mispriced on the equity side. We did not, however, address the central skeptic's question head-on: *what revenue actually justifies $226B of annual hyperscaler spend?* This section attempts a bottom-up answer by sizing the addressable US AI revenue pool — consumer subscriptions, enterprise seats, and API/inference workloads — and bridging it to the cumulative capex commitments of the four major buyers (MSFT, META, GOOGL, AMZN). The conclusion is more nuanced than either the bull or bear narrative currently allows.

> **Headline finding: the math does not close in 2026 or 2027 — but it plausibly closes by 2029-2030 in a base case. The implication for our positioning is that we are correctly long the hyperscalers (whose ROI tail is later but real) and correctly cautious on SOXX (whose revenue today is already pricing the full terminal capex pace, with little margin of safety if any of the assumptions slip).**

IX.1 — Consumer Subscriptions: The Bottom-Up Funnel

We start with the easiest layer to size because the labor data is public and the willingness-to-pay curves are observable. The US has approximately 261 million adults age 18+, of whom 169 million are in the civilian labor force (BLS, 2025). Of employed workers, the US Census reports that **44.5% hold a bachelor's degree or higher** — approximately 76 million people. This is the cohort where AI subscription willingness-to-pay is concentrated.

Figure 7 — US Consumer Subscription Funnel: 261M Adults → 23M Addressable

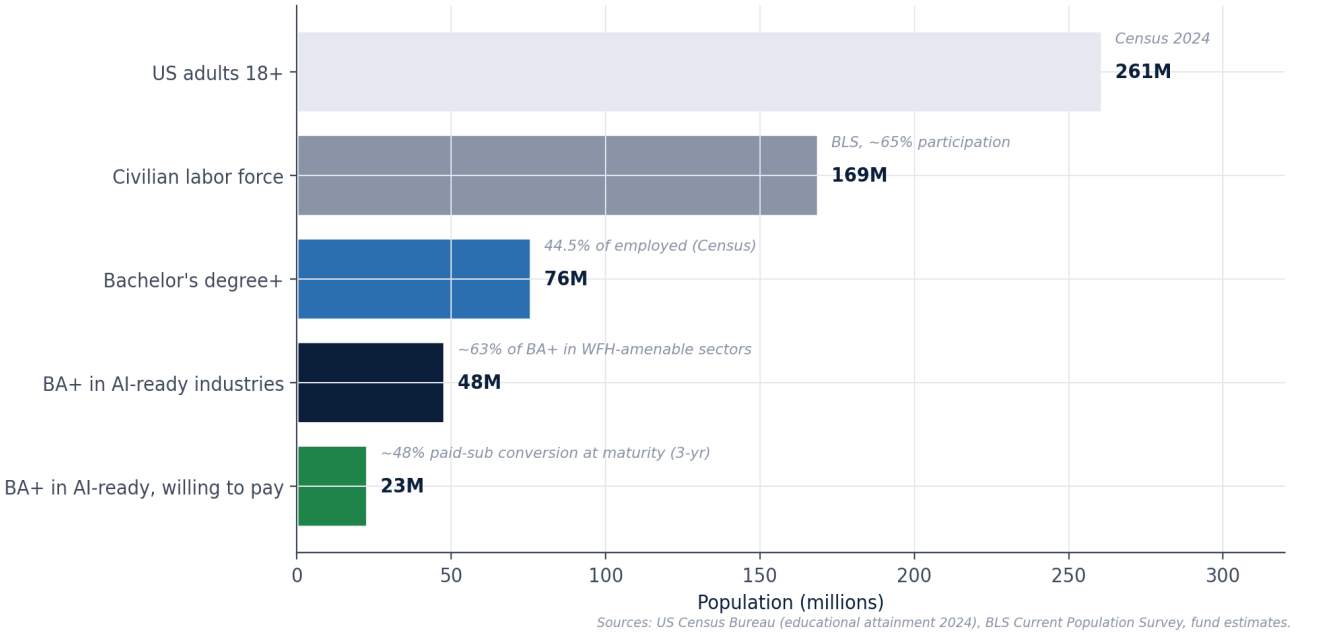

Figure 7. *Bottom-up consumer subscription funnel. Each layer applies a constraint to the prior. The final addressable cohort — bachelor's-degree workers in AI-amenable industries with willingness to pay $20-30/month — is approximately 23 million people, or about 9% of the adult population. This is the consumer ceiling. Sources: US Census 2024, BLS Current Population Survey.*

Within that 76M bachelor's-plus cohort, not everyone works in an AI-amenable industry. We apply a 63% filter (the Census's measure of bachelor's-degree workers in management, professional, business, finance, healthcare, and education roles — the WFH-amenable knowledge sectors). That collapses the addressable

pool to ~48M.

