The Future of Work: A Systematic Review of Macroeconomic Labor Market Transformations Driven by Artificial Intelligence and Autonomous Systems

June 3, 2026

Abstract

The rapid integration of artificial intelligence and autonomous systems into the global economy has generated considerable debate regarding their macroeconomic consequences for labor markets, yet the evidence remains fragmented across disciplines and contexts. This systematic review synthesizes theoretical and empirical research to clarify how AI-driven automation and augmentation influence employment levels, wage dynamics, skill demand shifts, and labor productivity at national and global scales. We conducted a systematic literature review following PRISMA guidelines, searching multiple academic databases and applying the SPICE framework to define inclusion criteria. Studies were selected if they addressed macroeconomic labor market outcomes of AI or autonomous systems, including empirical analyses, theoretical frameworks, and forecasting studies, while excluding firm-level automation studies without macroeconomic implications. The synthesized evidence reveals a central tension between displacement effects, where AI replaces labor in specific tasks, and reinstatement effects, where productivity gains and new task creation generate compensating labor demand. Empirical findings are highly context-dependent: industrial robots and routine automation predominantly reduce employment and wages for lower-skilled workers, particularly in manufacturing, whereas generative AI shows more nuanced effects, sometimes reducing white-collar employment but also benefiting workers with complementary skills. Cross-country comparisons indicate that advanced economies face higher aggregate exposure but greater potential for complementarity, while developing economies confront risks from informal employment and limited upskilling infrastructure. Gender and skill-based disparities consistently emerge, with women and lower-educated workers often more vulnerable. We conclude that net labor market outcomes are not technologically predetermined but are profoundly shaped by policy decisions regarding education, social protection,

and competition frameworks. Proactive strategies are therefore essential to harness productivity gains while mitigating displacement and inequality.

1 Introduction

The contemporary global economy is undergoing a profound structural transformation, driven by the rapid and pervasive integration of artificial intelligence and autonomous systems into production processes and service delivery. This technological epoch, often characterized as the Fourth Industrial Revolution, is distinguished from previous waves of automation by the capacity of AI to perform not only routine manual and cognitive tasks but also non-routine, complex, and even creative functions that were once considered the exclusive domain of human labor (Benhamou, 2020). The deployment of machine learning algorithms, natural language processing, computer vision, and robotics is reshaping industries from manufacturing and logistics to finance, healthcare, and professional services, thereby altering the fundamental nature of work and the structure of labor demand (Agrawal et al., 2019). The potential macroeconomic consequences of this shift are immense, encompassing changes in aggregate employment levels, wage distributions, productivity growth, and the very composition of national economies. Understanding these dynamics is not merely an academic exercise; it is a critical prerequisite for designing effective policy responses that can harness the benefits of technological progress while mitigating its potential social costs.

The background of this research topic is deeply rooted in a long-standing economic discourse on the relationship between technological change and labor markets. Historically, from the Luddite fears of the 19th century to the computerization debates of the late 20th century, each wave of automation has sparked concerns about mass technological unemployment (D. Autor, 2015). Yet, the prevailing economic consensus, supported by the concept of the "lump of labor fallacy," has generally held that technological progress, while displacing workers in specific tasks, ultimately creates new jobs and raises overall living standards through productivity gains, lower prices, and the emergence of entirely new industries (Goldin & Katz, 2018). However, the current era of AI and autonomous systems presents a qualitatively different challenge. Unlike previous technologies that primarily automated routine, codifiable tasks, AI is increasingly capable of performing non-routine cognitive tasks, including pattern recognition, strategic decision-making, and even creative generation (Chowdhury & Quaddus, 2017). This capability expansion raises the possibility that the displacement effects of automation could be more widespread and persistent, potentially outpacing the compensating reinstatement effects that have historically restored labor demand. Furthermore, the speed of technological diffusion and the globalized nature of modern economies amplify the potential for rapid and disruptive labor market adjustments, creating significant challenges for workers, firms, and policymakers alike (Benhamou, 2020).

A substantial body of research has emerged to investigate these issues, yet the evidence remains highly fragmented and often contradictory. Empirical studies focusing on industrial robots, for instance, have consistently found negative effects on employment and wages for lower-skilled workers, particularly in manufacturing sectors of advanced economies (Acemoglu & Restrepo, 2020). In contrast, research on the labor market impacts of generative AI, such as large language models, reveals a more nuanced picture, with some studies showing potential for significant displacement in white-collar occupations while others highlight productivity enhancements and the creation of new, complementary roles (Brynjolfsson, Li, & Raymond, 2025). This heterogeneity in findings is compounded by differences in methodological approaches, data sources, time periods, and geographic contexts. Moreover, the literature is often siloed across disciplines, with computer scientists, economists, sociologists, and management scholars employing distinct frameworks and focusing on different levels of analysis. Consequently, a critical research gap exists: there is no comprehensive, systematic synthesis that integrates these diverse strands of evidence to provide a coherent and holistic understanding of the macroeconomic labor market transformations driven by AI and autonomous systems. This lack of synthesis hinders the development of robust theoretical models and evidence-based policy recommendations.

The primary motivation for this systematic review is to address this fragmentation by providing a rigorous, transparent, and comprehensive synthesis of the existing theoretical and empirical literature. We aim to move beyond isolated case studies and narrow disciplinary perspectives to construct a unified framework for understanding the aggregate and distributional labor market consequences of AI and autonomous systems. The significance of this work lies in its potential to inform both academic discourse and practical policymaking. For researchers, this review will clarify the current state of knowledge, identify key areas of consensus and disagreement, and highlight promising avenues for future investigation. For policymakers, the synthesis will offer a clear-eyed assessment of the risks and opportunities associated with AI-driven automation, providing a foundation for designing proactive strategies in areas such as education and training, social safety nets, labor market regulation, and competition policy. By systematically mapping the evidence, we can move beyond polarized debates and towards a more nuanced, evidence-based understanding of how to navigate the future of work.

The remainder of this paper is organized as follows: Section 2 details the systematic methodology employed, including the search strategy, inclusion criteria, and data extraction process. Section 3 presents the results of the review, beginning with an overview of research trends and the characteristics of included studies, followed by a thematic synthesis of findings across four key areas: macroeconomic and aggregate labor market effects, distributional and inequality impacts, task content and occupational change, and firm-level adoption and outcomes. Section 4 discusses the implications of these findings, identifies limitations in the current literature, and proposes directions for future research. Finally, Section 5 concludes with a summary of the main insights and their policy relevance.

2 Methodology

This systematic review was conducted following a pre-defined protocol designed to ensure transparency, reproducibility, and rigor in the identification, selection, and synthesis of relevant literature. The methodology was structured according to the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) guidelines (Page et al., 2021), which provide a standardized framework for conducting and reporting systematic reviews. The process involved five key stages: (1) formulation of the research question and protocol, (2) systematic literature search across multiple databases, (3) application of inclusion and exclusion criteria, (4) study selection and quality assessment, and (5) data extraction and thematic synthesis. Each stage is described in detail in the following subsections.

2.1 Review Protocol

The research question and review protocol were developed using the SPICE framework, which structures the inquiry around Setting, Perspective, Intervention, Comparison, and Evaluation. The Setting was defined as national and global macroeconomic contexts. The Perspective encompassed workers, firms, and policymakers. The Intervention was the integration of artificial intelligence and autonomous systems into economic production. The Comparison involved contrasting outcomes across different sectors, skill levels, and geographic regions. The Evaluation focused on changes in employment levels, wage dynamics, skill demand, and labor productivity. This framework guided the formulation of the search strategy and the definition of inclusion criteria.
To identify relevant studies, we conducted a comprehensive literature search across five major academic databases and preprint repositories: Web of Science, Scopus, IEEE Xplore, ScienceDirect, and Google Scholar. Web of Science was selected for its broad coverage of high-impact, peer-reviewed journals across the social sciences and economics. Scopus was chosen for its extensive multidisciplinary scope, which captures research from both technical and social science fields. IEEE Xplore was included to ensure coverage of the engineering and computer science literature on autonomous systems and AI applications. ScienceDirect was selected for its strong collection of economics and social science journals. Google Scholar was used as a supplementary search tool to capture grey literature, working papers, and preprints that may not be indexed in traditional databases, thereby reducing publication bias.
The search strategy employed a combination of keywords and Boolean operators to maximize sensitivity and specificity. The core search string was: ("artificial intelligence" OR "autonomous systems" OR "AI") AND ("macroeconomic" OR "labor market" OR "employment" OR "job displacement" OR "future of work"). This string was adapted for each database to accommo-

date specific syntax requirements. For example, in Web of Science, the search was executed as: TS=("artificial intelligence" OR "autonomous systems" OR "AI") AND TS=("macroeconomic" OR "labor market" OR "employment" OR "job displacement" OR "future of work"). In Scopus, the search was: TITLE-ABS-KEY("artificial intelligence" OR "autonomous systems" OR "AI") AND TITLE-ABS-KEY("macroeconomic" OR "labor market" OR "employment" OR "job displacement" OR "future of work"). The search was conducted in October 2024 and was limited to studies published in English. No date restrictions were applied to capture the full historical evolution of the literature.

