Artificial Intelligence, Firm Heterogeneity, and Labor Market Adjustment: Evidence from Service and Industrial Tech Sectors in Italy

Ebrima K. Ceesay [a], Mamadou Salieu Jallow [b], Cosimo Magazzino [c], Alasana Gitteh [d], Baseedy Bojang [e]

[a] *Postdoctoral Research fellow, University of Insubria, Italy*
[b] *West African Science Service Center on Climate Change and Adapted Land Use, UCAD, Senegal*
[c] *Department of Management, Finance and Technology, LUM University "Giuseppe Degennaro", Casamassima, Italy*
[a, e, d] *School of Business and Public Administration, Department of Economics and Finance, University of The Gambia,*

Keywords
Artificial Intelligence; Labor Market Outcomes; Skill-Biased Technological Change; Firm-Level Panel Data; Italian Economy

Abstract
This study assesses how changes in AI patent stock affect employment, wages, productivity, and labor cost shares. It compares pooled estimates with firm-specific effects for UniCredit and Zerynth in Italy, representing the banking services and industrial tech sectors, respectively. Using a firm-level panel dataset from 2005 to 2024, the study applies task-based and skill-biased technological change frameworks to analyze how exposure to Artificial Intelligence (AI), proxied by AI patent stock, affects labor-market dynamics in the Italian economy. Our empirical strategy used fixed-effects regressions. The evidence points to a dual trajectory in how firms adjust to technological change. In pooled models, AI-related innovation has a positive and statistically significant effect on labor-market outcomes. However, when firms are analyzed separately, the effects diverge. For UniCredit, AI patent stock reduces employment and labor cost shares while raising productivity growth, with wage effects remaining small or insignificant. On the other hand, for Zerynth, AI patent stock produces consistently negative and significant effects on employment, wages, and labor cost shares. Overall, the findings highlight strong task-level substitution but heterogeneous firm-level outcomes, underscoring the need for targeted reskilling and labor-augmenting innovation policies in Italy's digital transition.

Citation (APA)
Ceesay, E. K., Jallow, M. S., Magazzino, C., Gitteh, A., & Bojang, B. (2026). Artificial Intelligence, Firm Heterogeneity, and Labor Market Adjustment: Evidence from Service and Industrial Tech Sectors in Italy . *Muslim Business and Economics Review, 5*(1), (90-114)
https://doi.org/10.56529/mber.v5i1.591

Submitted : 30 March 2026
Revised : 18 May 2026
Accepted : 15 June 2026
Published : 29 June 2026

Corresponding Author : jallow.m@edu.wascal.org

1. Introduction

Artificial Intelligence (AI) has moved from a specialized research domain to a core component of firms' production, organization, and competitive strategies. Its economic relevance is no longer limited to frontier technology companies. AI now affects credit allocation, fraud detection, customer interaction, predictive maintenance, industrial Internet of Things (IoT), logistics, software development, and managerial decision-making. For labor markets, this matters because AI differs from earlier waves of automation. It not only replaces routine manual work; it can also perform prediction, classification, language-processing, and analytical tasks that were previously associated with clerical, professional, and technical occupations. This gives AI the characteristics of a general-purpose technology whose impact depends on complementary investment, organizational redesign, worker skills, and the institutional environment in which firms operate (Brynjolfsson & McAfee, 2014; Brynjolfsson *et al*., 2018; Filippucci *et al*., 2024).

The central economic issue is therefore not whether AI affects work, but how its effects are distributed across firms, sectors, and types of workers. Task-based models show that automation simultaneously produces displacement effects, when machines substitute for workers in specific tasks, and reinstatement effects, when new tasks and occupations are created around technology (Acemoglu & Restrepo, 2018, 2019). Recent evidence on AI exposure suggests that the same technology may operate as either a substitute or a complement, depending on occupational content, digital capabilities, and firm-level absorptive capacity (Pizzinelli *et al*., 2023). These mechanisms make aggregate predictions insufficient. Employment may rise even when labor's cost share falls; wages may increase for AI-complementary workers while routine jobs are compressed; productivity gains may be delayed because firms must reorganize workflows before innovation becomes measurable in performance indicators.

Italy offers a relevant setting for studying these mechanisms. The country combines advanced manufacturing districts, internationally integrated services, a large base of small and medium-sized enterprises, persistent productivity weaknesses, and marked territorial disparities in digital capabilities. Recent official statistics confirm that AI adoption is accelerating but remains uneven: the share of Italian firms with at least 10 employees using AI rose from 5.0% in 2023 to 8.2% in 2024 and 16.4% in 2025, while adoption among large firms reached 53.1% in 2025 and the gap between large firms and SMEs widened (ISTAT(Istituto Nazionale di Statistica), 2025a, 2025b)). Bank of Italy evidence similarly shows that AI adoption among Italian firms is still concentrated in larger and more knowledge-intensive organizations, with productivity and profitability gains accompanied by occupational reallocation rather than uniform job creation (Bencivelli *et al*., 2025). This combination of rapid diffusion and structural heterogeneity makes Italy an informative case for assessing whether AI generates inclusive labor-market upgrading or reinforces existing divides.

The policy relevance of the topic has increased sharply. In the European Union (EU), the AI Act entered into force in August 2024 and established a regulatory framework to promote trustworthy AI while managing risks associated with high-risk systems, ensuring transparency, promoting governance, and ensuring human oversight (European Parliament and Council, 2024; European Commission, 2024). For Italy, the policy challenge is broader than legal compliance. AI policy must also address the distribution of productivity gains, the protection and retraining of workers exposed to task substitution, the diffusion of digital capabilities among SMEs, and the risk that large firms capture most of the benefits from AI-enabled capital deepening. A firm-level analysis is therefore useful because policy instruments do not operate on abstract technologies; they operate through firms that differ in resources, task structures, labor composition, and innovation strategies.

This study contributes to this debate by examining how changes in AI patent stock affect employment, wages, productivity, and labor cost shares in two Italian firms: UniCredit and Zerynth. The sample is deliberately comparative rather than representative. UniCredit is a large incumbent in financial services, where AI is mainly adopted to improve risk analytics, customer service, compliance, fraud detection, and back-office efficiency. Zerynth is an industrial technology firm operating in AI-driven IoT and automation, where AI is more closely embedded in products, platforms, and industrial process innovation. This contrast allows the paper to compare two distinct AI pathways: AI as an adoption and efficiency technology in a mature service-sector firm, and AI as a production and scaling technology in an industrial-tech firm. The two-firm design should therefore be read as a focused firm-level comparison that generates micro-evidence on heterogeneity, not as a population-average estimate for all Italian firms.

