Academic Journal of Education
ISSN: 3080-9940 (Print), ISSN: 3080-9959 (Online) | Vol. 2, No. 3, 2026

Artificial Intelligence Adoption, Labor Market Inequality, and Firm Transformation: A Systematic Narrative Review

Bo Wang, Chi Gong *

Graduate School of Global Business, Kyonggi University, Suwon 16227, South Korea

*** Corresponding author:** Chi Gong

Abstract: This article provides a systematic narrative review of recent literature on artificial intelligence adoption, labor market inequality, and firm transformation. Drawing on theoretical models, firm-level empirical research, quasi-experimental policy studies, and international organization reports, the paper examines how artificial intelligence changes economic activity through three connected mechanisms: task reallocation, productivity complementarity, and organizational capability formation. The review argues that the economic consequences of artificial intelligence cannot be understood through a simple substitution-versus-complementarity distinction. Instead, artificial intelligence reshapes the boundary between labor and capital, changes the relative value of different tasks, and creates uneven firm-level capacities to convert digital tools into innovation, sustainability performance, and competitive advantage. The literature suggests that aggregate productivity effects may be meaningful but more modest than optimistic public forecasts imply; wage inequality effects are ambiguous and depend on task exposure, adoption intensity, worker skill, and institutional context; and firm-level benefits are concentrated among organizations with complementary assets, data capability, and absorptive capacity. The article contributes by integrating macroeconomic, labor-market, and corporate-sustainability perspectives into a unified review framework. It concludes by identifying research gaps concerning developing economies, small and medium-sized enterprises, cross-border value chains, long-term employment adjustment, and the governance of AI-enabled inequality.

Keywords: Artificial Intelligence; Generative AI; Labor Market; Inequality; Firm Innovation; ESG; Green Innovation; Task Model; Global Business.

1. Introduction

Artificial intelligence (AI) has moved from a specialized technological field into a general business and policy concern. The rapid diffusion of generative AI, especially large language models (LLMs), has intensified debate over whether AI will raise productivity, replace labor, reduce knowledge barriers, or deepen social and economic inequality. For global business, these questions are not abstract. Multinational firms, export-oriented manufacturers, platform companies, and service providers increasingly depend on digital tools to coordinate production, manage customer relationships, write code, analyze markets, and meet environmental, social, and governance requirements. AI adoption therefore affects not only firms' internal efficiency but also the distribution of economic gains across workers, firms, sectors, and countries. [1-2]

The literature reviewed in this article shows that AI should be treated as a task-transforming technology rather than a single-purpose tool. In the task-based view, production consists of many tasks that can be performed by workers of different skills or by capital. AI changes economic outcomes when it automates some tasks, complements workers in other tasks, raises the productivity of capital, or creates new tasks that did not previously exist. This distinction is important because different types of technological change have different implications for productivity, wages, employment, and firm strategy. A technology that improves workers' performance in existing tasks may raise wages, while a technology that allows capital to perform tasks previously done by workers may weaken labor demand. A technology that creates valuable new tasks may expand employment, but a technology that creates socially harmful or rent-seeking tasks may increase measured activity without improving welfare.

The public discussion of AI often oscillates between two simplified claims. One claim is that AI will generate a productivity boom comparable to electricity or the internet. The other claim is that AI will cause large-scale unemployment and a sharp decline in the bargaining power of workers. The reviewed literature suggests a more cautious interpretation. Acemoglu's macroeconomic analysis argues that likely productivity gains from AI may be nontrivial but modest when measured through task-level cost savings and the share of tasks actually affected. At the same time, labor-market studies show that the distributional effects of AI may differ from earlier automation because AI exposure is not concentrated only in routine middle-skill occupations. Generative AI affects professional, clerical, analytical, creative, and managerial work, which means that its inequality effects may operate through new channels.

This review also considers firm-level and sustainability-oriented studies. Babina et al. connect AI investment to firm growth and product innovation, while Shen et al. use a refined LLM-based measurement strategy to examine corporate ESG performance. Mijit et al. evaluate China's Artificial Intelligence Innovation and Development Pilot Zones and find positive effects on firms' green innovation. These studies suggest that AI adoption is not merely a labor-saving process. It can also support innovation, operational efficiency, environmental upgrading, and strategic differentiation. However, the benefits are uneven. Large, technology-intensive, capital-intensive, or data-rich firms appear better positioned to capture the gains, raising the possibility of a digital ESG divide.

The purpose of this paper is to synthesize these debates into a coherent literature review for scholars of global business and technology management. It asks three questions. First, what mechanisms does the literature identify through which AI changes labor demand and inequality? Second, what evidence exists regarding AI's effects on firm growth, innovation, ESG performance, and green transformation? Third, what research gaps remain for scholars in global business, especially in relation to developing economies, cross-border production, and policy governance? The article does not conduct new empirical analysis. Instead, it organizes recent research into an integrated conceptual framework and identifies directions for future research. [3]

2. Review Method and Scope

This paper adopts a systematic narrative review approach. A systematic narrative review differs from a meta-analysis because it does not statistically combine effect sizes across comparable studies. It is suitable when the literature includes diverse theoretical models, empirical designs, institutional reports, and policy studies. The reviewed literature covers macroeconomic modelling, task-based labor economics, AI exposure measurement, wage and wealth inequality, generative AI in workplaces, firm innovation, corporate ESG performance, green innovation policy, and modern causal-inference methods such as staggered difference-in-differences and shift-share designs.