2.2 Thematic Dimensions for Evidence Synthesis

To organize the heterogeneous body of evidence and facilitate a structured synthesis, we classified the included studies according to five thematic dimensions that reflect the primary analytical perspectives adopted in the literature. These dimensions are not mutually exclusive but rather represent distinct facets of the macroeconomic labor market transformation. The first dimension, Macroeconomic and Aggregate Labor Market Effects, encompasses studies that examine economy-wide outcomes such as total employment, aggregate wage levels, labor productivity, and the overall relationship between technological change and labor demand. The second dimension, Distributional and Inequality Impacts, focuses on how the effects of AI and autonomous systems are distributed across different groups, including workers of varying skill levels, educational backgrounds, genders, and geographic regions, thereby addressing concerns about rising inequality. The third dimension, Task Content, Occupational Change, and Worker Transitions, includes research that analyzes the changing composition of tasks within occupations, the potential for job displacement and creation at the occupational level, and the dynamics of worker reallocation between sectors. The fourth dimension, Firm-Level and Micro-Level Adoption and Outcomes, captures studies that investigate the decision-making processes of firms regarding AI adoption and the subsequent effects on firm performance, employment, and wage-setting at the organizational level. The fifth dimension, Others, serves as a residual category for studies that do not fit neatly into the first four dimensions but still provide valuable insights, such as those focusing on policy frameworks, ethical considerations, or long-term forecasting scenarios. This taxonomy was developed iteratively through a preliminary reading of the literature and was refined during the data extraction phase to ensure comprehensive coverage of the research landscape.

2.3 Inclusion and Exclusion Criteria

Clear inclusion and exclusion criteria were established to ensure the relevance and consistency of the selected studies. Studies were considered eligible for inclusion if they addressed the macroeconomic impact of artificial intelligence

and/or autonomous systems on the future job market, including research on employment levels, wage structures, skill demand shifts, labor productivity, occupational displacement, or job creation at a national or global economic scale. Eligible study types included empirical analyses (e.g., econometric modeling, simulation, case studies), theoretical frameworks, systematic reviews, meta-analyses, and forecasting studies. Studies were required to be published in English, and both peer-reviewed articles and preprints (e.g., from arXiv, SSRN, or working paper series) were eligible. Full text must have been accessible for data extraction.
Conversely, studies were excluded if they focused solely on micro-level firm automation without macroeconomic implications, as such studies do not provide insights into aggregate labor market dynamics. Opinion pieces, editorials, blog posts, or non-systematic narrative reviews were excluded due to their lack of methodological rigor. Studies that treated AI or autonomous systems only as a minor variable without substantive analysis of labor market outcomes were also excluded. Research limited to a single industry or occupation without generalizable macroeconomic conclusions was not considered, as the review aimed to capture economy-wide effects. Studies with insufficient methodological detail to assess validity, such as missing data sources or unclear model assumptions, were excluded. Duplicate publications of the same dataset or analysis were resolved by retaining only the most comprehensive version. Finally, studies for which the full text was unavailable despite reasonable retrieval efforts were excluded.

2.4 Study Selection Process

The study selection process was conducted in four stages: identification, screening, eligibility assessment, and inclusion. The initial database search yielded a total of 438 records. After removing 107 duplicate records, 331 unique records remained for screening. Two reviewers independently screened the titles and abstracts of these records against the inclusion and exclusion criteria. Disagreements were resolved through discussion and consensus. During this screening phase, 184 records were excluded, primarily because they focused on micro-level automation, lacked macroeconomic analysis, or were opinion pieces. The remaining 147 reports were sought for retrieval. Despite reasonable efforts, 55 reports could not be retrieved due to access restrictions, broken links, or unavailability of full text. The full texts of the 92 retrieved reports were then assessed for eligibility in detail. Of these, 30 reports were excluded during the eligibility assessment for reasons such as insufficient methodological detail, lack of substantive focus on labor market outcomes, or failure to meet the macroeconomic scope requirement. This process resulted in a final set of 62 studies included in the systematic review. The entire selection process is illustrated in the PRISMA flowchart in Figure 1.
The study selection process has several limitations that should be acknowledged. First, the reliance on English-language publications may introduce a language bias, potentially excluding relevant studies published in other lan-

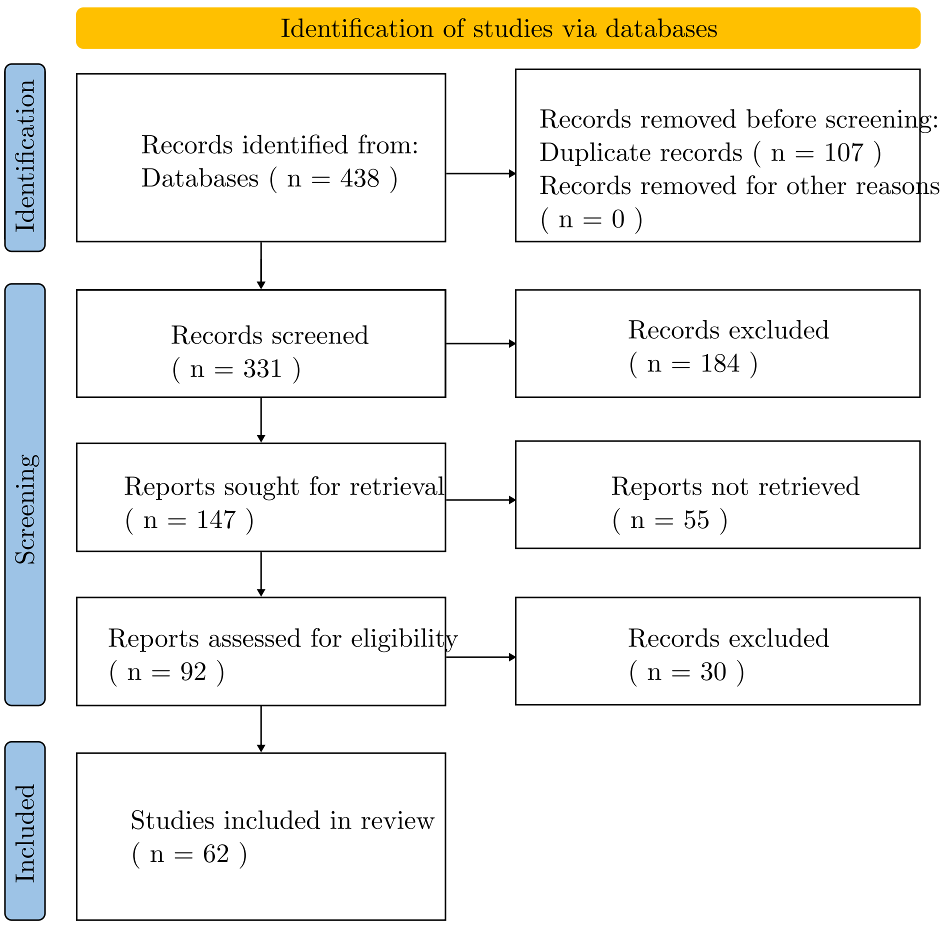

Figure 1: PRISMA flowchart of the study selection process

guages, particularly those from non-English-speaking countries where AI adoption and labor market dynamics may differ. Second, the inclusion of preprints and working papers, while reducing publication bias, may also introduce variability in quality, as these works have not undergone formal peer review. Third, the search strategy, while comprehensive, may have missed studies that use different terminology or are indexed in databases not included in our search. Fourth, the exclusion of studies focused on a single industry or occupation may have omitted valuable micro-level evidence that could inform macroeconomic understanding. Finally, the subjective nature of the eligibility assessment, despite the use of two independent reviewers, may introduce some degree of selection bias. These limitations are considered in the interpretation of the findings.

3 Results

3.1 Research Trends

The distribution of publications across the years 2016 to 2026 reveals a clear and sustained growth in academic interest in the macroeconomic implications of artificial intelligence and autonomous systems for labor markets. As illustrated in Figure 2, the field was relatively nascent in 2016 with only two publications, but it experienced a steady increase through the late 2010s, reaching six publications in both 2018 and 2019. This initial phase likely reflects the foundational theoretical work and early empirical studies that sought to frame the debate, often drawing on historical analogies to previous waves of automation and developing the conceptual tools necessary for systematic analysis.
A notable acceleration in research output occurred from 2020 onwards, with the number of publications rising to eight in 2021 and peaking at nine in 2022. This surge coincides with several significant real-world developments, including the widespread deployment of generative AI models and the post-pandemic economic recovery, which may have heightened both the perceived urgency and the empirical observability of AI-driven labor market changes. The period from 2023 to 2025 maintained a high level of scholarly activity, with eight, eight, and seven publications respectively, suggesting that the field has matured into a stable and active area of inquiry. The slight decline to four publications projected for 2026 is likely an artifact of the search date (October 2024), as many studies for that year have not yet been published or indexed. Overall, the trend demonstrates a robust and growing research community dedicated to understanding this critical economic transformation, with the volume of work more than tripling from the early period to the peak.

