ECONOMIA ITALIANA

Fondata da Mario Arcelli

Intelligenza artificiale: da innovazione tecnologica a crescita economica inclusiva

2026/2

LUISS
CASMEF Centro Arcelli per gli Studi Monetari e Finanziari

UNIVERSITÀ CATTOLICA del Sacro Cuore
CESPEM
Centro Studi di Politica economica e monetaria "Mario Arcelli"

Economia Italiana

Fondata da Mario Arcelli

COMITATO SCIENTIFICO

(Editorial board)

CO-EDITORS

FRANCESCO NUCCI - Sapienza, Università di Roma
PAOLA PROFETA - Università Bocconi
PIETRO REICHLIN - Luiss Guido Carli

MEMBRI DEL COMITATO *(Associate Editors)*

Carlotta Berti Ceroni
Università di Bologna

Massimo Bordignon
Università Cattolica del Sacro Cuore

Lorenzo Codogno
London School of Economics and Political Science

Giuseppe De Arcangelis
Sapienza, Università di Roma

Stefano Fantacone
Centro Europa Ricerche

Emma Galli
Sapienza, Università di Roma

Paolo Giordani
LUISS Guido Carli

Giorgia Giovannetti
Università di Firenze

Enrico Giovannini
Università di Roma "Tor Vergata"

Andrea Montanino
Cassa Depositi e Prestiti

Salvatore Nisticò
Sapienza, Università di Roma

Alessandro Pandimiglio
Università degli Studi "Gabriele d'Annunzio" Chieti - Pescara

Alberto Petrucci
Luiss Guido Carli

Gustavo Piga
Università di Roma "Tor Vergata"

Beniamino Quintieri
Università di Roma "Tor Vergata"

Fabiano Schivardi
LUISS Guido Carli

Marco Spallone
Università degli Studi "Gabriele d'Annunzio" Chieti - Pescara

Francesco Timpano
Università Cattolica del Sacro Cuore

Mario Tirelli
Università Roma Tre

Giovanna Vallanti
LUISS Guido Carli

DIRETTORE RESPONSABILE *(Editor in Chief)*: Giovanni Parrillo

---

ADVISORY BOARD

PRESIDENTE *(President)*

PAOLO GUERRIERI - Sapienza, Università di Roma e PSIA, SciencesPo

CONSIGLIO *(Members)*

Federico Arcelli, Center for International Governance Innovation
Riccardo Barbieri, Tesoro
Mariano Bella, Confcommercio
Piero Cipollone, Banca Centrale Europea
Stefano Costa, Istat
Giorgio Di Giorgio, Editrice Minerva Bancaria
Dario Focarelli, Ania
Riccardo Gabrielli, Deloitte
Mauro Micillo, Intesa Sanpaolo
Stefano Micossi, Assonime
Francesco Minotti, "Luigi Luzzatti" s. c. p. a.
Roberto Monducci, Osservatorio Imprese Estere
Fabiola Riccardini, Istat
Alberto Tosti, Sara Assicurazioni
Marco Vulpiani, Deloitte

Economia italiana

Fondata da Mario Arcelli

Intelligenza artificiale: da innovazione tecnologica a crescita economica inclusiva

numero 2/2026
Pubblicazione quadrimestrale
Roma

ECONOMIA ITALIANA
Rivista quadrimestrale fondata nel 1979 da Mario Arcelli

DIRETTORE RESPONSABILE
Giovanni Parrillo, Editrice Minerva Bancaria

COMITATO DI REDAZIONE
Simona D'Amico *(coordinamento editoriale)*
Francesco Baldi
Nicola Borri
Giuseppina Gianfreda
Rita Mascolo
Guido Traficante

(Pubblicità inferiore al 70%)

Autorizzazione Tribunale di Roma n. 43/1991

ISSN: 0392-775X

Gli articoli firmati o siglati rispecchiano soltanto il pensiero dell'Autore e non impegnano la Direzione della Rivista.
I *saggi* della parte monografica sono a invito o pervengono a seguito di call for papers e sono valutati dall'editor del numero.
I *contributi* vengono valutati anonimamente da due referee individuati dagli editor o dalla direzione e redazione con il supporto dei membri del Comitato Scientifico.
Le *rubriche* sono sottoposte al vaglio della direzione/redazione.

Finito di stampare nel mese di settembre 2026

www.economiaitaliana.org - redazione@economiaitaliana.org

Editrice Minerva Bancaria srl

DIREZIONE E REDAZIONE Largo Luigi Antonelli, 27 – 00145 Roma

AMMINISTRAZIONE EDITRICE MINERVA BANCARIA S.r.l.
presso Ptsclas, Viale di Villa Massimo, 29
00161 - Roma
amministrazione@editriceminervabancaria.it

Segui Editrice Minerva Bancaria su:

Sommario

Intelligenza artificiale: da innovazione tecnologica a crescita economica inclusiva

EDITORIALE

5 Intelligenza artificiale: da innovazione tecnologica a crescita economica inclusiva
A. Bonfiglioli, G. Gancia, P. Profeta

SAGGI

13 Intelligenza artificiale e produttività: canali, evidenze empiriche e sfide per l'Italia
F. Calvino, A. Ughi

109 Consequences of AI-Driven Chatbots in Economic Domains
S. Galletta

153 Workers and Generative AI: Evidence from Italy
M. Anelli, I. Colantone, A. Gallego, P. Stanig

227 Exposure to AI, Risk Sharing and Within-Household Inequality
E. Colombo, F. Trentini

279 AI and Workplace Welfare Plans
M. Caselli, E. Cristofoletti, A. Fracasso

INTERVENTI

323 L'Intelligenza Artificiale nei mercati finanziari italiani: dall'analisi all'azione
C. Di Noia

RUBRICHE

345 Oltre Hormuz. Sicurezza energetica e nuove dipendenze
M. Parco

RECENSIONI

357 R. Brunetta (a cura di) *Corpi Intermedi, società civile ed economia sociale. Per un governo delle transizioni*
M. Rotili

AI and Workplace Welfare Plans

Mauro Caselli *
Enrico Cristofoletti **
Andrea Fracasso *

Abstract

Using the 2022 wave of the Italian Permanent Census of Enterprises, this work studies whether firms that adopt artificial intelligence (AI) are also more likely to invest in workplace well-being and occupational health and safety plans. Relying on a direct survey-based measure of AI adoption, we estimate the association between AI adoption and worker welfare measures through inverse probability weighted regression adjustment to address selection-related issues. We find that AI-adopting firms are 15 and 14 percentage points more likely than non-adopting firms to introduce well-being and health and safety plans, respectively. The association with well-being plans is positive for all firm groups analysed and is broadly uniform across them. In contrast, the as-

* School of International Studies & Dep. of Economics and Management, University of Trento, Italy, mauro.caselli@unitn.it, andrea.fracasso@unitn.it

** School of International Studies ,University of Trento, Italy, enrico.cristofoletti@unitn.it

The authors would like to acknowledge the financial support provided by the European Union's Horizon Europe research and innovation programme through the SkillAbility project under Grant Agreement No 101177783. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them. The authors are solely responsible for any errors.

sociation with health and safety plans is more heterogeneous. More precisely, it is weaker for firms with stronger territorial commitment, larger size, hiring constraints, and in industrial sectors. The evidence suggests that firms tend to complement AI adoption with worker welfare measures, although this differs across welfare domains due to the organisational constraints associated with the introduction of H&S plans.

Sintesi - IA e piani di welfare aziendale

Utilizzando l'edizione 2022 del Censimento permanente delle imprese dell'ISTAT, questo lavoro studia se le imprese che adottano l'intelligenza artificiale (IA) siano anche più propense a investire in piani per il benessere dei lavoratori e per la salute e sicurezza sul lavoro. Facendo ricorso a una misura diretta dell'adozione dell'IA basata su specifiche domande dell'indagine, l'analisi stima il grado di associazione tra adozione dell'IA e misure di welfare dei lavoratori mediante il metodo dell'inverse probability weighted regression adjustment, al fine di affrontare i problemi di selezione. I risultati indicano che le imprese che adottano l'IA hanno, rispettivamente, una probabilità superiore di 15 e 14 punti percentuali rispetto alle imprese che non lo fanno, di introdurre piani per il benessere e piani per la salute e sicurezza sul lavoro. L'associazione con i piani per il benessere lavorativo è positiva per tutti i gruppi di imprese analizzati e risulta sostanzialmente uniforme. Al contrario, l'associazione con i piani per la salute e sicurezza sul lavoro presenta una maggiore eterogeneità, essendo essa più debole per le imprese caratterizzate da un maggiore radicamento territoriale, da dimensioni più elevate, da vincoli nelle assunzioni e per quelle appartenenti ai settori industriali. L'evidenza suggerisce che le imprese tendono ad accompagnare l'adozione dell'IA con misure di welfare rivolte ai lavoratori, sebbene tale relazione differisca tra le diverse dimensioni del welfare a causa dei maggiori vincoli organizzativi associati all'introduzione di piani per la salute e sicurezza sul lavoro.

JEL Classification: O33; J28; J81.
Parole chiave: IA; Benessere lavorativo; Salute e sicurezza dei lavoratori; Welfare dei lavoratori.
Keywords: AI; Worker well-being; Worker health and safety; Worker welfare.

