UNIVERSITÀ DEGLI STUDI
DI MODENA E REGGIO EMILIA

Dottorato di ricerca in Lavoro, sviluppo e innovazione

Ciclo XXXVIII

Artificial Intelligence and changes in labour market structure

Candidato: Giuseppe Molinari

Relatore (Tutor): Prof.ssa Chiara Strozzi

Correlatore (Co-Tutor): Prof. Michele Cantarella

Coordinatore del Corso di Dottorato: Prof.ssa Ylenia Curzi

Declaration

Chapter 1 was co-authored with Chiara Strozzi. Chapters 2 and 3 were co-authored with Michele Cantarella and Chiara Strozzi.

Abstract

It is now widely recognized that artificial intelligence can reshape work, the economy, and their future trajectories. While public debates have mainly focused on whether AI substitutes for human labour, its effects extend beyond, influencing labour market structures, working conditions, the distribution of economic power between firms and workers, and required abilities and skills. This PhD thesis adopts this broader perspective and analyses the impact of artificial intelligence on labour markets in Europe and the United States along three complementary dimensions. First, a structural "objective" dimension, concerning changes in occupational composition and employment shares associated with exposure to AI. Second, an "objective" institutional and distributive dimension, focusing on transformations in firms' market power on the labour demand side, measured through monopsony. Third, a "subjective" dimension, centred on changes in job content—namely abilities, skills, and knowledge—and how these evolve with advances in AI. Each chapter addresses one of these dimensions through a distinct yet conceptually integrated empirical analysis. The first chapter investigates the relationship between occupational exposure to AI and changes in the structure of the European labour market across 25 countries from 2011 to 2021, using AI exposure measures from the literature and data from the European Labour Force Survey (LFS-EU). Results show that employment shares increased in occupations more exposed to AI, suggesting a complementary rather than substitutive effect of technology during the period considered. Stratification by education reveals employment growth in low- and high-skill occupations and a decline in medium-skill jobs, consistent with theories of Routine Biased Technological Change and job polarization. Younger workers experience relative employment gains, potentially reflecting greater adaptability to AI-related skills. Country-level analyses highlight substantial heterogeneity, pointing to the role of national institutions, education systems, and the pace of AI diffusion. The second chapter, co-authored with Michele Cantarella and Chiara Strozzi, examines the relationship between monopsony power and occupational exposure to AI in 26 European countries between 2011 and 2020. Using LFS data on wages and employment and Google Patents to construct

an AI exposure index, monopsony power is measured through the wage elasticity of labour supply. Results show a marked decline in labour supply elasticity over time, indicating increasing monopsony power. This trend appears largely independent of AI exposure, which plays a limited role. Stratified analyses indicate that monopsony power is strongest in low-wage occupations, followed by high-wage ones, while middle-wage occupations exhibit weaker monopsony. A rising trend is also observed in low-education occupations. The third chapter, co-authored with Michele Cantarella and Chiara Strozzi, focuses on how AI reshapes the subjective content of work rather than task structures. Using longitudinal O*NET data from 2011 to 2025 and two novel AI exposure measures—based on AI research topics and LLM benchmark performance—we construct exposure indices at the occupation–requirement–year level. Results show that within occupations, AI exposure is positively associated with increases in importance-weighted requirement levels across Abilities, Skills, and Knowledge. At the occupation level, however, higher AI exposure is associated with a decline in overall requirement levels, driven mainly by reductions in Abilities. This divergence suggests that AI strengthens specific complementary requirements while contributing to an erosion of average occupational requirements.

Italian translation

È ormai ampiamente riconosciuto che l'Intelligenza Artificiale rappresenti una tecnologia capace di rimodellare il lavoro e di influenzare l'economia e le sue traiettorie future. Mentre il dibattito pubblico tende spesso a concentrarsi su come e in quale misura l'IA possa sostituire il lavoro umano, i suoi effetti vanno ben oltre, influenzando la struttura del mercato del lavoro, le condizioni di lavoro, la distribuzione del potere economico tra imprese e lavoratori e le abilità e skill richieste ai lavoratori.

Questa tesi di dottorato si inserisce in tale dibattito analizzando l'impatto dell'Intelligenza Artificiale sul mercato del lavoro in Europa e negli Stati Uniti lungo tre dimensioni complementari. La prima riguarda i cambiamenti nella composizione occupazionale associati all'esposizione delle occupazioni all'IA. La seconda è una dimensione istituzionale e distributiva, relativa alle trasformazioni del potere di monopsonio. La terza è una dimensione soggettiva, incentrata sui cambiamenti nel contenuto del lavoro — in termini di abilità, skill e conoscenze — e sul modo in cui questi evolvono con i progressi dell'IA.

Il primo capitolo analizza la relazione tra l'esposizione occupazionale all'IA e i cambiamenti nella struttura del mercato del lavoro europeo in 25 paesi nel periodo 2011–2021, utilizzando dati della Labour Force Survey europea (LFS-EU). I risultati mostrano che le quote occupazionali sono aumentate nelle occupazioni più esposte all'IA, suggerendo un effetto complementare, piuttosto che sostitutivo, della tecnologia. La stratificazione per livello di istruzione evidenzia una crescita dell'occupazione nelle professioni a bassa e alta qualificazione e una riduzione in quelle a qualificazione intermedia, in linea con le teorie della job polarization. I lavoratori più giovani registrano una crescita relativa dell'occupazione, probabilmente grazie a una maggiore adattabilità nell'acquisizione di competenze compatibili con l'IA. Le analisi a livello nazionale mettono inoltre in luce una forte eterogeneità tra paesi.

Il secondo capitolo, scritto in collaborazione con Michele Cantarella e Chiara Strozzi, esamina la relazione tra il potere di monopsonio e l'esposizione occupazionale all'IA in 26 paesi europei tra il 2011 e il 2020. Utilizzando dati LFS su salari e occupazione e dati di Google Patents per costruire un indice di esposizione all'IA, il potere di monopsonio viene misurato attraverso l'elasticità dell'offerta di lavoro rispetto al salario. I risultati mostrano una marcata diminuzione di tale elasticità nel tempo, indicando un aumento del potere di monopsonio. Questo andamento risulta in larga parte indipendente dall'esposizione all'IA, che svolge un ruolo limitato. Le analisi stratificate indicano che il potere di monopsonio è più forte nelle occupazioni a basso salario, seguito da quelle ad alto salario, mentre le occupazioni a salario medio presentano un monopsonio più debole. Si osserva inoltre un trend crescente di monopsonio nelle occupazioni a basso livello di istruzione.

Il terzo capitolo, scritto in collaborazione con Michele Cantarella e Chiara Strozzi, si concentra su come l'IA stia rimodellando il contenuto del lavoro piuttosto che le strutture delle mansioni. Utilizzando dati longitudinali O*NET dal 2011 al 2025 e due nuove misure di esposizione all'IA — basate sui temi della ricerca in IA e sulle prestazioni dei LLM nei benchmark — costruiamo indici di esposizione a livello occupazione–abilità/skill/conoscenze–anno. I risultati mostrano che, all'interno delle occupazioni, l'esposizione all'IA è positivamente associata a un aumento dei livelli dei requisiti, ponderati per importanza, nelle dimensioni di abilità, skill e conoscenze. A livello occupazionale, tuttavia, una maggiore esposizione all'IA è associata a una riduzione dei livelli complessivi delle caratteristiche richieste, trainata principalmente da una diminuzione del livello medio di abilità. Questa divergenza suggerisce che l'IA rafforza requisiti specifici e complementari, contribuendo al contempo a un'erosione delle abilità, skill e conoscenze medie delle occupazioni.

Quando sarà, felicità

Contents

1 Artificial Intelligence and Labour Market Dynamics in Europe **8**

1.1 Introduction 9

1.2 Literature review 10

1.3 Data and methodology 13

1.3.1 Labour market data 13

1.3.2 AI exposure 13

1.4 Results 16

1.4.1 Descriptive analysis 16

1.4.2 Empirical analysis 20

1.4.3 Pooled results 20

1.4.4 Results by country 24

1.5 Conclusion 26

Appendices **31**

1.A AI Scores for Occupations at the Three-Digit Level 31

1.B Change in employment vs. exposure to AI 36

2 Still too early? Artificial Intelligence and Monopsony Power **42**

2.1 Introduction 43

2.2 Literature review 45

2.2.1 Monopsony and power market 45

2.2.2 AI measurement 47

2.3 Data 50

2.3.1 Labour market data 50

2.3.2 AI exposure 51

2.3.3 Construction of our database 51

2.4 Empirical methodology 52

2.5 Descriptive Evidence 55

2.6 Results 59

2.7 Conclusions 68
2.8 Acknowledgements 69

Appendices **77**
2.A Derivation of equation 2.4 77
2.B Descriptive evidence 81
2.B.1 Annual evolution of AI exposure by profession 81
2.B.2 Yearly labour supply elasticities for the nine major groups of ISCO-08 82

3 Workers Without Quality? Enrichment or Impoverishment of Abilities, Knowledge, and Skills in the Age of Artificial Intelligence **91**
3.1 Introduction 92
3.2 Literature review 95
3.3 Data 98
3.3.1 Worker requirements data 99
3.3.2 AI scores 99
3.3.3 Labour market data 100
3.4 Method 100
3.4.1 Worker requirements 100
3.4.2 Data Integration and Harmonization 106
3.5 Empirical Strategy and Preliminary Evidence 106
3.5.1 Descriptive Evidence and Stylized Facts 108
3.6 Results 117
3.6.1 Results by occupation–worker requirements level 117
3.6.2 Results at the occupational level 120
3.7 Conclusion 121

Appendices **128**
3.A Worker Requirements Described by O*NET 128
3.B Benchmarks Used in the AI Exposure Measure 131
3.C Examples of Boolean Queries Used to Search AI-Related Publications 133
3.D Top 10 Semantic Matches Between Benchmarks and Worker Requirements 136
3.E Top Five and Bottom Five Worker Requirements by Exposure to the AI Topics Index 139
3.F Top Five and Bottom Five Worker Requirements by Exposure to the AI LLM Index 141

Chapter 1

Artificial Intelligence and Labour Market Dynamics in Europe

Abstract

In this paper, we explore the relationship between occupational exposure to Artificial Intelligence (AI) and changes in employment composition and occupational employment shares in the European labour market, using data for 25 European countries over the period 2011–2021. We combine AI exposure scores from the existing literature with microdata from the European Labour Force Survey (LFS-EU). Our findings indicate that employment shares have increased in occupations more exposed to AI, suggesting that, during the period considered, technological change has had a complementary rather than substitutive effect on human labour.

When dividing observations by education terciles, we find employment growth in both low-skill and high-skill occupations, alongside a decline in medium-skill jobs. This pattern is consistent with theories of Routine-Biased Technological Change and job polarization. Analyses by age terciles show that younger workers have experienced a relative increase in employment shares, possibly reflecting their greater adaptability in acquiring AI-compatible skills.

Finally, country-level analyses reveal substantial heterogeneity in the magnitude and composition of these effects. This variation suggests that differences in the diffusion of AI technologies, educational attainment, and national labour market institutions and policies play a crucial role in shaping employment dynamics across European countries.

Keywords: *Artificial Intelligence, Occupational Exposure, Employment Dynamics, Labour Market Polarization, Routine-Biased Technological Change, Skills and Education, European Labour Markets*

JEL codes: O33, O30, J21, J24, J23

1.1 Introduction

The relationship between technological change and transformations in the labour market has long been a central topic in economic debate. Each major technological leap has prompted scholars to reflect on and reassess its implications for the economy and employment. Interest in this issue extends well beyond academia, becoming a matter of public debate and raising recurring concerns such as: "Is the future of work at risk?" Similar questions have emerged in every phase of technological change. Even historical episodes such as the Luddite movement — whose members are often portrayed as "machine breakers" — have been interpreted by technology historian David Noble as early manifestations of a latent conflict between technological progress and labour.

For years, the debate has opposed those who foresee a "jobless future" to those who argue that digitalization — here understood as the diffusion of multiple technological families, ranging from robotics and Industry 4.0 to Artificial Intelligence and the Internet of Things — will instead usher in a new "golden age" of work (Nübler, 2016). This debate is particularly salient today, in the era of Artificial Intelligence, which is often presented as a transformative force capable of redefining the economy, work, and even modes of thinking (OECD, 2024). One feature that distinguishes AI from previous technological waves is the speed at which it can learn from errors and improve performance over short periods of time. This raises a natural question: will AI soon be able to perform not only routine tasks, but also those requiring creativity and higher-order cognitive skills? In other words, are we approaching the scenario imagined in twentieth-century science fiction, such as that envisioned by Philip K. Dick?

Despite its relevance, this topic must be approached with caution. As Acemoglu (2021) emphasizes, defining what Artificial Intelligence is capable of doing is fundamentally different from predicting what it will ultimately achieve. As noted by Autor et al. (2022), past technological revolutions have tended to transform the nature of work rather than produce a persistent decline in overall employment. Nevertheless, Artificial Intelligence is already affecting labour demand and supply in ways that differ from earlier technologies. In particular, AI is capable of performing cognitive tasks (Felten et al., 2018; Tolan et al., 2021), potentially affecting even high-skill occupations that were previously considered relatively insulated from technological substitution. Importantly, such influence does not necessarily imply displacement. As this paper will show, AI may also complement human labour, leading to employment growth in certain occupations. At the same time, the scale and scope of AI-driven innovation call for a reassessment of prevailing trends in employment composition and occupational dynamics.

Given that artificial intelligence can perform tasks requiring cognitive skills — and may therefore affect high-skill jobs — does its diffusion alter existing labour market trends? Or does it instead reinforce complementarities that favour these occupations? How does the adoption of AI interact with the process of job polarization that has been identified as a defining feature of employment dynamics in Europe? These are the questions this paper seeks to address.

The remainder of the paper is organized as follows. Section 2 reviews the relevant literature, Section 3 describes the data and methodology, Section 4 presents the empirical results, and Section 5 concludes. The Appendix reports additional information on AI exposure scores and regression results.

1.2 Literature review

It is now well established in the literature that Artificial Intelligence represents a distinct advancement compared to previous families of technologies. AI qualifies as a general-purpose technology with potential applications across a wide range of sectors (OECD, 2019). As highlighted by Abrardi et al. (2022), the ability of machine-learning-based systems to acquire knowledge with varying degrees of autonomy distinguishes AI from earlier digital technologies. Similarly, Agrawal et al. (2019) emphasize that AI can generate predictions from unstructured data and increasingly perform tasks that, until recently, were the exclusive domain of human labour. According to Brynjolfsson et al. (2017), Artificial Intelligence is specifically designed to self-improve through machine learning mechanisms, further amplifying its transformative potential. More recently, attention has increasingly focused on generative AI, including natural language processing tools such as ChatGPT and other and other generative AI tools. These technologies are capable of producing text, images, videos, music, and other forms of content in response to specific prompts, with potentially far-reaching implications for work and society (Baldassarre et al., 2023; Zarifhonarvar, 2023).

These developments have led several authors to question whether Artificial

Intelligence could permanently reduce labour demand. Addressing this issue requires recognizing that AI is not solely a substitute for human labour but can also act as a complement, reshaping the task content of occupations. Moreover, AI not only enables the automation of existing tasks but also contributes to the creation of new ones (Gmyrek et al., 2023). This perspective is formalized by Acemoglu and Restrepo (2018, 2019), who identify three channels through which AI affects labour demand: a displacement effect, whereby labour demand declines as tasks are automated; a productivity effect, whereby automation lowers costs and increases demand for complementary non-automated tasks; and a reinstatement effect, whereby new tasks and occupations emerge, particularly benefiting workers whose skills complement AI. The net impact on employment depends on the relative strength of these mechanisms, even though, as shown by Acemoglu (2024), productivity gains may also exacerbate inequality. Empirical evidence to date suggests that the aggregate employment effects of AI remain limited. For instance, Acemoglu et al. (2022), using data on online vacancies, find no statistically significant effects of AI adoption on employment in the United States between 2010 and 2017 at either the industry or occupational level. Similarly, Felten et al. (2019), introducing the AIOI index, do not detect employment effects of AI in the U.S. over the period 2010–2016.

Comparable findings emerge for Europe. Georgieff and Hyee (2021) report positive but statistically insignificant effects of AI exposure on aggregate employment across a sample of OECD countries, exploiting variation in occupational exposure. Using firm-level survey data from the financial and manufacturing sectors, Lane et al. (2023) also conclude that AI adoption does not significantly affect employment. Moreover, qualitative evidence from Milanez (2023), based on 100 case studies across eight OECD countries, suggests that job reorganization is currently more prevalent than outright job displacement.

According to OECD (2023) and Lane and Saint-Martin (2021), the absence of strong aggregate employment effects can be attributed to several factors: the high costs of AI adoption, firms' gradual adjustment of employment levels, the still limited role of AI within overall automation processes, and the difficulty of capturing task creation through exposure-based measures.

Despite the ambiguity surrounding aggregate employment effects, it is clear that Artificial Intelligence substantially expands the set of tasks exposed to technological change. A key novelty of AI lies in its rapid progress in performing non-routine cognitive tasks, which are predominantly associated with high-skill occupations. In constructing an index of occupational exposure to AI, Tolan et al. (2021) identify fourteen cognitive skills that fall within AI's expanding capabilities, ranging from memory and sensorimotor interaction to comprehension, communication, and emotional self-regulation. Their analysis combines task information from datasets such as PIAAC, O*NET, and EWCS with indicators of technological progress drawn from sources including Papers with Code. A similar focus on cognitive skills characterizes the work of Felten et al. (2021, 2018, 2019), who map advances in AI applications—such as image recognition, translation, and strategic reasoning—onto occupational skill requirements using O*NET data. Their findings highlight the speed and breadth of recent AI progress. Likewise, Lassébie and Quintini (2022) examine AI's ability to perform cognitive activities such as advising, scheduling, and communication, estimating that 28% of employment across OECD countries is concentrated in occupations facing high automation risk, many of which are classified as high-skill.

These findings suggest that high-skill workers are among the most exposed to Artificial Intelligence. Crucially, however, exposure should not be equated with substitution. As emphasized by Pizzinelli et al. (2023), AI exposure may also be associated with positive employment effects when technology complements human labour. Consistent with this view, Albanesi et al. (2025), using Felten's exposure scores for a sample of European countries, show that highly exposed occupations—such as mathematicians, finance professionals, software developers, scientists, and senior officials—experienced employment growth between 2011 and 2019. This evidence suggests that, at least in the current phase, many high-skill occupations are complementary to AI.

Overall, the impact of artificial intelligence on labour demand remains an open question. These considerations are embedded in a broader debate on how technological change interacts with the distribution of skills and tasks in the labour market. Within the Skill-Biased Technological Change (SBTC) framework, new technologies disproportionately benefit high-skill workers due to complementarities between advanced skills and ICT, while low-skilled workers tend to be substituted, generating a "skill bias" in the evolution of labour demand. Seminal contributions by Katz and Murphy (1992) and Autor et al. (1998) document this pattern for the United States, motivating the SBTC paradigm. However, the inability of this framework to fully account for observed employment and wage dynamics across countries and periods has led to alternative explanations.

In contrast, the Routine-Biased Technological Change (RBTC) hypothesis emphasizes the automation of routine tasks, which are disproportionately performed in middle-skill occupations. As ICT capital becomes cheaper, machines increasingly substitute for routine labour, reducing demand in the middle of the wage distribution.

Influential contributions by Autor et al. (2006); Goos and Manning (2007); Goos et al. (2014), building on the task-based framework of Autor et al. (2003), document the emergence of job polarization since the 1990s.

Given the distinctive task capabilities introduced by Artificial Intelligence, a reassessment of employment dynamics through the lens of occupational exposure is warranted. Building on the seminal contribution of Albanesi et al. (2025), this paper aims to contribute to this literature by examining how AI exposure relates to changes in employment shares and occupational composition in Europe.

1.3 Data and methodology

Our analysis examines the relationship between changes in employment shares between 2011 and 2021 and occupational exposure to artificial intelligence. It draws on two main data sources, described in detail below: the European Labour Force Survey and a set of widely used AI exposure scores drawn from the existing literature.

1.3.1 Labour market data

We utilize the Labour Force Survey (EU-LFS) annual microdata for the period between 2011 and 2021. The LFS is a large household sample survey providing quarterly results on labour participation of individuals aged 15 and over and on those outside the labour force, and it includes data for all countries within the European Union. Specifically, for our study, we use data for 25 European countries. We employ the LFS to calculate the change in employment share over the selected period. Additionally, we use data on educational level and age provided by the database to further specify our examination. Another data point we use is related to the economic sector of employment. Unlike Albanesi et al. (2025), who reference only 6 sectors, our research is conducted across all NACE sectors present in the sample. Each analysis is conducted at the occupational-sectoral level using the ISCO 3-digit classification.

1.3.2 AI exposure

Following a review of the literature, we select three established measures of occupational exposure to artificial intelligence to study its impact on employment: the JRC AI score, the Felten AI score, and the Webb AI score. The first measure is developed by the European Commission's Joint Research Centre (JRC) in Seville and is based on the work of Tolan et al. (2021). The authors propose a comprehensive framework that explicitly focuses on Artificial Intelligence, distinguishing it from earlier approaches to robotization, digitalization, and general technological change. Unlike technologies that do not rely on AI (such as self-checkout machines), this framework links occupational tasks to underlying cognitive abilities, which are then mapped to performance indicators across different AI domains.

The framework draws on data from the European Working Conditions Survey (EWCS), the Survey of Adult Skills (PIAAC), and the O*NET occupational database. In particular, 59 generic tasks are mapped to 14 cognitive abilities identified in the cognitive science literature. These abilities are subsequently linked to 328 AI evaluation tasks derived from benchmarking initiatives, challenges, competitions, and scientific studies. To capture the intensity of current research and development in different AI techniques, the authors construct a research intensity indicator, which serves as a proxy for the short- and medium-term potential impact of AI by identifying domains where research efforts are most concentrated. Compared to earlier task-based approaches, such as Autor et al. (2003), this framework introduces an intermediate layer of cognitive abilities between tasks and AI technologies, allowing for a more granular assessment of occupational exposure. This structure enables the identification of tasks and cognitive abilities that are more or less likely to be affected by AI, as well as the construction of a single occupation-level AI exposure score. Importantly, the framework does not measure past technological progress directly, but rather captures future-oriented trends by exploiting the distribution of benchmarks across AI domains. This approach allows for regular updates as new benchmarks emerge and provides a forward-looking measure of AI exposure. By focusing on cognitive abilities rather than skills, the JRC score offers a detailed perspective on how Artificial Intelligence may transform work tasks across occupations. The resulting AI exposure score is standardized.

Another measure employed in this study is developed by Felten et al. (2021), building on their earlier contributions (Felten et al., 2018, 2019). The authors introduce a measure of occupational exposure to Artificial Intelligence, termed the AI Occupational Exposure (AIOE). This measure links widely used and well-developed AI applications to workplace abilities and occupations. Its relevance is illustrated through several extensions, including measures of AI Industry Exposure (AIIE) and AI Geographic Exposure (AIGE), as well as firm-level indicators of AI exposure. The AIOE is empirically validated, and its potential applications in strategy, management, and innovation are discussed. The underlying datasets are made publicly available for researchers, policymakers, and practitioners.

The construction of the AIOE relies on mapping common AI applications to occupational abilities using a crowdsourced dataset. Specifically, the authors identify a set of well-established AI applications based on categories defined by the Electronic Frontier Foundation (EFF) AI Progress Measurement project and link these applications to data on occupational abilities from the Occupational Information Network (O*NET), developed by the U.S. Department of Labor. Exposure is first measured at the ability level and then aggregated to the occupation level, yielding a measure of potential exposure to Artificial Intelligence that remains agnostic as to whether AI substitutes for or complements human labour.

Conceptually, the framework distinguishes between "general" Artificial Intelligence—referring to fully autonomous systems, which do not yet exist—and "narrow" Artificial Intelligence, which encompasses machine-learning-based techniques for pattern recognition and prediction. The AIOE focuses on the latter by identifying specific applications of machine learning and linking them to the abilities required in each occupation. This approach isolates exposure to distinct AI functionalities and produces an aggregate measure of occupational exposure to Artificial Intelligence based on the ability composition of jobs.

Finally, the third measure employed in this study is proposed by Webb (2020). This approach measures occupational exposure to Artificial Intelligence by exploiting textual overlap between patent data and job task descriptions. Specifically, Webb analyzes verb–noun pairs extracted from patent titles in Google Patents Public Data and from task descriptions in the O*NET database.

To identify exposure to Artificial Intelligence, Webb first selects a subset of patents relevant to AI by searching for keywords such as "neural network" in patent titles and abstracts. From this set of patents, all titles are extracted and decomposed into verb–noun pairs (e.g., diagnose–disease, recognize–aircraft). The frequency of each pair, and of semantically similar pairs, is then calculated across the full corpus of AI-related patent titles.

The same procedure is applied to occupations. Each occupation is described by a set of tasks expressed in free-form text (e.g., "Interpret tests to diagnose patients' conditions"). Verb–noun pairs are extracted from these task descriptions, typically yielding between 20 and 40 pairs per occupation. Each pair is assigned a weight corresponding to the relative frequency of similar pairs in the AI patent corpus. Occupational exposure to Artificial Intelligence is then computed as the weighted average of these scores across all tasks, where weights reflect the importance of each task within the occupation.

While the three AI exposure measures differ in their methodological approaches, some commonalities emerge. Webb's and Felten's measures rely exclusively on task and ability data from O*NET, whereas the JRC measure by Tolan et al. (2021) additionally incorporates European data from PIAAC and the European Working Conditions Survey (EWCS). From a conceptual standpoint, both Felten's and Tolan's frameworks introduce an intermediate layer of cognitive abilities between tasks and AI advancements, whereas Webb's approach focuses directly on the alignment between machine-learning applications and occupational tasks through textual similarity.

Finally, the AI exposure measures are merged with labour market data. Both the Felten and Webb measures are originally defined using the occ1990d occupational classification developed by Autor and Dorn (2013). As no direct mapping exists between this classification and ISCO-08, occupations are first converted to the SOC classification following the procedure outlined by Dorn. Subsequently, SOC occupations are mapped to the ISCO-08 classification at the three-digit level using the crosswalks provided by International Labour Organization (2010).

1.4 Results

In this section, we provide descriptive statistics of the employment measures and results, first pooled and then by country.

1.4.1 Descriptive analysis

This section presents descriptive statistics for the three measures of Artificial Intelligence exposure across European countries in our sample.

Table 1.1 reports summary statistics for the three AI exposure indices described in the previous section, namely those developed by Tolan et al. (2021), Felten et al. (2021), and Webb (2020). The JRC measure by Tolan et al. is available for 119 occupations, whereas the Felten and Webb indices cover 126 occupations. To ensure comparability across measures, we restrict the sample to the 119 occupations common to all three indices, excluding the additional seven occupations covered only by the Felten and Webb measures.

The indices proposed by Tolan et al. and Felten et al. display similar distributions, with mean values of 0.50 and 0.485, respectively, and standard deviations of 0.292 and 0.272. In contrast, the Webb index exhibits a slightly higher mean (0.544) and a lower standard deviation (0.24). Maximum values are not reported, as all three measures are normalized and expressed in percentiles.

