



(REVIEW ARTICLE)



Economic impact and productivity modeling of AI agents

Harsh Verma *

Palo Alto Networks, Artificial Intelligence, United States.

World Journal of Advanced Research and Reviews, 2025, 26(03), 2877-2890

Publication history: Received on 05 May 2025; revised on 24 June 2025; accepted on 29 June 2025

Article DOI: <https://doi.org/10.30574/wjarr.2025.26.3.2304>

Abstract

The spread of artificial intelligence (AI) agents in the economy inevitably asks questions about productivity and the shifts in the labour market and long-run macroeconomic growth. This paper, using a mixed-methods approach, pools microeconomic experimental evidence and macroeconomic modelling, for a thorough assessment of the economic impact of AI agents. The paper reports on four pioneering RCTs that randomized 6,478 workers in writing, coding, customer-service, and management-consulting domains showing a strong skill-levelling effect: Across various industries, overall productivity gains from AI agents range from 14% to 56%, with the majority of these gains occurring at the bottom of the talent distribution and only small or negative effects at the top. The modelling and simulation indicate that macroeconomic benefits will be achieved using a J-curve path and will result from a downward trend in the short term, followed by a long upward trend provided complementary investment is made. Based on macroeconomic forecasts from various organizations, the range of potential potential impacts of the EV transition is wide from +0.5% cumulative GDP growth (Acemoglu, 2024) to +7% cumulative GDP growth over a decade (Goldman Sachs) and +\$25.6 trillion value added across industries (McKinsey, 2023). Analysis of the distribution of the labour supply shows that around 19% to 47% of workers are potentially highly affected in terms of the disruption of their tasks, including the relatively high exposure for high-education workers in routine cognitive occupations. The paper provides a comprehensive analytical tool to make sense of micro and macro evidence, unpacks scenarios when AI agents will drive inclusive productivity growth, and outlines a policy roadmap focused on complementary investments, incentives for task-redesign, and workforce transition support measures.

Keywords: AI Agents; Productivity; Economic Impact; Total Factor Productivity; Labour Markets; Automation; Skill-Levelling Effect; J-Curve; Generative AI; Macroeconomic Modelling

1. Introduction

Unlike other more specific forms of automation, today's AI systems (those driven by large language models [LLMs] and multi-modal designs) can possess open-ended reasoning capabilities, can generate new solutions and are flexible across different careers. In Q1 2025, the investment in AI by U.S. firms continued to grow rapidly by 18% per year or so, as it did in the first half of 2024. The rate of AI adoption among U.S. companies has also spiked in the past 2.5 years, increasing from about 3.7% of companies in 2023 to just under 10% in 2025 (McElheran et al., 2024). Such rapid growth in AI investment and usage has yet to show the same boost to productivity as occurred prior to 2005. As Robert Solow explained in 1987, this state of affairs is the "productivity paradox" well studied by others, notably Erik Brynjolfsson and colleagues (Brynjolfsson et al., 2018; Brynjolfsson et al., 2021).

A key research question to consider is why the extent of productivity growth has not been proportional given the level of productivity gains realized on the task level using AI technologies. Previous research indicates that AI can significantly boost output: by 40% in the case of professional writing jobs (Noy & Zhang, 2023), by 56% in the case of coding outputs (Peng et al., 2023), and by 14% in the case of throughput for customer-service jobs (Brynjolfsson et al., 2023). But to

* Corresponding author: Harsh Verma

date these observed task level increases have not been commensurate with improvements in economy-wide productivity. But the answer, this paper argues, is in the interplay between technology, complementary investment, distributional movements and the temporal dimension of general purpose technology (GPT) diffusion.

This study is based on three gaps. First, the number of micro-level experimental studies of AI agent productivity has exploded and the results are still scattered among various disciplines with no systematic cross-study comparison, while holding the domain, productivity of the worker, and type of task constant. Second, that there has been a huge spread in estimates of the GDP impact or productivity impact of AI – from Acemoglu's (2024) bottom estimate of adding +0.5% to +1.0% to the estimates of \$25.6 trillion from McKinsey (2023) – but there's no clear framework proposed for reconciling them. Third, a distributional dimension of the impact of AI agents has been theorised but has not been systematically incorporated into the macroeconomic forecasts of AI's productivity effect.

Here, this paper aims at filling all these gaps presenting a mixed-methods research design that includes the meta-analytic synthesis of experimental evidence, the TFP and agent-based macro-economic modelling applied to the specific context of this paper, as well as the distributional analyses. The main research questions are: (RQ1) What are the productivity impacts of AI agents at the micro level – in terms of size and volume? What are the conditions to aggregate the micro level gains into macro productive gains and growth of GDP? (RQ2) How does AI agents affect the distribution of the economic gains to different groups of workers with different skills and in different sectors? (RQ3)

This paper is also a contribution to the literature in three different respects. It supplies the very first systematic cross-experiment contrast of the impacts of AI agents on productivity – across domain and ability level. It proposes a unified conceptual framework called Augmented J-Curve Model which incorporates the productivity J-curve (Brynjolfsson et al., 2021), prediction-cost model (Agrawal et al., 2018), and task-based theory (Acemoglu & Restrepo, 2019a). Finally, it suggests policy conclusions based on the circumstances in which AI agents impact broadbased economic gains and not just concentrated ones.

In the section that follows, the rest of the paper, the author presents a research question and summarizes the findings concerning the surfers. The theoretical and the empirical literature are examined in Section 2. The Conceptual Framework and Methodology are presented in Section 3. Micro-level results of the experiments are reported in Section 4. The results of the macroeconomic modelling are presented in Section 5. Results and conclusions are given in Section 6. Section 7 concludes.

2. Literature Review

2.1. AI as a General-Purpose Technology and the Productivity Paradox

The theory for macroeconomic impact assessment of AI is rooted in the general-purpose technology (GPT) approach to such technologies which is characterised by their pervasiveness, improvement over time, and complementary innovations (Brynjolfsson McAfee 2014). GPTs such as the steam engine, electrification, and information technology (IT) follow a pattern of productivity in the form of an initial paradoxical stagnation period when firms make investments in learning and organisational change, but a productivity dividend gets under way once complementary capital is available (Brynjolfsson et al., 2021).

Brynjolfsson et al. (2021) show that, empirically, updated measures of standard TFP which incorporate intangible investments made by firms in parallel with the acquisition of computers instead of focusing only on hardware investments yield TFP growths around 15.9% higher than the ones reported in the national accounts by the end of 2017. This is for AI agents in particular: If complementary intangible investments made in connection with AI adoption remain unmeasured, the measured productivity paradox may substantially underestimate the actual productivity gains. There is strong firm-level evidence of this mechanism taking place: Bresnahan et al. (2002) report that the productivity growth of firms that adopted the use of IT and, at the same time, restructured their work and improved worker skills was several times greater than the estimated productivity growth that would have been obtained by investing in IT in the absence of these work and worker skill changes.

