The limitations of AI and their implications for the economy

Public debate on AI oscillates between dismissal and alarmism, but both extremes stem from the same misunderstanding: misreading what the technology actually does. Pattern recognition and imitation do not amount to capacity for reasoning or creativity, and that distinction will shape AI's ultimate impact on jobs and productivity.

Wolfgang Münchau

Abstract: Assessments of AI's economic impact are often distorted by a flawed understanding of intelligence — one that conflates pattern recognition with reasoning, and imitation with creativity. AI excels at the former, defining both its power and its limits: it can outperform humans in structured, repetitive cognitive tasks while remaining incapable of original inquiry or genuine creativity. Early productivity data support this distinction. In the United States, a 10 percentage point increase in AI adoption between 2019 and 2025 is associated with a cumulative productivity gain of 2.9 percentage points, meaningful but far from transformative. The more consequential divergence, however, is geographic. In 2026, 43% of U.S. employees used AI compared with 32% in the EU, while adoption gaps within Europe continue to follow familiar divides between Northern and Central Europe on the one hand, and Southern and Eastern Europe on the other, rooted more in management practices than in access to technology. Europe's combination of strong employment protection, precautionary regulation, and hierarchical corporate cultures is likely to slow the labour reallocation required for AI-driven productivity gains to materialise. The distributional effects also challenge conventional assumptions. Evidence suggests that AI may disproportionately augment lower-educated

workers while putting pressure on structured mid-level roles, reversing the pattern observed during earlier technological transitions. AI will destroy some jobs and create others, but countries with more rigid labour markets are likely to absorb the transition more slowly, deferring the productivity dividend rather than avoiding the disruption itself.

Introduction

If you have the wrong ideas about human intelligence, you are guaranteed to not understand artificial intelligence either, let alone assess its economic impact. The way we thought about intelligence in the past was lazy and wrong. We compressed a multidimensional phenomenon into a one-dimensional linear scale. Then we developed a convenient metric, the IQ test, to measure it. Financial economists did the same when they developed the fateful "value-at-risk" metrics that played a role in the Global Financial Crisis by giving bankers a false sense of security.

Our overconfidence about human intelligence is the reason behind many of the misjudgements about artificial intelligence. Eric Schmidt, the former CEO of Google, believes that AI will become "as smart as the smartest mathematician, physicist, artist, writer, thinker, politician". Right now, the discourse about AI straddles the entire spectrum from complacency to alarmism. The Luddites of this world, which includes a majority of my own profession of journalism but also many economists, have persistently underestimated the impact of AI. The European financial media tend to associate the word "AI" with the word "bubble". Jason Furman, Barack Obama's chief economic adviser, and now a Harvard economics professor, made a revealing comment recently: "A few months ago the discourse was about whether or not AI was a bubble. Now it's shifted to whether or not we're about to enter a dramatically new era."

Really? Anybody who has seriously engaged with AI and its underlying technology has known for a long time that AI would not be a bubble. Some market valuations may have been too optimistic, but that's a different story.

Perhaps the most common misjudgement, to which the tech industry is prone to, is getting the direction of the forecast right, but timing all wrong. One person who took this to an extreme was the head of Citrini Research, a financial analyst, who published a fictional AI Armageddon scenario that briefly sent Wall Street into a downward spiral – until they realised it was just a story.

In that story, entitled the "2028 Global Intelligence Crisis"[1] the stock market almost collapsed. It was a dystopian tale of how by 2028 AI takes over a growing proportion of white-collar jobs. Unemployment rises, people spend less money on services. They no longer go on holidays and can no longer afford their mortgages. The financial system collapses. As brilliant as the story is, I think it is wrong – because it is too optimistic about what AI can do. Before we discuss what AI will do to the economy, it is worth reflecting on what AI is good at, and what it is not.

What AI can do — and what it cannot

The technology behind modern AI is the neural network. What it does is probabilistic pattern recognition on a large scale. Pattern recognition also forms an under-appreciated part of human intelligence. We think of chess grandmasters as very intelligent. A lot of their skill is memorisation or a large number of moves combined with superior visual pattern recognition. IQ tests get these concepts mixed up. AI can do logical reasoning by applying the rules of a logical system in a consistent way. Human logical reasoning, by contrast, works differently. It includes, at least to some extent, memorisation and pattern recognition.

"Anybody who has seriously engaged with AI and its underlying technology has known for a long time that AI would not be a bubble."

“ AI is also not truly creative but excels at imitation. ”

The mathematician Terence Tao recently reported that AI managed to find a mistake in one of his proofs. It also managed to come up with the proof of a conjecture in discrete mathematics all by itself. But it does not think like a mathematician. On that point, Schmidt is completely wrong. AI is good at answering questions, but not good at asking them. Some of the most important contributions to mathematics started with interesting theories, like Fermat’s Last Theorem, the Goldbach conjecture or the Riemann Hypothesis – all three related to number theory.

AI is also not truly creative but excels at imitation. AI improves a lot of people’s writing. But AI is incapable of original writing.