Table 1.1: **Descriptive Statistics of AI measures**

Variables	**Obs**	**Mean**	**Std. Dev.**
JRC	119	0.500	0.292
Felten	126	0.485	0.272
Webb	126	0.544	0.240

Notes: Descriptive statistics of AI measures across occupations (unweighted).

Table 1.2 reports Spearman rank correlations among the different AI exposure measures, which are useful for assessing the extent to which these indices capture similar dimensions of Artificial Intelligence exposure by construction.

As shown in the table, the correlation between the JRC measure and the Felten index is strong and statistically significant (with a p-value equal to zero, reported in parentheses). A positive but weaker correlation is also observed between the JRC index and the Webb index, which is likewise statistically significant. By contrast, the correlation between the Felten and Webb indices is not statistically significant, as indicated by a high p-value, and we therefore fail to reject the null hypothesis of no association between these two measures.

Table 1.2: **Spearman's rank correlation coefficients**

	(1)	**(2)**	**(3)**
(1) JRC	1.000		
(2) Felten	0.625	1.000	
	(0.00)		
(3) Webb	0.392	0.007	1.000
	(0.00)	(0.9438)	

Notes: Spearman's rank correlations. P-values are reported in parentheses below the correlation coefficients. The null hypothesis tests independence between variables.

Finally, to further illustrate the differences across the three AI exposure measures, we report tables listing the occupations with the highest and lowest values for each index.

As shown in these tables, there is substantial overlap between the JRC and Felten measures, both at the top and bottom of the distribution. More generally, the two approaches yield similar patterns: occupations with higher exposure values largely correspond to those classified as high-skill by the International Labour Organization (ILO), whereas occupations with lower exposure values are predominantly low-skill.

In contrast, Webb's measure exhibits a markedly different pattern. For example, the occupation "University and higher education teachers" appears among the least exposed occupations according to Webb, while it is classified among the most exposed in Felten's measure. This divergence highlights a fundamental difference in index construction, reflecting the distinct methodological approaches used to link Artificial Intelligence to occupational tasks. As such, these differences should be taken into account when interpreting empirical results based on alternative AI exposure measures.

Table 1.3: **AI scores of bottom five occupations (JRC)**

ISCO-08	**Occupation**	**JRC**
952	Street vendors (excluding food)	0.000
912	Vehicle, window, laundry and other hand cleaning workers	0.009
941	Food preparation assistants	0.017
911	Domestic, hotel and office cleaners and helpers	0.026
513	Waiters and bartenders	0.034

Notes: Bottom five occupations ranked by JRC AI exposure score according to the ISCO-08 classification.

Table 1.4: **AI scores of top five occupations (JRC)**

ISCO-08	**Occupation**	**JRC AI score**
215	Electrotechnology engineers	1.000
252	Database and network professionals	0.992
251	Software and applications developers and analysts	0.983
214	Engineering professionals (excluding electrotechnology)	0.975
212	Mathematicians, actuaries and statisticians	0.966

Notes: Top five occupations ranked by JRC AI exposure score according to the ISCO-08 classification.

Table 1.5: **AI scores of bottom five occupations (Felten)**

ISCO-08	Occupation	Felten AI score
931	Mining and construction labourers	0.000
911	Domestic, hotel and office cleaners and helpers	0.012
912	Vehicle, window, laundry and other hand cleaning workers	0.013
713	Painters, building structure cleaners and related trades	0.030
932	Manufacturing labourers	0.045

Notes: Bottom five occupations ranked by Felten AI exposure score according to the ISCO-08 classification.

Table 1.6: **AI scores of top five occupations (Felten)**

ISCO-08	Occupation	Felten AI score
212	Mathematicians, actuaries and statisticians	1.000
241	Finance professionals	0.980
261	Legal professionals	0.979
242	Administration professionals	0.896
231	University and higher education teachers	0.888

Notes: Top five occupations ranked by Felten AI exposure score according to the ISCO-08 classification.

Table 1.7: **AI scores of bottom five occupations (Webb)**

ISCO-08	Occupation	Webb AI score
231	University and higher education teachers	0.000
513	Waiters and bartenders	0.056
412	Secretaries (general)	0.082
524	Other sales workers	0.105
521	Street and market salespersons	0.108

Notes: Bottom five occupations ranked by Webb AI exposure score according to the ISCO-08 classification.

Table 1.8: **AI scores of top five occupations (Webb)**

ISCO-08	Occupation	Webb AI score
612	Animal producers	1.000
613	Mixed crop and animal producers	0.990
211	Physical and earth science professionals	0.931
312	Mining, manufacturing and construction supervisors	0.912
611	Market gardeners and crop growers	0.901

Notes: Top five occupations ranked by Webb AI exposure score according to the ISCO-08 classification.

1.4.2 Empirical analysis

We follow the model and analysis procedure outlined in Albanesi et al. (2025) to calculate the impact of Artificial Intelligence measures on the change in employment share. Unfortunately, the 2021 LFS income data are not available; therefore, we will focus solely on employment variations. We estimate an equation of this type:

$$y^s_{o,c} = \alpha_c + \alpha^s + \beta_c X^s_{o,c} + \varepsilon^s_{o,c} \tag{1.1}$$

where the dependent variable $y^s_{o,c}$ is the change in the employment share of sector occupation so in country c during the 2011-2021 period. The change in the employment share is measured as a percentage change relative to the midpoint of a cell's share of overall employment between 2011 and 2021, winsorised at the top and bottom 1%. $X^s_{o,c}$ are the measures of potential exposure of the sector-occupation to AI. All the measures assess how much the tasks that make up a given job can also be performed by artificial intelligence. They do not indicate whether AI is complementary to or a substitute for human labour by design. Therefore, we interpret the coefficient β_c as an indicator of whether AI substitutes for or complements human labour. . Specifically, a positive value of the coefficient indicates that more exposed occupations grow, suggesting that AI is complementary to human labour. Conversely, a negative value indicates that occupations decrease during the selected period, implying that AI substitutes for human labour.

1.4.3 Pooled results

We first examine the results for the pooled sample of 25 countries. The estimates are reported in Table 1.9. Overall, we find a positive association between changes in employment shares and occupational exposure to Artificial Intelligence. Occupations characterized by higher AI exposure tend to experience stronger employment growth, a result that holds across all three AI measures, although the coefficient associated with Webb's index is not statistically significant.

Using the JRC AI exposure measure, a shift from the 25th to the 50th percentile of the AI exposure distribution is associated with an increase in sector–occupation employment shares of 6.67%. When using Felten's measure, the same percentile shift corresponds to an increase of approximately 1%. These findings suggest that, at the current stage of adoption, Artificial Intelligence tends to complement rather than substitute exposed occupations, for instance by supporting decision-making processes or facilitating human–machine collaboration (Alekseeva et al., 2019).

While these results are consistent with much of the existing literature (Acemoglu et al., 2022; Lane et al., 2023), we further explore heterogeneity across occupations to assess whether high-skill jobs are particularly vulnerable to displacement, as suggested by Georgieff and Hyee (2021). To this end, we split sector–occupation cells within each country by education and age terciles, defined using the 2021 distribution—the final year of our sample and the closest to the present. Results by education terciles are reported in columns 2–4 of Table 8, while results by age terciles are shown in columns 5–7.

A further disaggregation by income terciles would provide additional insights; however, income information is not available in the 2021 Labour Force Survey. We therefore leave this analysis to future research.

As shown in Figure 1.1, from an educational perspective both the JRC and Felten indices yield results consistent with the job polarization literature. The impact of Artificial Intelligence appears to favour both low-skill and high-skill occupations, while failing to offset the decline in medium-skill jobs. At first glance, this finding may seem at odds with the view that Artificial Intelligence predominantly affects high-skill professions. However, a closer examination suggests that Artificial Intelligence—despite the lack of a shared definition even across the studies underlying the three AI measures—often operates in conjunction with other technological families (Lorenz et al., 2023), potentially reinforcing substitution processes in specific occupations.

As highlighted by OECD evidence, Artificial Intelligence tends to have an additive effect with respect to other technologies, such as industrial robots, which are widely adopted in manufacturing—a sector where employment is concentrated in the middle of the skill distribution. Moreover, our sample includes countries with relatively low levels of AI adoption (OECD, 2024), which may dampen the observable effects.

Table 1.9: **Change in employment vs. exposure to AI. Pooled sample (2011–2021)**

	Change in employment share						
	All	**Low Educ**	**Med Educ**	**High Educ**	**Young**	**Core**	**Old**
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(a) JRC	0.267***	0.738***	-0.510***	0.393***	0.405***	0.164**	0.004
	(0.061)	(0.129)	(0.119)	(0.078)	(0.068)	(0.055)	(0.072)
(b) Felten	0.0369***	1.055***	-0.378***	0.407***	0.514***	0.315***	0.125*
	(0.052)	(0.081)	(0.105)	(0.080)	(0.068)	(0.050)	(0.071)
(c) Webb	-0.021	-0.055	-0.083**	0.143*	-0.064	-0.073	-0.273**
	(0.069)	(0.152)	(0.139)	(0.084)	(0.103)	(0.089)	(0.117)

Notes: Linear regression estimates. Robust standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Each observation is an ISCO 3-digit occupation × sector cell and observations are weighted by the cell's average labour supply. The dependent variable is the within-country change in the employment share from 2011 to 2021, winsorized at the top and bottom 1 percent. The sample includes 25 European countries. Columns (2)–(4) report results for cells in the lower, middle, and upper tercile of the country's education distribution in 2021. Columns (5)–(7) report results for cells in the lower, middle, and upper tercile of the country's worker age distribution in 2021.

Finally, the analysis covers the period 2011–2021. As noted by Agrawal et al. (2019), major advances in machine-learning-based AI have occurred after 2021, suggesting that the timing of our sample may partly explain these results.

Overall, the evidence indicates that Artificial Intelligence currently complements both low-skill and high-skill occupations, while contributing to the displacement of routine, middle-skill jobs, thereby reinforcing the process of job polarization.

Turning to age heterogeneity, the estimated coefficient for the youngest tercile exceeds that of the pooled sample. As suggested by Green and Lamby (2023), if Artificial Intelligence adoption entails skill adjustments, younger workers may benefit disproportionately due to their greater willingness or ability to adapt their skill sets. Consistent with this interpretation, the estimated effects decline monotonically across age terciles, with the strongest impact observed among younger workers and progressively weaker effects for middle-aged and older workers. Notably, when using Webb's AI exposure measure, the estimated effect for the oldest tercile is negative, suggesting that workers in this group are more likely to be displaced by Artificial Intelligence.

a) by skill terciles

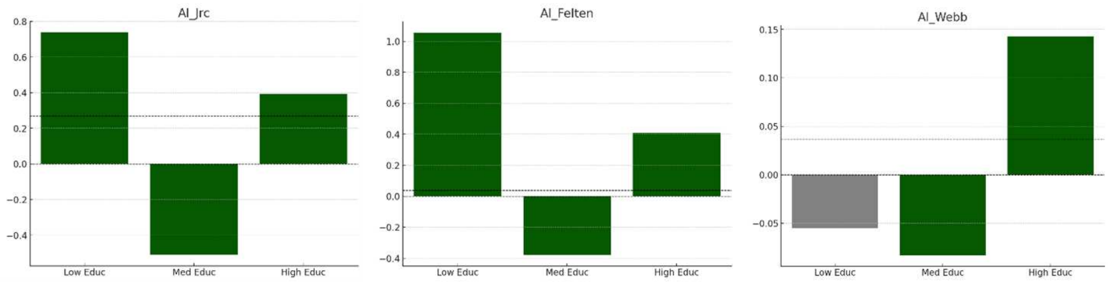

b) by age terciles

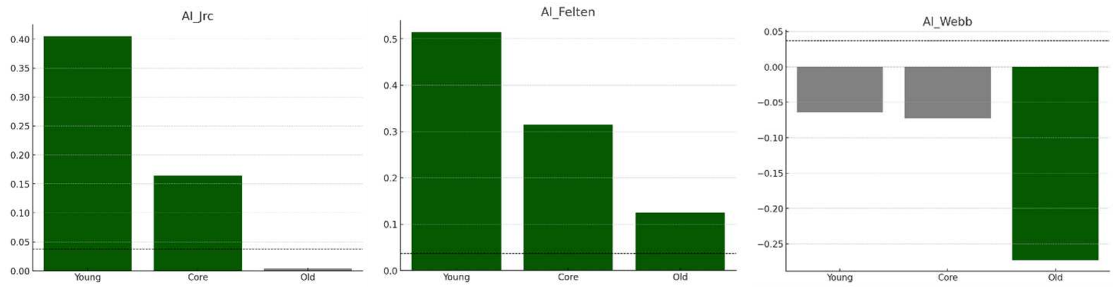

Figure 1.1: **Exposure to AI and changes in employment shares by skill terciles and age terciles**

Notes: Regression coefficients measuring the effect of exposure to artificial intelligence on changes in employment share (see Table 1.9). Each observation is an ISCO 3-digit occupation × sector cell. Observations are weighted by cells' average labour supply. Sample: 25 European countries, 2011–2021. The coefficient for the whole sample is displayed by the horizontal dotted line. Bars display the coefficient estimated for the subsample of cells whose average educational attainment is in the lower, middle, and upper tercile of the education distribution (top row) and whose workers' average age is in the lower, middle, and upper tercile of the workers' age distribution (bottom row). Coefficients statistically significant at least at the 10% level are plotted in green, otherwise in grey. CZ, NL, LU do not provide information regarding age; therefore, the sample differs between education and age.

1.4.4 Results by country

Turning to the country-level analysis, we observe substantial heterogeneity across countries for all three measures of Artificial Intelligence exposure. Focusing first on the JRC score, Figure 2 shows a wide dispersion of estimated coefficients, ranging from strongly positive values in countries such as Estonia and Luxembourg to negative values in Greece and Hungary, while the coefficient for Romania is not statistically significant. With few exceptions, the association between AI exposure and changes in employment shares remains positive across most countries. The pooled estimate lies approximately at the center of the distribution. For Italy, the coefficient is negative but not statistically significant, although it becomes significant once the sample is split by education terciles.

Similar patterns emerge when using Felten's measure. Again, there is considerable cross-country heterogeneity. Estonia, Luxembourg, and Cyprus exhibit the largest positive coefficients, whereas Greece, Hungary, Croatia, Norway, Italy, and Switzerland display negative values, although the coefficients for Hungary, Croatia, and Norway are not statistically significant. In the case of Italy, the coefficient is negative and statistically significant, suggesting that Italian employment dynamics diverge from those observed in most other countries. The pooled estimate remains positive and statistically significant and is located in the lower-middle part of the distribution.

When repeating the analysis using Webb's score, the degree of heterogeneity across countries is confirmed, but the number of statistically insignificant coefficients increases. Luxembourg and Estonia continue to display the highest coefficients, both exceeding one, while Greece, Norway, Hungary, and Croatia again exhibit negative values. For many countries, including Italy, the coefficients are not statistically significant.

To aid interpretation, we focus on a set of geographically and economically heterogeneous countries. Italy provides a particularly informative case. Using the JRC and Felten scores, Italy exhibits a negative coefficient, while the coefficient based on Webb's measure is slightly positive; however, only the Felten-based estimate is statistically significant. When the analysis is disaggregated by education terciles, all three measures yield consistent results that contrast with the pooled sample: employment growth is concentrated in medium-skill occupations, while both low- and high-skill occupations display negative coefficients. This pattern—albeit with some variation—is shared by other countries grouped under the "Southern Europe" category and may reflect the structural characteristics of economies that rank relatively low in terms of digital technology diffusion (Albanesi et al., 2025).

Considering France and Germany, both countries exhibit positive coefficients when using the JRC and Felten measures, with slightly larger effects under the latter. In contrast, when using Webb's score, the coefficient is negative for France and positive for Germany. Once again, disaggregation by education terciles confirms the job polarization hypothesis when using the first two measures, whereas Webb's score suggests a substitutive effect of Artificial Intelligence even for low-skill occupations.

Finally, similar patterns emerge among Scandinavian countries. Finland displays the highest coefficients, followed by Denmark and Sweden, while Norway shows negative coefficients in two out of three specifications. Interpretation in this case is complicated by the lack of statistical significance for several estimates, including those obtained after splitting the sample by education terciles.

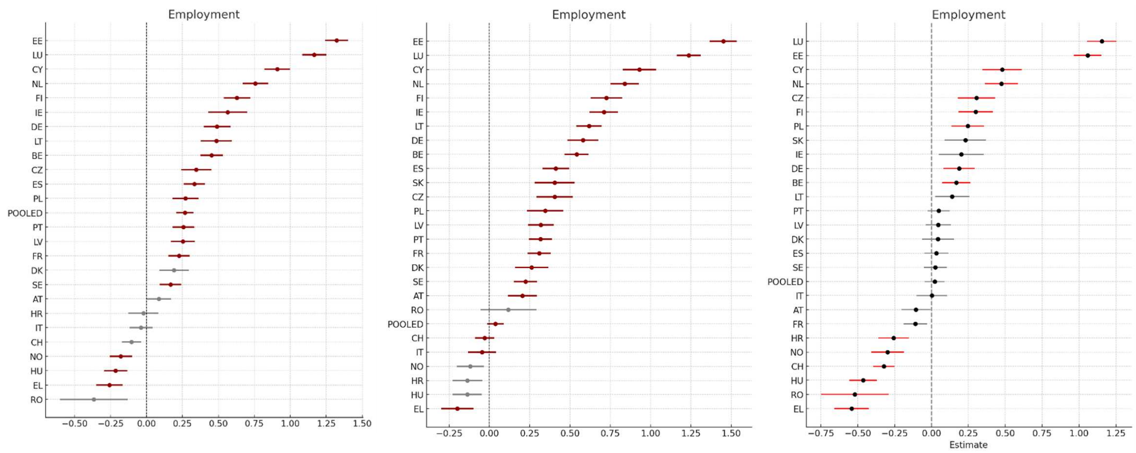

Figure 1.2: **Exposure to AI score and change in employment shares by countries.**

Notes: β_c coefficients from employment share regressions, shown jointly within each graph.

1.5 Conclusion

In this paper, we examine the impact of Artificial Intelligence on employment dynamics in a sample of 25 European countries, both at the aggregate and country level, over the period 2011–2021. To this end, we employ three occupational AI exposure measures drawn from the literature: the JRC score proposed by Tolan et al. (2021), which uniquely relies on European datasets in its construction; the measure developed by Felten et al. (2021); and the index introduced by Webb (2020). These measures capture the extent to which occupations are exposed to Artificial Intelligence in European labour markets. As emphasized throughout the paper, exposure should not be interpreted as synonymous with substitution, as it may also reflect complementarities between human labour and AI.

Our results indicate a positive relationship between employment growth and occupational AI exposure when using the JRC and Felten measures, whereas the association is slightly negative—but not statistically significant—when using Webb's index. When disaggregating by educational terciles, we find evidence that Artificial Intelligence has reinforced the process of job polarization in Europe during the period considered (Torrejón Pérez et al., 2023). Employment growth is observed at both the lower and upper ends of the educational distribution, while medium-education occupations experience a decline. One possible explanation is that Artificial Intelligence acts as an enabling technology for other technological families, such as industrial robots (OECD, 2024), thereby intensifying the displacement of routine tasks.

From an age perspective, the results suggest that Artificial Intelligence exhibits an "age-biased" pattern, favouring workers in the youngest age tercile. This finding is consistent with the idea that younger workers may be better positioned to adapt their skills in response to rapid technological change. Country-level analyses reveal substantial heterogeneity across Europe. For many countries, all three measures point to a positive association between AI exposure and employment changes. When further disaggregating by educational terciles, systematic patterns emerge across groups of countries sharing similar economic or geographical characteristics. In Southern European economies such as Italy, Spain, and Greece, the results diverge from the pooled sample, potentially reflecting lower levels of technological diffusion. In contrast, for core European economies such as Germany and France, the job polarization hypothesis is confirmed.

Overall, this heterogeneity highlights the importance of country-specific factors—such as the pace of Artificial Intelligence adoption, educational attainment, and national labour market institutions and policies—in shaping the employment effects of AI across Europe.

Finally, our findings should be interpreted with caution. Artificial Intelligence is a rapidly evolving technology, and during the period analyzed it was still in a relatively early stage of diffusion. As more recent data become available, future research will be able to assess whether the patterns identified in this study persist or evolve as AI technologies mature and become more widely adopted.

Bibliography

Abrardi, L., Cambini, C., and Rondi, L. (2022). Artificial Intelligence, firms and consumer behavior: A survey. *Journal of Economic Surveys*, 36(4):969–991.

Acemoglu, D. (2021). Harms of AI. CEPR Discussion Paper 16524, Centre for Economic Policy Research.

Acemoglu, D. (2024). The simple macroeconomics of AI. NBER Working Paper 32487, National Bureau of Economic Research.

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

Acemoglu, D. and Restrepo, P. (2018). Artificial Intelligence, automation and work. NBER Working Paper 24196, National Bureau of Economic Research.

Acemoglu, D. and Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30.

Agrawal, A., Gans, J. S., and Goldfarb, A. (2019). Artificial Intelligence: The ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives*, 33(2):31–50.

Albanesi, S., Dias da Silva, A., Jimeno, J. F., Lamo, A., and Wabitsch, A. (2025). New technologies and jobs in europe. *Economic Policy*, 40(121):71–139.

Alekseeva, L., Azar, J., Gine, M., Samila, S., and Taska, B. (2019). The demand for AI skills in the labour market. Ssrn scholarly paper, SSRN, Rochester, NY.

Autor, D., Chin, C., Salomons, A. M., and Seegmiller, B. (2022). New frontiers: The origins and content of new work, 1940–2018. NBER Working Paper 30389, National Bureau of Economic Research.

Autor, D., Katz, L. F., and Krueger, A. (1998). Computing inequality: Have computers changed the labour market? *The Quarterly Journal of Economics*, 113(4):1169–1213.

Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review*, 103(5):1553–1597.

Autor, D. H., Katz, L. F., and Kearney, M. S. (2006). The polarization of the U.S. labour market. *American Economic Review*, 96(2):189–194.

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

Baldassarre, M. T., Caivano, D., Fernandez Nieto, B., Gigante, D., and Ragone, A. (2023). The social impact of generative AI: An analysis on ChatGPT. In *Proceedings of the 2023 ACM Conference on Information Technology for Social Good*, pages 363–373. ACM.

Brynjolfsson, E., Rock, D., and Syverson, C. (2017). Artificial Intelligence and the modern productivity paradox: A clash of expectations and statistics. NBER Working Paper 24001, National Bureau of Economic Research.

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.

Felten, E. W., Raj, M., and Seamans, R. (2018). A method to link advances in Artificial Intelligence to occupational abilities. *AEA Papers and Proceedings*, 108:54–57.

Felten, E. W., Raj, M., and Seamans, R. (2019). The occupational impact of Artificial Intelligence: Labour, skills, and polarization. Ssrn scholarly paper, SSRN, Rochester, NY.

Georgieff, A. and Hyee, R. (2021). Artificial Intelligence and employment: New cross-country evidence. Oecd social, employment and migration working papers, OECD.

Gmyrek, P., Berg, J., and Bescond, D. (2023). Generative AI and jobs: A global analysis of potential effects on job quantity and quality. Ssrn scholarly paper, SSRN, Rochester, NY.

Goos, M. and Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in britain. *The Review of Economics and Statistics*, 89(1):118–133.

Goos, M., Manning, A., and Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, 104(8):2509–2526.

Green, A. and Lamby, L. (2023). The supply, demand and characteristics of the AI workforce across OECD countries. OECD Social, Employment and Migration Working Papers 287, OECD Publishing, Paris.

International Labour Organization (2010). International standard classification of occupations (ISCO-08). Ilo classification, International Labour Organization.

Katz, L. F. and Murphy, K. M. (1992). Changes in relative wages, 1963–1987: Supply and demand factors. *The Quarterly Journal of Economics*, 107(1):35–78.

Lane, M. and Saint-Martin, A. (2021). The impact of Artificial Intelligence on the labour market: What do we know so far? OECD Social, Employment and Migration Working Papers 256, OECD Publishing, Paris.

Lane, M., Williams, M., and Broecke, S. (2023). The impact of AI on the workplace: Main findings from the OECD AI surveys of employers and workers. Oecd report, OECD, Paris.

Lassébie, J. and Quintini, G. (2022). What skills and abilities can automation technologies replicate and what does it mean for workers?: New evidence. Oecd report, OECD, Paris.

Lorenz, P., Perset, K., and Berryhill, J. (2023). Initial policy considerations for generative Artificial Intelligence. OECD Artificial Intelligence Papers 1, OECD Publishing.

Milanez, A. (2023). The impact of AI on the workplace: Evidence from OECD case studies of AI implementation. OECD Social, Employment and Migration Working Papers 289, OECD Publishing, Paris.

Nübler, I. (2016). New technologies: a jobless future or golden age of job creation? Research Department Working Paper 13, International Labour Office, Geneva.

OECD (2019). *OECD Employment Outlook 2019: The Future of Work*. OECD Publishing, Paris.

OECD (2023). *OECD Employment Outlook 2023: Artificial Intelligence and the Labour Market*. Organisation for Economic Co-operation and Development, Paris.

OECD (2024). *OECD Digital Economy Outlook 2024 (Volume 1): Embracing the Technology Frontier*. Organisation for Economic Co-operation and Development, Paris.

Pizzinelli, C., Panton, A. J., Mendes Tavares, M., Cazzaniga, M., and Li, L. (2023). Labor market exposure to AI: Cross-country differences and distributional implications. IMF Working Paper 216, International Monetary Fund.

Tolan, S., Pesole, A., Martínez-Plumed, F., Fernández-Macías, E., Hernández-Orallo, J., and Gómez, E. (2021). Measuring the occupational impact of AI: Tasks, cognitive abilities and AI benchmarks. *Journal of Artificial Intelligence Research*, 71:191–236.

Torrejón Pérez, S., Hurley, J., Fernández-Macías, E., and Staffa, E. (2023). Employment shifts in europe from 1997 to 2021: From job upgrading to polarisation. JRC Working Papers on Labour, Education and Technology 2023-05, Joint Research Centre.

Webb, M. (2020). The impact of Artificial Intelligence on the labor market. Ssrn electronic journal, SSRN.

Zarifhonarvar, A. (2023). Economics of ChatGPT: A labour market view on the occupational impact of Artificial Intelligence. Ssrn scholarly paper, SSRN, Rochester, NY.