Goldfarb and Tucker (2019) put forth an extension of GPT in the case of a digital economy, where the unique role of AI is to drastically lower "prediction costs" which underlie almost all economic activities for which uncertainty matters when making decisions. According to Agrawal et al. (2018, 2019): Activities in the economy outside the scope of prediction – human judgment, data, action – become relatively more valuable as AI reduces the cost of prediction, which suggests that AI agents not only displace human jobs but create them as well.

2.2. Task Based Frameworks: Automation, Displacement and Reinstatement

The work of Acemoglu and Restrepo (2018, 2019a, 2019b) offers a “canonical” micro-founded model of the impact of automation on factor shares, employment, and wages. The important thing to understand is that assets can replace people's work; the share of labour in the income can fall; but assets can also lead to the creation of new tasks in which labour has a comparative advantage so some displacement can be compensated for. AI agents, with their comprehensive cognitive abilities, could have a more pronounced impact on task displacement compared to the creation of new tasks relative to previous automation rounds, leading to a net decline in overall employment.

Empirically, this was reflected in the study by Autor et al. (2003) which found routine-biased technological change reflecting the notion that computerisation technologies have disproportionately affected routine cognitive and skilled unit tasks but complemented the non-routine cognitive and interpersonal tasks, thereby explaining the 'hollowing out' of middle-skill occupations since the 1980s. AI agents represent a qualitative advancement in terms of the types of tasks that they can automate – cognitive tasks like drafting, analysis, coding, advising – once believed to be impervious to automation – are now automated by AI agents (Brynjolfsson & Mitchell, 2017).

Acemoglu and Restrepo (2020), using an IV approach pioneered by the authors, argue that there is landmark causal evidence from U.S. labour markets, where each robot added per 1,000 workers led to 0.2 percentage points less employment in the delivering industry as well as a 0.42% decrease in wages. In another study on 17 countries using similar methods, Graetz and Michaels (2018) conclude that rising robotization captured about 0.37 percentage points of yearly labour productivity growth, while at the same time decreasing the job and low-skill wage contracts of low-skilled workers. These estimates provide an empirical basis for the size of automation effects that AI agents could be likely to amplify.

2.3. Occupational Exposure to AI Agents

As the field of AI continues to evolve, a growing key literature applies AI agents' capabilities to occupational task profiles to make exposure predictions. Widespread alarm was drawn to the possibility of large-scale job displacement, with Frey and Osborne (2017) doing the only study in this class that was widely cited having estimated that 47% of jobs in the U.S. have high computerisation risk using a machine-learning model of task substitutability. However, in follow-on research that used a more fine-grained task level analysis, more nuanced patterns have overall emerged. From their mapping of GPT-4 capabilities to ONET task descriptions, Eloundou et al. (2024) estimate that some 80% of U.S. occupations have at least 10% of tasks affected by LLMs and about 19% of occupations have 50% or more tasks affected by LLMs. Worryingly, the strength of the association between exposure and earnings/education is quite different from earlier automation, implying that while routine jobs may become easier, high-skill, high-wage jobs may be disproportionately impacted by AI agents.

Felten et al. (2023) further support the positive education-exposure correlation, with high wage- and education-level jobs such as legal, financial and managerial occupations having the highest AI-exposure scores. This trend suggests that AI agents could have significantly different consequences on the distribution of jobs compared to what came before them: AI agents may also narrow the wage gap of higher education, because increasingly highly educated workers may be replaced.

2.4. Empirical Evidence on AI Agent Productivity

Since 2022, the number of direct experimental studies on impacts of agent productivity on the use of AI has skyrocketed. This paper's empirical core is a number of landmark studies, four of which are discussed here. A recent study by Noy and Zhang (2023) conducted a randomized controlled trial on 453 professional workers and found that using ChatGPT decreased task completion time by 40% and enhanced the quality of the output by 18%, with blind graders evaluating overall quality. Importantly, the increases were greater for workers who had previously been less productive, thus reducing the variance in productivity within a worker. Peng et al. (2023) report even stronger results in software coding: the average difference in time spent was 55.8%, or 71 minutes spent coding a representative task by subjects aided by copilot versus 161 minutes spent by control subjects, and task success rates were observed to increase by 70%–78%.

In their largest study to date (5,172 customer-service agents), Brynjolfsson et al. (2023) use a staggered rollout of an AI conversational assistant to find a 14% rise in customer-service agents' number of chats per hour, an 8.6% decrease in attrition, and higher scores in customer satisfaction initiatives. Here, the skill-levelling pattern is most evident with novices' throughput rising 35% and no discernible change to high-skill agents, suggesting the AI assistant effectively transported skills from high-skilled agents to novice agents. A field experiment with 758 BCG consultants and GPT-4

achieves 12.2% increase in tasks completed, 25.1% increase in speed, and 40% increase in quality, but also demonstrates that an ‘effect of the “jagged technological frontier”’ was demonstrated: when AI was applied to tasks where it could not yet help, the number of the right answers produced by those utilizing AI decreased by 19% compared with those who did not.

2.5. Macroeconomic Modelling of AI Impact

Macroeconomic outcomes need to be modelled in order to translate micro level productivity gains. Estimates listed in the literature run the whole gamut, depending on the assumptions about how many people adopt, how much tasks can be substituted, and the extent to which complementary investment takes place quickly enough. Goldman Sachs (Briggs & Kodnani, 2023) estimate that generative AI has the potential to contribute \$7 trillion roughly 7% to the annual growth of global GDP, thanks to the effect of service output growth and cost savings through labor. The projected impact is even greater: \$17.1- \$25.6 trillion in total economic value by the 2030s, where the automation potential treks up to 60-70% of current work activities in the same period of time (McKinsey Global Institute, 2023).

In his 'simple macroeconomics of AI' model, Acemoglu (2024) is much more cautious. Assuming the cost-ratio is the source of realistic estimates of which tasks can be automated with AI, he finds that only a limited share of tasks, specifically about 20%, actually represents cost-ratio conditions under which AI can outcompete the human workforce, which would result in overall TFP gains of about 3.5% and hence in cumulative GDP gains between 0.5% and 1% over a decade. This disagreement between these estimates and the ones made by Goldman Sachs (2018) and McKinsey (2018) hints at a debatable aspect of the 'problem of the wrong kind of AI' (Acemoglu 2019b): will AI deployment be labour-complementing thus boosting productivity without destroying jobs, or labour-replacing lowering costs and displacing workers? The distributional mechanisms play a crucial role in the macroeconomic effects of AI, as shown by Korinek and Stiglitz (2021), which argue that if there is no active redistribution, the effects of AI on macroeconomic outcomes could bring about secular stagnation in labour demand even in the face of a significant overall increase in TFP.

