Tomorrow’s AI Won’t Meet Today’s World

AI could reshape the pathways to economic growth. Those looking to grow will adapt to that. Predictions of our economic futures should account for both.

Three years ago, workers in Kenya made international headlines for the conditions they experienced while labeling data for OpenAI. These workers spent their days sorting through hate speech, sexual abuse and violent images — some of the most disturbing material on the internet — for around $1.35 an hour. The labels they produced were then used to help train ChatGPT to identify toxic or disallowed content.

In 2023, much of the outrage was focused — rightly — on the working conditions of these data labelers. But revisited in 2026, this story takes on another dark cast.

Amid a new wave of AI anxiety, people have begun to fear that the broader category of IT service jobs may disappear entirely. Under this view, by helping to train ChatGPT, these workers may have inadvertently been helping to automate away the very kind of tradable service work they performed, narrowing Kenya’s path to prosperity.

But how much evidence is there that AI-related displacement has actually begun in developing countries? And how should we forecast what comes next when the workers, firms and countries affected by AI will themselves adapt, improvise and change in response?

Climbing the Growth Ladder

Historically, manufacturing growth has been the path for countries out of poverty. In Japan, South Korea, Taiwan and, most spectacularly, China, the expansion of export manufacturing pulled millions of people from the farms to better-paying jobs in factories. Under the pressure of international competition, East Asian firms quickly converged to the technological frontier.

But, outside of a few bright spots like Vietnam and Bangladesh, developing countries have largely failed to replicate the success of East Asia. In India, the manufacturing share of GDP has never exceeded 20% and has even declined since 2010. In sub-Saharan Africa, manufacturing makes up around 10% of GDP. New analysis suggests that over the past 30 years, manufacturing productivity in the poorest countries has stubbornly remained far below the global frontier.

Enter IT services.

The arrival of cheap broadband across the world made the outsourcing of information technology services viable for Western companies, enabling them to take advantage of cheaper labor in developing countries. In 1999, India deregulated its telecom sector, accelerating the adoption of high-speed internet and allowing new firms like Infosys and Wipro to become highly profitable. A new sector — in India, “business process outsourcing” (BPO) — was born.

For countries struggling to get growth off the ground, BPO seemed like a lifeline. Like the export manufacturing jobs of yore, export services held out the promise of generating significant foreign exchange revenue while also creating a large number of better-paying jobs. In 2021, the World Bank published a report extolling “the promise of services-led development,” highlighting their exposure to “innovation that improves labor productivity.”

Over the past few decades, countries like the Philippines and Brazil have developed sizable IT sectors, exporting billions of dollars of services a year and employing close to 2 million workers apiece.[1] But, globally, the clear leader of the pack was India, whose IT services sector employed 5.8 million people and exported $200 billion in 2025:

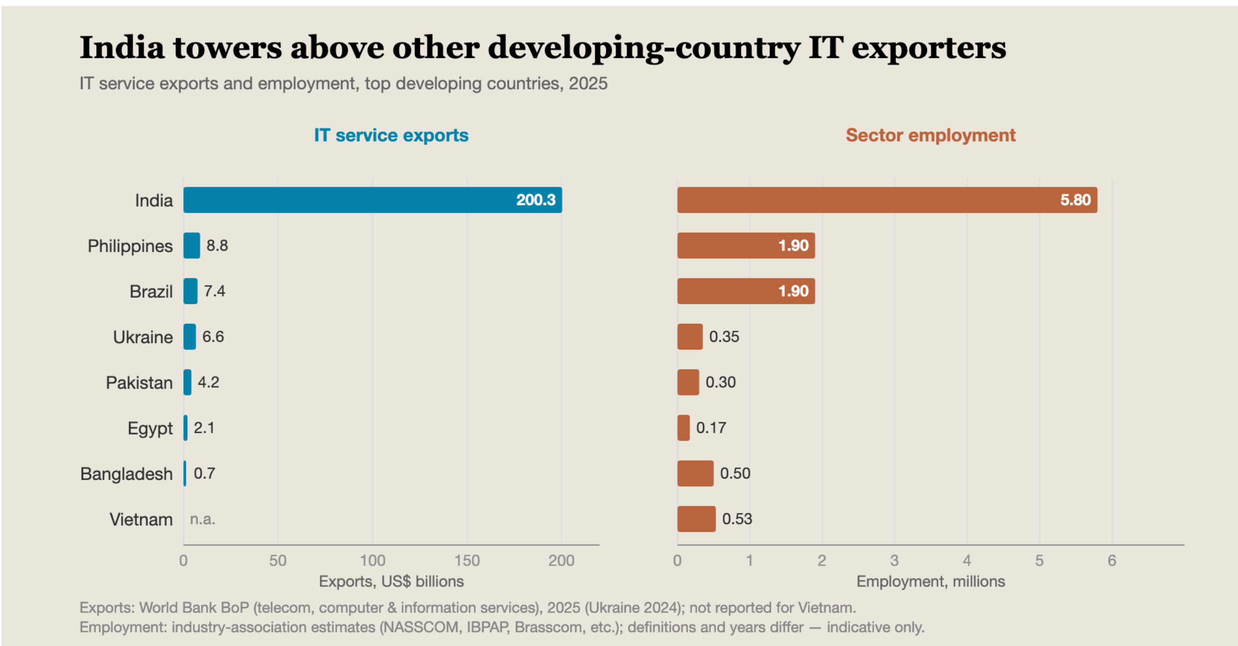

The fear of AI replacing human jobs is now ubiquitous. In America, attention has become fixed on the disruption of entry-level jobs in software development and customer service — the proverbial canaries in the coal mine, the work that should (at least in theory) be most easily automatable by AI coding agents. An influential paper by economists Erik Brynjolfsson, Bharat Chandar and Ruyu Chen found that early career workers (22-25-year-olds, in the U.S.) in the most AI-exposed sectors have seen a 19% relative decline in their employment since 2022.

Naturally, attention has started to shift to AI’s implications for the 6.8 billion people who live in low- and middle-income countries. Thus far, AI has not had much of an impact, positive or negative, on the 2.5 billion people who still depend on smallholder farming. In manufacturing,

[1] The OpenAI story was relatively unusual in featuring African workers; Kenya’s sector employed just 36,000 people in 2025.

the evidence for impact is also limited, though perhaps AI-induced advances in industrial robotics are coming down the pike. It is in service exports that the threat of AI-related job displacement looms the largest.

The Canaries in the Coal Mine

It has now been four years since the launch of ChatGPT. Agents are able to do much of the work of customer service; indeed, corporate support lines have been increasingly replaced with a ChatGPT or Claude wrapper. So how much is automation already displacing service exports in the developing world?

We can take a page from the recent research on AI-exposed displacement by Brynjolfsson et al. in the United States and look at how employment is evolving for the people we'd expect to be most exposed to AI displacement: entry-level tech workers.[2]

Perhaps the best available data from any developing country comes from India's Periodic Labour Force Survey (PLFS), a high-frequency look at trends in employment and wages.[3] We can use what economists call a "triple difference" to pinpoint the effect of AI exposure. First, we take employment for young workers (those 15-24 years old) in AI-exposed industries and subtract the trend for older workers (those 25-39) in the same industries. Then we can do the same comparison of young versus old in industries largely insulated from AI. Last, by taking the *difference* between these two differences, we can isolate the shock to entry-level work in the industries most exposed to AI.

[2] It is likely that older tech workers have acquired both more "soft skills" and more specialized technical skills, making them more difficult to replace with AI.
[3] The PLFS's methodology was revised in 2025, making backwards comparisons tricky. Nonetheless, it's the best data that we have.