Appendix

1.A AI Scores for Occupations at the Three-Digit Level

Table 1.A.1: AI Scores for Occupations at the Three-Digit Level

ISCO-08 code	Occupation	JRC AI score	Felten AI score	Webb AI score
111	Legislators and senior officials	0.559	0.790	0.486
112	Managing directors and chief executives	0.542	0.783	0.413
121	Business services and administration managers	0.754	0.823	0.696
122	Sales, marketing and development managers	0.780	0.865	0.782
132	Manufacturing, mining, construction, and distribution managers	0.737	0.714	0.686
133	Information and communications technology services	0.788	0.785	0.686
134	Professional services managers	0.610	0.789	0.706
141	Hotel and restaurant managers	0.111	0.538	0.686
142	Retail and wholesale trade managers	0.365	0.637	0.686
143	Other services managers	0.398	0.729	0.686
211	Physical and earth science professionals	0.907	0.804	0.931
212	Mathematicians, actuaries and statisticians	0.966	1.000	0.896
213	Life science professionals	0.915	0.702	0.751
214	Engineering professionals (excluding electrotechnology)	0.975	0.814	0.901
215	Electrotechnology engineers	1.000	0.733	0.881
216	Architects, planners, surveyors and designers	0.873	0.705	0.815

Continued on next page

ISCO-08 code	Occupation	JRC AI score	Felten AI score	Webb AI score
221	Medical doctors	0.712	0.689	0.533
222	Nursing and midwifery professionals	0.238	0.564	0.789
224	Paramedical practitioners	0.340	0.537	0.236
225	Veterinarians	0.432	0.485	0.861
226	Other health professionals	0.450	0.619	0.565
231	University and higher education teachers	0.898	0.888	0.000
232	Vocational education teachers	0.763	0.591	0.277
233	Secondary education teachers	0.881	0.835	0.430
234	Primary school and early childhood teachers	0.704	0.677	0.420
235	Other teaching professionals	0.594	0.753	0.525
241	Finance professionals	0.941	0.980	0.719
242	Administration professionals	0.831	0.896	0.562
243	Sales, marketing and public relations professionals	0.814	0.840	0.337
251	Software and applications developers and analysts	0.983	0.843	0.863
252	Database and network professionals	0.992	0.787	0.840
261	Legal professionals	0.729	0.979	0.799
262	Librarians, archivists and curators	0.661	0.679	0.341
263	Social and religious professionals	0.627	0.861	0.713
264	Authors, journalists and linguists	0.805	0.821	0.307
265	Creative and performing artists	0.424	0.515	0.327
311	Physical and engineering science technicians	0.950	0.547	0.859
312	Mining, manufacturing and construction supervisors	0.331	0.440	0.912
313	Process control technicians	0.890	0.397	0.833
314	Life science technicians and related associates	0.923	0.481	0.840
315	Ship and aircraft controllers and technicians	0.652	0.556	0.531
321	Medical and pharmaceutical technicians	0.669	0.485	0.658
322	Nursing and midwifery associate professionals	0.458	0.465	0.323
324	Veterinary technicians and assistants	0.297	0.331	0.338
325	Other health associate professionals	0.483	0.471	0.393
331	Financial and mathematical associate professionals	0.932	0.841	0.604
332	Sales and purchasing agents and brokers	0.585	0.860	0.610
333	Business services agents	0.644	0.784	0.631

Continued on next page

ISCO-08 code	Occupation	JRC AI score	Felten AI score	Webb AI score
334	Administrative and specialised secretaries	0.823	0.777	0.175
335	Regulatory government associate professionals	0.509	0.632	0.557
341	Legal, social and religious associate professionals	0.525	0.730	0.361
342	Sports and fitness workers	0.246	0.405	0.453
343	Artistic, cultural and culinary associate professionals	0.602	0.435	0.502
351	Information and communications technology operators	0.958	0.684	0.802
352	Telecommunications and broadcasting technicians	0.848	0.561	0.874
411	General office clerks	0.796	0.885	0.164
412	Secretaries (general)	0.771	0.794	0.082
413	Keyboard operators	0.839	0.615	0.517
421	Tellers, money collectors and related clerks	0.382	0.635	0.317
422	Client information workers	0.407	0.766	0.331
431	Numerical clerks	0.856	0.880	0.181
432	Material-recording and transport clerks	0.636	0.499	0.531
441	Other clerical support workers	0.695	0.489	0.307
511	Travel attendants, conductors and guides	0.373	0.500	0.216
512	Cooks	0.161	0.320	0.215
513	Waiters and bartenders	0.034	0.262	0.056
514	Hairdressers, beauticians and related workers	0.077	0.401	0.215
515	Building and housekeeping supervisors	0.136	0.230	0.615
516	Other personal services workers	0.356	0.432	0.270
521	Street and market salespersons	0.102	0.624	0.108
522	Shop salespersons	0.085	0.538	0.470
523	Cashiers and ticket clerks	0.059	0.464	0.379
524	Other sales workers	0.144	0.408	0.105
531	Child care workers and teachers' aides	0.280	0.539	0.143
532	Personal care workers in health services	0.196	0.335	0.243
541	Protective services workers	0.288	0.374	0.611
611	Market gardeners and crop growers	0.475	0.209	0.901
612	Animal producers	0.467	0.258	1.000
613	Mixed crop and animal producers	0.865	0.207	0.990
621	Forestry and related workers	0.678	0.203	0.784

Continued on next page

ISCO-08 code	Occupation	JRC AI score	Felten AI score	Webb AI score
622	Fishery workers, hunters and trappers	0.094	0.216	0.900
711	Building and related trades in construction	0.169	0.128	0.565
712	Building finishers and related trades workers	0.390	0.071	0.486
713	Painters, building structure cleaners and related trades	0.204	0.030	0.584
721	Sheet and structural metal workers, moulders	0.492	0.087	0.439
722	Blacksmiths, toolmakers and related trades workers	0.500	0.233	0.738
723	Machinery mechanics and repairers	0.551	0.202	0.593
731	Handicraft workers	0.305	0.287	0.493
732	Printing trades workers	0.746	0.419	0.803
741	Electrical equipment installers and repairers	0.577	0.261	0.647
742	Electronics and telecommunications installers and repairers	0.686	0.339	0.702
751	Food processing and related trades workers	0.127	0.209	0.520
752	Wood treaters, cabinet-makers and related trades	0.568	0.126	0.734
753	Garment and related trades workers	0.348	0.198	0.593
754	Other craft and related workers	0.271	0.308	0.644
811	Mining and mineral processing plant operators	0.314	0.134	0.665
812	Metal processing and finishing plant operators	0.619	0.123	0.716
813	Chemical and photographic products plant and machine operators	0.517	0.263	0.668
814	Rubber, plastic and paper products machine operators	0.229	0.141	0.628
815	Textile, fur and leather products machine operators	0.323	0.130	0.676
816	Food and related products machine operators	0.441	0.266	0.547
817	Wood processing and papermaking plant operators	0.534	0.140	0.645
818	Other stationary plant and machine operators	0.415	0.196	0.547
821	Assemblers	0.263	0.288	0.539
831	Locomotive engine drivers and related workers	0.721	0.308	0.669
832	Car, van and motorcycle drivers	0.255	0.317	0.318
833	Heavy truck and bus drivers	0.221	0.377	0.496
834	Mobile plant operators	0.153	0.214	0.715
835	Ships' deck crews and related workers	0.178	0.152	0.564
911	Domestic, hotel and office cleaners and helpers	0.026	0.012	0.111

Continued on next page

ISCO-08 code	Occupation	JRC AI score	Felten AI score	Webb AI score
912	Vehicle, window, laundry and other hand cleaning workers	0.009	0.013	0.320
921	Agricultural, forestry and fishery labourers	0.051	0.052	0.799
931	Mining and construction labourers	0.212	0.000	0.404
932	Manufacturing labourers	0.186	0.045	0.680
933	Transport and storage labourers	0.069	0.214	0.262
941	Food preparation assistants	0.017	0.126	0.123
952	Street vendors (excluding food)	0.000	0.795	0.123
961	Refuse workers	0.042	0.218	0.512
962	Other elementary workers	0.119	0.235	0.397

1.B Change in employment vs. exposure to AI

Table 1.B.1: AI score vs. change in employment share for countries (2011–2021)

	(1) Base	(2) Country FE	(3) Low ed.	(4) Med ed.	(5) High ed.	(6) Young	(7) Core	(8) Old
Panel A: JRC								
JRC	0.267***							
	(0.061)							
JRC×AT		0.085	0.743***	-1.948***	0.179*	0.381***	0.037	0.006
		(0.086)	(0.123)	(0.162)	(0.101)	(0.084)	(0.070)	(0.109)
JRC×BE		0.452***	0.269**	-1.893***	0.773***	0.551***	0.211***	0.037
		(0.078)	(0.118)	(0.152)	(0.081)	(0.100)	(0.066)	(0.070)
JRC×CH		-0.105	0.567***	-1.396***	-0.182***	0.541***	0.189***	0.217**
		(0.068)	(0.111)	(0.131)	(0.069)	(0.072)	(0.064)	(0.088)
JRC×CY		0.909***	0.359*	-1.410***	1.358***	0.619***	0.343***	0.057
		(0.090)	(0.204)	(0.171)	(0.070)	(0.111)	(0.089)	(0.174)
JRC×CZ		0.346***	1.858***	-1.361***	-0.394***			
		(0.105)	(0.105)	(0.148)	(0.092)			
JRC×DE		0.490***	0.918***	-1.263***	1.266***	0.603***	0.110	0.007
		(0.093)	(0.090)	(0.137)	(0.074)	(0.086)	(0.073)	(0.096)
JRC×DK		0.191*	0.714***	-1.054***	1.433***	0.594***	0.128**	0.082
		(0.102)	(0.094)	(0.110)	(0.087)	(0.074)	(0.061)	(0.078)
JRC×EE		1.322***	1.482***	-2.143***	1.735***	0.392**	0.363***	-0.246***
		(0.081)	(0.104)	(0.179)	(0.075)	(0.173)	(0.080)	(0.089)
JRC×EL		-0.258***	-2.332***	-0.126	-0.151*	0.566***	0.393***	0.158
		(0.092)	(0.619)	(0.168)	(0.089)	(0.104)	(0.093)	(0.209)
JRC×ES		0.332***	-4.583***	0.213	0.507***	0.144*	0.216***	0.130
		(0.074)	(0.404)	(0.172)	(0.072)	(0.082)	(0.058)	(0.083)
JRC×FI		0.629***	0.636***	-2.260***	1.055***	0.457***	-0.013	0.601***
		(0.091)	(0.117)	(0.171)	(0.072)	(0.119)	(0.080)	(0.094)
JRC×FR		0.226***	0.329***	-1.182***	0.338***	0.452***	0.325***	0.296***
		(0.074)	(0.111)	(0.171)	(0.089)	(0.142)	(0.086)	(0.071)
JRC×HR		-0.022	-0.088	0.492***	-0.358***	0.523***	0.130	0.067
		(0.104)	(0.579)	(0.178)	(0.102)	(0.110)	(0.124)	(0.297)
JRC×HU		-0.215***	1.058***	-1.403***	-0.241***	0.516***	0.246***	0.301***
		(0.081)	(0.132)	(0.151)	(0.086)	(0.077)	(0.073)	(0.095)
JRC×IE		0.565***	-0.415	-1.279***	1.132***	-0.305	-0.130	0.074

Continued on next page

Table 1.B.1: AI score vs. change in employment share for countries (2011–2021) (continued)

	(1) Base	(2) Country FE	(3) Low ed.	(4) Med ed.	(5) High ed.	(6) Young	(7) Core	(8) Old
		(0.136)	(0.298)	(0.141)	(0.071)	(0.1315)	(0.123)	(0.102)
JRC×IT		-0.039	-4.784***	0.859***	-0.916***	0.001	0.126**	0.277***
		(0.080)	(0.428)	(0.114)	(0.091)	(0.070)	(0.062)	(0.060)
JRC×LT		0.486***	0.204	-1.748***	0.828***	1.057***	0.742***	0.288***
		(0.109)	(0.372)	(0.172)	(0.076)	(0.118)	(0.078)	(0.085)
JRC×LU		1.166***	0.903***	-0.876***	1.574***			
		(0.083)	(0.109)	(0.149)	(0.063)			
JRC×LV		0.252***	0.632***	-1.228***	0.111	0.727***	0.454***	0.307**
		(0.083)	(0.116)	(0.208)	(0.094)	(0.105)	(0.092)	(0.120)
JRC×NL		0.757***	0.174	-1.094***	1.696***			
		(0.088)	(0.108)	(0.124)	(0.066)			
JRC×NO		-0.179**	0.135	-0.337***	-0.239	0.644***	0.188***	0.151*
		(0.078)	(0.110)	(0.105)	(0.080)	(0.166)	(0.067)	(0.089)
JRC×PL		0.271***	1.418***	-2.296***	-0.161***	-0.128	0.101	-0.538***
		(0.091)	(0.354)	(0.165)	(0.082)	(0.101)	(0.120)	(0.140)
JRC×PT		0.257***	-6.213***	0.954***	-0.574*	0.724***	0.423***	0.827***
		(0.076)	(0.800)	(0.132)	(0.116)	(0.110)	(0.077)	(0.082)
JRC×RO		-0.366	-2.125***	0.523	-0.026***	0.093	0.137	0.929
		(0.237)	(0.397)	(0.527)	(0.084)	(0.270)	(0.300)	(0.753)
JRC×SE								
Panel B: Felten								
Felten	0.036***							
	(0.050)							
Felten×AT		0.204**	1.049***	-1.727***	0.241**	0.518***	0.176**	0.006
		(0.090)	(0.086)	(0.360)	(0.101)	(0.112)	(0.086)	(0.109)
Felten×BE		0.541***	0.562***	-2.017***	0.815***	0.651***	0.319***	-0.037
		(0.074)	(0.100)	(0.126)	(0.075)	(0.103)	(0.063)	(0.070)
Felten×CH		-0.003	0.849***	-1.386***	-0.117*	0.657***	0.343***	-0.217**
		(0.064)	(0.105)	(0.126)	(0.068)	(0.080)	(0.060)	(0.088)
Felten×CY		0.931***	0.686***	-1.412***	1.400***	0.663***	0.427***	0.057
		(0.103)	(0.152)	(0.158)	(0.067)	(0.111)	(0.094)	(0.174)
Felten×CZ		0.405***	2.228***	-1.319***	-0.371***			

Continued on next page

Table 1.B.1: AI score vs. change in employment share for countries (2011–2021) (continued)

	(1) Base	(2) Country FE	(3) Low ed.	(4) Med ed.	(5) High ed.	(6) Young	(7) Core	(8) Old
		(0.112)	(0.096)	(0.152)	(0.087)			
Felten×DE		0.581***	1.180***	-1.170***	1.353***	0.694***	0.249***	-0.007
		(0.096)	(0.087)	(0.150)	(0.075)	(0.100)	(0.081)	(0.096)
Felten×DK		0.262**	0.931***	-0.978***	1.489***	0.708***	0.249***	0.082
		(0.103)	(0.095)	(0.097)	(0.079)	(0.084)	(0.057)	(0.078)
Felten×EE		1.451***	1.902***	-2.108***	1.792***	0.428**	0.408***	-0.152*
		(0.084)	(0.092)	(0.163)	(0.069)	(0.206)	(0.089)	(0.089)
Felten×EL		-0.199***	-3.238***	-0.135	-0.106	0.643***	0.533***	0.158
		(0.101)	(0.676)	(0.223)	(0.099)	(0.180)	(0.105)	(0.209)
Felten×ES		0.412***	-3.340***	0.358*	0.579***	0.197**	0.347***	0.130
		(0.082)	(0.666)	(0.215)	(0.070)	(0.094)	(0.067)	(0.083)
Felten×FI		0.726***	0.863***	-2.384***	1.140***	0.486***	0.087	-0.601***
		(0.099)	(0.143)	(0.166)	(0.075)	(0.149)	(0.085)	(0.094)
Felten×FR		0.309***	0.531***	-1.156***	0.392***	0.406**	0.403**	0.339***
		(0.072)	(0.118)	(0.198)	(0.086)	(0.168)	(0.090)	(0.073)
Felten×HR		-0.137	1.060***	0.740***	-0.328***	0.621***	0.418***	-0.067
		(0.092)	(0.256)	(0.185)	(0.099)	(0.119)	(0.088)	(0.297)
Felten×HU		-0.138	1.409***	-1.453***	-0.227***	0.571***	0.369***	0.301***
		(0.090)	(0.091)	(0.130)	(0.083)	(0.088)	(0.074)	(0.095)
Felten×IE		0.710***	-0.108	-1.360***	1.180***	-0.267	-0.068	0.034
		(0.089)	(0.161)	(0.172)	(0.073)	(0.091)	(0.82)	(0.088)
Felten×IT		-0.046	-4.960***	1.051***	-0.895***	0.086**	0.265***	0.277***
		(0.088)	(0.521)	(0.098)	(0.092)	(0.078)	(0.062)	(0.060)
Felten×LT		0.618***	0.778***	-1.769***	0.844***	0.873**	0.642***	0.253***
		(0.078)	(0.171)	(0.168)	(0.073)	(0.128)	(0.111)	(0.095)
Felten×LU		1.236***	1.105***	-0.821***	1.589***			
		(0.075)	(0.113)	(0.142)	(0.060)			
Felten×LV		0.319***	0.921***	-1.228***	0.151*	0.759***	0.528***	0.300**
		(0.80)	(0.126)	(0.241)	(0.088)	(0.134)	(0.097)	(0.120)
Felten×NL		0.839***	0.376***	-1.021***	1.744***			
		(0.088)	(0.101)	(0.109)	(0.060)			
Felten×NO		-0.119	0.242**	-0.204*	-0.214***	0.805***	0.323***	-0.150*
		(0.084)	(0.110)	(0.120)	(0.073)	(0.047)	(0.093)	(0.089)
Felten×PL		0.346***	2.245***	-2.380***	-0.154*	-0.186**	0.247***	-0.434***

Continued on next page

Table 1.B.1: AI score vs. change in employment share for countries (2011–2021) (continued)

	(1) Base	(2) Country FE	(3) Low ed.	(4) Med ed.	(5) High ed.	(6) Young	(7) Core	(8) Old
		(0.112)	(0.086)	(0.170)	(0.081)	(0.095)	(0.071)	(0.114)
Felten×RO		0.117	-1.973	1.669***	0.929	0.488***	0.698***	-0.929
		(0.173)	(1.352)	(0.197)	(0.753)	(0.134)	(0.120)	(0.753)
Felten×SE		0.224***	0.912***	-1.536***	-0.008	0.886***	0.479***	-0.075
		(0.072)	(0.073)	(0.144)	(0.082)	(0.079)	(0.062)	(0.062)
Panel C: Webb								
WEBB	0.021							
	(0.070)							
WEBB×AT		-0.106	0.170	-1.500***	-0.071	0.037	-0.099	-0.182
		(0.099)	(0.126)	(0.421)	(0.095)	(0.134)	(0.127)	(0.136)
WEBB×BE		0.167*	-0.409***	-1.716***	0.629***	0.179	-0.039	-0.262***
		(0.096)	(0.119)	(0.155)	(0.087)	(0.133)	(0.098)	(0.095)
WEBB×CH		-0.323***	-0.014	-1.437***	-0.469***	0.212**	-0.012	-0.477***
		(0.073)	(0.126)	(0.140)	(0.072)	(0.101)	(0.094)	(0.100)
WEBB×CY		0.478***	-0.821***	-1.158***	1.319***	0.291**	-0.021	-0.410**
		(0.134)	(0.168)	(0.200)	(0.089)	(0.138)	(0.124)	(0.165)
WEBB×CZ		0.305**	1.234***	-1.643***	-0.768***			
		(0.126)	(0.103)	(0.164)	(0.116)			
WEBB×DE		0.187*	0.282**	-1.492***	1.137***	0.101	-0.199**	-0.279***
		(0.107)	(0.119)	(0.127)	(0.083)	(0.104)	(0.088)	(0.080)
WEBB×DK		0.043	0.248**	-1.110***	1.507***	0.326***	-0.064	-0.053
		(0.108)	(0.110)	(0.149)	(0.084)	(0.108)	(0.091)	(0.095)
WEBB×EE		1.058***	0.828***	-2.314***	1.695***	0.085	0.038	-0.410***
		(0.094)	(0.112)	(0.162)	(0.075)	(0.201)	(0.097)	(0.089)
WEBB×EL		-0.542***	-1.840***	-0.106	-0.498***	0.091	0.089	-0.312
		(0.117)	(0.238)	(0.163)	(0.128)	(0.188)	(0.152)	(0.191)
WEBB×ES		0.032	-3.153***	0.460***	0.187**	-0.323***	-0.028	-0.235**
		(0.082)	(0.405)	(0.126)	(0.074)	(0.112)	(0.085)	(0.104)
WEBB×FI		0.300**	-0.038	-2.074***	0.946***	-0.134	-0.266**	-0.919***
		(0.117)	(0.149)	(0.166)	(0.081)	(0.149)	(0.114)	(0.116)
WEBB×FR		-0.110	-0.339***	-1.135***	0.033	0.067	-0.011	0.128
		(0.079)	(0.114)	(0.143)	(0.091)	(0.129)	(0.085)	(0.086)
WEBB×HR		-0.258**	-0.593*	0.341**	-0.735***	0.102	-0.023	-0.194

Continued on next page

Table 1.B.1: AI score vs. change in employment share for countries (2011–2021) (continued)

	(1) Base	(2) Country FE	(3) Low ed.	(4) Med ed.	(5) High ed.	(6) Young	(7) Core	(8) Old
		(0.103)	(0.334)	(0.162)	(0.123)	(0.135)	(0.141)	(0.309)
WEBB×HU		-0.464***	0.209	-1.112***	-0.589***	0.143	-0.015	0.150
		(0.094)	(0.140)	(0.133)	(0.100)	(0.107)	(0.101)	(0.107)
WEBB×IE		0.202	-0.960***	-1.194***	0.900***	-0.441***	-0.086	0.033
		(0.153)	(0.183)	(0.149)	(0.110)	(0.101)	(0.090)	(0.080)
WEBB×IT		0.002	-3.679***	1.120***	-1.467***	0.853***	0.675***	0.227**
		(0.103)	(0.314)	(0.115)	(0.126)	(0.145)	(0.104)	(0.104)
WEBB×LT		0.140	-0.268	-1.437***	0.628***	0.436	0.532	0.717
		(0.117)	(0.217)	(0.157)	(0.093)	(0.124)	(0.109)	(0.104)
WEBB×LU		1.154***	0.308**	-0.944***	1.615***			
		(0.101)	(0.143)	(0.173)	(0.074)			
WEBB×LV		0.045	-0.059	-1.334***	-0.120	0.253*	0.144	0.198
		(0.086)	(0.129)	(0.217)	(0.106)	(0.145)	(0.124)	(0.126)
WEBB×NL		0.473***	-0.533***	-1.148***	1.712***			
		(0.112)	(0.141)	(0.127)	(0.074)			
WEBB×NO		-0.298***	-0.375***	-0.228	-0.318***	0.200	0.071	-0.328***
		(0.111)	(0.139)	(0.195)	(0.092)	(0.260)	(0.140)	(0.126)
WEBB×PL		0.245**	0.831***	-2.434***	-0.569***	0.392***	-0.260**	-0.643***
		(0.112)	(0.270)	(0.162)	(0.091)	(0.106)	(0.126)	(0.110)
WEBB×PT		0.049	-3.453***	0.762***	-0.460***	0.428***	0.047	0.563***
		(0.075)	(0.524)	(0.127)	(0.107)	(0.130)	(0.099)	(0.135)
WEBB×RO		-0.521**	-1.804***	0.495	-1.116***	-0.271	-0.049	-1.007*
		(0.229)	(0.384)	(0.404)	(0.142)	(0.260)	(0.287)	(0.612)
WEBB×SE		0.026	0.185*	-1.738***	-0.318***	0.448***	0.165*	-0.303***
		(0.078)	(0.104)	(0.166)	(0.092)	(0.102)	(0.092)	(0.083)

[flushleft]

Notes: Linear regression. Robust standard errors in parentheses. *p<0.1, **p<0.05, ***p<0.01. Each observation is an ISCO 3-digit occupation × sector cell. Observations are weighted by cells' average labour supply. Dependent variable: within-country cell's change in employment shares from 2011 to 2021, winsorized at the top and bottom 1 percent. The sub-sample in columns (3)–(5) consists of sector-occupation cells whose average educational attainment is in the lower, middle, and upper tercile, respectively, of the country's education distribution in 2021. The sub-samples in columns (6)–(8) consist of sector-occupation cells whose workers' age was in the lower, middle, and upper tercile, respectively, of the country's workers' age distribution in 2021. CZ, NL, LU do not provide information regarding age in columns (6)–(8); therefore, it was not possible to proceed with the regression analysis.

Chapter 2

Still too early? Artificial Intelligence and Monopsony Power

Abstract

We examine the relationship between monopsony power and occupational exposure to artificial intelligence across 26 European countries over the period 2011–2020. Utilizing data from the Labour Force Survey (LFS) to extract wage and employment information, and Google Patents to construct an index measuring the exposure of occupations to artificial intelligence, we assess monopsony power via the wage elasticity of labor supply. Our results indicate a marked decline in this elasticity over time, suggesting an increase in firms' monopsony power. Notably, this trend appears independent of variations in AI exposure, which exert only limited influence on labor supply elasticity during the study period. Stratified analyses by wage further reveal that occupations within the lowest wage group experience the most pronounced monopsony power, followed by those in the highest wage group, while the medium wage group exhibits the weakest monopsony power. Analyses by educational terciles suggest that low-education occupations display an increasing trend in monopsony power.

Keywords: *Monopsony power, Artificial intelligence, Labor supply elasticity, Occupational exposure, Wage inequality, European labor markets*

JEL codes: O33, J42, J23, J31, J24

2.1 Introduction

It has been roughly a century since Joan Robinson's seminal work, *The Economics of Imperfect Competition*, criticized monopsony power for suppressing workers' incomes, generating inefficient levels of employment, and ultimately creating social inefficiencies. Despite the predictions of competitive market theory, reality has diverged from these expectations. Recent studies have documented emerging market imperfections (Autor et al., 2023; Dube et al., 2022) as well as an increase in monopsony power (Ashenfelter et al., 2021; Berger et al., 2022) in both Western and Eastern countries (Brooks et al., 2021). This issue is crucial to examine because it may underpin several significant trends observed in recent years, including rising inequality (Cairó and Sim, 2020; Impullitti and Rendahl, 2025), deteriorating worker well-being, and a declining labour share (Bellocchi et al., 2023; Bergholt et al., 2022). Moreover, increasing firm concentration (Calligaris et al., 2025; Covarrubias et al., 2020; Grullon et al., 2019) and the rapid growth of enterprises (Autor et al., 2020) further exacerbate these challenges.

Monopsony can be defined as the ability of employers to pay workers less than the competitive wage, acknowledging that the appropriate competitive wage benchmark may vary depending on the underlying assumptions (Azar and Marinescu, 2024). In monopsonistic labour markets, employers exploit their bargaining power to suppress wages and depress employment, thereby accumulating excess rents (OECD, 2022). More specifically, employers reduce labour demand to lower costs and maximize profits by paying workers below their marginal productivity, a phenomenon extensively documented in the literature (Boal and Ransom, 1997; Manning, 2003). In competitive markets, wages would typically align with workers' marginal productivity, leading to higher levels of employment and more equitable compensation.

This trend in the labour market intersects with the broader impacts of technological progress on the economy. As demonstrated by Kurz (2017), technological progress can significantly increase firms' market power, and evidence from the digital economy confirms that such advancements directly influence market power (OECD, 2022). Today, this issue is even more salient given the disruptive potential of artificial intelligence (Agrawal et al., 2019; Păvăloaia and Necula, 2023). In this context, AI has raised serious concerns about rising inequality and the erosion of workers' wages, ultimately resulting in a larger capital share at the expense of the labour share (Acemoglu, 2021). Several studies have examined the relationship between artificial intelligence and market power, as well as that between monopsony and various technological sectors. For instance, (Gans, 2024) focuses on three key segments, training data, input data, and AI predictions, and argues that the proper functioning of data markets, which allow data to be traded across firm boundaries, is crucial for the emergence and persistence of market power. Similarly, (Azar et al., 2023) investigate the broader impact of automation on monopsony, arguing that, due to higher marginal labour costs, monopsonistic firms have stronger incentives to automate than wage-taking firms. Their findings, based on data from U.S. commuting zones, indicate that areas with greater exposure to industrial robots experience significantly larger reductions in both employment and wages in highly concentrated labour markets.

