

The Macroeconomic Impact of Generative AI: Evidence and Projections

Tilka Devi

Assistant Professor, Department of Economics ML&JNL College.

Email: meenusingh220790@gmail.com

Article: Received: 01/06/2026, Returned: 15/06/2026, Accepted: 25/06/2026, Published:30/06/2026.

D.O.I. <https://doi.org/10.5281/zenodo.22827113>



@ 2026 The Author(s). This is an Open Access article/ Journal distributed under the terms of the Creative Commons Attribution 4.0 International which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are properly credited. (<https://creativecommons.org/licenses/by/4.0/>)

Abstract: Generative artificial intelligence (gen-AI) has diffused faster than any previous general-purpose technology, yet its effect on aggregate output remains contested. Forecasts of its ten-year contribution to GDP range from below 1% to around 7%, while early cross-country estimates suggest that measured macroeconomic effects have so far been close to zero. This paper addresses the resulting "micro-macro gap": the disconnect between large, well-identified productivity gains at the task level (commonly 14–56%) and negligible movements in national productivity statistics. We synthesise experimental, firm-level and macroeconomic evidence published between 2023 and 2026, and develop a transparent task-based growth-accounting framework in the tradition of Hulten (1978) and Acemoglu (2024). The framework decomposes aggregate gains into task exposure, adoption, and cost savings, and adds an institutional-readiness parameter that allows projections to differ across advanced, emerging and low-income economies. Calibrated to published parameters, the model reproduces the conservative benchmark of roughly 0.66% cumulative TFP growth for advanced economies, and yields cumulative ten-year GDP gains of 0.99–4.86% for advanced economies, 0.50–2.43% for emerging markets and 0.22–1.05% for low-income countries across three scenarios. The results imply that gen-AI is more likely to widen than to narrow cross-country income gaps unless adoption capacity is deliberately built. We argue that the micro-macro gap is explained mainly by slow diffusion, the need for complementary organisational investment, and mismeasurement of quality improvements, and we propose a cross-country panel design to test these channels as data accumulate.

Keywords: Generative artificial intelligence; productivity; economic growth; labour markets; general-purpose technology; emerging economies; growth accounting.

1. Introduction

The public release of large language models in late 2022 marked the beginning of a new phase in the economics of automation. Earlier waves of automation primarily substituted for routine manual and routine cognitive tasks. Generative AI, by contrast, performs non-routine cognitive work such as writing, summarising, coding, translating and analysing, which has historically been the domain of educated, high-wage workers. Adoption has been unusually rapid: survey evidence indicates that by late 2024 nearly 40% of U.S. adults aged 18–64 had used generative AI, and about 23% of employed respondents had used it for work in the previous week (Bick, Blandin and Deming, 2026).

This rapid diffusion has produced an equally rapid proliferation of forecasts, and these forecasts disagree sharply. Goldman Sachs (Briggs and Kodnani, 2023) projected that generative AI could raise global GDP by around 7% over a decade and lift annual labour productivity growth by roughly 1.5 percentage points. Acemoglu (2024), using a task-based model, concluded that the total factor productivity (TFP)

gain over ten years is unlikely to exceed about 0.66%, implying modest GDP effects and no reduction in labour income inequality. The Penn Wharton

Budget Model (2025) takes an intermediate position, estimating that AI will raise productivity and GDP by 1.5% by 2035, nearly 3% by 2055 and 3.7% by 2075, with the boost to annual growth peaking in the early 2030s before fading.

Meanwhile, the observed data remain ambiguous. Controlled experiments document substantial productivity gains for individual workers, but a 2026 study covering 67 countries between 2022 and 2025 estimates that gen-AI has so far added only about 0.008% to GDP in the average country (*Structural Change and Economic Dynamics*, 2026). U.S. nonfarm business productivity grew strongly in 2025, rising at a 4.9% annualised rate in the third quarter, but whether AI contributed to this remains unclear (Man Group, 2026). This situation echoes Robert Solow's 1987 observation that the computer age was visible everywhere except in the productivity statistics.

