

“AI” Already Means Nothing

COMMENTARY
November 2025

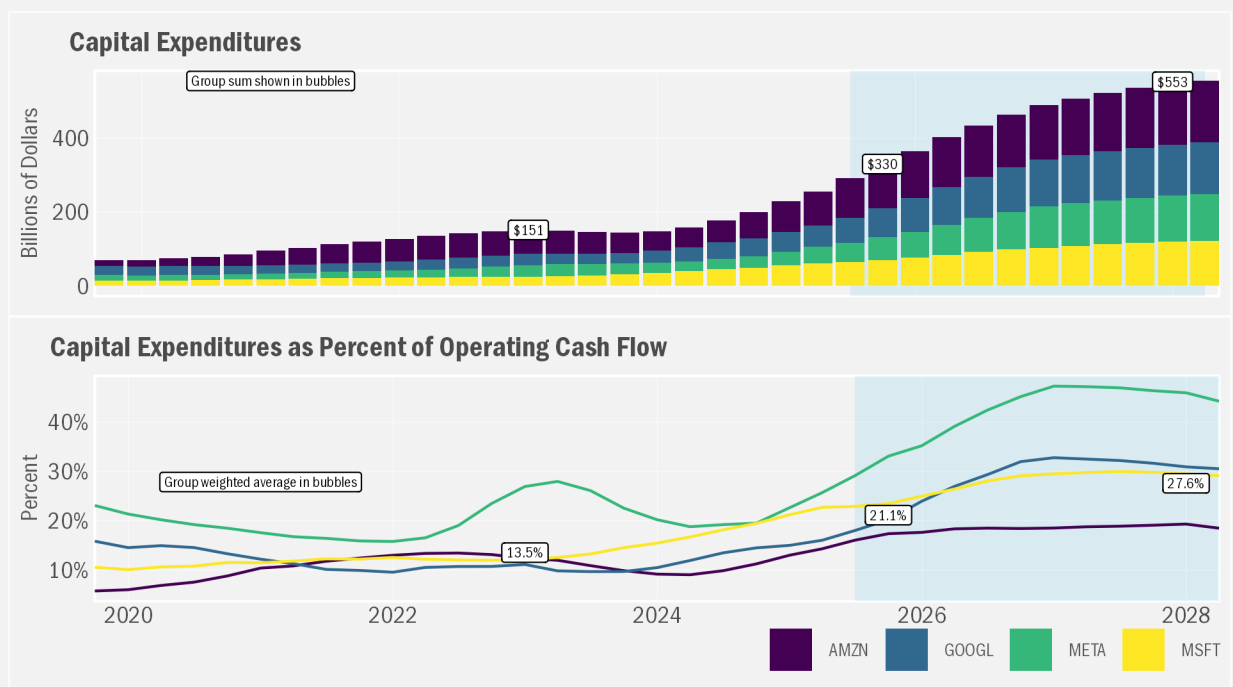
Seems not since the Tech Bubble of 2000-01 have we experienced such a rapid shift from “meaningful advance in technology” to rampant hype bordering on hysteria. Now ubiquitous, “AI” dominates the empty boilerplate phraseology of executives and pundits—however hollowly—seeking to convey techno-sophistication and endless promise. Finding the memetic expansion unprecedented and becoming ever more lacking in reasonable basis, we wouldn’t care so much were generative artificial intelligence (genAI, or just AI) not such a powerful market and macroeconomic force. But we find such boosterism is...slowly...losing ground to reality:

- Growing actual costs for building and operating AI infrastructure vastly outpace downstream AI revenue
- Parabolic ramps in expected revenue ignore the tech’s inherent flaws and are becoming more heavily dependent on presumed capabilities unlikely to materialize in the medium term from present technologies
- Enthusiasm for AI’s yet-unrealized potential arguably has boosted portions of the market well beyond reasonable valuations
- Concerned about excess investor enthusiasm for the AI theme, we think an immediate reckoning is not near

Search for Meaning

We’re kind of shocked how quickly “AI” has become a meaningless, throwaway concept, one particularly fond of which seem to be public company executives, investment strategist sorts and your average Reddit poster. Everyone’s gonna “do AI” and make loads of dough! Given time-in-the-market and technological maturity (much debate happens here, as we will discuss in as bit), though, we wonder why we’ve so little detail regarding actual services-related revenue being derived from AI technologies (excluding hardware sales and those from the folks renting out servers to do the required math). The relevance of course is that there’s no sense for the hardware to exist without those end-use dollars. For the half-trillion or more invested so far in infrastructure and considering plans for trillions more through decade end, we should have better insight into realistic top-line growth.

