

A Cold Shower for the AI Mania

di Raghuram G. Rajan

Although generative AI tools have improved rapidly and now outperform humans across many tasks, the market's current euphoria may not be justified. With AI firms increasingly resorting to debt financing, it is worth pausing to consider all the things that could go wrong.

CHICAGO—AI tools will undoubtedly transform the nature of work. Large language models can already generate referee reports on my own research papers that rival those by human referees. Unlike humans, who are always pressed for time, an LLM “knows” or can access much more of the literature in an instant, and often exhibits fewer biases. AI points out my analytical weaknesses, checks proofs, and makes suggestions for improvement. Only rarely are human reports better, typically because they connect the dots and offer new insights.

Nonetheless, the market euphoria around AI has become worrisome, especially given the extent of large-scale debt issuance by the sector. It is therefore worth considering where in the AI supply chain things could go wrong.

The supply chain starts with producers and designers of AI infrastructure: firms like TSMC and Samsung, which fabricate chips; Nvidia, which designs them; and Cisco, which provides connectivity. Then come the hyperscalers like Amazon, Google, and Microsoft. They are building data centers both for the use of their own AI models and in order to sell compute (processing power) to others. In addition to the hyperscalers are more specialized companies like Equinix (data centers) and, of course, Anthropic and OpenAI, the developers of foundational LLMs.

Finally, there are the individual and corporate end users of AI services. Individual use is growing fast, and corporate use in some areas (software development and customer support) is exploding.

But most large businesses, while experimenting intensely, have yet to implement end-to-end uses. Many still need to organize their historical data to train AI for their own purposes, and to restructure their traditional operations so that AI can be deployed to improve with experience. Moreover, many firms rightly worry about data security, AI errors, and hallucinations that could destroy their brand image. Still, as less conservative younger companies find more AI uses, they will put competitive pressure on older, larger firms to change.

The AI rollout could nevertheless get interrupted in a number of ways, generating risk for debt-funded players. For instance, if graphics processing units, CPUs, and memory chips become faster and more energy efficient, the equipment filling existing data centers could depreciate rapidly, making it harder for them to amortize their costs. And LLMs, which have become extraordinarily capable based on what is essentially next-word prediction, could plateau until some new technique emerges.

For now, AI labs are investing massive sums to train newer, larger models, on the assumption that the first model to reach some magic point where it becomes self-improving will rule the AI world, and reap enormous profits. But this scenario seems implausible. Even if there is such a point, competitors could still match the first mover's model (including by hiring away key employees to obtain technical trade secrets).

So far, no AI model seems to have gained a sustained advantage. Unless Gemini (Google), Claude (Anthropic), and ChatGPT (OpenAI) can eventually differentiate themselves by appealing to specific user segments (or by merging or colluding), it is hard to see where the profits justifying their enormous training investments will come from.

Moreover, although politicians have been largely standing on the sidelines so far, policy interventions to address AI risks and concerns are inevitable. Since data centers consume tremendous amounts of power—driving up the power price for everyone—state and local governments will be under increased political pressure to limit their

construction. In Indiana, for example, multiple counties recently [proclaimed](#) a moratorium on data-center construction.

Projections into next year already suggest that hardware makers and data centers will be unable to supply enough US compute. And as shortages of compute mount, end users will have more reasons to delay implementation. You cannot reorganize all your operations around AI if you have good reason to worry about the reliability of access or reasonable pricing in the future.

Worse, whereas broader use may take longer than many expect, malevolent use by hackers and deepfakers, as well as unsupervised use by children, is growing rapidly. It is not difficult to imagine disaster scenarios—such as a deadly cyber incident, gross data misuse by AI agents, or poorly trained AI models advising children to commit acts of violence against themselves or others (something that has already happened). The chorus demanding regulation and more liability for AI models will only grow louder. The risks posed by rogue AI could even prompt a sorely needed dialogue among major powers, perhaps leading to some kind of AI Geneva Convention.

Perhaps the most important trigger for political intervention would be massive AI-related job losses. Fearful of the political or social backlash, even firms that are inclined to adopt AI may be hesitant to shed redundant employees outside of a recession, reducing any gains from AI deployment and diffusion.

Given all these uncertainties, it is far from clear how widely and quickly AI will be rolled out, and who will profit. Hardware manufacturers and designers seemed well positioned, given the tremendous demand for compute. But if data-center construction is interrupted, that could shift profits to hyperscalers and AI labs. They might reduce the amount of compute dedicated to training better models, which gives them only fleeting advantages, and shift to selling the compute they have sewn up to firms using their already capable models. Such shifts are also likely if model capabilities plateau. Regulation might also force modelers to spend more effort on improving the training and safety of existing models, building broader public trust.

The good news is that a more limited, careful AI rollout could give firms more time to find labor-augmenting (as opposed to labor-displacing) uses, and governments and workers more time to adjust. The bad news is that euphoric visions of quick exceptional

profits could be unfounded, a particular problem for AI firms that have to make unforgiving debt payments. AI advances will likely pay off eventually. But not every provider will profit, or even survive.