OXFORD JOURNAL
OF STUDENT
SCHOLARSHIP

The Productivity Paradox Revisited: Artificial Intelligence, Labour Market Restructuring, and the Inequality Trap

Syed Ali Husain Rizvi
rizvisyedalihusain@gmail.com

ABSTRACT

The most widely cited figure in the AI-and-jobs debate is a net employment gain. This paper argues that numbers are real and, taken alone, misleading. Using a cross-sectoral synthesis of data from the IMF, McKinsey Global Institute, World Economic Forum, OECD, and peer-reviewed academic sources, three structural problems are identified that the aggregate figure cannot show: a credential barrier that locks displaced workers out of the new roles AI creates; a geographic concentration of AI investment that widens the gap between advanced and developing economies regardless of domestic policy; and a capital-labour distribution asymmetry that routes productivity gains toward asset owners rather than workers. Investment banking is examined as an illustrative sectoral example showing how automation compresses an industry's entry-level pipeline without eliminating it. The analysis concludes that the central policy failure is not a shortage of reskilling programmes but a mismatch between the pace of AI adoption and the capacity of educational institutions to respond. The analysis bears particular weight for India, where 90% of the workforce is informally employed and where the distance between AI's aggregate promise and its distributional reality is greatest. Key limitations, including reliance on projections and cross-source comparability constraints, are discussed explicitly.

Keywords

Artificial intelligence; labour economics; productivity paradox; economic inequality; general purpose technology; India; investment banking

1. INTRODUCTION

I did not set out to write a paper about inequality. I wanted to understand something more immediate: why everyone around me had started using AI, but almost nobody seemed to be thinking seriously about what that meant for their futures. In my home city of Lucknow in 2025, the shift showed up in small, concrete ways. A friend submitting ChatGPT-generated essays. My English teacher setting assignments that had the particular flatness of machine-generated text, something I spotted immediately, being a regular user myself. People outsourcing not just tasks, but thinking. The dependence had become ordinary quickly enough that most people had stopped registering it.

That observation raised a harder question. If AI was reshaping how educated, relatively privileged people in a mid-sized Indian city worked and learned, what was it doing to the hundreds of millions of workers with far fewer options and less room to adapt? The question pushed me toward the economics literature, and the economics literature pushed me toward a deeper problem: the gap between what the aggregate data says and what it actually means.

This paper addresses three research questions: First, do aggregate AI employment projections accurately represent the distributional outcomes for workers across different skill levels and geographies? Second, how does AI reshape the structure of entry-level labour market participation, using investment banking as an illustrative case? Third, what policy interventions can address the distributional mechanisms rather than just their symptoms? The paper makes three arguments corresponding to these questions. First, the net employment gain from AI is real but structurally inaccessible to the workers who most need it, the displaced workers do not hold the qualifications that emerging roles require. Second, the geographic concentration of AI investment creates an asymmetry that disadvantages developing economies regardless of their own policy choices. Third, AI's productivity gains flow primarily to capital rather than labour, which means the economic growth it generates will deepen inequality unless redistribution mechanisms are built in from the beginning.

The headline numbers look superficially reassuring. The WEF Future of Jobs Report 2020 projected 85 million jobs displaced by AI by 2025, and 97 million new roles emerging, a net gain of 12 million positions.[15] Economists reach for the internet as a precedent: predicted to hollow out industries, it instead built e-commerce, social media, digital marketing, and app development from nothing. AI, the argument goes, will do the same.

I am not sure the analogy holds in the way people think it does. It holds in the aggregate and breaks down in the specifics, and the specifics are where the real harm lives. The internet did create new industries. It also produced two decades of wage polarisation, the gig economy, and a concentration of wealth in a handful of technology companies that now exercise more economic power than many nation-states. If AI replicates the internet's trajectory faithfully, that is a less comfortable prediction than it sounds.