Figure 1: “Hyperscaler” Capital Expenditures



Trailing four-quarter data. Reported data as of 09.30.25. Estimated data through 03.31.28. SOURCE: SRCM using data from Bloomberg

All About Tomorrow

The typical response is that we should just wait and see. Thing is, we have been waiting and seeing. And for as many times as we've been impressed by AI tech—from the near-magical AI-generated “podcast” that discussed a commentary I wrote to the far more practical application of AI tech for searches across internal documentation—we have yet to be impressed with actual sales data.

“But wait!”, you might say, “OpenAI has secured more than a trillion dollars in deals!” That can't be for nothing! True. But those deals—increasingly circular (seller invests in buyer of goods/services from seller) as so many of them and others like them in the AI space are—relate mostly to securing **the capacity to do AI math**. They mostly do not relate to any sources of **reasons to do the AI math**. But the latter is all we really should care about.

And here we're certainly not the only ones with growing concerns that the tendency to provide dead-end results is an unavoidable defect of large language model-based technologies. It's now well understood (among those who care to pay attention) that AI technologies: 1) do not “understand” the world in any manner a human might conceptualize that word, 2) do not contextualize the world in any other sense than by consolidating descriptions thereof given some manner of relevant historical example in text, picture, illustration or video, 3) provide responses to prompts that one should never presume have any factual basis beyond a potentially irrelevant contextual basis given the underlying computational method.

The phrase “stochastic parrot” is overused and probably, at this point, too trite when considering more recent advances in overlay technologies used to accommodate underlying flaws, but it's nonetheless indicative of the core problem. I asked CoPilot to explain why, which I summarize as: stochastic means “based on probability” and refers to the fact that **LLMs use statistics to predict responses**. They therefore **don't in any manner have any “understanding” of their responses**—much as a parrot likely doesn't when mimicking human speech—which as a result may have no factual basis in reality.

Note the emphasis on “may”. Perhaps responses to prompts are most of the time completely accurate and on point. But we think it matters greatly that the expected hit rate is south of 100. And that's because even in non-mission-critical situations, it matters, given the seemingly growing tendency to offload work to an LLM.

For example, I asked Microsoft's CoPilot app to review this document for, “clarity, comprehensiveness and accuracy.” After otherwise generally praising the text, excepting sarcasm, editorialization and structure, it noted that, “The document occasionally repeats entire paragraphs, which may confuse readers or suggest editing oversights.” When asked to what it was referring, it highlighted the introduction above, stating that it both came at the beginning of the document and, “again in the early section before the ‘Tight Ropes’”. Hrm. After I noted that the paragraph is not repeated, just exists in a text box at the beginning of the document, it thanked me for the clarification and agreed.

Such clinical text reviews represent in our view the lowest bar to clear for AI tech. And, yet, in our experience it often comically fails to clear that hurdle. From queries on docs we're creating to AI results from Internet searches, we regularly find very obvious mistakes in the results. Cure cancer? AI tech can't even reliably cure Word files.

Seems the answer, at present at least, is to throw more money around. We would presume the zillion-dollar pay packages being reported for AI scientists/engineers/etc. are meant to solve the parrot problem. Not sure what to think about the monies being allocated to infrastructure under the presumption either that the flaw will be solved or that users won't care to notice. Shown in Figure 1, the “hyperscalers” have spent an estimated \$330 billion on capex over the 12 months ended 09.30.25 (representing ~21% of revenue and ~52% of operating cash flow) a figure expected to exceed \$550 billion by the end of 2027 (~27% of revenue and ~67% of CF). Much of that spend—Perhaps half? Depends on the source of the anecdote and the location of the AI complex...but a large portion—represents revenue for Nvidia (NVDA). The chipmaker has seen its trailing revenue jump nearly 7-fold to an estimated \$185 billion for the

four fiscal quarters ended 10.31.2025 since 04.30.23 after revenue trailed off post crypto-mining craze, over which time its market cap soared ~14x to **more than \$5 trillion**, now accounting for more than 7% of U.S. market cap.