This paper makes three contributions. First, it provides a structured synthesis of the rapidly growing evidence base, organised by level of analysis (task, firm, economy and global). Second, it develops a transparent and reproducible projection framework that separates task exposure, adoption, cost savings and institutional readiness, allowing readers to see precisely which assumptions drive divergent forecasts. Third, it extends the analysis beyond the United States, where most existing evidence originates, to emerging and low-income economies, and it proposes an empirical design to test the channels behind the micro–macro gap.

The paper is guided by four research questions:

- RQ1:** What does the current empirical evidence show about the productivity effects of generative AI at the task, firm and macroeconomic levels?
- RQ2:** Why have large micro-level gains not yet appeared in aggregate productivity statistics?
- RQ3:** What range of GDP effects is plausible over a ten-year horizon under transparent assumptions?
- RQ4:** How are these effects likely to differ between advanced, emerging and low-income economies?

The remainder of the paper is structured as follows. Section 2 reviews the literature and identifies research gaps. Section 3 describes the methodology. Section 4 presents key facts and evidence. Section 5 reports the projection results. Section 6 discusses implications and policy, and Section 7 concludes.

2. Literature Review

2.1 Generative AI as a General-Purpose Technology

Economic historians describe general-purpose technologies (GPTs) as innovations that are pervasive, improve over time and spawn complementary innovations. Electricity and information technology are canonical examples. David (1990) showed that electrification took several decades to raise manufacturing productivity because factories had to be physically redesigned around electric motors. Brynjolfsson and Hitt (2000) found that firms combining IT investment with organisational change achieved much larger productivity gains than those investing in technology alone. Building on this, Brynjolfsson, Rock and Syverson (2021) proposed the "productivity J-curve": because firms initially divert resources into intangible investments (training, process redesign, data infrastructure) that are not captured as output, measured productivity may first fall and only later rise. Generative AI displays GPT characteristics, suggesting that a lag between adoption and measured gains is to be expected rather than treated as evidence of failure.

2.2 Task-Based Models of Automation

The dominant theoretical framework is the task-based approach developed by Acemoglu and Restrepo (2018, 2019). Production is viewed as a set of tasks that can be performed by labour or capital. Technology affects the economy through a displacement effect (machines take over tasks), a productivity effect (lower costs raise demand for remaining labour) and a reinstatement effect (new tasks are created for humans). Autor et al. (2024) show empirically that the majority of current U.S. employment is in job titles that did not exist in 1940, highlighting the importance of new task creation.

The net effect of AI on wages and employment therefore depends on the balance between displacement and reinstatement, which cannot be determined by exposure measures alone.

A related strand measures occupational exposure. Felten, Raj and Seamans (2021) linked AI capabilities to occupational abilities. Eloundou et al. (2023) estimated that around 80% of the U.S. workforce could have at least 10% of their work tasks affected by large language models, and around 19% could see at least half of their tasks affected. The IMF (Cazzaniga et al., 2024) estimated that roughly 40% of global employment is exposed to AI, rising to about 60% in advanced economies, compared with 40% in emerging markets and 26% in low-income countries. However, as several authors emphasise, exposure is not adoption, adoption is not productivity, and productivity gains are not necessarily captured as higher output or wages.

2.3 Micro-Level Experimental and Firm Evidence

A growing body of well-identified studies documents large task-level gains. In a customer-support setting, access to an AI assistant raised issues resolved per hour by about 14–15% on average, with far larger gains for novice workers (Brynjolfsson, Li and Raymond, 2025). Noy and Zhang (2023) found that ChatGPT substantially reduced the time needed for professional writing tasks while improving quality. Peng et al. (2023) reported that software developers using GitHub Copilot completed a programming task about 56% faster. Dell'Acqua et al. (2023) introduced the concept of the "jagged technological frontier": consultants using GPT-4 performed better on tasks within AI's capability frontier, but worse on tasks just outside it. A recurring finding is that gains are largest for less experienced or lower-performing workers, suggesting AI may compress within-occupation skill differences.

Firm-level evidence is more mixed. A survey of nearly 750 corporate executives by the Federal Reserve Bank of Atlanta (2026) found that more than half of firms had already invested in AI and that labour productivity gains were positive and expected to strengthen in 2026, with the largest effects in high-skill services and finance. Importantly, perceived productivity gains exceeded measured gains, consistent with a delay in revenue realisation. Other work documents an "AI productivity paradox" in which substantial AI investment does not consistently translate into improved firm performance, with misaligned organisational structures, weak data infrastructure and talent shortages cited as explanations (McElheran et al., 2025).

