

The Labor Market Impact of Generative AI: A Critical Survey of the Empirical Literature¹

Emanuela Carbonara

University of Bologna, Department of Sociology and Business Law

Enrico Santarelli

University of Bologna, Department of Economics

Abstract

This article surveys the emerging empirical literature on the impact of artificial intelligence (AI) – and particularly generative AI – on labor market outcomes. Despite widespread public debate, the evidence remains inconclusive and highly heterogeneous. Drawing on both individual studies and three recent meta-analyses, the paper identifies four broad channels of impact: employment, wages, productivity, and inequality. It highlights contrasts in methodology, geographical coverage, sectoral focus, and definitions of AI, emphasizing the need to distinguish between task-based automation and general-purpose technologies such as large language models. Across the literature, no consistent evidence has yet emerged of systemic labor market disruption attributable to generative AI. While displacement occurs in certain occupations – especially within digital freelance markets – other sectors exhibit productivity gains or wage premia. Recent meta-analytic studies confirm the absence of statistically significant average effects of AI adoption on employment or wages and detect no robust signs of publication bias. These findings reflect the still-emergent diffusion of AI and its uneven adoption across industries and countries. The article concludes by identifying priorities for future research.

Keywords: Generative Artificial Intelligence, Labor Market, Technological Change, Empirical Literature

Riassunto. *L'impatto dell'IA generativa sul mercato del lavoro: Una rassegna critica della letteratura empirica*

Questo articolo esamina la letteratura empirica emergente sugli effetti dell'intelligenza artificiale (IA) – e in particolare dell'IA generativa – sui risultati occupazionali. Nonostante l'ampio dibattito pubblico, le evidenze disponibili restano inconcludenti ed eterogenee. Basandosi su studi individuali e tre recenti meta-analisi, il contributo identifica quattro principali categorie di effetti: occupazione, salari, produttività e disuguaglianze. Vengono messi in luce contrasti nelle metodologie, nei contesti geografici e settoriali, nonché nelle definizioni stesse di IA, sottolineando l'importanza di distinguere tra automazione basata sui compiti e tecnologie generaliste come i modelli linguistici di grandi dimensioni. Complessivamente, non emergono evidenze coerenti di una disgregazione sistemica dei mercati del lavoro imputabile all'IA generativa. Se da un lato si registrano effetti di sostituzione in alcune professioni – soprattutto nei mercati digitali freelance – dall'altro si osservano guadagni di produttività o premi salariali in determinati settori. Le più recenti evidenze meta-analitiche confermano l'assenza di effetti medi statisticamente significativi dell'adozione dell'IA su occupazione e salari, e non rilevano segnali robusti di bias di pubblicazione. Questi risultati riflettono la natura ancora emergente della diffusione dell'IA e la sua adozione disomogenea. L'articolo si chiude delineando le priorità per la ricerca futura.

Parole chiave: Intelligenza artificiale generativa, mercato del lavoro, cambiamento tecnologico, letteratura empirica

DOI: 10.32049/RTSA.2025.4.04

1. Introduction

The rise of generative AI (*GenAI*) has renewed interest in how technology affects labor markets. Unlike earlier automation, which focused on routine tasks, *GenAI* can create original content – text, images, code, and even decisions – raising key questions about the future of

¹ Finanziato dall'Unione europea – Next Generation EU, Missione 4 Componente 1 CUP J53D23004880006.

work (Brynjolfsson, Mitchell and Rock, 2018; Hatzius *et al.*, 2023).

GenAI refers to a subset of machine-based systems that, given a set of objectives and data inputs, generate outputs that are not pre-programmed but rather created through adaptive learning or symbolic reasoning (Floridi, 2023). In contrast to narrow AI systems engineered for specific tasks – such as robotic arms, recommendation algorithms, or fraud detection – *GenAI* includes large language models (LLMs), image synthesis systems, and general-purpose agents capable of performing a broad set of functions across domains. From this perspective, *GenAI* may be interpreted as a potential General Purpose Technology (GPT) (Bresnahan and Trajtenberg, 1995) or as an *Invention in the Methods of Invention (IMI)* (Cockburn, Henderson and Scott, 2019), with implications not only for productivity, but also for organizational structures and employment dynamics (Bianchini, Müller and Pelletier, 2022; Damioli *et al.*, 2025)².

Although the transformative potential of generative AI is widely acknowledged, the empirical evidence on its actual impact on labor markets is still limited and no clear consensus has emerged. While some studies point to positive effects such as task reallocation, upskilling, and productivity improvements (Yang, 2022), others raise concerns about job displacement, wage polarization, and the decline of middle-skill occupations (Acemoglu and Restrepo, 2019b; Felten, Raj and Seamans, 2019; Wu *et al.*, 2024).

This article complements existing reviews (Biggi and Giuliani, 2021; Montobbio *et al.*, 2024b) by focusing specifically on AI in its generative, stand-alone form³. The aim is twofold: first, to give an account of the theoretical models and the empirical strategies employed in recent studies; and second, to provide original insights through a meta-analytic lens, which will provide quite striking results. While we will discuss that these results may reflect methodological and definitional heterogeneity across the studies analyzed, they also suggest that the nature of the aggregate effects of *GenAI* on labor markets is still unclear. This lack of consensus calls for prudence in policy design, emphasizing the need for adaptive and

² The term *generative AI* has gained prominence with the emergence of large-scale foundation models such as *ChatGPT*, *DALL·E*, and *GitHub Copilot*, which produce novel content based on learned data representations. For an overview, see Bommasani *et al.* (2021).

³ We therefore choose not to deal with robotics and process automation.

evidence-based approaches.

Several explanations are possible. First, adoption varies significantly across sectors, firms, and countries (OECD, 2023), which makes it difficult to generalize the available evidence. Second, measurement issues – such as relying on proxies for exposure or patent-based indicators – can make it hard to isolate causal effects. It is also possible that the nature of technological change is less disruptive than expected, advancing gradually or in ways that complement existing processes (Acemoglu and Restrepo, 2020). Finally, it may simply be too early to detect any impact: generative AI is a phenomenon still in its early stages, and its full impact may be yet to become visible.

The article is organized as follows. Section 2 clarifies key definitional boundaries and distinguishes *GenAI* from other forms of automation and digital transformation. Section 3 reviews the main theoretical frameworks linking technological change to labor market dynamics. Section 4 surveys the empirical literature on the effects of *GenAI* on employment, wages, productivity, and occupational structures. Section 5 presents and interprets the findings of meta-analyses. Section 6 concludes by discussing potential explanations for the heterogeneity in the evidence and reflecting on the implications for policy and research.

2. Definitional Challenges and Conceptual Approaches

The proliferation of studies on *artificial intelligence* and the future of work has brought renewed urgency to a longstanding challenge in empirical research: the lack of a shared, operational definition of AI. Without conceptual clarity, comparisons across studies and consistency in policy advice are compromised. This challenge is compounded when the discussion moves from AI in general to *GenAI systems* – large-scale models capable of producing novel content through self-supervised learning and adaptive inference. A conceptual distinction is therefore necessary between narrow, application-specific forms of AI and the new class of generalizable, *stand-alone generative systems*. In what follows, we argue that this distinction is not merely taxonomic, but structurally significant for

understanding how AI affects labor markets.