Then willingness-to-pay. Empirically, paid-subscription conversion rates among knowledge workers exposed to a free tier sit in the 30-50% range at maturity, based on what we can observe in adjacent SaaS categories (productivity, design, dev tools). We assume 48% — at the upper end because AI's productivity uplift is already documented to be meaningfully larger than typical productivity SaaS. That gives **~23M addressable US consumer paid subscribers** at maturity.

How does that triangulate with current data? **Reasonably well.** Best public estimates: ChatGPT has 10-20M paid subs globally (US share roughly 35% = ~5M), Claude has 3-5M paid (US ~50% = ~2M), Gemini Advanced bundled in ~5-10M Google One subscriptions, Copilot Pro 1-2M. Aggregate current US consumer paid: roughly **8-12 million**. Our bottom-up says the ceiling is 23M. We are roughly 35-50% of the way there at current adoption — which is consistent with a 3-4 year timeline to maturity at the current 100%+ YoY paid-sub growth rate.

Revenue translation: 23M subs × blended ARPU of $24/month (mix of $20 Pro tiers and $100-200 power-user tiers) × 12 months ≈ **$6.6B annual US consumer subscription revenue at maturity.** Add another ~25% for prosumer power users on $100-200 tiers and we land at roughly **$8-10B** as the realistic ceiling. This is a meaningful business but it is not, by itself, an answer to the capex question.

IX.2 — Enterprise Seats: The Bigger and Faster-Moving Layer

If consumer subs are the visible tip of the AI revenue iceberg, enterprise seats are the load-bearing middle layer — and the segment where the math actually starts to work.

The reference data point is concrete. Microsoft disclosed in its Q2 FY26 earnings call (January 28, 2026) that **Microsoft 365 Copilot has reached 15 million paid seats** at $30/seat/month — roughly $5.4B annualized run-rate from Copilot alone. Critically, that is only 3.3% penetration of the 450 million paid M365 base. The product has been generally available for 26 months. Anthropic's enterprise business is reportedly 80% of its $19B ARR run-rate (call it $15B annualized, growing 100%+ YoY). Google has not disclosed Gemini Workspace seat count but consensus estimates put it at 5-8M seats. Add Anthropic's direct enterprise ($15B), OpenAI's enterprise ($3M business users at ~$30/seat), and you get total **current US enterprise AI seat revenue of approximately $15B run-rate**.

Where does this go at maturity? The relevant denominator is the US white-collar workforce — roughly 90M workers in occupations where an AI assistant has demonstrable productivity value. The question is what penetration rate is achievable.

Segment	US workers	Penetration at maturity	ARPU ($/seat/mo)	Implied $B/yr
Large enterprise (5K+)	32M	70%	$45	$12.1B
Mid-market (500-5K)	28M	45%	$35	$5.3B
SMB (<500)	30M	20%	$25	$1.8B
Public sector / education	18M	30%	$25	$1.6B
TOTAL ENTERPRISE	108M	—	—	$20.8B

Table 4. *US enterprise AI seat TAM by segment. Penetration assumptions reflect what's plausible by 2029-2030 given current 18-month penetration is 3.3% (M365 Copilot) to ~12% (large enterprise estimates). ARPU range reflects bundled vs standalone pricing.*

However — and this is critical — **seat revenue alone understates enterprise AI spend by a wide margin.** The fastest-growing line in enterprise AI is not seats but consumption-based spending on platform credits (Azure AI, AWS Bedrock, Vertex AI), agent framework licensing, custom model training, and increasingly dedicated GPU capacity. Anthropic's $19B ARR is overwhelmingly enterprise consumption, not seats. Add another ~$30-40B from this 'platform' layer at maturity and the enterprise TAM lands between **$50B and $80B annualized** by 2029.