1. Introduction

The increasing diffusion of digital technologies, particularly artificial intelligence (AI), is reshaping the organisation of work in fundamental ways. Much of the existing literature has focused on the impact of these technologies on employment, wages, productivity, and task allocation within firms (Acemoglu et al., 2022; Aldasoro et al., 2026; Bonfiglioli et al., 2024; Brynjolfsson et al., 2025; Eloundou et al., 2023; Felten et al., 2021; Manning and Aguirre, 2026; Jiang et al., 2025). However, considerably less attention has been devoted to how firms adapt the conditions under which work is carried out. A key and still unresolved question in this respect is whether the adoption of AI leads firms to strengthen or weaken their commitment to protecting workers, in particular by investing in occupational health and safety monitoring systems or broader well-being practices. As technological innovations can be a source of workers' concerns about job security, bargaining power and hold-up opportunities (Zirar et al., 2023; Golgeci et al., 2025), it is possible that adopting firms find it convenient to accompany technology adoption with reassuring measures promoting workers' well-being (WB, hereafter) and health and safety (H&S). However, there might be firm-specific incentives and constraints that discourage firms from providing welfare-improving initiatives while adopting AI, including limited organisational and managerial capacity and workforce composition.

This paper investigates whether the adoption of AI and other digital technologies in Italy between 2020 and 2022 is associated with firms investing in worker WB and H&S plans and practices. Drawing on data from the 2022 wave of the Italian Permanent Census of Enterprises, we distinguish between specific AI adoption and investments in more broadly defined AI application

areas. Another feature of this paper is the distinction between two different dimensions of worker welfare: workplace WB plans and occupational H&S plans. While both can be interpreted as investments in worker welfare, they differ substantially in their organisational foundations. WB plans are generally more discretionary and managerial in nature, whereas H&S plans are more closely linked to formal organisational structures, operational procedures, and risk-management systems. Distinguishing between these two dimensions allows us to investigate whether AI adoption is associated with different forms of worker-oriented investments. In what follows, for the sake of brevity, we will use the expression 'worker welfare' to include both WB and H&S plans and practices.

The estimation of the causal relationship between the adoption of AI technologies, on the one hand, and measures of worker welfare, on the other hand, is marred with selection issues and omitted variable problems. The disparities in investment in worker welfare can be attributed to the divergent capabilities, resources and inclinations of companies in regard to innovation. For example, companies with the capacity to innovate may also allocate a greater proportion of their investment towards WB and H&S plans, while those lacking this capacity may devote a smaller proportion of resources and efforts. Thus, OLS estimates are likely to be biased. The methodological approach adopted in this work, namely the inverse probability regression adjustment (IPWRA), serves to mitigate the potential bias arising from observable confounders affecting both firms' decisions. Furthermore, IPWRA offers protection against misspecification in either the model for the probability of treatment or the model for the outcome. The inclusion of several control variables and various fixed effects in the estimations serves to strengthen the approach. Nevertheless, given that issues associated with unobserved omitted

variables cannot be entirely ruled out, we interpret our findings with care and refrain from making strong causal claims.

In addition, in this work we explore how the relationship between AI adoption and worker welfare varies across firms. Firms differ not only in their likelihood of investing in digital technologies and worker welfare-oriented practices, but also in the constraints and incentives they face in adopting all of them at the same time. For example, firms with limited financial resources may face a trade-off between introducing new technologies and expanding expensive worker WB plans. Larger firms, instead, may be more prone to introduce worker welfare measures, independently from their technological investments. Moreover, occupational H&S systems and AI-based technologies may both entail challenging organisational interventions because their effective implementation requires integration with operations, reorganisation of routines, redistribution of responsibilities, and new human resource management practices (Uhrenholdt Madsen et al., 2020; Valtonen et al., 2025): this may represent a constraint that makes it more difficult for certain companies to adopt them jointly. By examining heterogeneity across territorial commitment, firm size, hiring constraints, and sector, the paper provides new evidence on the conditions under which AI adoption is associated with further investments in worker welfare.

In a nutshell, our main findings are that AI is positively associated with investments in both worker WB and H&S. In particular, AI-adopting firms are 15 and 14 percentage points more likely to introduce WB and H&S plans, respectively. This suggests that firms frequently accompany technology adoption with worker-related supportive measures. Nevertheless, WB and H&S prove to be different as they respond to different types of incentives. Indeed, we find evidence of a heterogeneous relationship between the introduction of

AI and H&S plans, but not with WB plans. We argue that this heterogeneity is at least in part driven by the greater need of formal organisational structures for investment in H&S plans relative to WB plans.

Our paper contributes to two related strands of the literature. The first strand examines how AI and digital technologies affect workers and how firms adapt their organisational and human resource management practices in response to technological change. Existing studies provide mixed evidence on the consequences of AI adoption for workers' WB. While Nazareno and Schiff (2021) document effects on stress, job satisfaction, and health, Schulz et al. (2026) show that the relationship between AI adoption and job satisfaction may be non-linear, whereas Giuntella et al. (2025) find no sizeable negative effects on workers' WB and mental health. Other contributions emphasize that workers' reactions to technological change depend on organisational conditions and human resource management practices. For instance, Caselli et al. (2025) show that the relationship between technological innovation and workers' job insecurity is moderated by the presence of workforce-oriented human resource strategies. Similarly, Zirar et al. (2023) show that workers' trust in AI depends on the skills that facilitate human-AI coexistence, while Kim and Lee (2025) find that worker welfare provisions mitigate the increase in perceived job insecurity associated with AI adoption. More broadly, Golgeci et al. (2025) and Valtonen et al. (2025) stress the importance of organisational processes in supporting the integration of AI within firms. While this literature has mainly focused on workers' perceptions and labour-market outcomes, considerably less attention has been devoted to whether firms accompany AI adoption with concrete investments in worker welfare. One exception is Chen et al. (2024), who find that AI adoption in Chinese firms is associated with stronger investment in environmental governance and so-

cial responsibility, which often includes worker-related initiatives.[1] However, whether firms adopting AI are also more likely to introduce workplace WB and occupational H&S plans remains largely an open empirical question. By distinguishing between these two dimensions of worker welfare, our paper contributes to filling this gap.

The second strand of the literature concerns firms' incentives to invest in worker welfare (Kitzmueller and Shimshack, 2012). Beyond ethical considerations, firms may find it profitable to provide worker welfare measures for several reasons. Such investments may improve firms' reputation among consumers and other stakeholders (Kitzmueller and Shimshack, 2012; Elfenbein and McManus, 2010), reduce labour costs by increasing firms' attractiveness to workers (Tonin and Vlassopoulos, 2015; Kitzmueller and Shimshack, 2012), and enhance job satisfaction and productivity (Cassar and Meier, 2018; Nielsen et al., 2017; Judge et al., 2001). These mechanisms are particularly relevant in the context of AI adoption, as firms may use worker welfare measures to facilitate organisational change, reassure employees, and support the integration of new technologies into the workplace.

The remainder of the article proceeds as follows. Section 2 describes the dependent and independent variables, as well as the conditioning variables and the controls used in our analysis. Section 3 introduces the model and the empirical method adopted in this work, that is the inverse probability regression adjustment. Section 4 presents the main results and some robustness checks. Section 5 provides some concluding remarks.

1 Notably, neither we nor Chen et al. (2024) focus on the relationship between AI and worker welfare that stems from the adoption of AI applications specifically designed to monitor and reduce workers' injuries. We also do not restrict our attention to welfare measures concerning ethical principles and guidelines for AI adoption itself, as in Bonsón et al. (2021).

2. Data

The data used in our empirical estimation come from the Italian Permanent Census of Enterprises (Census hereafter) by the Italian Institute for Statistics (ISTAT), and in particular the wave carried out in 2022.

Every three years since 2019, ISTAT has gathered qualitative information on the Italian productive system, surveying Italian firms on several topics ranging from the adoption of novel technologies to firm organisation, from internationalisation to ecological and social practices adopted by the organisations. The Census is representative of the universe of the Italian enterprises as it covers all the firms with more than 20 employees and the remaining companies are sampled according to a stratified design. The literature using the Census as the main source of information is extremely vast given its structure and reliability.

The 2022 wave of the Census explores AI and other digital technologies by asking precise questions about firms' adoption over the three-year period 2019-2022. Our focal explanatory variable, defined as *AI adoption*, is based on the response to the question: "During the three-year period 2020-2022, did the enterprise use the following enabling technologies to innovate processes, goods, and services? Artificial intelligence." The response takes a value of 1 if the enterprise adopted the technology, and 0 otherwise.

Consistently with recent literature (Bahoo et al., 2023), ISTAT defines AI as "the ability of a machine to demonstrate the capacity to establish logical connections, learn concepts and processes, and plan activities. In the business sector, the most common AI solutions involve the monitoring of places or people, process automation, interaction with human beings, and support for

their activities" (ISTAT, 2022 p. 15, authors translation).[2]

The possibility of exploring a precise survey-based measure of AI adoption at the firm level constitutes an advantage for this work with respect to previous studies that had to resort either to indirect proxies of adoption (such as workers' potential exposure) or to broader measures of investment in AI-related technologies covering very broad areas of application. Montresor and Vezzani (2023) and Nucci et al. (2023), for instance, analyse older waves of the Census and measure AI adoption with a dummy variable equal to 1 if the firm invested in at least one of the technological areas that ISTAT defines as 'Areas of application of artificial intelligence' (i.e., immersive technologies, elaboration and analysis of big data, advanced automation, collaborative robots, and smart systems). The rationale is that firms investing in AI are likely to report investments in these broader areas, as they encompass a wide range of potential AI applications. However, these categories may also include other technologies (e.g., advanced robots or big data analytics not based on AI), potentially leading to an overestimation of AI diffusion and biasing the estimates. For a robustness check, we present the results using this different measure of AI adoption based on investments in AI application areas and constructed exactly as in Montresor and Vezzani (2023) and Nucci et al. (2023), defined as *AI application areas*. While the overall results are similar to those obtained using the first and more precise measure of AI adoption, some differences indicate an incomplete overlap between the two measures. The possibility of using a precise survey-based measure and an alternative proxy allows us to explore empirically the consequences of different solutions to operationalise the concept of AI adoption, and this represents another contribution to the literature.