2.1. Institutional Definitions and the Emphasis on Autonomy

Institutional actors have made recent efforts to codify the boundaries of what constitutes an *AI system*. Article 3 (Definition 1) of the *European Union's AI Act – Regulation (EU) 2024/1689* – offers one of the most detailed legal definitions, stating that an *AI system* is «a machine-based system designed to operate with varying levels of autonomy, that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers from the input it receives how to generate outputs such as predictions, content, recommendations, or decisions». The OECD (2024) shares a similar view, representing *AI systems* as computational agents whose outputs are at least partially inferred – rather than deterministically coded – and capable of shaping human or digital environments through interaction⁴.

These definitions matter not only for regulation, but also for research design. Most notably, they allow for a distinction between *embedded* and *stand-alone AI systems*. An *embedded system* refers to AI used as a component within a broader technological or organizational infrastructure (e.g., a predictive module in *customer relationship management systems*); a *stand-alone system* operates independently and is visible in its effects (e.g., a *generative text model accessible through an application programming interface or a graphical user interface*). This distinction is relevant to labor market dynamics, as *stand-alone systems* directly interact with user input and can substitute or augment complex cognitive labor across occupations and sectors (Bommasani *et al.*, 2021; Bianchini, Müller and Pelletier, 2022; OECD, 2024)⁵.

⁴ See OECD (2024). These institutional definitions emphasize functional autonomy and adaptive inference, in contrast to earlier definitions focused on rule-based decision trees or embedded software components.

⁵ The notion of *stand-alone* is found also in the *European Union's AI Act*, Art. 3.1(b), which distinguishes between *embedded* and *non-embedded AI systems*. The distinction is vital for understanding which systems exert direct pressure on labor inputs, as opposed to those operating in the background of larger technical infrastructures.

2.2. From Task Exposure to Generative Capacity

The dominant empirical tradition in studies of AI and employment has been *task-based exposure modeling* (Frey and Osborne, 2017). These models estimate the probability that occupational tasks can be automated given technological feasibility and current system capabilities, often relying on static databases such as *O*NET* or *ESCO*. While useful in early mapping exercises, these models suffer from two main limitations.

First, they treat automation as a process in which machines entirely replace human tasks. This framework draws on the automation literature of the 1990s and early 2000s, centered on *industrial robotics* and *process control systems* (Acemoglu and Restrepo, 2019a; Dosi *et al.*, 2021). However, it is poorly equipped to capture the generative capability of recent systems, which often complement, rather than displace, labor by expanding the set of tasks or enabling higher-order functions (Wang, Gao and Agarwal, 2024).

Second, *task-exposure models* rely on *ex-ante expert assessments* or *technical benchmarks* rather than empirical data on actual AI implementation or usage. As a result, the links between technologies and jobs are often based on assumptions, which can overestimate or misclassify the risk of automation (Montobbio *et al.*, 2024a; Wu *et al.*, 2024). Studies using adoption or investment data suggest that AI technologies are often deployed to increase efficiency and support decision rather than for full automation, particularly in knowledge-intensive sectors (Yang, 2022).

Crucially, *task-exposure models* see AI as a collection of separate tools, while *GenAI* – because it can be easily applied to new tasks across many contexts – should be viewed as part of a broader technological shift.

2.3. Generative AI as GPT and Invention in the Methods of Invention (IMI): Theoretical Relevance

A better way to define *GenAI* comes from the literature on *General Purpose Technologies*

(*GPTs*). As first formulated by Bresnahan and Trajtenberg (1995), a *GPT* is characterized by three attributes: pervasiveness across sectors, continuous technical improvement, and a strong potential for generating complementary innovations. *GenAI* fits these criteria. Its functionality extends from language and image generation to code writing and data analysis; it is improving at an exponential pace via scaling laws; and it enables new practical uses and workflows across multiple industries (Cockburn, Henderson and Scott, 2019; Grashof and Kopka, 2023)⁶.

Some scholars have gone further by proposing that *GenAI* should be considered not merely as a *GPT* but as an *Invention in the Methods of Invention (IMI)*, i.e., a *meta-technology* that accelerates the rate of innovation itself (Cockburn, Henderson and Scott, 2019; Bianchini, Müller and Pelletier, 2022). *GenAI systems* reduce the marginal cost of experimentation in product development, customer interaction, and content creation. This has direct implications for the redefinition of labor, particularly in creative, managerial, and scientific domains where the production frontier is shaped less by repetition and more by novelty (Acemoglu and Restrepo, 2020).

Modern AI's generative abilities merit separate analysis in labor market studies. While automation in the past targeted routine tasks only, *GenAI* challenges the idea that cognitive and creative work is safe from machine competition, requiring updates to research models and measurement methods.

3. Theoretical Perspectives on AI and Labor Market Dynamics

Since the 1800s, economists have analyzed the impact of technology on jobs and wages. Nowadays, the arrival of *GenAI* has made this debate more pressing, forcing us to consider which types of jobs can be replaced by machines, and which ones can be enhanced or supported by them, being complementary to AI.

⁶ Bommasani *et al.* (2021) and Bianchini, Müller and Pelletier (2022) provide compelling empirical support for viewing large-scale generative systems as infrastructure-like technologies with generalizable utility.

This section reviews the main lines of economic thought relevant to the analysis of AI and labor, with particular attention to the conditions under which technological innovation becomes either disruptive or labor-augmenting.

3.1. Technological Change: From Classical Intuition to Structural Disruption

David Ricardo (1821) acknowledged that new machinery could displace workers in the short term but held that such displacements would eventually be offset by rising demand and new production opportunities⁷. Joseph Schumpeter (1911) deepened this dynamic view by emphasizing «creative destruction»: the process whereby old industries and job structures are destroyed and replaced by new ones, with net employment gains driven by innovation-led expansion.

This optimistic line has been challenged by John Maynard Keynes (1931, p. 364) – who warned of «technological unemployment» caused by the economy's inability to redeploy labor at the pace required by labor-saving innovations – and, in contemporary economics, by models that treat automation as a force capable of reducing the share of labor in national income (Acemoglu and Restrepo, 2019a).

Whether these effects materialize depends crucially on which types of tasks are affected, the elasticity of substitution between capital and labor, and the macro-level feedback loops between productivity, prices, and demand. The arrival of *GenAI* reactivates this debate by introducing a class of technologies that, unlike previous ones, target non-routine cognitive tasks, such as writing, analysis, and design – tasks traditionally considered complements rather than substitutes to skilled labor (Autor, 2015).

⁷ Ricardo's early reflections on machinery were notably ambivalent. In his 1821 edition, he revised earlier optimism and acknowledged that capital accumulation could reduce labor demand in the short term – a view later echoed in modern models of technological displacement.