The inclusion logic is thematic rather than database-driven. Studies were selected for their relevance to AI adoption, labor-market inequality, firm transformation, sustainability outcomes, or causal research design. Duplicate or near-duplicate versions of the same research were treated as one intellectual contribution where appropriate. Methodological papers by Callaway and Sant'Anna, de Chaisemartin and D'Haultfoeuille, Goodman-Bacon, and Borusyak, Hull, and Jaravel are discussed not because they are directly about AI, but because they shape the credibility of empirical AI studies that rely on staggered policy adoption, occupational exposure, or regional variation.

The review proceeded in three steps. First, each document was classified by its primary level of analysis: macroeconomic, labor-market, firm-level, policy-level, or methodological. Second, the main causal mechanism was identified: task automation, task complementarity, capital deepening, product innovation, operational efficiency, ESG upgrading, or measurement improvement. Third, the findings were compared across studies to identify convergence, disagreement, and unresolved questions. This procedure is appropriate for a literature in which AI adoption is still recent and where many influential contributions are working papers or policy reports rather than long-established journal articles.

A limitation of this review is that it is organized as a focused thematic synthesis rather than an exhaustive bibliometric review. Several influential studies examine the United States, OECD countries, or Chinese listed firms, which means that conclusions may not transfer directly to small firms, informal labor markets, or lower-income economies. These limitations are not weaknesses of the paper if they are stated clearly; rather, they help define the boundary of the argument.

Table 1. Thematic structure of the reviewed literature

Theme	Representative sources	Main contribution
Task reallocation	Acemoglu; Acemoglu, Kong & Restrepo; Acemoglu & Loebbing	AI shifts tasks between labor and capital; effects depend on comparative advantage and task complexity.
Productivity complementarity	Brynjolfsson, Li & Raymond; Georgieff	AI can raise worker productivity and may reduce performance dispersion within occupations.
Inequality trade-off	Rockall, Tavares & Pizzinelli; Cazzaniga et al.; Acemoglu	Wage inequality may fall in some settings, while wealth inequality and capital income may rise.
Firm capability	Babina et al.; Shen et al.	AI adoption creates value when combined with data, talent, routines, and absorptive capacity.
Green transformation	Mijit et al.; Shen et al.	AI can improve green innovation and ESG performance, but effects are stronger for capable firms.
Research design	Callaway & Sant'Anna; de Chaisemartin & D'Haultfoeuille; Goodman-Bacon; Borusyak et al.	Modern causal methods are necessary for staggered adoption, heterogeneous effects, and exposure designs.

3. Conceptual Foundations: AI as a Task-Transforming Technology

The strongest theoretical foundation in the reviewed literature is the task model of technological change. Acemoglu, Kong, and Restrepo explain that production requires a set of tasks, and each task can be assigned to workers with different skills or to capital. The key idea is comparative advantage: workers and machines are not perfect substitutes across all activities. Some tasks are better performed by highly skilled labor, some by lower-skilled labor, and some by machines or algorithms. Technology affects labor demand by changing how tasks are allocated and by changing the productivity of the factors that perform them.

The task model is useful because it avoids a common mistake in AI discussion: treating AI as if it affects all workers in the same way. AI may automate a task previously done by a worker, but it may also make the same worker more productive. For example, an AI writing assistant can substitute for routine drafting in some contexts, but it can complement a professional who uses it to generate alternatives, summarize information, or accelerate communication. Similarly, AI-based prediction systems may replace some monitoring tasks while increasing demand for workers who can interpret, verify, and govern algorithmic outputs. [4-5]

Acemoglu and Loebbing extend this discussion through a model of automation and polarization. Their model emphasizes that automation does not necessarily begin with the lowest-skill tasks. It may focus on intermediate-complexity tasks because these tasks are profitable to automate: they are structured enough for machines to perform, yet expensive enough in labor terms to make automation attractive. This insight helps explain why earlier waves of automation were associated with employment and wage polarization. Middle-skill routine jobs were more exposed than both low-wage service jobs and high-skill abstract jobs.

Generative AI complicates the earlier polarization story. Unlike industrial robots or traditional software, generative AI can affect text, code, image, legal, educational, marketing, and managerial tasks. This means that exposure extends into occupations previously considered relatively protected by education and professional status. However, exposure does not equal displacement. A task may be exposed because AI can assist with it, not because AI can fully replace the worker. The central question is therefore whether AI reduces the demand for human labor in exposed tasks or raises the productivity and value of workers who remain in those tasks.

Acemoglu's macroeconomic contribution adds an important discipline to this debate. He argues that aggregate productivity effects should be estimated by asking what share of tasks is affected and how large the task-level cost savings are. Under this logic, spectacular productivity forecasts may be overstated if they assume that AI can quickly transform a large fraction of economically important tasks. Many tasks are context-dependent, hard to learn, and difficult to evaluate with objective performance metrics. For these tasks, AI may be helpful but not transformative in the short run.

The implication for global business is that AI adoption should be analyzed at the task, firm, and institutional levels simultaneously. A multinational firm may adopt AI in customer service, logistics, translation, compliance, product design, and financial forecasting, but the productivity gains from each application depend on task structure, data quality, managerial routines, regulatory constraints, and worker adaptation. A task-based perspective therefore connects micro-level adoption decisions to macro-level outcomes.