Nonetheless, the direct relationship between artificial intelligence and monopsony remains relatively underexplored. This paper aims to investigate this link by examining how AI has influenced the elasticity of the labour supply curve, and, by extension, monopsony power, across 16 European countries between 2011 and 2020. Our contribution builds on Manning (2003) model, which posits that worker mobility depends on wage levels, thereby measuring monopsony power through labour supply elasticity. To calculate this measure, we extend the methodology of (Bachmann et al., 2022), who investigated the extent to which workers performing different tasks are exposed to varying degrees of monopsony power over the past three decades. In our approach, we incorporate an interaction between our key variables and an index of occupational exposure to artificial intelligence, computed according to the procedure outlined by Webb (2020), with the crucial modification that our index is time-varying.

The primary contribution of our paper is that it provides an overview of how artificial intelligence has affected the elasticity of labour supply, and its individual components, namely, the separation rates from employment and non-employment and the recruitment rate from non-employment, across European countries between 2011 and 2020. Our results indicate that during this period labour supply elasticity has decreased, suggesting an increase in monopsony power. However, the direct impact of artificial intelligence on this trend appears to be limited, which may imply that the period examined is "still too early" (OECD, 2023) to capture the full disruptive effects of AI. To further investigate the differential impact of artificial intelligence across skill levels, we subdivide the observations into income terciles and educational terciles. The literature suggests that AI predominantly affects high-skill and non-routine occupations (Green, 2024; Lane and Saint-Martin, 2021; Lassébie and Quintini, 2022), which we then expect to be reflected in the highest income and education terciles. Our findings suggest that low-skill occupations experience significantly higher monopsony power, which increases over time across both educational and income dimensions. Once again, the influence of artificial intelligence on these patterns remains modest.

It is important to note that the period under consideration predates the significant AI boom observed in later years, so future research should revisit this analysis as more recent data becomes available.

The remainder of this paper is organized as follows. Section 2 reviews the most pertinent literature on monopsony and AI exposure measurements. Section 3 details the data, and Section 4 details the methodology used. Section 4 presents the descriptive evidence, while Section 5 discusses the results. Finally, Section 6 concludes.

2.2 Literature review

2.2.1 Monopsony and power market

Over the past two decades, economists have increasingly recognized that labour markets often deviate from the frictionless perfect competition model, enabling employers to pay wages below the marginal productivity of workers and, consequently, fostering monopsonistic power (Boal and Ransom, 1997; Manning, 2003, 2021) . As a result, many firms are able to lower wages without triggering mass resignations.

This phenomenon not only has significant implications for income distribution and worker conditions, but also has important theoretical consequences. Card (2022) notes that traditional models treated wages merely as "market outcomes," largely overlooking firm-level wage-setting policies. Building on earlier insights from Diamond (1971); Mortensen (1970), he argues that recent theoretical and empirical advances,such as on-the-job search models (Burdett and Mortensen, 1998; Pissarides, 1985) and preferences for different employers (Chamberlin, 1973; Kaldor, 1934), demonstrate that real-world labour markets diverge from perfect competition. Card emphasizes that the "monopsony in motion" perspective (Manning, 2003) elucidates how wage frictions can emerge even in the presence of many employers. Much like the shift from perfect competition to monopolistic competition in product markets (Dixit and Stiglitz, 1977; Spence, 1976), contemporary economists agree that a limited number of employers, by increasing their market power, can unilaterally set wages.

A canonical approach to measuring this power, following Burdett and Mortensen (1998), is to compute separation and recruitment rates in order to derive the elasticity of labour demand with respect to wages. Crucially, a firm's monopsonistic power is closely linked to workers responsiveness to wage changes: the less elastic the labour supply, the more likely workers are to remain employed despite wage variations, thereby exacerbating the wage markdown (Manning, 2003). Empirical studies, ranging from linked employer-employee datasets to analyses of online job postings, consistently confirm significant wage-setting power across diverse labour market contexts when using this elasticity methodology (Ashenfelter et al., 2010, 2021; Dube et al., 2020). For example,Sokolova and Sorensen (2021), in a meta-analysis of over one thousand estimates drawn from 53 global studies, conclude that, on average, marginal wages are approximately 7% lower than they would be in a perfectly competitive environment. Complementary evidence from Azar et al. (2019) indicates that the elasticity of labour applications with respect to wages is roughly 0.43 and tends to be lower in more concentrated labour markets, suggesting reduced competition. Similarly, Yeh et al. (2022) report that the vast majority of U.S. manufacturing plants operate under monopsonistic conditions and that, since the early 2000s, the U.S. manufacturing labour market has become increasingly monopsonistic. They interpret the gap between the output elasticity of labour and labour's revenue share as clear evidence of substantial market power. Moreover, the literature now makes it evident that monopsony power contributes to wage inequality in multiple dimensions, among similar workers at different employers (Song et al., 2019) , both between and within firms (Zwysen, 2024a), and particularly affecting women who are generally less inclined to change jobs (Harkness et al., 2023). This dynamic reduces workers' bargaining power (Caldwell and Danieli, 2024) and disproportionately impacts jobs experiencing significant labour shortages (Eurofound et al., 2023a,b).

In the context of job polarization, the decline in employment among workers in the middle of the skill distribution (Autor et al., 2003; Bachmann et al., 2022; Goos and Manning, 2007), demonstrates that workers in non-routine cognitive occupations exhibit a less elastic labour supply compared to those in routine manual roles. This finding implies that, as job polarization intensifies, middle-skill positions may be particularly vulnerable to monopsonistic exploitation, since firms are able to maintain larger wage markdowns in segments where alternative employment opportunities are limited.

Simultaneously, the rapid emergence of advanced technologies, including artificial intelligence (AI) and various machine learning applications, has catalyzed extensive debate about their impact on labour market power (Agrawal et al., 2019; Gans, 2024). On one hand, AI-driven analytical tools, such as predictive algorithms and recommendation systems, can reduce workers' search costs by making job opportunities more visible, potentially eroding firms' wage-setting power (Manning, 2021; Zwysen, 2024b). On the other hand, the adoption of proprietary AI-based screening or performance monitoring systems can reinforce an employer's position by increasing barriers to worker mobility and enabling more tailored wage offers that minimize turnover (Calvano et al., 2023; Yeh et al., 2022).

Recent theoretical research further elucidates these dynamics. For instance, Hagiu and Wright (2023) examine "data feedback loops" to demonstrate how improvements in data, collected through AI-enabled learning processes, can generate a self-reinforcing market power for an incumbent, even though the precise form of the data-enabled learning function may occasionally allow a lagging firm to catch up. Other studies have argued that the degree of market power granted by data critically depends on whether firms share or trade training data across organizational boundaries, a factor that can either mitigate or reinforce the data advantage(Gans, 2024; Jones and Tonetti, 2020). Turning to the digital realm, recent research on online labour markets (OLMs), such as that documented by (Duch-Brown et al., 2023), shows that platform policies and design choices substantially influence the elasticities of both labour demand and labour supply. For example, when platforms require employers to signal their willingness to pay based on the required level of expertise, this exogenous policy shift can lead to increased competition among workers. However, the same digital interfaces and contractual modalities, such as hourly versus fixed-wage projects, can also facilitate employer monopsony by reducing worker mobility (Chen and Horton, 2016; Dube et al., 2020). Duch-Brown et al. (2021)report that AI-related projects on platforms such as PeoplePerHour tend to exhibit significantly higher labour demand and lower labour supply compared to non-AI projects, with workers on AI projects earning 3.0–3.2% higher wages.

In summary, the evidence indicates that monopsonistic power remains pervasive in contemporary labour markets, driven both by traditional factors, such as limited alternative employment options and search frictions, and by new technological dynamics, including AI-enabled data feedback loops. This convergence of theoretical and empirical insights necessitates a reassessment of labour market policies, which may include the implementation of sector-specific antitrust measures (Azar and Marinescu, 2024) and the strengthening of collective bargaining mechanisms (Benmelech et al., 2018). Such measures are essential to ensure that the benefits of technological innovation are distributed more equitably rather than being captured by a few dominant employers.

2.2.2 AI measurement

Various studies have sought to measure or predict which tasks artificial intelligence is capable of performing in order to assess the occupational impact that this technology may have on the labour market.

Different approaches have been developed to estimate AI exposure across occupations, relying on various methodologies and data sources. Brynjolfsson et al.

(2018) employ human annotators to evaluate 2,069 detailed work activities in the ONET database, specifically assessing their potential to be performed by machine learning. This approach provides a structured estimation of AI's capabilities and limitations, offering insight into which tasks may be most susceptible to automation. Webb (2020) takes a different approach by evaluating AI exposure through textual analysis of patents and occupational descriptions. By extracting verb-noun pairs from patent titles and comparing them with those in O*NET job descriptions, this method assigns an exposure score to each occupation based on the frequency of similar pairs in AI-related patents, providing an indication of how AI technologies could be applied to job tasks.

Tolan et al. (2021) develop an AI score, which measures occupational exposure by linking AI developments to cognitive abilities required in different jobs. This approach uses data from the European Working Conditions Survey (EWCS), the Programme for the International Assessment of Adult Competencies (PIAAC), and O*NET, mapping 59 generic tasks to 14 cognitive abilities derived from cognitive science. These abilities are then connected to AI benchmarks to estimate research intensity in different AI fields, providing insight into the short- and medium-term impact of AI on occupations.

Felten et al. (2021) introduce the AI Occupational Exposure (AIOE) index, which connects ten common AI applications, such as image recognition and text generation, with 52 occupational abilities, including oral comprehension and inductive reasoning. This measure is constructed using a crowdsourced matrix that quantifies the relationship between AI applications and occupational skills, weighting occupations by the prevalence and importance of these abilities. The underlying data is drawn from ONET and the Electronic Frontier Foundation's AI Progress Measurement Project.

Based on AIOE, Pizzinelli et al. (2023) extends the analysis by incorporating a measure of AI complementarity using work contexts and job zones in O*NET. Work contexts capture the physical and social conditions affecting AI adoption, while job zones categorize occupations by their required education, training, and experience. This approach considers that occupations with extensive training requirements may be more capable of integrating AI tools into their workflows, highlighting the role of skill adaptation in mitigating displacement risks.

Briggs and Kodnani (2023) investigate the potential automation effects of generative AI by analyzing job tasks using O*NET and ESCO. They classify 13 out of 39 work activities as susceptible to AI automation and assume that AI can complete tasks up to level four on O*NET's seven-point complexity scale. Their results indicate that approximately two-thirds of U.S. occupations exhibit some degree of

AI exposure, with 25% of total work tasks potentially automatable. The highest levels of exposure are found in administrative (46%) and legal (44%) professions, while physically intensive jobs such as construction (6%) and maintenance (4%) demonstrate minimal exposure. A similar analysis conducted for Europe using the ISCO classification system confirms comparable trends.

Eloundou et al. (2023) propose an exposure rubric to assess the impact of large language models (LLMs) on job tasks. They define exposure as the extent to which LLMs can reduce the time required to perform a given task by at least 50%. Their measure incorporates both human annotations and GPT-4-generated evaluations, categorizing exposure into three levels: the lower bound of exposure, tasks requiring complementary tools, and the upper bound capturing maximal AI exposure. This framework allows for a more detailed estimation of how LLMs may influence occupational structures.

Handa et al. (2025) contribute to the literature by empirically analyzing AI adoption patterns using large-scale data on human-AI interactions. Mapping these interactions to occupational categories in O*NET, the study identifies sectors where AI is already being utilized and where it is likely to expand. The findings reveal that AI use is most prevalent in software engineering, content creation, and data analysis, while occupations requiring physical manipulation of the environment show minimal AI adoption. Approximately 36% of occupations exhibit AI use in at least 25% of their tasks, but only 4% rely on AI for 75% or more of their tasks. The study also examines the relationship between AI usage, wage levels, and educational barriers, revealing that AI adoption peaks in mid-to-high-wage occupations, while both very high-paying professions (e.g., physicians) and low-wage jobs (e.g., restaurant workers) exhibit lower usage. Finally, it differentiates between AI as a tool for automation (43% of interactions) and as a tool for augmentation (57%), underscoring its dual role in enhancing efficiency and serving as a collabourative partner.

Taken together, these studies provide a multidimensional view of AI's impact on labour markets. While some methodologies (Brynjolfsson, Webb, Tolan, Felten) focus on structured linkages between AI capabilities and occupational requirements, others (Pizzinelli, Briggs and Kodnani, Eloundou, Handa et al.) analyze real-world AI usage and diffusion. This comprehensive approach enables a deeper understanding of which occupations are most exposed to AI and which have the highest potential for adaptation and complementarity with emerging technologies.

2.3 Data

To analyze how exposure to artificial intelligence affects labour supply elasticity and, consequently, monopsony power, we focus on Europe and provide empirical evidence for 26 European countries: Austria, Belgium, Cyprus, Czech Republic, Germany, Denmark, Estonia, Greece, Spain, France, Croatia, Hungary, Ireland, Italy, Lithuania, Luxembourg, Latvia, Netherlands, Poland, Portugal, Romania, Slovakia. This paper differs from most of the existing literature by integrating a key model widely used to estimate monopsony power in the labour market (Manning, 2003) with an established framework for measuring AI exposure in specific occupations (Webb, 2020), which we adapt into a time-variant version for this study.

Our unit of analysis is an occupation-region-year cell. Occupations are categorized according to the International Standard Classification of Occupations (ISCO) at the three-digit level. Each country is further disaggregated at the NUTS-1 regional classification level. The analysis covers the period from 2011 to 2020.

2.3.1 Labour market data

We use the EU Labour Force Survey (EU-LFS) annual microdata for the period 2011–2020, which provides detailed cross-country information on labour force composition. Our objective is to calculate four key variables necessary to estimate the elasticity of labour supply. All measures are aggregated at the ISCO-08 three-digit level, by year and NUTS-1 region.

Specifically, we measure the separation rate from employed to employed, defined as the growth rate of workers who actively search for and obtain a new job relative to the total population, and the separation rate from unemployed to employed, defined as the growth rate of unemployed individuals transitioning into employment relative to the total population. Additionally, we compute the overall recruitment rate, which is calculated as the growth rate of total employed individuals relative to the total population. These measures allow us to analyze labour market dynamics and assess the responsiveness of employment transitions to economic conditions. We then calculate wages for each occupation-region-year cell using the LFS, which provides income data in centiles, by computing the average of all observations within that specific cell. Regarding wage data, the Labour Force Survey (LFS) does not provide absolute wage values; it only provides decile information. Consequently, for each cell, we compute the mean wage percentile to indicate its relative position within the overall wage distribution.[1]

[1] For example, *Primary school and early childhood teachers* in the "Isole" region of Italy in 2011

2.3.2 AI exposure

To measure the exposure of occupations to artificial intelligence, we follow the method proposed by Webb (2020). In this paper, the author measures occupational exposure to technology by analyzing patents and job tasks. He extracts verb-noun pairs from patent titles, which describe technological capabilities, and from O*NET occupation task descriptions, which define job functions. Using dependency parsing and lexical categorization from the Wordnet (Fellbaum, 1998) database, he matches job tasks with technological capabilities. The exposure score for an occupation is computed as a weighted average of task-level scores, reflecting the frequency and importance of relevant patent-related verb-noun pairs. This method captures how closely artificial intelligence aligns with occupational tasks over time.

We further develop this measure by employing a time-variant version, which provides an AI exposure value for each year from 2011 to 2020. In addition to canonical verb-noun combinations, we incorporate action nominals—nouns derived from verbs that retain eventive meaning, such as *detection*, *management*, or *inspection*. These forms often appear in both patents and task descriptions as nominal references to actions without explicit verbs. To identify them, we rely on Wordnet to filter nouns that belong to semantic categories associated with actions (e.g., administrative or social processes), and then trace their connections to related verb forms. Although syntactically distinct, these nouns are semantically close to verbs and contribute meaningfully to describing technological capabilities and occupational functions. Their inclusion allows us to capture a broader range of task–technology alignments, particularly in contexts where actions are embedded in nominal structures rather than expressed through verb phrases.

2.3.3 Construction of our database

To empirically assess the impact of artificial intelligence on monopsony power, we merge labour market data with measures of technological exposure. This integration is done across the year and three-digit ISCO occupations dimensions. The AI exposure measure, derived from Google Patents, is originally classified under the Standard Occupational Classification (SOC) system. Since our employment microdata (EU-LFS) follows the ISCO classification, we align the two using crosswalks and correspondence tables from Hardy et al. (2018). These mappings are performed at the four-digit ISCO level before aggregating to three-digit occupations. When an ISCO occupation corresponds to multiple SOC occupations, we compute the average

are positioned at the 34th percentile.

AI exposure score across ISCO categories. To obtain three-digit ISCO occupations, we drop the last digit from the four-digit classification and take the mean across matching occupations. Given that the AI exposure scores were constructed for the US labour market, we rely on the implicit assumption that occupational tasks are equivalent across US and EU occupations. This assumption is reasonable because, at an aggregate level, the tasks that make up a given occupation should not vary significantly across different labour markets. The frequency of our data sources varies, but for consistency, we use annual labour force composition data from the EU-LFS. Unlike Webb, we develop a time-variant AI exposure measure that identifies AI capabilities using patents.

2.4 Empirical methodology

To calculate the overall wage elasticity of labour supply, we follow the approach of Bachmann et al. (2022), which builds on the work of Burdett and Mortensen (1998) and Manning (2003). In short, the law of motion for labour supply to the firm can be expressed as:

$$L_t = R_e(w_t) + R_n(w_t) + [1 - S_e(w_t) - S_n(w_t)]L_{t-1} \tag{2.1}$$

Where R represents the recruitment rate and S represents the separation rate, while the subscript e denotes transitions to employment and n denotes transitions to nonemployment. In the steady state, the recruitment rate is equal to the separation rate, so:

$$L(w) = \frac{R_e(w) + R_n(w)}{S_e(w) + S_n(w)} \tag{2.2}$$

from this, by making the appropriate rearrangements[2]:

$$\epsilon_{LW} = -\big(\theta_Z\big)\,\epsilon^n_{SW} - (1-\theta_Z)\,\epsilon^e_{SW} + \big(\theta_N\big)\,\epsilon^n_{RW} + (1-\theta_N)\,\epsilon^e_{RW} \tag{2.3}$$

This implies that the long-term elasticity of labour supply to the individual firm, ϵ_{LW}, is the difference of a weighted average between the wage elasticities of recruitment from employment (ε^e_{RW}) and from nonemployment (ε^n_{RW}), and the wage elasticities of the separation rates to employment (ε^e_{SW}) and to nonemployment (ε^n_{SW}), where the weights are, unlike the aforementioned article, given by θ_N, the share of recruits hired from non-employment, and θ_Z, the share of separations to

[2] See Bachmann et al. (2022) for further details.

non-employment.

Following Hirsch (2010), but taking into account that our weight has a different meaning, we arrive at the following equation:

$$\epsilon_{LW} = -\big(1 + \theta_N\big)\,\epsilon^{n}_{SW} - (1 - \theta_N)\,\epsilon^{e}_{SW} - \epsilon_{\theta,W} \tag{2.4}$$

The term $\epsilon_{\theta W}$ represents the wage elasticity of the share of recruits hired from nonemployment, while θ_N denotes the overall share of hires coming from nonemployment. The four components that determine the wage elasticity of labour supply to the firm are the wage elasticity of the separation rate to employment, computed as the growth rate of individuals who were searching for a job in the previous period, the wage elasticity of the separation rate to nonemployment, measured as the growth rate of unemployed individuals in the sample, the wage elasticity of the share of recruits from employment, calculated as the growth rate of employed individuals, and the overall share of recruits originating from nonemployment. This strategy allows us to isolate different effects driven by wage variations. For separation rates, for example, a higher separation rate elasticity implies that workers are more likely to leave their jobs in response to wage changes. This reduces firms' monopsony power, as they cannot arbitrarily lower wages without risking worker turnover. These separation rate elasticities are weighted by θ_N, the share of hires coming from nonemployment, to account for their relative impact on the overall wage elasticity of labour supply. Conversely, a high elasticity of the hiring rate indicates that firms are better able to attract workers by offering higher wages, drawing them away from competing firms. The overall effect on labour supply will depend on the interaction among these factors, which provides a comprehensive measure of market power: as labour supply elasticity decreases, monopsony power increases because workers are less willing to leave their jobs. Consequently, firms gain greater bargaining power and are able to keep wages lower.

To estimate the components of Equation 3, we proceed as follows. As a first step, we estimate a linear regression model in which the dependent variable is s^{ρ}, a dove ρ can take the value n or e, which imply, respectively, a transition toward nonemployment or unemployment. Recall that we operate at the occupation, region, and year level.

$$s^{\rho}(or,t) = \beta_0 + \beta_1\,W(or,t) + \mathbf{X}(or,t)\,\boldsymbol{\gamma} + \alpha_{\text{year}(t)} + \alpha_{\text{country}(or)} + \varepsilon(or,t) \tag{2.5}$$

Our main explanatory variable is W is the wage expressed in centiles. We also include a set of controls capturing worker characteristics (age, gender, educational attainment), firm size categories, and sector. We include year and country fixed effects to control for unobserved time-varying shocks common to all countries and for time-invariant differences across countries. We cluster the standard errors by occupation and region to allow for arbitrary correlation of the errors within each occupation–region cell. The coefficient β_1 is therefore interpreted as the elasticity considering that both the dependent variable and the independent variable are expressed in percentage terms.

Similarly, for the wage elasticity of the share of recruits hired from nonemployment, we estimate the following equation:

$$r^{\mathrm{n}}(or,t) = \beta_0 + \beta_1\, W(or,t) + \mathbf{X}(or,t)\,\boldsymbol{\gamma} + \alpha_{\mathrm{year}(t)} + \alpha_{\mathrm{country}(or)} + \varepsilon(or,t) \tag{2.6}$$

where $r^{\mathrm{n}}(or,t)$ is interpreted as the recruitment rate from nonemployment. In this case as well, the coefficient β_1 is interpreted as an elasticity.

As a second step, we must study the impact that artificial intelligence has on this degree of monopsony. As an indicator of AI exposure, we use the Webb index described in Section 3, inserting it into the equation to be estimated. This interaction term, which is constructed as time-varying, allows us to assess how the increased introduction of artificial intelligence has impacted monopsony. To do this, in both Equation (5) and Equation (6) we add both the index measuring AI exposure and an interaction between it and W, the salary expressed in centiles. The equations we estimate then become as follows:

$$\begin{aligned} s^{\rho}(or,t) = \beta_0 &+ \beta_1\, W(or,t) + \beta_2\, AI(or,t) \\ &+ \beta_3 \Big[W(or,t) \times AI(or,t)\Big] + \mathbf{X}(or,t)\,\boldsymbol{\gamma} \\ &+ \alpha_{\mathrm{year}}(t) + \alpha_{\mathrm{country}}(or) + \varepsilon(or,t). \end{aligned} \tag{2.7}$$

$$\begin{aligned} r^{\mathrm{n}}(or,t) = \beta_0 &+ \beta_1\, W(or,t) + \beta_2\, AI(or,t) \\ &+ \beta_3 \Big[W(or,t) \times AI(or,t)\Big] + \mathbf{X}(or,t)\,\boldsymbol{\gamma} \\ &+ \alpha_{\mathrm{year}}(t) + \alpha_{\mathrm{country}}(or) + \varepsilon(or,t). \end{aligned} \tag{2.8}$$

where the coefficient β_1 is now also influenced by β_3, i.e. the coefficient of the interaction term between wages and AI exposure.

2.5 Descriptive Evidence

Table 1 and 2 provides some details about the reference sample. We recall that the 84,126 cells composing the sample are organized at the employment-region-year level[3] and the occupations are classified at the three-digit level. The statistics suggest an average salary lower than the median salary[4]. The share of separations from employment is 3,3%, which implies that every year, on average, 3,3% of the total working population at the regional level moves from one job to another, while 6,52% becomes unemployed (this is, in fact, the value of the share of separations to non-employment). At the same time, 7,7% of the nonemployed are hired from nonemployment. Table 1 and 2 are further supplemented with additional information that, on average for each occupation-region-year cell, provides the mean age, the percentage of women, the percentage of educational attainment achieved, the proportion of firms of a certain size, and the number of firms for each 1-digit NACE code. Tables 3 and 4, by contrast, give an idea of how the AI exposure index is structured, showing the detailed distribution of its measure by occupation. In particular, Table 3 provides the ranking of the five occupations most exposed to artificial intelligence for each year[5]. As can be seen, there are two occupations that stand out as most exposed, namely "University and higher education teachers" and "Mathematicians, Actuaries, and Statisticians. This aligns with the idea that artificial intelligence can handle not only repetitive tasks, but also those requiring higher cognitive skills (Eloundou et al., 2023; Felten et al., 2021; Georgieff and Hyee, 2021; Tolan et al., 2021). Table 4, by enumerating the five least exposed occupations for each year, also indicates that artificial intelligence has a smaller impact on professions with a stronger social component (Fernández-Macías and Bisello, 2022) or on low-skill service occupations (Autor and Dorn, 2013), namely jobs that involve assisting or caring for others. In Appendix B, there is also a table where the ranking is produced overall (i.e., without splitting by year) along with additional statistics to help understand the temporal evolution of AI exposure.

[3]For example, Mobile plant operators in Italy, specifically in the "South" region in the year 2019.

[4]We recall that we only have information on salaries in centiles.

[5]We highlight the year because, by construction, it is crucial to understand how AI exposure varies over time. Indeed, the index is built starting from how patents filed in a specific year affect the tasks that make up an occupation.

Table 2.5.1: **Descriptive Statistics (Part I)**

Variable	**Mean**	**SD**	**Min**	**Max**
AI exposure	0.102	0.074	0.008	0.97
Wage	4.38	2.141	0.0008	10
Share of separations to employment	0.0334	0.0634	0	1
Share of separations to nonemployment	0.0652	0.1187	0	1
Share of recruitments from nonemployment	0.077	0.0951	0	1
Age (average)	44.34	19.51	15	89
Female (avg)	0.4098	0.3036	0	1
Share of non-primary education	0.0051	0.0312	0	1
Share of primary education	0.0423	0.1127	0	1
Share of lower secondary education	0.1348	0.1743	0	1
Share of upper secondary education	0.4219	0.2839	0	1
Share of post-secondary non-tertiary	0.0440	0.0967	0	1
Share of short-cycle tertiary	0.1302	0.2366	0	1
Share of bachelor's degree	0.0882	0.1593	0	1
Share of master's degree	0.1181	0.2234	0	1
Share of doctoral degree	0.0105	0.0604	0	1
Share in firms of 0-9 employees	0.2155	0.1961	0	1
Share in firms of 10-19 employees	0.1134	0.1144	0	1
Share in firms of 20-49 employees	0.1293	0.1230	0	1
Share in firms of 50+ employees	0.3670	0.2588	0	1

Source:: Authors' calculations.
Notes: The number of observations is 84,132. These observations are organized at the occupation-region-year level. Occupations are classified at the 3-digit level based on ISCO08, while regions follow the NUTS 1-digit classification.