3.2. The Substitution-Reinstatement Spectrum

Acemoglu and Restrepo (2019b) introduce the idea of task-based production functions in which automation displaces certain labor tasks (*substitution effect*) while simultaneously creating demand for new, higher-productivity tasks (*reinstatement effect*). According to this approach, the net impact of AI on employment depends not only on the scope of automation, but also on the economic system's ability to shift labor toward tasks that either complement AI or emerge as new forms of work.

This model offers a corrective to earlier automation paradigms that assumed technological change to be either uniformly labor-saving or uniformly productivity-enhancing. It shows, for instance, that reinstatement effects are more likely when AI complements rather than replaces human judgment – as may be the case in hybrid decision-making environments (e.g., legal, medical, or editorial domains). However, when the new tasks created by AI require specialized skills or capital-intensive infrastructure, the transition can increase inequality or lead to job polarization (Felten, Raj and Seamans, 2019; Wu *et al.*, 2024)⁸.

The effect of AI on wages and jobs is not driven solely by what technology can do. It also depends on how it is used, who adopts it, and on the broader economic and institutional context in which it is introduced. Institutional variables – e.g. collective bargaining power, training systems, and industrial policy – play a role in mediating outcomes. In Casey's (2024) extension of the Acemoglu-Restrepo framework, a Leontief production function is combined with wage bargaining dynamics, leading to the result that labor-saving technical change causes persistent unemployment when bargaining power is strong, but reallocation channels are weak.

⁸ For empirical evidence on wage polarization linked to technological change, see Autor, Mindell and Reynolds (2022). For a theoretical formalization of unemployment under directed technical change, see Casey (2024).

3.3. Sectoral and Innovation-Based Approaches

Beyond task-based models, a complementary line of inquiry draws from innovation economics. Dosi *et al.* (2021) distinguish between *embodied* technological change – such as the replacement of machinery – and *disembodied* change, such as investment in research and development (R&D). Their analysis shows that the employment effects of innovation differ significantly depending on whether the locus of innovation is in upstream (R&D-intensive) or downstream (production-intensive) sectors. Employment growth is more likely in upstream sectors due to the innovation-driven expansion of value-added activities, while downstream sectors face stronger substitution pressures.

This sectoral heterogeneity is particularly relevant for *GenAI*, which is often deployed as a digital infrastructure in upstream activities (e.g., product design, market analysis, customer engagement). In such settings, AI may support employment by increasing task productivity and demand elasticity. But in downstream applications – such as automated content generation or translation – the same technologies may displace workers directly without clear short-term compensatory channels (Felten, Raj and Seamans, 2019; Domini *et al.*, 2022)⁹.

Moreover, the effects of AI may depend on firms' strategic behavior and their position in the innovation ecosystem. Bianchini, Müller and Pelletier (2022) argue that AI constitutes an *innovation-enabling technology*, capable of amplifying the rate and direction of subsequent innovations. When firms use *GenAI* to accelerate product development cycles or tailor services, they may generate complementary employment effects – particularly in roles that involve oversight, refinement, or integration.

3.4. Toward a Behavioral-Structural Synthesis?

Finally, some recent models try to connect behavioral and structural perspectives in the

⁹ Domini *et al.* (2022) find that while automation and AI investments slightly increase average wages, they do not significantly affect within-firm wage inequality – suggesting limited distributive disruption but also limited progressive effects.

analysis of AI's impact. For instance, Grashof and Kopka (2023) suggest that AI's influence depends on organizational size, absorptive capacity, and market power. Larger firms are more likely to internalize the benefits of AI, possibly reinforcing concentration and reducing competitive pressure in labor markets. In this sense, the debate over AI and employment is also a debate about industrial structure, competition policy, and the diffusion of technological capabilities.

These approaches suggest that there is not a unique theoretical expectation for how AI affects labor. Rather, the outcome depends on a complex interaction of substitution and complementarity, firm strategy and sectoral dynamics, institutional context and policy design. This requires empirical investigations that are capable of disentangling these mechanisms and of using proper identification strategies.

4. Mapping the Empirical Landscape

This section analyzes the empirical literature on the relationship between AI and labor. It classifies existing studies into four main categories: methods of empirical research, types of labor outcomes examined, units of analysis (such as firms or sectors), and time horizons (short- vs. long-term effects).

4.1 Methodological Lenses: From Exposure Indices to Natural Experiments

The first category of classification concerns the methodologies. Much of the early empirical work relies on *task-exposure indices*, following the seminal work by Frey and Osborne (2017), which estimates the automation potential of occupations based on *O*NET* descriptors¹⁰. Subsequent contributions include Felten, Raj and Seamans (2019) *AI*

¹⁰ The Frey and Osborne (2017) study used Gaussian process classifiers to predict computerization probabilities based on 70 occupations, then extrapolated to 702 *O*NET*-coded occupations. This extrapolation has been criticized for high false-positive rates and limited construct validity – especially regarding service-sector and manual roles.

*Occupational Impact Scores*¹¹ and Cockburn, Henderson and Scott (2019) *patent-based measure of AI exposure*. Although these studies do not measure actual labor market effects, they are relevant because they provide a starting point for understanding which occupations and tasks are most susceptible to AI-driven change, thereby representing the basis of subsequent empirical research.

Later studies often exploit timing differences in AI exposure as a natural experiment – for instance, language-intensive occupations being affected earlier by the release of *ChatGPT*. Using these differences across occupations or regions, they construct treatment and control groups, allowing researchers to compare outcomes before and after exposure, helping isolate the effect of AI. This approach provides a clearer view of labor market effects, though typically in a narrower, short-term context, as it focuses on specific events and captures early impacts rather than long-term adjustments.

4.2 Labor Market Outcomes: From Employment and Wages to Inequality

The literature explores several dimensions of labor market transformation under the influence of *GenAI*. We classify the empirical literature on market outcomes into four categories: (a) employment effects, (b) wage impacts, (c) productivity and value creation, and (d) inequality.

This classification allows us to distinguish between substitution and reallocation mechanisms, between wage and income shocks and skill premia, and between redistributive and efficiency-related outcomes.

¹¹ Felten, Raj and Seamans (2019) introduced the *AI Occupational Exposure* score using text similarity between AI capabilities and *O*NET* work activities. Unlike Frey and Osborne (2017), their index allows temporal updating but still does not distinguish between AI types (e.g., *GenAI* vs. symbolic AI).

4.2.1 Employment Effects: Displacement, Reallocation, and Stability

Several studies on employment effects study whether *GenAI* leads to job losses, job creation, or occupational shifts. This includes not only net employment outcomes, but also the reorganization of tasks within organizations and the demand for new or transformed occupations.

One of the most robust findings of labor displacement is found in Teutloff *et al.* (2025), who study the impact of *ChatGPT*'s release on *Upwork*, a global freelancing platform. The authors identify a large drop in the demand for services and average earnings of freelancers in writing, translation, and other language-intensive categories. In contrast, programmers and data analysts were not hit by a statistically significant change. Interestingly, there is no evidence of a decrease in the pre-treatment period (i.e., the period before the introduction of *ChatGPT*), which reinforces the plausibility of a *GenAI*-induced shock¹². Their results suggest that *GenAI* adoption can cause rapid substitution (in this case platform-mediated). This is especially true when tasks can be broken into discrete components, are performed digitally, and rely on cognitive routines that can be replicated by AI systems.