2.4 Macroeconomic Estimates

Macroeconomic estimates are polarised. Optimistic projections (Briggs and Kodnani, 2023; McKinsey Global Institute, 2023) extrapolate from exposure and experimental gains and assume broad adoption. McKinsey estimated that generative AI could add USD 2.6–4.4 trillion in value annually across the use cases it studied. Conservative projections (Acemoglu, 2024) apply Hulten's theorem, which states that the aggregate TFP effect of a task-level improvement equals the GDP share of affected tasks multiplied by the average cost saving. Because the share of tasks that can be profitably automated within a decade is modest, the aggregate effect is small. Intermediate estimates from the OECD (Filippucci, Gal and Schief, 2024) and Penn Wharton (2025) sit between these poles, and stress that results depend heavily on assumptions about adoption speed and sectoral reallocation.

2.5 Distributional and Cross-Country Effects

Distributional research suggests that the gains will be uneven. Acemoglu (2024) predicts a widening gap between capital and labour income. At the global level, the IMF (Cazzaniga et al., 2024) finds that advanced economies face greater exposure but are also better placed to benefit through complementarity. An IMF working paper using a multi-sector dynamic general-equilibrium model (IMF WP/25/76, 2025) concludes that AI will exacerbate cross-country income inequality, with growth effects in advanced economies potentially more than double those in low-income countries, and suggests that AI productivity gains concentrated in non-tradable sectors could weaken the role of exchange-rate adjustment. By contrast, a 2026 World Bank report emphasises that jobs in low- and middle-income

countries are less at risk of automation (4.5% versus 14.2% in high-income countries), while the share of jobs that could be meaningfully augmented is similar (16.2% versus 18.7%). This implies that developing economies may face less displacement but also limited capacity to capture gains because of deficits in electricity, connectivity, data and skills.

2.6 Research Gaps

The review identifies five gaps that motivate this paper:

- Gap 1 – The micro–macro disconnect.** Large task-level gains coexist with negligible measured aggregate effects, and the mechanisms behind this gap have not been systematically decomposed.
- Gap 2 – Scarcity of macro data.** Because few years of post-adoption data exist, most macro studies rely on analogies with earlier technologies, case studies or task-based simulations rather than observed outcomes.
- Gap 3 – Geographic concentration.** Experimental and firm-level evidence comes overwhelmingly from the United States and Europe, leaving emerging and low-income economies underexplored.
- Gap 4 – Opaque forecast assumptions.** Published projections differ by an order of magnitude, but the assumptions responsible for these differences are rarely presented in a comparable, decomposed form.
- Gap 5 – Role of complementary capabilities.** Institutional readiness, digital infrastructure and organisational capital are widely cited as determinants of AI returns but are seldom incorporated explicitly into macro projections.

3. Methodology

3.1 Research Design

The study adopts a mixed design with three components: (i) a structured narrative synthesis of empirical studies published between 2023 and 2026; (ii) a calibrated task-based growth-accounting projection; and (iii) a proposed econometric specification for testing the model's channels using cross-country panel data. Components (i) and (ii) produce the results in this paper; component (iii) is presented as a framework for future empirical work, since the post-2022 time series is not yet long enough for robust inference.

3.2 Evidence Synthesis

Sources were identified through Google Scholar, SSRN, NBER, arXiv and the publication databases of the IMF, OECD, World Bank and regional Federal Reserve banks, using combinations of the terms "generative AI", "large language models", "productivity", "GDP", "labour market" and "economic growth". Inclusion criteria were: (a) publication between 2023 and 2026, or foundational theoretical work; (b) a quantitative estimate of productivity, output or employment effects; and (c) a transparent methodology. Studies were classified by level of analysis (task, firm, macro, global) and by method (randomised experiment, quasi-experiment, survey, simulation or general-equilibrium model).