I gained some hands-on experience with AI when I developed an AI system by myself, using available open-source technologies, to create a filtered newswire in my specific area of interest, the European economy. The system I developed goes through a long list of publicly-available newsfeeds in most European countries, filters, translates, classifies, and then summarises. This is a non-trivial endeavour even for AI as it combines several different functions. The end result is a much more filtered list than anything else I find on the internet. It gets complicated things right. But it also gets easy things wrong. For example, on one occasion it told me that Pedro Sánchez was the prime minister of Israel. And this is not because it has a sense of humour.

To understand AI at a certain level of depth is important for the discussion about the impact of AI on productivity. If we treat it as a black box technology, we will most likely get it wrong.

Productivity gains do not imply human replacement

In some sectors, AI has already started to replace humans. I would not advise young people to become advertising copy writers or event photographers. When the AI company Anthropic introduced the Claude Cowork platform with legal plugins, it took over the work of junior lawyers and legal assistants. It made the senior lawyers more productive. It got rid of some legal assistance, but not of lawyers themselves.

But be careful with generalising predictions, like those of Schmidt. Even Geoffrey Hinton, the godfather of the neural network, the technology behind AI, committed an error of hubris. In 2016, he predicted that within five years, AI would replace all radiologists. Back then, AI had become very good at detecting tumours on X-rays or other medical images. Hinton must have thought that looking at pictures is what radiologists do all day. His prediction was completely wrong. According to the Royal College of Radiologists, the UK had 3,318 consultant radiologists in 2014. By 2024, that number had grown to 4,923.

Hinton fell for a classic “lump of labour” or “fixed-sized pie” fallacy. In our specific case, the fallacy consists of the idea that if AI takes half of our jobs, the rest of us have to share the remaining half between us. This is essentially what the Citrini report did too.

So, instead of asking how AI will affect the economy, we should start by asking: What effect will it have on existing white-collar employment? Over what period? And will it eventually add more jobs than it destroys?

The AI software developer and computer scientist Francois Chollet once remarked: “I'm so old I remember when fully autonomous cars were going to be ready for mass deployment by late 2017.” In 2017, people predicted that it would happen by the mid-2020s. The self-driving car will be one of the biggest technological innovations of all time. The robotaxi industry has made impressive progress, but they are still at stage four of the five stages of autonomous driving.

> “Even if AI destroys millions of white-collar jobs, it does not mean that it destroys white-collar employment if new industries spring up.”

By September of last year, Waymo robotaxis had clocked 127 million miles collecting data without a driver in the car. When they reach stage five, they will be ready for a bigger rollout. They will still lack human intuition. They will not be perfect, but they will be better, on average, than humans because they don't drive angry, they don't speed, don't drink alcohol, and they don't send text messages whilst driving or fall asleep. When we get to the final stage, Waymo will have been collecting data for some twenty years. And when the self-driving cars and taxis are eventually deployed, at that point we will no longer need taxi drivers.

Even if AI destroys millions of white-collar jobs, it does not mean that it destroys white-collar employment if new industries spring up. Ivan Yotzov et al. [2] presented research results in an NBER paper that would outwardly point towards a contradiction. The survey found that executives predict AI would boost productivity at their companies by 1.4% and cut employment 0.7%. Employees, by contrast, thought it would raise employment by 0.5%. Both statements can be simultaneously true if existing companies cut their workforce, and new companies hire.

Early evidence on productivity and adoption

AI has already had an economic impact, and this has started to show in productivity data. It's still early days. Alexander Bick et al. [3] analysed U.S. and European business surveys to find the impact of AI on productivity. They found that for the period 2019-2025, a 10 percentage point increase in the AI adoption rate in the U.S. would lead to an increase in cumulated productivity of 2.9 percentage points. The authors issued a warning about reading these statistics, but the data seem consistent with the result that the impact of AI on productivity is real, but nothing like what the Citrini scare story would suggest. One of the caveats I would add is that the study related to the period from 2019 to 2025. The really big economic events still lie ahead – AI-based manufacturing in particular. This is an example of where the past may not be a reliable guide to the future.

The report also throws light on the difference in AI adoption rates between the U.S. and the EU, and within the EU. They found that in 2026, 43% of employees in the U.S. used AI against 32% in the EU. The U.S. also has a nearly double adoption rate of AI for manufacturing production. Unsurprisingly, the U.S. leads Europe.

But perhaps the bigger story is the gap between the European AI adopters and the laggards. The UK is at the top, with an adoption of rate of 36%, still behind the U.S., whereas Italy is at the bottom with 26%. The overall picture is a classic European north-south and east-west divide.

But why is this so? It is unsurprising that the U.S. is ahead of Europe. Modern AI is an American technology. The Europeans did themselves no favour by starting to regulate before they even had an AI industry.