2 This definition, together with the time period the survey refers to, excludes firms' investment in new forms of Generative AI as these latter started spreading in Italy only after the year 2023.

Regarding the dependent variables associated with worker welfare, we use two proxies built again on the basis of the Census questionnaire. The first variable captures whether the firm implemented plans for monitoring and increasing workers' WB during the years 2021 and 2022 (*Well-being plans*). Although ISTAT does not provide a precise definition of WB, according to the literature it can be interpreted as the "overall quality of an employee's experience and functioning at work" (Grant et al., 2007 p. 52). In practice, the Census survey seems to consider a number of human resource measures that improve workers' WB: the provision of opportunities for social interaction at work, the enhancement of competencies through training and development programmes, support for career-related activities, as well as employee-related surveys, and forms of collective representation.

The second variable related to worker welfare is built using a question about whether the firm adopted a plan to monitor worker H&S during the years 2021-2022 (*Health and safety plans*). This kind of activity involves identifying and applying a set of procedures to manage risks and systematically improve occupational H&S. For example, firms can implement occupational H&S management systems that comply with the ISO 45001 standard (ISO, 2018). It should be noted that there are potentially two forms of investments in H&S plans in the Italian context: mandatory measures to ensure compliance with Law *D.Lgs.* 81/2008,[3], and more discretionary measures, such that the implementation of a safety management system. Our variable refers to this second type of workers' H&S measures.

WB and H&S measures exhibit heterogeneous diffusion across firms: only 15% of the companies in the sample adopted worker WB plans, whereas circa 40% of them implemented some plans for monitoring H&S during the period of interest.

3 The law requires implementing some measures concerning health and worker safety (e.g., risk evaluation).

In our empirical exercise, we not only estimate the direct relationship between the adoption of AI and measures to strengthen worker welfare, but also look at factors that can moderate this relationship.

The analysis considers whether the commitment of firms to the territory in which they operate may influence the relationship between AI adoption and its efforts to introduce measures that improve worker welfare. We build a dummy variable that takes value 1 if the firms had collaborations with schools and associations, and/or performed philanthropic activities over the 2021-2022 period (*Territorial commitment*). At the theoretical level, how firms' commitment to the territory interacts with AI in shaping the introduction of measures for worker welfare is *a priori* unclear. On the one hand, firms adopting AI may be willing to invest in both worker welfare and activities external to the firm to show that technological upgrading is not financed by shifting the adjustment costs onto workers; hence, plans to improve worker welfare may be the internal counterpart of the external commitment of the firm. Moreover, firms adopting AI technologies and committed to the territory may pursue an inclusive form of digitalisation that includes worker welfare measures and that adds to the other efforts to run a socially sustainable business. On the other hand, external social action can crowd out internal measures devoted to worker welfare because the latter are less visible and more recurrent in cost than external initiatives.[4] Moreover, firms that introduce AI and digital technologies to save labour may cooperate with schools and other local actors as a recruitment device, and then lower the resources directed towards incumbent workers.[5]

4 A survey of 932 human resources leaders and 2036 employees conducted by Achievers in October 2021 found that only 25% of the employees acknowledged their firms' efforts to support their well-being, whereas almost half of leaders acknowledged the implementation of measures to support employee wellbeing and conditions (Achievers, 2022).

5 Muehlemann (2025) explores whether German firms train their existing workforce or acquire necessary skills

Furthermore, we explore the heterogeneity of the relationship across firm size. We build a dummy variable that takes value 1 if the company has 50 employees or more (*Medium/large*). Also in this case, two contrasting cases are possible. Larger AI-adopting firms may have stronger incentives to stabilise their skilled workforce via new welfare-related measures, but it is also possible that smaller firms have a relatively stronger preference to use these welfare measures, rather than wage increases, to please their workers. Furthermore, it is possible that not all small firms are capable of introducing these technologies and worker welfare plans at once, and one of them may crowd out the other. It is also plausible that this crowding out effect is stronger for the welfare measures that are more expensive and discretionary. In addition, it should be considered that process innovation (Edquist et al., 2001) often needs to be accompanied by organisational innovations in order to be effective (Hervas-Oliver et al., 2016; Reichstein and Salter, 2006), especially in small firms with a less formal organisational structure.

We then examine whether the relationship between AI adoption and worker welfare is moderated by firms' difficulties in hiring new personnel. In particular, we focus on whether the firm faces challenges in finding workers with sufficient cross-competencies (*Cross-competencies constraint*) or technical competencies (*Technical competencies constraint*) during 2021-2022. These two dummy variables can help to capture the possibility that firms face a trade-off between hiring externally and investing on training incumbent workers. While firms investing in AI may seek to reassure incumbent workers by supporting their welfare, those aiming to replace existing workers with newly skilled hires may have weaker incentives to invest in welfare measures.

externally. He finds that AI adoption is associated with an increase in new apprenticeships (not training decision) and with a more modest increase in continuing training concentrated on the high-skilled employees.

However, the interpretation of these two dummy variables requires caution. On the one hand, a firm reporting difficulties in hiring competent workers implicitly indicates that it has been seeking specific skills through external recruitment. On the other hand, firms facing such difficulties represent a subset of those attempting to hire, and their inability to find suitable candidates may reflect firm-specific constraints rather than broader human resource management practices. Furthermore, the need to recruit new workers may signal a misalignment between the firm's existing skill base and the competencies required to fully integrate AI into the organisation. This condition may keep firms in an early stage of AI adoption, in which organisational changes have not yet been implemented, thereby reducing the likelihood of introducing worker welfare plans that require more complex organisational arrangements (Huang et al., 2023; Zhu et al., 2006; Cooper and Zmud, 1990).

Finally, we investigate whether firms in the manufacturing sector (*Industry*) exhibit different patterns compared to firms in other sectors. The ways in which AI is integrated into production processes vary across industries, potentially leading to different effects on firms' labour demand. For instance, firms in the industrial sector are more likely to adopt robotics, technologies often associated with labour substitution (Acemoglu and Restrepo, 2020), whereas firms in service sectors are more likely to invest in big data analytics, which tend to enhance productivity and increase labour demand (Ladeira et al., 2024). Within our framework, a reduction in labour demand lowers the incentive to invest in worker welfare practices, while the opposite holds when technology complements labour. A complementary explanation concerns specifically H&S practices. To the extent that the adoption of AI reduces workers' exposure to certain risks and improves safety (Patel et al., 2018), the marginal incentive to invest in additional H&S measures may decline. This,

in turn, may weaken the relationship between AI adoption and H&S plans.[6]

All the above mediating variables also enter the specification by themselves, as they can have a direct effect on the probability of adopting measures to enhance worker welfare. Similar to previous studies focusing on AI and innovation (Montresor and Vezzani, 2023), our estimates also include a number of additional controls. We include a dichotomous variable for firms' access to public incentives during the period 2021-2022 (*Public Incentives*). We also account for whether firms borrow from one or more banks in 2022 (*Lending*).[7] Following Montresor and Vezzani (2023), we build a proxy for financial availability with a dummy variable that is equal to 1 if the firm identifies either self-financing or equity (through initial public offering or own means) as its main source of funding in 2022 (*Internal financing*). Furthermore, we include a control for family ownership (*Family firm*) as the application of AI solutions in family businesses is characterised by specific strengths and challenges (Kumar and Ratten, 2024). In order to avoid confounding effects associated with other innovation-related variables, we also control for investment in R&D: the dichotomous variable *High R&D* assumes value 1 if the firm declares a high investment intensity in R&D in the years 2020-2022.[8] By the same token, we control for investments in digital technologies that are not areas of application of AI (*Non-AI application areas*). The non-AI application areas include internet based technologies (i.e., ultra-broadband fibre optic internet connection, mobile internet connection 4G-5G, Internet of Things (IoT)) and other digital technologies, such as 3D printers, simu-

6 We acknowledge that, due to data limitations, we are unable to distinguish between AI-powered technologies explicitly designed to reduce work-related injuries and more general AI applications (Shah and Mishra, 2024; European Agency for Safety and Health at Work, 2023).

7 The inclusion of this control forces us to remove all financial institutions from the sample, which is a standard approach in studies using this firm-level dataset.

8 The other response options available to firms in the questionnaire are medium intensity, low intensity, and no investment.

lation between interconnected machines, blockchain for product or process certification, and cybersecurity.

Finally, we control for unobserved regional and sectoral heterogeneity by including province-level (NUTS3) and industry (2-digit NACE) fixed effects in the model. These controls account for spatial and industry-specific factors that may influence both AI adoption and worker welfare measures.

Table 1 **Descriptive Statistics**

	Mean	Std. dev.	Min	Max
Main explanatory variables				
AI adoption	0.0154	0.1232	0	1
AI application areas	0.0943	0.2922	0	1
Outcome variables				
Well-being plans	0.1588	0.3655	0	1
Health and safety plans	0.3871	0.4871	0	1
Conditioning variables and controls				
Territorial commitment	0.1531	0.3601	0	1
Medium/large	0.1507	0.3578	0	1
Technical competencies constraint	0.3490	0.4767	0	1
Cross-competencies constraint	0.1912	0.3932	0	1
Industry	0.3193	0.4762	0	1
Public Incentives	0.2496	0.4328	0	1
Lending	0.3241	0.4681	0	1
Internal financing	0.7734	0.4186	0	1
Family firm	0.7243	0.4469	0	1
High R&D	0.0499	0.2178	0	1
Non-AI application areas	0.5270	0.4993	0	1

Notes: The table shows mean, standard deviation, maximum, and minimum for the reported variables. The total number of observations is 172202.

Table 1 provides basic descriptive statistics for all the variables listed above, divided into main explanatory variables, outcome variables, and conditioning variables and controls.