The originality of the study lies in combining firm-level labor outcomes with a forward-looking measure of AI innovation. Much of the Italian literature relies on occupational exposure indices, robot-density indicators, sectoral digitalization measures, or aggregate employment trends. These approaches are valuable, but they cannot fully identify how innovation accumulated inside specific firms is associated with internal labor-market adjustment. This paper instead constructs a firm-year panel for 2005–2024 and links employment, wages, productivity, labor share, profits, revenues, and AI-patent indicators. AI patent stock is used as a proxy for cumulative technological capability and innovation trajectory, following the innovation-economics literature that treats patents as observable traces of knowledge accumulation, even while recognizing that not every patent becomes a commercial application and not every AI application is patented (Webb, 2020; Felten *et al*., 2021; Giczy *et al*., 2022).

The data strategy is designed to match this research question. For UniCredit, financial and labor-market indicators are obtained from BankFocus, which provides standardized annual information for banking institutions. For Zerynth, firm-level information is drawn from Orbis Intellectual Property (Orbis IP), complemented where necessary by public disclosures to preserve continuity across the sample period. Patent-level data are extracted from Orbis IP and classified as AI-related using a hybrid approach that combines IPC(International Patent Classification)/CPC (Cooperative Patent Classification) codes, AI-specific keywords, and citation links to AI-relevant inventions. The resulting panel connects innovation inputs to labor-market and performance outcomes over two decades, spanning the pre-deep-learning phase, the expansion of machine-learning applications, and the recent acceleration of generative and data-intensive AI.

The empirical methodology is appropriate to the study's objective because it combines several layers of analysis rather than relying on a single specification. Fixed-effects regressions control for time-invariant differences between the two firms and isolate within-firm associations between lagged AI patent stock and labor outcomes. The use of lagged patent stock reflects the delay between innovation activity and observable labor-market adjustment. Given the narrow two-firm structure of the sample, the estimates are interpreted cautiously as firm-level longitudinal evidence rather than definitive causal parameters.

The research's expected contribution is fourfold. First, it adds firm-level evidence to a literature that often studies AI through aggregate, occupational, or sectoral exposure measures. Second, it shows that AI's labor-market effects are heterogeneous even within the same national economy: the direction and magnitude of employment, wage, productivity, and labor-share responses depend on whether AI is used to reorganize a mature service operation or to scale industrial-technology capabilities. Third, it links AI patent stock to labor cost shares, thereby connecting the innovation literature with debates on the distribution of value between labor and capital. Fourth, it provides policy-relevant evidence for Italy's digital transition by distinguishing between labor-augmenting and labor-substituting channels and by identifying where reskilling, worker protection, innovation incentives, and complementary investment policies are most likely to matter.

The novelty of the results is that the pooled evidence and the firm-specific evidence do not tell the same story. Pooled estimates suggest that AI-related innovation is associated with higher wages, employment, and productivity, but also with a lower labor cost share. Firm-level estimates reveal sharper divergence. In UniCredit, AI patent stock is associated with moderate labor adjustment in a mature financial services setting, while profitability plays a central role in wage, productivity, and labor share dynamics. In Zerynth, AI patent stock is associated with short-run reductions in employment, wages, and productivity, consistent with the disruption and reorganization costs faced by scaling technology firms. These findings support the view that AI reallocates tasks and value rather than producing a uniform employment shock (Acemoglu *et al*., 2022).

The policy implications follow directly from this heterogeneity. A single AI policy cannot fit all firms. In service sectors, policy should focus on transition support for routine cognitive workers, internal mobility, worker retraining, transparency in algorithmic management, and mechanisms that allow productivity gains to be shared through wages and career progression. In industrial technology, policy should support AI-complementary skills, R&D, and data infrastructure, and the creation of high-value technical occupations while monitoring possible short-run displacement. Across the Italian economy, the results point to the need for SME-oriented digital upgrading, targeted fiscal incentives that combine technology adoption with human-capital investment, and labor-market observatories capable of tracking AI's effects at the firm and occupation levels.

The remainder of the paper is organized as follows. Section 2 reviews the literature on AI, automation, task-based technological change, and labor-market outcomes, with particular attention to Italy and sectoral heterogeneity. Section 3 presents the theoretical framework, the construction of the AI patent-stock variable, the data sources, and the econometric specifications. Section 4 describes the dataset and reports the empirical results, including descriptive statistics, baseline fixed-effects estimates, firm-level comparisons, and robustness checks. Section 5 discusses the conclusions and develops policy recommendations for an AI transition that strengthens productivity while limiting labor-market polarization.

2. Literature review

The impact of AI and automation on labor markets is neither uniform nor predetermined; it is shaped by national economic structures, institutional settings, and firm-level characteristics. Italy, with its distinctive blend of specialized manufacturing, small and medium-sized enterprises (SMEs), regional disparities, and persistent productivity challenges, offers a revealing context for examining how AI adoption differentially influences employment, wages, and skill demand. This review synthesizes theoretical frameworks, empirical findings, and sectoral analyses relevant to AI's labor-market effects in Italy, focusing on the interplay among technology, tasks, and firm heterogeneity.

2.1. Theoretical Foundations: Task-Based Models and Routine-Biased Technological Change

Contemporary analysis of technological change is grounded in task-based models, notably advanced by Acemoglu & Restrepo (2017, 2018, 2019). In this framework, production is conceptualized as a continuum of tasks that can be performed by either labor or capital. Technological progress does not simply eliminate jobs; it reallocates tasks between humans and machines, simultaneously displacing labor in automatable tasks and creating new, more complex tasks that require human intervention. This dynamic generates both substitution and reinstatement effects, with net employment outcomes depending on the balance between the two. A closely related concept, Routine-Biased Technological Change (RBTC), provides a compelling explanation for structural shifts observed in advanced economies (Autor *et al*., 2003). RBTC posits that technologies such as AI are particularly adept at replacing labor in routine tasks – whether cognitive or manual – that are rule-based, predictable, and codifiable. This leads to job polarization—a decline in middle-skill, routine-intensive occupations alongside growth in high-skill analytical jobs and low-skill manual service roles (Autor & Dorn, 2013). For Italy, where middle-skill manufacturing and administrative roles remain prevalent, the RBTC framework suggests heightened vulnerability to polarization, especially in industrial districts characterized by precision and repeatability.

2.2. Empirical Evidence: Automation and AI in the Italian Context

Empirical research on automation in Italy has predominantly focused on manufacturing, the sector most exposed to routine task substitution. Barbieri *et al*. (2020) find that exposure to automation technologies – including industrial robots and machine-learning systems – is negatively associated with employment, particularly for low-skilled workers performing routine production tasks. These effects are more pronounced in the industrialized north, suggesting that regional economic structures may amplify technological displacement and potentially widen Italy's longstanding north–south divide. More recent studies have begun to disentangle the specific effects of AI from broader automation. Vivarelli *et al*. (2023), using firm-level data from Italian manufacturing, find that robot adoption tends to reduce low-skill employment, while AI technologies more often complement high-skill technical roles. This highlights the differentiated labor-market impacts of various technologies. However, the Italian literature remains limited in its use of granular firm-level data. Most studies rely on sectoral or occupational exposure indices, which, while informative, cannot capture within-firm dynamics such as wage structures, investment in innovation, or internal task reallocation.