3.3 Projection Model

Following Hulten (1978) and Acemoglu (2024), the aggregate TFP effect of gen-AI over horizon h in country group g is approximated as:

$$\Delta TFP(g) = E(g) \times A(s) \times R(g) \times S(s) \quad (1)$$

where $E(g)$ is the share of economic tasks (weighted by GDP) exposed to gen-AI; $A(s)$ is the fraction of exposed tasks that are adopted and profitably transformed within ten years in scenario s ; $R(g)$ is an institutional-readiness factor between 0 and 1 that scales adoption according to infrastructure, skills and access to technology; and $S(s)$ is the average cost saving on transformed tasks. Because TFP gains raise the return to capital and stimulate investment, the GDP effect exceeds the TFP effect. We approximate this with a capital-deepening multiplier:

$$\Delta GDP(g) = \Delta TFP(g) / (1 - \alpha), \quad \alpha = 1/3 \quad (2)$$

where α is the capital share of income, which yields a multiplier of 1.5. This value is consistent with Acemoglu's (2024) estimate that a 0.66% TFP gain translates into a GDP increase of roughly 0.93–1.16%.

3.4 Calibration

Parameters are drawn from published sources, as summarised in Table 1. The advanced-economy task exposure of 20% follows Acemoglu (2024). Exposure for emerging and low-income economies is scaled by the ratio of IMF employment-exposure estimates (40/60 and 26/60 respectively). Readiness factors are set at 1.0 for advanced economies, 0.75 for emerging markets and 0.5 for low-income countries, informed by IMF AI Preparedness Index differences and World Bank evidence on infrastructure deficits. Three adoption and cost-saving scenarios capture uncertainty: a Conservative scenario that replicates Acemoglu (2024), a Moderate scenario, and an Accelerated scenario consistent with the upper range of labour-cost savings reported in the literature.

Table 1. Model parameters and sources

Parameter	Advanced	Emerging	Low-income	Source / rationale
Task exposure E	20.0%	13.3%	8.7%	Acemoglu (2024); scaled by IMF exposure ratios (Cazzaniga et al., 2024)
Readiness R	1.00	0.75	0.50	IMF AI Preparedness Index; World Bank (2026)
Adoption A: Conservative / Moderate / Accelerated	23% / 40% / 60%	same	same	Acemoglu (2024) for lower bound; author scenarios
Cost saving S: Conservative / Moderate / Accelerated	14.4% / 20% / 27%	same	same	Acemoglu (2024); experimental range
Capital share α	1/3	1/3	1/3	Standard growth-accounting value

*Note: Exposure for emerging and low-income groups = 20% $\times$ (IMF group exposure $\div$ 60%).
Scenario parameters are illustrative assumptions intended to span the range of published estimates.*

3.5 Proposed Empirical Specification

To test the channels in Equation (1) as data accumulate, we propose the following country–year panel regression:

$$\Delta \ln LP(i, t) = \beta_1 \cdot AI(i, t) + \beta_2 \cdot AI(i, t) \times READY(i) + \beta_3 \cdot AI(i, t) \times EXP(i) + \gamma \cdot X(i, t) + \mu(i) + \tau(t) + \varepsilon(i, t) \quad (3)$$

where $\Delta \ln LP$ is labour productivity growth in country i and year t ; AI is a gen-AI adoption proxy (for example, per-capita traffic to gen-AI platforms, enterprise adoption survey rates or AI-related job postings); $READY$ is a pre-2022 readiness index; EXP is occupational exposure; X contains controls (capital deepening, human capital, trade openness, business-cycle indicators); and μ and τ are country and year fixed effects. The coefficient β_2 tests whether readiness amplifies AI's productivity effect (Gap 5), while lagged AI terms can test for J-curve dynamics (Gap 1). To address the endogeneity of adoption, pre-existing internet penetration or English-language proficiency interacted with the global release of major models could serve as instruments.

3.6 Limitations

The projection model is a partial-equilibrium approximation. It does not capture new task creation, product quality improvements, sectoral reallocation, or cross-border spillovers, all of which could raise or lower actual effects. Readiness factors and scenario parameters are stylised and should be interpreted as sensitivity bounds rather than point forecasts. The evidence synthesis is narrative rather than a formal meta-analysis, and is subject to publication bias towards positive experimental findings.