4. AI, Productivity, and Macroeconomic Expectations

The first major theme in the literature concerns the size of AI's productivity effects. Optimistic forecasts often assume that AI will transform a large share of tasks and generate major cost savings. The reviewed literature is more conservative. Acemoglu's task-based macroeconomic analysis suggests that AI's contribution to total factor productivity may be meaningful but modest over a ten-year horizon. The argument is not that AI is unimportant. Rather, it is that macroeconomic transformation requires broad, reliable, and economically valuable task-level improvements. Early evidence from highly structured or easy-to-evaluate tasks may not generalize to complex organizational work.

This caution is significant for business strategy. Firms may overinvest in AI if they assume that any deployment will automatically create productivity gains. The literature suggests that productivity depends on complementarity between AI tools and organizational processes. AI can reduce the cost of information processing, prediction, writing, coding, and monitoring, but these savings become economically meaningful only when firms redesign workflows, train employees, and develop governance systems. Without such complementary changes, AI may produce local efficiency gains but limited aggregate performance improvement.

Brynjolfsson, Li, and Raymond provide important workplace evidence. Their study of generative AI at work shows that AI assistance can improve productivity, particularly for less experienced or lower-performing workers. This finding supports the idea that AI may compress performance differences within certain occupations. If an AI system provides suggestions based on high-quality examples, inexperienced workers may learn faster and make fewer errors. This mechanism differs from classic skill-biased technological change, where new technology often increases the returns to already highly skilled workers.

However, productivity improvement at the worker level does not automatically imply equal distribution of benefits. Firms may capture a large part of the surplus, especially when AI tools are owned by capital and deployed through employer-controlled platforms. Acemoglu emphasizes that AI may widen the gap between labor income and capital income even if some workers become more productive. Rockall, Tavares, and Pizzinelli similarly show that wealth inequality may increase because higher-income households are more likely to own capital and benefit from AI-driven returns.

Auer, Kopfer, and Sveda's analysis of generative AI exposure contributes to this discussion by modelling exposure and substitution possibilities. The key insight is that exposure must be separated from economic impact. A task can be technically exposed to AI but remain organizationally difficult to automate because of legal responsibility, customer trust, tacit knowledge, or integration costs. Conversely, a task with moderate technical exposure may be strongly affected if it is repeated at scale and embedded in a standardized workflow.

The literature therefore supports a balanced conclusion: AI is likely to raise productivity in specific tasks and firms, but aggregate gains depend on diffusion, organizational adaptation, and the value of affected tasks. For global business scholars, this means that AI should be studied as a process of capability building rather than a one-time technology shock. The firms that gain most may be those that combine AI with data infrastructure, process redesign, human capital, and strategic experimentation.

5. AI and Labor-Market Inequality

The second major theme concerns inequality. Earlier automation technologies were often associated with declining demand for routine middle-skill labor, wage polarization, and a falling labor share. The reviewed AI literature suggests that generative AI may produce a different distributional pattern. Because AI affects cognitive and language-based tasks, it may influence both high-wage and middle-wage occupations. Cazzaniga et al. argue that AI has the potential to reshape the global economy especially through labor markets, but the direction of inequality depends on whether AI substitutes for or complements workers.

Georgieff's OECD study examines AI exposure and wage inequality across nineteen OECD countries. The results provide no clear evidence that AI increased wage inequality between occupations during the studied period. At the same time, the paper reports some evidence that AI may be associated with lower wage inequality within occupations. This is consistent with workplace evidence suggesting that AI can help lower-performing workers improve more than high performers. If AI codifies good practices and makes them accessible, it may reduce performance dispersion among workers doing similar jobs.

This possibility should not be overstated. A reduction in within-occupation wage inequality does not necessarily mean that AI is egalitarian. Workers who cannot adapt to AI tools may exit occupations, creating a selection effect. Firms may also use AI to monitor workers more closely, standardize work, or reduce autonomy. In such cases, measured wage inequality might decline while job quality or bargaining power worsens. The literature therefore requires a broader view of inequality, including employment security, skill acquisition, surveillance, work intensity, and access to career progression.

Rockall, Tavares, and Pizzinelli provide a particularly important distinction between wage inequality and wealth inequality. Their calibrated task-based model suggests that AI could reduce wage inequality if it disrupts high-income tasks, but it may still increase wealth inequality because high-income households are more likely to own capital and benefit from AI-related returns. This distinction matters for policy. A narrow focus on wages could miss the distributional consequences of capital ownership, firm profits, and asset markets.

Acemoglu's analysis also warns that AI may increase the capital share. If AI systems substitute for labor in economically valuable tasks, the owners of AI capital may capture a larger share of income. Even when AI complements workers, the distribution of gains depends on labor-market institutions, bargaining power, intellectual property arrangements, and the availability of alternative employment. In weakly regulated labor markets, productivity gains may not translate into proportional wage gains.

He et al. extend the inequality discussion to the urban-rural dimension. This line of research highlights an important gap: AI may interact with regional inequality, digital infrastructure, educational resources, and industrial composition. Urban regions with better data infrastructure and more AI-intensive industries may benefit earlier, while rural workers and firms may face slower adoption or displacement without equivalent opportunities. For global business, this suggests that AI inequality is not only occupational but also spatial and infrastructural.