Table 2.5.2: **Descriptive Statistics (Part II)**

Variable	**Mean**	**SD**	**Min**	**Max**
Share in agriculture, forestry, fishing	0.0478	0.1854	0	1
Share in mining and quarrying	0.0054	0.0442	0	1
Share in manufacturing	0.1865	0.2843	0	1
Share in electricity, gas, steam	0.0092	0.0422	0	1
Share in water supply, sewerage	0.0132	0.0702	0	1
Share in construction	0.0597	0.1619	0	1
Share in wholesale, retail trade	0.1054	0.2041	0	1
Share in transportation	0.0624	0.1731	0	1
Share in accommodation, food services	0.0477	0.1684	0	1
Share in publishing, media	0.0295	0.1031	0	1
Share in IT, telecom, consulting	0.0245	0.0964	0	1
Share in financial activities	0.0075	0.0443	0	1
Share in real estate	0.0642	0.1538	0	1
Share in professional services	0.0370	0.0954	0	1
Share in administrative and support services	0.0771	0.1684	0	1
Share in public administration	0.0710	0.2019	0	1
Share in education	0.0805	0.1977	0	1
Share in health and social work	0.0304	0.1111	0	1
Share in arts, sports, recreation	0.0324	0.1140	0	1
Share in other services	0.0055	0.0376	0	1
Share in household activities	0.0014	0.0182	0	1

Source: Authors' calculations.
Notes: The number of observations is 84,132. Observations are at the occupation–region–year level. Occupations are classified at the 3-digit ISCO-08 level, and regions follow the NUTS-1 classification.

Table 2.5.3: **AI Exposure of the Top Five Occupations by Year**

Rank	Occupation	Year	Score
1	University and higher education teachers (231)	2018	0.9713
2	University and higher education teachers (231)	2019	0.9272
3	University and higher education teachers (231)	2017	0.8341
4	Mathematicians, Actuaries and Statisticians (212)	2018	0.6137
5	Mathematicians, Actuaries and Statisticians (212)	2017	0.5871

Source: Authors' calculations.
Notes: Top five 3-digit occupations by AI exposure. ISCO-08 occupation codes are reported in parentheses.

Table 2.5.4: **AI exposure of bottom five occupations for year**

Rank	Occupation	Year	Score
1	Street Vendors (952)	2015	0.00837
2	Information and Communications Technology Service managers (133)	2014	0.0088
3	Street and Market Salespersons (521)	2015	0.009803
4	Information and Communications Technology Service managers (133)	2013	0.0101
5	Food Preparation Assistants (941)	2012	0.01156

Source: Authors' calculations.
Notes: 3-digit top five occupations by AI exposure (ISCO 08 classification in brackets).

2.6 Results

Table 5 presents the initial results, where the two columns report, respectively, the estimates without control variables and with control variables. The decision to include the results for the control variables is made to highlight the changes once certain variables are accounted for. Specifically, column 2 includes control variables such as age, education level, firm size, gender, and economic sector. The sample consists of 84,132 observations and is constructed using data from the labour Force Survey. Each observation corresponds to a specific ISCO 3-digit occupation, a given year ranging from 2011 to 2020, a country, and a NUTS 1-digit code. Recall that our wage data are expressed in terms of the wage percentile distribution. Specifically, for each cell, we report the average position of wages within the 1 to 100 percentile range. Consequently, the elasticity we analyze does not capture the effect of a 1% increase in wages in absolute terms, but rather the effect of moving up by one percentile—for example, shifting from the 49th to the 50th percentile. This approach offers a relative measure of wage changes within the overall distribution, providing a nuanced understanding of labour market dynamics. As shown in the first column, the elasticity coefficient is negative for all three selected variables, although it is close to zero for both separation rates to non-employment and hiring probability from non-employment. Consequently, as wages increase (or decrease), separation rate to employment and non-employment decrease (increase) and similarly, the recruitment rate from non-employment. The highest level of elasticity is observed in the component related to the separation rate to employment. In particular, a 1% increase in the percentile income distribution is associated with an approximately 4.5% reduction in the number of individuals within a given region and occupation who change jobs each year. The negative coefficient on wages for the separation rate to employment suggests that higher wages are associated with a lower probability of workers leaving their current employment for opportunities elsewhere. This finding is consistent with classic efficiency-wage theories (Katz, 1986; Shapiro and Stiglitz, 1984) and job-search models (Mortensen, 1970) , suggesting that higher wages serve as an effective retention mechanism since employees are less inclined to leave when their compensation is competitive in the market. Similarly, the negative relationship between an increase of one wage percentile and the separation rate to non-employment indicates that higher wages discourage transitions out of the labour market. In economic terms, as wages rise, the opportunity cost of leaving employment increases, thereby reducing the incentive for workers to exit and become non-employed or inactive. Furthermore, the negative coefficient observed for the hiring probability from non-employment indicates that as we move higher in the wage percentile distribution, the likelihood of hiring from the available pool decreases. This can be interpreted in several ways. On one hand, higher wages may prompt firms to revise their recruitment strategies by imposing stricter hiring criteria or increasingly relying on internal candidates rather than attracting external applicants (Berger et al., 2022). On the other hand, elevated wages might reduce the pool of candidates who are both willing and qualified to work at those levels, as the opportunity cost of changing jobs rises (Song et al., 2019). When considering the model that incorporates control variables such as education, sociodemographic characteristics, firm size, and economic sector, the results show some differences. The elasticity of separation toward employment remains negative and statistically significant, with a one-percentile increase in the wage distribution leading to a 3.62% reduction in the number of workers switching employers. In contrast, the effect on separations toward non-employment becomes negligible and statistically insignificant after the inclusion of controls. This indicates that the negative relationship observed in the uncontrolled model for separation toward non-employment likely resulted from omitted variable bias, and that including these controls provides a more accurate picture of the underlying variability. Finally, the elasticity for the recruitment rate from non-employment remains similar, decreasing slightly by 0.03. When considering the overall elasticity, we find a positive value of 3.4%. As discussed above, the main contributing component appears to be the separation rate to employment

Table 6 shows the results when wages are interacted with an occupation's exposure to artificial intelligence. This specification tests whether wages impact labour market transitions differently according to the level of AI exposure. It also assesses how occupational AI exposure influences the responsiveness of separation and hiring outcomes to changes in wage percentiles. The outcomes for separation rates and hiring probabilities remain similar to those observed in previous models while accounting for differences in AI exposure across occupations. The effect of the AI exposure index captures, on average, how occupations more heavily exposed to artificial intelligence may differ in their separation patterns. As explained in Section 3, an interaction term between wages and AI exposure has also been included.

Similarly to the Table 5, we conduct two separate regressions, one without controls and one with controls. The wage coefficients remain similar to those in the previous table, differing only slightly in magnitude. For the separation rate toward employment in the model without controls, the AI exposure index has a significant effect corresponding to a 26% reduction in job transitions. This indicates that as AI exposure increases, transitions from one job to another decline markedly. When controls are added, the coefficient on the AI exposure index decreases substantially but remains negative and statistically significant. One possible explanation for the negative coefficient is that occupations with higher AI exposure undergo changes in task composition. (Acemoglu and Restrepo, 2018) and (Autor, 2022) show that such modifications can lead to shifts in job assignments within the same occupational category because firms adopt AI to varying degrees. Some firms substantially alter the tasks associated with an occupation, while others do not, or even trigger a reorganization of tasks that results in shifts from one occupation to another (Acemoglu, 2024). The coefficient on the interaction term is positive and of smaller magnitude, suggesting that higher occupational AI exposure reduces the elasticity of the separation rate toward employment , thereby reducing overall mobility. In other words, workers in occupations with greater AI exposure tend to remain in their jobs, thereby reinforcing monopsony power. For separations toward non-employment and hiring probability, the regression coefficients exhibit only minor changes in magnitude while retaining their sign as reported in Table 5, regardless of whether control variables are included. In the regression with controls, the coefficients for AI exposure are not statistically significant for both variables, as is the interaction term for the separation rate toward non-employment. In contrast, the interaction term for the hiring probability from non-employment is statistically significant, though its magnitude is negligible and close to zero. This stage of regressions also seems to confirm that the main impact on elasticity comes from the separation rate to non-employment component, while the effect of AI is present but does not appear to be decisive. This is further confirmed by the fact that total elasticity, which now also depends on the interaction term — and thus on AI exposure — does not vary substantially across exposure levels: indeed, elasticity ranges from 1.91% to 1.81% depending on the degree of exposure.

Once the individual components have been analyzed, we now examine the overall labour supply as calculated in Equation 1. Figure 1 shows the evolution of this variable over time. Three distinct curves are displayed, corresponding to a regression specification that includes an interaction term with AI exposure. These curves represent an exposure value of 0, the mean value (0.102 in our sample), and a high exposure value (0.5). Our analysis is based on the labour market framework established by Manning (2003)and Burdett and Mortensen (1998), which posits that workers actively pursue higher wages and that their mobility depends on the wage differential. Monopsony power is measured by the wage elasticity of labour supply to the firm; a lower elasticity indicates greater firm power because workers are less responsive to wage changes, whereas a higher elasticity suggests lower monopsony power as wage increases prompt higher worker mobility.

From Figure 1, a sharp decline in labour supply elasticity is observed between 2011 and 2013, after which the value stabilizes at levels close to 1 until 2016. From that point onward, the curves experience a further decline, progressively converging toward 0. The three curves demonstrate that interacting wages with AI exposure at the two intensities (0.102 and 0.500) produces only slight differences compared to the case without interaction (AI exposure = 0). This suggests that, on average, the degree of AI exposure does not drastically alter labour supply elasticity. However, during the early period (2011–2013), the curve associated with higher AI exposure (green) deviates slightly more than the others, indicating a potentially stronger AI-related effect at the beginning of the observation period. These descriptive results, while not implying causality, open the way for future analyses aimed at exploring the causal dynamics underlying the observed patterns.

Thus, the graph indicates that between 2011 and 2020, monopsony power has increased significantly, as evidenced by the decline in elasticity. Nonetheless, the impact of artificial intelligence does not appear to be very pronounced: the three curves nearly overlap, with the curve corresponding to higher AI exposure (green) being only marginally lower than the others. This finding implies that even the occupations most exposed to AI follow a consolidated pattern driven mainly by structural factors (Elder et al., 2023)).

We further refine the analysis by subdividing observations each year into both education and wage terciles, with Figure 2 and 3 illustrating the evolution of labour supply elasticity across these two dimensions.

Figure 2 presents labour supply elasticity calculated by dividing the data, on a yearly basis, into education terciles. Figure 3 follows the same approach using wage terciles.

If artificial intelligence primarily impacts non-routine, cognitive tasks (Felten et al., 2021; Tolan et al., 2021), one would expect that professions characterized by higher education and wages would display a more pronounced downward trend in elasticity compared to the other groups. However, the results are ambiguous. Between 2011 and 2013, the elasticity for high-education occupations is even negative. This result is not surprising, as we should recall that what is being analysed here is a macroeconomic elasticity rather than a firm-level one. . This negative elasticity may result from short-term constraints such as strong contractual agreements, institutional rigidities, measurement issues, or limited variation within certain subgroups. From 2014 onward, elasticity is positive and follows a slightly decreasing trend, indicating that even those with higher educational profiles are experiencing greater monopsony power (Goolsbee and Syverson, 2021) . In contrast, the profile for low-education occupations is much clearer: these jobs have experienced a marked decline in elasticity from 2011 to 2020, with elasticity values nearly approaching zero. This confirms some observations made by Ashenfelter and Jurajda (2022) that low-skilled occupations are more vulnerable to monopsony power. Meanwhile, the pattern for medium-education occupations is less defined, with some periods (between 2012 and 2015) showing an increase in elasticity, followed by declines until 2019 and a subsequent rise in 2020. For the economic dimension, only the curves corresponding to average AI exposure are reported. This decision is based on the observation that profiles for lower or higher levels of AI exposure closely mirror the average, implying that AI exposure itself does not alter the evolution of labour supply elasticity. From this perspective, low-skill occupations still appear to be the most affected. It should be recalled, however, that the analysis refers to the years between 2011 and 2020.

Figure 3, which categorizes observations by wage terciles, provides a particularly distinct picture. The medium-wage curve consistently lies above the high-wage curve, which in turn is above the low-wage curve. Notably, the low-wage group shows negative elasticity throughout the period. In this context, negative elasticity indicates that as wages increase, rather than encouraging worker mobility, they are associated with even lower responsiveness. This may be due to the scarcity of alternative opportunities or market rigidities that keep low-wage workers effectively "locked in" to their positions. Such a scenario is consistent with theories of monopsony power, wherein firms benefit when workers exhibit low mobility in response to wage increases.

Overall, the evidence indicates that medium-wage occupations, typically comprising routine jobs (Autor, 2015; Autor and Dorn, 2013), are less susceptible to rising monopsony power, thereby supporting the job polarization hypothesis (Goos et al., 2014) . In contrast, high-wage professions exhibit a modest downward trend in elasticity over time. For low-wage occupations, the pattern is even more striking: elasticity declines up to 2014, suggesting an increase in monopsony power, but begins to rise thereafter. However, as these values remain negative, the labour market continues to display low worker responsiveness, likely due to persistent structural frictions and limited alternative employment opportunities.

In summary, although the observed patterns across different wage and education groups provide some insight into the evolution of labour supply elasticity, the impact of AI appears ambiguous. The modest differences across AI exposure levels suggest that other structural factors are primarily driving labour supply responsiveness. In the appendix, granular results are presented, showing the evolution of elasticity for each of the nine ISCO groups. These findings further highlight the variability in labour supply responsiveness across occupations and offer additional insights into the distribution of monopsony power.

threeparttable booktabs array

Table 2.6.1: **Labour Supply Elasticity**

	Without controls	**With controls**
Separation rate to employment		
Wage (ε^e_{sw})	-4.53***	-3.62***
	(1.18)	(0.99)
Observations	84,312	84,132
Separation rate to non-employment		
Wage (ε^n_{sw})	-0.02*	0.003
	(0.013)	(0.016)
Observations	84,132	84,132
Hiring probability from non-employment		
Wage ($\varepsilon_{\theta w}$)	-0.09***	-0.06***
	(0.02)	(0.017)
Observations	84,132	84,132
Share of hires from employment (θ)	0.077	0.077
Firm-level labour supply elasticity (ε_{LW})	4.30	3.40

Source: Authors' calculations.
Notes: Clustered standard errors at the ISCO-08 3-digit and NUTS-1 level are reported in parentheses. Controls include age, education, gender, firm size, and economic sector.
Significance levels: $^*p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

Table 2.6.2: **Labour Supply Elasticity by AI Exposure**

	Without controls	**With controls**
Separation rate to employment		
Wage (ε^{e}_{sw})	-5.94***	-2.06*
	(1.88)	(1.06)
AI exposure	-26.24**	-6.74*
	(6.79)	(3.55)
Wage $\times$ AI	0.37***	0.11**
	(0.02)	(0.05)
Observations	84,132	84,132
Separation rate to non-employment		
Wage (ε^{n}_{sw})	-0.04**	0.055***
	(0.02)	(0.001)
AI exposure	-0.125	0.09
	(0.078)	(0.06)
Wage $\times$ AI	0.002	-0.001
	(0.011)	(0.001)
Observations	84,132	84,132
Hiring probability from non-employment		
Wage ($\varepsilon_{\theta w}$)	-0.12***	-0.07**
	(0.03)	(0.019)
AI exposure	-0.36**	-0.12
	(0.144)	(0.079)
Wage $\times$ AI	0.01**	0.002*
	(0.02)	(0.11)
Observations	84,132	84,132
Share of hires from employment (θ)	0.077	0.077
Firm-level labour supply elasticity (ε_{LW}), min	5.65	1.91
Firm-level labour supply elasticity (ε_{LW}), max	5.61	1.81

Source: Authors' calculations.
Notes: Clustered standard errors at the ISCO-08 3-digit and NUTS-1 level are reported in parentheses. Controls include age, education, gender, firm size, and economic sector.
Significance levels: *$p < 0.05$; **$p < 0.01$; ***$p < 0.001$.

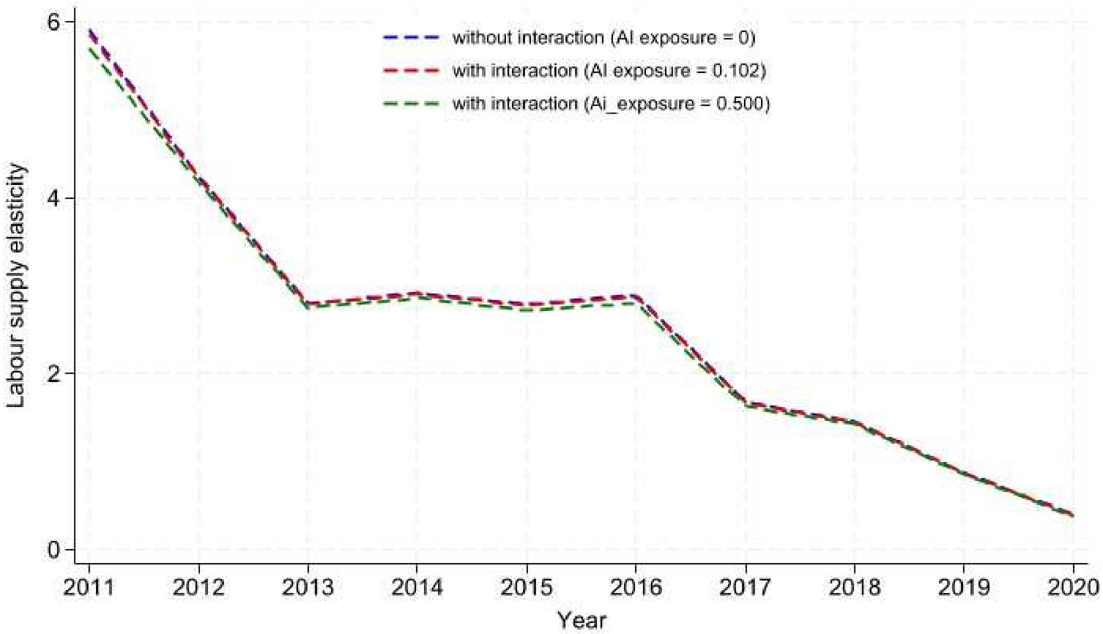

Figure 2.6.1: **Yearly Labour Supply Elasticities**

Source: Authors calculation.

Notes: The estimates are derived from the same specification as in Table 6. The three curves represent: the blue one showing elasticity with zero AI exposure, the red one with a medium AI exposure (equal to 0.102), and the green one with high AI exposure (equal to 0.500).

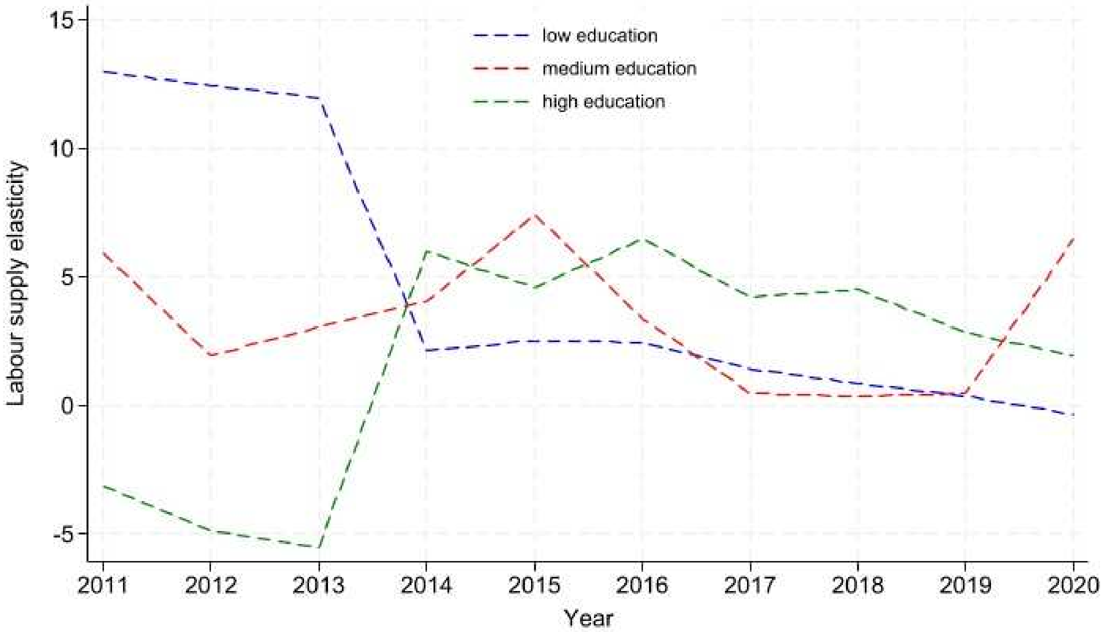

Figure 2.6.2: **Yearly Labour Supply Elasticities by educational terciles**

Source: Authors calculation.

Notes: The estimates are derived from the same specification as in Table $\tilde{6}$, but the observations for each year have been divided into terciles based on the education level values. The three curves represent: the blue one showing elasticity with zero AI exposure, the red one with a medium AI exposure (equal to 0.102), and the green one with high AI exposure (equal to 0.500).

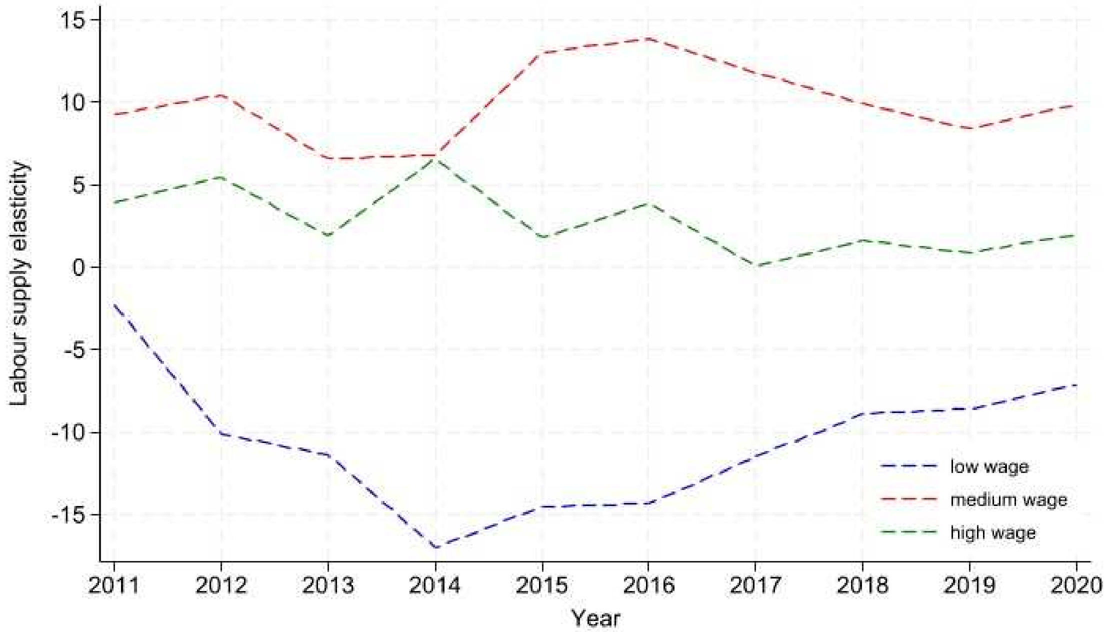

Figure 2.6.3: **Yearly Labour Supply Elasticities by wage terciles**

Source: Authors calculation.

Notes: The estimates are derived from the same specification as in Table 6 but the observations for each year have been divided into terciles based on the wage level value. The three curves represent: the blue one showing elasticity with zero AI exposure, the red one with a medium AI exposure (equal to 0.102), and the green one with high AI exposure (equal to 0.500).

2.7 Conclusions

This paper provides compelling evidence that over the period from 2011 to 2020, labor market dynamics in the 26 European countries under examination have been characterized by a notable decline in the elasticity of labor supply with respect to wages, a trend that directly implicates an increase in firm monopsony power. Among the most salient results, the analysis reveals that low-wage occupations are particularly vulnerable; these groups exhibit the lowest responsiveness to wage changes, a pattern consistent with established theories of job polarization and structural market rigidities. In contrast, medium-wage occupations demonstrate comparatively higher mobility, while high-wage segments show only modest shifts in elasticity over time. From the educational perspective, the results show that low-education occupations are once again those most affected by monopsony power over the years. Considering that AI technologies are generally thought to primarily target cognitive tasks and occupations, these results seem to confirm that the years between 2011 and 2020 are still too early to observe any substantial effects of AI.

Moreover, when examining the interaction between wage dynamics and artificial intelligence (AI) exposure, our findings indicate that while AI-related factors do play a role, their overall impact is relatively minor in altering the underlying patterns. This suggests that, at least for the period studied, enduring structural factors are the primary drivers behind the observed changes in labor supply elasticity. The modest effect of AI exposure, despite significant technological developments in recent years, raises important questions about the pace at which emerging technologies are integrated into productive processes and their capacity to reshape traditional employment dynamics.

Taken together, these results underscore the critical importance of understanding labor market segmentation and the differential impacts of monopsony power. They also highlight the need for continued research in light of the rapidly evolving role of technology in the workplace. As AI and other innovations become more pervasive, future studies should reexamine these relationships with more recent data, to assess whether the trends identified here persist, intensify, or undergo a structural break.

2.8 Acknowledgements

We thank Michael Webb for valuable advice in writing the code to calculate the AI score. All remaining errors are our own.

Bibliography

Acemoglu, D. (2021). Harms of AI. NBER Working Paper 29247, National Bureau of Economic Research.

Acemoglu, D. (2024). The simple macroeconomics of AI. NBER Working Paper 32487, National Bureau of Economic Research.

Acemoglu, D. and Restrepo, P. (2018). Artificial Intelligence, automation, and work. In Agrawal, A., Gans, J., and Goldfarb, A., editors, *The Economics of Artificial Intelligence: An Agenda*, pages 197–236. National Bureau of Economic Research.

Agrawal, A., Gans, J., and Goldfarb, A. (2019). *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Ashenfelter, O., Farber, H., and Ransom, M. (2010). Labor market monopsony. *Journal of Labor Economics*, 28(2):203–210.

Ashenfelter, O. C., Card, D., Farber, H. S., and Ransom, M. (2021). Monopsony in the labor market: New empirical results and new public policies. NBER Working Paper 29522, National Bureau of Economic Research.

Ashenfelter, O. C. and Jurajda, (2022). Minimum wages, wages, and price pass-through: The case of McDonald's restaurants. *Journal of Labor Economics*, 40(S1):179–201.

Autor, D. H. (2015). Why are there still so many jobs? the history and future of workplace automation. *Journal of Economic Perspectives*, 29(3):3–30.

Autor, D. H. (2022). The labor market impacts of technological change: From unbridled enthusiasm to qualified optimism to vast uncertainty. NBER Working Paper 30074, National Bureau of Economic Research.

Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review*, 103(5):1553–1597.

Autor, D. H., Dorn, D., Katz, L. F., Patterson, C., and Van Reenen, J. (2020). The fall of the labor share and the rise of superstar firms. *The Quarterly Journal of Economics*, 135(2):645–709.

Autor, D. H., Dube, A., and McGrew, A. (2023). The unexpected compression: Competition at work in the low wage labor market. NBER Working Paper 31010, National Bureau of Economic Research.

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

Azar, J., Chugunova, M., Keller, K., and Samila, S. (2023). Monopsony and automation. Research Paper 23-21, Max Planck Institute for Innovation & Competition.

Azar, J., Marinescu, I., and Steinbaum, M. (2019). Measuring labor market power two ways. 109:317–321.

Azar, J. and Marinescu, I. E. (2024). Monopsony power in the labor market. SSRN Working Paper 5041452, Social Science Research Network.

Bachmann, R., Demir, G., and Frings, H. (2022). Labor market polarization, job tasks, and monopsony power. *Journal of Human Resources*, 57(S):11–49.

Bellocchi, A., Marin, G., and Travaglini, G. (2023). The labor share puzzle: Empirical evidence for European countries. *Economic Modelling*, 124(C):106327.

Benmelech, E., Bergman, N., and Kim, H. (2018). Strong employers and weak employees: How does employer concentration affect wages? NBER Working Paper 24307, National Bureau of Economic Research.

Berger, D., Herkenhoff, K., and Mongey, S. (2022). Labor market power. *American Economic Review*, 112(4):1147–1193.

Bergholt, D., Furlanetto, F., and Maffei-Faccioli, N. (2022). The decline of the labor share: New empirical evidence. *American Economic Journal: Macroeconomics*, 14(3):163–198.

Boal, W. M. and Ransom, M. R. (1997). Monopsony in the labor market. *Journal of Economic Literature*, 35(1):86–112.

Briggs, J. and Kodnani, D. (2023). The potentially large effects of Artificial Intelligence on economic growth. Global Economics Analyst Report 26 March 2023, Goldman Sachs.

Brooks, W. J., Kaboski, J. P., Li, Y. A., and Qian, W. (2021). Exploitation of labor? classical monopsony power and labor's share. *Journal of Development Economics*, 150(C):102627.

Brynjolfsson, E., Mitchell, T., and Rock, D. (2018). What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings*, 108:43–47.

Burdett, K. and Mortensen, D. T. (1998). Wage differentials, employer size, and unemployment. *International Economic Review*, 39(2):257–273.

Cairó, I. and Sim, J. W. (2020). Market power, inequality, and financial instability. Finance and Economics Discussion Series 2020-057, Board of Governors of the Federal Reserve System.

Caldwell, S. and Danieli, O. (2024). Outside options in the labour market. *The Review of Economic Studies*, 91(6):3286–3315.

Calligaris, S., Criscuolo, C., Lyon, J. D., Greppi, A., and Pallanch, O. (2025). Concentration and business dynamics in product markets. OECD Science, Technology and Industry Working Papers 2025/04, OECD Publishing.

Calvano, E., Calzolari, G., Denicolò, V., and Pastorello, S. (2023). Artificial intelligence, algorithmic recommendations and competition. CEPR Discussion Paper 18176, Centre for Economic Policy Research.

Card, D. (2022). Who set your wage? 112(4):1075–1090.

Chamberlin, E. H. (1973). *The Theory of Monopolistic Competition: A Re-orientation of the Theory of Value*. Harvard University Press, Cambridge, MA, 8 edition.

Chen, D. L. and Horton, J. J. (2016). Are online labor markets spot markets for tasks?: A field experiment on the behavioral response to wage cuts. TSE Working Paper 16-675, Toulouse School of Economics.

Covarrubias, M., Gutiérrez, G., and Philippon, T. (2020). From good to bad concentration? US industries over the past 30 years. In *NBER Macroeconomics Annual*, volume 34 of *NBER Macroeconomics Annual*, pages 1–46. University of Chicago Press.

Diamond, P. A. (1971). A model of price adjustment. *Journal of Economic Theory*, 3(2):156–168.

Dixit, A. K. and Stiglitz, J. E. (1977). Monopolistic competition and optimum product diversity. *The American Economic Review*, 67(3):297–308.

Dube, A., Jacobs, J., Naidu, S., and Suri, S. (2020). Monopsony in online labor markets. *American Economic Review: Insights*, 2(1):33–46.

Dube, A., Naidu, S., and Reich, A. D. (2022). Power and dignity in the low-wage labor market: Theory and evidence from wal-mart workers. NBER Working Paper 30441, National Bureau of Economic Research.

Duch-Brown, N., Gomez-Herrera, E., Mueller-Langer, F., and Tolan, S. (2021). Market power and artificial intelligence work on online labour markets. Research Paper 22-01, Max Planck Institute for Innovation & Competition.

Duch-Brown, N., Grzybowski, L., Romahn, A., and Verboven, F. (2023). Evaluating the impact of online market integration—evidence from the EU portable PC market. *American Economic Journal: Microeconomics*, 15(4):268–305.

Elder, T. E., Haider, S. J., and Orr, C. (2023). The evolution of the wage elasticity of labor supply over time. IZA Discussion Paper 16393, Institute of Labor Economics (IZA).

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 2303.10130, arXiv.

Eurofound , Parent-Thirion, A., Mackova, B., Ivaškaitė-Tamošiūnė, V., and Cabrita, J. (2023a). *Job quality of COVID-19 pandemic essential workers*. Publications Office of the European Union.

Eurofound , Weber, T., Adăscăliței, D., Isusi, I., Durán, J., Hahne, A. S., Spielhofer, T., and Hadjivassiliou, K. (2023b). *Measures to tackle labour shortages: lessons for future policy*. Publications Office of the European Union.

Fellbaum, C. (1998). *WordNet 1. 6: An electronic lexical database*. Bradford Books, Cambridge, MA.

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.

Fernández-Macías, E. and Bisello, M. (2022). A comprehensive taxonomy of tasks for assessing the impact of new technologies on work. *Social Indicators Research: An International and Interdisciplinary Journal for Quality-of-Life Measurement*, 159(2):821–841.

Gans, J. S. (2024). Market power in artificial intelligence. SSRN Working Paper 4763294, Social Science Research Network.

Georgieff, A. and Hyee, R. (2021). Artificial Intelligence and employment: New cross-country evidence. OECD Social, Employment and Migration Working Papers 265, OECD Publishing.

Goolsbee, A. and Syverson, C. (2021). Monopsony power in higher education: A tale of two tracks. In Hurst, E. and Kahn, L. B., editors, *Wage Dynamics in the 21st Century*, NBER Chapters, pages 257–290. National Bureau of Economic Research.

Goos, M. and Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in britain. *The Review of Economics and Statistics*, 89(1):118–133.

Goos, M., Manning, A., and Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, 104(8):2509–2526.

Green, A. (2024). Artificial Intelligence and the changing demand for skills in the labour market. OECD Artificial Intelligence Papers 14, OECD Publishing.

Grullon, G., Larkin, Y., and Michaely, R. (2019). Are US industries becoming more concentrated? *Review of Finance*, 23(4):697–743.

Hagiu, A. and Wright, J. (2023). Data-enabled learning, network effects, and competitive advantage. *The RAND Journal of Economics*, 54(4):638–667.

Handa, K., Tamkin, A., McCain, M., Huang, S., Durmus, E., Heck, S., Mueller, J., Hong, J., Ritchie, S., Belonax, T., Troy, K. K., Amodei, D., Kaplan, J., Clark, J., and Ganguli, D. (2025). Which economic tasks are performed with ai? evidence from millions of claude conversations. arXiv preprint.

Hardy, W., Keister, R., and Lewandowski, P. (2018). Educational upgrading, structural change and the task composition of jobs in europe. *The Economics of Transition*, 26(2):201–231.

Harkness, S., Popova, D., and Avram, S. (2023). Gender differences in job mobility and pay progression in the UK. ISER Working Paper Series 2023-02, Institute for Social and Economic Research.

Impullitti, G. and Rendahl, P. (2025). Market power, growth, and wealth inequality. Discussion Paper 2025/01, Centre for Finance, Credit and Macroeconomics, University of Nottingham.

Jones, C. I. and Tonetti, C. (2020). Nonrivalry and the economics of data. *American Economic Review*, 110(9):2819–2858.

Kaldor, N. (1934). Mrs. Robinson's Economics of Imperfect Competition. *Economica, New Series*, 1(3):335–341.

Katz, L. F. (1986). Efficiency wage theories: A partial evaluation. In *NBER Macroeconomics Annual*, volume 1, pages 235–290. University of Chicago Press, Chicago, IL.

Kurz, M. (2017). On the formation of capital and wealth: IT, monopoly power and rising inequality. SSRN Working Paper 3014361, Social Science Research Network.

Lane, M. and Saint-Martin, A. (2021). The impact of Artificial Intelligence on the labour market: What do we know so far? OECD Social, Employment and Migration Working Papers 256, OECD Publishing.

Lassébie, J. and Quintini, G. (2022). What skills and abilities can automation technologies replicate and what does it mean for workers?: New evidence. OECD Social, Employment and Migration Working Papers 282, OECD Publishing, Paris.

Manning, A. (2003). *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton University Press.

Manning, A. (2021). Monopsony in labor markets: A review. *ILR Review*, 74(1):3–26.

Mortensen, D. T. (1970). Job search, the duration of unemployment, and the phillips curve. *American Economic Review*, 60(5):847–862.

OECD (2022). *OECD Employment Outlook 2022: Building Back More Inclusive Labour Markets*. OECD Employment Outlook. OECD.

OECD (2023). *OECD Employment Outlook 2023: Artificial Intelligence and the Labour Market*. OECD Publishing, Paris.

Pissarides, C. A. (1985). Short-run equilibrium dynamics of unemployment, vacancies, and real wages. *American Economic Review*, 75(4):676–690.

Pizzinelli, C., Panton, A. J., Mendes Tavares, M., Cazzaniga, M., and Li, L. (2023). Labor market exposure to AI: Cross-country differences and distributional implications. IMF Working Paper 216, International Monetary Fund.

Păvăloaia, V.-D. and Necula, S.-C. (2023). Artificial intelligence as a disruptive technology—a systematic literature review. *Electronics*, 12(5):1102.

Shapiro, C. and Stiglitz, J. E. (1984). Equilibrium unemployment as a worker discipline device. *American Economic Review*, 74(3):433–444.

Sokolova, A. and Sorensen, T. (2021). Monopsony in labor markets: A meta-analysis. 74(1):27–55. Publisher: Cornell University, ILR School.

Song, J., Price, D. J., Guvenen, F., Bloom, N., and von Wachter, T. (2019). Firming up inequality. *The Quarterly Journal of Economics*, 134(1):1–50.

Spence, M. (1976). Product selection, fixed costs, and monopolistic competition. *The Review of Economic Studies*, 43(2):217–235.

Tolan, S., Pesole, A., Martínez-Plumed, F., Fernández-Macías, E., Hernández-Orallo, J., and Gómez, E. (2021). Measuring the occupational impact of AI: Tasks, cognitive abilities and AI benchmarks. Technical report.

Webb, M. (2020). The impact of Artificial Intelligence on the labor market. Ssrn electronic journal, SSRN.

Yeh, C., Macaluso, C., and Hershbein, B. (2022). Monopsony in the us labor market. *American Economic Review*, 112(7):2099–2138.

Zwysen, W. (2024a). Monopsony and non-competitive labour markets: Workers' weakening bargaining position. Technical report, ETUI Research Paper - Report 2024.10.

Zwysen, W. (2024b). Monopsony and non-competitive labour markets: Workers' weakening bargaining position. Technical report, ETUI Research Paper - Report 2024.10.

Appendix

2.A Derivation of equation 2.4

We define

$$\theta_N = \frac{R_n}{R_e + R_n} \quad \implies \quad R_n = \frac{\theta_N}{1-\theta_N} R_e.$$

Taking logs,

$$\log R_n = \log\Big(\frac{\theta_N}{1-\theta_N}\Big) + \log R_e = \log(\theta_N) - \log\big(1-\theta_N\big) + \log R_e.$$

Differentiating with respect to w gives

$$\frac{R_{n\prime}}{R_n} = \frac{\theta_N'}{\theta_N} + \frac{\theta_N'}{1-\theta_N} + \frac{R_E'}{R_E} = \theta_N' \, \frac{1}{\theta_N(1-\theta_N)} + \frac{R_E'}{R_E}.$$

Hence,

$$\frac{R_N'}{R_N} = \frac{\theta_N'}{\theta_N(1-\theta_N)} + \frac{R_E'}{R_e}.$$

By definition, the wage elasticity of R_N is

$$\epsilon_R^n = \frac{\frac{R\prime_N}{R_N}}{\frac{w'}{w}}, \quad \text{and similarly} \quad \epsilon_R^e = \frac{\frac{R_E'}{R_E}}{\frac{w'}{w}}.$$

So,

$$\frac{\frac{R^{n\prime}}{R^n}}{\frac{w'}{w}} = \frac{\frac{\theta^{N\prime}}{\theta^N(1-\theta^N)}}{\frac{w'}{w}} + \frac{\frac{R^{e\prime}}{R^e}}{\frac{w'}{w}} \quad \implies \quad \epsilon_R^n = w \, \frac{\theta_N'}{\theta_N(1-\theta_N)} + \epsilon_R^e.$$

Rearranging to express everything in terms of ϵ_R^e:

$$\boxed{\epsilon_R^e = \epsilon_R^n - w \, \frac{\theta_N'}{\theta_N(1-\theta_N)} \, .}$$

Let $\phi\Big(\frac{x}{w}\Big)$ be the probability that a nonemployed worker, whose outside option is w, accepts a job offering x. Let $F(x)$ be the distribution of wage offers. The separation rate to nonemployment of a firm paying wage w can then be expressed as

$$s^n(w) = \lambda^n \int_{\underline{w}}^{\overline{w}} \phi\Big(\frac{x}{w}\Big)\, dF(x), \quad \text{with derivative} \quad \frac{d\, s^n(w)}{d\, w} = -\, \lambda^n \int_{\underline{w}}^{\overline{w}} \phi'\Big(\frac{x}{w}\Big)\, \frac{x}{w^2}\, dF(x).$$

The firm's number of recruits from nonemployment is

$$R^n(w) = \lambda^n \int_{\underline{w}}^{\overline{w}} \phi\Big(\frac{x}{w}\Big)\, dF(x), \quad \text{with derivative} \quad \frac{d\, R^n(w)}{d\, w} = \lambda^n \int_{\underline{w}}^{\overline{w}} \phi'\Big(\frac{x}{w}\Big)\, \frac{x}{w^2}\, dF(x).$$

Then, we have:

$$\begin{aligned}
\int_{w}^{w} \epsilon_{sw}^n(x)\, s^n(x)\, L(x)\, dF(x) &= \int_{w}^{w} \frac{d\, s^n(x)}{dx}\, x\, s^n(x)\, L(x)\, dF(x) \\
&= \int_{w}^{w} \Big(-\, \lambda^n\Big) \int_{x} \phi'\Big(\frac{z}{x}\Big)\, \frac{z}{x^2}\, dF(z)\; x\, L(x)\, dF(x) \\
&= -\, \lambda^n \int_{w}^{w} \int_{x} \phi'\Big(\frac{z}{x}\Big)\, \frac{z\, L(x)}{x}\, dF(z)\, dF(x) \\
&= -\, \int_{w}^{w} \frac{d\, R^n(x)}{dx}\, x\, dF(x).
\end{aligned}$$

Thus,

$$\int_{w}^{w} \epsilon_{sw}^n(x)\, s^n(x)\, L(x)\, dF(x) = -\, \int_{w}^{w} \epsilon_{Rw}^n(x)\, R^n(x)\, dF(x).$$

Note that in steady state, for the aggregate economy, it holds that

$$s^n(x)L(x) = \theta_Z S(x) \quad \text{(for separations to nonemployment)}$$

and

$$R^n(x) = \theta_N R(x) \quad \text{(for hirings from nonemployment)}.$$

Hence,

$$\theta_Z \int_{w}^{w} \epsilon_{Sw}^n(x)\, S(x)\, dF(x) = -\, \theta_n \int_{w}^{w} \epsilon_{Rw}^n(x)\, R(x)\, dF(x).$$

Dividing both sides by $-\, \theta_N$ yields:

$$\int_{w}^{w} \epsilon_{Rw}^n(x)\, R(x)\, dF(x) = -\frac{\theta_Z}{\theta_N} \int_{w}^{w} \epsilon_{Sw}^n(x)\, S(x)\, dF(x).$$

Defining average elasticities, we can write:

$$\epsilon_S^n = \frac{\int \epsilon_{Sw}^n(x)\, S(x)\, dF(x)}{\int S(x)\, dF(x)}, \quad \text{and} \quad \epsilon_r^n = \frac{\int \epsilon_{Rw}^n(x)\, R(x)\, dF(x)}{\int R(x)\, dF(x)}.$$

Thus,

$$\epsilon_R^n = -\frac{\theta_Z}{\theta_N}\, \epsilon_S^n.$$

Finally, considering the overall labor market elasticity we have:

$$\epsilon_{L_w} = (1-\theta_n)\, \epsilon_{R_w}^e + \theta_n\, \epsilon_{R_w}^n - (1-\theta_n)\, \epsilon_{S_w}^e - \theta_n\, \epsilon_{R_w}^n.$$

We start from the initial equation:

$$\epsilon_{L_w} = (1-\theta_n)\, \epsilon_{R_w}^e + \theta_n\, \epsilon_{R_w}^n - (1-\theta_n)\, \epsilon_{s_w}^e - \theta_n\, \epsilon_{s_w}^n.$$

We then make the following substitutions:

$$\epsilon_{R_w}^n = -\frac{\theta_z}{\theta_n}\, \epsilon_s^n,$$

and

$$\epsilon_{R_w}^e = \epsilon_R^n - w\, \frac{\theta^{N\prime}}{\theta^N(1-\theta^N)},$$

with

$$\epsilon_R^n = -\frac{\theta_Z}{\theta_N}\, \epsilon_s^n.$$

Thus,

$$\epsilon_{R_w}^e = -\frac{\theta_z}{\theta_n}\, \epsilon_s^n - w\, \frac{\theta^{N\prime}}{\theta^N(1-\theta^N)}.$$

Substituting these into the initial equation gives:

$$\begin{aligned}\epsilon_{L_w} = {} & (1-\theta_n)\left[-\frac{\theta_z}{\theta_n}\, \epsilon_s^n - w\, \frac{\theta^{N\prime}}{\theta^N(1-\theta^N)}\right] + \theta_n \left[-\frac{\theta_z}{\theta_n}\, \epsilon_s^n\right] \\ & - (1-\theta_n)\, \epsilon_{s_w}^e - \theta_n\, \epsilon_{s_w}^n.\end{aligned}$$

Expanding the terms:

$$\epsilon_{L_w} = -\frac{(1-\theta_n)\, \theta_z}{\theta_n}\, \epsilon_s^n - (1-\theta_n)\, w\, \frac{\theta^{N\prime}}{\theta^N(1-\theta^N)} - \theta_z\, \epsilon_s^n - (1-\theta_n)\, \epsilon_{s_w}^e - \theta_n\, \epsilon_{s_w}^n.$$

Grouping the ϵ_s^n terms:

$$-\frac{(1-\theta_n)\,\theta_z}{\theta_n}\,\epsilon_s^n - \theta_z\,\epsilon_s^n = -\theta_z\,\epsilon_s^n\left(\frac{1-\theta_n}{\theta_n}+1\right) = -\frac{\theta_z}{\theta_n}\,\epsilon_s^n.$$

Thus, we obtain:

$$\epsilon_{L_w} = -\frac{\theta_z}{\theta_n}\,\epsilon_s^n - (1-\theta_n)\,w\,\frac{\theta^{N\prime}}{\theta^N(1-\theta^N)} - (1-\theta_n)\,\epsilon_{s_w}^e - \theta_n\,\epsilon_{s_w}^n.$$

Now, imposing the additional conditions

$$\theta_Z = \theta_N, \quad \epsilon_{s_w}^n = \epsilon_s^n,$$

the equation becomes:

$$\epsilon_{L_w} = -\epsilon_s^n - (1-\theta_n)\,w\,\frac{\theta_n'}{\theta_n} - (1-\theta_n)\,\epsilon_{s_w}^e - \theta_n\,\epsilon_s^n.$$

Combining the ϵ_s^n terms gives:

$$-\epsilon_s^n - \theta_n\,\epsilon_s^n = -(1+\theta_n)\,\epsilon_s^n.$$

Therefore, defining

$$\epsilon_{\theta,w} = w\,\frac{\theta_n'}{\theta_n},$$

the final expression is:

$$\boxed{\epsilon_{L_w} = -\big(1+\theta_n\big)\,\epsilon_s^n - \epsilon_{\theta,w} - (1-\theta_n)\,\epsilon_{s_w}^e.}$$

2.B Descriptive evidence

2.B.1 Annual evolution of AI exposure by profession

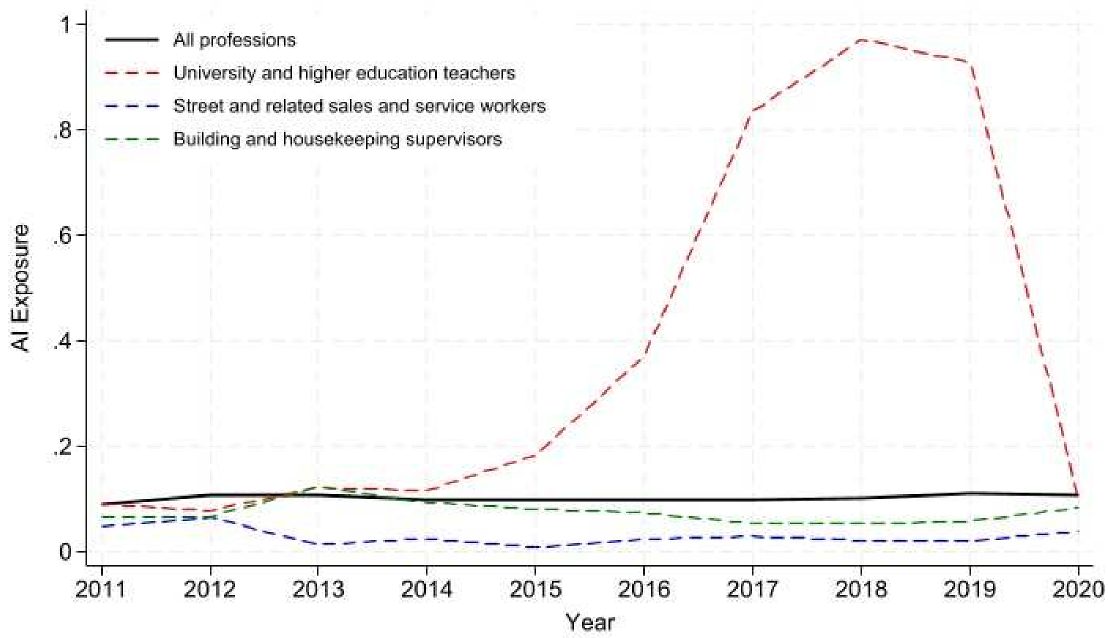

Figure 2.B.1: **Annual evolution of AI exposure by profession**

Source: Authors' calculation.

Notes: The graph shows how the Webb index we used to measure AI exposure evolves for three different professions (represented by dashed lines) and for all professions (represented by a solid black line).).

2.B.2 Yearly labour supply elasticities for the nine major groups of ISCO-08

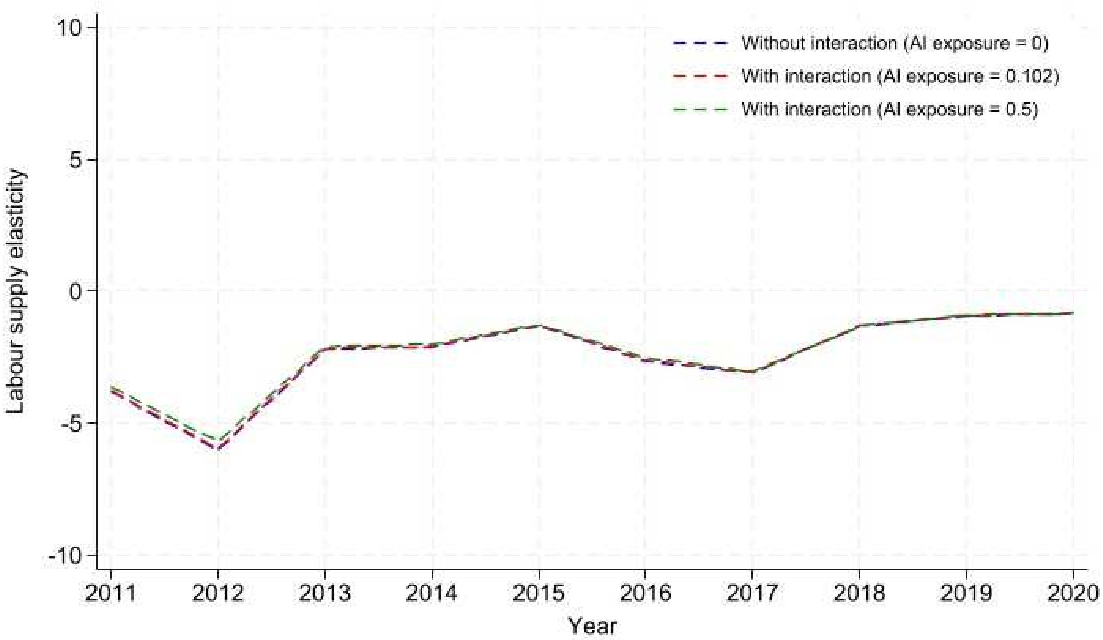

Figure 2.B.2: **Labour supply elasticity by major groups 1**

Source: Authors' calculation.

Notes: Labour supply elasticity is computed following the guidelines outlined in the paper for each of the occupations belonging to Major Group 1 of the ISCO-08 1-digit classification. Then, an annual average is calculated by weighting each 3-digit occupation according to its employment share.