2.3. Sectoral Heterogeneity: Services versus Industry

The labor-market impact of AI varies across sectors due to differences in task structures. In data-rich service industries, AI automates routine cognitive tasks such as algorithmic trading, fraud detection, and customer service operations. In banking, this can reduce demand for routine clerical roles while increasing demand for high-skill positions in data analytics, risk modeling, and compliance. Tourism and retail similarly deploy AI for dynamic pricing, inventory management, and personalized marketing, potentially automating routine booking or checkout tasks while creating new roles in digital platform management and customer-experience design. The net employment effect in services depends on whether AI is deployed primarily for cost reduction or for service enhancement and innovation. Firms that produce AI technologies – such as industrial IoT and automation solutions – exhibit a different dynamic. Here, AI is embedded directly into products and manufacturing processes (e.g., predictive maintenance, edge computing). This can drive expansion in high-skill R&D and engineering employment within the technology firm itself, even as its products may automate tasks in client firms. This sector exemplifies the reinstatement effect, where innovation creates new, complex tasks and occupational niches. The contrasting experiences of a service-sector adopter (UniCredit) and an industrial tech producer (Zerynth) illustrate this sectoral duality and provide a micro-level lens to examine these divergent pathways.

2.4. Skills, Polarization, and the SME Dilemma

The shift in labor demand precipitated by RBTC increases the premium on skills that complement AI: high-level cognitive abilities, digital literacy, and socio-emotional competencies. Colombo *et al*. (2019), analyzing online job postings in Italy, confirm that demand for digital and analytical skills is rising, while occupations reliant on routine tasks face higher automation risk. This trend poses a challenge for Italy's education and training systems; a mismatch between skill supply and demand risks exacerbating wage inequality and polarization. Crucially, the translation of technological potential into labor-market outcomes is mediated by firm-level characteristics. Firm size is particularly salient in Italy's SME-dominated landscape. Smaller firms often lack the financial resources, internal R&D capabilities, and managerial expertise to adopt and integrate AI effectively. This can lead to a dualistic trajectory in which large firms harness AI for innovation and productivity gains, while many SMEs either delay adoption or implement narrow, cost-focused automation that may suppress employment (Bottero & Schiaffi, 2022). This "SME dilemma" underscores that policies designed for large corporations may be ineffective for most Italian firms, highlighting the need for targeted support mechanisms.

2.5. Synthesis and Research Gap

The existing literature establishes that AI and automation are reshaping Italian labor markets through task displacement, skill-biased demand shifts, and job polarization, with effects varying by region and sector. However, a significant gap remains: few studies exploit detailed, firm-level longitudinal data to link internal innovation activity – proxied by AI-specific patent stocks – directly to employment, wage, and productivity outcomes. By constructing such a dataset and applying a task-based theoretical framework to the cases of UniCredit and Zerynth, this study aims to bridge this gap. It seeks to provide nuanced evidence on whether AI acts primarily as a labor-displacing force in traditional sectors or as a labor-augmenting driver in innovation-intensive environments, thereby informing more targeted policy in Italy's digital transition.

3. Research Methodology: Theoretical Framework and Econometric Specification

3.1. Theoretical Framework: A Task-Based Model with AI Patents

Building on Acemoglu and Restrepo's (2018) task-based framework, we conceptualize firm production[1] as a continuum of tasks $\tau \epsilon [0,1]$. Each task can be performed by either labor $L(\tau)$ or AI capital $K_{AI}(\tau)$, where AI capital is derived from the firm's accumulated AI patent stock. In our adaptation, AI capital is operationalized through a firm's accumulated stock of AI-related patents, serving as a proxy for its innovation capacity and technological trajectory.

[1] This formulation adapts Acemoglu and Restrepo's (2018) task based model by operationalizing AI capital through firm level AI patent stocks. The use of patent based measures follows established approaches in innovation economics, where accumulated patents serve as proxies for technological capability and knowledge capital.

The production function for the firm f is specified as:

$$Y_f = \left(\int_0^1 \left[\gamma_L(\tau)L_f(\tau)^{\frac{\sigma-1}{\sigma}} + \gamma_{AI}\left(\tau\right)\left(\theta_f AI\_Patent\_Stock_f\right)^{\frac{\sigma-1}{\sigma}}\right]d\tau\right)^{\frac{\sigma}{\sigma-1}}$$

where:

Y_f: Output of firm f

σ: Elasticity of substitution between labor and AI capital in performing tasks

$\gamma_L(\tau)$,$\gamma_{AI}(\tau)$: Task-specific productivity parameters

θ_f: Firm's specific absorptive capacity (ability to transform patents into productive AI applications)

$AI\_Patent\_Stock_f$: Firm's cumulative AI-related patent stock

The framework generates several key theoretical predictions. First, AI patent stock is expected to increase automation of routine tasks $\tau \epsilon R$, thereby reducing labor demand in those occupations. At the same time, AI patents enable the creation of new, more complex tasks $\tau \epsilon N$, which raises overall productivity. Moreover, AI tends to complement high-skill non-routine cognitive work while substituting for middle-skill routine tasks, producing a clear skill-biased pattern. Taken together, these mechanisms align with Routine-Biased Technological Change (Autor *et al*., 2003) and with Aghion *et al*.'s (2021) distinction between labor-displacing and labor-augmenting innovation in these two firms. Existing Italian studies on AI and the labor market predominantly rely on occupation and sector-level measures of AI exposure, combined with aggregate indicators of employment and wages.

3.2. Research Design and Data Construction

Following established methodologies (Webb, 2020; Felten *et al*., 2018, 2021), we identify AI-related patents using a multi-step approach:

- Data Source: Firm-level patent data were extracted from Orbis Intellectual Property (Orbis IP) from 2005 to 2024.
- Classification: Patents were filtered using relevant IPC/CPC codes (G06N, G06Q, G06F) and AI-specific keywords ("neural network," "deep learning," "computer vision").
- Citation Networks: Patents citing seminal AI papers or previous AI patents were included to capture indirect innovation.
- Patent Stock Construction: For each firm f in year t:

$$AI\_Patent\_Stock_{f,t} = \sum_{s=0}^{\infty} (1-\delta)^{s}.\, AI\_Patent\_Stock_{f,t-s}$$

where:

δ: Knowledge depreciation rate

Patents Grantedf,t denotes the number of AI patents granted in year t (zero in years without a grant), and δ=0.15 is the annual knowledge depreciation rate. Following established practice in the innovation literature (w, Jaffe, and Trajtenberg 2005; Bloom, Schankerman, and Van Reenen 2013), we adopt a depreciation rate of 15%. This declining balance approach—widely known as the Perpetual Inventory Method (PIM) or permanent inventory method—allows past innovative output to depreciate gradually while incorporating new patent grants each year.

3.3. Econometric Specification

To analyze the relationship between AI innovation and labor outcomes, we estimate four baseline fixed-effects models. All models include firm fixed effects α_f to control for time-invariant heterogeneity and year fixed effects δ_t to capture macroeconomic trends. AI patent stock is lagged one period to mitigate simultaneity bias.