3. Empirical Model

Although we use a framework that can potentially identify average treatment effects, we refrain from making strong causal claims given data limitations that prevent us from exploiting a longitudinal dimension or including a richer set of controls. Accordingly, we refer to our empirical strategy as estimating the average relationship between AI adoption on firm-level worker welfare measures. We further examine heterogeneity by conditioning on observable characteristics, thereby exploring conditional relationships. Indeed, while our empirical approach helps mitigate selection bias, it does not fully address all potential sources of endogeneity.

Let $Y_i(1), Y_i(0)$ be the potential outcomes when firm i is, respectively, treated and not treated. The average treatment effect is then defined as: $ATE \equiv E[Y_i(1) - Y_i(0)]$. When X_i^C is a conditioning variable, the conditional average treatment effect is then calculated as $CATE(x_i^C) \equiv E[Y_i(1) - Y_i(0) | X_i^C = x_i^C]$ as shown by Abrevaya et al. (2015).[9]

Let D_i be a variable indicating whether the firm i has been treated or not, the variable Y_i is the observed outcome and the term $\mathbf{X_i}$ a vector of exogenous covariates. Under unconfoundedness (i.e., $Y_i(1), Y_i(0) \perp D_i | \mathbf{X_i}$),

9 When the conditioning variable is continuous, many scholars estimate the conditional average treatment effect function that maps each point x_i^C in the range of X_i^C to the corresponding $CATE(x_i^C)$. See Fan et al. (2022) and Abrevaya et al. (2015) for an in-depth discussion.

it holds that (Cerulli, 2015):

$$ATE = E[E[Y_i|D_i = 1, \mathbf{X}_i] - E[Y_i|D_i = 0, \mathbf{X}_i]] \quad (1)$$

As anticipated in the Introduction, the estimation of the relationship between AI adoption and measures of worker welfare is marred by selection issues and omitted variable issues. In order to take into account the problem of self-selection into treatment, we use the inverse probability regression adjustment estimator, introduced by Wooldridge (2007).[10] The IPWRA estimator combines inverse probability weighting with regression adjustment. Specifically, it uses weighted regression coefficients, where the weights are given by the estimated inverse probabilities of treatment, to compute the average values of treatment-specific predicted outcomes, thereby extending a standard regression adjustment approach. In a simple regression adjustment framework, the outcome equation (i.e., $E(Y_i|D_i = j, \mathbf{X}_i), j = 0,1$) is estimated via linear regression, and the average treatment effect is obtained by averaging predicted outcomes across the sample. By contrast, the IPWRA approach weights the regression by the propensity score, which is estimated in a first-stage model of treatment assignment. Importantly, the IPWRA estimator is 'doubly robust', meaning that it remains consistent if either the propensity score model or the outcome equation is correctly specified.

When implementing this method, we use logistic regressions to estimate the outcome equation as well as the propensity score. All controls and conditioning variables appear in both equations. Since our conditioning variables are dichotomous, we recover $CATE(x_i^C)$ by splitting the sample according to the conditioning variable and by calculating the ATE on the resulting

10 For the implementation, we follow the code by Jann (2019), with some modifications.

subsamples (Abrevaya et al., 2015). Full-sample ATE are the estimates obtained on the entire sample.

We also present some robustness checks using an alternative estimation procedure developed in Fan et al. (2022) to obtain conditional average treatment effects. This procedure is more flexible in that it does not rely on any parametric assumptions, but rather employs a machine learning approach. However, it presents two limitations. First, as often happens in machine learning contexts, the choice of the regressors to include in the model is based on data and increasing prediction accuracy is prioritised, even at the expense of increasing the estimation bias (Hastie et al., 2009). Furthermore, the procedure does not easily handle a high number of fixed effects. Hence, we consider IPWRA our preferred estimator.

Under the assumption of selection on observables and in the absence of additional endogeneity issues, both approaches would identify causal effects. Admittedly, we cannot be certain that these assumptions hold in the present setting. First, we cannot rule out the presence of unobserved confounding factors. For example, firms characterised by stronger managerial capabilities, greater technological readiness, or a more worker-oriented organisational culture may be simultaneously more likely to adopt AI technologies and to invest in worker welfare measures. Although our specifications include a rich set of controls related to innovation activities, digital investments, financial conditions, and other firm characteristics, these variables may not fully capture all relevant dimensions of organisational quality and managerial sophistication. If such factors remain unobserved, the estimated associations may be upward biased. Second, reverse causality cannot be excluded. AI adoption is measured over the period 2020-2022, while worker welfare plans refer to 2021-2022, implying substantial temporal overlap. Consequently, the observed relation-

ship may not only reflect firms introducing welfare measures after adopting AI, but also the possibility that welfare-oriented firms are better positioned to adopt AI technologies. For instance, firms investing in worker welfare may benefit from a more cooperative and adaptable workforce, facilitating the introduction of new technologies, or they may adopt productivity-enhancing technologies in order to sustain the costs associated with welfare provision. For these reasons, we interpret the estimates as associations and conditional associations at x_i^C, rather than as causal effects.

4. Results

4.1. Main Results

Table 2 reports the IPWRA estimates for the entire sample. Overall, our estimates suggest the existence of complementarities between firms' AI investment and the adoption of worker welfare measures. In particular, AI investment is strongly associated with the efforts of firms to introduce measures improving worker WB and H&S. On average, adopting firms have approximately 15 and 14 percentage points higher probability of introducing WB and H&S plans, respectively, compared to non-adopters.

The positive association between AI adoption and worker welfare measures is consistent with several alternative mechanisms. One possibility is that firms adopting AI introduce worker welfare measures to reassure employees and facilitate the acceptance of technological change, or to signal social responsibility towards workers and external stakeholders. Another possibility is re-

verse causality: firms that have already invested in worker welfare may benefit from a more cooperative and adaptable workforce, making the adoption of AI technologies easier and more effective. Relatedly, welfare-oriented firms may have stronger incentives to adopt productivity-enhancing technologies in order to sustain the costs associated with welfare provision. Although our specifications control for several firm characteristics, innovation activities, digital investments, and financial conditions that may partly capture this latter channel, we cannot rule it out completely. A third possibility is that both AI adoption and worker welfare investments are jointly driven by underlying managerial capabilities, technological readiness, or organisational sophistication. While some of these dimensions are partially captured by our controls, they may still represent a source of unobserved heterogeneity.

Table 2 **Associations between WB and H&S plans and AI Adoption**

	(1) WB	(2) H&S
AI adoption	0.1518***	0.1400***
	(0.0133)	(0.0159)
Observations	164161	166372

Notes: The table presents the IPWRA estimates of the association between AI adoption and WB and H&S plans. The row 'Observations' indicates the number of observations in the sample used in the estimates. Sample size varies as the method requires that each observation has propensity score >0 and <1. All the observations not satisfying this condition are removed.
*** p < 0.01, ** p < 0.05, * p < 0.1.

The available data do not allow us to distinguish among these mechanisms and the results should therefore be interpreted as documenting associations rather than identifying specific causal channels. Nevertheless, the heteroge-

neity analysis presented below provides useful insights into the conditions under which AI adoption and worker welfare measures are more strongly associated. As such, it offers some suggestive evidence that can help assess the plausibility of alternative interpretations. Table 3 and Table 4 report estimates of the conditional associations between AI adoption and WB and H&S plans, respectively.

All conditional associations for WB plans are positive: the data reject the possibility of negative relationship between AI adoption and WB measures for any group of firms. In addition, none of the differences are statistically significant: the significant association between the adoption of AI and of WB plans is uniform and it does not change with the conditioning variables under consideration. One possible explanation is that WB plans are generally discretionary, and yet are associated with relatively low implementation costs, high visibility, and strong scalability across different types of firms. This increases their attractiveness as instruments for firms seeking to implement or signal worker-oriented practices across heterogeneous organisational contexts.

Table 3 **Conditional associations between WB plans and AI Adoption**

	(1) *Territorial commitment*	(2) *Medium/ large*	(3) *Technical competencies constraint*	(4) *Cross competencies constraint*	(5) *Industry*
AI adoption × Conditioning var., Low	0.1501*** (0.0165)	0.158*** (0.0156)	0.1471*** (0.015)	0.1406*** (0.0136)	0.1617*** (0.0156)
AI adoption × Conditioning var., High	0.1402*** (0.0197)	0.1487*** (0.0183)	0.1493*** (0.0176)	0.1775*** (0.0215)	0.1526*** (0.0216)
Difference, High — Low	-0.0099 (0.0259)	-0.0093 (0.0240)	0.0022 (0.0232)	0.037 (0.0255)	-0.0091 (0.0267)
Observations, Low	128750	135799	92504	126094	104768
Observations, High	18889	19575	52773	26217	47566

Notes: The table presents the IPWRA estimates of the conditional associations between AI adoption and WB plans for low and high values of the conditioning variables, reported in each column header. Row 'Difference, High – Low' displays the differences between 'AI adoption x Conditioning var., High' and 'AI adoption x Conditioning var., Low'. The rows 'Observations, Low' and 'Observations, High' indicate the size of the sample used for estimating the conditional association at 0 and 1, respectively. The method requires that each observation has propensity score > 0 and < 1. Therefore, the observations not satisfying the condition are removed and the size of each subsample may not be equal to that of the whole sample. *** p < 0.01, ** p < 0.05, * p < 0.1.

In contrast, the results for H&S point to the presence of greater firm-level heterogeneity. This signals that the two types of welfare measures are different and do not respond to the same incentives and constraints. To anticipate and summarise the detailed discussion that follows, we argue that these differences reflect the twofold nature of H&S plans. On the one hand, they aim to improve worker welfare and help firms to ensure a positive co-existence between AI and workers; on the other hand, the implementation of H&S plans requires some adjustment to the internal organisation of firms and it is influenced by legal and risk-management considerations. These factors make the

adoption of H&S plans depend on a number of firm-specific characteristics.

