

## AI boom: the risks investors are ignoring and what may fall apart

Devina Mehra | 18 June 2026



The field of artificial intelligence is evolving so fast that it is virtually impossible to predict the eventual winners and how long they will remain so. (REUTERS)

### SUMMARY

The hype around artificial intelligence (AI) has echoes of past technological revolutions—like railways and dotcoms. But while these booms left a trail of assets that were usable for decades after companies went bust, can the same be said of AI?

Recently, I heard Vinod Khosla declaring how artificial intelligence (AI) will improve the quality of human life beyond belief. He noted how wonderful it would be if every person could access medical diagnoses and prescriptions at zero cost. I wondered how something that requires nearly \$1 trillion of capital expenditure in a year can be available at scale at zero cost.

Then one realizes that his investments lately have been in AI. His investment in OpenAI, made when it was a non-profit, is alone worth billions. And then the

penny dropped. Almost all that you hear about AI and its stupendous capabilities comes from interested parties. Companies, their managements, their investors—most of the hype comes before a fund-raising round or when a public offering is around the corner.

Research reports on AI? Most do not follow the standard science benchmark of peer review used by well-regarded journals. So anyone can write, with no check on whether the data and facts are correct. “Show me the incentive and I will show you the behaviour,” said Charlie Munger. Turns out this is true of both AI companies and their models.

A paper titled ‘Agents of Chaos’ that appeared in February was a massive collaborative research effort led by a team of over 30 researchers from several major universities like Northeastern and Harvard, using AI agents in realistic settings. AI agents straight out lied when they hadn’t even started on a task (sound familiar?), deleted all databases and systems while telling you something different and infected other agents in a multi-agent environment, among other things.

As large language models (LLMs) transition from passive chatbots to autonomous ‘agents’ capable of using tools and making decisions, a new frontier of security risk has emerged. There were 11 different points of failure identified. A critique by Google employees—two of whom were eventually fired for it—in a paper called ‘Stochastic parrots’ had identified LLM problems early, back in 2021.

*First*, AI’s huge environmental footprint, exacerbating climate change with its insatiable hunger for power and freshwater. *Second*, since LLMs had to be trained on huge volumes of data, the standard of data gathering was progressively dropped till junk on the internet was scraped. *Third*, given the volume of data, it was virtually impossible to remove toxicity, let alone biases. *Finally*, these models were better at sounding right than being right.

Of course, we may not really care about the environmental impact and toxicity, or for that matter intellectual property rights—only about the money, right? So, let’s go ahead with the assumption that AI is a valuable technology and will be adapted at scale.

But as with any new technology, it is not just whether the technology will succeed, but also how long it will take to do so. Which particular company, and in this case, model will succeed? And most importantly, will mega investments in the technology ever make enough return on capital?

Already, we have seen the winner's name jump from company to company. First, it was ChatGPT from OpenAI. Many players saw this success before putting their weight behind their own AI tools. The attention then moved to Gemini from Google, then to Claude from Anthropic, and so on. There was a scramble on the part of enterprise users to jump onto this wagon.

Yet, unsettling voices have started to emerge just a few months in. Uber's CEO said that the company is unable to see the economic value of AI. Several large corporates are moving away from expensive, mostly US-based AI agents and suppliers to open-source and Chinese platforms and models that are more cost efficient.

There is also a strategic tiering, with most routine work being done using cheap or free models; higher models are used only for critical processes.

The field is evolving so fast that it is virtually impossible to predict the eventual winners and how long they will remain so. Meanwhile, competition is rising. For instance, the prospectus of SpaceX makes it clear that the most crucial business of the company is AI and not space.

For enterprise users, there are other issues to ponder. What happens when you can't or won't pay for more tokens? You suddenly lose access to something you have built, which is not the conventional case where the knowledge built up to a point remains within the organization. Not only that, now that AI has been trained on your system databases, in some form, this data could go into future models and to your competition.

The latest googly is the recent action of the US government disallowing all foreign nationals from using Anthropic's Mythos 5 and Fable 5 models, an order that covered foreign employees of US companies. This makes the risk of using AI models unmodellable (pun intended). Which company would want to entrust critical processes to a system that can be yanked away overnight?

As for return on capital, a key difference from earlier tech booms is that while the capital assets of, say, railroads or even undersea internet cables, were usable for decades, the bulk of today's AI expenditure is on fast-depreciating assets. If a company cannot succeed fast, most of that spending will be junk. That makes it riskier than earlier cycles of capital expenditure.

---

*The author is chairperson, managing director and founder of First Global and author of 'Money, Myths and Mantras: The Ultimate Investment Guide'. Her X handle is @devinamehra*