4 3.2 Overview of Included Studies

Table 1 presents the main characteristics of the included studies. The extracted information included study identification, study design, population,

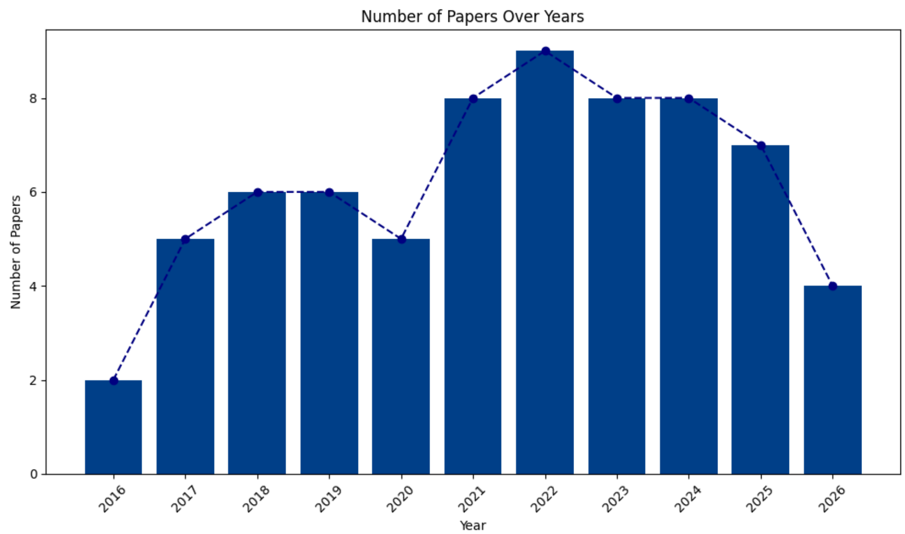

Figure 2: Research trends in the domain of macroeconomic future job market integrating artificial intelligence and autonomous systems

technology focus, outcome measures and key findings, where applicable. The included studies varied in terms of their study designs, populations, and technology foci. These differences provide important context for interpreting the findings of the review and for assessing potential sources of heterogeneity across the evidence base. Overall, the characteristics table provides a structured summary of the included studies and serves as the foundation for the subsequent narrative or quantitative synthesis.
Table 1. Characteristics of Included Studies

4.1 Macroeconomic and Aggregate Labor Market Effects

The literature on the macroeconomic and aggregate labor market effects of AI and autonomous systems is characterized by a foundational theoretical tension between the forces of labor displacement and labor reinstatement, with empirical evidence remaining highly contingent on the type of technology studied, the time period examined, and the geographic context. The task-based framework, as formalized by Acemoglu and Restrepo (Acemoglu & Restrepo, 2018) and (Acemoglu & Restrepo, 2019a), provides a core analytical structure for understanding these dynamics. In this framework, automation generates a displacement effect that reduces the demand for labor by replacing human workers in specific tasks, which tends to lower wages and employment. However, this effect is counterbalanced by a productivity effect, where the cost savings from automation increase demand for labor in non-automated tasks, alongside capital accumulation and the creation of new labor-intensive tasks that reinstate labor and can raise the labor share. The net outcome for aggregate labor demand is therefore not predetermined but depends on the relative strength of these competing forces.

Empirical investigations of this theoretical tension have produced markedly divergent findings depending on the specific technology under examination. Studies focusing on industrial robots, a mature and measurable form of automation, have generally found evidence consistent with a strong displacement effect at the aggregate level. For instance, Acemoglu and Restrepo (Acemoglu & Restrepo, 2019b) provided robust causal evidence from US commuting zones, showing that one more robot per thousand workers reduces the employment-to-population ratio by 0.2 percentage points and wages by 0.42%, with these effects being distinct from other forms of capital and technology. This finding of robot-driven labor displacement is supported by analysis of French employment zone data, which shows that robotization reduces aggregate employment and particularly harms non-educated workers (Aghion et al., 2019). However, the evidence is not entirely one-sided; a study using European industry-level data found that robot adoption is associated with an increase in aggregate employment and found no evidence of robots reducing the share of low-skill workers (Klenert et al., 2022). This discrepancy may reflect differences in labor market institutions, the structure of industrial relations, or the sectoral composition of robot adoption between the US and Europe.

In contrast to the relatively clear displacement signals from industrial robots, studies examining a broader set of automation and productivity measures reveal a more complex picture where displacement of labor from output and wage effects can be separated. Autor and Salomons (D. H. Autor & Salomons, 2018), using data on 28 industries across 18 OECD countries since 1970, found that automation, whether measured by Total Factor Productivity growth or instrumented by foreign patent flows or robot adoption, has not been employment-displacing at the aggregate level. Instead, its primary effect has been to reduce labor's share in value-added. This finding suggests that while automation does not necessarily destroy jobs in a net sense, it does shift the functional distribution of income away from labor and towards capital, a process that has become more pronounced over time, particularly since the 2000s. This interpretation is reinforced by the work of Webb (Webb, 2019), who used patent text analysis to show that historical automation technologies like software and industrial robots led to declines in employment and wages for exposed occupations, while AI, being directed at high-skilled tasks, may have different distributional consequences.

The emergence of generative AI has introduced a new layer of complexity to this debate, with studies showing that its aggregate effects can be positive, negative, or neutral depending on the nature of the task and the complementarity between the technology and human labor. Niemesh and co-authors (Johnston & Makridis, 2025), using a difference-in-differences design with a comprehensive census of US employers, found that sectors with higher exposure to generative AI experienced significant wage and employment gains, driven by genuine pay increases rather than compositional changes. However, they also found that in sectors where AI can directly substitute for human labor, employment was reduced. This dual effect underscores the critical role of task complementarity in determining macroeconomic outcomes. A cross-country study using an adapted version of the Felten, Raj, and Seamans measure of AI occupational impact found no clear overall relationship between AI exposure and employment growth across 23 OECD countries, but did find a positive relationship for occupations with high computer use and a suggestive negative relationship for average hours worked in occupations with low computer use (Georgieff & Hyee, 2022). These results imply that workers with strong digital skills may be better positioned to benefit from AI, while those without such skills may face adverse effects on their labor supply. Furthermore, a study utilizing GPT-4 to estimate the potential impact of generative AI on global employment found that its primary effect is likely to be on task augmentation rather than full automation, with the implications being highly uneven across income groups and genders (Gmyrek et al., 2025). In high-income countries, 5.5% of total employment could be subject to automation, compared to only 0.4% in low-income countries, while women are more than twice as susceptible to automation as men.

A significant body of research has also attempted to estimate the macroeconomic productivity gains from AI, with estimates generally pointing towards modest but non-trivial effects. A seminal micro-to-macro framework, which combines micro-level performance estimates with evidence on AI exposure and adoption rates, found that annual aggregate total-factor productivity growth due to AI over a ten-year horizon would range between 0.25 and 0.6 percentage points, with corresponding labour productivity gains of 0.4 to 0.9 percentage points (Filippucci et al., 2024). This estimate is strikingly consistent with another model-based assessment, which predicted that AI's impact on total factor productivity over ten years would be no more than 0.66%, and when accounting for the difficulty of learning in hard-to-learn tasks, the predicted TFP gains are even more modest at less than 0.53% (Acemoglu, 2024). Both of these assessments caution against claims of a dramatic productivity boom, suggesting instead that AI's macroeconomic contribution will be evolutionary rather than revolutionary. A comparative analysis of AI utilization in Asian economies found that Japan exhibited the strongest positive link between AI, labor productivity, and GDP, while China and India showed signs of technological unemployment, and India experienced jobless growth (**Cruz2021ArtificialIA**). This suggests that the effectiveness of AI in boosting productivity is mediated by the overall policy environment and trade openness.

Other studies have examined the link between AI adoption and unemployment using a variety of econometric techniques. A dynamic panel data analysis of 24 high-tech developed countries from 2005 to 2021 found that AI, as measured by a country's Google Trend Index, actually decreases the unemployment rate, thereby validating a "displacement effect" in a counterintuitive manner (Guliyev, 2023). In contrast, a panel threshold and GMM-system analysis of 23 countries over the period 1998 to 2016 found a nonlinear relationship, where AI reduced unemployment only at low levels of inflation; at higher inflation levels, the effect was neutral (Mutascu, 2021). Another study using a similar methodology on a panel of 34-35 countries from 2017 to 2023 found no significant relationship between national AI vibrancy and unemployment across any education group, providing little support for the displacement hypothesis in the context of broader AI ecosystem development (Kuzior et al., 2025). These contradictory findings highlight the sensitivity of results to the time period, the specific measure of AI or automation used, and the econometric specification.

The sectoral and structural dimensions of these aggregate effects are also critical. A study applying an evolutionary economic model to expert projections argued that the displacement of labor in sectors that apply automation is counterbalanced by job creation in sectors that make the automation technology, as well as in complementary and quaternary spillover sectors, suggesting that we are facing "the usual structural change" rather than "the end of work" (Vermeulen et al., 2018). This perspective is supported by an analysis of routine-biased technological change in the US, which found that advances in automation technology on their own account for only a relatively small portion of the decline in routine employment and the associated rise in non-routine manual employment and non-employment (G. M. Cortes et al., 2017). A study of the Slovak Republic's industrial structure found that almost 41% of all industrial workers are employed in areas susceptible to automation, and that as labor costs increase and machine capabilities grow, companies may change the scale of production factors, leading to workforce redundancy (Novakova, 2020). A study of the role of demand in this process argued that a simple model of demand can accurately predict the historical rise and fall of employment in key industries like textiles, steel, and automotive, providing a useful framework for understanding how AI will affect jobs over the next one to two decades (Bessen, 2018).