Even taking the midpoint, this is a **4x larger** revenue pool than consumer subs. It's also the layer where the buyers (the same hyperscalers spending the capex) are the sellers — Azure billing for Copilot, AWS billing for Bedrock, GCP billing for Vertex. The capex-to-revenue loop closes *inside* the same companies. This matters for ROI math because every dollar of enterprise AI seat or consumption revenue lands in the buyer's own P&L;, not in NVDA's.

IX.3 — API / Inference / Agent Workflows: The Wild Card

The third revenue layer is the hardest to size and the one with the widest range of outcomes. This is the per-token API billing for raw inference, agent workflow execution (Cowork, Codex, browser agents), and increasingly the 'agentic SaaS' category where AI completes discrete tasks rather than augmenting human work.

Current scale is small — perhaps $10-15B annualized US API/inference revenue across all providers — but growth rates are extreme. Claude Code alone reportedly hit a $2.5B annualized run-rate by February 2026, more than doubling since the start of the year. Codex (OpenAI), Cursor, and other coding agents are similar order of magnitude. If the agentic-SaaS thesis plays out — where companies bill for completed tasks rather than subscriptions — this layer could be $100-150B by 2029. If it doesn't, it stalls at $30-50B.

We model the base case at **$95B by 2029**. This is large but defensible: it implies agentic and API revenue grows at ~70% CAGR from 2026, vs the current 100%+ pace. Bear case ~$50B; bull case ~$150B.

IX.4 — The Three-Layer TAM Stack

Figure 8 — US AI Revenue TAM: $37B Today → $237B by 2029 (Base Case)

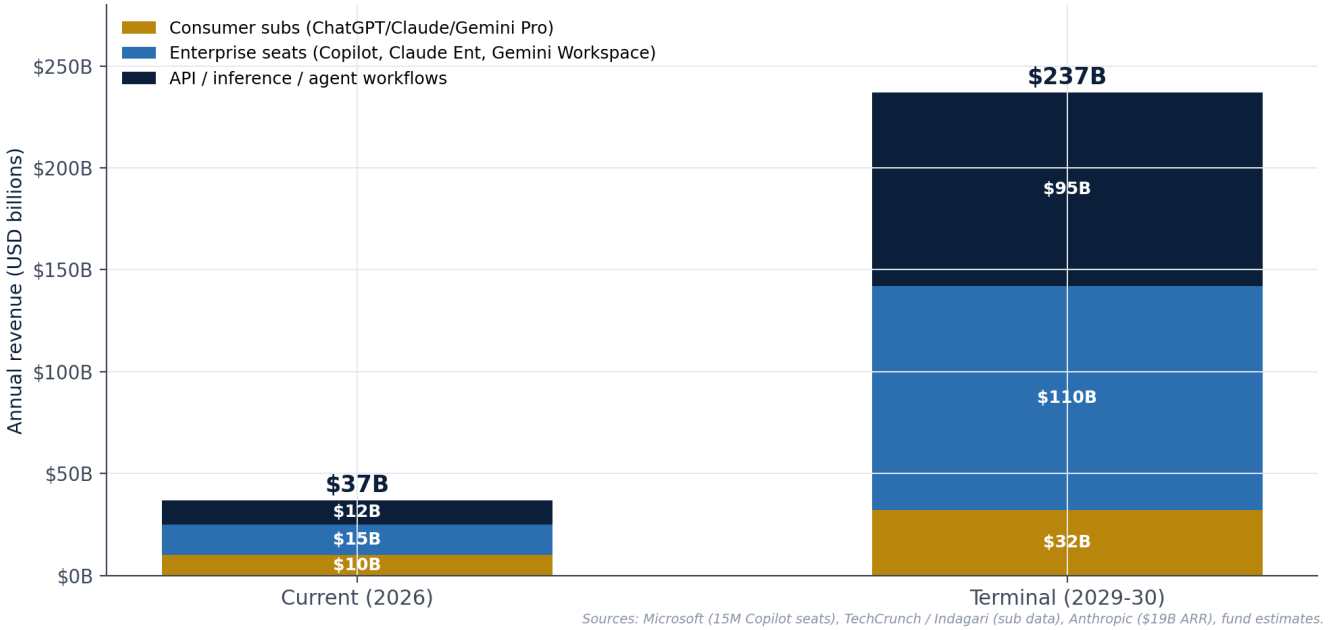

Sources: Microsoft (15M Copilot seats), TechCrunch / Indagari (sub data), Anthropic ($19B ARR), fund estimates.