Bonney, Erhard and Stern (2024) adopt a firm-level perspective, using survey data from U.S. businesses to assess how employers respond to AI. They find that only 5% of firms report any change in employment due to AI, while over 27% report the replacement of specific tasks. Interestingly, more firms anticipate future increases in employment than reductions. This suggests that while AI may alter job roles, it does not necessarily lead to job losses. These findings contrast with Teutloff *et al.* (2025), suggesting that the impact of *GenAI* on employment may be conditional on the organizational context. While some tasks may be replaced, this does not necessarily lead to job destruction when firms choose to reorganize work internally.

These findings can be compared to earlier automation waves. For instance, Graetz and

¹² Teutloff *et al.* (2025) use a difference-in-differences design over 1.2 million contracts on Upwork to estimate *GenAI* exposure. They find that in the five months after *ChatGPT*'s release, treated occupations experienced up to a 20% reduction in earnings, with no similar trend in placebo periods. The effect is most pronounced in writing, editing, and translation – occupations most substitutable by LLMs.

Michaels (2018), using panel data from 17 countries between 1993 and 2007, found that robot adoption increased productivity and value added without reducing total employment, although job losses did occur among low-skilled manufacturing workers. This confirms that general-purpose automation technologies – while potentially disruptive at the sectoral level – do not necessarily generate large-scale unemployment, especially when counterbalanced by demand-side effects and task reallocation¹³.

Stephany and Teutloff (2025) confirm these patterns in a related freelancing study. They observe a reallocation of demand across occupational categories, with increased competition in high-*GenAI*-exposed segments and relative stability or even growth in marketing, graphic design, and software development. Moreover, they document a shift in demand from high-wage to low-wage countries, suggesting a global redistribution of freelance labor rather than a net contraction. This finding adds a geographic dimension to the employment effects of *GenAI*, mediated by platform design and pricing structures¹⁴.

4.2.2 Wage Impacts: Polarization, Premiums, and Erosion

The impact of AI on wages – namely, which workers gain and which lose – depends not only on how exposed different occupations are to AI technologies (as discussed in Section 4.1, with reference to Frey and Osborne, 2017, and Felten, Raj and Seamans, 2019), but also on the specific tasks involved and on institutional factors such as wage-setting mechanisms, labor protections, and employment arrangements.

Teutloff *et al.* (2025), observe a direct wage erosion for content producers following the release of *ChatGPT*. The earnings drop is large (7–20%) and sustained, suggesting income

¹³ Graetz and Michaels (2018) use IFR robot data merged with sectoral labor market outcomes across 17 developed countries. They find that robot adoption raised labor productivity and GDP growth without reducing aggregate employment, though it did displace low-skilled workers in manufacturing.

¹⁴ Stephany and Teutloff (2024) also use data from online freelancing platforms, including Upwork, but with more detailed occupational segmentation. They show that translation, copywriting, and transcription experienced declines in both demand and average remuneration, while categories such as illustration, software development, and virtual assistance remained relatively stable or showed greater resilience.

substitution by machine-generated outputs.

In contrast, Pouliakas, Santangelo and Dupire (2025) find strong wage premia for AI developers in the EU labor market. Using survey data across 29 countries, they estimate that workers in AI-specific programming roles earn significantly more than other programmers and professionals, even when controlling for education, experience, and performance-based pay schemes. Premia persist even after decomposing wage differentials by task content and worker control and discretion, indicating that the premium is not entirely explained by observable characteristics. This finding is consistent with models of temporary skill scarcity and signaling, which may decline as AI competences diffuse – unless new layers of technical specialization emerge¹⁵.

4.2.3 Productivity and Value Creation: Limited but Emerging Signals

The impact of *GenAI* on firm-level productivity and value creation remains difficult to measure, given the lack of detailed input-output data and the relatively recent introduction of *GenAI* in production processes. However, we can point to some indirect evidence.

Bonney, Erhard and Stern (2024) report that over 15% of firms using AI expect both future gains in efficiency and stable employment levels. This suggests that they believe productivity improvements will come from the reorganization of work within the existing workforce. This finding is particularly evident in the service sectors. Such expectations are consistent with the theoretical predictions illustrated above, namely that *GenAI* enhances task efficiency and human-AI complementarity, particularly for knowledge-intensive firms with sufficient absorptive capacity (Cockburn, Henderson and Scott, 2019; Acemoglu and Restrepo, 2020).

However, recent evidence suggests that the short-term effects of generative AI on productivity are still limited. For example, a 2024 survey by the Federal Reserve Bank of St. Louis found that *GenAI* allows time to save of approximately 5.4% of work hours, equivalent

¹⁵ Pouliakas, Santangelo and Dupire (2025) find that AI programmers earn, on average, 16% more than their non-AI counterparts, controlling for demographic and job-level characteristics.

to only 1.1% of total work time (Bick, Blandin and Deming, 2025). Similarly, experimental evidence from Noy and Zhang (2023) shows that the use of AI may increase the productivity for individual workers depending on tasks, but the extent of these effects is narrow. Moreover, AI adoption is still limited overall: according to McKinsey & Company (2025), only 1% of firms identify as “mature” adopters.

Brynjolfsson, Rock and Syverson (2021) offer a theoretical explanation for such limited impact in the short run. In their model, AI follows a *productivity J-curve*, similar to other general-purpose technologies introduced in the past. Early stages of adoption show slow or even declining productivity, as firms must first invest in complementary assets such as skills, infrastructure, and process reorganization. Substantial productivity gains tend to materialize only in later stages, once these intangible investments begin to pay off.¹⁶

Going back to the freelance market, Stephany and Teutloff (2024) note that some freelancers are reorienting their profiles to AI *integration* or *prompt engineering* suggesting new value-creating tasks emerging within the *GenAI* ecosystem.. While not productivity metrics in the narrow sense, such repositioning signals a possible positive-sum trajectory for certain occupational groups that actively align with *GenAI* capabilities.

4.2.4 Inequality: Compression, Segmentation, or Polarization?

Finally, the literature offers diverging findings on inequality. Yuan, Sun and Chen (2025) provide empirical evidence on within-firm inequality, showing that AI adoption in Chinese manufacturing firms is associated with a reduction in wage dispersion. The compression is stronger in firms with weaker managerial power and larger pools of skilled labor, suggesting that AI may enable tighter performance monitoring and reduce managerial rents¹⁷. These

¹⁶ Brynjolfsson, Rock and Syverson (2021) formalize the idea of a productivity J-curve: new general-purpose technologies initially reduce measured productivity due to the lag between adoption and complementary organizational investments.

¹⁷ Yuan, Sun and Chen (2025) use a panel of Chinese A-share manufacturing firms (2011–2019), and measure within-firm inequality as the executive–employee pay ratio. By instrumenting firm-level AI adoption with U.S. robot exposure, they find a robust reduction in the wage gap associated with AI adoption.

results support the idea that AI can flatten organizational hierarchies under certain governance regimes.