Humlum and Vestergaard's work on generative AI and labor markets adds another important point: adoption may be rapid at the level of awareness but uneven in practical intensity. Workers may experiment with AI tools, yet the effect on wages and employment depends on whether firms incorporate the tools into formal workflows and whether workers develop complementary skills. The phrase 'still waters, rapid currents' captures this tension: visible labor-market outcomes may appear stable while underlying task practices change quickly.

Taken together, the literature does not support a simple claim that AI will either increase or reduce inequality. The most defensible conclusion is conditional. AI may reduce performance gaps within some occupations, disrupt high-income cognitive tasks, and help less experienced workers. At the same time, it may increase capital income, concentrate gains among AI-ready firms and workers, widen regional divides, and intensify monitoring. Future research should therefore distinguish at least four types of inequality: between occupations, within occupations, between labor and capital, and between regions or firms.

6. Firm-Level AI Adoption, Growth, and Product Innovation

A third theme concerns the firm-level consequences of AI adoption. Babina et al. examine AI, firm growth, and product innovation. Their study is important because it shifts attention from general exposure to actual organizational investment. Firms do not benefit from AI simply because AI exists in the external environment. They benefit when they build capabilities, hire relevant talent, integrate AI into production or service processes, and use it to generate new products or improve existing offerings. [6-8]

The firm-level literature suggests that AI can support growth through several mechanisms. First, AI can reduce information-processing costs. Firms can analyze customer data, forecast demand, screen opportunities, and automate routine decision support. Second, AI can accelerate experimentation. In product development, generative AI can support ideation, prototyping, code generation, design variation, and market testing. Third, AI can improve coordination across departments and supply chains by providing faster access to knowledge and predictive insights. These mechanisms are especially relevant for global business because international firms face complex information environments, heterogeneous markets, and cross-border coordination challenges.

However, the benefits of AI adoption are likely to be uneven. The resource-based view of the firm suggests that competitive advantage depends on valuable, rare, inimitable, and organizationally embedded resources. AI tools may be widely available, but the ability to use them effectively depends on complementary assets such as proprietary data, skilled employees, digital infrastructure, managerial routines, and organizational learning. This is why the same AI technology can produce different results across firms.

The technology-organization-environment framework offers another useful lens. Technology factors include data quality, model reliability, cybersecurity, and system compatibility. Organizational factors include management support, employee skills, financial resources, and willingness to redesign processes. Environmental factors include regulation, competition, customer expectations, and industry norms. Shen et al. explicitly connect AI adoption and ESG outcomes to these theoretical perspectives, emphasizing that complementary assets and absorptive capacity influence whether firms can realize AI's sustainability potential.

AI adoption may also reshape market competition. Firms with strong AI capability may scale faster, personalize products more effectively, and reduce marginal costs. This could increase productivity and consumer welfare, but it could also intensify concentration if leading firms accumulate data advantages and attract scarce AI talent. Smaller firms may use cloud-based AI tools to overcome capability gaps, but they may also become dependent on platform providers. This creates a strategic tension between democratization and concentration.

For global business research, firm-level AI adoption should therefore be examined not only as technology investment but as organizational transformation. Important questions include how firms govern AI use across subsidiaries, how they protect data across jurisdictions, how they train employees in different cultural contexts, and how AI changes global value chain coordination. Current literature provides strong starting points, but more cross-country and industry-comparative evidence is needed.

7. Sectoral and Global Business Implications

Different sectors are likely to experience AI transformation through different channels. In finance and professional services, AI may automate document processing, risk screening, customer support, and compliance monitoring. In manufacturing, AI may improve quality inspection, predictive maintenance, supply-chain planning, and energy efficiency. In retail and marketing, AI may support demand forecasting, personalized recommendation, content generation, and customer segmentation. In education and training services, AI may support tutoring, feedback, and curriculum design. These sectoral differences matter because the same aggregate label of AI adoption hides very different task structures.

For multinational enterprises, AI may change the geography of work. Some service tasks previously offshored because of labor-cost differences may become automatable through generative AI. This could reduce the advantage of low-cost service locations. At the same time, AI-enabled translation, coordination, and market analysis may make it easier for smaller firms to internationalize. The net effect on globalization is therefore uncertain. AI may centralize some knowledge work while decentralizing other forms of market access.

Global value chains may also become more data-intensive. Lead firms may require suppliers to provide real-time production, labor, environmental, and logistics data. AI can then be used to monitor compliance, optimize inventory, and predict disruption. This may improve resilience and sustainability, but it can also increase power asymmetry between lead firms and suppliers. Suppliers lacking digital capability may be excluded from high-value chains even if their production costs are competitive.

The literature on firm innovation and ESG suggests that AI can help firms move from reactive compliance to proactive strategy. Instead of reporting sustainability outcomes after the fact, firms can use AI to predict risks, identify process improvements, and design greener products. However, this requires high-quality data and responsible governance. Poorly measured or poorly governed AI may simply produce faster reporting without better substance.

A global business research agenda should therefore examine AI adoption across the full chain of value creation: upstream suppliers, focal firms, downstream distributors, customers, regulators, and capital markets. Such a perspective would make it possible to study how AI redistributes value not only within firms but also between firms and across countries.

8. AI, ESG Performance, and Green Innovation

The fourth theme concerns sustainability. AI is increasingly discussed as a tool for environmental management, social responsibility, and governance improvement. Shen et al. examine AI adoption and corporate ESG performance using a refined large language model to identify genuine applied AI technologies in corporate disclosures. Their measurement strategy is important because firms may mention AI rhetorically without actually adopting it. By distinguishing authentic AI application from vague digital language, the study improves the credibility of firm-level AI measurement.