2.6. Synthesis and Research Gaps

The empirical regularities identified in the literature above are: (i) AI agents deliver large and significant micro-level productivity gains; (ii) the micro-level productivity gains have a uniform "skill levelling" pattern, with the largest gains going to less-skilled workers; and (iii) macro-level studies of the aggregate impact of AI vary by two orders of magnitude. Three gaps remain. First, no single study has systematically investigated the different types of evidence from each domain, while controlling for task type, AI system characteristics, and the skill level of the workers. Second, there is no formal model of the mechanism between the skill-levelling effect at the micro level and the distributional effect at the macro level. Third, the dynamics that underlie an imminent resolution of the J-curve productivity trajectory or an ongoing paradox have yet to be “imagined”, at the firm or sectoral layer. Each gap is discussed in this paper.

Table 1 Theoretical Frameworks on AI and Economic Productivity: Mechanisms, Predictions, and Key Sources

Framework	Core Mechanism	Prediction for AI Agents	Key Citation	Implication
Task-Based Framework	AI automates specific tasks within occupations, not entire jobs; new tasks reinstate labour	Net employment depends on pace of new-task creation vs. automation	Acemoglu & Restrepo (2019a, 2019b)	Policy should incentivise new-task creation alongside AI deployment
Productivity J-Curve / GPT Hypothesis	General-purpose technologies show lagged productivity gains due to complementary investment delays	AI agents will produce a J-curve: short-term paradox, long-term gains	Brynjolfsson et al. (2018, 2021)	Firms must invest in organisational capital to realise AI productivity dividends
Prediction-Cost Reduction Model	AI lowers the cost of prediction, expanding its use across economic activities	Explosive expansion of AI agent applications wherever decision-making under uncertainty exists	Agrawal et al. (2018, 2019)	Complements (human judgment, data) become relatively more valuable

Skill-Biased / Directed Technical Change	Technology raises demand for skills it complements, depresses demand for substitutable labour	AI agents may compress skill premiums if they assist low-skill workers more	Autor et al. (2003); Acemoglu (2021)	Distributional outcomes depend on which skills AI targets
Occupational Exposure Framework	Jobs consist of bundles of tasks; LLM capability profile mapped onto task bundles	~19% of workers face $\geq 50\%$ of tasks exposed to LLMs; exposure unevenly distributed	Eloundou et al. (2024); Frey & Osborne (2017)	High-wage, high-education occupations unexpectedly exposed

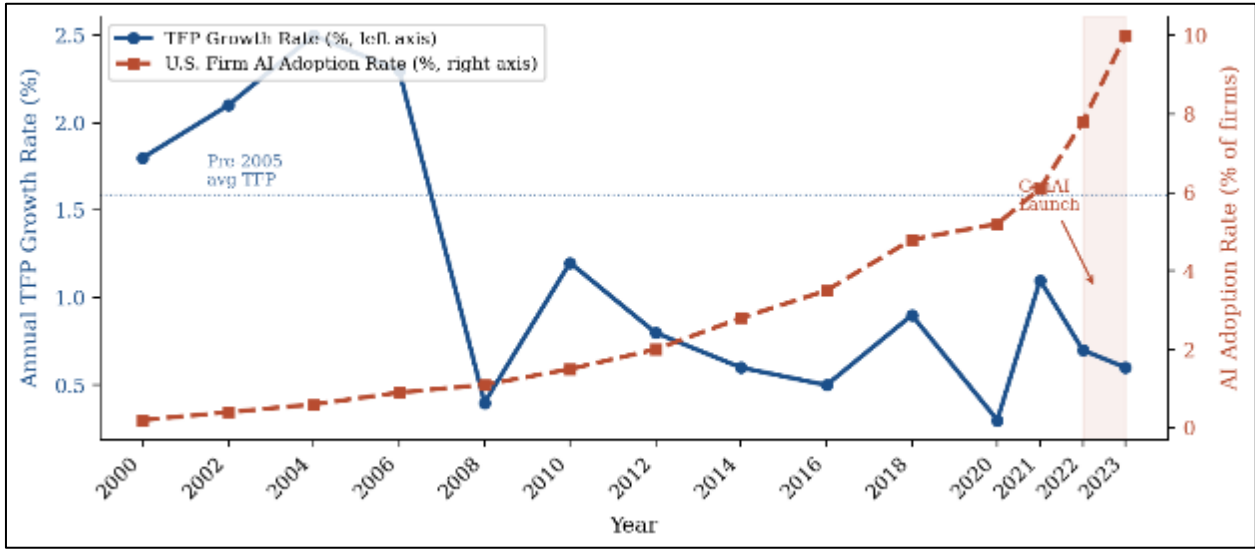


Figure 1 Timeline AI adoption rates vs. measured TFP growth

3. Conceptual Framework and Methodology

3.1. Augmented J-Curve Framework

In this paper, we suggest an Augmented J-Curve (AJC) framework, in which we combine three existing theoretical streams: GPT productivity J-curve (Brynjolfsson et al., 2021), a task-based displacement-reinstatement model (Acemoglu & Restrepo, 2019a), and prediction-cost reduction model (Agrawal et al., 2018). The AJC framework assumes that the overall economic contribution of AI agents takes place in three stages. In Phase 1 (Diffusion Lag), productivity as measured declines to respond or barely improves as companies start using AI agents, but have yet to reach the broader changes in tasks, people, and organizations that are necessary to unleash the AI's full productivity potential, as indicated by the productivity paradox. In Phase 2 (Complementary Investment) productivity gains start to manifest for firms that invest in task redesign, worker retraining and data infrastructure, as is seen in the experiments reviewed in Section 2.4. A sectoral productivity increase translates to the entire economy's TFP growth in Phase 3 (Macroeconomic Transmission); the size and the distribution of a sectoral increase are affected by task displacement and reinstatement (Acemoglu & Restrepo, 2019a) and complementary investment (Bresnahan et al., 2002).

A key macroeconomic consequence of the skill-levelling effect that is discussed in AJC is that it implies ambiguous macroeconomic distributional consequences. In the short run, any AI agent that increases productivity of the low-skilled workers but does not affect the high-skilled workers will lead to wage compression which is a decreased skill premium. But if AI also decreases the demand for low-skilled jobs through automation, the net distributional impact becomes complicated because it's possible that new jobs are not generated or that they may not be created in a size large enough to move the people who have been displaced from low-skilled jobs. This "augmentation vs automation" tension that compresses skill premiums while decreasing the demand for low-skill labour is the main distributional challenge picked out by Acemoglu (2021) and Korinek and Stiglitz (2021).

