Tomorrow's Affairs
Analysis of today
Assessment of tomorrow

By: Raghuram G. Rajan

Can a slowdown benefit the AI industry?



I have a friend who recently decided to take a year off from his role as a tenured mathematics professor to study AI safety.

With OpenAI releasing 722 manuscripts claiming solutions to 377 previously unresolved mathematics questions, and with AI models getting only better, he is worried about whether his profession has a future. If a brilliant mathematician can be displaced so easily, why are the rest of us not worried?

Part of the answer may be that, although AI may solve a variant of the Navier–Stokes conjecture, my bank has yet to create a reliable chatbot capable of resolving a funds-transfer issue without me angrily demanding to speak to a human.

AI labs are using recursive self-improvement techniques to furnish their unreleased models with even more miraculous research capabilities, and early adopters in some industries do have bots negotiating contracts and reconciling invoices with deliveries.

But adoption by mainstream businesses is moving slowly overall, and the AI capabilities they can use are still modest.

Why, then, are so many AI firms using so much compute to develop more sophisticated models whose capabilities only a few can leverage fully?

AI capability race

Apart from the thrill of creating a new, faster form of intelligence, a common belief is that the first lab to create a fully self-improving model will benefit from its exponentially improving capabilities and pull permanently ahead of the competition.

Yet it seems unlikely that competitors would not follow quickly. The leader could get a month or two of advantage, albeit with exponential capability growth, but fast followers may well discover techniques that allow them to leapfrog to the front.

That is precisely what has been happening in the AI capability race, with model leadership positions changing regularly.

> It seems implausible that any winner will dominate forever

It seems implausible that any winner will dominate forever—especially if they can hire each other's workers—unless they literally disable or destroy their competitors. Hopefully, this is no one's intent.

Another motive may simply be that none of the top AI labs' messianic leaders trusts each other to handle the power that comes with extraordinary AI capabilities.

Each is motivated not just by a desire to win but also by a sense that only they can handle the power well.

Fear of China catching up

Another source of momentum is the fear of China catching up. Chinese AI models have consistently stayed only a few months behind the leading US models.

As with the US capabilities race, it is unlikely that they will fall much further behind unless the US can mount a more effective campaign to deny Chinese access to the most advanced chips, both in China and around the world.

> The enormous investment of compute in model training puts pressure on AI labs to show more revenue—the sooner the better

Regardless of the reason, the enormous investment of compute in model training puts pressure on AI labs to show more revenue—the sooner the better.

That is why Anthropic has launched a joint venture, Ode, to push the use of AI in non-tech companies, using embedded Anthropic

engineers.

That means some banks may have vastly superior chatbots soon, putting pressure on all banks to follow suit. How long before customer-service jobs join those of math professors?

Society is not ready for a large-scale AI rollout

There are, of course, multiple ways to implement AI. You can simply replace existing jobs, or you can redesign work and skills training so that existing employees' capabilities are augmented enormously.

The former will produce revenues faster but result in social carnage; the latter will take more time and thought, but preserve jobs and stability.

> Businesses that are being targeted for increased AI use are in the best position to implement employee-centric AI

Businesses that are being targeted for increased AI use are in the best position to implement employee-centric AI, but they need to take more responsibility for their employees' future, even as the government provides more incentives for corporations to resist the short-term profits obtainable through layoffs.

There are many workable ideas—such as taxes on corporate token use and credits for employee training—but few governments have them in place.

The conclusion seems unavoidable: Society is not ready for a large-scale AI rollout.

A collective AI slowdown

Yet we may now be seeing a push to slow the tremendous build-out of compute.

Many political leaders in the United States are demanding a pause on new data-center approvals, echoing the concerns of local communities and calling attention to the broader consequences of an uncontrolled AI rollout.

Faced with the prospect of fewer computation resources being available for monetizing AI over the medium term, investors will be forced to limit allocations of costly resources to model development.



A construction slowdown also raises the cost of AI usage, which will change the calculus behind decisions to replace employees

A construction slowdown also raises the cost of AI usage, which will change the calculus behind decisions to replace employees.

It allows the supply of chips and memory to catch up, so that costs don't rise inexorably into the future.

And it allows AI regulation to catch up so that AI lab founders are not the only ones determining how much existential risk—or even risks to the integrity of the financial system—the world should tolerate.

The intent of a pause should not be to stop AI development or usage, but rather to buy time to implement the technology more deliberately.

The race with China remains manageable. US labs will not stop all development, and given what we have seen so far, the risks of anyone

obtaining a permanent advantage are small.

Besides, the Chinese are not insensitive to AI risk and may see some merit in slowing down their own development if they see the US doing so.

A collective AI slowdown might even jump-start overdue discussions between China and the US on how to harness AI for the benefit of humankind, while addressing the risk of us unwittingly becoming its serfs.

Raghuram G. Rajan, a former governor of the Reserve Bank of India and chief economist of the International Monetary Fund, is Professor of Finance at the University of Chicago Booth School of Business.