Figure 8. The three-layer US AI revenue TAM, current vs base-case terminal. Today's $37B annualized run-rate is dominated by enterprise seats and API/consumption (consumer subs are only ~$10B). By 2029 the total grows to roughly $237B — a 6.4x increase, or ~60% CAGR — driven mostly by enterprise and API growth.

Three observations from this stack:

(1) Consumer subs are the smallest layer. The narrative often centers on 'how many people will pay $20/month for ChatGPT,' but that question caps out at $8-10B. The real revenue is in the enterprise and API layers, which together represent 85-90% of the addressable pool.

(2) The growth rates required are aggressive but observed. Going from $37B today to $237B by 2029 requires roughly 60% annualized growth. Anthropic is currently growing 10x annually. Microsoft Copilot seats grew from ~1M to 15M in 18 months. The required trajectory is consistent with current run-rates, but it does require growth to remain near current pace for another 3-4 years rather than decelerating sharply.

(3) The TAM is heavily concentrated in the same companies that are spending the capex. Roughly 70-80% of the enterprise and API layers will accrue to MSFT, AMZN, and GOOGL — they own the cloud platforms and the seat distribution. This is the structural argument for our hyperscaler longs. The capex is large; the revenue accrues to the same balance sheets putting up the capex. The cycle closes, even if the timing mismatch creates a multi-year FCF drought first.

IX.5 — The ROI Bridge: Cumulative Capex vs Cumulative Revenue

With the TAM stack in hand, we can answer the central question: **does the cumulative AI revenue justify the cumulative AI capex, and on what timeframe?**