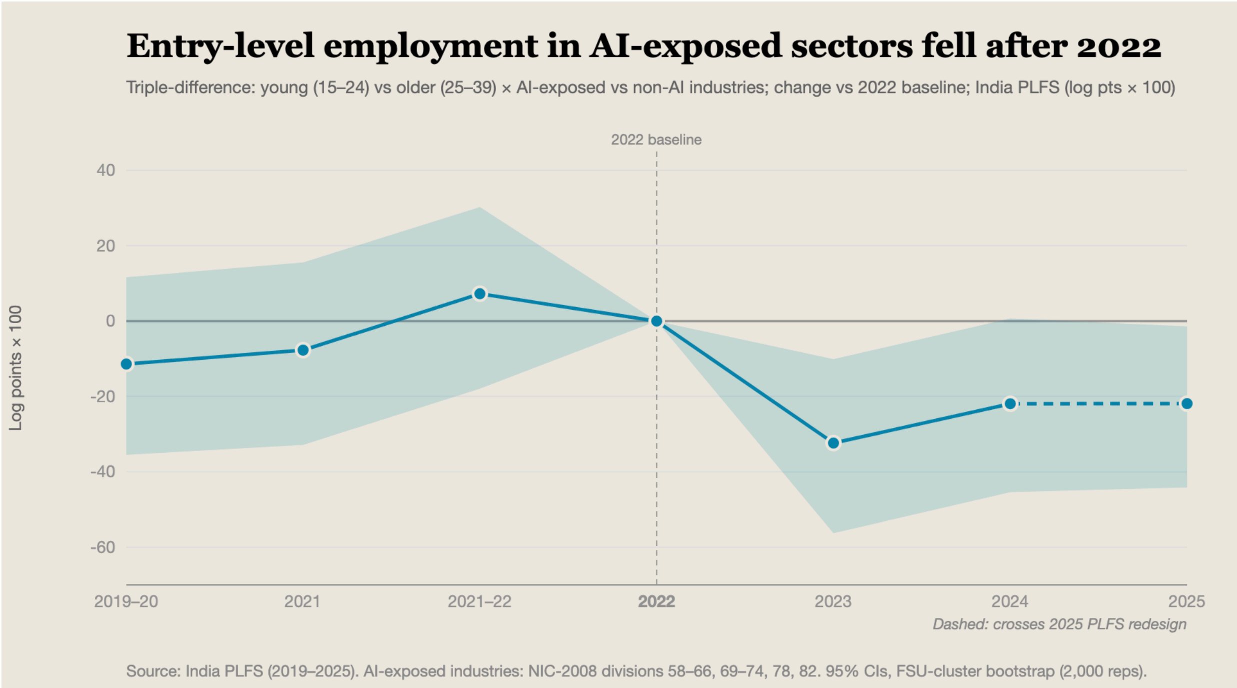

The chart above shows this triple difference on employment, with confidence intervals marked by the vertical bars. For the three years before ChatGPT, there was little gap between entry-level and older workers in the most AI-exposed and non-exposed industries. But after 2022 — the year ChatGPT was released to the public — relative employment for the youngest, AI-exposed workers suddenly fell 32%. Relative employment has not recovered since, remaining around 20% below the 2022 baseline in the latest 2025 snapshot.