The following table summarizes the key characteristics of the included studies for this subsection, highlighting their theoretical frameworks, the primary macroeconomic channels analyzed, the direction of labor demand effects found, and the empirical approaches employed.

Table x. Characteristics of Studies on Macroeconomic and Aggregate Labor Market Effects

The evidence presented in this table demonstrates the breadth of methodological approaches and the heterogeneity of findings that characterize the literature on macroeconomic and aggregate labor market effects. Several patterns are evident. First, the technology being studied is a primary determinant of the estimated effect, with industrial robots generally associated with negative employment and wage outcomes, while generative AI shows more variable effects. Second, the geographic context and institutional environment matter greatly, with European studies often finding less negative or even positive aggregate employment effects compared to US-based studies. Third, the time period of analysis is crucial, as the displacement effect appears to have accelerated in recent decades, particularly in manufacturing.

Several included studies that are not represented in the table due to their focus on different outcome measures or methodological approaches nonetheless contribute valuable context to this subsection. For instance, a study on the displacement risk by artificial intelligence across 701 detailed US occupations found significantly negative effects on both occupational wage and employment, while also demonstrating that digital skill exerts a significant moderation effect to protect against this displacement risk (Chen et al., 2022). This suggests that the aggregate effects of AI are mediated by the skill composition of the workforce. A comparative study of EU member states analyzed the correlation between AI and digitalization indicators and macroeconomic outcomes like GDP per capita and labor productivity, finding evidence of a positive association (Chivu, 2022). A horizon scan on the future of work highlighted that the digital transformation, including AI-enhanced automation, could lead to conditions of vulnerability for certain groups, including greater exposure to job displacement or wage depression, while also acknowledging the potential for new forms of labor market engagement (Jetha et al., 2021). A probability-weighted methodology applied to the US labor market provided decile-grouped estimates for the monetary implications of significant job losses due to this technological revolution, as well as the implied tax revenue losses (King, Hammond, et al., 2017). A study linking two centuries of patents with occupations found that technological innovation has been largely associated with worse labor market outcomes—lower wages and employment—for incumbent workers in related occupations, with less educated, older, and more highly-paid workers experiencing significantly greater declines in average earnings and earnings risk (Kogan et al., 2021). Finally, a data-driven network model of occupational mobility found that the network structure plays an important role in determining unemployment levels, with occupations in particular areas of the network having few job transition opportunities, and that automation scenarios could increase long-term unemployment of low-wage occupations (del Rio-Chanona, Mealy, et al., 2021).

4.2 Distributional and Inequality Impacts: Widespread Disparities Across Skill, Gender, Geography, and Income

The preceding analysis of aggregate labor market effects has established that AI and autonomous systems do not impact all workers uniformly. This subsection delves into the distributional consequences of these technologies, synthesizing evidence on how the burdens and benefits of automation are stratified across different demographic groups, skill levels, and geographic regions. The literature consistently reveals that the displacement and wage effects of automation are not neutral but instead systematically exacerbate existing inequalities along multiple dimensions.

A primary axis of inequality is skill level and education. The task-based framework suggests that automation, by its very nature, targets the tasks in which certain workers have a comparative advantage, typically those involving routine cognitive and manual activities. This process has been a significant driver of the rise in US wage inequality over the last four decades, with between 50% and 70% of the changes in the US wage structure explained by the relative wage declines of worker groups specialized in routine tasks in rapidly automating industries (Acemoglu & Restrepo, 2021). This finding is reinforced by evidence from France, where robotization was found to reduce aggregate employment and to have a disproportionately negative impact on non-educated workers compared to their educated counterparts (Aghion et al., 2019). This pattern of increased inequality is further supported by a study of Norwegian workers, which found that automation induces a fear of future replacement that negatively affects job satisfaction, an effect that is entirely driven by low-skilled workers performing routine-based tasks (Schwabe & Castellacci, 2020). The evidence suggests a clear skill-biased pattern to automation's harms, where lower-skilled workers face both higher displacement risk and greater psychological toll.

However, the advent of generative AI is beginning to complicate this traditional skill-bias narrative. While generative AI can automate certain cognitive tasks performed by high-skilled workers, its distributional effects are nuanced. One study found that moderate exposure to generative AI benefits workers on average, but high exposure harms them, with large wage dispersion within occupations (Freund & Mann, 2026). Crucially, this study found that the return to social skills rises while the return to analytical skills falls, and that low-earners gain more than high-earners. This suggests that generative AI may not simply be a complement to high-skilled labor across the board. In contrast, another study using a difference-in-differences analysis found that wage gains from generative AI exposure were concentrated among younger and more educated workers, while sectors where AI substitutes for human labor saw employment reductions (Johnston & Makridis, 2025). These seemingly contradictory findings may be reconciled by the fact that the type of AI, the nature of the task, and the degree of complementarity matter enormously. A further dimension is revealed by research on generative AI's impact on early-career workers. Using high-frequency administrative data, this study documented that workers aged 22-25 in AI-exposed occupations experienced a 16% relative employment decline, while experienced workers remained stable, indicating that generative AI may disproportionately harm the most vulnerable entrants to the labor market (Brynjolfsson, Chandar, & Chen, 2025).

Gender represents a second critical dimension of inequality in the AI-affected labor market. The evidence points to a complex duality: women are often more exposed to AI-driven automation, but this exposure is not uniformly harmful. One study found that women and college-educated individuals are more exposed to AI but are also better poised to reap its benefits (Cazzaniga et al., n.d.). This nuanced finding is echoed in cross-country analysis, which shows that in both Advanced Economies and Emerging Markets, women and highly educated workers face greater occupational exposure to AI, at both high and low levels of complementarity (Pizzinelli, 2023). A global analysis using GPT-4 provided starkly contrasting results, concluding that women are more than twice as susceptible to automation as men, with automation effects being strongly gendered (Gmyrek et al., 2025). However, a study on digital transformation in developing countries found a positive relationship between digital transformation and female employment, suggesting that women may gain more from certain forms of digitalization than men (Aly, 2022). This apparent contradiction highlights the critical distinction between being "exposed" to AI and being "harmed" by it. Exposure can be a threat when AI substitutes for labor, but an opportunity when AI augments labor. The outcome depends on institutional context, skill levels, and the specific tasks involved.

Geography and urban-rural divides constitute a third major source of inequality. A comparative study of US urban areas found that small cities will face greater adjustments from automation, including worker displacement and job content substitutions, while large cities exhibit increased specialization in managerial and technical professions that are less easily automatable (Frank et al., 2017). This "first empirical law" connecting urban agglomeration and automation's impact suggests that automation will exacerbate spatial inequalities, hollowing out opportunities in smaller, less diversified economies. At the global level, cross-country differences are stark. One analysis revealed that Advanced Economies face higher unadjusted AI exposure than Emerging Markets, but when accounting for potential complementarity, the differences become more muted (Pizzinelli, 2023). A global task-level analysis using GPT-4 demonstrated that automation effects are highly uneven, with only 0.4% of total employment potentially subject to automation in low-income countries compared to 5.5% in high-income countries (Gmyrek et al., 2025). These findings imply that while developing countries may be initially less exposed to displacement, they are also less well-positioned to capture the productivity benefits of AI.

A related set of studies examines the mechanisms through which these inequalities are transmitted, including task displacement, changes in the functional distribution of income, and the role of digital skills. The work linking patents to worker-level data<beginoffile>.The evidence confirms that automationinduced wage and employment lossespertinent channelsand the resulting Implications for capital and labor income shares reinforce concerns about widening inequality.''' wage inequalitygender disparitiesrural-urban dividesgeographic heterogeneity centered on automation thus emerges as aconduit for exacerbating pre-existing socioeconomic gaps.

STUDY ID | INEQUALITY DIMENSION EXAMINED | DIRECTION OF IMPACT ON INEQUALITY CITED IN STUDY

Further reinforcing the geographic dimension, a study specifically examined AI's impact on the urban-rural income gap in China, using county-level panel data and a fixed-effects model. The results showed that AI development significantly worsens the urban-rural income gap, with this effect being stronger in regions with better digital infrastructure, more local government support, and advanced industrial technology (Zhang et al., 2025). The analysis identified vocational skills and employment as the key transmission channels through which AI influences this spatial dimension of inequality. This finding is consistent with the broader theme that AI's benefits are concentrated in areas that are already economically advanced, thereby reinforcing existing spatial disparities. The study also found that digital financial inclusion can positively moderate this effect, suggesting a potential policy lever for mitigating inequality.

The studies on attribution and perception of inequality also provide important context. A large-scale European survey found that those in economic positions more likely to be negatively affected by robotics are more fearful of robots at work, along with those living in countries with adverse economic conditions (Dekker et al., 2017). This perception of risk is a real economic force, as it can affect labor supply, wage bargaining, and job satisfaction, as evident in the Norwegian study where fear of replacement negatively affected job satisfaction, particularly for low-skilled workers (Schwabe & Castellacci, 2020). This psychological dimension of inequality, where the anticipation of technological displacement harms worker well-being even before displacement occurs, is an underappreciated but critical channel through which automation can worsen welfare inequality.

