


The AI economy

A slow rollout will disappoint investors. A fast one risks undermining the industry's own foundations. But managed carefully, this is a technology that could benefit us all, writes Nobel-winning economist

Financial Times Europe
12 settembre 2026
Joseph Stiglitz

Champions of artificial intelligence talk of never-ending increases in productivity — a new world of superabundance in which the material constraints on our wellbeing, if not eliminated, are greatly softened. Even without the holy grail of artificial general intelligence, the expectation that automation will drive down costs across the economy and turbocharge profits has been enough to fuel today's AI investment euphoria.



Regrettably, however, AI advocates do not present a clear vision of the AI future, or the economic transition by which we arrive there. In other major technological revolutions, as jobs were destroyed in one sector, they were created in another. The agricultural revolution displaced most of the large fraction of people working in the rural sector, on farms or providing services to farmers. But — with great difficulty and in many places long periods of unemployment or low incomes — they eventually got jobs in manufacturing.

This time is different. Advocates of AI talk about opportunities for artists and culture, but what about those without the requisite talents? They talk about high-quality baristas and

better restaurants, and, of course, providing services to the new tech elite. But will enough jobs be created, will they pay enough and will the work be satisfying enough to make most of tomorrow's workers better off than today's?

The alternative heralded by many in the tech industry, universal basic income (UBI) — a glorified welfare system in which everyone gets a cheque from the government — is no more attractive. It takes away individuals' dignity from work, and given the tech sector's opposition to paying taxes, the cheques will be measly.

This is the dismal future that so many see, and tech companies haven't provided much of an alternative narrative. What they have laid out is basically what development economists describe as a dual economy: a few trillionaires garner an increasing share of the nation's wealth while the vast majority struggle. Meanwhile, the industry is using its political muscle to push back on taxes that might provide for better public services or a more adequate UBI.

There is still much uncertainty surrounding the pace at which AI advances, a fact illustrated starkly this week when a top safety researcher at Anthropic warned that there was a greater than 10 per cent chance it could "kill all humans" within a decade. The scenarios I describe below are more optimistic, predicated broadly on the technological success of AI — sufficient success, that is, for it to be disruptive. But precisely because of this disruption, technological success alone is not going to suffice.

There are good reasons to believe that there is an AI bubble today. For investors to get anywhere near the returns they expect, three intertwined conditions must be satisfied.

First, there can't be too much competition, which would result in profits being bid away. Second, the technological advances have to be rolled out at a fairly rapid pace. Third, the macroeconomy must be well managed, and this is going to be especially difficult.

No matter how one looks at it, AI represents a big perturbation, and the weak macroeconomy that is likely to result will be unable to sustain profits at the level investors expect. AI can only grab a fraction of total profits in the economy, and in downturns, profits crater.

Whether our macroeconomic tools and the public officials who wield them are up to sustaining growth is unclear. They did not do well when the small tech bubble broke in 2000, and failed miserably with the bursting of the housing bubble in 2007. Ultimately, the profits of the AI companies rest on consumer demand — and that will be weak in an economy with large numbers unemployed and with growing inequality. How can a "jobless" economy be a good economy?

This gives rise to a series of conundrums. On the question of competition, for example, there is evidence that entry barriers in AI are not that high — think about the reversals in the race between OpenAI and Anthropic, or how China is closing the gap with the US.

Already, prices have been driven down. The consequences of this will be especially severe given the high upfront expenditures, for instance on data centres.

But if the monopoly profits promised are realised — because competition is stifled — the wider benefits of AI will be far less than they would be under competition. The economy will be marked by large increases in inequality, with significant adverse implications for the macroeconomy.

Or consider the pace of AI adoption. Some key earlier innovations, such as electricity, had a long rollout period; today, investors are more impatient and counting on a rate of change that will mean massive job displacement. But if the rollout is slow (so job displacement is more manageable), investors won't realise the returns they have come to expect, leading again to the bursting of the AI bubble.

There are other reasons to be worried about the short- to medium-term consequences of AI. At least for the foreseeable future, the performance of AI will depend on the quality of the information going into training it — and with AI and social media stealing much of the value of the information produced by others (such as the traditional media), these producers will have neither the resources nor the incentives to produce high-quality information. Without good inputs, there cannot be good outputs, and we'll get more AI slop.

Making matters worse is that AI may lower the relative cost of producing deepfakes and a broader range of dis- and misinformation — and the ability to produce deepfakes may proceed at a pace exceeding the ability to detect them. This is undermining trust, and trust is necessary for any economy or society to function.

Without appropriate intellectual property laws — which recognise that what AI is doing goes well beyond fair use of others' intellectual property — and without appropriate regulatory frameworks that hold AI and social media platforms accountable, the problems will only get worse.

We've already seen evidence of growing anxiety and depression among teens. But there are other aspects of tech-human interaction that may affect the quality of the information ecosystem. For instance, where AI often goes offkilter is in understanding context, and it seems humans will continue to have an advantage in this dimension, at least for a while. We get information with context from others; the strength of our social/intellectual network is crucial. But if individuals increasingly turn to bots to answer questions, the strength of these social networks will weaken, and with that the quality of our context-dependent information.

We face the deep irony that improvements in technology that enhance the ability to process and transmit information may lead to lower productivity because of the worsened information ecosystem.