Investment banking serves as the paper's illustrative sectoral example because it is the sector I have studied most carefully, and because it shows the broader dynamic in an unusually legible way. AI is not destroying finance. It is compressing the entry-level pipeline, changing who gets access to high-skill financial careers, and doing so in ways that most industry commentary has not yet taken seriously. The analysis draws on IMF, WEF, OECD, McKinsey Global Institute, Goldman Sachs Research, and peer-reviewed academic work, with sustained attention to India given the stakes for a country with 90% informal employment and one of the largest youth populations in the world.

2. METHODLOGY

This paper uses a structured cross-sectoral literature synthesis. Sources were identified through searches of Google Scholar, the IMF eLibrary, WEF Publications, OECD iLibrary, McKinsey Global Institute, and SSRN, using keyword clusters including 'AI labour market impact', 'generative AI productivity', 'automation employment inequality', and 'AI India economy'. The search was conducted between January and March 2025, with a primary time window of 2022–2025 to capture the generative AI era. Approximately 55 sources were identified; 16 are cited based on relevance to the three research questions.

Sources were selected on three criteria: institutional credibility, recency, and data specificity. Sources that disaggregate findings by sector, geography, skill level, or demographic group were preferred over aggregate-only estimates. Key inputs include the WEF Future of Jobs Reports 2020 and 2023, IMF working papers on AI and inequality, McKinsey Global Institute's 2023 analysis of generative AI's economic potential, OECD publications on AI as a general purpose technology, Goldman Sachs sector research, Acemoglu and Restrepo's automation-augmentation framework, PwC's 2025 Global AI Jobs Barometer, and the 2025 METR study on AI's effect on experienced software developers.

Where sources conflict on productivity projections — which they do, often — the range is presented and the methodological reasons for divergence examined rather than selecting a single figure. Scepticism runs through the analysis regarding studies measuring AI's productivity effect over short time horizons or in controlled settings: the J-curve pattern in general purpose technology adoption is well-documented, and early-phase transition costs suppress measured productivity before long-run gains materialise. India appears throughout not because Indian data is uniquely reliable, but because India concentrates the conditions under which AI's distributional risks are most severe.

3. RESULTS AND DISCUSSION

3.1 AI as a General Purpose Technology: what the framework explains and what it skips

The most useful theoretical starting point is the General Purpose Technology framework developed by Lipsey and Carlaw, which identifies technologies that reorganise entire economic systems rather than individual industries.[2] GPTs diffuse across multiple sectors, improve continuously over time, and generate waves of complementary innovation. Steam power, electricity, and the internet all fit. Generative AI fits too, probably more dramatically than any of its predecessors.

The pace of improvement alone is worth sitting with. GPT-4 outperformed 90% of human test-takers on the U.S. bar exam; GPT-3.5, one year earlier, scored at the 10th percentile.[3] AI's context window grew from the equivalent of 7.5 pages of text in 2020 to nearly 300 pages by late 2023.[3] These are not incremental gains. They represent a technology improving faster than the institutions around it can track.

The GPT framework, though, carries an implicit optimism that deserves pushback. It emphasises long-run gains while treating transition costs as a temporary inconvenience. Acemoglu and Restrepo's distinction between automation and augmentation technologies is more honest about the distributional stakes.[4] Automation displaces workers from tasks they currently perform; augmentation expands what workers can accomplish, raising demand for their labour. Most historical GPTs have done both, and the balance between the two has determined whether technological transitions reduced or deepened inequality.

My reading of the current evidence: generative AI is more automating than augmenting at the task level, particularly for routine cognitive work. Junior analysts, paralegals, data processors, customer service representatives, these workers are being displaced, not enhanced. The augmentation gains are real, but they tend to go to workers who already have advanced skills. This is not a neutral outcome. It is a mechanism for compounding existing advantages.