Model 1: Wage Effects

This model evaluates whether AI adoption contributes to wage polarization by raising wages for AI-complementary workers while compressing wages for routine workers.

$$In(W_{f,t}) = \beta_0 + \beta_1 \log(AI\_Patent\_Stock)_{f,t,t-1}) + \beta_2 \log(X_{f,t}) + \alpha_f + \delta_t + \epsilon_f$$

where:

$E_{f,t}$,$W_{f,t}$: Total employment and total wages at firm f in year t,the two firms are listed above

$AI\_Patent\_Stock_{f,t-1}$: lagged AI patent stock to address simultaneity bias

$X_{f,t}$: for other control variables

α_f, δ_t:firms and year fixed effects.

Model 2: Employment Effects

This model examines whether AI-related innovation behaves as a labor-saving or labor-augmenting technology. It tests whether increases in AI patent stock reduce employment in routine-intensive firms (UniCredit) and increase employment in innovation-intensive firms (Zerynth). This task-based FE model is as follows:

$$\ln(_{Ef,t}) = \beta_0 + \beta_1 \log(AI\_Patent\_Stock_{f,t,t-1}) + \beta_2 \log(X_{f,t}) + \alpha_f + \delta_t + \epsilon_f$$

Here, β_1 is the percentage change in employment associated with a one-unit increase in AI patents.

Model 3: Labor Cost Share (Skill Composition Proxy)

Because direct occupational skill data are unavailable at the firm level, the labor cost share is used as a proxy for skill composition. A declining labor share may indicate automation of routine tasks, while a rising labor share may reflect increased hiring of high-skill workers.

$$\ln(labor\_Cost\_Share_{f,t}) = \beta_0 + \beta_1 \log(AI\_Patent\_Stock_{f,t,t-1}) + \beta_2 \log(X_{f,t}) + \alpha_f + \delta_t + \epsilon_f$$

Model 4: Productivity Effects

$$\ln(Productivity_{f,t}) = \beta_0 + \beta_1 \log(AI\_Patent\_Stock_{f,t,t-1}) + \beta_2 \log(X_{f,t}) + \alpha_f + \delta_t + \epsilon_f$$

Productivity (revenue per employee) is used as a second proxy for skill composition. AI-complementary firms typically experience productivity gains as high-skill workers leverage AI tools. Furthermore, to measure productivity, we used revenue per employee for output efficiency, and labor cost share for input efficiency assessment.

Identification Strategy

The identification of the causal effect of AI innovation on labor-market outcomes relies on a multi-layered strategy designed to address endogeneity, simultaneity bias, and unobserved heterogeneity. First, we lag AI patent stock by one period to capture the realistic delay between innovation activities and observable labor-market adjustments, thereby reducing concerns about reverse causality. Second, all specifications include firm fixed effects to control for time-invariant unobserved characteristics – such as managerial quality, corporate culture, and long-term innovation strategies – and year fixed effects to absorb macroeconomic shocks, regulatory changes, and sector-wide technological trends. Finally, the comparative design – contrasting UniCredit (service sector) with Zerynth (industrial technology sector) – enables us to isolate sector-specific technological trajectories while controlling for common time effects, thereby strengthening causal inference regarding whether AI operates as labor-displacing or labor-augmenting across different economic contexts.

3.5. Data Collection

The empirical analysis is based on a firm-year panel constructed from multiple sources to ensure coverage of financial, employment, and innovation metrics. For

UniCredit, annual financial and labor market data – including employment counts, wage bills, revenue, and capital stock – were obtained from BankFocus for the period 2005–2024. For Zerynth, corresponding financial and employment data were sourced from Orbis Intellectual Property (Orbis IP), supplemented with verified public disclosures to ensure continuity in early years. Patent-level data for both firms were extracted from Orbis IP, covering all granted patents from 2005 onward, with AI-related patents identified using a hybrid method of IPC/CPC code filtering (G06N, G06Q, G06F), keyword screening, and citation-based network analysis. These datasets were merged using firm identifiers and annual timestamps, yielding a balanced panel that integrates innovation inputs with labor-market outcomes, providing the foundation for econometric analysis of AI's firm-level effects.

The study integrates data from two primary sources to ensure consistency and completeness. Financial and labor market indicators for UniCredit were obtained from BankFocus, which provides standardized annual data for banking institutions, including employment, revenue, wage bills, profits, and capital structure. For Zerynth, all firm-level financial, employment, and innovation data were sourced from Orbis Intellectual Property (Orbis IP), which offers coverage for technology firms not always included in banking-specific databases. Patent data for both firms—including application years, grant years, IPC/CPC classifications, and citations—were also drawn from Orbis IP. The use of Orbis IP for Zerynth ensures comparability in patent measurement across firms, while BankFocus provides high-quality, audited financials for UniCredit, maintaining reliability across sectors.

4. Results and Discussion

Analysis of the summary statistics for the 2005–2024 firm year panel indicates distinct trends in labor market outcomes and innovation activity. The mean logged wage growth is 10.80 (SD = 0.64), while employment growth averages 7.01 (SD = 4.80) and productivity growth averages 11.53 (SD = 0.86). Labor cost share exhibits a negative mean (–0.73, SD = 0.25), indicating that labor's share in total costs declines on average across the sample. AI patent stock growth averages 3.63 (SD = 3.31), with a range from 0 to 7.60, reflecting substantial variation in innovation intensity. Negative skewness and kurtosis values for wages and labor cost share suggest distributions concentrated at higher values with occasional downward shocks, consistent with automation-driven adjustments. The positive mean AI patent stock alongside a negative mean labor cost share aligns with capital-deepening narratives, in which innovation correlates with a shift in cost structure away from labor.

Table 1: Descriptive statistics.

Variable	Mean	Std. Dev.	Min.	Max.	Skewness	Kurtosis
Wage	10.80	0.64	9.62	11.62	-0.15	-1.59
Employment	7.01	4.80	0.69	12.07	-0.03	-1.95
Productivity	11.53	0.86	10.13	12.82	-0.01	-1.69
Labor Cost Share	-0.73	0.25	-1.27	-0.51	-0.55	-1.17
Patent Stock	3.63	3.31	0.00	7.60	0.14	-1.85

Correlation analysis reveals strong positive associations between AI patent stock and wages (0.892), employment (0.869), and productivity (0.890), while labor cost share correlates negatively with innovation (−0.790), underscoring the dual role of AI as both a complement to high-skill labor and a substitute for routine tasks.

Table 2: Correlation Matrix.

Variable	Wage	Employment	Productivity	Labor Share	Patent
Wage	1.000	0.950	0.989	-0.856	0.892
Employment	0.950	1.000	0.963	-0.895	0.869
Productivity	0.989	0.963***	1.000	-0.924	0.890
Labor Share	-0.86	-0.895***	-0.924***	1.000	-0.790
Patent	0.892	0.869	0.890	-0.790	1.000

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.
Source: Authors' evaluation using Stata 16.