4. Key Facts and Evidence

4.1 Task-Level Productivity Evidence

Table 2 summarises the most widely cited micro-level studies. Across settings, gen-AI consistently raises speed and often quality on tasks within its capability frontier.

Table 2. Selected micro-level evidence on generative AI productivity

Study	Setting	Method	Main finding
Brynjolfsson, Li & Raymond (2025)	Customer support agents	Staggered field rollout	~14–15% more issues resolved per hour; largest gains for novices
Noy & Zhang (2023)	Professional writing tasks	Online experiment	Large reduction in time taken; higher output quality
Peng et al. (2023)	Software developers (GitHub Copilot)	Controlled experiment	Task completed ~56% faster
Dell'Acqua et al. (2023)	Management consultants (GPT-4)	Field experiment	Better performance inside AI frontier; worse outside it ("jagged frontier")
Atlanta Fed (2026)	~750 corporate executives	Survey	Positive gains, strongest in high-skill services and finance; perceived > measured

4.2 Macroeconomic Projections

Table 3. Published macroeconomic estimates of generative AI

Source	Scope	Estimated effect
Briggs & Kodnani / Goldman Sachs (2023)	Global, 10 years	~7% rise in global GDP; ~1.5 pp faster annual productivity growth
McKinsey Global Institute (2023)	Global, use-case based	USD 2.6–4.4 trillion in annual value
Acemoglu (2024)	United States, 10 years	TFP $\leq$ ~0.66%; GDP ~0.93–1.16%
Penn Wharton Budget Model (2025)	United States	GDP +1.5% by 2035, ~3% by 2055, 3.7% by 2075
Structural Change & Econ. Dynamics (2026)	67 countries, 2022–2025	~0.008% GDP increase in average country so far
IMF WP/25/76 (2025)	Global DSGE model	Advanced-economy growth effect more than double that of low-income countries

4.3 Exposure and Labour Market Facts

- Around 40% of global employment is exposed to AI: about 60% in advanced economies, 40% in emerging markets and 26% in low-income countries (Cazzaniga et al., 2024).
- Approximately 80% of U.S. workers could have at least 10% of their tasks affected by large language models, and 19% could have at least half affected (Eloundou et al., 2023).
- Roughly two-thirds of current jobs in the U.S. and Europe are exposed to some degree of AI automation (Briggs and Kodnani, 2023).
- Only 4.5% of jobs in low- and middle-income countries are at risk of automation by gen-AI, compared with 14.2% in high-income countries, while 16.2% and 18.7% respectively could see meaningful productivity gains (World Bank, 2026).
- Nearly 40% of U.S. adults aged 18–64 had used gen-AI by late 2024 (Bick, Blandin and Deming, 2026).
- Executive surveys show little evidence so far of near-term aggregate employment declines attributable to AI (Federal Reserve Bank of Atlanta, 2026).

5. Results

5.1 Scenario Projections

Table 4 reports the projections from Equations (1) and (2). As a validation check, the Conservative scenario for advanced economies reproduces Acemoglu's (2024) benchmark TFP gain of 0.66%, with an implied GDP gain of 0.99%, inside his reported range of 0.93–1.16%.

Table 4. Projected ten-year effects of generative AI by country group and scenario

Country group	Scenario	Effective Adoption (A×R)	Cumulative TFP gain	Cumulative GDP gain	Avg. annual GDP growth boost
Advanced	Conservative	23.0%	0.66%	0.99%	0.10 pp
Advanced	Moderate	40.0%	1.60%	2.40%	0.24 pp
Advanced	Accelerated	60.0%	3.24%	4.86%	0.49 pp
Emerging	Conservative	17.3%	0.33%	0.50%	0.05 pp
Emerging	Moderate	30.0%	0.80%	1.20%	0.12 pp
Emerging	Accelerated	45.0%	1.62%	2.43%	0.24 pp
Low-income	Conservative	11.5%	0.14%	0.22%	0.02 pp
Low-income	Moderate	20.0%	0.35%	0.52%	0.05 pp
Low-income	Accelerated	30.0%	0.70%	1.05%	0.11 pp

Note: Author calculations from Equations (1)– (2) using Table 1 parameters. Values are illustrative projections, not observed outcomes




Figure 1. Projected cumulative ten-year GDP gains from generative AI by country group and scenario

5.2 Key Findings

Finding 1: Plausible effects are meaningful but not transformative within a decade. Even in the Accelerated scenario, advanced economies gain about 0.5 percentage points of annual GDP growth, well below the most optimistic industry forecasts. The Moderate scenario, around 2.4% cumulative GDP for advanced economies, is close to the intermediate estimates of Penn Wharton (2025).