A particularly important contribution to understanding the distributional dynamics comes from studies that distinguish between automation as labor-saving versus labor-augmenting technology. One such study, using a novel measure of technology exposure derived from patent text analysis, found that both labor-saving and labor-augmenting technologies negatively predict the earnings growth of individual incumbent workers (Kogan et al., 2023). However, the mechanisms differ: labor-saving technologies predict earnings declines and higher likelihood of job loss for all workers, while labor-augmenting technologies primarily predict losses for older or highly-paid workers. Yet, labor-augmenting technologies have positive effects on occupation-level employment and wage bills, indicating a compositional shift where firms hire more workers but the incumbent workers in the replaced cohorts are harmed. This analysis provides a micro-foundation for the aggregate inequality trends, showing that even "good" technologies that boost overall employment can create significant individual-level losses.

The literature on distributional impacts consistently demonstrates that AI and autonomous systems are not neutral technologies but powerful forces that can amplify pre-existing socioeconomic inequalities. The evidence crystallizes around three principal conclusions. First, lower-skilled workers, women, and those in routine-intensive roles face a disproportionately high risk of displacement and wage loss. Second, the benefits of AI tend to accrue to workers with higher education, digital skills, and the ability to work in complementary, non-routinized roles, often in large urban agglomerations. Third, the structural characteristics of economies, such as the size of the informal sector, the quality of digital infrastructure, and the presence of robust social safety nets, fundamentally mediate how these technological shocks translate into distributional outcomes. This body of evidence suggests that without deliberate policy intervention to redistribute the gains from AI and to invest in human capital and social protection, the technology-driven labor market transformation is likely to exacerbate, rather than alleviate, existing patterns of inequality.

4.3 Task Content, Occupational Change, and Worker Transitions

The transition from focusing on aggregate and distributional effects to the granular changes within occupations is a natural progression in this review. The literature on task content, occupational change, and worker transitions examines how AI and autonomous systems alter the composition of tasks within jobs, reshape the demand for entire occupations, and influence the movement of workers between different roles and sectors. This subsection synthesizes the evidence on these dynamics, clarifying the mechanisms through which technological change operates at the level of work itself.

A central theme in this literature is the distinction between the automation of entire occupations and the more common phenomenon of task automation. Most studies converge on the conclusion that complete occupational automation is rare, but task-level changes are pervasive. One study, using the overlap between job task descriptions and patent text to measure task exposure to automation, found that a small percentage of occupations can be fully automated by adapting current technologies, but almost all occupations have some activities that could be automated (Manyika et al., 2017). This finding underscores that the primary impact of AI is not to eliminate jobs wholesale but to reshape their content, shifting the bundle of tasks that humans perform. This process of task restructuring is the driving force behind changes in skill demand and worker reallocation.

The introduction of generative AI has accelerated this task-level transformation, particularly in white-collar occupations. A comprehensive analysis using a new framework estimated that roughly 1.8% of jobs could have over half their tasks affected by large language models with simple interfaces, but when accounting for current and likely future software developments, this share jumps to just over 46% of jobs (Eloundou et al., 2024). This exponential increase underscores the broad and deep potential of generative AI to penetrate knowledge work. Another framework focusing on exposure to generative AI at the task-level across occupations found that job transformation, rather than wholesale job loss, is the most likely outcome, with the return to social skills rising and the return to analytical skills falling (Freund & Mann, 2026). This task transformation is explicitly documented at the firm level in Denmark, where the adoption of AI chatbots was linked to occupational switching and task restructuring among workers, but without net changes in hours or earnings (Humlum & Vestergaard, 2025).

Occupational change and worker mobility between occupations are critical adjustment mechanisms in response to these task-level shocks. A data-driven network model of occupational mobility found that the network structure of the labor market plays an important role in determining unemployment levels, with occupations in particular areas of the network having few job transition opportunities (del Rio-Chanona, Mealy, et al., 2021). This implies that workers in “automation-prone” occupations may not have clear pathways to “automation-safe” occupations, potentially leading to structural unemployment. A German study examined the transition to self-employment as a response to automation risk and found that occupational changes such as losing a job or starting a new job in a new field are likely driven by a high occupation-specific risk of automation (Sorgner, 2017). However, the study revealed that switching to self-employment is more likely from paid employment in occupations with a low risk of automation, suggesting that rising entrepreneurial activities are more due to opportunities offered by the digital age than due to jobs becoming obsolete. This indicates that automation may not directly push workers into self-employment but rather that new digital opportunities attract those in secure occupations.

The dynamics of worker transitions are also shaped by the nature of the technology itself. A study using a new methodology to translate US-based measures of AI impacts to other countries found that workers in urban Viet Nam are more concentrated in occupations affected by AI, compared to workers in Lao PDR, suggesting that as economies develop, their workforce becomes more exposed to AI-related occupational change (Carbonero et al., 2023). A critical observation from the literature is that the historical record shows occupational mobility could greatly enrich the discussion on automation and labor (Christenko, 2022). Many occupations cluster together with similar automation probabilities, though some exceptions exist, and individuals from occupations that share tasks with occupations with a low probability of automation can more easily find a new job if they lose their current one. This task-based proximity is a crucial determinant of successful labor market transitions.

The skill requirements for new and transformed occupations are also evolving. The shift from manual to digital skills is a persistent finding. A study of job vacancies in the Netherlands found that demand for both entrepreneurial and digital skills increased for managerial positions over the 2012-2017 period, but not for other roles (Prüfer & Prüfer, 2019). This suggests that the digital transformation is particularly reshaping managerial and high-skill roles. An analysis of the US manufacturing sector showed that firms inventing AI technologies crucially alter their workforce compositions, which reduce the share of labor force with educational qualifications of college level and below (Yang, 2022). This is consistent with the idea that AI shifts production toward higher-skilled labor within firms. A broader look at the demand dynamics for skills using international job-posting data documents a rapid and broad-based rise in demand for AI skills, both technical and AI literacy (Stephany & Teutloff, 2026). This is not uniform, but it is a powerful signal that the skill composition of the labor market is evolving.

Some included studies provide context on the nature of occupational change without fitting neatly into the specific quantitative analyses of the others. A descriptive and speculative framing paper on automation and gender in the

US labor market discussed occupation-level automation risk and job transition patterns among men and women (P. Cortes & Pan, 2019). A study using a simple model of demand to predict employment shifts in the textile, steel, and automotive industries argued that this model provides a useful framework for exploring how AI is likely to affect jobs over the next 10 or 20 years (Bessen, 2018). Finally, an analysis of the National Strategy on AI in India found that the adoption of new technologies is likely to occur in niches in the manufacturing and services sectors, limiting the potential for large-scale occupational change to address societal inequalities (Hammer & Karmakar, 2021).
The following table summarizes the key characteristics of the studies discussed in this subsection, highlighting their focus on specific aspects of task content, occupational change, or worker transitions.

5 Discussion

The synthesis of evidence presented in this review reveals a complex and contingent landscape of macroeconomic labor market transformations driven by artificial intelligence and autonomous systems. Taken together, the findings from the 62 included studies converge on several core insights while also exposing significant areas of disagreement and unresolved tension. The central narrative that emerges is one of a fundamental restructuring of work, where the net aggregate effects on employment and wages are not technologically predetermined but are profoundly shaped by the type of technology deployed, the institutional and policy context in which it is adopted, and the pre-existing structure of the labor market. This discussion section interprets these collective findings, explores their theoretical and practical implications, acknowledges the limitations of the review process, and proposes directions for future research that can address the gaps and inconsistencies identified in the literature.
A consistent pattern that emerges across studies is the critical distinction between automation technologies that substitute for human labor and those that augment it. The evidence on industrial robots, for instance, predominantly points to displacement effects that reduce employment and wages for lower-skilled workers, particularly in manufacturing sectors of advanced economies (Acemoglu & Restrepo, 2019b) (Aghion et al., 2019). In contrast, the literature on generative AI reveals a more nuanced picture where the technology can both substitute for and complement human labor, depending on the task, the worker's skill level, and the degree of exposure (Freund & Mann, 2026) (Johnston & Makridis, 2025). This duality is not merely a technical detail but a fundamental determinant of macroeconomic outcomes. The task-based framework, as formalized by Acemoglu and Restrepo (Acemoglu & Restrepo, 2018) (Acemoglu & Restrepo, 2019a), provides a powerful lens for understanding this tension, positing that automation creates a displacement effect that reduces labor demand, counterbalanced by a productivity effect and the creation of new tasks that reinstate labor. The net outcome depends on the rela-

tive strength of these forces, which the empirical evidence shows varies significantly across technologies, sectors, and time periods. This suggests that theoretical models must move beyond simple dichotomies of "automation" versus "augmentation" and instead incorporate the granular, task-level heterogeneity that characterizes real-world technological change.