Cost/Profit Chasm

Importantly, Nvidia revenue is mostly not from AI. Rather, the chipmaker generates revenue mostly from the means to “do” AI. Genuine genAI-driven revenue from user-focused products and services remains broadly under-reported, with guesstimates based on anecdotes still coming in below \$100 billion on an annualized basis [at least by means of our searches the data covering those services/products that are enabled solely via the deployment of AI technology and excluding AI overlays of existing services/products and excluding revenue generated for providing the means (servers, etc.) for AI training and inference]. To be fair, though, it’s tough to gauge AI-specific revenue and revenue that’s enhanced through the addition of AI technology. For example, how does one consider Microsoft’s addition of CoPilot AI technology to its Office 365 suite? The company upped the annual fee for Microsoft 365 Personal and Microsoft 365 Family users by \$3 per month to \$99.99 and \$129.99, respectively, but does not otherwise charge them specifically for “basic” CoPilot, use of which is limited. A Microsoft 365 Premium license for \$199.99 ups CoPilot limits and users can add an unlimited version of CoPilot for \$20 per month. Corporate users can add Microsoft 365 CoPilot for \$30 per month, but an existing Microsoft 365 plan is required. Importantly, Microsoft is not yet breaking out revenue for CoPilot.

This quarter’s earnings season so far has brought little in the way of additional detail, so we have the sense that the mismatch between capex and revenue to pay back that investment is unlikely to improve materially over the short term. And that’s before we consider the expected ongoing capex ramp (see Figure 1), which more and more seems to us a grand exercise in blind faith. The reference to past bubbles is easy here. We still use some train tracks laid down more than a century ago. But without trains full of sold or sellable cargo, train tracks are useless. And we’ve still yet to fill the optical fiber strung a quarter-century ago. Without the paid-for data traffic to fill that fiber and busy the connectors and switches that move them all around, the whole of the optical space infrastructure otherwise would be useless.

Worse in the current context, those analogies fail to reflect the relatively delicate nature of AI infrastructure. Unlike rails and fiber optic cable, the chips used for AI “training” (building the models) and “inference” (using the models) don’t last anywhere near as long. Estimates suggest that the latest chips used in training will be useful at their current tasks for little more than three years. So, at least the semiconductor portion of the AI infrastructure cycle must be assumed to be set on a repeat cycle. Not only that, but each increment in AI capability so far has proved vastly more expensive from the required-infrastructure and time-to-completion standpoints, further exacerbating the mismatch between development/operational costs and revenue.

Meta (META) CEO Mark Zuckerberg remains unfazed by such concerns, having commented earlier in the year that his company wants to make sure it’s not underinvesting. Hard to see that happening at present. Capex for 2025 is now expected to range between \$70 billion and \$72 billion. On next year’s budgets, Meta CFO Susan Li stated on the call, “our current expectation is that capex dollar growth will be **notably larger in 2026 than 2025** [our emphasis]. We also anticipate total expenses will grow at a significantly faster percentage rate in 2026 than 2025, with growth primarily driven by infrastructure costs, including incremental cloud expenses and depreciation.”

META shares have sunk 13.7% since the call, perhaps a hint that investors may be turning a more skeptical eye toward “more equals more” spending on AI. Part of that may be a growing realization that AI technology may be of a radically different sort when it comes to profitability dynamics. While investors have grown to expect that costs for the provision of Internet-based services tend to fall, often dramatically, as use scales, such may not be the case with AI tools. As one means to make them more accurate and/or useful, inference providers have developed “reasoning” models, which (crudely speaking) iterate over a series of answers to steer the eventual result toward factual and contextual precision.