Finding 2: Adoption, not exposure, is the binding constraint. Moving from the Conservative to the Accelerated scenario raises GDP gains roughly fivefold. Because exposure is held constant across scenarios, the difference comes entirely from adoption and cost savings. This explains why published forecasts diverge so widely: they embed very different assumptions about how quickly firms reorganise around the technology.

Finding 3: Gen-AI is likely to widen cross-country income gaps. In every scenario, gains in advanced economies are about twice those in emerging markets and more than four times those in low-income countries. This is consistent with IMF WP/25/76 (2025). Low-income countries face less displacement, but lower exposure combined with weaker readiness limits their gains.

Finding 4: Readiness is a powerful policy lever. Raising the readiness factor of emerging markets from 0.75 to 1.0 would increase their projected gains by one-third (for example, from 1.20% to 1.60% in the Moderate scenario), narrowing but not closing the gap with advanced economies.

5.3 Explaining the Micro–Macro Gap

Combining the synthesis and the model, we identify four channels that together explain why large task-level gains have not yet appeared in macro statistics:

- Diffusion lag:** In the model, the aggregate effect scales with the share of tasks actually transformed. With only a few years of adoption, realised A is far below its ten-year value, so a near-zero measured effect (e.g., 0.008% of GDP) is consistent with substantial eventual gains.
- Complementary investment (J-curve):** Firms must redesign workflows, retrain staff and build data systems. These intangible investments absorb resources before output rises (Brynjolfsson, Rock and Syverson, 2021).
- Measurement lag:** Executives report larger gains than revenue data show, as improvements in quality, speed and organisational capacity translate into revenue only gradually (Federal Reserve Bank of Atlanta, 2026).
- Uneven task coverage:** The jagged frontier means gains on some tasks are partly offset by errors or verification costs on others (Dell'Acqua et al., 2023), so time saved on individual tasks does not translate one-for-one into output.

6. Discussion and Policy Implications

The results suggest that debates framed as "boom versus hype" are misplaced. Generative AI is likely to raise output measurably, but the size and timing of the gains depend primarily on choices made by firms and governments about adoption and complementary investment, rather than on the technology's raw capabilities. This has several implications.

For advanced economies, the priority is converting pilots into scaled, reorganised production processes. Policies that support management practices, worker training and data infrastructure are likely to have higher returns than subsidies for AI purchases alone. Because gains accrue disproportionately to capital and high-skill workers, attention to labour income distribution, including tax and transfer design and reskilling programmes, will be important.

For emerging and low-income economies, the central challenge is readiness. Reliable electricity, broadband access, digital skills, and access to affordable models adapted to local languages determine whether these countries can capture augmentation gains. The World Bank's (2026) framework of adopting available tools, adapting them to local conditions and advancing gradually towards frontier development offers a practical sequence. There is also a risk that AI reduces the cost advantage of offshore services, threatening export-oriented service sectors such as business-process outsourcing.

For statistical agencies, the findings highlight a need to improve the measurement of intangible investment, service quality and AI usage. Without better data, the micro–macro gap will remain difficult to resolve, and policy will continue to rely on assumption-driven forecasts.

For future research, the specification in Equation (3) offers a path to test whether readiness amplifies AI's productivity effect and whether gains follow a J-curve. Country-level case studies of emerging economies, analysis of AI's effect on trade in services and exchange rates, and modelling of agentic AI systems that execute entire workflows rather than individual tasks are especially promising directions.