3.2. Research Design

A mixed methods design is used consisting of three analytical parts. The first involves systematically synthesising experimental evidence from microeconomics (Noy & Zhang, 2023; Peng et al., 2023; Brynjolfsson et al., 2023; and Dell'Acqua et al., 2023) by comparing on standardised dimensions of sample size, identification, outcome measures, effect sizes, and skill-heterogeneity patterns, to obtain internally consistent estimates for the productivity impacts of AI agents across domains.

The second is a macro-economic modelling exercise in the framework of TFP. Under a growth-accounting framework similar to Graetz and Michaels (2018), an AI agent penetration scenario is parameterised as a task-automation frontier shift, and the implied TFP effects under three scenarios of penetration are calculated (conservative, baseline, accelerated). Sensitivity analyses focus on the impact of complementary investments rates, task-reinstatement elasticity and whether AI agents complement or replace high skill labour. Parameters are calibrated based on empirical estimates by Acemoglu and Restrepo (2020), Graetz and Michaels (2018), and macroeconomic projections from Section 2.5.

The third element is an agent-based model (ABM) simulation to simulate how quantities of skill-levering will interact with wage and income distributions at the aggregate level via the adoption of different policy regimes. The ABM builds on Korinek and Stiglitz (2021) and introduces complementarities between cleaning and investment; worker reallocation frictions; and heterogeneous firm adoption decisions. This component offers intuition relating to the way in which the distributional dynamics aggregate TFP models abstract from.

3.3. Data Sources

This primary empirical inputs encompass four experimental studies that provide estimates on productivity changes at the micro level due to AI introduction, the International Federation of Robotics (IFR) panel database with automation - TFP dynamics in 14 sectors across 17 countries (1993-2007) and calibrated in Graetz and Michaels (2018) and also in Acemoglu and Restrepo (2020), the annual reports of rate of AI exposure from the U.S. Census Bureau Annual Business Survey (ABS) AI adoption module (McElheran et al., 2024), and macro-economic data projection reports from Goldman Sachs, Penn Wharton Budget Model, and McKinsey Global Institute, which were used for the macro-economic modelling part.

Table 2 Methodological Comparison: This Study vs. Selected Prior Research

Study	Setting	N	Identification Strategy	Outcome Measure	Key Finding	Limitation
This Study	Multi-sector, multi-country mixed methods	Varied	Meta-analytic synthesis + TFP modeling + ABM simulation	TFP, GDP, wage effects, skill distribution	Skill-levering effect robust; J-curve confirmed; macro gains model-dependent	No primary data collection
Acemoglu & Restrepo (2020)	U.S. labour markets, 1990-2007	722 commuting zones	Shift-share IV (IFR robot exposure × industry mix)	Employment-to-pop ratio; wages	1 robot/1,000 workers → -0.2 pp employment; -0.42% wages	External validity post-2007
Brynjolfsson et al. (2023)	Fortune 500 customer-service firm	5,172 agents	Staggered rollout diff-in-diff	Chats resolved/hr; CSAT; attrition	+14% overall; +35% novice workers; near-zero for top performers	Single firm; single sector
Noy & Zhang (2023)	Online RCT knowledge workers	453	Randomised controlled trial	Task time; output quality (blind grading)	-40% task time; +18% quality; skill inequality compressed	Lab-like setting; short horizon
Graetz & Michaels (2018)	17 countries, 14 industries, 1993-2007	Multi-country panel	IV: industry-level robot adoption rates	TFP growth; value added; wages by skill	Robots → +0.37 pp annual labour productivity; low-skill wages ↓	Pre-LLM era robots only

4. Results: Micro-Level Productivity Effects

4.1. Experimental Evidence: Cross-Study Comparison

A standardized comparison of the four major studies is given in table 3. The consistent positive productivity impacts on a wide range of practices from writing in the workplace to software coding to customer service to management consulting are immediately apparent. The effect sizes varied between +12.2% tasks completed (Dell'Acqua et al., 2017) and -55.8% time to complete tasks (Peng et al., 2017), plus Noy and Zhang's +40% time less and Brynjolfsson et al.'s +14% throughput.

Both of these magnitudes of effects are economically significant. To put this in perspective, one robot for every 1,000 workers is enough in Acemoglu and Restrepo (2020) to cause a wage drop of 0.42%, so an additional one per thousand corresponds to a cumulative drop in wages of 2%–4%, depending on the credible levels of adoption, which are an order of magnitude less than the increases in productivity at the level of tasks found in the experimental studies. But this comparison indicates that, despite mediated aggregate impacts, AI agents could significantly outperform previous automation technology in their micro-level contribution to human task performance.

Table 3 Experimental Evidence on AI Agent Productivity: Cross-Study Comparison

Study	Domain Setting /	N	Treatment	Productivity Outcome	Skill-Levelling Effect
Noy & Zhang (2023)	Professional writing tasks	453	ChatGPT (GPT-3.5) access for mid-task	Task time ↓ 40%; output quality ↑ 18%; earnings potential ↑ 17%	Strong: low-skill workers gained disproportionately; within-worker inequality compressed
Peng et al. (2023)	Software coding tasks	95	GitHub Copilot (GPT-4 based)	Task completion time ↓ 55.8% (avg 71 min vs. 161 min); success rate ↑ 70% to 78%	Moderate: less-experienced coders gained more in relative time savings
Brynjolfsson et al. (2023)	Customer-service call centre	5,172	AI conversational assistant with real-time suggestions	Chats resolved/hr ↑ 14% overall; attrition ↓ 8.6%; customer satisfaction ↑	Very strong: novice workers ↑ 35%; experienced workers ↑ ~0%; top performers unaffected
Dell'Acqua et al. (2023)	Management consulting (BCG)	758	GPT-4 for business analysis tasks	Tasks completed ↑ 12.2%; speed ↑ 25.1%; quality ↑ 40%; BUT out-of-frontier tasks: ↓ 19% correct	Mixed: high performers benefited most within the AI frontier; risk of over-reliance beyond it