First, we observe a negative moderating role of territorial commitment on the relationship between AI adoption and H&S plans. This negative moderating role may reflect a reallocation of financial and organisational resources. A strong commitment to projects external to the firm can crowd out resources available for discretionary H&S investments, which may be postponed or scaled back when firms face competing demands, such as simultaneous investment in AI and external initiatives. By contrast, WB initiatives are more directly linked to employee morale and adaptation during technological change and may also complement firms' efforts to signal social responsibility, making them less susceptible to such crowding-out effects.

Second, we find a weaker relationship for medium and large firms. One possible reason why larger firms seem not to introduce H&S plans after adopting AI solutions is the wider unconditional diffusion of H&S plans among these firms (64%) with respect to smaller firms (38%). The smaller estimated relationship between AI and H&S plans may thus reflect less room for additional adoption, not weaker interest. Indeed, differently from WB initiatives, the introduction of H&S plans requires that firms formalise their organisational practices by defining roles, responsibilities, and work routines. While larger firms have more structured management systems that allow them to deal with such organisational complexity, smaller firms are more likely to expand their managerial capabilities when AI adoption forces them to review their organisational practices (Hervas-Oliver et al., 2016; Edquist et al., 2001; Reichstein and Salter, 2006). In other words, smaller firms are more likely to introduce H&S plans when they also adopt AI solutions because the investment in both justifies the costly organisational changes that both of them require.

Third, we find that both cross- and technical competencies constraints

have a negative moderating role on the relationship between AI adoption and investment in H&S plans. AI adoption increases firms' demand for workers with such skills. Reliance on external hiring signals that incumbent workers are not yet equipped to manage the technological transition, and this may make the assimilation of AI incomplete. In this 'in-progress' phase, firms are less likely to introduce formal H&S procedures, which can be implemented at more mature stages of organisational adjustment so as to cover the entire workforce affected by the innovative AI solutions. Hence, an incomplete implementation may slow complementary H&S formalisation. However, it is important to point out that firms categorised as having low competencies constraints include both firms that can easily recruit workers with the required competencies and firms that are not actively hiring. Therefore, the weaker association observed among firms reporting hiring constraints should not be interpreted solely as the effect of skill shortages, but more broadly as reflecting the conditions under which firms manage the workforce implications of AI adoption.

Finally, estimates suggest that AI-adopting firms in the industrial sector are less likely to adopt worker H&S plans compared to those in other sectors. Following our previous discussion, we believe that this may reflect the differences in how AI affects the production processes across different sectors. In industrial settings, AI adoption often entails the integration of automated systems and robotics, which can reduce workers' direct exposure to certain physical risks. In this sense, AI adoption may itself enhance H&S conditions. As a result, the marginal need for additional H&S investments may be lower than in service sectors, where AI primarily reshapes work organisation and worker interactions without necessarily reducing exposure to risks.

Table 4 **Conditional associations between H&S plans and AI Adoption**

	(1) *Territorial commitment*	(2) *Medium/ large*	(3) *Technical competencies constraint*	(4) *Cross competencies constraint*	(5) *Industry*
AI adoption × Conditioning var., Low	0.1514*** (0.0188)	0.1605*** (0.0172)	0.1454*** (0.0176)	0.1361*** (0.0165)	0.1566*** (0.0185)
AI adoption × Conditioning var., High	-0.0016 (0.0197)	0.0229 (0.0224)	0.052*** (0.0164)	0.0398 (0.0265)	0.0862*** (0.0212)
Difference, High — Low	-0.153*** (0.0237)	-0.1377*** (0.0213)	-0.0934*** (0.0240)	-0.0962*** (0.0311)	-0.0705** (0.0282)
Observations, Low	135147	139323	97322	130968	109187
Observations, High	15320	18977	49604	22834	47542

Notes: The table presents the IPWRA estimates of the conditional associations between AI adoption and H&S plans for low and high values of the conditioning variables, reported in each column header. Row 'Difference, High – Low' displays the differences between 'AI adoption x Conditioning var., High' and 'AI adoption x Conditioning var., Low'. The rows 'Observations, Low' and 'Observations, High' indicate the size of the sample use for estimating the conditional association at 0 and 1, respectively. The method requires that each observation has propensity score > 0 and <1. Therefore, the observations not satisfying the condition are removed and the size of each subsample may not be equal to that of the whole sample. *** p < 0.01, ** p < 0.05, * p < 0.1.

The Census survey allows us to build a different, broader measure of AI adoption that considers all those technologies that fall in one of the areas of application of AI. This is, for instance, the measure of AI used by Montresor and Vezzani (2023). The estimates based on investments in AI application areas as the main explanatory variable are reported in Table 5. They show that the estimated association between investment in AI application areas and worker welfare is positive and significant, and consistent with the results discussed above. The coefficients are almost halved. This attenuation may reflect measurement issues, but it also suggests that AI-specific adoption and broad-

er digital upgrading are associated with worker welfare practices in different ways.

Table 5 **Associations between WB and H&S plans and AI application areas**

	(1) WB	(2) H&S
AI application areas	0.0821***	0.0670***
	(0.0073)	(0.0093)
Observations	172166	172166

Notes: The table presents the IPWRA estimates of the association between investment in AI application areas and WB and H&S plans. The row 'Observations' indicates the number of observations in the sample used in the estimates. Sample size varies as the method requires that each observation has propensity score > 0 and < 1. All the observations not satisfying the condition are, therefore, removed. Robust standard errors are shown in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

Next, we study the conditional associations between investment in AI application areas and WB and H&S plans. The estimates are reported in Table 6 and Table 7, respectively.

In contrast to the results based on the specific measure of AI adoption, territorial commitment and belonging to the industrial sector significantly moderate the relationship between WB plans and investment in AI application areas. The positive moderating role of territorial commitment may be explained by the fact that these AI-related technologies include non-negligible automation components and adopting firms with a strong territorial commitment are more likely to provide visible forms of compensation to their workers. The negative moderating role associated with belonging to the industrial sector may reflect the nature of such investments, which often include automation technologies that substitute for labour. The resulting reduction in labour demand can weaken firms' incentives to invest in WB plans. However,

it is important to note that this broader proxy for AI investment may capture a range of technologies serving different purposes, such as efficiency gains, monitoring, connectivity, or automation, which further complicates the interpretation of these results.

Looking at the conditional associations between H&S plans and investments in AI application areas, it is possible to note that territorial commitment, firm size, and hiring constraints still play a moderating role and negatively condition the relationship between broadly defined investment in AI and the implementation of H&S plans. This is in line with the previous results and our expectations. On the other hand, belonging to the industrial sector no longer plays a significant moderating role. This result is somewhat surprising, given that industrial sector affiliation significantly moderates the relationship between investments in AI application areas and WB plans. However, as noted earlier, the broader proxy for AI investment may capture a wide range of technologies with different purposes, which may account for this discrepancy.

Although the results based on the two alternative measures of AI adoption are broadly consistent, the differences discussed above highlight the importance of clearly defining what constitutes AI, both at the survey design stage and in subsequent empirical analysis. The first and more precise measure of AI adoption does not capture all digital technologies that may incorporate AI-related features, whereas the second broader proxy does. When AI proxies reflect the joint adoption of multiple technologies, it becomes difficult to disentangle the relative contribution of AI from that of other technologies.

Table 6 **Conditional associations between WB plans and AI application areas**

	(1)	(2)	(3)	(4)	(5)
	Territorial commitment	*Medium/ large*	*Technical competencies constraint*	*Cross-competencies constraint*	*Industry*
AI application areas × Conditioning var., Low	0.0773*** (0.0086)	0.0787*** (0.0084)	0.0817*** (0.0092)	0.0777*** (0.0085)	0.0904*** (0.0092)
AI application areas × Conditioning var., High	0.1049*** (0.0106)	0.0946*** (0.0093)	0.0774*** (0.0091)	0.0964*** (0.0107)	0.0602*** (0.009)
Difference, High — Low	0.0277** (0.0136)	0.0159 (0.0125)	-0.0043 (0.0128)	0.0187 (0.0137)	-0.0303** (0.0128)
Observations, Low	145577	145905	111929	139251	117211
Observations, High	26286	25690	59921	32485	54766

Notes: The table presents the IPWRA estimates of the conditional associations between investment in AI application areas and WB plans for low and high values of the conditioning variables, reported in each column header. Row 'Difference, High – Low' displays the differences between 'AI adoption x Conditioning var., High' and 'AI adoption x Conditioning var., Low'. The rows 'Observations, Low' and 'Observations, High' indicate the size of the sample use for estimating the conditional association at 0 and 1, respectively. The method requires that each observation has propensity score > 0 and < 1. Therefore, the observations not satisfying the condition are removed and the size of each subsample may not be equal to that of the whole sample.
*** p < 0.01, ** p < 0.05, * p < 0.1.

Table 7 **Conditional associations between H&S plans and AI application areas**

	(1)	(2)	(3)	(4)	(5)
	Territorial commitment	*Medium/ large*	*Technical competencies constraint*	*Cross-competencies constraint*	*Industry*
AI application areas × Conditioning var., Low	0.0768*** (0.0108)	0.0711*** (0.0108)	0.0898*** (0.0127)	0.0752*** (0.0111)	0.0699*** (0.0114)
AI application areas × Conditioning var., High	0.0083 (0.0096)	0.042*** (0.0116)	0.0265** (0.0104)	0.0178* (0.0107)	0.0505*** (0.0133)
Difference, High — Low	-0.0685*** (0.0106)	-0.0291* (0.0158)	-0.0633*** (0.0163)	-0.0574*** (0.0154)	0.0194 (0.0175)
Observations, Low	145727	146225	111597	139251	117211
Observations, High	25820	25810	59934	32483	54828

Notes: The table presents the IPWRA estimates of the conditional associations between investment in AI application areas and H&S plans for low and high values of the conditioning variables, reported in each column header. Row 'Difference, High – Low' displays the differences between 'AI adoption x Conditioning var., High' and 'AI adoption x Conditioning var., Low'. The rows 'Observations, Low' and 'Observations, High' indicate the size of the sample use for estimating the conditional association at 0 and 1, respectively. The method requires that each observation has propensity score > 0 and < 1. Therefore, the observations not satisfying the condition are removed and the size of each subsample may not be equal to that of the whole sample.
*** p < 0.01, ** p < 0.05, * p < 0.1.