The implications of these findings for theory are substantial. First, the evidence challenges the historical optimism of the "lump of labor fallacy," which held that technological progress inevitably creates as many jobs as it destroys. While the literature does not support claims of mass technological unemployment in the aggregate, it does demonstrate that the adjustment process can be slow, painful, and highly unequal (Acemoglu & Restrepo, 2019a) (Kogan et al., 2021). The displacement effect has accelerated in recent decades, particularly in manufacturing, while the reinstatement effect has weakened, leading to slower employment growth and a declining labor share of income (D. H. Autor & Salomons, 2018). This suggests that the historical pattern of rapid job creation in new sectors may not automatically repeat itself with AI, especially if the new tasks created are few in number or require skills that displaced workers do not possess. Second, the findings underscore the need for theoretical models to incorporate distributional dynamics more explicitly. The evidence consistently shows that AI and automation exacerbate inequalities along multiple dimensions, including skill, gender, geography, and income (Acemoglu & Restrepo, 2021) (Gmyrek et al., 2025) (Zhang et al., 2025). Standard macroeconomic models that focus solely on aggregate outcomes may therefore miss the most consequential social impacts of these technologies. Third, the literature highlights the importance of institutional context. The divergent findings between US and European studies on the employment effects of robots, for example, suggest that labor market institutions, social safety nets, and industrial relations systems play a crucial mediating role (Acemoglu & Restrepo, 2019b) (Klenert et al., 2022). Theoretical frameworks must therefore be sensitive to these institutional variations rather than assuming universal technological effects.

The practical implications of this synthesis are equally profound, particularly for policymakers, educators, and business leaders. For policymakers, the evidence strongly suggests that proactive intervention is essential to harness the productivity gains of AI while mitigating its disruptive effects. The finding that AI's macroeconomic productivity contributions are modest, estimated at 0.25 to 0.6 percentage points of annual total-factor productivity growth (Filippucci et al., 2024), implies that claims of a dramatic productivity boom are overstated. This should temper expectations and focus policy attention on the distributional challenges rather than on aggregate growth alone. Specific policy levers that emerge from the literature include investments in education and training, particularly in digital skills and AI literacy, which have been shown to moderate the negative effects of displacement risk (Chen et al., 2022) (Stephany & Teutloff, 2026). Social protection systems, including unemployment insurance and active labor market policies, are critical for supporting workers during transitions, especially given the evidence that occupational mobility networks are often limited and that displaced workers may face long-term unemployment (del Rio-Chanona, Mealy, et al., 2021). Competition policy also matters, as the concentration of AI capabilities in a small number of large firms could exacerbate inequality and reduce the diffusion of productivity gains (Aghion et al., 2019). Furthermore, the evidence on gender and geographic disparities calls for targeted interventions to ensure that women and workers in smaller cities or developing economies are not left behind (Frank et al., 2017) (Gmyrek et al., 2025). For educators, the rising demand for both technical AI skills and complementary human skills, such as social and entrepreneurial skills, suggests that curricula must evolve to prepare students for a labor market where human-machine collaboration is the norm (Prüfer & Prüfer, 2019) (Stephany & Teutloff, 2026). For business leaders, the findings on firm-level adoption indicate that AI can boost productivity and employment in some contexts, but that the benefits are not automatic and depend on strategic implementation, workforce training, and a clear understanding of where AI complements rather than substitutes for human labor (Cui et al., 2026) (Yu & Qi, 2024).

Despite the rigor of the systematic review methodology, several limitations must be acknowledged. First, the search strategy, while comprehensive across five major databases, may have introduced a language bias by restricting inclusion to English-language publications. This is a significant concern given that AI adoption and labor market dynamics in non-English-speaking countries, particularly in Asia, Africa, and Latin America, may differ substantially from those in the Anglophone world. The exclusion of studies in languages such as Chinese, Spanish, or French may have omitted valuable evidence on how these technologies are reshaping labor markets in diverse institutional and cultural contexts. Second, the inclusion of preprints and working papers, while reducing publication bias, introduces variability in quality, as these works have not undergone formal peer review. Some of the findings reported in preprints may be revised or contradicted in later peer-reviewed versions. Third, the exclusion of studies focused on a single industry or occupation without generalizable macroeconomic conclusions may have omitted micro-level evidence that could inform macroeconomic understanding. For example, detailed case studies of AI adoption in specific sectors like healthcare or finance could provide insights into the mechanisms of task restructuring that are not captured in broader econometric analyses. Fourth, the subjective nature of the eligibility assessment, despite the use of two independent reviewers, may introduce some degree of selection bias. The interpretation of what constitutes a “macroeconomic implication” is inherently subjective, and different reviewers might have made different inclusion decisions. Fifth, the review is limited by the quality and scope of the primary studies themselves. Many of the included studies rely on observational data and correlational analyses, making it difficult to establish causal relationships. The use of instrumental variables, difference-in-differences, and other quasi-experimental methods in some studies strengthens causal inference, but these methods are not immune to threats from omitted variables or measurement error. Finally, the rapid pace of technological change means that the literature reviewed here, which includes studies published up to October 2024, may already be outdated in some respects. The capabilities of generative AI, in particular, are evolving so quickly that the labor market impacts documented in studies from 2023 or 2024 may not fully capture the effects of more advanced systems.

The limitations of the review process point to several promising directions for future research. First, there is a clear need for more cross-country comparative studies that can systematically examine how institutional factors, such as labor market regulations, education systems, and social protection frameworks, mediate the labor market impacts of AI. The existing evidence is heavily skewed towards the United States and a small number of European countries, with far less attention paid to developing economies, where the informal sector is large and upskilling infrastructure is limited (Hammer & Karmakar, 2021) (Pizzinelli, 2023). Future research should explore how AI adoption affects labor markets in these contexts, where the dynamics of displacement and reinstatement may operate very differently. Second, the literature would benefit from more longitudinal studies that track the same workers and firms over time to understand the dynamics of adjustment. Cross-sectional studies can identify correlations but cannot reveal the processes of occupational switching, skill upgrading, or wage recovery that occur over years. The few longitudinal studies included in this review, such as those using administrative data from Denmark (Humlum & Vestergaard, 2025) or the US (Brynjolfsson, Chandar, & Chen, 2025), provide valuable insights and should be replicated in other settings. Third, there is a need for more research on the specific mechanisms through which AI affects labor market outcomes. While the task-based framework is theoretically appealing, empirical measures of task exposure to AI are still crude and often rely on expert judgment or patent text analysis. Future research should develop more granular, data-driven measures of task automation and augmentation that can be linked to individual-level outcomes. Fourth, the intersection of AI with other structural changes, such as globalization, demographic shifts, and the green transition, is understudied. The labor market impacts of AI do not occur in a vacuum, and understanding how these forces interact is crucial for policy design. Fifth, future research should explore the role of firm-level strategies and market structure in shaping aggregate outcomes. The evidence suggests that AI adoption is concentrated in certain firms and sectors, and that the competitive dynamics of product and labor markets influence how the gains from AI are distributed (Damioli et al., 2021) (Eisfeldt et al., 2023). More research is needed on how market power, patenting strategies, and firm organization affect the translation of AI-driven productivity gains into employment and wage outcomes. Finally, there is a need for more research on the psychological and social dimensions of AI-driven labor market change. The evidence on fear of automation and its effects on job satisfaction and well-being (Dekker et al., 2017) (Schwabe & Castellacci, 2020) suggests that the anticipation of displacement can have real economic consequences, even in the absence of actual job loss. Future research should explore how these perceptions are formed, how they affect labor supply and wage bargaining, and how policy interventions can mitigate their negative effects.

6 Conclusion

This systematic review synthesized a fragmented body of theoretical and empirical research to clarify how artificial intelligence and autonomous systems reshape macroeconomic labor markets. Our analysis confirms that the net aggregate effects on employment and wages are not technologically predetermined but are profoundly contingent on the type of technology, the institutional context, and the pre-existing structure of the labor market. The central tension between displacement and reinstatement effects, as formalized in task-based frameworks, provides a robust analytical lens, yet the empirical evidence reveals that this tension plays out differently across industrial robots, generative AI, and other automation technologies. The review contributes to the field by demonstrating that the most consequential impacts of AI are distributional, systematically exacerbating inequalities along skill, gender, geographic, and income dimensions, rather than producing uniform aggregate outcomes.

The practical implications of these findings are clear: proactive policy intervention is essential to harness the productivity gains of AI while mitigating its disruptive social costs. Investments in digital skills and AI literacy, robust social protection systems, and targeted support for vulnerable groups, including women, early-career workers, and those in smaller cities or developing economies, are critical levers for shaping a more equitable future of work. The theoretical implication is that macroeconomic models must move beyond aggregate analyses to incorporate the granular, task-level heterogeneity and institutional mediation that characterize real-world technological change. Future research should prioritize cross-country comparative studies, longitudinal tracking of worker transitions, and the development of more granular measures of task automation and augmentation. The intersection of AI with globalization, demographic shifts, and the green transition also warrants urgent investigation. Ultimately, the future of work is not a technological destiny but a policy choice, and the evidence synthesized here provides a foundation for making that choice wisely.

References

Benhamou, S. (2020). Artificial intelligence and the future of work. *Revue d'économie industrielle.*

Agrawal, A., Gans, J., & Goldfarb, A. (2019). *The economics of artificial intelligence: An agenda* (tech. rep.). books.google.com.