The correlation matrix reveals consistent relationships among key variables, offering preliminary insights into AI's role in firm-level dynamics. AI patent stock exhibits a strong positive correlation with wages (0.892), employment (0.869), and productivity (0.890), suggesting that firms with higher AI innovation intensity experience concurrent growth in these outcomes. Productivity and wages are nearly perfectly correlated (0.989), reinforcing the skill-biased technological change narrative where efficiency gains translate into higher compensation. Labor cost share is negatively correlated with productivity (−0.924) and innovation (−0.790), indicating that as firms innovate and become more productive, labor's share in total costs declines – a pattern consistent with capital deepening and automation. Employment and productivity also correlate positively (0.963), yet this relationship is moderated by sector, as shown in firm-specific analyses where UniCredit and Zerynth exhibit divergent employment responses to innovation. These correlations provide an empirical foundation for the regression models, highlighting the intertwined nature of innovation, productivity, wages, and labor share in Italy's evolving firm landscape.

a. Distributional impacts of the variables: Descriptive evidence.

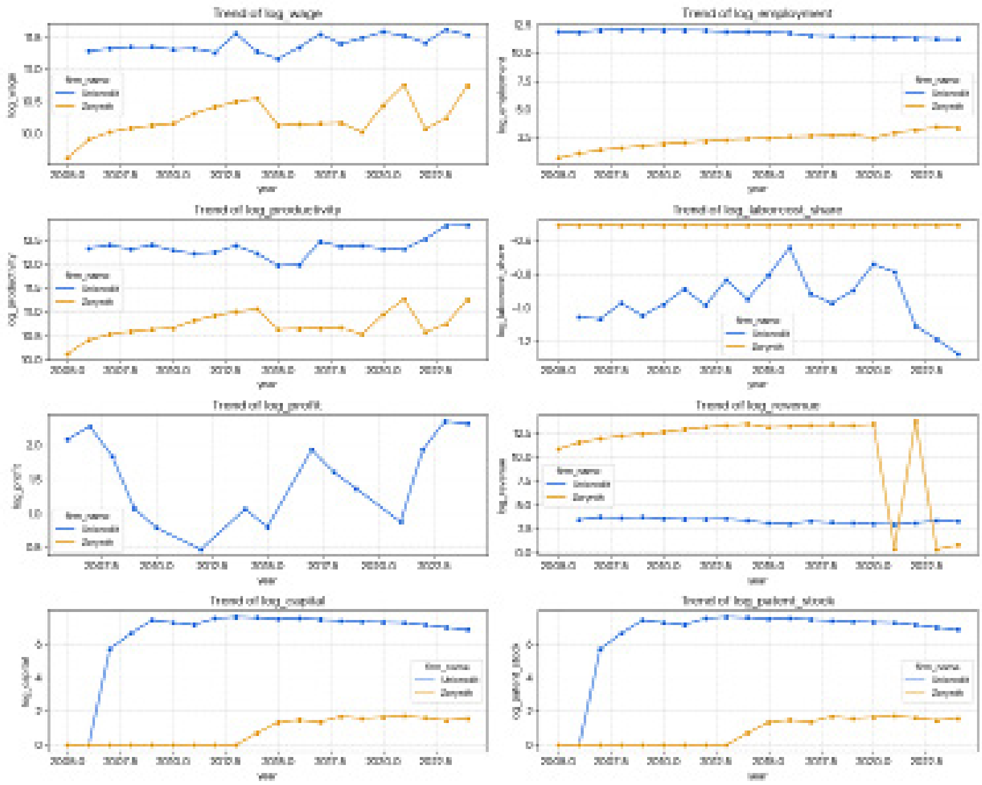

Figure 1: Trend graphs for UniCredit and Zerynth.

The comparative trend analysis between UniCredit and Zerynth from 2005 to 2024 highlights divergent trajectories in firm-level dynamics, reflective of their respective sectors. UniCredit, representing the mature financial sector, displays stable, upward trends in logged wages, employment, and productivity, alongside steady growth in AI patent stock—evidence of incremental, efficiency-focused innovation. In contrast, Zerynth, operating in the high-growth industrial tech sector, exhibits greater volatility, particularly in profit and labor cost share metrics, with a flatter trajectory in patent accumulation that suggests different R&D and scaling patterns. Symmetry histograms further reveal a distinctly bimodal structure across variables, shaped by the contrasting lifecycles of the two firms. UniCredit's distributions cluster around gradually declining employment and moderate productivity gains, consistent with a digitalizing yet mature institution. Zerynth's distributions skew toward rapid expansion of innovation and fluctuating labor demand, characteristic of a scaling tech firm. These patterns underscore the sector-specific nature of AI adoption and its heterogeneous labor-market impacts: while AI in finance appears to operate as a labor-substituting force that compresses cost shares, in the tech sector, it follows a more volatile, labor-complementing trajectory, where innovation coincides with hiring and restructuring. This descriptive evidence sets the stage for rigorous econometric testing of whether AI functions as a displacement or reinstatement force at the firm level.

The baseline fixed-effects regressions reveal a dualistic impact of AI patent stock on labor outcomes. In pooled models, AI innovation shows a positive and statistically significant relationship with wages (β = 0.2256), employment (β = 0.1775), and productivity (β = 0.2421), while exerting a small but significant negative effect on labor cost share (β = −0.0164). These results align with Skill-Biased Technical Change: AI increases average wages and productivity but also reallocates value away from labor. However, firm-specific analyses uncover stark heterogeneity. For UniCredit, AI patent stock raises employment (β = 0.0426) but has negligible effects on wages and productivity, whereas profit significantly boosts wages and productivity while reducing employment and labor share. For Zerynth, AI patent stock significantly reduces employment (β = −0.3060), wages (β = −0.3309), and productivity (β = −0.3309), reflecting short-run displacement and reorganization costs typical of scaling tech firms. Together, these baseline estimates underscore that AI's labor-market impact is not uniform but contingent on sectoral context, firm strategy, and financial health.

Table 3: Four Baseline Regression Models Results (Panel FE)

Model	Dept. var.	Independent	Coefficient	P Value	Sige	N	R^2	Prob > F
Model 1	Wage	Patent	0.2256	0.0000	***	40	0.22	0.0093
		Profit	2.3669	0.0000	***			
Model 2	Employment	Patent	0.1775	0.0000	***	40	0.19	0.0188
		Profit	2.2883	0.0000	***			
Model 3	Labor Share	Patent	−0.0164	0.0004	***	40	0.48	0.0000
		Profit	−0.2825	0.0000	***			
Model 4	Productivity	Patent	0.2421	0.0000	***	40	0.32	0.0008
		Profit	2.6493	0.0000	***			

Notes: *** p < 0.01, ** p < 0.05, * p < 0.10.
Source: Authors' evaluation using Stata 16.