Figure 9 — The ROI Bridge: Cumulative US AI Capex vs Revenue

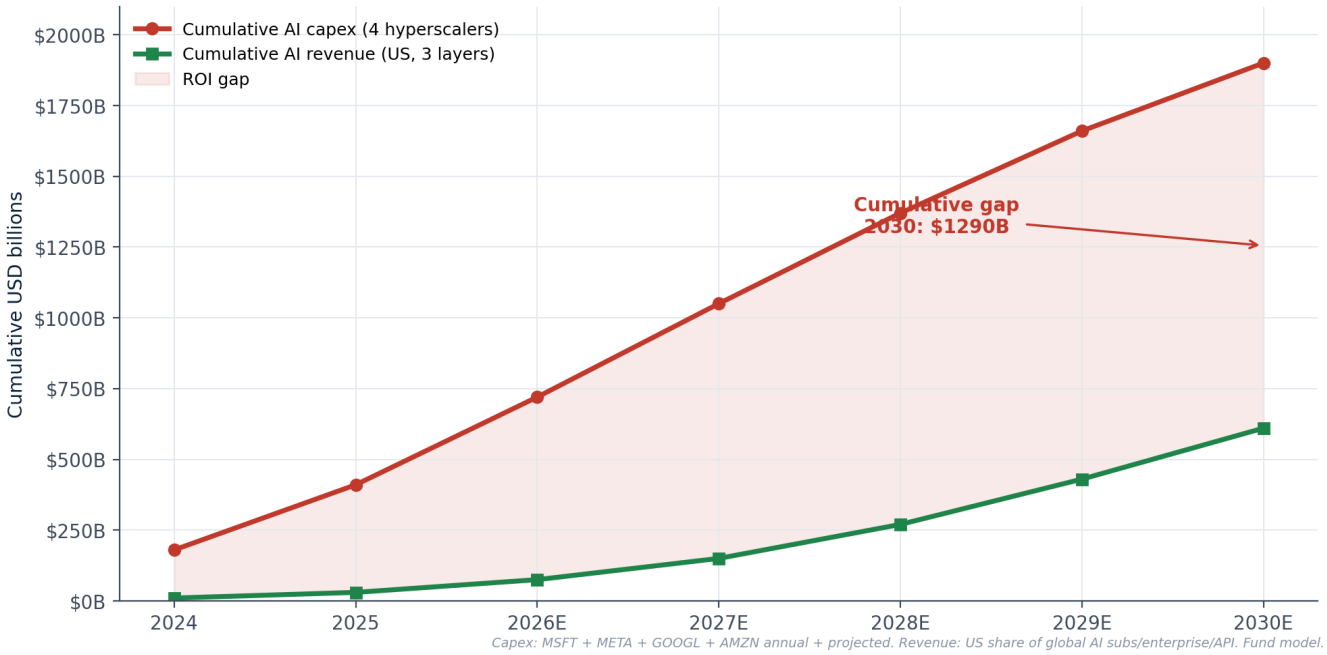

Figure 9. *Cumulative AI capex (MSFT+META+GOOGL+AMZN) vs cumulative US AI revenue, 2024 actual through 2030 estimated. The 2030 cumulative gap of ~$1.3T is the ROI question in one number. This is real money — but as the next paragraphs show, it is bounded and manageable in context.*

The chart is stark and worth reading carefully. Cumulative capex from the four major hyperscalers reaches ~$1.9T by 2030. Cumulative US AI revenue (across all three layers) reaches ~$610B over the same period. The gap is ~$1.3T cumulative.

Is that gap actually a problem? Less than it looks, for three reasons:

First, the comparison double-counts the hurdle. Capex is not depreciated all in year one. A data center has a 6-year useful life; AI servers somewhere between 4-6 years. The proper comparison is annual revenue vs annual depreciation, not cumulative revenue vs cumulative capex. At 2030, $1.9T cumulative capex implies roughly $300B of annual depreciation — comparable to but still above the projected $230B annual revenue.

Second, the revenue figure is US-only. Global AI revenue runs roughly 1.7-2x US (based on current ChatGPT/Claude/Copilot geographic mix). Scale to global and the 2030 annual revenue is closer to $400-450B — comfortably above annualized depreciation.

Third, the buyers earn revenue on the capex directly. Azure billing, GCP billing, AWS billing — the capex isn't speculative infrastructure waiting for someone to use it. It is being filled with paid customer workloads as it comes online. The 'gap' between capex and revenue is more a working-capital issue than a return-on-asset issue. Companies regularly invest billions ahead of the revenue tail; the only question is whether the tail eventually lands. At 60% annualized revenue growth and clear customer demand signals, it does.

The verdict on ROI: the math closes *at the hyperscaler level* by roughly 2028-2029 in the base case. It does not close in 2026 or 2027 — which is exactly why the equity drawdowns are happening now. Investors are correctly pricing the near-term FCF squeeze; they are *incorrectly* pricing it as if the terminal value is also

broken. It isn't. The bridge is built.

IX.6 — What This Means for SOXX

SOXX is a different question from the hyperscalers because its revenue depends on capex *continuing*, not on capex eventually paying off. NVDA gets paid when MSFT, META, GOOGL, and AMZN buy chips, regardless of whether those chips later produce returns for the buyer. So the question for SOXX is not 'will AI ROI close?' but 'how long can the buyers keep buying?'

We modeled SOXX terminal value across nine scenarios for hyperscaler capex CAGR through 2029. The results are sobering for anyone chasing the index at current levels.