Their findings suggest a robust positive relationship between AI adoption and ESG performance, driven mainly by green innovation and improved internal control quality. This mechanism is theoretically plausible. AI can help firms monitor emissions, optimize energy use, detect operational inefficiencies, improve supply-chain transparency, and support compliance reporting. It can also improve governance by strengthening internal controls, risk detection, and information processing. However, the effects are stronger among large and technology-intensive firms, which indicates that AI-enabled ESG benefits require complementary capabilities. [9-12]

Mijit et al. provide policy-level evidence from China's Artificial Intelligence Innovation and Development Pilot Zones. Using a heterogeneity-robust staggered difference-in-differences design, they find that the policy significantly enhances the quantity and quality of firms' green innovation. The mechanism appears to operate through improved operational efficiency. This finding is useful because it links AI policy, firm operations, and environmental innovation in a causal framework. It also shows why modern DID methods matter: policy adoption across cities and years can produce biased estimates if heterogeneous treatment effects are ignored.

The ESG literature also raises a digital divide problem. If AI improves sustainability performance mainly for large or technology-intensive firms, then AI may widen the gap between leading firms and lagging firms. Large firms can hire AI specialists, build data platforms, and absorb compliance costs. Small and medium-sized enterprises may lack the resources to implement AI for environmental monitoring or ESG disclosure. In global supply chains, this could create new barriers for suppliers in developing economies if buyers expect AI-enabled traceability and sustainability reporting.

AI also has environmental costs. Although much of the emerging sustainability literature focuses on AI's positive contribution to green innovation, a balanced review should acknowledge that AI systems consume energy, require computing infrastructure, and may increase demand for data centers. The sustainability effect of AI is therefore not automatically positive. It depends on whether efficiency gains and green innovation outweigh energy use and rebound effects. Future research should compare AI's direct environmental footprint with its indirect contribution to resource efficiency and low-carbon innovation.

For global business, the central insight is that AI may become part of corporate sustainability strategy. Firms that integrate AI with ESG governance may improve reporting quality, identify operational risks, and develop greener products. Yet these gains depend on institutional incentives, disclosure standards, stakeholder pressure, and internal capability. AI should therefore be studied as a socio-technical sustainability tool rather than a purely technological solution.

9. Methodological Development in AI Impact Research

The reviewed literature also demonstrates the importance of research design. Because AI adoption is not randomly assigned, empirical studies face serious endogeneity problems. More productive firms may adopt AI earlier, richer cities may receive AI policy support, and high-wage occupations may be both more exposed to AI and more likely to experience unrelated wage changes. Without credible identification, researchers may confuse correlation with causation.

Difference-in-differences (DID) designs are common in policy-oriented AI studies. Traditional two-way fixed effects models compare treated and untreated units before and after policy implementation. However, de Chaisemartin and D'Haultfoeuille, Goodman-Bacon, and Callaway and Sant'Anna show that conventional DID estimators can be biased when treatment timing varies and treatment effects are heterogeneous. This is highly relevant for AI policy because cities, firms, or industries often adopt AI-related policies at different times. A policy may also affect early adopters differently from later adopters.

Mijit et al.' use of a heterogeneity-robust staggered DID model reflects this methodological progress. In AI research, this matters because treatment effects are unlikely to be constant. Large firms may respond quickly to AI pilot zones, while smaller firms may respond slowly. Capital-intensive industries may use AI for process optimization, while labor-intensive industries may use it for quality control or monitoring. A robust DID design can better capture these differences.

Shift-share designs are also relevant when researchers use exposure measures based on occupational, industrial, or regional composition. Borusyak, Hull, and Jaravel clarify the logic of quasi-experimental shift-share instruments. In AI research, exposure is often constructed by combining the share of workers in an occupation with an occupation-level AI exposure score. This can be powerful, but it requires careful assumptions. Researchers must ask whether variation in exposure is plausibly exogenous or whether it reflects pre-existing economic trends.

Measurement is another methodological challenge. AI adoption can be measured through patents, job postings, earnings-call transcripts, annual reports, software purchases, employee surveys, or direct platform usage. Each measure has strengths and weaknesses. Patent data may capture innovation but miss routine adoption. Job postings may capture demand for AI talent but not actual implementation. Corporate disclosures may contain rhetorical inflation. Platform usage data may be precise but difficult to access. Shen et al.' LLM-based classification strategy is an important response to this problem because it attempts to separate genuine applied AI from superficial mentions.

Future AI research should combine stronger measurement with stronger identification. Ideally, studies would distinguish AI exposure, AI adoption, AI intensity, and AI capability. Exposure means that tasks could be affected by AI. Adoption means that an organization has introduced AI tools. Intensity means the degree of use across workflows. Capability means the organizational ability to generate value from AI. Treating these concepts as the same variable risks producing misleading conclusions.

10. An Integrated Framework

Based on the reviewed literature, this paper proposes an integrated framework linking AI adoption to economic and organizational outcomes through three pathways. The first pathway is task reallocation. AI changes which tasks are performed by workers and which are performed by capital. This pathway directly affects labor demand, wages, and the labor share. The second pathway is productivity complementarity. AI raises worker productivity when it helps workers perform tasks better, faster, or with fewer errors. This pathway may reduce performance gaps in some occupations but may also increase returns to workers who can use AI effectively. The third pathway is capability formation. Firms convert AI into growth, product innovation, ESG improvement, and green innovation when they combine AI with data, human capital, managerial routines, and institutional support.