By contrast, Pouliakas, Santangelo and Dupire (2025) find evidence of occupational wage segmentation between AI-exposed and non-AI workers, with high task autonomy and problem-solving linked to larger wage gaps. While not directly measuring within-firm inequality, the results imply a rising skill premium and potential polarization, at least in the short term.

Stephany and Teutloff (2024) add an inter-country dimension to the inequality debate. They find that demand for freelancers increasingly shifts to low-income countries post-*GenAI*, suggesting a possible compression of global wage gaps – though driven by downward earnings pressure in advanced economies rather than upward convergence.

Further evidence on labor displacement in online markets comes from Hui, Sundararajan and Tang (2024), who implement a randomized field experiment on a major global freelancing platform. The study evaluates how *GenAI* affects hiring outcomes by posting matched job listings – one AI-exposed (e.g., writing) and one less exposed (e.g., coding) – across different time periods. The authors find a 15% decrease in the likelihood of hiring freelancers for AI-exposed jobs in the months following the release of *ChatGPT*, with no comparable decline in control categories. Importantly, these effects are concentrated in the short term and in low- to medium-skill projects, suggesting that the impact of *GenAI* is especially acute in task environments characterized by low switching costs, high substitutability, and algorithmic mediation. The platform-mediated nature of the experiment reinforces the observation that online labor markets act as early indicators of technological disruption, often amplifying substitution pressures due to their price transparency and client-side experimentation incentives.

4.3 Scope, Gaps, and Emerging Patterns

The recent empirical contributions reveal an evolving field. While the methodologies

employed vary widely – from task-exposure indices to quasi-experimental designs – some convergences are beginning to appear, particularly regarding the type of tasks and occupations where *GenAI* seems to matter most.

The most robust signals of labor displacement arise in language-intensive, low-entry-barrier occupations, especially those in digital labor markets. Both Teutloff *et al.* (2025) and Stephany and Teutloff (2024) find that, following the release of *ChatGPT*, freelance writers, translators, and editors experienced substantial declines in both the demand for their services and their earnings. These occupations are particularly exposed not only because their outputs are easily replicable by large language models, but also because the platforms that mediate their employment are designed for rapid, often one-shot, and typically unrepeated matching, with high sensitivity to price. In such contexts – deprived of reputational concerns and focused on routine tasks – *GenAI* functions as an immediate substitute, particularly for clients who prioritize speed and cost over interpretive precision or creative insight.

Yet this picture of substitution is far from universal. In more structured organizational environments, such as the firms studied by Bonney, Erhard and Stern (2024), *GenAI* adoption has not (at least yet) led to visible changes in employment. On the contrary, a slight majority of AI-using firms anticipate future employment increases, and most report no change in occupation despite task-level adjustments. This suggests a reorganization of work rather than its reduction and supports the view that technological displacement is often mediated by firm strategy, skill availability, and complementarities with existing workflows.

Regarding wage, the picture is equally variegated. For workers performing tasks that are complementary to *GenAI* – such as AI developers, infrastructure engineers, or prompt specialists – the technology seems to generate premiums, at least in the short term. The study by Pouliakas, Santangelo and Dupire (2025) demonstrates that even when accounting for traditional wage determinants such as education and experience, AI-specific roles guarantee significantly higher compensation. However, such premium is not distributed evenly across tasks or geographical zones. Moreover, as the number of AI-skilled workers increases, these premia may erode, particularly in the case of tasks that are standardized, interchangeable, and widely distributed. Thus, while *GenAI* may currently reward technical expertise as a scarce

asset, its distributive impact over time is difficult to predict.

One must recall the temporality of this literature, which reflects the early stages of *GenAI* adoption – often within the first year following the public release of tools like *ChatGPT*. As Bonney, Erhard and Stern (2024) note, fewer than 10% of surveyed firms reported using AI in a systematic way as of early 2024. The absence of large, economy-wide effects may therefore reflect not insignificance, but immaturity: the unfolding of a technological transition whose labor implications will become more visible as adoption deepens, models improve, and firms reorganize. In this sense, the current evidence should be read as the illustration of a potential trend, offering not conclusive patterns but indicating the most likely path forward.

5. Meta-Analytic Perspectives on AI, Employment, and Productivity

Individual studies may differ in methods, data, and scopes. Thus, it is not easy to spot empirical regularities and general trends. Meta-analysis, a quantitative methodology largely employed in the hard sciences but not so commonly accepted in the social sciences, addresses this issue by combining results from multiple studies to estimate overall effects – such as AI's impact on jobs or productivity – revealing patterns not clearly emerging from 'qualitative' comparison of single papers.

Meta-analysis is a useful tool to assess the average effects of a phenomenon, accounting for variation across context, highlighting differences between sectors, countries, or skill groups in the process. Equally important is its role in addressing the so-called *publication bias*, namely the tendency of journals to disproportionately publish studies with large or statistically significant findings or privileging the publication of certain results rather than others. This bias can make certain effects appear stronger or more consistent than they actually are. Meta-analytic techniques offer tools to detect and correct for such distortions, helping to

build a more accurate picture of the phenomenon under study¹⁸.

In this section, we examine three recent meta-analyses relevant to the core questions raised by this paper, namely: what are the effects of automation and AI on employment, wages, and productivity? The first study, by Guarascio, Piccirillo and Reljic (2024), takes into account results from papers dealing with the impact of the adoption of industrial robots on employment and wages. The second, by Coupé and Wu (2025), analyzes empirical results from studies on *GenAI* and productivity. The third, by Carbonara, Santarelli and Tripathi (2025), focuses on the labor market outcomes from adoption of AI as a potential general-purpose technology. Together, these studies represent the most comprehensive attempts to date to systematize the empirical knowledge in this rapidly evolving field.

5.1 Robots, Employment, and Wages

Guarascio, Piccirillo and Reljic (2024) conduct a meta-analysis of 30 empirical studies on the impact of robot adoption on employment and wages across developed economies. They find that robotization has had a modest but statistically significant negative effect on the employment of low- and medium-skilled workers. In contrast, high-skilled employment appears largely unaffected, and in some cases even positively impacted. These findings are consistent with the *task-biased* model of technological change, in which machines substitute for routine or codifiable tasks while complementing more abstract or analytical functions.

Importantly, the authors find considerable variation across countries and sectors. For example, manufacturing sectors in countries with stronger labor protections exhibit smaller displacement effects, suggesting that institutional context mediates technological impact. Furthermore, more recent studies tend to report smaller effects than less recent ones.

Besides, Guarascio, Piccirillo and Reljic (2024) find evidence of publication bias: studies with larger or more negative effects have a greater likelihood of being published, especially

¹⁸ See Stanley and Doucouliagos (2012). For a methodological overview of publication bias in meta-analysis, see Rothstein, Sutton and Borenstein (2005).

in top journals. Once this bias is accounted for, the average displacement effect becomes slightly smaller, although still negative and statistically significant. The fact that robotization is not found to be conducing to mass unemployment but only to increased labor market segmentation calls more for Active Labor Market Policies (e.g., targeted retraining programs) than for strong policy actions to mitigate its effects.