But those iterations require more compute, effectively reducing the potential margin on the request (leaving the potential revenue question aside for the moment). Same for pictures and video. For each additional complication, compute time rises, often dramatically. While we have learned from company reports that such costs are coming down via efficiencies gained from dynamic model choice and otherwise, comparison to, for example, the provision of classic search results is not a favorable one for the delivery of responses generated from inference. That is, a simple “search” request given to a classic search engine is handled in a radically different fashion when that request is entered into an AI chat, with costs (again, depending on an evolving mix of requested output and model choice) still strikingly higher for the latter.

It remains an open question as to what long-term-profitable AI product/service models will look like. And here is among the places where, for us, the “AI means nothing” truly presents. AI “agents” were meant to have been the next big thing already this year, paving the path toward expansive profits. The idea is that an agent could operate in some manner of autonomous fashion to produce desired outcomes in various situations. Often cited is an agent that could book an entire vacation for a family, not that anyone in their right mind would allow such a thing to occur. We tend to think that such agents must ultimately be more rules-based than AI and therefore are likely to prove little more than extensions of long-available algorithms. An exploration of analyst reports suggests many companies are reporting reasonably large revenue streams from “agentic” platforms. But it seems even this term is being diluted to include AI chat overlays onto existing customer service-oriented platforms (zzzzz) and, well, folks just calling it “agentic” because they know they’re supposed to.

Too Critical Too Soon?

To be clear though, while we believe we might maintain a better understanding of the inherent limitations of large language model-based artificial intelligence platforms than the average AI booster, as we said earlier, we continue to be amazed (if not as much anymore) by the inherent magic in the tech. Haven’t mentioned AI in the audio-visual space yet. Here, too, the journey from prompt to output can seem close to mystical.

And while you will never see such output on these pages or in our podcasts unless it is obviously identified as such, we admittedly are utilizing AI tools more often for querying our own content and external resources. Outside of otherwise normal (ehem...acknowledging that we remain in a relatively high inflation regime) year-on-year cost increases, though, **we are not paying more for the AI add-ons**. That “costs up, revenue flat” scenario we think ultimately will prevail across a wide swath of the AI space, likely leaving broad provision of such services only to the largest, most profitable of the players (who also will gobble up any minnows that manage to break the surface). But we don’t think that the provision of AI tools will radically alter tech-space profitability except for, perhaps, to the downside. Important in that consideration is our present understanding that there is little demonstrably and defensibly unique about any of the large language models or any individual company’s ability to utilize the output of LLMs for some manner of product/service provision. We suppose patents could end up allowing service providers to develop moats, but there’s little hint yet of efforts to lay sole claim to the AI tool space. So, we expect rising costs and competitive pressures to limit profitability for the near term across the AI space.

Of Course They Will

Those pressures may lead to more creative application of AI. Which brings us, of course, to the now seemingly inevitable dip-into-the-dumpster for any once fresh Internet-tangent tech: porn and ads. Just a few weeks ago, OpenAI CEO Sam Altman indirectly highlighted the intention to open its ChatGPT service to more adult-oriented fare, though he remains cagey on the topic of in-chat advertising. Meta’s Zuckerberg, on the other hand, provided a glimpse of a glorious AI future on the earnings call: “I think that these models will also improve modernization and all of the different ways that we’ve talked about so far in terms of improving engagement, improving advertising, helping advertisers

engage. I mean, there's the one opportunity that we just, we usually talk about on these calls, but hasn't, hasn't come up as much here is just the ability to make it so that advertisers are increasingly just going to be able to give us a business objective and give us a credit card or bank account and like have the AI system basically figure out everything else that's necessary, including video or different types of creative that might resonate with different people that are personalized in different ways, finding who the right customers are all of the capabilities that we're building, I think go towards improving all of these different things. So, I'm quite optimistic about that."