Why is the adoption rate so different amongst EU countries? This dates back to a

> “For the period from 2019-2025, a 10 percentage point increase in the AI adoption rate in the U.S. would lead to an increase in cumulated productivity of 2.9 percentage points.”

well-known issue that has been extensively studied by economists: the low adoption rate of computers and internet-technology in European companies. This is a story primarily about management practice, as Luis Garicano reminded us in a recent article. [4]

The reason why Italy became a laggard in adopting ICT technologies was outdated management practices. I have been observing similar phenomena in recent years in Germany, especially in the car industry. The hierarchies in the car industry imposed old views, steeped in 20th century engineering and management traditions, a view that was oblivious to new technological trends.

It is not just the car industry. The chief executive of Rheinmetall, the German defence company, ridiculed Ukrainian drones on the grounds that the biggest manufacturers were housewives with 3d printers. I would surmise that his company will probably not be at forefront of the development of AI-powered drones. Italy started its economic decline around the beginning of the century. German productivity growth has been stagnating since 2018. Germany's economic institutes see potential growth at 0.1% from 2030 onwards. One of the reasons for this decline is Europe's structural attachment to old companies and old technologies. Europe was once, but no longer is, a place for entrepreneurs.

Another not yet widely recognised factor that will limit the impact of AI on productivity is laws that prevent companies from dismissing workers and replacing them with AI. This is already happening in China. The Intermediate People's Court of Hangzhou has ruled against a company that dismissed a worker and replaced him with AI. The case was brought by an employee of a financial technology firm, whose job it was to assess the quality of large-language models. I am not aware yet of any such cases in Europe, but this is almost surely only a matter of time. Given the strong labour protection rights of European workers, it is hard to imagine that large European companies would manage to improve their productivity by shedding workers. The best shot they will have is to make their existing workforce more productive.

There will be no such protection in the U.S. JP Morgan rolled out what it called "human redeployment plans" over the next few years, with a targeted headcount reduction of 10% in operations and account services. Meta, the company that owns Facebook and Instagram, has already announced a 10% cut in its workforce, some 8,000 jobs, not so much because of productivity gains through AI, but to offset the cost of building AI systems. Content moderators are amongst the staff the company is willing to let go.

The data sets we have at this moment do not give us a complete picture, and they probably underestimate the gap between the U.S. and Europe by a wide margin. I cannot think of any European companies planning job cuts to make room for AI. The main impact from AI will come from new companies with AI-era business models, or existing companies that expand. The more interesting question is what happens if these AI-driven firms begin to outcompete traditional companies with large workforces.

Who gains and who loses from AI

AI will be different in one respect from previous technological innovation. The main impact of shocks will not hit the lower end of the income scale – but the middle part. It could make poorly educated people more productive, but it might squeeze those in the middle. Cruces et. al. [5] write in a recent NBER paper about a survey amongst 1,174

> "The reason why mid-level clerical jobs are most affected by AI is that these jobs are structured and repetitive – thus prone to being taken over by AI."

adults aged between 25-45, showing that AI benefits people with lower education more than those with higher education.

We have to be careful here not to draw the wrong inference. AI will itself change education. Memorisation will count for less than critical and original thinking, tasks at which AI does not excel. The reason why mid-level clerical jobs are most affected by AI is that these jobs are structured and repetitive – thus prone to being taken over by AI. The bank employee who spends their time checking on mortgage applications will almost surely be replaced. The chief executive of the bank will not. Manual workers are safe from AI right now, but China is already busy developing humanoid robots that can perform human tasks. They can even run marathons. We should probably not rush to conclusions as to who will be most affected.

But I am confident to predict that AI will be the most consequential technological development of our age. The industrial revolution forced humans to move from rural areas into cities. It destroyed many agricultural jobs but added more industrial ones. The invention of electricity ended the job of the street-lighter, who would walk around the streets at dusk to illuminate our neighbourhoods. They were the car mechanics and legal assistants of their era. AI, too, will destroy jobs. But its main economic effect is that it will make our economies more productive over time.

But it will cause disruption on the way- although nothing like what Citrini or Schmidt are predicting. I am much less worried about the economic effects than the political ones. The industrial revolution drove workers into the cities, but where do the workers displaced by AI go? And what political parties will they vote for?

Countries with flexible labour markets, like the U.S., and to a lesser extent the UK, will get through this economic and social transition faster than countries with a high degree of employment protection like France and Germany. The latter will probably avoid some of the disruptions in the short run, but the quid pro quo is that they will not benefit from the productivity gains that come from AI later.

I think the optimal policy response will be to let this change happen, and at the same time to encourage companies and their employees to diversify into new sectors. My personal hunch is that Europeans will not follow this advice.

Notes

[1] https://www.citriniresearch.com/p/2028gic?hide_intro_popup=true

[2] https://www.nber.org/system/files/working_papers/w34836/w34836.pdf

[3] https://www.brookings.edu/wp-content/uploads/2026/03/6_Bick-et-al_unembargoed.pdf

[4] https://www.siliconcontinent.com/p/what-explains-heterogeneity-in-ai

[5] https://www.nber.org/papers/w34851

Wolfgang Münchau. Managing Director of Eurointelligence LTD