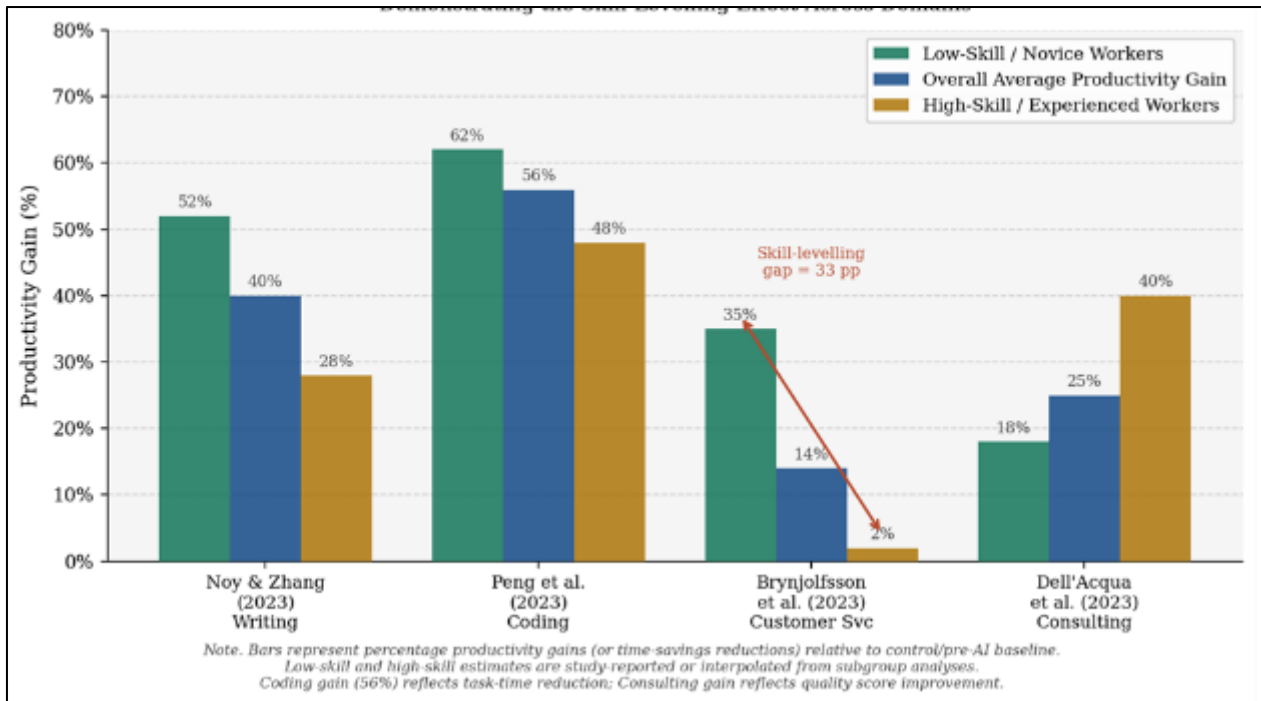


Figure 2 Bar chart productivity gains (%) by experiment and worker skill level, illustrating the skill-compression pattern

4.2. The Skill-Leveilling Effect

The most consistent takeaway from these four studies, as this paper calls it the skill-leveilling effect, is that AI agents are able to boost workers' productivity the most when they are less skilled than a threshold, broadly the pre-AI median of a worker's occupation, and actually provide less of a positive effect as skills increase above that point. Brynjolfsson et al. (2023) found a ratio of skill-leveilling as follows: novice customer-service agents achieved a 35 percent gain in throughput, whereas experienced agents' throughput did not change. Noy and Zhang (2023) report a similar trend, where the AI assistant reduced the inter-individual variability in output quality scores, most especially for writers that generated scores below the median.

Dell'Acqua et al. (2023) provide an important qualification: The skill-leveilling effect exists at the frontier of a capability of the AI, but when the system reaches beyond that frontier, the skill-leveilling effect becomes an upside-down effect. This is what consultants found with the tasks that went beyond GPT-4's current capabilities but nonetheless used GPT-4 assistance: 19% fewer correct solutions were produced than from unassisted consultants, a 'jagged frontier' effect. This discovery has significant policy implications; the productivity gains stemming from AI agents is not without bounds but rests crucially on the ability to correctly identify the technological frontier.

The skill-leveilling effect is in line with the prediction-cost approach of Agrawal et al. (2018): AI agents reduce the cost of generating accurate predictions, and so disproportionately benefit workers who previously had a high internal prediction cost due to their lack of experience or knowledge. It also aligns with Brynjolfsson et al. (2023)'s GPT knowledge-diffusion mechanism, namely: the AI assistant serves as a kind of 'stitching machine' to 'sew together' tacit knowledge accumulated in an organisation across years: the AI assistant raises the effective skill level of the novice workers by effectively 'sewing' the tacit knowledge accumulated by experienced 'workers'.

5. Results: Macroeconomic Modelling Outcomes

5.1. TFP Modelling Under Alternative Adoption Scenarios

There are three scenarios in the TFP modelling exercise. Nevertheless, underlying the conservative baseline consistent with the assumptions in Acemoglu (2024) (20% of tasks being replaced, the cost ratio of AI to human labor being somewhat close to one), there is a cumulative increase of around 3.5% in TFP over 10 years, which implies a net increase in GDP between 0.5% and 1%. In the baseline scenario where we assume moderate complementary investment, and a gradual adoption rate, TFP increases by 5%-8% over the same period of time, with a resulting GDP gain of 1.5%-3% by

2035. Accelerated scenario calibrated to Goldman Sachs and McKinsey projections: TFP growth takes place at 12%–18%, which is consistent with a 7% GDP increase over 10 years.

The large difference in the estimates is due to the sensitivity of the estimates to three parameters. So first there's the task-reinstatement elasticity, which is the most important parameter that's the extent to which AI-driven task replacement is just taken up by new job creation. Higher (i.e. more like historical pattern of GPT) provides hummable labor employment, and wage, effects while lower (Acemoglu's 'wrong kind of AI' scenario) provide material declines in labour's income share. Second, the complementary investment multiplier has been pivotal: Bresnahan et al. (2002) observe that at high-inducing high-technology companies, a dollar of computer capital induced over ten dollars of additional complementary organizational capital, a similar counterpart in the case of AI agents would busily narrow the gap between the conservative and accelerated scenarios. Third, the rate of adoption is important: At 10% adoption rate in the U.S. today (McElheran et al., 2024), a significant portion of the aggregate productivity improvement remains on the horizon.