Going to editorialize to the extreme here, but we can't say we share that optimism. Indeed, much of the early hype around agents relates to helping customers buy things (read: steer them toward the highest ad bidder), an impact that remains well removed, to put it mildly, from using AI to cure cancer and nowhere close to the far dreamier "artificial general intelligence" or AGI. That terminology once meant a level of autonomous capability meeting or exceeding that of a human but has since become yet another moniker without meaning.

Case in point, productivity software and cloud services vendor Microsoft (MSFT) just reached a deal with OpenAI to restructure an earlier ownership/tech-sharing agreement. A sticking point in the deal was the definition of AGI, the achievement of which—as determined by the board of OpenAI—was meant to have enabled OpenAI to limit/terminate Microsoft's exclusive access to OpenAI's technology. Per reporting on the matter, OpenAI had suggested a financial benchmark—\$100 million in profits—to represent the trigger. Arbitrary and irrelevant we presume Microsoft countered, wanting the restriction removed instead. The two now have agreed to have a yet-unnamed panel of experts define AGI and determine when it's been reached. No one knows what AGI is. They think they will know when they see it. But we remain very far from the metaphorical corner around which AGI exists.

A Bubble to Burst?

In our view, the gaps between "what AI could be" and "what AI is" and between "what AI costs" and "what AI makes" are chasms unlikely to be spanned by funding alone and over a timeline approaching that which we see embedded in the AI investment theme. Placing those observations against the weight of the AI trade in the U.S. stock market and concerns that we are in some manner of bubble are both warranted and worrying.

Countering those fears, many don't find the could-be/is now gap as stark as we do. And, of course, AI boosters have a retort. For example, further down the line, LLMs will give way to "world models", which will (by their conjectures) greatly expand the capabilities and applications of AI tech. World models are variously thought of as digital representations of real-world scenarios that should enable AI to "think" in an environment based on "rules" rather than, as LLMs presently do, simply predicting text outputs. Seems the basis for world model projections are the great successes of mostly game-playing supercomputers that have mastered Go and chess. And meaningful progress is likely to be made in world models that incorporate physics for use in autonomous piloting in all its forms. But those "worlds" are incredibly narrow relative to the actual world—which humans navigate within a vastly wider set of set physical rules (e.g., gravity) and using ever-expanding and -evolving knowledge and experience gain through recognition, cognition, interaction and repetition. Relevant rules and contextualizations within both sets are likely too incomprehensively—and therefore potentially incompatibly—numerous, to resolve into a computer model. Given precedent, then, we won't be surprised to find the concept of world models diluted into something vacuously narrow at some point in the future.

So, we expect that the cost/benefit chasm will remain wide for some time to come. Investor reaction to Meta's and Microsoft's Q3 results leave us thinking investor suspicion and scrutiny are on the rise. But a look at shares in Alphabet (GOOG/GOOGL) and Amazon (AMZN) since those two provided their calendar Q3 results leaves us skeptical that a crash is nigh. But we do expect one to come, left wondering now at the onset of and eventual pace at which an alternative reality sets in. It may be that the present opacity regarding potential revenue from AI products/services is

the very thing that keeps hopes high for AI companies and their shares and that will continue to provide support for stocks over the near and even medium term. As actual revenue figures become more widely available, however, more realistic long-term outlooks will follow. And as growth expectations potentially wane, so too may the now high-flying shares in AI-adjacent companies.

And that may include shares in Nvidia (NVDA). While we might think that the chipmaker will retain its technological superiority in the dedicated chip space, perhaps forever, we remain convinced that infrastructure budgets will stabilize, perhaps even begin to drift lower next year and beyond as the pressure to show paybacks on the spending ramps. And the denouement of this chip cycle we think will have an effect on NVDA shares similar to that of the bust of the crypto mining craze. At a mere 29 times fiscal 2028 (ending 01/31/2028) full-year revenue, according to Bloomberg, the stock seems cheap, no? But that multiple presumes a 55 percent year-over-year jump in current quarter revenue, a 34 percent gain next fiscal year and another 20 percent or so the next to \$322 billion. Meantime, profitability is set to at worst be flat, as earnings soar 45 percent, 43 percent and 23 percent in that same period order, to \$7.04 per share. It would seem announced plans for AI infrastructure might just support those estimates. We expect, however, that many of those plans won't come to fruition.