4.2. Robustness checks

Although IPWRA is a doubly robust estimation method, it is highly parametric and this is particularly restrictive for the conditional estimations. For this reason, we provide alternative estimates using a non-parametric proce-

dure suggested by Fan et al. (2022), in particular their full sample estimator.[11]

The key equation of the procedure, identical to the one used in the augmented regression adjustment, is the following one:

$$ATE = E\left[\frac{D_i(Y_i - E[Y_i|D_i = 1, \mathbf{X_i}])}{\pi(x_i)} + E[Y_i|D_i = 1, \mathbf{X_i}] - \right.$$

$$\left. \frac{(1 - D_i)(Y_i - E[Y_i|D_i = 0, \mathbf{X_i}])}{1 - \pi(x_i)} + E[Y_i|D_i = 0, \mathbf{X_i}]\right]$$

where $\pi(x_i)$ is the propensity score. The procedure consists of estimating the outcome equation and the propensity score by machine learning methods, and then calculating the $CATE(x_i^C)$ using a local linear polynomial estimator (Hastie et al., 2009). In our case, the latter is equivalent to calculating the sample average of equation above when the conditional variable is equal to either 0 or 1. Fan et al. (2022) show that, under some regularity conditions, it is possible to consistently estimate both ATE and $CATE(x_i^C)$ with a non-parametric procedure based on random forest estimations of both the propensity score and the outcome equation.

To implement this method, we use random forests to estimate both the propensity score and the outcome equation. A random forest combines the predictions of many regression trees, each estimated on a bootstrapped sample of the data. By averaging across trees, it reduces the high variance typically associated with single-tree estimators (Hastie et al., 2009).[12]

---

11 For the implementation, we use the hdcate package (Fan and Hong, 2022).

12 Each regression tree divides the covariate space into a set of regions defined by combinations of regressors. Within each region, the conditional expectation of the outcome is estimated by the sample average of the outcome variable. The partition is constructed recursively by evaluating candidate splits across regressors and selecting those that minimise the sum of squared errors (Hastie et al., 2009). To reduce the computational burden of tree construction, we apply target encoding to province and industry fixed effects. In practice, this means replacing each categorical variable with a numerical variable equal to the average value of the outcome within each category (Clerici and Nobani, 2026; Borrohou et al., 2025).

As testing the statistical significance of conditional associations in this setting is computationally demanding, we focus on whether the random forest point estimates are consistent with our main results and draw conclusions regarding their robustness. Tables A1-A6 in the Appendix report estimates based on the Fan et al. (2022) method. The point estimates are broadly consistent with those obtained using IPWRA, suggesting that our results are robust to this alternative approach.

5. Conclusions

This paper examines whether the adoption of AI is associated with firms' investments in worker well-being and occupational health and safety plans. We draw on firm-level data from the 2022 wave of the Italian Permanent Census of Enterprises.

Our results indicate a positive association between AI adoption and both dimensions of worker welfare. Relative to non-adopters, firms adopting AI are approximately 15 and 14 percentage points more likely to introduce WB and H&S plans, respectively. This suggests that AI-related technological upgrading is accompanied by worker-oriented measures. This may reflect alternative mechanisms at work. Firms may try to reassure incumbent employees or to pre-empt resistance, bridging labour-relations management with welfare-enhancing intents, thereby linking labour-relations management with welfare-enhancing initiatives. Moreover, as AI adoption may attract scrutiny from external stakeholders, the introduction of visible worker-oriented plans may also serve as a signalling device. Furthermore, firms that have already invested in worker welfare may face a more cooperative and adaptable work-

force, and this facilitates the introduction of productivity-enhancing technologies that help the companies to sustain the costs of the welfare provisions. Finally, a positive association may derive from the fact that welfare-oriented firms and firms with higher managerial capabilities and technological readiness have stronger incentives and lower constraints to adopt both productivity-enhancing technologies and welfare-enhancing initiatives.

The analysis shows that the two welfare domains should not be treated as equivalent. The association between AI adoption and WB plans is relatively stable across firms, whereas the relationship with H&S plans is markedly heterogeneous. In particular, it is weaker among firms with stronger territorial commitment, larger firms, firms facing hiring constraints in technical and cross-functional competencies, and firms operating in the industrial sector. This pattern suggests that H&S plans are shaped not only by considerations related to worker support and signalling concerns, but also by the organisational complexity involved in integrating both AI and H&S practices into production processes.

We acknowledge that our findings should be interpreted with caution. Although our empirical strategy addresses potential selection bias through a doubly robust estimation method (IPWRA), and the results are broadly robust to both an alternative measure of AI investment based on investments in at least one AI application area, and an alternative estimation approach based on machine learning methods, the available data do not allow us to rule out omitted variables, reverse causality, or other sources of endogeneity. Accordingly, our findings should be interpreted as associations and not in terms of treatment effects.

It is also worth noting that the analysis refers to Italian firms over the period 2020–2022, thus capturing a phase of AI diffusion that predates the

widespread adoption of generative AI. This does not necessarily constitute a limitation, insofar as the type of AI considered is clearly defined. Relatedly, while the alternative proxy for AI investment, commonly used in the literature, yields broadly consistent results, some differences emerge. This highlights the importance of clearly defining AI measures in empirical analyses.

Overall, this study suggests that firms often accompany AI adoption with measures aimed at supporting workers, although the form of such support varies across welfare domains. In some cases, AI adoption and H&S plans may even coincide, particularly when AI applications contribute to improving workplace safety. However, this is not always the case, as firms introduce AI across a wide range of activities, including logistics, customer portfolio management, preventive maintenance, and the development of intelligent systems.

Future research could more directly test the mechanisms underlying these patterns, for instance by distinguishing between compliance-driven and discretionary safety investments, or by exploiting longitudinal data to assess whether worker welfare measures precede, accompany, or follow the adoption of AI within firms.

Finally, an important caveat is that our analysis refers to the period 2020-2022 and therefore predates the widespread diffusion of Generative AI. While the forms of AI considered in this paper are primarily related to process automation and intelligent systems, Generative AI is likely to affect a broader range of cognitive and knowledge-intensive tasks. As a result, it may alter firms' incentives to invest in worker WB measures more substantially than in the past. In particular, because Generative AI may increase workers' concerns about job displacement, changes in task content, and the future value of their skills, firms may face stronger incentives to accompany its adoption

with measures aimed at supporting employee WB, engagement, and acceptance of technological change. By contrast, there is less reason to expect a comparable effect on occupational H&S plans, as Generative AI does not directly modify workers' exposure to physical workplace risks in the way that many process-oriented AI applications and automation technologies can. Future research should therefore examine whether the diffusion of Generative AI strengthens the relationship between AI adoption and worker WB initiatives while leaving the relationship with H&S measures largely unchanged.

References

[1] Abrevaya, J., Hsu, Y.-C., and Lieli, R. P. (2015). Estimating Conditional Average Treatment Effects. *Journal of Business & Economic Statistics*, 33(4):485–505.

[2] Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340.

[3] Acemoglu, D. and Restrepo, P. (2020). Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy*, 128(6):2188–2244.

[4] Achievers (2022). Empowering employee wellbeing in the new world of work. Technical report, Achievers Workforce Institute.

[5] Aldasoro, I., Gambacorta, L., Pal, R., Revoltella, D., Weiss, C., and Wolski, M. (2026). Ai adoption, productivity and employment: Evidence from european firms. BIS Working Papers 1325, Bank for International Settlements.

[6] Bahoo, S., Cucculelli, M., and Qamar, D. (2023). Artificial intelligence and corporate innovation: A review and research agenda. *Technological Forecasting and Social Change*, 188:122264.

[7] Bonfiglioli, A., Crinò, R., Gancia, G., and Papadakis, I. (2024). Artificial intelligence and jobs: evidence from US commuting zones. *Economic Policy*, 40(121):145–194.

[8] Bonsòn, E., Lavorato, D., Lamboglia, R., and Mancini, D. (2021). Artificial intelligence activities and ethical approaches in leading listed companies in the European Union. *International Journal of Accounting Information Systems*, 43:100535.

[9] Borrohou, S., Fissoune, R., and Badir, H. (2025). The role of data transformation in modern analytics: A comprehensive survey. *Journal of Computer Languages*, 84:101329.

[10] Brynjolfsson, E., Li, D., and Raymond, L. R. (2025). Generative AI at Work. *The Quarterly Journal of Economics*, 140(2):889–942.

[11] Caselli, M., Fracasso, A., Marcolin, A., and Scicchitano, S. (2025). Technological innovations and workers' job insecurity: the moderating role of human resource strategies. *Economia e Politica Industriale: Journal of Industrial and Business Economics*, 52(1):153–176.

[12] Cassar, L. and Meier, S. (2018). Nonmonetary Incentives and the Implications of Work as a Source of Meaning. *Journal of Economic Perspectives*, 32(3):215–238.

[13] Cerulli, G. (2015). *Econometric Evaluation of Socio-Economic Programs*, volume 49 of *Advanced Studies in Theoretical and Applied Econometrics*. Springer Berlin Heidelberg, Berlin, Heidelberg.