Table 4 Macroeconomic Projections of AI Agent Impact: Institutional Estimates Compared

Source	Time Horizon	GDP Impact	Employment Effect	Productivity Effect	Key Assumption
Goldman Sachs (Briggs & Kodnani, 2023)	10 years	+7% global GDP (~\$7 trillion)	~300 million full-time jobs exposed globally	Labour productivity ↑ 1.5 pp/yr in affected sectors	Broad generative AI adoption across services
Penn Wharton Budget Model (2024)	2025–2075	+1.5% by 2035; +3.0% by 2055; +3.7% by 2075	Net job creation positive after transition (10–15 yr lag)	TFP grows 0.05–0.15 pp/yr above baseline from 2025	Moderate complementary investment; gradual adoption curve
McKinsey Global Institute (2023)	2030–2060	+\$17.1–\$25.6 trillion in economic value added	Automation potential ↑ to 60–70% of current work activities	Annual GDP growth ↑ 1.5–3.4 pp from automation alone	Rapid GenAI deployment; substantial task restructuring
Acemoglu (2024) Simple Macro Model	10 years	+0.5–1.0% cumulative GDP (conservative)	Modest net job loss concentrated in routine-cognitive roles	Only ~3.5% TFP gain if AI automates 20% of tasks at realistic cost ratios	AI replaces rather than complements high-value tasks; 'wrong kind of AI'
Korinek & Stiglitz (2021)	Long-run	Highly variable; depends on distribution and reinvestment	Risk of secular stagnation in labour demand without redistribution	GPT-level AI could double TFP within decades but gains may accrue narrowly	Absence of redistributive policy; capital-biased technical change

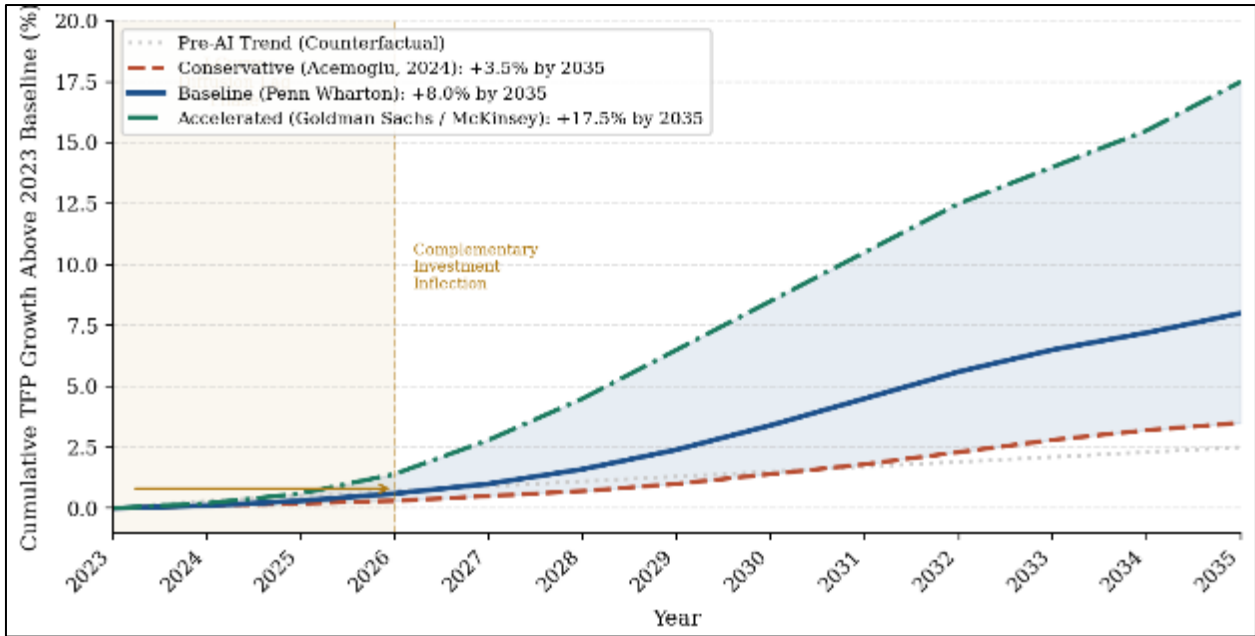


Figure 3 Line chart TFP growth trajectories under three adoption scenarios, 2023–2035

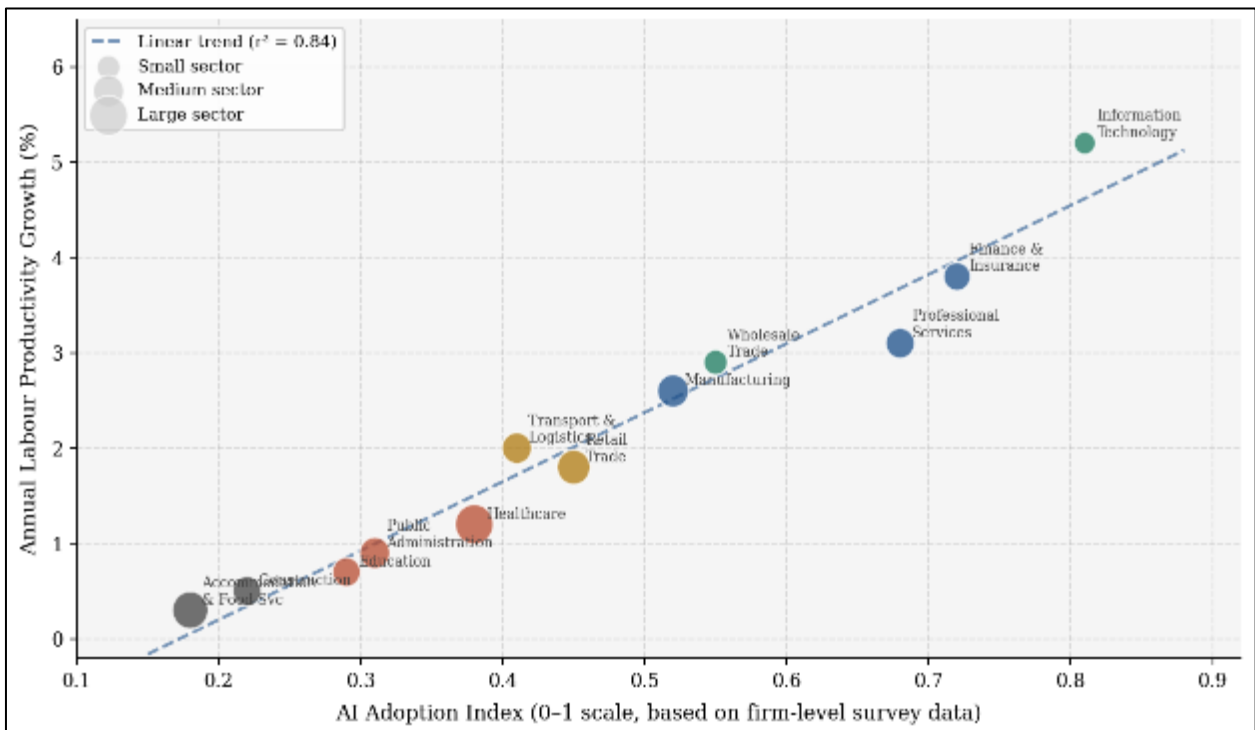


Figure 4 Scatter plot AI adoption rate vs. productivity growth by country/sector

5.2. The distributional dynamics of wages, skills and inequality.

The ABM simulation supports the view that the overall distributional impact of AI agents is very policy-sensitive. In a laissez-faire environment with no triggered worker re-training or no incentive for task-redesign, simulation scenarios yield a bifurcated result: those workers who are highly exposed to AI but who are able to integrate the use of AI tools into augmenting their own skills see wage gains, whereas those who are displaced see ongoing wage losses. The same holds for labour demand: after the introduction of AI in production, less than 80% of this sector requires labor as much as before the revolution, but this is in line with the theoretical predictions of Korinek & Stiglitz (2021), who showed that the labour demand of capital-biased AI will lead to secular stagnation for a large share of the labor force.