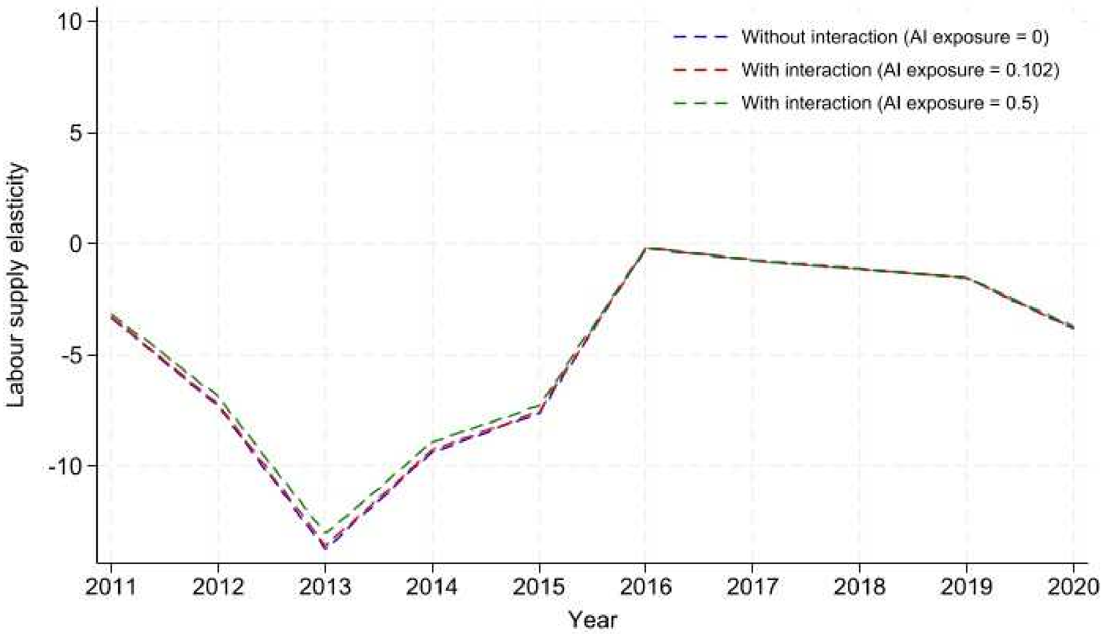

Figure 2.B.3: **Annual evolution of AI exposure by major groups 2**

Source: Authors' calculation.

Notes: The graph shows how the Webb index we used to measure AI exposure evolves for three different professions (represented by dashed lines) and for all professions (represented by a solid black line) within Major Group 2.

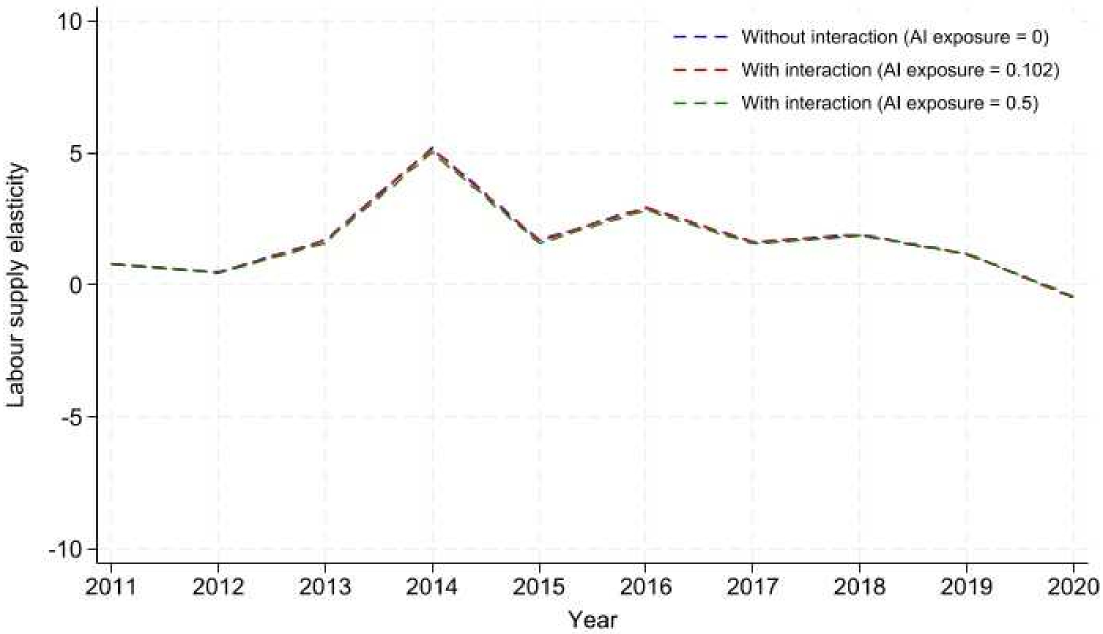

Figure 2.B.4: **Annual evolution of AI exposure by major groups 3**

Source: Authors' calculation.

Notes: The graph shows how the Webb index we used to measure AI exposure evolves for three different professions (represented by dashed lines) and for all professions (represented by a solid black line) within Major Group 3.

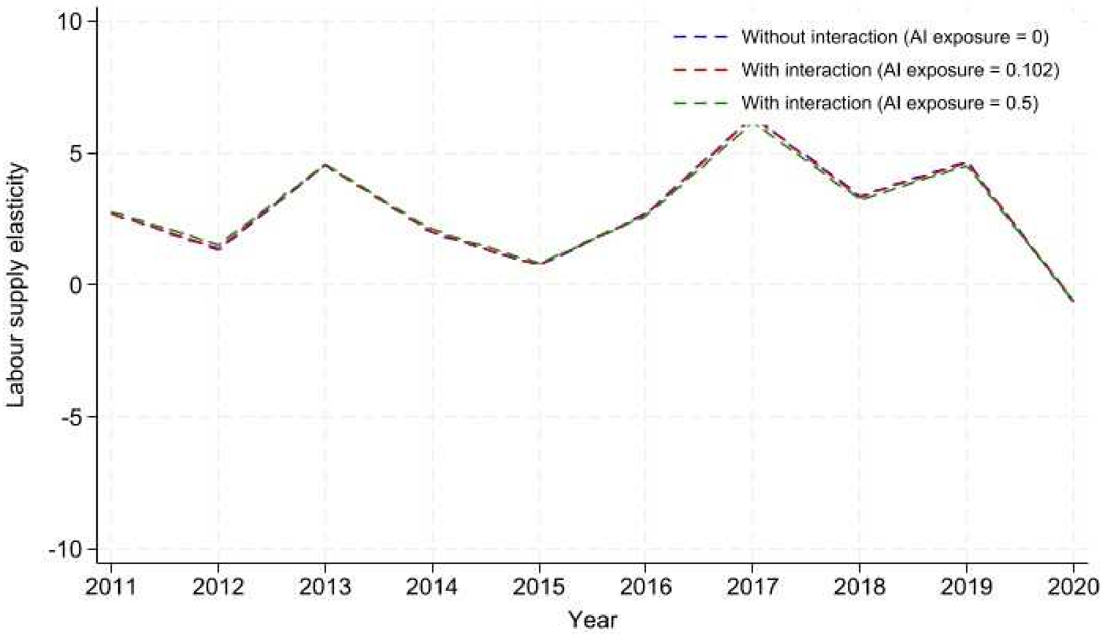

Figure 2.B.5: **Annual evolution of AI exposure by major groups 4**

Source: Authors' calculation.

Notes: The graph shows how the Webb index we used to measure AI exposure evolves for three different professions (represented by dashed lines) and for all professions (represented by a solid black line) within Major Group 4.

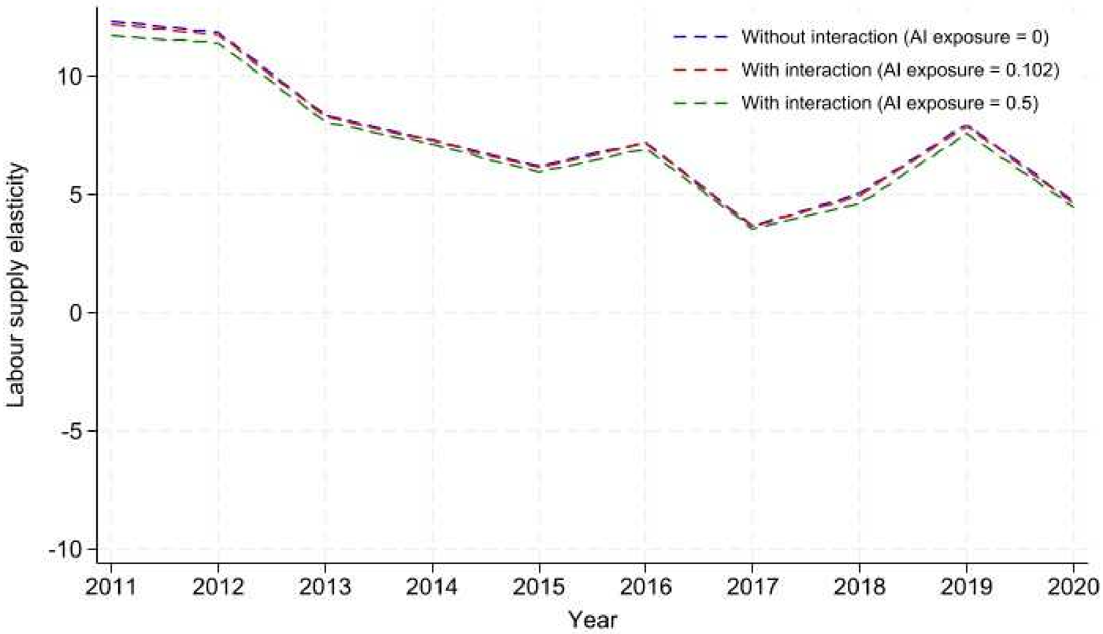

Figure 2.B.6: **Annual evolution of AI exposure by major groups 5**

Source: Authors' calculation.

Notes: The graph shows how the Webb index we used to measure AI exposure evolves for three different professions (represented by dashed lines) and for all professions (represented by a solid black line) within Major Group 5.

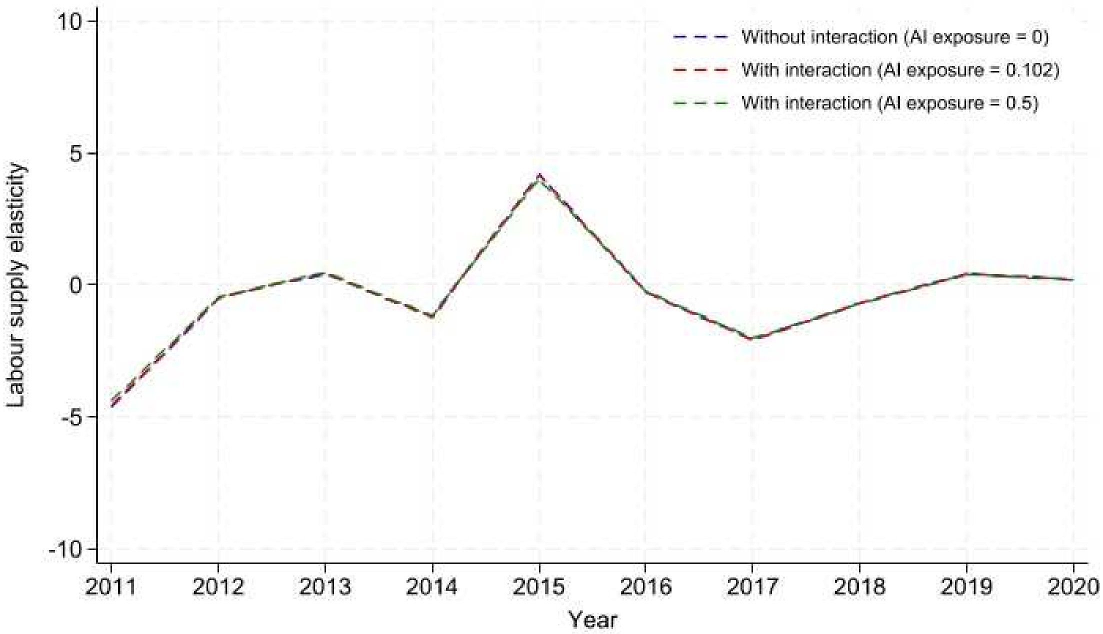

Figure 2.B.7: **Annual evolution of AI exposure by major groups 6**

Source: Authors' calculation.

Notes: The graph shows how the Webb index we used to measure AI exposure evolves for three different professions (represented by dashed lines) and for all professions (represented by a solid black line) within Major Group 6.

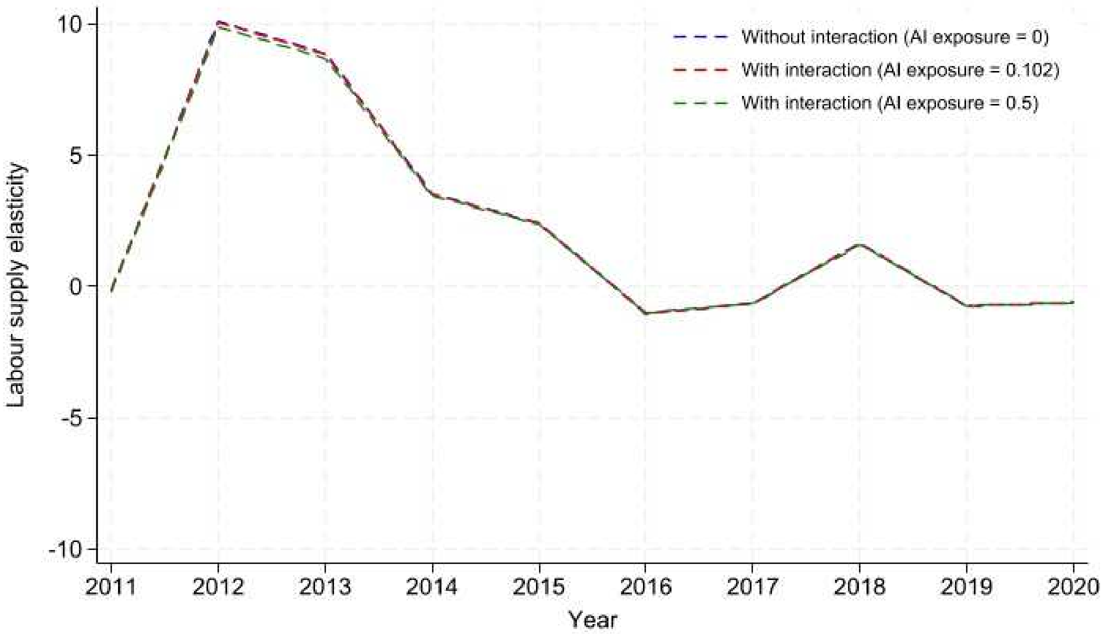

Figure 2.B.8: **Annual evolution of AI exposure by major groups 7**

Source: Authors' calculation.

Notes: The graph shows how the Webb index we used to measure AI exposure evolves for three different professions (represented by dashed lines) and for all professions (represented by a solid black line) within Major Group 7.

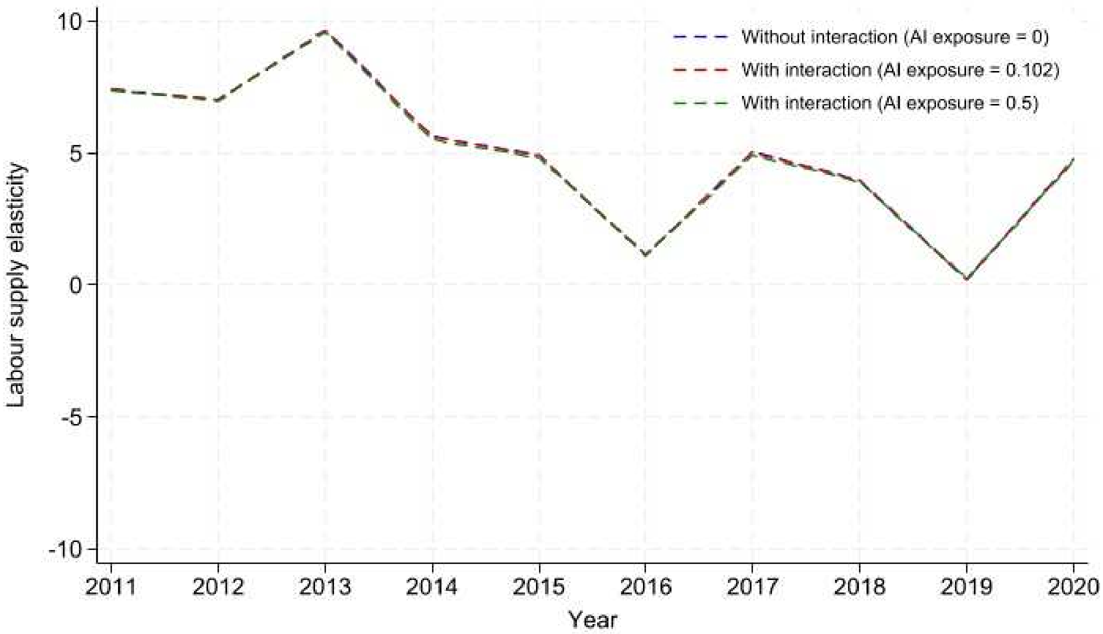

Figure 2.B.9: **Annual evolution of AI exposure by major groups 8**

Source: Authors' calculation.

Notes: The graph shows how the Webb index we used to measure AI exposure evolves for three different professions (represented by dashed lines) and for all professions (represented by a solid black line) within Major Group 8.

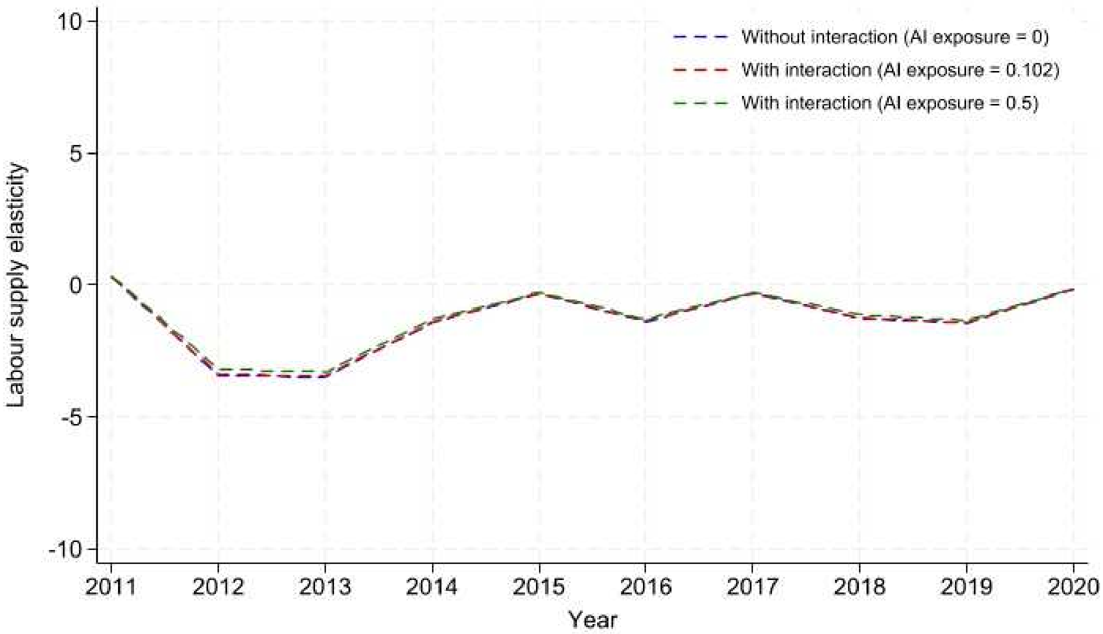

Figure 2.B.10: **Annual evolution of AI exposure by major groups 9**

Source: Authors' calculation.

Notes: The graph shows how the Webb index we used to measure AI exposure evolves for three different professions (represented by dashed lines) and for all professions (represented by a solid black line) within Major Group 9.

Chapter 3

Workers Without Quality? Enrichment or Impoverishment of Abilities, Knowledge, and Skills in the Age of Artificial Intelligence

Abstract

This paper investigates how artificial intelligence (AI) is reshaping the subjective content of work— abilities, skills, and knowledge— rather than task structures. Using longitudinal O*NET data from 2011 to 2025 combined with two novel AI exposure measures, one based on AI research topics and one on LLM benchmark performance, we construct worker-requirement–level indices of AI exposure for occupation–requirement–year observations. Our empirical analysis produces two main findings. First, at the worker-requirement–occupation level, AI exposure is positively associated with increases in importance-weighted requirement levels: within occupations, the requirements most exposed to AI tend to intensify over time. This pattern is broad-based across Abilities, Skills, and Knowledge. Second, when aggregating to the occupation level, the relationship reverses: occupations with higher AI exposure experience a net decline in overall requirement levels, driven primarily by reductions in Abilities. This divergence suggests that AI enriches specific complementary requirements while simultaneously contributing to an erosion of average requirement levels across occupations. The paper highlights the need for new methodologies

integrating task-based and skill-based perspectives to capture the subjective evolution of work in the age of AI and calls for further research on long-run shifts in worker requirement clusters and the transformative role of emerging LLM technologies.
Keywords: *Artificial Intelligence; Occupational Requirements; Skills and Abilities; O*NET; Large Language Models; Technological Change*
JEL codes: J24, O33, J23, C81

3.1 Introduction

In the debate on the transformations and the future of work, in addition to the dimension that concerns employment levels, considerable discussion also revolves around the evolution of the occupational structure and on the changing skill and labour-force requirements in Europe. The debate has taken various directions; one of the most central concerns whether contemporary labour markets are undergoing processes of deskilling, upskilling, or reskilling. This debate, which captures broader long-term transformations in advanced economies, is closely intertwined with the diffusion of new technologies, widely regarded as one of the key forces reshaping the organisation of work. It is also a debate that resonates strongly today, at a moment often described as a new technological revolution capable of redefining the contours of labour. We refer here, of course, to the rapid development and diffusion of Artificial Intelligence (AI), a domain in which the relationship between skills and technological change has become a dominant theme in policy and academic discussions.

The notion of deskilling itself has a long intellectual lineage. In general terms, deskilling can be understood as the process through which skills previously required in production or service work become redundant due to organisational or technological change, allowing tasks once performed by skilled workers to be executed by semi-skilled or unskilled labour. Deskilling can also occur at the individual level, as workers become less proficient over time due to changes in job content, occupational shifts, underemployment, or prolonged detachment from the labour market.

Although a full reconstruction of the history of economic thought lies beyond the scope of this introduction, the classic reference remains Braverman (1974), who identified the progressive erosion of skill in factory and clerical work as a consequence of the consolidation of control over production by monopolistic capital. Following this line, several scholars have shown how the separation between the conception and execution of work, exemplified in the principles of Scientific Management introduced by Taylor (1911) has systematically reduced workers' discretion and concentrated technical knowledge in managerial and engineering roles. Modern production systems, especially those relying on automation, have often been designed precisely to fragment complex tasks into simpler routines, reducing the need for craft-based expertise while reinforcing managerial authority over the labour process. The experiences of the nineteenth and early twentieth centuries therefore led Braverman, Marglin (1974), Thompson (1983) to argue that a major purpose of successive waves of technical change was to expand the division of labour and simplify artisanal tasks by decomposing them into highly standardised components. In their view, the key principle was to render production increasingly independent from workers' embodied knowledge and skills.

Counterbalancing this position, another line of thought highlights instead that one of the most significant trends shaping contemporary labour markets is upskilling (Acemoglu, 2002; Katz and Murphy, 1992). Upskilling may be understood either as a growing demand for occupations with high cognitive content or requiring high-level skills, or more generally as an increase in the skill requirements across all occupations. This is partly because simpler tasks can increasingly be performed by machines, allowing human workers to concentrate on more complex cognitive or social skills. As is evident, the question of whether labour markets are experiencing deskilling, upskilling, or reskilling can thus be approached from multiple angles: it may refer to an overall shift in the demand for different categories of jobs, or to changes in the complexity of the tasks performed by humans relative to machines.

The relevance of these themes is evident in the work of two recent Nobel laureates in economics, Acemoglu and Johnson (2023) and Mokyr (2002), who, albeit from different perspectives, have explored the complex interplay between technology, skills, and economic transformation.

Today, with the widespread diffusion of AI, these debates re-emerge with renewed urgency. As with previous general-purpose technologies (Calvino et al., 2025), the initial wave of discussion has often focused on whether AI will threaten employment levels or create new opportunities. Yet the deeper question concerns its impact on the structure of work and on human skill requirements. As Agrawal et al. (2019) or Acemoglu and Restrepo (2019) have argued, AI should be understood not merely as a labour-saving innovation but as a general-purpose prediction technology capable of reshaping economic organisation as a whole. This interpretation aligns with recent OECD (2023) analyses describing AI as a transformative force with the potential to redefine production systems across sectors. Ultimately, each technological shift recombines the relationship between humans and machines. Understanding how AI alters this relationship, and with it the distribution and content of skills, thus remains a central challenge for contemporary research, particularly because, as multiple studies have shown (Felten et al., 2021, 2019; Georgieff and Hyee, 2021; Tolan et al., 2021; Webb, 2019), AI is a technology that not only affects blue-collar occupations but is also capable of performing tasks that were previously considered the exclusive domain of high-skill or cognitively intensive professions. Building on these premises, we intervene in this debate by providing empirical evidence on how and to what extent Artificial Intelligence is affecting worker requirements[1], where by this term we refer to the set of Abilities, Knowledge, and Skills as defined in O*NET[2].

Our general objective is to examine whether, by studying the evolution of these three dimensions between 2011 and 2025 in the United States and regressing them on a measure of exposure to Artificial Intelligence, we are able to identify a potential impoverishment or enrichment attributable to AI. Compared to the broader debate outlined above, our contribution does not evaluate how demand for specific occupational profiles characterised by particular skills has changed. Instead, our empirical strategy allows us to observe how these dimensions vary within and across occupations, without assuming ex ante that a profession is cognitive, technical, or otherwise.

Our interest lies in capturing the evolution of the subjective content of work, rather than its objective task structure. For this reason, we focus on worker requirements rather than on the complexity or simplification of tasks, which would reflect a more job-oriented perspective, in order to emphasise precisely this subjective dimension. This first contribution can later be related to the evolution of objective tasks as documented in the existing literature.

In other words, we address the following question: within each occupation, are the worker requirements most exposed to AI also those that gained or lost the most centrality between 2011 and 2025? A related question concerns the extent to which AI influences these worker requirements. We reformulate the model as well to examine whether the occupations most exposed to AI are also those in which worker requirements changed the most between 2011 and 2025.

Our contribution relies on an empirical analysis combining several data sources. O*NET provides detailed information on the content of worker requirements. The CPS supplies information on the characteristics of labour market occupations in the United States. Two additional datasets allow us to measure exposure of worker

[1]We use the term *worker requirements* directly from O*NET, although we introduce a simplification: in the O*NET Content Model, *Abilities* are included under *Worker Characteristics*. We nonetheless adopt the consolidated term "worker requirements" for analytical clarity.

[2]https://www.onetcenter.org/content.html

requirements to AI: a topics-based measure capturing research intensity in specific AI domains, and a dataset on the evolution of benchmark performance of large language models, which captures a different dimension of AI development. These data sources are harmonised in order to construct a comprehensive analytical framework.

Our aim is to fill a gap in the literature on the impact of AI on labour market structures and worker requirements, while highlighting a subjective dimension that is often overlooked. Our findings indicate that, at present, AI has a positive, albeit modest, effect on worker requirements within occupations, contributing to a slight enrichment of Abilities, Knowledge, and Skills at the micro level. Conversely, at a more macro occupational level, the data suggest a mild impoverishment in professions with higher exposure to AI. This divergence suggests that AI enriches specific complementary requirements while simultaneously contributing to an erosion of average requirement levels across occupations The rest of the paper is organised as follows. Section 2 reviews the relevant literature. Section 3 describes the data used in our empirical analysis. Section 4 explains the methodology, while Section 5 outlines the empirical strategy and presents the first set of results. Section 6 discusses the main findings, and Section 7 concludes. The Appendix provides additional details on specific steps of the analysis, the construction of the dataset, and further descriptive evidence.