5.2 Generative AI and Productivity Gains

The second meta-analysis, by Coupé and Wu (2025), presents a comprehensive synthesis of the productivity effects of generative artificial intelligence (*GenAI*). They gather data for over 1,000 effect-size estimates from 45 studies that assess how *GenAI* tools affect output, task efficiency, or work quality in different settings, from coding and writing to design and customer support.

The authors show that *GenAI* adoption is associated with an average productivity gain of approximately 17%, with the largest gains observed in quantitative and repetitive tasks benefiting from *GenAI*'s speed and pattern-recognition capabilities. In contrast, tasks requiring subtle judgment, originality, or contextual understanding see smaller or negligible improvements. Interestingly, the effects are strongest in experimental or pilot settings, suggesting that the observed gains may reflect short-term efficiency boosts that have not yet been tested in sustained real-world conditions.

Coupé and Wu (2025) also conduct a careful examination of publication bias, finding evidence of bias toward large and positive estimates, especially in less recent studies. This means that the studies reporting large productivity gains from *GenAI* are also those most likely to be published. When this bias is statistically accounted for, the average productivity gain falls to around 7–10%, suggesting that the early hype around *GenAI* may have inflated expectations.

From an economic perspective, the meta-analysis conducted by Coupé and Wu's (2025) suggests that *GenAI* can significantly improve task performance, particularly in structured or

rule-based domains where outputs are well-defined and time-sensitive. However, the authors emphasize that most of the evidence they deal with comes from short-term, experimental settings, and recommend being careful in generalizing these results to broader settings. While the meta-analysis does not assess employment outcomes directly, the potential for productivity gains raises important questions about whether such improvements will complement human labor or, in some contexts, substitute for it – a question that remains open and requires further study, as we argued in Section 4 above.

5.3 GenAI and Employment

Carbonara, Santarelli and Tripathi (2025) conduct a meta-analysis of 19 studies published since 2018, covering a broad range of labor market outcomes including employment, wages, productivity, and task reallocation. In what follows, the focus is on their findings related to employment and wages, which are most directly comparable to the results presented by Guarascio, Piccirillo and Reljic (2024), while also offering a complementary perspective to the productivity-focused analysis by Coupé and Wu (2025).

Carbonara, Santarelli and Tripathi (2025) adopt a deliberately inclusive definition of AI, encompassing both general-purpose machine learning and more recent generative models, and find no consistent or statistically significant average effect of AI adoption on net employment levels. This result holds across different modeling strategies, time periods, and study types¹⁹. Moreover, the authors do not find robust evidence of publication bias. Unlike the other two meta-analyses, here data show a rather symmetric distribution of results, with both positive and negative estimates well represented. This suggests that the distribution of findings in the literature analyzed in this paper is relatively balanced, which may lead to the conclusion that, so far, there is no clear evidence of a systematic bias in how results are produced or published.

¹⁹ Estimates derive from a range of econometric models (including, among others, linear regression, difference-in-differences, and instrumental variables), covering time periods from the early 2000s to 2023. Study types vary by unit of analysis, including firm-level, regional, and individual-level data.

This null result, however, does not imply that AI has no effect in practice. Rather, it reflects the fact that AI's impact on employment is complex, multifaceted, and still evolving and subject to change. As discussed in Section 4, such heterogeneity possibly reflects differences in AI application, workforce characteristics, and organizational response. It also highlights the fact that adoption remains limited, and substantial transformations may still lie ahead.

In this sense, the findings in Carbonara, Santarelli and Tripathi (2025) offer a balanced empirical assessment of widespread concerns about large-scale employment losses due to *GenAI*. For instance, Hatzius *et al.* (2023) conclude that *GenAI* could replace around 300 million full-time jobs globally. On a similar vein, an International Monetary Fund (IMF) report warns of “profound concerns” about labor disruption and rising inequality (IMF, 2024a).²⁰ Yet, other analyses offer a more favorable view. Again, the IMF, in a different report, notes that in advanced economies, AI is more likely to complement than replace existing jobs and may boost productivity and wages over time (IMF, 2024b). Amidst these contrasting views, policies addressing *GenAI* and its use should be anticipatory and evidence-based, aimed at monitoring real transformations in work content, supporting workforce transitions, and addressing distributional risks as they emerge.

In conclusion, the three meta-analyses discussed above offer a view that is more balanced than the current debate. The effects of automation and AI are real but not uniform, measurable but often modest, and depend highly on sector and type of occupation.

In light of the literature surveyed in Section 4, these findings reinforce the idea that we are still in a transitional phase, where local shocks coexist with system-level inertia. To be precise: the absence of strong empirical evidence in one sense or the other does not dismiss concerns. As mentioned above, the challenge for economists and policymakers alike is not only to measure impact, but to shape it – through policies that can steer technological change toward more inclusive and sustainable trajectories.

²⁰ Media outlets such as the Financial Times highlights the risk of mass layoffs in the tech sector and warns that white-collar roles are being rapidly affected in technology and media, suggesting that «the great AI jobs disruption is under way». See Financial Times, 2025.

6. Interpretation and Policy Implications

Despite the rapidly expanding body of literature on AI and labor market outcomes, no clear consensus has yet emerged. The findings reviewed in this paper point in multiple directions: some studies report displacement, others augmentation; some suggest wage compression, others polarization.

The first possible explanation is that there are persistent measurement issues. Studies vary considerably in how they define AI, particularly *GenAI* – ranging from machine learning systems to large language models and generative tools²¹. Similarly, labor market outcomes are not consistently measured: some papers examine net employment change, others focus on task substitution, wage dispersion, or rising inequality²². These definitional inconsistencies make it difficult to compare findings or identify robust empirical regularities.

On the other hand, the extent of these effects depends on lags in adoption. While AI technologies evolve rapidly, their use and diffusion across firms and sectors is gradual and possibly uneven.

Moreover, firm-level heterogeneity plays a crucial role. The capacity to adopt, absorb, and benefit from AI varies widely across firms, depending on their size, their endowment of human capital, and digital maturity. Larger firms may hold the technical capabilities to use AI to increase efficiency, whereas smaller firms may lag behind (Autor, Mindell and Reynolds, 2022).

Finally, current research is constrained by data limitations. Many studies rely on proxies – such as occupation-level exposure scores, online job ads, or survey self-reports – rather than direct measures of AI adoption and use. Moreover, many datasets lack the granularity needed to trace intra-firm dynamics, task-level reallocation, or the temporal unfolding of impacts (Arntz, Gregory and Zierahn, 2016).

Taken together, these challenges suggest that the mixed results found in the literature may stem less from genuine differences in how AI affects the labor market and more from these

²¹ See Brynjolfsson and Mitchell (2017) and Frey and Osborne (2017).

²² See Carbonara, Santarelli and Tripathi (2025).

underlying measurement and data problems. Recognizing this limitation is crucial: only by improving the comparability of concepts and indicators can future research move toward a clearer and more reliable understanding of how AI is reshaping work and employment.