These pathways interact. Task reallocation can increase firm productivity but reduce labor bargaining power. Productivity complementarity can raise worker output but may still benefit capital owners if firms capture most of the surplus. Capability formation can support green innovation but may concentrate benefits among large firms. Therefore, AI's overall impact depends on the balance among labor substitution, labor complementarity, capital ownership, and organizational capability.

The framework also highlights the role of institutions. Education systems influence whether workers can adapt to AI. Labor-market regulation influences whether productivity gains are shared. Competition policy influences whether AI advantages become market concentration. ESG disclosure rules influence whether firms use AI for genuine sustainability improvement or symbolic reporting. Industrial policy influences whether AI adoption spreads beyond leading firms and regions. These institutional factors are especially important in global business because firms operate across different regulatory and cultural environments.

A global business perspective adds two dimensions that are less developed in the reviewed literature. The first is cross-border value chains. AI may allow headquarters to monitor suppliers, forecast demand, and automate coordination, but it may also shift power toward lead firms with superior data infrastructure. The second is international inequality. Countries with strong digital infrastructure, AI talent, and capital markets may capture more gains, while countries dependent on routine service outsourcing or low-cost manufacturing may face new competitive pressure. AI therefore has implications for both firm strategy and the international distribution of value. [13-15]

11. Critical Evaluation of the Existing Literature

Although the reviewed literature has developed rapidly, it remains marked by several tensions. The first tension concerns the difference between technical capability and economic adoption. Many AI exposure studies begin from what AI could do, but business impact depends on whether firms actually deploy the tool, whether employees use it regularly, and whether the tool is embedded into revenue-generating or cost-saving processes. A high exposure score may therefore exaggerate short-term disruption if organizational adoption is slow. Conversely, a low exposure score may understate change in firms that use AI creatively to redesign work beyond formally measured tasks.

The second tension concerns short-term measurement versus long-term transformation. Workplace experiments and early adoption studies can identify immediate productivity effects, but they may miss later changes in job design, occupational entry, promotion pathways, and organizational hierarchy. For example, if AI helps junior workers perform better, firms may initially become more productive. Over time, however, firms may hire fewer junior workers, change training systems, or compress career ladders. The long-term distributional effect may therefore differ from the short-term productivity effect.

The third tension concerns average effects versus heterogeneity. The literature repeatedly shows that AI effects vary across workers, firms, industries, and regions. Yet many public claims about AI still rely on average exposure or average productivity estimates. This is analytically weak because AI is a general-purpose technology whose value depends on context. A more rigorous approach should ask where AI is adopted, by whom, for which tasks, under what governance conditions, and with what complementary investment.

The fourth tension concerns productivity and welfare. Measured productivity may rise even when social welfare does not improve proportionally. Acemoglu notes that some AI-created tasks may have negative social value, such as manipulation-oriented algorithmic design. Similarly, AI may increase the efficiency of advertising, monitoring, or compliance paperwork without necessarily improving human well-being. Business research should therefore avoid equating all AI-enabled output with desirable economic progress.

The fifth tension concerns the unequal visibility of firms. Listed firms, large firms, and technology-intensive firms are easier to study because they disclose more information and leave more digital traces. This creates an evidence bias. The firms most visible in data may also be the firms most capable of using AI. As a result, the literature may overestimate the ease of adoption for ordinary firms and underestimate the adjustment difficulties faced by small suppliers, local service firms, and firms in less digitized industries.

12. Research Gaps and Future Directions

Several gaps remain in the literature. First, more evidence is needed from developing economies and emerging markets. Much of the current literature focuses on the United States, OECD countries, or Chinese listed firms. These contexts are important, but they do not represent the full global economy. In many developing economies, firms face weaker digital infrastructure, lower access to AI talent, and different labor-market institutions. The same AI tool may therefore produce different productivity and inequality effects.

Second, small and medium-sized enterprises require more attention. Large firms dominate existing AI adoption studies because they disclose more data and have more visible AI investments. Yet SMEs are central to employment and supply chains. Research should examine whether cloud-based AI lowers entry barriers or whether it increases dependence on large platform providers. This question is especially relevant for export-oriented SMEs that must meet international standards for quality, delivery, and ESG reporting.

Third, future research should examine long-term labor adjustment. Early studies may capture short-term productivity gains or exposure patterns, but labor-market effects can unfold slowly. Workers may change tasks within occupations, firms may redesign jobs, and education systems may adjust curricula. Longitudinal research is needed to determine whether AI creates durable new tasks or mainly reorganizes existing work.

Fourth, researchers should study AI governance inside firms. Many studies measure whether firms adopt AI, but fewer examine how firms manage risks such as bias, data privacy, hallucination, cybersecurity, accountability, and employee deskilling. Governance quality may determine whether AI adoption improves performance or creates new vulnerabilities. In global firms, governance is even more complex because AI systems must comply with different legal regimes and cultural expectations.

Fifth, the relationship between AI and ESG should be studied more critically. Current evidence suggests positive associations, but there is a risk of overclaiming. AI may improve green innovation and internal control, but it may also increase energy consumption, enable greenwashing, or shift environmental burdens to data-center regions. Future studies should measure both positive and negative environmental channels.