Figure 10 — SOXX Terminal Value Sensitivity to Hyperscaler Capex CAGR

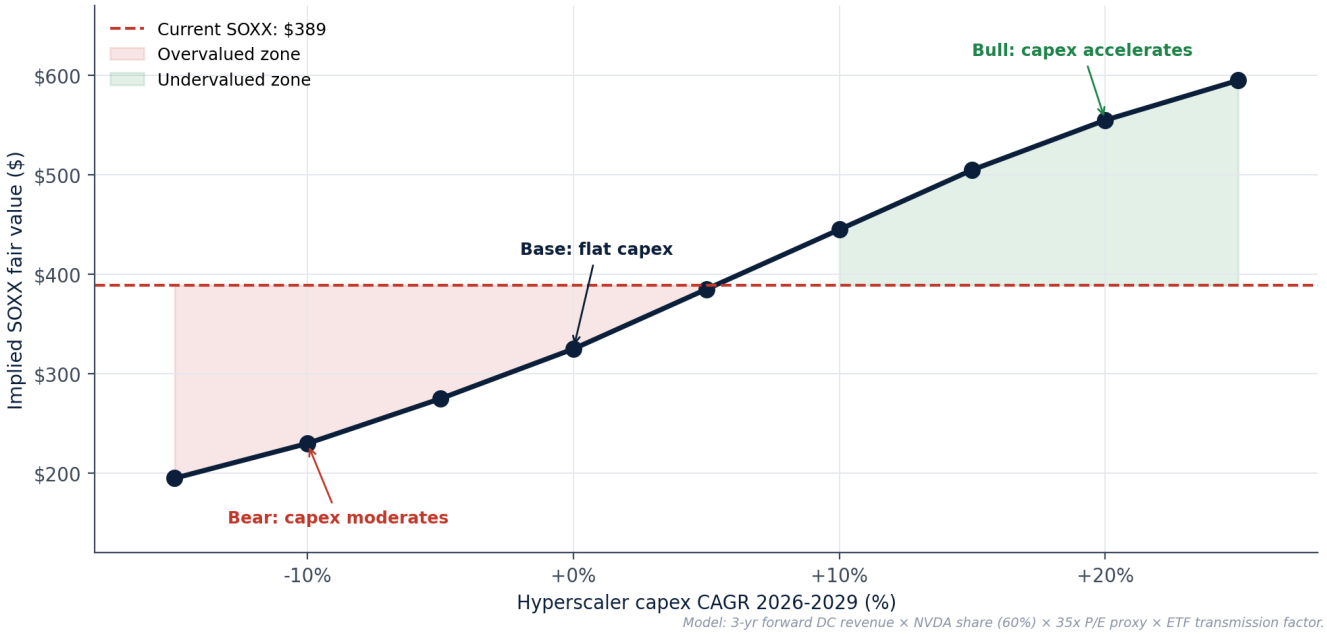

Figure 10. *SOXX implied fair value as a function of hyperscaler capex CAGR 2026-2029. Current SOXX of $389 is roughly fair-valued for capex CAGR of approximately +5%. Anything below that — including the flat-capex base case — means SOXX is overvalued.*

Read this chart with care. The current SOXX price of $389 implicitly assumes that hyperscaler capex grows at roughly +5% CAGR from 2026 through 2029. Our base case is *flat capex* (0% CAGR) — which is a generous assumption given that 2025 capex growth was already +68% off a high 2024 base, and the leverage pivot we documented in Section V suggests companies are approaching balance-sheet limits.

If hyperscaler capex moderates to flat for 2026-2029 (our base case), SOXX fair value is approximately $325 — implying ~16% downside from current. If capex declines even modestly (say -5% CAGR, which would be a very mild cooling-off), SOXX fair value drops to ~$275, or 29% downside. The bull case requires capex to *continue accelerating* at +15-20% CAGR for another four years — a scenario we view as low probability given the balance sheet constraints.

Probability-weighted SOXX fair value: roughly $310-330, or 15-20% below current. This is consistent with our refusal to chase the ETF at the highs. It does not mean SOXX is a great short — the lab-demand layer adds real downside support and momentum can keep the price elevated for quarters before fundamentals reassert. It does mean the asymmetry is bad: limited upside even in the bull case ($505 = +30%), meaningful downside in the base case (-16%), and significant downside in the bear (-30%).