Model 1: Wages, AI Patent Stock, Profit, and the SBTC Mechanism

The wage model shows that AI patent stock has a highly significant and positive impact on wages ($\beta = 0.2256$, $p < 0.01$). Interpreted elastically, a 10% increase in AI patent stock raises average wages by approximately 2.2%, a result fully aligned with the Skill-Biased Technical Change (SBTC) framework. In the Italian labor market – where digitalization progresses unevenly across sectors – this finding indicates that AI-intensive firms reward workers with higher pay as they reorganize production around more complex, knowledge-intensive and AI-complementary tasks. The positive and significant coefficient on profit strengthens this interpretation. With a coefficient of $\beta = 2.3669$, the model suggests that a 10% increase in profit is associated with a 23.669% rise in wages, highlighting the crucial role of financial capacity in shaping compensation dynamics. While profit does not directly 'buy' patents, higher profitability expands a firm's capacity to invest in AI-related R&D, increasing the likelihood of generating new patentable innovations in subsequent periods. This dynamic reinforces the wage-enhancing effect of AI patent stock, as financially strong firms are better positioned to undertake the complementary investments required for sustained technological upgrading. Profitable firms are better positioned to convert innovation-driven productivity gains into higher wages, invest in skill upgrading, and maintain competitive pay structures. Together, these results show that wage growth in AI-adopting Italian firms is driven by a dual mechanism: innovation increases the demand for skilled labor, while profitability provides the fiscal space to reward that labor accordingly.

Model 2: Employment, AI Patent Stock, Profit, and Displacement–Reinstatement Balance

The employment model shows that AI patent stock has a positive and highly significant effect on the employment growth rate ($\beta = 0.1775$, $p < 0.01$), indicating that AI innovation is associated with net job creation rather than displacement across the panel. Interpreted elastically, a 10% increase in AI patent stock raises the growth rate of employment by approximately 1.78%, reflecting the reinstatement effect whereby AI generates new tasks, complementary occupations, and expanded labor demand. This pattern is evident in firms that integrate AI into scalable production processes, where innovation stimulates organizational expansion rather than contraction. Profit also enters positively and significantly, with a coefficient of $\beta = 2.2883$, implying that a 10% rise in the growth rate of profit increases the growth rate of employment by roughly 22.883%. This large elasticity underscores the central role of financial strength in shaping labor outcomes: profitable firms are better positioned to absorb the adjustment costs of technological change, maintain workforce stability, and hire additional workers as AI systems are deployed.

These results contrast with the UniCredit-specific model, which showed pockets of labor substitution consistent with automation in finance. When the two firms are pooled, however, the fixed-effects estimates reveal that innovation in Italy tends to complement labor on average, particularly in high-growth or technology-intensive environments. The lagged structure of the AI patent variable is essential for capturing this dynamic. Employment effects do not materialize immediately; instead, they emerge after firms have reorganized production, retrained workers, and integrated AI tools into their operational routines. This temporal pattern reinforces the idea that innovation today shapes labor demand in the next period, highlighting the importance of modeling AI adoption as a gradual, path-dependent process rather than an instantaneous shock.

Model 3: Labor Cost Share, AI Patent Stock, Profit, and Distributional Pressures

The labor cost share model presents a nuanced distributional story. AI patent stock has a small but significant negative effect on labor's share of total costs ($\beta = -0.0164$, $p < 0.01$), indicating that as firms deepen their AI capabilities, labor becomes relatively less central in the cost structure. Interpreted elastically, a 10% increase in the growth rate of profit corresponds to an estimated 2.825% reduction in labor cost share, suggesting that as firms become more profitable, they tend to rely relatively more on capital-intensive processes or efficiency-enhancing technologies rather than expanding labor inputs proportionally. This stands in contrast to the positive wage effects observed elsewhere and highlights a deeper structural tension: while innovation (via AI patent stock) pushes labor's share downward through automation and capital deepening, rising profitability reinforces this decline by enabling firms to scale production without commensurate increases in labor expenditure. The lagged AI patent variable again underscores that these distributional adjustments unfold gradually, reflecting the time it takes for firms to reorganize their production mix, adopt new technologies, and recalibrate their skill composition.

Model 4: Productivity, AI Patent Stock, and Profit

The productivity model demonstrates that both AI patent stock and profit exert large, positive, and highly significant effects on firm-level efficiency. The coefficient on AI patent stock ($\beta = 0.2421$, $p < 0.01$) implies that a 10% increase in AI patent stock raises productivity by approximately 2.4%, once the technology has been absorbed into the firm's operational routines. This effect captures the post-adjustment phase of AI adoption: after an initial period of organizational disruption – where workflows are redesigned, employees adapt to new tools, and complementary investments are made – productivity begins to rise sharply. The lagged structure of the model is essential here, as it reflects the realistic delay between innovation and measurable performance gains. Profit amplifies this trajectory. With a coefficient of $\beta = 2.6493$ ($p < 0.01$), the model indicates that a 10% increase in profit is associated with a 26.49% increase in productivity. This exceptionally strong elasticity underscores the central role of financial capacity in unlocking the full benefits of AI. Profitable firms are better equipped to invest in the complementary assets that make AI effective – advanced training, software integration, data infrastructure, and organizational restructuring. In this sense, profit does not merely correlate with productivity; it magnifies the productivity-enhancing effects of AI by enabling firms to undertake the costly adjustments required for successful technological adoption.

Firm-Level Comparison (UniCredit vs Zerynth)

The sectoral and strategic differences between UniCredit and Zerynth give rise to markedly divergent impacts of AI innovation on their respective labor markets. For UniCredit, a large incumbent in the financial sector, AI patent stock exhibits a small but positive effect on employment ($\beta = 0.0426$, $p < 0.01$), suggesting mild job creation consistent with the generation of new high-skill roles in areas such as data analytics and digital banking. In contrast, profit exerts a negative effect on employment ($\beta = -0.0562$, $p < 0.10$), pointing to employment-saving restructuring driven by efficiency pressures. Wages at UniCredit are largely unaffected by AI patents but respond positively to profit ($\beta = 0.0931$, $p < 0.10$), indicating that compensation is shaped more by financial performance than by innovation intensity. Labor cost share declines with profit ($\beta = -0.1393$, $p < 0.05$), signaling a shift toward capital-intensive operations even as wages rise. Productivity responds positively to profit ($\beta = 0.2327$, $p < 0.01$) but not significantly to AI patents, reflecting the bank's ability to streamline processes using existing profitability rather than frontier innovation alone.

Table 4: Estimate results for Unicredit (Finance sector).

Dependent Variable	Coefficient (Patent)	P-Value	Coefficient (Profit)	P-Value
Employment	0.0426***	0.0017	-0.0562*	0.0699
Wage	0.0081	0.6283	0.0931*	0.0569
Labor Share	-0.0029	0.8896	-0.1393**	0.0277
Productivity	0.0111	0.6581	0.2327***	0.0048

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.
Source: Authors' evaluation using Stata 16.