Sixth, more interdisciplinary work is needed. Economics provides strong models of tasks, wages, and productivity. Management research explains capabilities, strategy, and organizational change. Information systems research contributes knowledge about digital implementation. Sustainability research evaluates environmental and governance outcomes. AI's business impact crosses all these fields, so future literature should integrate them rather than treating AI as a narrow technical variable.

13. Managerial and Policy Implications

For managers, the first implication is that AI adoption should begin with task diagnosis rather than tool enthusiasm. Firms should identify which tasks are repetitive, information-intensive, error-prone, or creativity-supporting, and then decide whether AI should automate, assist, monitor, or redesign those tasks. A task inventory can prevent firms from adopting AI as a symbolic digital transformation project without clear operational value.

The second managerial implication is that AI requires complementary human capital. The literature does not support the idea that AI simply eliminates the need for skilled workers. Instead, it changes the type of skill that matters. Workers need the ability to formulate prompts, evaluate outputs, understand domain context, protect sensitive data, and know when human judgment should override algorithmic suggestions. Training should therefore focus not only on tool operation but also on critical evaluation and responsible use.

The third implication is that firms should govern AI adoption explicitly. Internal policies should clarify acceptable use, data protection, accountability for AI-assisted decisions, documentation requirements, and human review procedures. This is especially important for firms using AI in hiring, credit assessment, customer profiling, legal compliance, ESG reporting, or cross-border data processing. Without governance, AI may create hidden operational and reputational risks.

The fourth implication concerns inclusive implementation. If AI helps some workers more than others, firms should avoid letting adoption deepen internal inequality. Training, access, and support should be distributed broadly rather than limited to technical departments or high-status employees. Managers should monitor whether AI changes workloads, promotion opportunities, and worker autonomy. Productivity gains are more sustainable when employees perceive AI as a support for better work rather than as a surveillance or replacement device.

For policymakers, the first implication is that education and training systems must adapt to task-level change. AI literacy should include data awareness, algorithmic limitations, ethical reasoning, and domain-specific application. Policies that only subsidize technology purchase may fail if workers and managers lack the skills to use AI productively. Public training programs, university-industry collaboration, and lifelong learning systems are therefore central to inclusive AI diffusion.

The second policy implication is that competition and industrial policy matter. If AI advantages accumulate among a small group of large firms, productivity gains may be accompanied by market concentration. Policymakers should support data access, digital infrastructure, and AI services for SMEs while also monitoring platform dependence and anticompetitive behavior. In global value chains, support for supplier digital upgrading may be necessary to prevent exclusion from ESG-sensitive markets.

The third policy implication concerns distribution. AI may reduce some wage gaps while increasing capital income and wealth inequality. Tax policy, social insurance, labor standards, and worker bargaining institutions will influence whether AI gains are broadly shared. The reviewed literature therefore suggests that AI policy should not be limited to innovation promotion. It should also include mechanisms for adjustment, redistribution, and protection against harmful uses. [16-17]

The fourth policy implication is environmental. Governments can encourage AI applications that improve energy efficiency, green innovation, and emissions monitoring, but they should also measure the energy cost of AI infrastructure. A balanced policy approach would support AI for green transformation while setting standards for transparent environmental accounting in data centers and AI-intensive industries.

14. Proposed Future Research Agenda for Global Business

A useful future research agenda can be organized around four levels. At the worker level, scholars should examine how AI changes skills, autonomy, career progression, and wage bargaining across occupations. At the firm level, researchers should study how AI capabilities are built, governed, and converted into innovation or ESG outcomes. At the industry level, studies should compare sectors where AI mainly automates routine knowledge work with sectors where it supports complex coordination or product development. At the international level, scholars should investigate how AI redistributes value across countries and value-chain positions.

One promising direction is comparative research on AI adoption in global value chains. Researchers could examine whether lead firms use AI to strengthen supplier monitoring, reduce transaction costs, or shift compliance burdens to suppliers. Another direction is the study of AI and service offshoring. Generative AI may substitute for some outsourced language and analytical tasks, but it may also create new demand for human review, localization, and domain-specific judgment. The effect on countries specializing in business-process outsourcing remains uncertain.

A second promising direction is the relationship between AI and entrepreneurial internationalization. AI tools may help small firms translate materials, analyze foreign markets, generate marketing content, and manage customer communication. This could lower barriers to exporting and cross-border digital entrepreneurship. However, small firms may also face risks related to data dependence, model errors, and platform lock-in. Empirical work should therefore distinguish empowerment from dependency.

A third direction concerns AI and institutional distance. Multinational firms operate across countries with different privacy rules, labor norms, ESG standards, and attitudes toward automation. AI governance strategies that work in one country may not be acceptable in another. Research could examine how firms adapt AI systems to different regulatory and cultural contexts, and whether global AI policies become centralized or locally customized.

A fourth direction concerns responsible AI as a source of competitive advantage. Firms that can demonstrate transparent, accountable, and sustainable AI use may gain trust from customers, regulators, investors, and employees. This is especially relevant in ESG-sensitive industries. Future research could test whether responsible AI governance improves reputation, lowers compliance risk, or supports long-term innovation performance.

Finally, future research should use mixed methods. Quantitative methods can estimate productivity, wage, and innovation effects, but qualitative research can reveal how workers and managers actually experience AI adoption. Case studies, interviews, and process tracing can uncover mechanisms that are invisible in large datasets. A mature global business literature on AI will likely require both credible causal inference and rich organizational understanding.