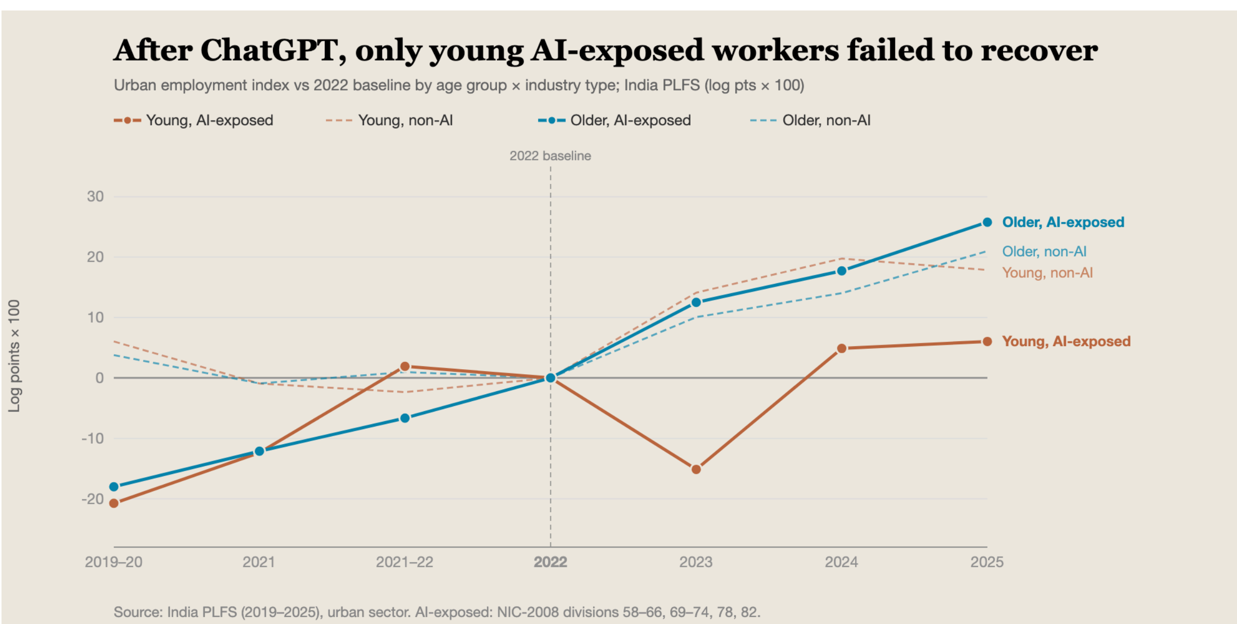

We should be careful about overreading this result. It's possible that the decline reflects structural forces that were occurring at the same time as the rise of generative AI. For instance, some Indian IT firms overhired during the pandemic and froze graduate hiring in 2023.

The absence of a robust employment recovery — amid rhetoric from the BPO companies themselves about increasing AI adoption — gives a hint that the effect may be real. Incidentally, this echoes the result in the latest revision of the work by Brynjolfsson and his coauthors, where the 19% drawdown in employment for AI-exposed young workers has shown no sign of disappearing.

But we should not take this as a sign that all the jobs are going away. This data is from one sector, in one country, among those workers we expect to be most affected. There is (as yet) no evidence of a generalized downturn across *all* age categories or AI-exposed sectors. Even this result is barely statistically distinguishable from zero.

How and when can we draw clearer conclusions? It may be longer than we'd like — in most other developing countries, data comes with far more of a lag.

Flying Blind Since Takeoff

"Gentlemen, you have come 60 days too late. The depression is over." Herbert Hoover, June 1930[4]

In 1930, the Depression was very much *not* over. But how was Hoover to know? This was before Kuznets' "National Income" report was presented in 1934, the precursor to modern national income accounts and GDP. It was also before reliable monthly unemployment statistics were available, as the modern household survey did not begin until 1940.

Many of today's leaders, particularly in low- and middle-income countries, find themselves in a similar position to Hoover in the 1930s — entering a period of economic transformation blind to how their economies are actually changing in real time. India's PLFS is a remarkable effort, and one for which there is no equivalent in other countries that might be similarly vulnerable to an IT exports shock.

Take the Philippines, another country with a significant service export sector. Their annual survey of business and industry is released more than one year after the period it studies. This makes it nearly useless in analyzing high-frequency labor force changes in response to changes in AI capabilities. It might not be until Fable 8 is released that we can fully disentangle the impact of Fable 5 (and today's other frontier models) on the Philippines IT sector's labor force. There simply is no real-time tracking data.[5]

---

[4] Exact wording is uncertain, but he most likely did say something along these lines https://quoteinvestigator.com/2015/10/12/depression/.

[5] There is a growing body of research exploring new forms of measurement, e.g., machine-learning enabled nowcasting, satellite data, mobile phone data and automated classification. It remains to be seen

With the attention and money now flowing to AI economics research, we are likely to be bombarded with similar estimates to our own, from hot-off-the-press working paper abstracts that tell us how AI has displaced some category of worker relative to another or boosted productivity in some company relative to another.

So it will also be tempting to overread the data we do have — to dash off a Twitter thread that launches into the coming employment apocalypse facing India's young IT workers. But each result is only one piece of the puzzle. This result — and others like it — don't tell us about overall impacts.