3.2 Productivity projections: why the disagreement matters more than the number

Productivity estimates for AI vary so much that picking any single figure without context is actively misleading. McKinsey Global Institute projects generative AI could add 0.1 to 0.6 percentage points annually to global productivity growth between 2023 and 2040.[1] Goldman Sachs projects a 15% improvement in labour productivity in advanced economies on full adoption.[6] Daron Acemoglu at

MIT, whose scepticism is well-founded, projects a 0.7% total productivity increase in the U.S. over the next decade.[5] These estimates do not overlap much.

The disagreement is not noise. It reflects genuinely different assumptions about what productivity means, what time horizon matters, and who is being measured. McKinsey and Goldman Sachs are estimating potential gains under favourable conditions. Acemoglu is estimating likely gains given realistic adoption constraints, skill gaps, and the organisational friction that every large-scale technology transition produces. Both are right on their own terms. The practical implication: productivity gains from AI will be real but uneven, concentrated in sectors with strong digital infrastructure and the management capacity to absorb new tools effectively.

Healthcare studies project productivity improvements of 30–40% through generative AI, driven by diagnostics, drug discovery, and administrative automation.[7] Financial services could capture $200–340 billion in annual value.[1] These are sectors with sophisticated institutions and existing digital infrastructure. The contrast with India's vast informal economy captures the distributional problem precisely: the sectors that will capture the most from AI are not where most workers in developing economies actually work. Figure 1 illustrates the variation across sectors. It is important to note that these estimates are not directly comparable, they use different definitions of 'productivity improvement', different time horizons, and different methodological bases.

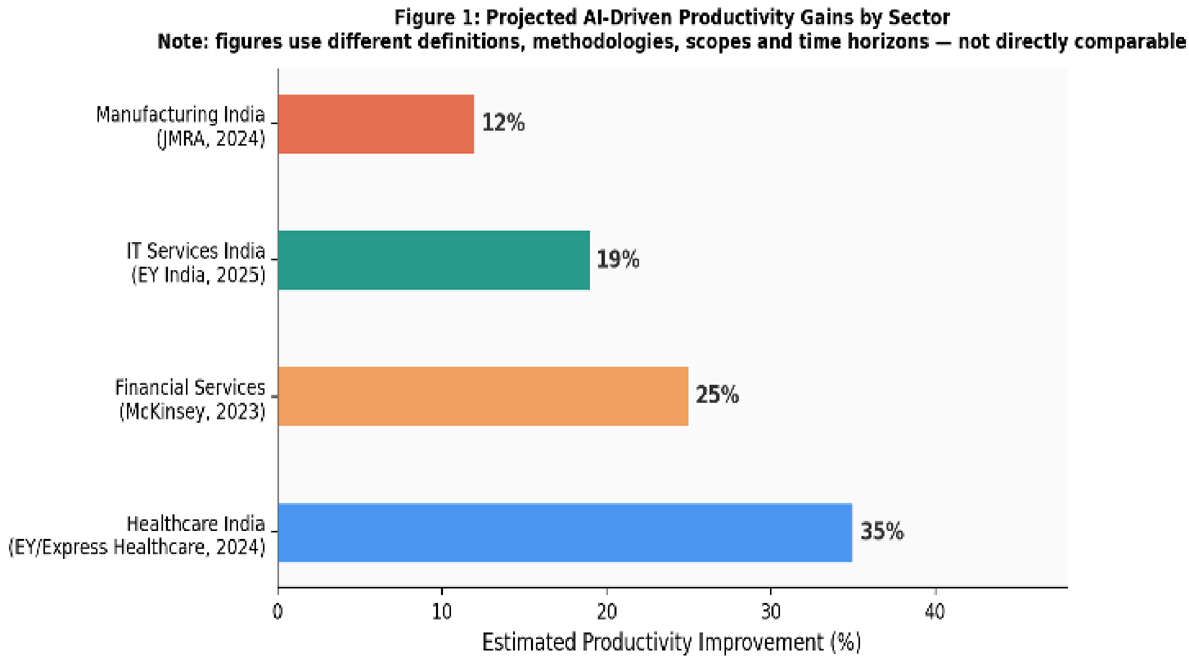

Figure 1: Projected AI-driven productivity gains by sector. Note: figures are not directly comparable, they reflect different definitions, methodologies, scopes, and time horizons. Sources: Healthcare: EY/Express Healthcare, 2024; Financial Services: McKinsey, 2023; IT Services India: EY India, 2025; Manufacturing India: JMRA, 2024.