7. Conclusion

This paper examined the macroeconomic impact of generative AI by synthesising task-level, firm-level and macroeconomic evidence and developing a transparent projection framework. The evidence shows that gen-AI produces large and well-identified productivity gains for individual workers, especially less experienced ones, but that these gains have not yet materially affected national productivity statistics. Our calibrated model indicates that plausible ten-year GDP gains range from about 1% to 5% for advanced economies, 0.5% to 2.4% for emerging markets and 0.2% to 1% for low-income countries. The divergence among existing forecasts is driven mainly by assumptions about adoption rather than exposure. The micro–macro gap is best understood as a combination of diffusion lags, complementary investment needs, measurement problems and uneven task coverage, and not as evidence that the technology lacks economic value. Without deliberate investment in readiness, generative AI is more likely to widen than narrow global income disparities. The magnitude of AI's economic impact is therefore not predetermined; it will be shaped by how quickly and how inclusively economies adapt.

References

- Acemoglu, D. (2024). The simple macroeconomics of AI. *NBER Working Paper* 32487; published in *Economic Policy* (2025).
- Acemoglu, D., & Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6), 1488–1542.
- Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30.
- Autor, D., Chin, C., Salomons, A., & Seegmiller, B. (2024). New frontiers: The origins and content of new work, 1940–2018. *Quarterly Journal of Economics*, 139(3), 1399–1465.
- Bick, A., Blandin, A., & Deming, D. (2026). The rapid adoption of generative AI. Working paper.
- Briggs, J., & Kodnani, D. (2023). The potentially large effects of artificial intelligence on economic growth. *Goldman Sachs Global Economics Analyst*, 26 March.

- Brynjolfsson, E., & Hitt, L. M. (2000). Beyond computation: Information technology, organizational transformation and business performance. *Journal of Economic Perspectives*, 14(4), 23–48.
- Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *Quarterly Journal of Economics*, 140(2), 889–942.
- Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1), 333–372.
- Cazzaniga, M., et al. (2024). Gen-AI: Artificial intelligence and the future of work. *IMF Staff Discussion Note* SDN/2024/001.
- David, P. A. (1990). The dynamo and the computer: An historical perspective on the modern productivity paradox. *American Economic Review*, 80(2), 355–361.
- Dell'Acqua, F., et al. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Harvard Business School Working Paper* 24-013.
- Eloundou, T., Manning, S., Mohapatra, P., & Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models. *arXiv* 2303.10130 (published in *Science*, 2024).
- Federal Reserve Bank of Atlanta. (2026). Artificial intelligence, productivity, and the workforce: Evidence from corporate executives. *Working Paper* 2026-04.
- Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence. *Strategic Management Journal*, 42(12), 2195–2217.
- Filippucci, F., Gal, P., & Schief, M. (2024). Miracle or myth? Assessing the macroeconomic productivity gains from artificial intelligence. *OECD Artificial Intelligence Papers*, No. 29.
- Hulten, C. R. (1978). Growth accounting with intermediate inputs. *Review of Economic Studies*, 45(3), 511–518.
- International Monetary Fund. (2025). The global impact of AI: Mind the gap. *IMF Working Paper* WP/25/76.
- Man Group. (2026). The productivity paradox: When will AI deliver? Man Group Insights, February.
- McElheran, K., et al. (2025). The rise of industrial AI in America: Microfoundations of the productivity J-curve(s). Working paper.
- McKinsey Global Institute. (2023). *The economic potential of generative AI: The next productivity frontier*. McKinsey & Company.
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192.
- Peng, S., Kalliamvakou, E., Cihon, P., & Demirel, M. (2023). The impact of AI on developer productivity: Evidence from GitHub Copilot. *arXiv* 2302.06590.
- Penn Wharton Budget Model. (2025). The projected impact of generative AI on future productivity growth. University of Pennsylvania, 8 September.
- Solow, R. (1987). We'd better watch out. *New York Times Book Review*, 12 July, p. 36.
- Structural Change and Economic Dynamics. (2026). The macroeconomic effects of generative AI [cross-country study of 67 countries, 2022–2025]. Elsevier. [Verify author names before submission.]
- World Bank. (2026). AI offers lifeline to developing economies in an era of weak growth. Press release and report, 4 August.

Declaration by Author (s): "I/We hereby declare that this manuscript is my/our original work, free from plagiarism, and that all sources and any use of Artificial Intelligence tools for content generation or editing have been fully disclosed and verified for accuracy." **Tilka Devi**