And then, there's only the small matter of projecting this out into the future. Even if we could perfectly understand economies today, simply extrapolating that out to meet the capabilities of tomorrow's models could lead us to highly misleading forecasts.

Not-So-Dumb Farmers and the Population Bomb That Wasn't

Many of the fears about AI's economic effects arise from projecting the effects of future technologies on a world that looks broadly similar to the one today. But these kinds of extrapolations have a long and checkered history, as we humans tend to be far more adaptable than we give ourselves credit for.

Perhaps the most infamous prediction in international development comes from Paul Ehrlich's 1968 bestseller "The Population Bomb." Extrapolating out observed trends in population growth and limited progress in food production, Ehrlich was alarmed by the possibility of overpopulation and mass starvation. The most optimistic scenario presented in "The Population Bomb," which required "a maturity of outlook and behavior in the United States that seems unlikely to develop," involved the death by starvation of half a billion people (one-fifth of the world's population).[6]

But Ehrlich's projections were based on static extrapolations of human behavior. He failed to anticipate the Green Revolution, which dramatically improved grain yields (particularly in India), and the Demographic Transition, in which fertility fell. Deaths from malnutrition and famines have actually decreased. Ehrlich bet against our oft-underappreciated ability to innovate and adapt, by imposing a rapidly growing population onto a future world whose productive possibilities changed too slowly.

Indeed, the Y2K episode,[7] the demand shock that supercharged India's IT export sector, is another example of predictions not holding up to adaptation. This time, before the clock struck

---

whether new approaches like these can leapfrog the need for traditional administrative and survey data sources.

[6] Paul Ehrlich, *The Population Bomb* (Sierra Club Ballantine Books), 80.

[7] Originally known as the "Millennium Bug"; many computer programs stored years using just two digits, leading to widespread fears that systems would mistake "00" for 1900 and crash en masse on January 1, 2000.

midnight on December 31, 1999, adaptation came in the form of coders in India inspecting and modifying millions of lines of code in order to avert disaster. In the end, apocalyptic predictions of financial, government and transport systems simultaneously crashing across the globe never came to pass, as we adapted to ensure they didn't.

Climate economics, too, has learned a similar lesson over recent decades. Early models of the agricultural damage from climate change often took projected temperatures, ran them through equations of crop yields at different temperatures and projected out damages. Predictably, the estimated climate damages in this early generation of predictions tended to be enormous. But these models baked in the critical assumption that farmers would not adapt to a changing climate.

There will still, of course, be large costs from climate change. And not all adaptations are successful: In U.S. agriculture, adaptations over the long run have mitigated at most half, but most likely none, of the damages from climate change. There are real constraints to adaptation — even where there is a will to adapt, there is not necessarily a way to do so.

AI may play out the same way as these previous challenges. There will likely be losers from adaptation. There may even be thresholds at which the ability to adapt runs out. But if you are guesstimating what the world will look like in five or 10 years' time, you may end up being very wrong if you neglect our future selves' ability to adapt. Tomorrow's AI will meet tomorrow's world, not today's.

This points us to one small but perhaps constructive suggestion for economists. For the past 20 years, causal inference — the ability to cleanly separate cause and effect — has been the sine qua non of academic economic research. By contrast, forecasting has been relegated to secondary or tertiary status, the domain of policy people and (shudder) bankers.

That division of labor made sense when we could reasonably expect that the world of tomorrow resembled the world of today. But if AI really is about to reshape the pathways to growth, the most useful thing economists can do is not to measure yesterday's shock ever more precisely but to get better at predicting the next one.

More accurate visions of our economic futures must account for both machines' capabilities and humans' adaptabilities. Just as humanity didn't sit around waiting to starve, business executives at Infosys are not sitting around waiting to be replaced. There might be fewer jobs to go around, or different types of jobs, but if there's anything they can do about it, the industry born out of adapting to a crisis will try to adapt once more.

Oliver Kim is a development economist working as a Research Fellow on Coefficient Giving's Global Growth Fund. He writes a Substack called Global Developments. Oliver Hanney is a contributing editor at In Development, and is the managing editor at VoxDev, where he hosts the Ideas in Development podcast.