Then there is the J-curve problem. Historical patterns from steam power, electrification, and the internet all show that GPT adoption initially suppresses measured productivity while firms absorb transition costs before long-run gains materialise. The 2025 METR study on AI's effect on software developers makes this concrete: experienced developers took 19% longer to complete certain tasks with AI assistance than without it.[13] The easy response is to call this a measurement problem. I think it is more informative than that. It suggests that AI's productivity effects depend critically on how deeply tools are integrated into actual workflows, and that in the absence of that integration, AI can generate friction rather than efficiency. For developing economies where the infrastructure for deep AI integration does not yet exist, this matters a lot.

3.3 Labour markets: what the WEF projections conceal

WEF labour market projections have shifted between reporting periods in ways that are analytically instructive, and that shift is itself part of the argument. The WEF Future of Jobs Report 2020, based on a survey of 300 companies, projected that by 2025, 85 million jobs would be displaced while 97 million new roles would emerge, a net gain of approximately 12 million positions. The WEF Future of Jobs Report 2023, based on a larger sample of 803 companies and projecting to 2027, revised these figures substantially: 83 million roles are expected to be eliminated against 69 million created, a net loss of approximately 14 million jobs. Both figures are employer survey expectations, not model estimates. Their divergence reflects both a revised sample and the observed acceleration of AI adoption between reporting periods. Figure 2 shows both projections side by side.

