

AI: the capex cycle survives the weekend's safety debate



Image source: iStockphoto.com/pingingz

Jakub Dubaniewicz

Senior Equity Analyst

jakub.dubaniewicz@syzgroup.com

A weekend debate over “recursive self-improvement,” the risk that AI could accelerate its own development to the point where humans struggle to stay meaningfully in control, prompted some in the industry to call for a more cautious pace of development. We think the correction in AI stocks on Monday was overdone. Investors may be conflating two separate questions: how fast frontier AI capability should advance, and how fast the world adopts the AI that already exists. The first is genuinely more uncertain. The second isn’t, and it’s the second that drives datacentre capex cycle.

What triggered the Monday’s AI pullback

The AI debate took a more consequential turn last weekend, as leading AI researchers and executives focused not just on how quickly AI is improving, but on the possibility that AI could increasingly help develop the next generation of AI itself. The discussion around recursive self-improvement has revived questions about whether the pace of development could become difficult to control, prompting calls from some in the industry for greater safeguards and potentially a slower pace of frontier development.

The market reaction was immediate, with Philadelphia Semiconductor sector index losing -5.6% on Monday, and many other AI-exposed hardware vendors outside the semiconductor sector suffered similar declines. We think that reaction was too broad: while the debate could eventually affect the pace of frontier model development, there is little evidence so far that it is changing the much more immediate investment cycle around AI deployment and datacentre infrastructure.

Why we think investors are conflating two separate questions

The debate over the weekend was about frontier development: should the pace of building more capable AI slow down, given the risk that humans lose meaningful control over increasingly autonomous systems? That is a real and unresolved question.

But the market priced it as if it were a different question: will companies keep buying the computing capacity to run the AI that already exists? Those are not the same question, and it is the AI capex investment outlook, not the frontier-safety debate, that drives stock prices today. The first has genuinely become more uncertain this week. The second has not moved at all.

Recursive self-improvement is not yet the autonomous intelligence explosion the phrase implies. AI today assists human researchers, writing and debugging code, generating training data, designing experiments, evaluating models, rather than independently redesigning its successor end to end, implementing the redesign, and iterating without a human in the loop. The distinction matters because it defines where the actual risk sits. Today, humans still set the research agenda and decide whether results are good enough to deploy, even as AI does more

of the underlying work. Recursive self-improvement, in its fuller form, would erode that human bottleneck: each generation of AI designs a better research process, that process produces a more capable model, and the more capable model designs a still-better process again. If the humans supervising that loop eventually stop being able to fully explain why the new system is more capable, oversight, not compute, becomes the binding constraint. That is a legitimate reason for caution at the frontier.

It is also a more complicated argument than “AI is dangerous, slow down.” AI development is a geopolitical and commercial competition as much as a technical one. If one lab or one country voluntarily slows while rivals keep pushing, the party exercising restraint simply cedes technological ground. A unilateral pause is not commercially or strategically sustainable in the way a coordinated, multi-party slowdown might be. Coordination of that kind requires agreement not just between companies but ultimately between governments, a materially higher bar. That is why we’d expect any deceleration in frontier development to be gradual and contested, at best.

Slower frontier progress does not mean slower infrastructure spending

The read-through to infrastructure spend is where we push back on Monday’s price action. The market’s implicit logic looks like this: if frontier development slows, fewer and models get trained; if fewer models are trained, fewer GPUs are needed; if less compute is needed, the datacentre investment cycle slows. We think that chain is too linear, because it treats the AI infrastructure build as a bet on the next frontier model when it is increasingly a bet on deployment.


The first phase of this investment cycle was dominated by training ever-larger frontier models. But the next (current) phase is broader: companies need capacity not just to train models but to run them, inference, networking, storage, power and cooling, on top of raw datacentre capacity. That build has its own physical constraints that are largely decoupled from the frontier debate. A datacentre takes years to plan, permit, construct and connect to the grid; semiconductor supply agreements and power procurement are similarly locked in years ahead. The capex decisions being made today are therefore bets on the pace of broader AI adoption across search, cloud, enterprise software, coding and customer service, not bets on how capable the next model will be.

The corporate evidence supports that framing. Hyperscaler capex guidance has moved up through the year, not down, with several management teams raising outlooks rather than trimming them. Nvidia’s own guidance points the same way, with commentary suggesting continued demand from hyperscalers and AI labs rather than an imminent slowdown. Directionally, nothing in recent guidance from the companies that are spending the money signals an air pocket. On the contrary, the current guidance implies that the combined AI capex of 6 top hyperscalers will reach \$1.5tn in 2027, a

70% increase YoY. Moreover, the magnitude of the pre-committed spend of \$2.7tn between these 6 companies suggests that their capex plans will extend well beyond next year.