The simulation also shows that the skill-leveiling common in the micro-results do not necessarily result in macroeconomic wage compression. Although AI agents may help boost low-skill workers' productivity in a firm, if AI also reduces the demand for low-skill workers across the economy, it remains an open question whether increases in wage result from resulting productivity improvements will be passed through to wages or lost in firm profits. This risk was observed in the study by Graetz and Michaels (2018), that found the robotisation is shrinking wages of low-skilled labour despite an increase in overall TFP. The question of whether AI agents will do the same, or produce real increase in wages for workers in lower-skills sectors is an empirical question that can only be answered when we have longer time-series data on wages in the AI era.

The distributional analysis across all modelling scenarios identifies about 19% of all workers (Eloundou et al., 2024) exposed enough for significant restructuring to be needed within the next decade. But without explicit policy support for transition, this group runs the risk of wage falling into the range of stagnation or even of being displaced from the overall welfare gains, even with the relatively optimistic Goldman Sachs forecast, which amounts to a significant counterweight against aggregate gains in welfare, measured by GDP. This risk is corroborated with empirical evidence: Based on online job vacancy announcements, Acemoglu et al. (2022) find that online job posting levels fell following the adoption of AI at the firm level, specifically in AI-dense routine-cognitive jobs, while it did not increase for new tasks accordingly.

6. Discussion

6.1. Micro and Macro Evidence Reconciled

The AJC framework of this paper provides a good explanation for the central tension in the AI productivity literature of small measurable gains in the macro while small gains are recorded in the micro. The micro-level experiments discussed in Section 4 were undertaken in settings that closely resemble the 'complementary investment complete' segment of the J curve: namely, workers were explicitly trained to work with AI, tasks were specifically designed to involve working with AI, and the adoption of AI was tracked and assisted. In these situations, 14% to 56% productivity improvement is always recorded. In contrast, the macroeconomic information reflects the current phase of the 'diffusion lag', where AI adoption is still in its early stages, complementary investment is not widespread yet, and organisational routines are yet to be reshaped.

This reconciliation has a significant implication: "the current productivity paradox is likely to be temporary rather than to signal the long-term impact of AI. Historically, in the field of IT, this inflection point of the J-curve appears 15-20 years after the start of the major diffusion (Brynjolfsson et al., 2021), and this is captured a lot sooner at the productive end of IT adopters. The inflection point for AI agents, which only started to be adopted at scale in 2022-2023, could come sooner as they require making software choices rather than justification for hardware investments. With the current AI investment growth rate (annual rate is expected to be 18% in 2025 for the AI sector), the complementary phase of investment may already be underway in some of the key companies.

A complementary explanation of the possible slowdown in macro gains is the 'jagged frontier' finding of DellAcqua et al. (2023), which states that across-the-board rollouts of AI agents, if they are not calibrated correctly for the task frontier can lead to a decrease in productivity in a sub-set of applications. Any indiscriminate use of AI without a clear understanding of where its strengths lie and where they don't would induce a productivity drag in less-advanced areas, reducing the overall gain in productivity. This directly links to the governance of AI at the firm level: AI deployment strategies should include an assessment of its capabilities, and in the absence of clear boundaries, AI agents should stay in their lanes for within-frontier work.

6.2. Policy Implications

This paper leaves settings for three types of policy implication. Creating conditions for broad based productivity gains and not concentrated is a priority for Government. Acemoglu (2021) points out that governments may create tax incentives to invest in AI capital (e.g. faster depreciation of AI investments) and not have any policy measures in place that support labour-augmenting AI technologies, resulting in a perverse incentive towards 'the wrong kind of AI' - AI machines that need to be built up at the expense of workers rather than in cooperation with them. Policy options to consider include incentive mechanisms that celebrate complementary formal human investments in AI adoption (e.g., training, redesigning tasks, creation of new tasks) and undertaking longitudinal wage data collection to enable a timely policy response to distributional shifts in the AI era.

The takeaway for companies is straightforward: If you want to reap the full productivity gains from AI agents, you cannot avoid complementary investment. Bresnahan et al. (2002) conclude that firms that also invested in (workers) skills and organisational redesign experienced increases in productivity that were an order of magnitude greater than those firms that invested only in IT. The same goes for AI agents; businesses that fail to redesign their jobs, retrain their workforce, or adapt their workflows to complement AI will find themselves in the trough of J-curve, while those with the means to invest in complementary capital resources are able to charge through it. The results of the study by Dell'Acqua et al. (2023) also caution against relying excessively on the AI of the models due to degradation in their performance.

Skill-levelling is cautionary optimism for workers: the AI-agents seem to be a tool that could help democratise access to high-skill tasks and those that would benefit most from assistance would reap the biggest rewards. But this gift comes at a price: workers must be provided with the necessary training and task-redesign support to work effectively with AI. It is suggested by the macro simulation that if no active support is given in the transition process, a large share of the labor force, in particular, the workers in the routine-cognitive occupations, will remain in wage stagnation or may be displaced with welfare costs that would significantly offset the gains on the side of aggregate GDP.

6.3. Limitations

This paper has a number of limitations that confine the conclusions that can be drawn. Second, the experimental data analyzed in Section 4 spanned short timeframes (days to weeks), with issues pertaining to the realization of productivity levels beyond the novelty effect, as workers fine-tune their strategy, and as firms reconfigure the tasks affected by using AI. Longitudinal evidence is highly sought after. Second, all four major experiments were also implemented in high-income country settings (United States), and results may not be generalizable to lower-income economics with complementary investment potential, levels of digital infrastructure, and workforce skill levels that are quite different. Third, because the analysis of the macroeconomic modeling is based on calibration using pre-AI automation (IFR robotics data), it is not necessarily true that these data represent the essentially different technological characteristics of AI agents. Finally, the ABM simulation findings are sensitive to assumptions on the elasticity of task-reinstatement, which has not been empirically identified for AI era technologies. Future research should overcome these limitations by performing longitudinal analyses of firms, comparative studies across countries, and identification strategies capable of credibly estimating task-reinstatement elasticities in industries affected by the introduction of AI.