On the other hand, public debate often runs disconnected from empirical evidence (Eubanks, 2018). Media narratives tend to highlight salient, high-profile cases – such as job losses in large companies or involving many workers – consequently overlooking the many contexts in which AI has minimal or even positive effects.

From a policy perspective, this tension calls for prudence. Premature interventions – such as AI taxes or deployment restrictions – risk curbing innovation or generating unintended distortions. At the same time, regulatory passivity may leave vulnerable workers exposed. As argued above, the challenge is to develop adaptive, evidence-based policies that can evolve with the technology. Such policies include improved monitoring systems, targeted reskilling programs, support for job transitions, and sector-specific social dialogue. Crucially, governments should invest in better data infrastructures to track AI adoption and its impacts more systematically (Carbonara and Santarelli, 2023).

Among the various policy instruments mentioned above, automation taxes have received much attention. Taxing firms that displace workers by implementing automation techniques leads them to internalize the social costs of technological change. However, such measures could discourage beneficial innovation (Acemoglu and Restrepo, 2019b).

Recent experimental evidence offers valuable insights. In a survey experiment with 2,000 U.S. entrepreneurs, Carbonara, Focacci and Santarelli (2024) find that a robot tax – a fiscal surcharge on firms using automation – significantly reduces the likelihood of a reduction in the workforce. By contrast, offering automation subsidies appears to encourage labor cuts. Notably, the robot tax has a larger behavioral effect than the equivalent reward, suggesting that penalties may be more effective than incentives in curbing automation-induced job losses. However, the authors also warn that such policies must be carefully designed to avoid negative innovation spillovers.

Beyond economic instruments, there are ethical and political economy concerns as well. The distributional implications of AI raise questions about justice and accountability.

Drawing on Rawlsian principles, some scholars argue that technological transitions should be governed in ways that protect the most vulnerable, and that the benefits of automation should be shared equitably across society (Eubanks, 2018). Others stress the need for inclusive policymaking and public engagement on which uses of AI are socially desirable or acceptable – particularly because labor is not just a means of income, but a source of identity and dignity (Suskind, 2020).

In sum, the goal should not be to prevent the adoption of AI, but to monitor it – through institutions and proper incentives – so that it aligns with broader goals of social resilience and economic inclusion.

Reference List

- Acemoglu D., Restrepo P. (2019a). Artificial intelligence, automation and growth. In Agrawal A., Gans J., Goldfarb A., eds., *The economics of artificial intelligence: An agenda*. Chicago: University of Chicago Press. DOI: 10.7208/chicago/9780226613475.003.0003.
- Acemoglu D., Restrepo P. (2019b). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33, 2: 3. DOI: 10.1257/jep.33.2.3.
- Acemoglu D., Restrepo P. (2020). The wrong kind of AI? Artificial intelligence and the future of labor demand. *Cambridge Journal of Regions, Economy and Society*, 13, 1: 1. DOI: 10.1093/cjres/rsz022.
- Arntz M., Gregory T., Zierahn U. (2016). The risk of automation for jobs in OECD countries: A comparative analysis. OECD Social, Employment and Migration Working Papers, 189. Paris: OECD Publishing. DOI: 10.1787/5jlz9h56dvq7-en.
- Autor D., Mindell D., Reynolds E. (2022). *The work of the future: Building better jobs in an age of intelligent machines*. Cambridge, MA: MIT Press. DOI: 10.7551/mitpress/14174.001.0001.
- Autor D.H. (2015). Why are there still so many jobs? The history and future of workplace

automation. *Journal of Economic Perspectives*, 29, 3: 3. DOI: 10.1257/jep.29.3.3.

Bianchini S., Müller M., Pelletier P. (2022). Artificial intelligence in science: An emerging general method of invention. *Research Policy*, 51, 10: 104604. DOI: 10.1016/j.respol.2022.104604.

Bick A., Blandin A., Deming D. (2025). The impact of generative AI on work productivity. Federal Reserve Bank of St. Louis. *St. Louis Federal On the Economy*, 27 February. Retrieved from: <https://www.stlouisfed.org/on-the-economy/2025/feb/impact-generative-ai-work-productivity> (23/12/2025).

Biggi G., Giuliani E. (2021). The noxious consequences of innovation: What do we know? *Industry and Innovation*, 28, 1: 1. DOI: 10.1080/13662716.2020.1726729.

Bommasani R., Hudson D.A., Adeli E., Altman R., Arora S., von Arx S., Bernstein M.S., Bohg J., Bosselut A., Brunskill E., Brynjolfsson E., et al. (2021). On the opportunities and risks of foundation models. *arXiv preprint*. DOI: 10.48550/arXiv.2108.07258.

Bonney J., Erhard L., Stern S. (2024). Early evidence on the labor market impacts of artificial intelligence: New insights from the Business Trends and Outlook Survey. *Economics Letters*, 235: 111769. DOI: 10.1016/j.econlet.2024.111769.

Bresnahan T.F., Trajtenberg M. (1995). General purpose technologies “engines of growth”? *Journal of Econometrics*, 65, 1: 83. DOI: 10.1016/0304-4076(94)01598-T.

Brynjolfsson E., Mitchell T. (2017). What can machine learning do? Workforce implications. *Science*, 358, 6370: 1530. DOI: 10.1126/science.aap8062.

Brynjolfsson E., Mitchell T., Rock D. (2018). What can machines learn and what does it mean for occupations and the economy? *AEA Papers and Proceedings*, 108, 5: 43. DOI: 10.1257/pandp.20181019.

Brynjolfsson E., Rock D., Syverson C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13, 1: 333. DOI: 10.1257/mac.20180367.

Carbonara E., Focacci C.N., Santarelli E. (2024). Mitigating the labor displacing effects of automation through a robot tax: Evidence from a survey experiment. *Economics of Innovation and New Technology*, 33, 8: 1145. DOI: 10.1080/10438599.2023.2293031.