Chart 1: Hyperscaler Capex to Grow +70% to \$1.5T+ in 2027


Capex \$ In Billions



Source: Company reports and JP Morgan estimates

Chart 2: Hyperscalers Have Already Committed over \$2.7tn In Future Spend

\$ in Billions



Source: Company reports

The efficiency argument cuts both ways

There's also a less obvious, two-sided effect worth flagging rather than glossing over. If safety concerns intensify, evaluation and red-teaming of increasingly autonomous models is itself compute-intensive; more capable systems require more testing and simulation before deployment, which argues for more compute demand, not less. Working the other way, if AI-assisted research genuinely accelerates algorithmic efficiency, researchers could eventually extract substantially more capability per unit of compute, which would reduce the compute needed to hit a given capability level.

Both effects are plausible and point in opposite directions, which is precisely why we don't think "is AI getting safer or more dangerous?" is the right investment question. The right question is narrower: will the growth in AI usage keep outrunning the rate of compute-efficiency improvement? On current evidence, utilisation levels, persistent hardware shortages, and the pace at which businesses are adopting existing models, we think it does, but this is the variable we'd watch most closely for a change of view.

Four signals that would change our view

1. A cut to hyperscaler capex guidance. This is the one that matters most: not cautious rhetoric about safety, but actual downward revisions from Microsoft, Amazon, Alphabet or Meta. That would be far stronger evidence of a structural change than any statement about the pace of frontier research.

2. Easing utilisation and capacity constraints. Today, we observe continued shortages of GPUs, networking equipment, power and datacentre capacity are a sign that underlying demand remains strong. A reversal of that would matter.

3. Deteriorating inference economics. If AI applications stop generating enough value to justify the compute they consume, demand growth slows even without any change at the frontier.

4. A step change in AI R&D productivity from recursive self-improvement. This is the hardest to observe in real time, and it could cut either way for infrastructure suppliers: a genuine acceleration could raise the value and volume of compute demanded, or it could make AI development dramatically more compute-efficient and reduce it.

None of these are visible yet. They remain questions about the future rather than evidence that the current investment cycle has broken.

Final thoughts

We read this correction as a repricing of the frontier-development debate, not a change to the earnings outlook for AI infrastructure. The weekend arguably changed the question from "how quickly can AI become more capable?" to "how quickly should we allow AI to become more capable?"—a genuinely important shift in the safety and regulatory conversation. But for investors, the more immediate question is simpler: are companies still building the infrastructure needed to deploy the AI that already exists? For now, the answer is clearly yes.

Welcome to Syzerland®

For further information

Banque Syz SA

Quai des Bergues 1
CH-1201 Geneva
T. +41 58 799 10 00
syzgroup.com

Jakub Dubaniewicz

Senior Equity Analyst
jakub.dubaniewicz@syzgroup.com

This marketing document has been issued by Bank Syz Ltd. It is not intended for distribution to, publication, provision or use by individuals or legal entities that are citizens of or reside in a state, country or jurisdiction in which applicable laws and regulations prohibit its distribution, publication, provision or use. It is not directed to any person or entity to whom it would be illegal to send such marketing material.

This document is intended for informational purposes only and should not be construed as an offer, solicitation or recommendation for the subscription, purchase, sale or safekeeping of any security or financial instrument or for the engagement in any other transaction, as the provision of any investment advice or service, or as a contractual document. Nothing in this document constitutes an investment, legal, tax or accounting advice or a representation that any investment or strategy is suitable or appropriate for an investor's particular and individual circumstances, nor does it constitute a personalized investment advice for any investor.

This document reflects the information, opinions and comments of Bank Syz Ltd. as of the date of its publication, which are subject to change without notice. The opinions and comments of the authors in this document reflect their current views and may not coincide with those of other Syz Group entities or third parties, which may have reached different conclusions. The market valuations, terms and calculations contained herein are estimates only. The information provided comes from sources deemed reliable, but Bank Syz Ltd. does not guarantee its completeness, accuracy, reliability and actuality. Past performance gives no indication of nor guarantees current or future results. Bank Syz Ltd. accepts no liability for any loss arising from the use of this document.