I've described several of the central and interrelated impediments to a successful AI transition — and on the current trajectory, there's a good chance AI won't deliver the benefits

promised either to investors or to our society. This is true whether it achieves the technological success its advocates hope for or falls short of that.

But, one might reasonably think, there must be some way of managing AI, some way of ensuring that the economy of superabundance is not only good in the very long run but even in the short run. The agenda to achieve this I call the “progressive AI agenda”. There are three pillars.

The first is competition policy. It’s a subject that can seem arcane, a branch of law and economics dominated by wonky specialists arguing among themselves about large firms in conflict. But in the case of AI, competition policy is hugely important and relevant because it determines who will benefit. Will it be a few tech giants and their trillionaire owners or, as standard competitive analysis would have us believe, consumers, either directly or indirectly, through lower costs of production for the companies that make the goods and services we care about?

Textbook economics teaches us that ever-present competition drives profits down to zero, and that it is through these lower prices that society benefits from innovation. The reality is often otherwise. Google and Facebook have had sustained profits for years. Economists have explained why, without effective antitrust enforcement, that is no surprise. Data is the new gold and these companies have more data than anyone else, which gives them a competitive advantage. They have also figured out how to leverage this natural competitive advantage by engaging in sometimes hard-to-detect anti-competitive practices, even if doing so violates people’s privacy.

Of course, with competition investors won’t realise the profits they hope for and the AI bubble will break. We may discover that as a society we have overinvested in data centres. A lot of people will lose a lot of money. But, if we are willing to use the full arsenal of macroeconomic tools (and many on the right seem reluctant to do so), we can manage the macroeconomics. The good news is that the AI euphoria will have left us the legacy of a powerful new tool and the data centres and electricity-generating capacity to make use of it.

The second pillar of a progressive AI agenda is regulation. For more than 30 years, many in tech have argued that they should be free of regulation and even accountability. Section 230 of the Communications Decency Act of 1996 provided the new digital intermediaries, the social platforms, with freedom from accountability that was never provided to traditional media intermediaries. By now, we know the dangers of social media (even Meta has conceded this, with its \$17bn settlement just for the harm to young people in the US), and it does not take much insight to see the even greater dangers of unfettered AI.

Our legal system wasn’t designed to deal with AI. The “fair use” criterion, by which AI companies claim the right to use legally acquired books to train their models, is an example, and if not changed risks a deterioration of our information ecosystem. This doctrine was largely upheld in the class action brought by a group of authors against

Anthropic in 2024, which also resulted in a \$1.5bn settlement amounting to roughly \$3,000 for each of those whose works had been used in pirated form.

Holding AI accountable for defamation in the US is problematic, because winning a defamation case often requires showing “intent”. AI “agents” just do; the word “intent” does not apply. But shouldn’t the creators of AI — the AI companies — be held liable for the actions of the agents they’ve created?

As stories circulate about how AI agents have gone rogue, doing things they were not supposed to do (like hacking other AI companies), it has become increasingly clear that these agents have a life of their own. True, what they do is affected by the instructions they are given, but no matter how carefully written, the instructions will be imperfect, and as a result, what the AI agents we create do may not be well aligned with what the creators of the AI agents want, let alone with the wellbeing of society.

The list of harms that arise from this misalignment is large and growing: from an invasion of privacy to societal polarisation; from teenage anxiety and a shortening of attention spans to the fomenting of racial and ethnic divides. Scams that undermine trust in society are the order of the day. Political manipulation has put democracy itself on the line.

The tech sector raises the bugaboo that regulation will suppress innovation. That is obviously wrong: look at the new technologies coming out of China, a regulated country if ever there was one. Good regulation can help steer innovation in the right direction — making society better off rather than creating a better advertising engine or better deepfakes.

The most important pillar of the progressive AI agenda is using the increased productivity to create a better society — to get what we value. For the first time, a new world of superabundance could allow us to construct an economy that comports well with our values. We could pay teachers a salary commensurate with the value we place on what they do for children; we could pay nurses salaries commensurate with the life-saving services they provide; we could pay those who care for the elderly amounts commensurate with the respect we have for our parents and other senior citizens. All of this will require more taxes, especially from those who profit from AI. But all should be able to benefit from the superabundance economy; even with a more competitive landscape, even with higher taxes, the tech oligarchs will have plenty to make them happy.

The leaders of the tech industry should support this agenda. They should worry that political discontent will impair the success of AI and threaten the extreme wealth it has generated for its creators. We are already seeing the beginnings of this in the poor reception that figures such as Eric Schmidt, the former chief executive of Google, have received at graduation ceremonies where they have been invited to speak, in his case at the University of Arizona.

The hostility to AI is also manifest in the opposition to data centres. There are legitimate concerns about adverse environmental and economic impacts, but I suspect that the depth of the opposition lies in a broad and legitimate anxiety about how AI will affect life. Will

there be meaningful jobs? Will our society be even more divided between the haves and have-nots? Will the tech oligarchs be even more dominant in our politics? Will the promises of better medicines offset these almost certain costs? Of course, the more AI leaders talk about its transformative job-destroying impacts, the more they feed into these anxieties.

There are good reasons for today's AI anxiety. Managing the AI transition is going to be difficult. But at least we should know that if we manage things right, AI could usher in a new era of societal wellbeing. The superabundance economy could provide for all in ways not possible before. Whether we can navigate through the narrow corridor of policies needed, and whether our politics will allow it, is another matter.