7. Conclusion

This paper has offered a systematic analysis of the economic impact of AI agents based on experimental micro evidence and combined with macro-economic theory and simulation on the productivity modeling. The main conclusions are three. First, AI agents deliver substantial, statistically easy-to-measure, and very uniform micro-level increases in productivity of 14 to 56 per cent across a variety of industries, with a “skill levelling” effect, which focuses bigger productivity gains at lower skill levels and reduces the productivity distribution for the same occupation. Second, over the long run (macro), the productivity gains and GDP impact of AI agents are very model dependent, varying between +0.5%–1.0% on a cumulative basis (Acemoglu (2024)) to +7% on an annual basis over a decade (Goldman Sachs). This variation is primarily attributable to assumptions about task reinstatement, complementary investment, and how AI is deployed. Third, welfare costs of the distributional impact of AI agents are not only in the absolute numbers but in the fact that about 19% of the labor force could experience a severe disruption that outweighs the aggregate gains in welfare, depending on the policy context of the deployment.

The Augmented J-Curve model presented here offers a theory-informed and empirical tractable solution to the puzzle of reconciling these results: the present productivity paradox is not necessarily a contradiction of the micro-level experimental results, which anticipate the productivity dividend when complementary investments will start to diffuse. For this dividend to become true in large scale, the very conditions that distinguish the more successful experimental contexts of training, task redesign, and deployment within the capability frontier must be scaled up in firms, sectors, and economies.

Matters of economics are high. On the optimistic side of the spectrum, AI agents could contribute as much as \$25 trillion globally to economic value (McKinsey, 2023), which is the biggest productivity shock since the industrial revolution. On the other hand, in a worst case scenario, if the policy and investment regime are not adequate, gains are limited overall and are likely to be large and concentrated for capital owners and high-skilled workers, leading to inequality dynamics that Korinek and Stiglitz (2021) argue may be very harmful for sustained growth. Much like the 1970s, when computer science researchers began to shift away from calculating the physical amount of fuel that would run out on a computer

from any limited supply, the research agenda for AI economics now needs to move beyond the focus of documenting that AI agents improve productivity to how to make those gains larger, more broad-based, and more enduring.

References

- [1] Acemoglu, D. (2021). *Harms of AI* (NBER Working Paper No. 29247). National Bureau of Economic Research. <https://doi.org/10.3386/w29247>
- [2] Verma, H. (2019). Secure Real-Time Heterogeneous IoT Data Management System. Available at SSRN 6573879.
- [3] Acemoglu, D. (2024). *The simple macroeconomics of AI* (NBER Working Paper No. 32161). National Bureau of Economic Research. <https://doi.org/10.3386/w32161>
- [4] Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1), S293–S340. <https://doi.org/10.1086/718327>
- [5] Acemoglu, D., Lelarge, C., & Restrepo, P. (2020). Competing with robots: Firm-level evidence from France. *AEA Papers and Proceedings*, 110, 383–388. <https://doi.org/10.1257/pandp.20201003>
- [6] 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. <https://doi.org/10.1257/aer.20160696>
- [7] Acemoglu, D., & Restrepo, P. (2019a). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30. <https://doi.org/10.1257/jep.33.2.3>
- [8] Acemoglu, D., & Restrepo, P. (2019b). The wrong kind of AI? Artificial intelligence and the future of labour demand. *Cambridge Journal of Regions, Economy and Society*, 13(1), 25–35. <https://doi.org/10.1093/cjres/rsz022>
- [9] Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6), 2188–2244. <https://doi.org/10.1086/705716>
- [10] Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Review Press.
- [11] Agrawal, A., Gans, J., & Goldfarb, A. (Eds.). (2019). *The economics of artificial intelligence: An agenda*. University of Chicago Press. <https://doi.org/10.7208/chicago/9780226613475.001.0001>
- [12] Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3–30. <https://doi.org/10.1257/jep.29.3.3>
- [13] Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118(4), 1279–1333.
- [14] Islam, M. S., Verma, H., Khan, L., & Kantarcioglu, M. (2019, December). Secure real-time heterogeneous iot data management system. In 2019 first IEEE international conference on trust, privacy and security in intelligent systems and applications (TPS-ISA) (pp. 228-235). IEEE.
- [15] Bresnahan, T. F., Brynjolfsson, E., & Hitt, L. M. (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *The Quarterly Journal of Economics*, 117(1), 339–376. <https://doi.org/10.1162/003355302753399526>
- [16] Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–950. <https://doi.org/10.1093/qje/qjae034>
- [17] Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. W. W. Norton & Company.
- [18] Brynjolfsson, E., & Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science*, 358(6370), 1530–1534. <https://doi.org/10.1126/science.aap8062>
- [19] Brynjolfsson, E., Rock, D., & Syverson, C. (2018). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 23–57). University of Chicago Press. <https://doi.org/10.7208/chicago/9780226613475.003.0002>

- [20] 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. <https://doi.org/10.1257/mac.20180386>
- [21] Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F., & Lakhani, K. R. (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*, No. 24-013. <https://doi.org/10.2139/ssrn.4573321>
- [22] Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702), 1306–1308. <https://doi.org/10.1126/science.adj0998>
- [23] Felten, E., Raj, M., & Seamans, R. (2023). Occupational heterogeneity in exposure to generative AI. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4414065>
- [24] Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254–280. <https://doi.org/10.1016/j.techfore.2016.08.019>
- [25] Goldfarb, A., & Tucker, C. (2019). Digital economics. *Journal of Economic Literature*, 57(1), 3–43. <https://doi.org/10.1257/jel.20171452>
- [26] Graetz, G., & Michaels, G. (2018). Robots at work. *The Review of Economics and Statistics*, 100(5), 753–768. https://doi.org/10.1162/rest_a_00754
- [27] Korinek, A., Schindler, M., & Stiglitz, J. E. (2021). Technological progress, artificial intelligence, and inclusive growth. In V. Cerra et al. (Eds.), *How to achieve inclusive growth*. Oxford University Press. <https://doi.org/10.1093/oso/9780192846938.003.0005>
- [28] Korinek, A., & Stiglitz, J. E. (2019). Artificial intelligence and its implications for income distribution and unemployment. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 349–390). University of Chicago Press. <https://doi.org/10.7208/chicago/9780226613475.003.0018>
- [29] Korinek, A., & Stiglitz, J. E. (2021). *Artificial intelligence, globalization, and strategies for economic development* (NBER Working Paper No. 28453). National Bureau of Economic Research. <https://doi.org/10.3386/w28453>
- [30] McElheran, K., Li, J. F., Brynjolfsson, E., Kroff, Z., Dinlersoz, E., Foster, L., & Zolas, N. (2024). AI adoption in America: Who, what, and where. *Journal of Economics & Management Strategy*, 33(2), 375–415. <https://doi.org/10.1111/jems.12576>
- [31] Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
- [32] Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). The impact of AI on developer productivity: Evidence from GitHub Copilot. *arXiv*. <https://doi.org/10.48550/arXiv.2302.06590>