Autor, D. (2015). Why are there still so many jobs? the history and future of workplace automation. *Journal of economic perspectives.*

Goldin, C., & Katz, L. (2018). The race between education and technology. *Inequality in the 21st Century.*

Chowdhury, M., & Quaddus, M. (2017). The impact of artificial intelligence on the future of work: A systematic review. *Technological Forecasting and Social Change, 2017.*

Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from us labor markets. *Journal of political economy.*

Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative ai at work. *The Quarterly Journal of Economics.*

Page, M., McKenzie, J., Bossuyt, P., et al. (2021). The prisma 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, *372*, n71.

Acemoglu, D., & Restrepo, P. (2018). Artificial intelligence, automation, and work. *The Economics of Artificial Intelligence: An Agenda.*

Acemoglu, D., & Restrepo, P. (2019a). Automation and new tasks: How technology displaces and reinstates labor. *Journal of economic perspectives.*

Acemoglu, D., & Restrepo, P. (2019b). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, *128*, 2188–2244.

Acemoglu, D., & Restrepo, P. (2021). Tasks, automation, and the rise in us wage inequality. *SSRN Electronic Journal.*

Acemoglu, D., Autor, D. H., Hazell, J., & Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, *40*, S293–S340.

Acemoglu, D. (2024). The simple macroeconomics of ai. *SSRN Electronic Journal.*

Aghion, P., Antonin, C., & Bunel, S. (2019). Artificial intelligence, growth and employment: The role of policy. *Economie et Statistique / Economics and Statistics.*

Aly, H. (2022). Digital transformation, development and productivity in developing countries: Is artificial intelligence a curse or a blessing? *Review of Economics and Political Science.*

Autor, D. H., & Salomons, A. (2018). Is automation labor-displacing? productivity growth, employment, and the labor share.

Berger, P., Cai, W., Qiu, L., & Shen, C. (2024). Employer and employee responses to generative ai: Early evidence. *Available at SSRN 4874061.*

Bessen, J. (2016). How computer automation affects occupations: Technology, jobs, and skills.

Bessen, J. (2018). Artificial intelligence and jobs: The role of demand, 291–307.

Brynjolfsson, E., Chandar, B., & Chen, R. (2025). *Canaries in the coal mine?: Six facts about the recent employment effects of artificial intelligence* (tech. rep.). digitaleconomy.stanford.edu.

Carbonero, F., Davies, J., Ernst, E., Fossen, F. M., Samaan, D. K., & Sorgner, A. (2023). The impact of artificial intelligence on labor markets in developing countries: A new method with an illustration for lao pdr and urban viet nam. *Journal of Evolutionary Economics*, 1–30.

Cazzaniga, M., Jaumotte, F., Li, L., Melina, G., Panton, A. J., Pizzinelli, C., Rockall, E., & Tavares, M. (n.d.). Gen-ai: Artificial intelligence and the future of work.

Chen, N., Li, Z., & Tang, B. (2022). Can digital skill protect against job displacement risk caused by artificial intelligence? empirical evidence from 701 detailed occupations. *PLOS ONE*, *17*.

Chivu, S. I. (2022). The macroeconomic impact of artificial intelligence. *International Journal of Sustainable Economies Management*.

Christenko, A. (2022). Automation and occupational mobility: A task and knowledge-based approach. *Technology in Society*.

Cortes, P., & Pan, J. (2019). Gender, occupational segregation, and automation. *Economics Studies at Brookings*.

Cortes, G. M., Jaimovich, N., Jaimovich, N., Siu, H. E., & Siu, H. E. (2017). Disappearing routine jobs: Who, how, and why? *Journal of Monetary Economics*, *91*, 69–87.

Cui, K., Demirer, M., Jaffe, S., Musolff, L., et al. (2026). The effects of generative ai on high-skilled work: Evidence from three field experiments with software developers. *Management Science*.

Damioli, G., Roy, V. V., & Vértesy, D. (2021). The impact of artificial intelligence on labor productivity. *Eurasian Business Review*, *11*, 1–25.

DeCanio, S. (2016). Robots and humans–complements or substitutes? *Journal of Macroeconomics*.

Dekker, F., Salomons, A., & Waal, J. (2017). Fear of robots at work: The role of economic self-interest. *Socio-Economic Review*.

Doubková, M., & Magdin, M. (2026). The impact of ai on the labour market. *Management amp; Marketing*.

Eisfeldt, A., Schubert, G., & Zhang, M. B. (2023). *Generative ai and firm values* (tech. rep.). National Bureau of Economic Research.

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). Gpts are gpts: Labor market impact potential of llms. *Science*.

Filippucci, F., Gal, P., & Schief, M. (2024). Miracle or myth? assessing the macroeconomic productivity gains from artificial intelligence. *OECD Artificial Intelligence Papers*.

Frank, M., Sun, L., Cebrian, M., Youn, H., & Rahwan, I. (2017). Small cities face greater impact from automation. *Journal of the Royal Society Interface*, *15*.

Freund, L., & Mann, L. (2026). *Job transformation, specialization, and the labor market effects of ai* (tech. rep.). econstor.eu.

Georgieff, A., & Hyee, R. (2022). Artificial intelligence and employment: New cross-country evidence. *Frontiers in Artificial Intelligence*, *5*.

Gmyrek, P., Berg, J., & Bescond, D. (2023). *Generative ai and jobs: A global analysis of potential effects on job quantity and quality* (tech. rep.). papers.ssrn.com.

Gmyrek, P., Berg, J., & Bescond, D. (2025). Generative ai and jobs: An analysis of potential effects on global employment. *Gospodarka Narodowa*.

Guliyev, H. (2023). Artificial intelligence and unemployment in high-tech developed countries: New insights from dynamic panel data model. *Research in Globalization*, *7*, 100140.

Hammer, A., & Karmakar, S. (2021). Automation, ai and the future of work in india. *Employee Relations: The International Journal*, *ahead-of-print*.

Hui, X., Reshef, O., & Zhou, L. (2024). The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market. *SSRN Electronic Journal*.

Humlum, A., & Vestergaard, E. (2025). *Large language models, small labor market effects* (tech. rep.). nber.org.

Jetha, A., Shamaee, A., Bonaccio, S., Gignac, M., Tucker, L. B., Tompa, E., Bültmann, U., Norman, C. D., Banks, C., & Smith, P. M. (2021). Fragmentation in the future of work: A horizon scan examining the impact of the changing nature of work on workers experiencing vulnerability. *American Journal of Industrial Medicine*, *64*, 649–666.

Johnston, A., & Makridis, C. (2025). The labor market effects of generative ai: A difference-in-differences analysis of ai exposure. *Available at SSRN 5375017*.

King, B., Hammond, T., et al. (2017). Disruptive technology: Economic consequences of artificial intelligence and the robotics revolution. *Journal of Strategic Information Systems*.

Klenert, D., Fernández-Macías, E., & Antón, J. (2022). Do robots really destroy jobs? evidence from europe. *Economic and Industrial Democracy*, *44*, 280–316.

Kogan, L., Papanikolaou, D., Schmidt, L., & Seegmiller, B. (2021). *Technology-skill complementarity and labor displacement: Evidence from linking two centuries of patents with occupations* (tech. rep.). bryanseegmiller.com.

Kogan, L., Papanikolaou, D., Schmidt, L., & Seegmiller, B. (2023). *Technology and labor displacement: Evidence from linking patents with worker-level data* (tech. rep.). nber.org.

Kuzior, A., Marszałek-Kotzur, I., Sansyzbayeva, K., & Lukács, E. (2025). From ai vibrancy to labour market outcomes: Testing displacement across education groups. *Economics amp; Sociology*.

Manyika, J., Chui, M., Miremadi, M., Bughin, J., et al. (2017). A future that works: Ai, automation, employment, and productivity. *McKinsey Global Institute Research, Tech. Rep*.

Mutascu, M. (2021). Artificial intelligence and unemployment: New insights. *Economic Analysis and Policy*.

Novakova, L. (2020). The impact of technology development on the future of the labour market in the slovak republic. *Technology in Society*.

Olaniyi, O., Ezeugwa, F., Okatta, C., et al. (2024). *Dynamics of the digital workforce: Assessing the interplay and impact of ai, automation, and employment policies* (tech. rep.). . . . Policies.

Pizzinelli, C. (2023). Labor market exposure to ai: Cross-country differences and distributional implications. *IMF Working Papers*.

Prüfer, J., & Prüfer, P. (2019). Data science for entrepreneurship research: Studying demand dynamics for entrepreneurial skills in the netherlands. *Small Business Economics*, *55*, 651–672.

Rawashdeh, A. (2025). The consequences of artificial intelligence: An investigation into the impact of ai on job displacement in accounting. *Journal of Science and Technology Policy Management*.

del Rio-Chanona, R., Mealy, P., et al. (2021). Occupational mobility and automation: A data-driven network model. *Journal of the Royal Society Interface*.

Schwabe, H., & Castellacci, F. (2020). Automation, workers' skills and job satisfaction. *Plos one*.

Somjai, S., Jermsittiparsert, K., & Chankoson, T. (2020). Determining the initial and subsequent impact of artificial intelligence adoption on economy: A macroeconomic survey from asean. *Journal of Intelligent Fuzzy Systems*, *39*, 5459–5474.

Sorgner, A. (2017). The automation of jobs: A threat for employment or a source of new entrepreneurial opportunities? *11*, 37–48.

Stephany, F., & Teutloff, O. (2026). Ai skills wanted: How ai technologies create demand for skilled workers. *Degrees of Change: What AI Means for Education*.