Still, while we may be worried about excess investor enthusiasm for AI-adjacent names we think an immediate reckoning is not upon us. And to be clear, one may never arrive. Perhaps we are totally wrong in our expectations for eventual remedy to LLM imperfections. Maybe AI-related revenue will soar from new services, existing product overlays and otherwise. AI infrastructure budgets might well grow even from here. AI-space valuations presently expect such soon to prove the case, however, and it's not unlikely the world will turn out otherwise.

A financial crisis of any sort we yet don't expect to see though. Given that the majority of potential excess capex so far has been bought using mega-cap Tech cash flow (not debt), we continue to believe that any negative stock reactions will come as a result of disappointment, rather than contagion. To the extent that AI infrastructure investors (including the Magnificent 7) shift further into debt (with the help of private credit providers), the potential for contagion may grow. But we think we remain well removed from any manner of financial crisis.

Whether from declining infrastructure spend, the collapse of a small player (or players) in the infrastructure space, disappointing actual AI revenue trends or a mix of the above (likely to be the case), however, a crisis of confidence may be in the cards. And that might mean a substantial decline in the U.S. equity market, given the present weight of Nvidia and the hyperscalers of just under 30% by market cap. While our strategies generally maintain exposure to the broader market, including many of the more extremely valued AI-related names, we generally are underweight those names relative to a passive exposure based on market capitalization. Any flight from AI stocks, therefore, is likely to negatively impact our portfolios, but we imagine to a lesser extent, given our generally lighter exposure to the trend. Just as was the case after the Tech Bubble, we tend to believe that funds that had been invested on-trend may shift to rather more unloved sections of the U.S. market and abroad. And that shift may find those stocks sporting characteristics we tend to appreciate gaining favor from investors increasingly weary and wary of the AI theme.

Important Information

Statera Asset Management is a dba of Signature Resources Capital Management, LLC (SRCM), which is a Registered Investment Advisor. Registration of an investment adviser does not imply any specific level of skill or training. The information contained herein has been prepared solely for informational purposes and is not an offer to buy or sell any security or to participate in any trading strategy. Any decision to utilize the services described herein should be made after reviewing such definitive investment management agreement and SRCM's Form ADV Part 2A and 2Bs and conducting such due diligence as the client deems necessary and consulting the client's own legal, accounting and tax advisors in order to make an independent determination of the suitability and consequences of SRCM services. Any portfolio with SRCM involves significant risk, including a complete loss of capital. The applicable definitive investment management agreement and Form ADV Part 2 contains a more thorough discussion of risk and conflict, which should be carefully reviewed prior to making any investment decision. Please contact your investment adviser representative to obtain a copy of Form ADV Part 2. All data presented herein is unaudited, subject to revision by SRCM, and is provided solely as a guide to current expectations.

"U.S. stocks" are represented by the S&P 500 Index measures the performance of the large-cap segment of the U.S. equity market.

The opinions expressed herein are those of SRCM as of the date of writing and are subject to change. The material is based on SRCM proprietary research and analysis of global markets and investing. The information and/or analysis contained in this material have been compiled, or arrived at, from sources believed to be reliable; however, SRCM does not make any representation as to their accuracy or completeness and does not accept liability for any loss arising from the use hereof. Some internally generated information may be considered theoretical in nature and is subject to inherent limitations associated thereby. Any market exposures referenced may or may not be represented in portfolios of clients of SRCM or its affiliates, and do not represent all securities purchased, sold or recommended for client accounts. The reader should not assume that any investments in market exposures identified or described were or will be profitable. The information in this material may contain projections or other forward-looking statements regarding future events, targets or expectations, and are current as of the date indicated. There is no assurance that such events or targets will be achieved. Thus, potential outcomes may be significantly different. This material is not intended as and should not be used to provide investment advice and is not an offer to sell a security or a solicitation or an offer, or a recommendation, to buy a security. Investors should consult with an advisor to determine the appropriate investment vehicle.

Publication: 11.03.25

2025-SRCM-44