Table 5: Estimate results for Zerynth (AI Tech sector).

Dependent Variable	Coefficient(Patent)	P-Value	Coefficient (Profit)	P-Value
Employment	-0.3060	0.0281	0.0125	0.2613
Wage	-0.3309	0.0369	-0.0077	0.5441
Productivity	-0.3309	0.0369	-0.0077	0.5441
Labor Share	-0.0000	0.5423	-0.0000	0.6622

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.
Source: Authors' evaluation using Stata 16.

Conversely, for Zerynth – an industrial-tech firm specializing in AI-driven IoT solutions – AI patent stock is associated with significant short-run reductions in employment ($\beta = -0.3060$, $p < 0.05$), wages ($\beta = -0.3309$, $p < 0.05$), and even productivity ($\beta = -0.3309$, $p < 0.05$). This pattern reflects the displacement and reorganization costs typical of scaling tech firms, where innovation initially disrupts workflows, compresses routine roles, and reallocates tasks before potential long-run gains materialize. Revenue growth shows no significant effect on labor outcomes, suggesting that Zerynth's expansion is currently decoupled from traditional employment and wage growth, possibly due to lean scaling, automation, or reliance on external contractors. Labor cost share remains statistically unchanged, indicating a stable cost structure amid technological flux.

Taken together, the comparison underscores a sector-specific duality: in service-based finance, AI acts mainly as a labor-augmenting force that supports employment and wage stability when coupled with profitability; in industrial tech, AI initially operates as a labor-substituting force that reduces headcount, depresses wages, and temporarily lowers productivity during intensive innovation phases. This heterogeneity highlights that the labor-market effects of AI are not technologically determined but are mediated by sectoral task structures, firm lifecycle stages, and strategic orientation.

To sum up, the findings reveal a dual pattern: while AI drives productivity and wage growth consistent with SBTC, it also reduces labor's share of total costs, signaling a capital-biased distribution of gains. Sectoral heterogeneity is pronounced: in UniCredit's financial services context, AI supports moderate employment growth and wage stability, largely mediated by profitability; in Zerynth's tech-intensive setting, AI initially acts as a disruptive, labor-substituting force, compressing employment, wages, and productivity during intensive innovation phases.

These results align with global evidence framing AI as a general-purpose technology that reallocates tasks rather than eliminating jobs outright (Acemoglu & Restrepo, 2019; Brynjolfsson *et al*., 2019), in contrast to more alarmist forecasts of mass displacement (Frey & Osborne, 2017). Within Italy, the findings resonate with observed trends of job polarization and skill mismatch (INAPP, 2023), yet advance the literature by directly linking these outcomes to firm-specific innovation metrics and strategic orientation.

5. Discussion

The empirical results reveal a clear divergence in how AI patent stock relates to labor outcomes in the two firms. For UniCredit, higher AI patent stock is associated with reductions in employment and labor cost shares alongside gains in productivity, while wage effects are modest. For Zerynth, AI patent stock is negatively associated with employment, wages, and labor cost shares, indicating a more pronounced labor-displacing pattern. These findings are consistent with task-based and routine-biased technological change frameworks, in which AI substitutes for routine tasks but may complement high-skill, non-routine work in specific organizational contexts (Acemoglu & Restrepo, 2018; Autor & Dorn, 2013).

The findings reveal a dual pattern: while AI drives productivity and wage growth consistent with Skill-Biased Technical Change (SBTC), it also reduces labor's share of total costs, signaling a capital-biased distribution of gains. Sectoral heterogeneity is pronounced: in UniCredit's financial services context, AI supports moderate employment growth and wage stability, largely mediated by profitability; in Zerynth's tech-intensive setting, AI initially acts as a disruptive, labor-substituting force, compressing employment, wages, and productivity during intensive innovation phases. Robustness checks using quadratic and dynamic panel specifications confirm these patterns are non-linear and persistent—early AI adoption raises wages and productivity, but marginal returns diminish as innovation accumulates, and past innovation exerts a long-tail influence on compensation and cost structures.

These results align with global evidence framing AI as a general-purpose technology that reallocates tasks rather than eliminating jobs outright (Acemoglu & Restrepo, 2019; Brynjolfsson *et al*., 2019), while contrasting with more alarmist forecasts of mass displacement (Frey & Osborne, 2017). Within Italy, the findings resonate with observed trends of job polarization and skill mismatch (INAPP, 2023), yet advance the literature by directly linking these outcomes to firm-specific innovation metrics and strategic orientation. Vivarelli *et al*. (2023) find that AI tends to complement high-skill technical roles. The decline in labor cost shares and employment at UniCredit and Zerynth is consistent with these patterns, although the magnitude and wage responses differ across sectors. At a broader level, Dalla Zuanna *et al*. (2024) and INAPP studies on AI occupational exposure and wage distribution suggest that AI-exposed occupations in Italy face both displacement risks and rising skill premia, a duality that our firm-level results also reflect.

This research contributes to the field in two key ways. First, it provides rare firm-level longitudinal evidence in the Italian context, moving beyond sectoral exposure indices to reveal how the same technology produces divergent labor outcomes depending on sectoral task structures and strategic adoption pathways. Second, it introduces AI patent stock as a forward-looking proxy for technological capability, offering a precise mechanism to assess how cumulative innovation activity translates into employment, wage, and productivity dynamics.

6. Conclusions and Policy Recommendations

This paper has examined how firm-level AI patent stock relates to employment, wages, productivity, and labor cost shares in two Italian firms operating in distinct sectors: UniCredit in banking services and Zerynth in industrial IoT. Using fixed-effects panel models over 2005–2024, we find that AI-related innovation is consistently associated with higher productivity but lower employment and labor cost shares in both firms. Wage effects are modest and mixed—slightly positive for UniCredit and more clearly negative for Zerynth. Taken together, these patterns suggest that AI functions primarily as a labor-displacing technology in our sample, consistent with task-based and routine-biased technological change frameworks, even as it delivers measurable productivity gains.

The sectoral contrast between UniCredit and Zerynth highlights that the labor-market consequences of AI are not technologically predetermined. Instead, they depend on task structure, organizational capabilities, and whether the firm is primarily an AI user (as in banking) or an AI producer (as in industrial IoT). Our results show that AI is reshaping Italy's labor mark*et al*ong two distinct trajectories: as a labor-augmenting, efficiency-enhancing force in service sectors, and as a labor-displacing, restructuring force in technology-intensive industrial firms. While AI innovation—proxied by AI patent stock—raises productivity and, in some cases, average wages, it simultaneously reduces labor's share of total costs, pointing to a capital-biased distribution of gains that may widen inequality.

The contrasting experiences of the two firms illustrate these dynamics. In banking, AI supports incremental digitalization and skill upgrading, yet profitability-driven restructuring can suppress employment growth. In industrial technology, AI drives product innovation and scaling, but often disrupts employment and wages during periods of rapid reorganization. These findings underscore the need for differentiated and proactive policy responses. In the service sector, policies should prioritize transition support, retraining, and wage protection for workers in routine roles. In the tech sector, policy should encourage skill development in AI-complementary fields and promote the creation of high-skill positions linked to AI deployment.