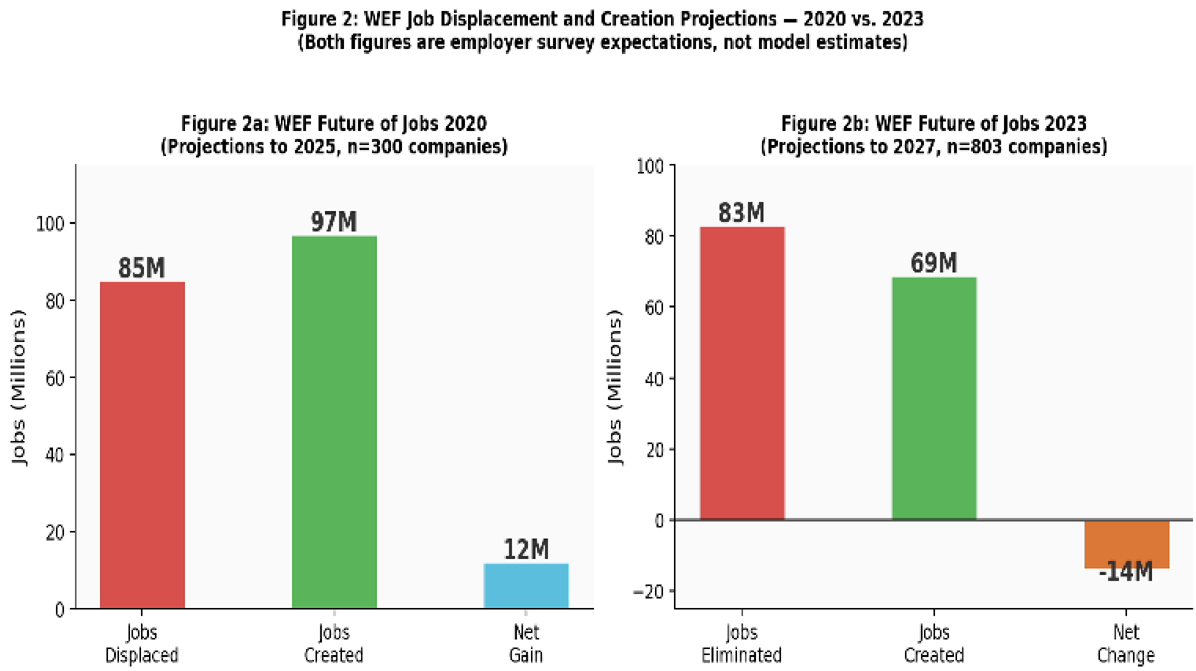

Figure 2: WEF job displacement and creation projections from the 2020 and 2023 Future of Jobs Reports. The 2020 report (n=300 companies) projected a net gain of 12 million positions by 2025; the 2023 report (n=803 companies) projects a net loss of 14 million by 2027. Both are employer survey expectations, not model estimates. Sources: WEF Future of Jobs 2020[15]; WEF Future of Jobs 2023.[10]

Whether the net figure is positive or negative is less analytically useful than examining what is underneath it. The first problem is the credential barrier. New AI roles, machine learning specialists, AI ethics officers, human-AI interaction designers, require advanced qualifications that the workers facing the highest displacement risk do not hold. The WEF 2023 report identifies data entry clerks, cashiers, and bank tellers as the fastest-declining occupational categories. PwC's 2025 Global AI Jobs Barometer reports that skill requirements in AI-exposed occupations are evolving 66% faster than in the least-exposed roles.[16] Workers with AI proficiency already earn a 25% wage premium.[10] These figures describe opportunity, but they also describe a moving target. By the time a displaced data entry worker in her mid-40s completes a retraining programme built for today's AI landscape, the landscape will have shifted again.

The second problem is gender. In the U.S., 58.87 million women hold positions with high automation exposure, compared to 48.62 million men.[11] This reflects the concentration of women in administrative, clerical, and customer service work, exactly the roles at highest risk. The new roles AI creates skew heavily toward STEM and technical fields where women remain underrepresented. The net employment figure, whatever its sign, will not be distributed proportionally.

The third problem is generational. Workers in their 40s and 50s who built careers around stable, routine cognitive work face the steepest retraining demands and the shortest remaining career horizons over which to recover the investment. Younger workers have more time, but only if educational institutions keep pace with what employers actually need, which in most countries they currently do not. The lag between a new AI application becoming commercially viable and its incorporation into mainstream curricula is measured in years, not months. That is not a minor inefficiency. It means each entering cohort starts their working life with preparation that is already out of date.

3.4 Investment banking: compression without elimination

I should be upfront about what this section is and is not. It is an illustrative sectoral example, not a formal case study with primary data collection. The claims that follow draw on publicly available reporting on AI deployments and conceptual reasoning about the structure of IB work. Primary data on analyst hiring trends since AI deployment would be needed to make stronger empirical claims; what this section offers instead is a plausible account of the mechanism, grounded in publicly documented technology adoption.

The traditional IB structure is a pyramid. Junior analysts do the foundational work: financial modelling, DCF valuations, comparable company analysis, pitch book assembly, data aggregation. These tasks are demanding, repetitive, and consume enormous time, a single pitch book can take 80 analyst-hours. They are also how the industry transmits knowledge: you learn how deals work by doing the groundwork of a hundred of them, and by making mistakes that senior bankers correct.

Goldman Sachs, JPMorgan, and Morgan Stanley have deployed generative AI tools that directly target this work. JPMorgan's LLM Suite is available to over 60,000 employees.[6] These are not pilots. They are production deployments at scale, and the economic logic is obvious: if AI does in two hours what a junior analyst does in forty, the business case for hiring fewer analysts is straightforward arithmetic.