[14] Chen, P., Chu, Z., and Zhao, M. (2024). The Road to corporate sustainability: The importance of artificial intelligence. *Technology in Society*, 76:102440.

[15] Clerici, F. and Nobani, N. (2026). Categorical variable encoding methods for tabular data: A benchmarking study. *International Journal of Data Science and Analytics*, 22(1):55.

[16] Cooper, R. B. and Zmud, R. W. (1990). Information Technology Implementation Research: A Technological Diffusion Approach. *Management Science*, 36(2):123–139.

[17] Edquist, C., Hommen, L., and McKelvey, M. (2001). *Innovation and Employment: Process versus Product Innovation*. Edward Elgar Publishing.

[18] Elfenbein, D. W. and McManus, B. (2010). A Greater Price for a Greater Good? Evidence That Consumers Pay More for Charity-Linked Products. *American Economic Journal: Economic Policy*, 2(2):28–60.

[19] Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2023). GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. *arXiv preprint arXiv:2303.10130*.

[20] European Agency for Safety and Health at Work (2023). Smart digital monitoring systems for occupational safety and health: Implementation at the workplace. Technical report, European Agency for Safety and Health at Work.

[21] Fan, Q. and Hong, H. (2022). Hdcate: Estimation of Conditional Average Treatment Effects with High-Dimensional Data.

[22] Fan, Q., Hsu, Y.-C., Lieli, R. P., and Zhang, Y. (2022). Estimation of Conditional Average Treatment Effects With High-Dimensional Data. *Journal of Business & Economic Statistics*, 40(1):313–327.

[23] Felten, E., Raj, M., and 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.

[24] Giuntella, O., Konig, J., and Stella, L. (2025). Artificial intelligence and the wellbeing of workers. *Scientific Reports*,15.

[25] Golgeci, I., Ritala, P., Arslan, A., McKenna, B., and Ali, I. (2025). Confronting and alleviating ai resistance in the workplace: An integrative review and a process framework. *Human Resource Management Review*, 35(2):101075.

[26] Grant, A. M., Christianson, M. K., and Price, R. H. (2007). Happiness, Health, or Relationships? Managerial Practices and Employee Well-Being Tradeoffs. *Academy of Management Perspectives*, 21(3):51–63.

[27] Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. New York,

NY: Springer.

[28] Hervas-Oliver, J.-L., Boronat-Moll, C., and Sempere-Ripoll, F. (2016). On Process Innovation Capabilities in SMEs: A Taxonomy of Process-Oriented Innovative SMEs. *Journal of Small Business Management*, 54(S1):113–134.

[29] Huang, K., Wang, K., Lee, P. K. C., and Yeung, A. C. L. (2023). The impact of industry 4.0 on supply chain capability and supply chain resilience: A dynamic resource-based view. *International Journal of Production Economics*, 262:108913.

[30] ISO (2018). ISO 45001:2018. https://www.iso.org/standard/63787.html.

[31] ISTAT (2022). Censimento 2022 -questionario con guida alla compilazione -imprese con 3-9 addetti.

[32] Jann, B. (2019). Influence functions for linear regression (with an application to regression adjustment). University of Bern Social Sciences Working Papers 32, University of Bern, Department of Social Sciences.

[33] Jiang, W., Park, J., Xiao, R. J., and Zhang, S. (2025). AI and the Extended Workday: Productivity, Contracting Efficiency, and Distribution of Rents. NBER Working Papers 33536, National Bureau of Economic Research, Inc.

[34] Judge, T. A., Thoresen, C. J., Bono, J. E., and Patton, G. K. (2001). The job satisfaction– job performance relationship: A qualitative and quantitative review. *Psychological Bulletin*, 127(3):376–407.

[35] Kim, B.-J. and Lee, J. (2025). AI adoption, employee depression and knowledge: How corporate social responsibility buffers psychological impact. *Journal of Innovation & Knowledge*, 10(6):100815.

[36] Kitzmueller, M. and Shimshack, J. (2012). Economic Perspectives on Corporate Social Responsibility. *Journal of Economic Literature*, 50(1):51–84.

[37] Kumar, D. and Ratten, V. (2024). Artificial intelligence and family businesses: A systematic literature review. *Journal of Family Business Mana-*

gement, 15(2):373–392.

[38] Ladeira, W. J., Santini, F. D. O., Rasul, T., Cheah, I., Elhajjar, S., Yasin, N., and Akhtar, S. (2024). Big data analytics and the use of artificial intelligence in the services industry: A meta-analysis. *The Service Industries Journal*, 44(15-16):1117–1144.

[39] Manning, S. J. and Aguirre, T. (2026). How Adaptable Are American Workers to AIInduced Job Displacement? Working Paper 34705, National Bureau of Economic Research.

[40] Montresor, S. and Vezzani, A. (2023). Digital technologies and eco-innovation. Evidence of the twin transition from Italian firms. *Industry and Innovation*, 30(7):766–800.

[41] Muehlemann, S. (2025). Artificial intelligence adoption and workplace training. *Journal of Economic Behavior & Organization*, 238:107206.

[42] Nazareno, L. and Schiff, D. S. (2021). The impact of automation and artificial intelligence on worker well-being. *Technology in Society*, 67:101679.

[43] Nielsen, K., Nielsen, M. B., Ogbonnaya, C., K¨ans¨al¨a, M., Saari, E., and Isaksson, K. (2017). Workplace resources to improve both employee well-being and performance: A systematic review and meta-analysis. *Work & Stress*, 31(2):101–120.

[44] Nucci, F., Puccioni, C., and Ricchi, O. (2023). Digital technologies and productivity: A firm-level investigation. *Economic Modelling*, 128:106524.

[45] Patel, P. C., Devaraj, S., Hicks, M. J., and Wornell, E. J. (2018). County-level job automation risk and health: Evidence from the United States. *Social Science & Medicine*, 202:54–60.

[46] Reichstein, T. and Salter, A. (2006). Investigating the sources of process innovation among UK manufacturing firms. *Industrial and Corporate Change*, 15(4):653–682.

[47] Schulz, C., Bendig, D., Braunche, A., and Kindermann, B. (2026). Curse or blessing: Investigating the influence of firms' artificial intelligence

adoption on employee job satisfaction. *Journal of Management Studies*, 63(2):561–595.

[48] Shah, I. A. and Mishra, S. (2024). Artificial Intelligence in Advancing Occupational Health and Safety: An Encapsulation of Developments. *Journal of Occupational Health*, 66(1):uiad017.

[49] Tonin, M. and Vlassopoulos, M. (2015). Corporate Philanthropy and Productivity: Evidence from an Online Real Effort Experiment. *Management Science*, 61(8):1795–1811.

[50] Uhrenholdt Madsen, C., Kirkegaard, M. L., Dyreborg, J., and Hasle, P. (2020). Making occupational health and safety management systems 'work': A realist review of the OHSAS 18001 standard. *Safety Science*, 129:104843.

[51] Valtonen, A., Saunila, M., Ukko, J., Treves, L., and Ritala, P. (2025). AI and employee wellbeing in the workplace: An empirical study. *Journal of Business Research*, 199:115584.

[52] Wooldridge, J. M. (2007). Inverse probability weighted estimation for general missing data problems. *Journal of Econometrics*, 141(2):1281–1301.

[53] Zhu, K., Kraemer, K. L., and Xu, S. (2006). The Process of Innovation Assimilation by Firms in Different Countries: A Technology Diffusion Perspective on E-Business. *Management Science*, 52(10):1557–1576.

[54] Zirar, A., Ali, S. I., and Islam, N. (2023). Worker and workplace Artificial Intelligence (AI) coexistence: Emerging themes and research agenda. *Technovation*, 124:102747.

Appendix

Table A1 **Estimates for WB and H&S plans, Fan et al. (2022) method**

	(1) WB	(2) H&S
AI adoption	0.1656	0.1376
Observations	164161	166372

Notes: The table presents estimates of the association between AI adoption and WB and H&S plans, obtained using the Fan et al. (2022) method. The row 'Observations' indicates the number of observations in the sample used in the estimates. The sample size varies to match the sample used in IPWRA.

Table A2 **Conditional estimates for WB plans, Fan et al. (2022) method**

	(1) *Territorial commitment*	(2) *Medium/ large*	(3) *Technical competencies constraint*	(4) *Cross-competencies constraint*	(5) *Industry*
AI adoption × Conditioning var., Low	0.1608	0.1641	0.1694	0.1647	0.1736
AI adoption × Conditioning var., High	0.1921	0.1739	0.1587	0.1692	0.1491
Difference, High — Low	0.0313	0.0098	-0.0106	0.0044	-0.0245
Observations	164161	164161	164161	164161	164161

Notes: The table presents estimates of the conditional associations between AI adoption and WB plans, obtained using the Fan et al. (2022) method, for low and high values of the conditioning variables, reported in each column header. Row 'Difference, High – Low' displays the differences between 'AI adoption x Conditioning var., High' and 'AI adoption x Conditioning var., Low'. The row 'Observations' indicates the number of observations in the sample used in the estimates.
*** p < 0.01, ** p < 0.05, * p < 0.1.

Table A3 **Conditional estimates for H&S plans, Fan et al. (2022) method**

	(1) *Territorial commitment*	(2) *Medium/ large*	(3) *Technical competencies constraint*	(4) *Cross-competencies constraint*	(5) *Industry*
AI adoption × Conditioning var., Low	0.1548	0.1500	0.1717	0.1454	0.1514
AI adoption × Conditioning var., High	0.0420	0.0684	0.0746	0.1047	0.1085
Difference, High — Low	-0.1128	-0.0816	-0.0971	-0.0408	-0.0429
Observations	166372	166372	166372	166372	166372

Notes: The table presents estimates of the conditional associations between AI adoption and H&S plans, obtained using the Fan et al. (2022) method, for low and high values of the conditioning variables, reported in each column header. Row 'Difference, High – Low' displays the differences between 'AI adoption x Conditioning var., High' and 'AI adoption x Conditioning var., Low'. The row 'Observations' indicates the number of observations in the sample used in the estimates.
*** p < 0.01, ** p < 0.05, * p < 0.1.