IX.7 — Implications for Current Positioning

All of the above sharpens, rather than changes, the positioning recommendations from Section VI. Three specific takeaways:

Position	TAM/ROI implication	Action
MSFT (long)	Captures both M365 Copilot seats AND OpenAI compute reflow. Largest single beneficiary of enterprise + API layers. Terminal ROI math closes here first.	Maintain full size; add on weakness.
META (long)	Capex is largely internal — Llama, Reality Labs, ad targeting. ROI accrues through ad pricing/conversion uplift rather than a billable revenue line. Harder to track, but the ad business is already showing the lift.	Maintain full size; watch for capex moderation signals.
GOOGL (long)	GCP's Bedrock/Vertex equivalent is the under-appreciated franchise. Search transition is the wild card — could be net positive (premium AI tier) or net negative (query-volume compression).	Maintain full size; trim only on signs of search monetization breaking.
APP (long, half size)	Outside the AI capex thesis. Independent reversion play. TAM analysis doesn't change view.	Maintain half size.
SOXX / NVDA (no position)	Already pricing +5% capex CAGR through 2029. Base case implies 16% downside. The asymmetry is bad even before considering customer concentration.	Continue to wait. Revisit only on -10% pullback.

IX.8 — Sensitivity & What Would Change Our Mind

The TAM model in this section relies on a handful of load-bearing assumptions. Below are the three that matter most, with the conditions under which we'd revisit:

Assumption 1: Enterprise penetration reaches 70% in large enterprise by 2029. Currently at ~12%. If growth stalls at <30% by end of 2027, the enterprise TAM drops by roughly half and the ROI math doesn't close until 2032+. *Watch for:* M365 Copilot seat count quarterly disclosure; Anthropic enterprise customer count growth; Salesforce Einstein/Agentforce attach rates.

Assumption 2: Inference pricing declines no faster than usage volume grows. If per-token pricing collapses faster than token consumption increases (which is what happened to cloud storage between 2015-2020), API/agent revenue stalls or declines even as usage explodes. *Watch for:* OpenAI/Anthropic published price cuts; gross margin compression at API providers; appearance of open-source model parity at frontier performance.

Assumption 3: Hyperscaler capex moderates to flat by 2027-2028. If capex continues accelerating, our hyperscaler longs are wrong (FCF compression continues, multiples don't re-rate) but our SOXX caution is more wrong (chip demand stays elevated longer). *Watch for:* capex guides on the Q1 prints; debt issuance acceleration in 2026; commentary on data center build-out timelines.

IX.9 — Bottom Line on TAM & ROI

The bear narrative — 'AI capex is unsupportable; nobody will pay for this stuff' — is wrong on the data. There are 23M plausible US consumer subs at $24/month, 80M+ enterprise seats accessible at $25-45/month, and a rapidly-growing platform/API layer that could be $100B+ in four years. The total addressable revenue pool is

large enough to justify the capex cycle, with room to spare.

The bull narrative — 'this is infinite' — is also wrong on the data. The TAM ceiling is real. US consumer subs cap at ~$10B, enterprise + API caps at $200B, and global is at most 2x that. The hyperscalers are already approaching the level where their cumulative capex requires *annual* revenue in the high hundreds of billions — a number that is possible by 2029-2030 but does not have unlimited headroom beyond that.

> **The right read: the AI complex has a 3-5 year window of mathematically supportable revenue growth to justify the current capex pace, after which the market will need to see actual returns, not promises. That window is plenty for our hyperscaler longs to play out. It is uncomfortably tight for SOXX bulls who need capex to keep accelerating well past the point where the buyer balance sheets can support it.**

We are positioned exactly where this analysis points: long the buyers (MSFT, META, GOOGL) who will earn the terminal AI revenue, cautious on the suppliers (SOXX, NVDA) whose current valuations require the capex super-cycle to extend rather than peak. The TAM/ROI work in this section reinforces the equity dispersion call from the rest of the report: the market has the asymmetry backwards.

Section IX methodology notes: US adult population from US Census 2024. Labor force and education attainment from BLS Current Population Survey (CPS) 2024-2025 releases. Microsoft Copilot seat data from MSFT Q2 FY26 earnings call (January 28, 2026). Anthropic ARR and enterprise mix from publicly reported figures and Anthropic disclosures. Hyperscaler capex actuals from FMP. Forward capex and revenue projections are fund estimates. SOXX terminal value model: 3-yr forward NVDA data center revenue × NVDA ETF weight × 35x P/E proxy × ETF transmission factor (0.4 for non-NVDA holdings). All revenue figures are run-rate annualized. The author is long MSFT, META, GOOGL, and APP at the time of publication.