The counter-argument that gets made here is the Bloomberg terminal analogy: when terminals arrived, they did not eliminate financial analysts, they shifted the work from data retrieval to interpretation. Spreadsheets did not eliminate accountants; they changed the job. AI will do the same, this argument goes, pushing junior analyst work further up the value chain. I think this is partially right and importantly wrong. Bloomberg terminals and spreadsheets automated discrete tasks. Generative AI can automate chains of reasoning: drafting an analysis, checking assumptions against comparable cases, stress-testing scenarios, flagging inconsistencies, formatting outputs. It automates the process, not just individual steps within it.

What this means in practice: not the elimination of junior roles, but their compression. Banks will hire fewer analysts per deal, each expected to operate at a higher level of abstraction from day one. Goldman Sachs Research estimates AI could raise financial sector labour productivity by up to 15% in the medium term.[6] What that figure does not capture is the cost to the next generation of analysts of developing professional judgement without the apprenticeship structure that historically produced it.

3.5 Inequality: why the aggregate obscures the distributional reality

The IMF identifies two mechanisms through which AI deepens inequality.[3] The first is the skill premium effect: AI augments high-skill workers more than low-skill workers, because high-skill workers are better positioned to use AI as a complement to what they already know. A consultant with twenty years of domain expertise who uses AI to synthesise research faster becomes dramatically more productive. A data entry clerk using AI to check their own work may be automating themselves out of a job. The same technology, applied to workers at different points of the skill distribution, produces opposite results.

The second mechanism is the capital income effect. AI raises the productivity of capital, the computing infrastructure, the models, the platforms, more than it raises the productivity of labour in aggregate. Productivity gains generated by AI flow disproportionately to the owners of AI capital: technology companies, investors, and the firms with the balance sheets to make large upfront AI investments. Workers capture some of this through wages, but a smaller share than they would if the productivity gains came from labour-complementary rather than capital-complementary technology.

UNCTAD data shows that 40% of global AI R&D investment is concentrated in just 100 companies, almost all headquartered in the U.S. or China.[1] Advanced economies could see more than double the GDP growth from AI compared to low-income countries.[3] These reflect the same underlying dynamic: AI's gains cluster around existing concentrations of capital, infrastructure, and skill.

For India, the numbers sit uncomfortably. The $1.25 billion IndiaAI Mission is a real commitment. But India's labour market is 90% informal employment, with a large rural workforce in roles with high automation exposure and educational institutions not yet producing AI-ready graduates at scale. The author's direct observation, conducting financial literacy workshops for over 300 villagers in Lakhimpur Kheri, confirmed that basic digital literacy, a prerequisite for any AI-adjacent skill, remains absent for large segments of the rural population. Without deliberate redistribution, not just reskilling, but fiscal policy, wage floors, and social protection for displaced workers, the net effect of AI on Indian inequality is probably negative even if the aggregate GDP effect is positive.

There is a counterargument worth taking seriously. Some evidence suggests AI may narrow certain skill-based gaps by raising the productivity floor for less experienced workers more than it raises the ceiling for experienced ones. The METR developer study points in this direction indirectly: experienced developers found AI assistance disruptive to established workflows, while less experienced workers may find it more useful as a scaffold. If this dynamic generalises, AI could be modestly equalising within skill groups even while widening gaps between them. But the burden of proof for an optimistic distributional story should be high, given the historical track record of technological transitions.

4. POLICY IMPLICATIONS

The policy literature on AI and labour markets has a familiar answer: reskilling, digital infrastructure, regulation, international cooperation. These are not wrong recommendations. They are just insufficient ones, because they go after symptoms rather than the mechanisms that produce the distributional problems in the first place. Treating reskilling as the primary response to AI displacement is a bit like handing someone an umbrella during a flood.