- Carbonara E., Santarelli E. (2023). Artificial intelligence and robots: A threat or an opportunity for SMEs and entrepreneurship? In Carbonara E., Tagliaventi M.R., eds., *SMEs in the digital era: Opportunities and challenges of the digital single market*. Cheltenham: Edward Elgar. DOI: 10.4337/9781803921648.00013.
- Carbonara E., Santarelli E., Tripathi I. (2025). Assessing the impact of AI on labor market outcomes: A meta-analysis. *SSRN Working Paper*. DOI: 10.2139/ssrn.5126345.
- Casey G. (2024). Unemployment and the direction of technical change. *European Economic Review*, 168: 104802. DOI: 10.1016/j.euroecorev.2024.104802.
- Cockburn I.M., Henderson R., Scott S. (2019). The impact of artificial intelligence on innovation: An expository analysis. In Agrawal A., Gans J., Goldfarb A., eds., *The economics of artificial intelligence: An agenda*. Chicago: University of Chicago Press. DOI: 10.7208/chicago/9780226613475.003.0004.
- Coupé T., Wu W. (2025). *The impact of generative AI on productivity: Results of an early meta-analysis*. Working Paper no. 9, Department of Economics and Finance, School of Business and Economics, University of Canterbury, Christchurch, New Zealand. Retrieved from: <https://repec.canterbury.ac.nz/cbt/econwp/2509.pdf> (23/12/2025).
- Damioli G., Van Roy V., Vertesy D., Vivarelli M. (2025). Is artificial intelligence leading to a new technological paradigm? *Structural Change and Economic Dynamics*, 72: 347. DOI: 10.1016/j.strueco.2024.12.006.
- Domini G., Grazzi M., Moschella D., Treibich T. (2022). For whom the bell tolls: The firm-level effects of automation on wage and gender inequality. *Research Policy*, 51, 7: 104533. DOI: 10.1016/j.respol.2022.104533.
- Dosi G., Piva M., Virgillito M.E., Vivarelli M. (2021). Embodied and disembodied technological change: The sectoral patterns of job-creation and job-destruction. *Research Policy*, 50, 4: 104199. DOI: 10.1016/j.respol.2021.104199.
- Eubanks V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. New York: Picador.
- Felten E.W., Raj M., Seamans R. (2019). The occupational impact of artificial intelligence: Labor, skills, and polarization. *SSRN Working Paper*. DOI: 10.2139/ssrn.3368605.

- Financial Times (2025). The great AI jobs disruption is under way. *Financial Times*, 23 May. Retrieved from: <https://www.ft.com/content/cb9ea970-e6de-4daf-aa9e-7a48d5e648c3> (23/12/2025).
- Floridi L. (2023). On the Brussels–Washington consensus about the legal definition of artificial intelligence. *Philosophy & Technology*, 36, 4: 1. DOI: 10.1007/s13347-023-00690-z.
- Frey C.B., Osborne M.A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114: 254. DOI: 10.1016/j.techfore.2016.08.019.
- Graetz G., Michaels G. (2018). Robots at work. *Review of Economics and Statistics*, 100, 5: 753. DOI: 10.1162/rest_a_00754.
- Grashof N., Kopka A. (2023). Artificial intelligence and radical innovation: An opportunity for all companies? *Small Business Economics*, 61, 2: 771. DOI: 10.1007/s11187-022-00698-3.
- Guarascio D., Piccirillo A., Reljic J. (2024). Will robots replace workers? Assessing the impact of robots on employment and wages with meta-analysis. *SSRN Working Paper*. DOI: 10.2139/ssrn.4750618.
- Hatzius J., Briggs J., Kodnani D., Pierdomenico G. (2023). *The Potentially Large Effects of Artificial Intelligence on Economic Growth*. Goldman Sachs Global Economics Paper, 26 March. Retrieved from: <https://www.gspublishing.com/content/research/en/reports/2023/03/27/d64e052b-0f6e-45d7-967b-d7be35fabd16.pdf> (23/12/2025).
- Hui X., Sundararajan A., Tang Z. (2024). The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market. *SSRN Working Paper*. DOI: 10.2139/ssrn.4760243.
- IMF - International Monetary Fund (2024a). Gen-AI: Artificial intelligence and the future of work. IMF Staff Discussion Note, SDN/2024/001. Washington DC: IMF. Retrieved from: <https://www.imf.org/-/media/files/publications/sdn/2024/english/sdnea2024001.pdf> (23/12/2025).
- IMF - International Monetary Fund (2024b). AI will transform the global economy. Let's

make sure it benefits humanity. *IMF Blog*, 14 January. Retrieved from: <https://www.imf.org/en/blogs/articles/2024/01/14/ai-will-transform-the-global-economy-lets-make-sure-it-benefits-humanity> (23/12/2025).

Keynes J.M. (1931). *Essays in persuasion*. London: Macmillan.

McKinsey & Company (2025). *Superagency in the workplace: Empowering people to unlock AI's full potential*. New York: McKinsey Global Institute. Retrieved from: <https://www.mckinsey.com/~/media/mckinsey/business%20functions/quantumblack/our%20insights/superagency%20in%20the%20workplace%20empowering%20people%20to%20unlock%20ais%20full%20potential%20at%20work/superagency-in-the-workplace-empowering-people-to-unlock-ais-full-potential-v4.pdf> (23/12/2025).

Montobbio F., Staccioli J., Virgillito M.E., Vivarelli M. (2024a). Labor-saving automation: A direct measure of occupational exposure. *World Economy*, 47, 1: 332. DOI: 10.1111/twec.13522.

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

Noy S., Zhang W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381, 6654: 187. DOI: 10.1126/science.adh2586.

OECD (2023). *OECD Employment Outlook 2023*. Paris: OECD Publishing.

OECD (2024). AI principles overview. Paris: OECD Publishing.

Pouliakas K., Santangelo G.D., Dupire M. (2025). The anatomy of an AI wage premium: Task content, autonomy and job design among European developers. *Eurasian Business Review*, 15, 2: 245. DOI: 10.1007/s40821-025-00302-0.

Ricardo D. (1821). *On the principles of political economy and taxation*. London: John Murray.

Rothstein H.R., Sutton A.J., Borenstein M., eds. (2005). *Publication bias in meta-analysis: Prevention, assessment and adjustments*. Chichester: Wiley.

Schumpeter J.A. (1911). *Theorie der wirtschaftlichen Entwicklung*. Leipzig: Duncker & Humblot.

Stanley T.D., Doucouliagos H. (2012). *Meta-regression analysis in economics and business*.

London: Routledge.

- Stephany F., Teutloff O. (2024). What is the price of a skill? The value of complementarity. *Research Policy*, 53, 1: 104898. DOI: 10.1016/j.respol.2023.104898.
- Susskind D. (2020). *A world without work: Technology, automation, and how we should respond*. London: Penguin.
- Teutloff O., Einsiedler J., Kässi O., Braesemann F., Mishkin P., del Rio-Chanona R.M. (2025). Winners and losers of generative AI: Early Evidence of Shifts in Freelancer Demand. *Journal of Economic Behavior & Organization*, 235: 106845. DOI: 10.1016/j.jebo.2024.106845.
- Wang W., Gao G., Agarwal R. (2024). Friend or foe? Teaming between artificial intelligence and workers with variation in experience. *Management Science*, 70, 9: 5753. DOI: 10.1287/mnsc.2021.00588.
- Wu Y., Lin Z., Zhang Q., Wang W. (2024). Artificial intelligence, wage dynamics, and inequality: Empirical evidence from Chinese listed firms. *International Review of Economics and Finance*, 96: 103739. DOI: 10.1016/j.iref.2024.103739.
- Yang C.H. (2022). How artificial intelligence technology affects productivity and employment: Firm-level evidence from Taiwan. *Research Policy*, 51, 6: 104536. DOI: 10.1016/j.respol.2022.104536.
- Yuan Y., Sun Y., Chen H. (2025). Does artificial intelligence affect firms' inner wage gap? Evidence from China's manufacturing sector. *Applied Economics*, 57, 19: 2365. DOI: 10.1080/00036846.2024.2324090.