Vermeulen, B., Kesselhut, J., Pyka, A., & Saviotti, P. (2018). The impact of automation on employment: Just the usual structural change? *Sustainability*, *10*, 1661–1688.

Webb, M. (2019). The impact of artificial intelligence on the labor market. *Economics of Innovation eJournal*.

Yang, C. (2022). How artificial intelligence technology affects productivity and employment: Firm-level evidence from taiwan. *Research Policy*.

Yu, J., & Qi, C. (2024). The impact of generative ai on employment and labor productivity. *Review of Business*.

Zhang, W., Han, J., Işık, C., Song, J., Liu, X., Dembińska, I., et al. (2025). Inequality in the digital economy: The impact of artificial intelligence on income gap—an empirical analysis based on county-level data from china. *Cities*.

	STUDY DESIGN
(Acemoglu & Restrepo, 2018)	task-based framework
(Acemoglu & Restrepo, 2019a)	Framework and empirical decomposition using industry-level data
(Acemoglu & Restrepo, 2019b)	Observational study using variation in exposure to robots across c
(Acemoglu & Restrepo, 2021)	Quantitative analysis using a conceptual framework and regression
(Acemoglu et al., 2022)	Observational study using establishment-level data on online vacan
(Acemoglu, 2024)	Task-based model of AI's effects, working through automation and
(Aghion et al., 2019)	Survey and empirical analysis
(Aly, 2022)	The paper uses the Digital Evolution Index (DEI) and feasible gen
(D. H. Autor & Salomons, 2018)	Empirical estimation using data on 28 industries for 18 OECD cou
(Berger et al., 2024)	Generalized difference-in-differences approach
(Bessen, 2016)	General occupational model analyzing detailed occupations since 1
(Bessen, 2018)	A simple model of demand is presented
(Brynjolfsson, Chandar, & Chen, 2025)	Observational study using high-frequency administrative data
(Carbonero et al., 2023)	New methodology to translate existing measures of AI impacts fro
(Cazzaniga et al., n.d.)	
(Chen et al., 2022)	fixed-effects modeling, heterogeneity analyzing and moderation eff
(Chivu, 2022)	Comparative study
(Christenko, 2022)	Task and knowledge-based occupational mobility network analysis
(P. Cortes & Pan, 2019)	Descriptive and speculative framing paper
(G. M. Cortes et al., 2017)	General neoclassical model of the labor market featuring endogeno
(**Cruz2021ArtificialIA**)	Multiple regression analysis
(Cui et al., 2026)	Randomized controlled trials (field experiments) at Microsoft, Acc
(Damioli et al., 2021)	Observational study using patent data and company-level analysis
(DeCanio, 2016)	Cross-sectional analysis using Houthakker's method
(Dekker et al., 2017)	Cross-sectional survey
(Doubková & Magdin, 2026)	
(Eisfeldt et al., 2023)	Artificial Minus Human portfolios
(Eloundou et al., 2024)	Framework for evaluating potential impacts of LLMs on work by c
(Filippucci et al., 2024)	Micro-to-macro framework combining micro-level performance gai
(Frank et al., 2017)	
(Freund & Mann, 2026)	General-equilibrium model
(Georgieff & Hyee, 2022)	Cross-country observational study
(Gmyrek et al., 2023)	
(Gmyrek et al., 2025)	Analiza globalna z wykorzystaniem modelu GPT-4 do szacowania
(Guliyev, 2023)	Dynamic panel data and GMM-system estimation
(Hammer & Karmakar, 2021)	Critical assessment of the National Strategy on AI through a secto
(Hui et al., 2024)	Observational study examining short-term effects of generative AI
(Humlum & Vestergaard, 2025)	difference-in-differences
(Jetha et al., 2021)	horizon scan
(Johnston & Makridis, 2025)	difference-in-differences analysis
(King, Hammond, et al., 2017)	probability-weighted methodology
(Klenert et al., 2022)	Observational study using industry-level data and fixed-effects tec
(Kogan et al., 2021)	Panel data analysis with textual analysis of patent documents and
(Kogan et al., 2023)	Observational study using textual analysis of patents and job task
(Kuzior et al., 2025)	Two-way fixed- and random-effects models, employing Box–Cox/lo
(Manyika et al., 2017)	
(Mutascu, 2021)	Panel threshold and GMM-system estimations
(Novakova, 2020)	scenario analysis using Cobb-Douglas production function
(Olaniyi et al., 2024)	Quantitative research design using a structured questionnaire
(Pizzinelli, 2023)	Cross-country analysis using worker-level microdata
(Prüfer & Prüfer, 2019)	Observational study using data science methods on a large dataset
(Rawashdeh, 2025)	Cross-sectional survey disseminated through Facebook, using snow
(del Rio-Chanona, Mealy, et al., 2021)	data-driven network model
(Schwabe & Castellacci, 2020)	Instrumental variables (IV) analysis using variation in the pace of
(Somjai et al., 2020)	Structural questionnaire survey, descriptive statistics, structural ec

	THEORETICAL FRAMEWORK / MODEL
(Acemoglu & Restrepo, 2018)	Task-based framework
(Acemoglu & Restrepo, 2019a)	Task-based framework: allocation of tasks to capital and labor (task cont
(Acemoglu & Restrepo, 2019b)	Theoretical model showing robots may reduce employment and wages
(Acemoglu, 2024)	Task-based model of AI's effects, working through automation and task c
(Aghion et al., 2019)	Survey and empirical analysis
(D. H. Autor & Salomons, 2018)	Acemoglu and Restrepo (2018b)
(Bessen, 2018)	A simple model of demand
(G. M. Cortes et al., 2017)	General neoclassical model of the labor market featuring endogenous par
(**Cruz2021ArtificialIA**)	Multiple regression analysis
(Filippucci et al., 2024)	Multi-sector general equilibrium model with input-output linkages
(Georgieff & Hyee, 2022)	Cross-country observational study
(Gmyrek et al., 2025)	GPT-4 model
(Guliyev, 2023)	Displacement effect
(Johnston & Makridis, 2025)	Difference-in-differences analysis
(Klenert et al., 2022)	Fixed-effects techniques
(Kuzior et al., 2025)	Two-way fixed- and random-effects models
(Mutascu, 2021)	Panel threshold and GMM-system estimations
(Novakova, 2020)	Cobb-Douglas production function
(Vermeulen et al., 2018)	Evolutionary economic model of multisectoral structural change united w
(Webb, 2019)	Method using overlap between job task descriptions and patent text to m

	PRIMARY EMPIRICAL METHODOLOGY / MODEL
(Acemoglu & Restrepo, 2018)	Labor share of national income
(Bessen, 2016)	Within-occupation wage inequality
(Brynjolfsson, Chandar, & Chen, 2025)	Employment by age/experience
(Cazzaniga et al., n.d.)	Gender and Education
(Chen et al., 2022)	Occupational wage and employment
(DeCanio, 2016)	Aggregate wages
(Dekker et al., 2017)	Economic self-interest and perceptions
(Frank et al., 2017)	Urban inequality (cities of different sizes)
(Freund & Mann, 2026)	Wage inequality (within-occupation and overall)
(Gmyrek et al., 2025)	Income level (between countries) and gender
(Hammer & Karmakar, 2021)	Structural inequalities in India
(Jetha et al., 2021)	Social and health inequities
(Acemoglu & Restrepo, 2021)	Wage inequality (education wage differentials)
(Johnston & Makridis, 2025)	Wage inequality and employment
(Kogan et al., 2021)	Earnings (wages and employment)
(Kogan et al., 2023)	Earnings
(Pizzinelli, 2023)	Earnings inequality (across earnings distribution)
(del Rio-Chanona, Mealy, et al., 2021)	Wage inequality (between low-wage and high-wage occupations)
(Schwabe & Castellacci, 2020)	Skill level (low-skilled vs high-skilled)
(Acemoglu, 2024)	Labor income inequality; capital-labor income gap
(Zhang et al., 2025)	Urban-rural income gap in China
(Aghion et al., 2019)	Employment by education level
(Aly, 2022)	Gender inequality (employment)
(D. H. Autor & Salomons, 2018)	Labor share of output

	ASPECT FOCUSED ON
(Brynjolfsson, Chandar, & Chen, 2025)	Employment by age/experience within AI-exposed occupations
(Christenko, 2022)	Occupational mobility to alternative jobs
(P. Cortes & Pan, 2019)	Occupation-level automation risk and transition patterns by gende
(Eloundou et al., 2024)	Potential exposure of occupations and tasks to LLMs
(Freund & Mann, 2026)	Job transformation through task content change
(Hui et al., 2024)	Task automation versus job automation
(Prüfer & Prüfer, 2019)	Demand for entrepreneurial and digital skills across professions
(del Rio-Chanona, Mealy, et al., 2021)	Network effects on long-term unemployment
(Sorgner, 2017)	Transitions to self-employment in response to automation risk
(Stephany & Teutloff, 2026)	Demand for AI skills (technical and literacy)
(Yang, 2022)	Workforce composition by educational qualifications
(Carbonero et al., 2023)	Occupational exposure to AI in developing economies
(Hammer & Karmakar, 2021)	Adoption and occupational change in India