4.1 Reskilling is necessary but not sufficient

The standard reskilling case assumes that once displaced workers are retrained, they can walk into the net employment gain. The credential barrier shows why that assumption does not hold: the new AI roles require advanced qualifications that short-cycle certification courses cannot provide quickly enough. The deeper problem is timing. The lag between an AI application becoming commercially viable and its incorporation into mainstream curricula is measured in years. By the time a programme is designed, funded, and deployed, it is training people for a job market that has already moved.

Finland's national AI education initiative trained over 1% of its population in AI fundamentals. That is genuinely impressive as a model of population-scale AI literacy, and worth studying. What it does not fix is the credential gap for the specific technical roles where AI job creation actually concentrates. Reskilling policy needs to be honest about what it can achieve. Competency-based certifications and short-cycle programmes are useful entry points to adjacent work, not substitutes for structural reform. For India, the immediate priority is getting AI literacy into secondary and vocational education now, before the current student cohort enters a labour market that will expect it.

4.2 Infrastructure investment must reach the informal economy

A 10% increase in broadband penetration raises GDP growth in developing nations by 1.4%.[14] That is a reasonable justification for infrastructure spending. But infrastructure that reaches India's major cities and stops at their outskirts does not solve the distributional problem, it intensifies it, giving AI's productivity benefits to the already-connected while leaving everyone else further behind. The IndiaAI Mission's computing infrastructure commitment is a real achievement. Its value depends entirely on whether connectivity actually reaches the roughly 600 million Indians who currently lack reliable access to it.

4.3 Regulation should target the transition, not just the endpoint

Most current AI regulation focuses on risks at the moment of deployment: algorithmic bias, transparency, high-stakes decision-making. These are legitimate concerns about the endpoint of AI adoption, not the transition period when distributional harm is most acute. Workers who lose roles to automation need protections now, not after a regulatory framework finishes being designed. Advance notice requirements before AI-driven workforce changes, portability of retraining benefits, anti-discrimination protections in AI-assisted hiring, these should be treated as urgent, not aspirational. India's forthcoming AI regulatory framework has a chance to build transition protections in from the start. That window will not stay open indefinitely.

4.4 The capital-labour problem requires fiscal tools

Reskilling and regulation can limit the labour market disruption from AI. What they cannot do is address the capital income effect, the mechanism through which AI raises returns to asset owners faster than returns to workers. That requires fiscal policy: progressive capital gains taxation on AI-generated returns, automation levies that fund social protection for displaced workers, public investment mechanisms that give citizens a stake in AI's productivity gains. These are politically difficult proposals. They are also, based on the distributional evidence, the only tools that directly address the mechanism producing inequality rather than just managing its downstream effects.

5. LIMITATIONS

Several limitations are worth making explicit. First, the synthesis relies substantially on institutional projections and employer survey expectations rather than measured outcomes; projections carry inherent uncertainty and may be revised as AI adoption evolves. Second, cross-source comparability is a significant constraint: the productivity estimates in Figure 1 use different definitions of 'productivity improvement', different baselines, different time horizons, and different geographic scopes, they are presented as illustrative orders of magnitude rather than a unified comparison. Third, India-specific observations are partly grounded in the author's direct experience, which provides authentic contextual grounding but is geographically limited and anecdotal. Fourth, the investment banking section makes conceptual arguments about pipeline compression that are not supported by primary sector data on analyst hiring trends since AI deployment, this is explicitly noted in that section. Fifth, the pace of AI development means estimates published in 2023–2024 may already understate current adoption levels.

6. CONCLUSION

I started with a simple question, would AI destroy more jobs than it creates?, and ended up somewhere more complicated. The answer, based on the best available evidence, is probably not, at least not in aggregate. But that answer, on its own, is considerably less reassuring than the people who cite it tend to suggest.

The net employment figures are real and structurally inaccessible to the workers who most need them. The productivity gains are real and geographically concentrated in economies that are already ahead. The capital income effects are real and systematically favour asset owners over wage earners. AI is not producing a catastrophe for the global labour market in aggregate. It is producing three compounding structural problems, credential exclusion, geographic concentration, and capital-labour asymmetry, that will worsen inequality within and between countries unless policy targets the mechanisms directly.