The negative association between AI patent stock and labor cost shares further indicates a shift in internal cost structures away from labor. To prevent this from translating into persistent wage stagnation or polarization, policymakers could consider wage-support measures and incentives for firms that invest in labor-augmenting AI applications—tools that enhance worker productivity rather than replace tasks entirely. For firms similar to Zerynth, innovation policies could be tied to commitments to internal skill development and high-skill job creation.

Finally, the productivity gains associated with AI patents highlight the importance of complementary investments in organizational capital and management practices. Industrial policies that support AI adoption in SMEs—through technical assistance, shared AI infrastructure, and improved access to finance—can help diffuse the benefits of AI beyond large firms, reducing the risk that Italy's digital transition exacerbates existing regional and firm-size disparities.

REFERENCES

Abrardi, L., Cambini, C., & Rondi, L. (2022). Artificial intelligence, firms and consumer behavior: A survey. *Journal of Economic Surveys*, *36*(4), 969–991. https://doi.org/10.1111/joes.12455

Acemoglu, D., & Restrepo, P. (2018). *Modeling automation* (NBER Working Paper No. 24321). National Bureau of Economic Research. http://www.nber.org/papers/w24321

Acemoglu, D., & Restrepo, P. (2019). Artificial intelligence, automation, and work. In *The economics of artificial intelligence* (pp. 197–236). University of Chicago Press. https://doi.org/10.7208/chicago/9780226613475.003.0008

Acemoglu, D., & Restrepo, P. (2020). The wrong kind of AI? Artificial intelligence and the future of labour demand. *Cambridge Journal of Regions, Economy and Society*, *13*(1), 25–35. https://doi.org/10.1093/cjres/rsz022

Aghion, P., Antonin, C., & Bunel, S. (2021). *The power of creative destruction: Economic upheaval and the wealth of nations*. Harvard University Press.

Autor, D., & Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the US labor market. *NBER Working Paper Series*, *15150*, 1689–1699.

Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, *118*(4), 1279–1333.

Barbieri, L., Mussida, C., Piva, M., & Vivarelli, M. (2020). Testing the employment and skill impact of new technologies. In *Handbook of labor, human resources and population economics* (pp. 1–27). Springer. https://doi.org/10.1007/978-3-319-57365-6_1-1

Bottero, M., & Schiaffi, S. (2022). *Firm heterogeneity and the impact of policies: Evidence from Italy* (Occasional Paper No. 674). Bank of Italy.

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

Calvino, F., & Fontanelli, L. (2026). Decoding AI: An early look at how French firms use AI. *Eurasian Business Review*, 51–93. https://doi.org/10.1007/s40821-026-00347-9

Cirillo, V., Evangelista, R., Guarascio, D., & Sostero, M. (2021). Digitalization, routineness and employment: An exploration on Italian task-based data. *Research Policy*, *50*(7), Article 104079. https://doi.org/10.1016/j.respol.2020.104079

Cirillo, V., Fanti, L., Mina, A., & Ricci, A. (2023). New digital technologies and firm performance in the Italian economy. *Industry and Innovation*, *30*(1), 159–188. https://doi.org/10.1080/13662716.2022.2055999

Colombo, E., Mercorio, F., & Mezzanzanica, M. (2019). AI meets labor market: Exploring the link between automation and skills. *Information Economics and Policy*, *47*, 27–37. https://doi.org/10.1016/j.infoecopol.2019.05.003

Costa, G. G., Pugliese, R., & Venier, F. (2026). Exploring artificial intelligence adoption among Italian firms: The AI readiness level. *International Journal of Business Information Systems*. Advance online publication. https://doi.org/10.1504/IJBIS.2025.10076399

Dalla Zuanna, A., Dottori, D., Gentili, E., & Lattanzio, S. (2024). An assessment of occupational exposure to artificial intelligence in Italy. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.5246679

Damioli, G., Van Roy, V., Vértesy, D., & Vivarelli, M. (2024). Drivers of employment dynamics of AI innovators. *Technological Forecasting and Social Change*, *201*, Article 123249. https://doi.org/10.1016/j.techfore.2024.123249

DiNardo, J., Fortin, N. M., & Lemieux, T. (1996). Labor market institutions and the distribution of wages, 1973–1992: A semiparametric approach. *Econometrica*, *64*(5), 1001. https://doi.org/10.2307/2171954

Felten, E. W., Raj, M., & Seamans, R. (2018). A method to link advances in artificial intelligence to occupational abilities. *AEA Papers and Proceedings*, *108*, 54–57. https://doi.org/10.1257/pandp.20181021

Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, *42*(12), 2195–2217. https://doi.org/10.1002/smj.3286

Filippucci, F., Gal, P., Jona-Lasinio, C., Leandro, A., & Nicoletti, G. (2024). *The impact of artificial intelligence on productivity, distribution and growth* (OECD Artificial Intelligence Papers No. 15). OECD Publishing.

Firpo, S., Fortin, N. M., & Lemieux, T. (2009). Unconditional quantile regressions. *Econometrica*, *77*(3), 953–973.

Firpo, S., Fortin, N. M., & Lemieux, T. (2018). Decomposing wage distributions using recentered influence function regressions. *Econometrics*, *6*(2), Article 28.

Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, *114*, 254–280.

Giczy, A. V., Pairolero, N. A., & Toole, A. A. (2022). Identifying artificial intelligence

(AI) invention: A novel AI patent dataset. *Journal of Technology Transfer*, *47*. https://doi.org/10.1007/s10961-021-09900-2

INAPP. (2023). *Rapporto annuale 2023: Lavoro e formazione*. Istituto Nazionale per l'Analisi delle Politiche Pubbliche (INAPP).

ISTAT. (2025a). *Imprese e ICT: Anno 2024*. Istituto Nazionale di Statistica.

ISTAT. (2025b). *Imprese e ICT: Anno 2025*. Istituto Nazionale di Statistica.

Manestra, S., & Morelli, A. P. (2013). *Questioni di economia e finanza* (Occasional Papers No. 727). Banca d'Italia. https://www.bancaditalia.it/pubblicazioni/qef/2022-0727/QEF_727_22.pdf

Misch, F., Park, B., Pizzinelli, C., & Sher, G. (2025). *Artificial intelligence and productivity in Europe* (IMF Working Paper No. WP/25/67). International Monetary Fund.

Montobbio, F., Staccioli, J., Virgillito, M. E., & Vivarelli, M. (2024). The empirics of technology, employment and occupations: Lessons learned and challenges ahead. *Journal of Economic Surveys*, *38*(5), 1622–1655. https://doi.org/10.1111/joes.12601

Webb, M. W. (2020). *Essays in the economics of artificial intelligence* [Doctoral dissertation, Stanford University].