Table A4 **Estimates for WB and H&S plans, AI application areas, Fan et al. (2022) method**

	(1) WB	(2) H&S
AI application areas	0.0846	0.0715
Observations	172166	172166

Notes: The table presents estimates of the association between investment in AI application areas and WB and H&S plans, obtained using the Fan et al. (2022) method. The row 'Observations' indicates the number of observations in the sample used in the estimates. The sample size varies to match the sample used in IPWRA.

Table A5 **Conditional estimates for WB plans, AI application areas, Fan et al. (2022) method**

	(1)	(2)	(3)	(4)	(5)
	Territorial commitment	*Medium/ large*	*Technical competencies constraint*	*Cross-competencies constraint*	*Industry*
AI application areas × Conditioning var., Low	0.0795	0.0817	0.0845	0.0832	0.0903
AI application areas × Conditioning var., High	0.1130	0.1015	0.0848	0.0906	0.0725
Difference, High — Low	0.0335	0.0199	0.0003	0.0074	-0.0179
Observations	172166	172166	172166	172166	172166

Notes: The table presents estimates of the conditional associations between investment in AI application areas and WB plans, obtained using the Fan et al. (2022) method, for low and high values of the conditioning variables, reported in each column header. Row 'Difference, High – Low' displays the differences between 'AI application areas x Conditioning var., High' and 'AI application areas x Conditioning var., Low'. The row 'Observations' indicates the number of observations in the sample used in the estimates. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A6 **Conditional estimates for H&S plans, AI application areas, Fan et al. (2022) method**

	(1) *Territorial commitment*	(2) *Medium/ large*	(3) *Technical competencies constraint*	(4) *Cross-competencies constraint*	(5) *Industry*
AI application areas × Conditioning var., Low	0.0835	0.0756	0.0941	0.0826	0.0752
AI application areas × Conditioning var., High	0.0055	0.0489	0.0295	0.0247	0.0636
Difference, High — Low	-0.0780	-0.0267	-0.0646	-0.0580	-0.0116
Observations	172166	172166	172166	172166	172166

Notes: The table presents estimates of the conditional associations between investment in AI application areas and H&S plans, obtained using the Fan et al. (2022) method, for low and high values of the conditioning variables, reported in each column header. Row 'Difference, High – Low' displays the differences between 'AI application areas x Conditioning var., High' and 'AI application areas x Conditioning var., Low'. The row 'Observations' indicates the number of observations in the sample used in the estimates. *** p < 0.01, ** p < 0.05, * p < 0.1.

PARTNER ISTITUZIONALI

CeSPEM
CENTRO STUDI DI POLITICA ECONOMICA E MONETARIA "MARIO ARCELLI"

UNIVERSITÀ CATTOLICA del Sacro Cuore

LUISS
CASMEF Centro Arcelli per gli Studi Monetari e Finanziari

aeb ASSONEBB

FONDAZIONE DI PIACENZA E VIGEVANO

CONFCOMMERCIO IMPRESE PER L'ITALIA

BUSINESS PARTNER

IMI CORPORATE & INVESTMENT BANKING
INTESA SANPAOLO

sara ti assicura

Deloitte.

SOSTENITORI

3D Works
Allianz Bank F.A.
Banca d'Italia
Banca Popolare del Cassinate
Banco Posta SGR
Confindustria Piacenza
CONSOB
EFPA - Italia
Fondazione Ave Verum
Fondazione Mario Arcelli
ISTAT
Istituto per il Credito Sportivo e Culturale
TIM
Unicredit
Ver Capital SGR

Per rinnovare o attivare un nuovo abbonamento
effettuare un **versamento** su:

c/c bancario n. 10187 Intesa Sanpaolo
Via Vittorio Veneto 108/b- 00187 ROMA
IBAN IT92 M030 6905 0361 0000 0010 187

intestato a: **Editrice Minerva Bancaria s.r.l.**

oppure inviare una **richiesta** a:
amministrazione@editriceminervabancaria.it

Condizioni di abbonamento ordinario per il 2027

	Rivista Bancaria Minerva Bancaria bimestrale	**Economia Italiana** quadrimestrale	**Rivista Bancaria Minerva Bancaria + Economia Italiana**
Canone Annuo Italia *(print)*	**€ 130,00**	**€ 100,00**	**€ 180,00**
Canone Annuo Estero *(print)*	**€ 185,00**	**€ 130,00**	**€ 260,00**
Abbonamento WEB	**€ 80,00**	**€ 70,00**	**€ 110,00**
Canone Annuo Italia *(print + web)*	**€ 170,00**	**€ 130,00**	**€ 260,00**
Canone Annuo Estero *(print + web)*	**€ 220,00**	**€ 160,00**	**€ 330,00**

L'abbonamento è per un anno solare e dà diritto a tutti i numeri usciti nell'anno.
L'Amministrazione non risponde degli eventuali disguidi postali.
I fascicoli non pervenuti dovranno essere richiesti alla pubblicazione del fascicolo successivo.
Decorso tale termine, i fascicoli disponibili saranno inviati contro rimessa del prezzo di copertina.

Prezzo del fascicolo in corso **€ 50,00 / € 10,00** digitale
Prezzo di un fascicolo arretrato (annata precedente) **€ 60,00 / € 10,00** digitale

Pubblicità

1 pagina **€ 1.500,00** - 1/2 pagina **€ 800,00**

Editrice Minerva Bancaria
COMITATO EDITORIALE STRATEGICO

PRESIDENTE
GIORGIO DI GIORGIO, Luiss Guido Carli

COMITATO
Claudio Chiacchierini, Università degli Studi di Milano Bicocca
Mario Comana, Luiss Guido Carli
Adriano De Maio, Università Link Campus
Raffaele Lener, Università degli Studi di Roma Tor Vergata
Marcello Martinez, Università della Campania
Giovanni Parrillo, Editrice Minerva Bancaria
Marco Tofanelli, Assoreti

ECONOMIA ITALIANA 2026/2

Intelligenza artificiale: da innovazione tecnologica a crescita economica inclusiva

Il volume, editor **Alessandra Bonfiglioli** (Università di Bergamo), **Gino Gancia** (Università di Milano Bicocca) e **Paola Profeta** (Università Bocconi), offre un'attualissima analisi degli effetti dell'Intelligenza Artificiale sul sistema economico.

Una nuova *general purpose technology* suscita previsioni estreme: un'accelerazione senza precedenti della crescita per alcuni, la sostituzione del lavoro umano per altri. Se l'IA sta già cambiando organizzazioni e compiti cognitivi, resta da capire, osservano gli editor, "*attraverso quali meccanismi questi cambiamenti possano tradursi in trasformazioni dell'intero sistema economico e sociale*".

Il filo conduttore dei cinque saggi è una domanda: **quali condizioni consentono all'IA di trasformarsi da innovazione tecnologica in crescita economica inclusiva?** Tre le questioni in gioco: gli effetti dell'IA sulla produttività, la redistribuzione di opportunità e rischi tra lavoratori, famiglie e imprese, il ruolo delle politiche pubbliche.

Flavio Calvino e Antonio Ughi leggono la diffusione dell'IA come un processo sistemico, i cui benefici dipendono da un ecosistema favorevole all'innovazione. **Sergio Galletta** ricostruisce le prime evidenze empiriche sull'IA generativa: la questione non è più il potenziale trasformativo dell'IA, ma le condizioni per una diffusione efficace e socialmente desiderabile. L'indagine di **Massimo Anelli, Italo Colantone, Aina Gallego e Piero Stanig** sui lavoratori italiani più esposti all'IA generativa rileva insieme ottimismo, cautela e una crescente domanda di politiche pubbliche. **Emilio Colombo e Francesco Trentini** propongono la famiglia, anziché il singolo lavoratore, come unità di analisi degli effetti distributivi: un meccanismo di assicurazione informale. Infine **Mauro Caselli, Enrico Cristofoletti e Andrea Fracasso** mostrano che le imprese italiane che adottano l'IA investono di più anche in benessere organizzativo, salute e sicurezza dei lavoratori.

Ne emerge un'immagine dell'IA diversa da quella del dibattito pubblico: "*non una tecnologia che produce automaticamente vincitori e vinti*", ma un insieme di innovazioni il cui impatto dipende dall'interazione tra progresso tecnico, organizzazione economica e istituzioni. Di qui le tre indicazioni di policy degli editor: **investimenti complementari, gestione della transizione occupazionale, governance dell'innovazione**. Come nelle rivoluzioni tecnologiche del passato, tutto dipenderà dalla capacità di imprese, lavoratori e istituzioni di riorganizzarsi.

Completano il numero l'*intervento* di **Carmine Di Noia**, con un'ampia ricerca (450 partecipanti, Authority incluse) sull'uso dell'IA nella finanza italiana, e la *rubrica* di **Massimiliano Parco** su sicurezza energetica, materie prime e nuove dipendenze.

ECONOMIA ITALIANA nasce nel 1979 per approfondire e allargare il dibattito sui nodi strutturali e i problemi dell'economia italiana, anche al fine di elaborare adeguate proposte strategiche e di *policy*. L'Editrice Minerva Bancaria è impegnata a portare avanti questa sfida e a fare di Economia Italiana il più vivace e aperto strumento di dialogo e riflessione tra accademici, *policy makers* ed esponenti di rilievo dei diversi settori produttivi del Paese.

EMB
EDITRICE MINERVA BANCARIA