Investment banking illustrates this in miniature. The industry is not disappearing. The entry-level pipeline is contracting. The skills expected on day one are rising. The apprenticeship model through which junior analysts historically developed their judgement is being undercut by the same tools that make senior analysts more productive. That is a small example of a pattern that runs across sectors: AI raises the floor of competence required to be economically viable while removing the scaffolding through which people have historically built toward it.

For India, the stakes are high in a specific way. A country with 90% informal employment, a large and young workforce, and educational institutions not yet producing AI-ready graduates at scale is precisely where the gap between AI's aggregate promise and its distributional reality is widest. The IndiaAI Mission is a genuine reason for optimism. It is not, by itself, enough.

Technology is not the problem, and it is not the solution. The question is whether governments, educational systems, and regulatory bodies can move fast enough to match AI's pace of adoption. The historical record on that question is not encouraging. Whether this time is different is not a technological question. It is a political one.

ACKNOWLEDGEMENTS

The author thanks Mrs. Swati Abhishek Mishra (MPhil & PhD, University of Cambridge and MIT; Former Professor, IIM Lucknow and IIM Ahmedabad) for her mentorship and detailed feedback throughout this research. Her grounding in development economics shaped the analytical framework in ways that are visible throughout the paper.

REFERENCES

[1] McKinsey Global Institute. (2023). Economic potential of generative AI. McKinsey & Company. https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier

[2] OECD. (2024). Is generative AI a general purpose technology? OECD Publications. https://www.oecd.org/en/publications/is-generative-ai-a-general-purpose-technology_704e2d12-en.html

[3] International Monetary Fund. (2024). AI will transform the global economy: let's make sure it benefits humanity. IMF Blog. https://www.imf.org/en/Blogs/Articles/2024/01/14/ai-will-transform-the-global-economy-lets-make-sure-it-benefits-humanity

[4] Acemoglu, D., & Restrepo, P. (2025). Displacement or complementarity? The labour market impact of AI. Harvard Business School Working Paper 25-039.

[5] Acemoglu, D. (2024). A new look at the economics of AI. MIT Sloan Management Review. https://mitsloan.mit.edu/ideas-made-to-matter/a-new-look-economics-ai

[6] Goldman Sachs Research. (2024). AI is showing very positive signs of eventually boosting GDP. Goldman Sachs Insights.

[7] Express Healthcare / EY. (2024). India's healthcare sectors eye 30–40% productivity gains with GenAI adoption. https://www.expresshealthcare.in/news/indias-healthcare-sectors-eye-30-40-percent-productivity-gains-with-genai-adoption-ey-report/448107/

[8] Journal of Management Research and Analysis. (2024). Impact of artificial intelligence on the Indian economy. JMRA.

[9] Spherical Insights. (2024). Top 10 AI spending countries 2025. https://www.sphericalinsights.com/blogs/top-10-artificial-intelligence-spending-countries-in-2025-statistics-and-facts-analysis-2024-to-2035

[10] World Economic Forum. (2023). The future of jobs report 2023. WEF. https://www.weforum.org/publications/the-future-of-jobs-report-2023/

[11] CNBC / UN. (2025). AI could affect 40% of jobs, widen inequality between nations.

[12] EY India. (2025). How much productivity can GenAI unlock in India? EY.

[13] METR. (2025). Measuring the impact of early-2025 AI on experienced OS developers. https://metr.org/blog/2025-07-10-early-2025-ai-experienced-os-dev-study/

[14] World Economic Forum. (2024). Digital inclusion and broadband access report. WEF.

[15] World Economic Forum. (2020). The future of jobs report 2020. WEF. https://www.weforum.org/publications/the-future-of-jobs-report-2020/

[16] PwC. (2025). 2025 global AI jobs barometer. PricewaterhouseCoopers.