15. Conclusion

This review has synthesized recent literature on AI adoption, labor-market inequality, and firm transformation. The central conclusion is that AI is best understood as a task-transforming and capability-dependent technology. Its effects are not predetermined by technical possibility alone. They depend on which tasks are affected, whether AI substitutes for or complements workers, who owns the relevant capital, how firms redesign workflows, and how institutions shape the distribution of gains.

The reviewed literature challenges both excessive optimism and excessive pessimism. AI may generate productivity gains, but aggregate effects may be more modest than public forecasts suggest. AI may help less experienced workers and reduce within-occupation performance gaps, but it may also increase capital income and firm-level concentration. AI may support ESG performance and green innovation, but benefits are likely to be uneven and may depend on complementary assets. The same technology can therefore produce different outcomes across workers, firms, industries, and countries.

For global business scholars, AI adoption should be studied as an international organizational transformation rather than a purely technological shock. Future research should connect task-level mechanisms with firm-level capabilities, sustainability governance, cross-border value chains, and inequality between regions and countries. For managers and policymakers, the key implication is that AI value depends on complementary investment: human capital, data governance, organizational learning, responsible deployment, and inclusive policy design. Without these complements, AI may deepen existing divides. With them, AI can become a tool for productivity, innovation, and more sustainable business transformation.

The reviewed studies also suggest a practical scholarly attitude toward AI: cautious openness. Caution is necessary because early evidence is incomplete, many effects are heterogeneous, and some benefits may be captured by capital owners or leading firms. Openness is equally necessary because AI can support learning, innovation, operational efficiency, and environmental upgrading when it is embedded in capable organizations and supportive institutions. The most valuable future studies will therefore avoid deterministic claims and instead examine the concrete conditions under which AI produces inclusive or exclusive outcomes. Such work can help managers, policymakers, and scholars understand not only whether AI matters, but how, for whom, and under what forms of governance it can contribute to sustainable global business development.

References

[1] Acemoglu, D. (2024). The simple macroeconomics of AI. Working paper prepared for Economic Policy.

[2] Acemoglu, D., Kong, F., & Restrepo, P. (2025). Tasks at work: Comparative advantage, technology and labor demand. NBER Working Paper No. 32872. National Bureau of Economic Research. https://www.nber.org/papers/w32872.

[3] Acemoglu, D., & Loebbing, J. (2024). Automation and polarization. Working paper.

[4] Auer, R., Köpfer, D., & Švéda, J. (2024). The rise of generative AI: Modelling exposure, substitution, and inequality effects on the US labour market. BIS Working Papers, No. 1207. Bank for International Settlements.

[5] Babina, T., Fedyk, A., He, A., & Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. Journal of Financial Economics, 151, 103745. https://doi.org/10. 1016/j. jfineco. 2023.103745.

[6] Borusyak, K., Hull, P., & Jaravel, X. (2022). Quasi-experimental shift-share research designs. Review of Economic Studies, 89(1), 181-213. https://doi.org/10. 1093/ restud/ rdab030.

[7] Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at work. Quarterly Journal of Economics, 140(2), 889-942. https: // doi.org/10.1093/qje/qjae044.

[8] Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. Journal of Econometrics, 225(2), 200-230. https://doi. org/ 10. 1016/ j. jeconom. 2020.12.001.

[9] Cazzaniga, M., Jaumotte, F., Li, L., Melina, G., Panton, A. J., Pizzinelli, C., Rockall, E., & Tavares, M. M. (2024). Gen-AI: Artificial intelligence and the future of work. IMF Staff Discussion Notes, 2024(001). International Monetary Fund.

[10] de Chaisemartin, C., & D'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. American Economic Review, 110(9), 2964-2996. https://doi. org/ 10.1257/aer.20181169.

[11] Georgieff, A. (2024). Artificial intelligence and wage inequality. OECD Artificial Intelligence Papers, No. 13. OECD Publishing.

[12] Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. Journal of Econometrics, 225(2), 254-277. https://doi.org/10.1016/j.jeconom.2021.03.014.

[13] He, H., Wang, Q., Huang, W., Yang, M., Ma, H., & Pang, H. (2026). Can artificial intelligence narrow the urban-rural income inequality? Evidence from a quasi-natural experiment in China. Sustainability, 18(10), 4785. https://doi. org/ 10. 3390/ su18104785.

[14] Humlum, A., & Vestergaard, E. (2025). Still waters, rapid currents: Generative AI and labor markets. NBER Working Paper No. 33777. National Bureau of Economic Research. https://www.nber.org/papers/w33777.

[15] Mijit, R., Hu, Q., Xu, J., & Ma, G. (2025). Greening through AI? The impact of Artificial Intelligence Innovation and Development Pilot Zones on green innovation in China. Energy Economics, 146, 108507. https://doi.org/ 10.1016/ j.eneco. 2025. 108507.

[16] Rockall, E., Tavares, M. M., & Pizzinelli, C. (2025). AI adoption and inequality. IMF Working Paper WP/25/68. International Monetary Fund.

[17] Shen, L., Li, Z., Liang, Y., Feng, Y., & Zhang, Z. (2025). Artificial intelligence adoption and corporate ESG performance: Evidence from a refined large language model. Frontiers in Artificial Intelligence, 8, 1691468. https://doi.org/ 10. 3389/frai.2025.1691468.