3.2 Literature review

The deskilling hypothesis, rooted in Braverman (1974) analysis of capitalist production, posits that technological and organisational innovations tend to fragment complex activities into simpler tasks, reducing worker autonomy and eroding the embodied knowledge that traditionally characterised skilled occupations. Economic studies reinforce this interpretation when showing how automation displaces routine activities through codification and standardisation, thereby diminishing the informational and decision-making content of certain jobs (Autor et al., 2003). In this sense, deskilling may operate not only through the simplification of technical tasks but also through the reduction of discretion, autonomy, and opportunities for informal skill accumulation.

A contrasting interpretation is offered by the literature on skill-biased technological change (SBTC), which maintains that technological progress systematically increases the demand for highly skilled workers. The central mechanism is that new technologies tend to complement cognitive, analytical, and problem-solving skills while substituting for routine-intensive tasks (Acemoglu, 2002; Katz and Murphy, 1992). This theory was initially advanced to explain the persistently high education and skill wage premiums observed since the 1980s, as well as evidence that computerisation and digitalisation disproportionately raise the productivity of more educated workers (Autor et al., 1998). Moreover, technological upgrading has been shown to account for a substantial share of the displacement of low-skilled labour in the United States (Berman et al., 1998), with similar patterns documented in several European economies undergoing rapid structural change (Esposito and Stehrer, 2009). SBTC also developed alongside the diffusion of ICT, as numerous studies demonstrated that the adoption of computer-based technologies is systematically associated with increases in the demand for skilled labour (Autor et al., 1998; Fernandez, 2001). Thus, under the SBTC framework, the idea is that as production processes become more complex, firms increasingly rely on workers with advanced qualifications and competences (Goldin and Katz, 2007; Hutter and Weber, 2017). In this view, the direction of technological change systematically favours workers endowed with stronger cognitive and problem-solving capabilities, leading to a generalised process of upskilling across the labour market. A third line of research challenges the notion of monotonic upgrading by proposing the employment polarisation hypothesis. Studies by Acemoglu and Autor (2011) and Goos et al. (2009, 2014)argue that digital technologies automate routine, middle-skill jobs while simultaneously increasing demand for both high-skilled analytical labour and low-skilled interpersonal services. The resulting pattern is characterised by job growth at the top and bottom of the occupational distribution and decline in the middle. The theoretical foundation of this literature is the task-based model first formalised by Autor et al. (2003) (ALM), which distinguishes between routine and non-routine manual, cognitive and analytical tasks. The ALM framework posits that computerisation primarily substitutes routine tasks—both manual and cognitive—while complementing abstract and interactive tasks. Subsequent studies have adapted and refined this taxonomy. Early applications in the United Kingdom and Germany (Goos and Manning, 2007; Spitz-Oener, 2006); followed the original five-task structure, whereas more recent work converged on a three-fold classification—abstract, routine and manual—introduced by Autor et al. (2006) and formalised by Autor and Dorn (2013) . Across these studies, the underlying mechanism is consistent: technology displaces routine-intensive jobs and supports both high-skill cognitive occupations and certain non-routine manual service roles. A second strand of work has expanded the task approach by proposing alternative frameworks that capture additional dimensions of work complexity. For example, Matthes et al. (2014) incorporate autonomy and define routine tasks in an abstract, evolving manner to account for the changing codifiability of work.

Fernández-Macías and Hurley (2017) add a social-interaction dimension that is considered resilient to automation. These alternative operationalisations highlight that the boundaries of "routine" work are not fixed but shaped by both technological feasibility and organisational choices. A key reason for the divergent conclusions across the three strands of literature lies in the measurement of skills. Advocates of the deskilling hypothesis typically examine changes within occupations, focusing on how the substance of work evolves over time. Research on SBTC and on polarisation generally studies changes between occupations, inferring skill trends from shifts in employment shares. Moreover, studies differ in the indicators used to proxy skills: some rely on wages or educational attainment, while others use task content, autonomy, prestige, or job satisfaction (Autor and Handel, 2013). Because these indicators capture different dimensions of work, they may suggest different patterns of change even when applied to the same labour markets. Recent evidence further complicates the polarisation narrative by showing that, around the early 2000s, the demand for several cognitive tasks began to decline. (Beaudry et al., 2016) document that following the tech bust, high-skilled workers increasingly moved down the occupational ladder, displacing less educated workers in lower-skill jobs. This reversal of cognitive task demand suggests that technological change can compress the occupational structure from the top, and not only hollow out the middle, thereby blurring the boundaries between upskilling and deskilling.

This conceptual heterogeneity becomes even more salient in the context of Artificial Intelligence. Unlike earlier waves of computerisation, which primarily automated routine, codifiable tasks, recent advances in AI—especially in machine learning and large language models—extend automation and augmentation into domains traditionally associated with high-skill cognitive and professional work. A growing body of research constructs task- and occupation-level measures of AI exposure, showing that prediction technologies, natural language processing and generative models increasingly overlap with tasks involving information synthesis, drafting, coding and decision support (Agrawal et al., 2019; Felten et al., 2021). Cross-country evidence based on PIAAC and related measures indicates that exposure to AI is highest in highly educated white-collar occupations, such as managers, business professionals and science and engineering professionals, while elementary and manual jobs remain comparatively less exposed (Georgieff and Hyee, 2021; OECD, 2023). Micro-level evidence is more nuanced: AI adoption is associated with reduced hiring in non-AI-intensive roles and with sector-specific displacement in some high-exposure occupations, but also with employment gains or better retention prospects for highly skilled workers, whereas low-skilled workers appear more vulnerable and harder to reallocate when their tasks are automated Bessen (2018); Milanez (2023).

The emerging literature on generative AI goes further, suggesting that high-wage occupations may be especially exposed to these new tools, with some studies emphasising cost savings and profitability gains, others highlighting complementarities with high-skill work and uneven diffusion across workers (Acemoglu, 2024; Auer et al., 2024; Eisfeldt et al., 2023; Eloundou et al., 2024; Humlum and Vestergaard, 2024). Autor (2024) interprets these developments as the emergence of an "inversion technology": rather than merely displacing middle-skill routine work, AI can, in principle, diffuse forms of expert judgment to a broader set of workers by providing real-time guidance and decision support. In this view, AI may simultaneously substitute for certain standardised cognitive routines and complement workers' higher-order capabilities, potentially upgrading segments of middle-skill work and reshaping the structure of expertise. Whether this results in deskilling or upskilling, and for whom, is ultimately an empirical question that requires looking beyond occupation-level employment shares to the evolution of worker requirements within jobs.

3.3 Data

In our study, we focus on the U.S. labor market to examine the impact of artificial intelligence on the evolution of worker requirements, namely Abilities, Knowledge, and Skills. To emphasize our interest in the *subjective* component of work, that is, characteristics pertaining to workers themselves rather than to jobs, our approach departs from the prevailing task-based perspective in the literature (Autor, 2013), which builds on the work of Acemoglu and Autor (2011), Acemoglu and Zilibotti (2001), Autor et al. (2006), Autor et al. (2008). By concentrating on this subjective dimension, our framework aims to illuminate the evolution of worker-specific attributes rather than variation in job tasks. A further step would involve linking skills and tasks; some theoretical advances in this direction have already been proposed by Rodrigues et al. (2021). Our unit of analysis is an occupation–worker-requirement cell. Occupations are classified according to the O*NET-SOC 2019 taxonomy, using the six-digit level of disaggregation. Worker requirements are drawn from the O*NET Content Model and include Abilities, Skills, and Knowledge. Each occupation is then matched with occupation-level data from the U.S. Bureau of Labor Statistics' Current Population Survey (CPS).

To construct our exposure indices, defined at the worker-requirement–occupation level, we use two datasets: AI Topics, which tracks the evolution of artificial intelligence research, and LLM Stats, which provides statistics on the development of major large language models available on the market, as well as trends in specific AI benchmarks. Our analysis covers the period from 2011 to 2025.

3.3.1 Worker requirements data

To extract data on worker requirements, we rely on National Center for O*NET Development (b). As stated by O*NET, "The Content Model was developed using research on job and organizational analysis. It embodies a view that reflects the character of occupations (via job-oriented descriptors) and people (via worker-oriented descriptors). The Content Model also allows occupational information to be applied across jobs, sectors, or industries (cross-occupational descriptors) and within occupations (occupation-specific descriptors). These descriptors are organized into six major domains, which enable the user to focus on areas of information that specify the key attributes and characteristics of workers and occupations."

In particular, we focus on the indicators *Abilities* (defined as "Enduring attributes of the individual that influence performance"), *Basic and Cross-Functional Skills* (respectively defined as "Developed capacities that facilitate learning or the more rapid acquisition of knowledge" and "Developed capacities that facilitate performance of activities that occur across jobs"), and *Knowledge* (defined as "Organized sets of principles and facts applying in general domains").

For each of these worker requirements, O*NET assigns to every occupation both an *importance* score and a *level* score. The former indicates "the degree of importance a particular descriptor is to the occupation," while the latter indicates "the degree, or point along a continuum, to which a particular descriptor is required or needed to perform the occupation" (National Center for O*NET Development, a).

3.3.2 AI scores

To compute our AI score, we rely on two datasets: *AI Topics* and *LLM Stats*. *AI Topics* [3]is a long-standing repository curated by the Association for the Advancement of Artificial Intelligence (AAAI), which compiles a vast range of material related to artificial intelligence research, applications, and contributors.[4] The platform aggregates a wide spectrum of AI-related content, including news items, blog posts, conference proceedings, journal articles, and other documents spanning the period from 1905 to 2019, collected automatically through the NewsFinder system (Buchanan et al.,

[3]See https://aitopics.org/
[4]See http://www.aaai.org/aitopics.

2013). This dataset has been employed in previous empirical analyses of technological change, such as Tolan et al. (2021). *LLM Stats* [5] is a dataset that tracks recent advances in Artificial Intelligence by compiling one of the most extensive collections of high-quality evaluations across multiple AI modalities, including language, vision, code, and reasoning. It enables systematic comparisons across the major models available on the market and provides benchmark-based assessments (see Appendix for details). The benchmark suite, covering, for instance, Mathematics, Vision, Language, General Intelligence, and Code, allows for direct comparisons among all principal models released by OpenAI, Amazon, DeepSeek, Anthropic, and others.

3.3.3 Labour market data

Finally, data used to describe the U.S. labor market are drawn from the Current Population Survey (CPS), sponsored jointly by the U.S. Census Bureau and the U.S. Bureau of Labor Statistics (BLS). The CPS is the primary source of labor force statistics for the United States and provides information for the period 2011–2025 considered in our analysis. We are particularly interested in employment shares and their evolution over time by occupation, which are available at the 6-digit SOC 2018 level. In addition, we extract a set of individual and occupational covariates used as controls, including industry, education, age, gender, and wage levels for each occupation.

3.4 Method

In this section, we describe how we construct a new, objective measure of the exposure of worker requirements to Artificial Intelligence by using two distinct databases and then aggregating it to the worker requirements dataset.

3.4.1 Worker requirements

Although O*NET is one of the main datasets used to analyse the U.S. labour market, worker skills, and the components of work, the literature still lacks studies that exploit O*NET in a genuinely temporal or longitudinal perspective. To construct our dataset, we begin by assembling, for each occupation, worker-requirement pair (Abilities, Skills, and Knowledge), the corresponding *importance* and *level* scores provided directly by O*NET. For instance, in our raw dataset, the occupation *Economist* is assigned an Importance and a Level value for each of the 51 Abilities defined in the

[5] See https://llm-stats.com/

Content Model, for each of the 36 Skills, and for each of the 33 Knowledge categories. This procedure is repeated for every year from 2011 to 2025, although it should be noted that some worker requirements begin to appear only in later years.

As a result, the occupation *Economist* features roughly 120 worker-requirement observations per year, multiplied by 15 years, yielding about 1,800 observations. Accounting for missing values, the final dataset comprises 1,645,906 observations organised in occupation–worker-requirement–year cells, with each cell containing an Importance and a Level score, both ranging from 0 to 100.

We then construct a key variable in which the Level score—our main measure for assessing whether a worker requirement becomes richer or poorer over time, is weighted by its Importance score. This approach is widely used in the literature since Autor et al. (2003). We focus on the Level score because it captures a qualitative dimension of human capability: it measures *how much* of a certain attribute an individual needs to perform a task, and therefore how artificial intelligence may affect that requirement. By weighting Level by Importance, we incorporate the extent to which each worker requirement matters for the occupation under consideration. Recent evidence, for instance, shows that artificial intelligence is associated with declines in Critical Thinking (Gerlich, 2025) or Memory (Zhai et al., 2024).

Given that O*NET is updated multiple times per year,[6] we average the Importance and Level scores within each calendar year to obtain a single yearly value. This choice also mitigates inconsistencies in the database, since some updates produce erratic or highly noisy scores.

A further methodological issue, discussed in more detail in the results section, concerns the nature of O*NET data. Importance and Level are based on survey questions administered to workers. For example, workers are asked: "If at least somewhat important, what level of information ordering is needed to perform your current job?" The difficulty lies in determining whether respondents answer with respect to their *human* capabilities alone, or with respect to a *human-AI hybrid*, for instance when a worker routinely relies on tools such as ChatGPT but implicitly attributes that capability to themselves. This conceptual ambiguity is difficult to resolve and raises broader philosophical questions about the nature of human in the age of AI (Galli, 2025).

Construction of the AI score from AI Topics

To assess how artificial intelligence contributes to the enrichment or impoverishment of worker requirements, we compute two distinct AI scores. Most of the existing

[6] See `https://www.onetcenter.org/dataUpdates.html` for a complete list of updates.

literature constructs indices that measure occupational exposure to AI through a task-based approach, mapping AI benchmarks or AI capabilities to tasks (Brynjolfsson et al., 2018; Eloundou et al., 2024; Gmyrek et al., 2023), or by relying on Google Patents to infer technological exposure (Webb, 2019). In contrast to these approaches, we directly map the relationship between worker requirements and artificial intelligence, following the line of work introduced by Felten et al. (2021) and Tolan et al. (2021). Both contributions avoid mapping AI benchmarks to tasks (i.e., the objective dimension of work): the former links AI capabilities directly to O*NET Abilities, while the latter uses cognitive abilities as an intermediate layer between tasks and benchmarks.

Unlike Felten et al. (2021), who focus exclusively on 52 Abilities and 10 AI benchmarks, we construct a broader mapping by linking each worker requirement to a specific AI value. To do so, we rely on *AI Topics*, from which we retrieve all publications related to each Ability, Knowledge, or Skill in O*NET.

We proceed as follows. For each worker requirement, we take its textual definition from O*NET and construct a query that captures all relevant concepts, manually expanding it to include synonyms and semantically related expressions. For example, for the Ability *Inductive Reasoning*, defined as "the ability to combine pieces of information to form general rules or conclusions," the query used is:

```
"inductive reasoning" OR ("combine" AND "information") OR
("form" AND "conclusions") OR "pattern recognition".
```

This allows us to retrieve, for each year between 2011 and 2020, the number of AI-related publications concerning that specific worker requirement.

The resulting publication counts are then globally standardised to a 0–1 scale. This global standardisation accounts for both temporal variation and absolute differences across worker requirements, as some requirements naturally generate far more research outputs than others. The outcome is a first AI index defined at the worker-requirement level: each worker requirement receives a unique AI value that remains constant across occupations.

To incorporate heterogeneity across occupations, this worker-requirement AI score is subsequently weighted by the O*NET Importance score of that requirement within each occupation. Thus, for instance, occupations in which *Inductive Reasoning* is less important receive a lower AI score than occupations where that same requirement plays a central role.

Table 3.4.1: **Construction of the AI Score from AI Topics**

Step	**Description**
1. Identify worker requirements	Select all Abilities, Skills, and Knowledge items from the O*NET Content Model.
2. Extract O*NET definitions	Retrieve the textual definition of each worker requirement to build a semantic search query.
3. Construct AI Topics query	For each worker requirement, construct a Boolean query including key terms, synonyms, and conceptual expansions (e.g., for *Inductive Reasoning*).
4. Retrieve yearly publication counts	Query AI Topics for each year (2011–2020) and extract the number of AI-related publications matching the query.
5. Global standardisation	Standardise publication counts to a 0–1 scale across all worker requirements and years to capture temporal trends and absolute differences in research volume.
6. Construct worker-requirement AI index	Assign each worker requirement a yearly AI score based on its standardised publication count.
7. Weight by occupational importance	Multiply the worker-requirement AI score by O*NET *Importance* values for each occupation, generating an occupation–worker-requirement–year AI exposure measure.

Notes: The procedure is applied to all worker requirements in the O*NET Content Model and repeated annually. The resulting index captures the intensity of AI-related research associated with each worker requirement over time.

Construction of the AI Score from LLM Stats

A second AI score is constructed using data from LLM Stats. This index is important not only as a robustness check but also because it captures a conceptually distinct dimension of artificial intelligence. While the AI Topics measure reflects developments in general artificial intelligence research, the LLM Stats index is built directly from the performance of large language models and therefore captures advances in machine learning systems that learn from data to improve their predictive and generative capabilities.

From LLM Stats we download all 383 available AI benchmarks, each of which reports the performance of major models on a task-specific evaluation metric. For each benchmark and each company, we select the highest score achieved among all models released by that company. This identifies the yearly technological frontier for each firm (e.g., OpenAI, Anthropic, and others). Since data are available only for the period 2023–2025, we compute, for each year, the average of these frontier values across all companies to obtain an annual benchmark-level performance measure.

Because all benchmark scores are already expressed on a 0–100 scale, the resulting dataset assigns to each benchmark a directly comparable value of machine-learning capability for each available year. The next step is to link each benchmark to a specific worker requirement. To link each benchmark to a specific worker requirement, we compare the textual descriptions of benchmarks with the O*NET definitions of Abilities, Skills, and Knowledge. We construct a semantic similarity measure linking textual descriptions of AI capabilities (benchmarks) to human abilities, skills, and knowledge described in the O*NET database. The goal is to obtain a continuous score capturing how closely each AI benchmark corresponds to a specific worker requirements, following the method presented by Montobbio et al. (2021, 2024). Let $B = \{b_1, \ldots, b_{N_B}\}$ denote the set of benchmarks and $O = \{o_1, \ldots, o_{N_O}\}$ the set of O*NET elements. For each benchmark b_i, we concatenate its name and description into a single text string $T(b_i)$, and analogously for each O*NET element we use the combined descriptive text $T(o_j)$. All texts undergo minimal normalization consisting of lowercasing and removal of irregular whitespace, with no stemming, lemmatization, or stopword filtering, since such operations tend to distort semantic representations.

Each text is then mapped into a dense semantic vector using a pre-trained Sentence-BERT model (Reimers and Gurevych, 2019), specifically the *all-MiniLM-L6-v2* architecture widely used in semantic textual similarity applications. Denoting the embedding of a text x as $E(x)$, we L2-normalize all vectors to obtain $\hat{E}(x) = E(x)/\|E(x)\|$. Semantic similarity between benchmark b_i and O*NET element o_j is computed as cosine similarity between their normalized embeddings:

$$S_{ij} = \hat{E}(b_i) \cdot \hat{E}(o_j) = \sum_{k=1}^{d} \hat{E}(b_i)_k \, \hat{E}(o_j)_k,$$

where d is the embedding dimension. Since vectors are normalized, $S_{ij} \in [-1, 1]$ and measures the conceptual proximity of the two textual descriptions. Higher values of S_{ij} indicate that the linguistic and semantic content of the benchmark closely resembles that of the corresponding human ability, skill, or knowledge item. The resulting mapping dataset consists of tuples (b_i, o_j, S_{ij}) for all j in the selected top-K set for each i. These scores provide a continuous measure of the conceptual alignment between AI benchmarks and worker requirements, and are accompanied by ancillary indicators such as shared lexical keywords, which we compute solely for interpretability and do not use in the construction of the exposure index. In particular, we do not incorporate the similarity score S_{ij} itself in the definition of AI exposure: once the mapping between benchmarks and worker requirements is established, the similarity measure is used only as a selection device, rather than as a weight in the aggregation. This ensures that the level of exposure of a given worker requirement to AI reflects the underlying capability of the associated benchmarks, rather than mechanically mirroring the scale of the similarity metric.

After constructing the benchmark–worker-requirement links, we take the AI score previously assigned to each benchmark and transfer it to the corresponding worker requirements via a simple crosswalk. When multiple benchmarks are associated with the same worker requirement, we aggregate them by averaging the benchmark-level AI scores to obtain a single AI value for each O*NET element.

As in the case of the AI Topics index, we then translate these worker-requirement values into an occupation-level measure. For each occupation, the AI score of every worker requirement is weighted by its O*NET Importance value, so that requirements that are more central to an occupation contribute proportionally more to its overall exposure. Finally, the resulting occupation–worker-requirement scores are standardised on a 0–1 scale to ensure comparability across occupations and across years.

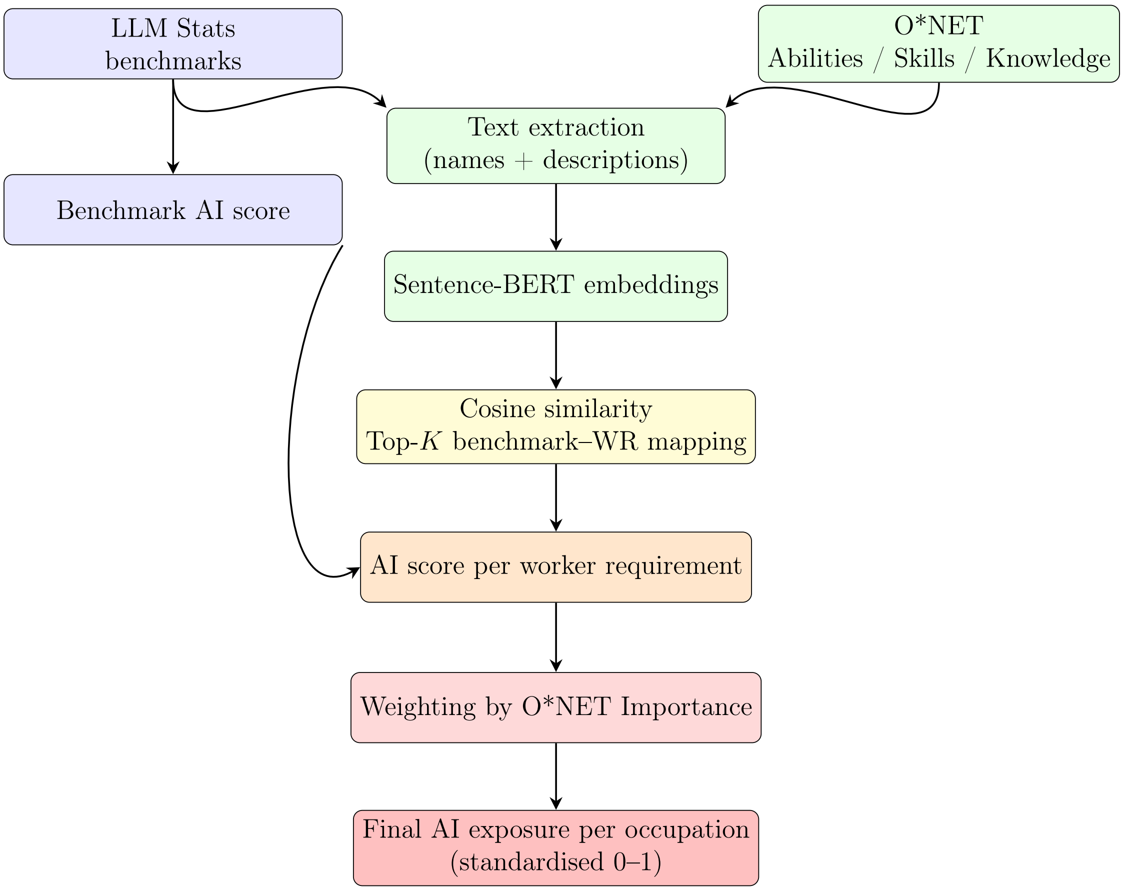

Figure 3.4.1: **Construction of the LLM Stats–based AI exposure index**

3.4.2 Data Integration and Harmonization

Before proceeding with the analysis, we merge the dataset structured at the *occupation–worker requirement–year* level with labor market information from the *Current Population Survey* (CPS). The CPS provides data on employment shares and changes, wages, industry sector, educational attainment, age, and racial composition of workers. Since CPS microdata are coded using the 2010 six-digit Standard Occupational Classification (SOC) system, whereas O*NET relies on the O*NET-SOC taxonomy, we harmonize the two sources by applying an official crosswalk between SOC 2010 and O*NET-SOC codes. This allows for a consistent integration of CPS variables into our occupation-level dataset.

3.5 Empirical Strategy and Preliminary Evidence

We now explore the relationship between occupational exposure to Artificial Intelligence and the evolution of Abilities, Knowledge, and Skills between 2011 and 2025.

To do so, we estimate the following regression at the worker–requirement–occupation level:

$$y_{wr,o} = \beta\, x_{wr,o} + \gamma_o + \varepsilon_{wr,o}, \tag{1}$$

where the dependent variable $y_{wr,o}$ is the change in the standardized and importance-weighted level of the worker requirement between 2011 and 2025, and $x_{wr,o}$ is the AI exposure index previously constructed, either using AI topics or LLM statistics. The term γ_o denotes occupation fixed effects, and standard errors are clustered at the worker-requirement level.

The change in each worker requirement is computed as the difference between its standardized, importance-weighted level in 2025 and 2011. When one of the two years was missing, we applied an iterative procedure: if the 2025 value was not available, we used the 2024 value, and so on; similarly, if the 2011 value was missing, we used the 2012 value. To avoid generating differences very close to zero, we excluded occupations whose most recent update of level or importance occurred before 2013. We also removed observations with a level equal to zero in either 2011 or 2025, since such requirements are unlikely to exhibit meaningful variation.

Finally, we rely on simple differences rather than growth rates. Percentage changes would inflate small variations for low-level worker requirements while underweighting equally meaningful changes for those starting from medium or high levels.

The variable $X_{wr,o}$ measures the potential exposure of each worker–requirement–occupation unit (wr, o) to AI, as described in Section 3. As already discussed, these measures capture the extent to which a given worker requirement is exposed to Artificial Intelligence within an occupation. We rely on two alternative measures: one based on AI topics and one based on LLM statistics. Since our focus is on how worker requirements evolve over time, these indices are not designed to capture the degree of complementarity or substitutability between AI and tasks or occupations, in the spirit of Pizzinelli et al. (2023). Rather, they are intended to measure a more subjective and intangible dimension of exposure.

The underlying idea is that the more exposed a worker requirement is to AI, the more it is at risk of becoming "impoverished", irrespective of whether the occupation as a whole is substituted by AI or not. In other words, there may be occupations that gain employment share over time but, internally, become less intensive in specific worker requirements. We therefore interpret the coefficient β in equation (8) as an index of enrichment or impoverishment of worker requirements. A positive value of β indicates that worker requirements with higher exposure to AI experience larger increases in their standardized, importance-weighted level between 2011 and 2025.

Conversely, a negative β implies that worker requirements that are more exposed to AI tend to experience a decline in their level over time.

In addition to the worker–requirement–level specification, we also estimate a macro-micro-occupation version of the regression, obtained by aggregating all variables at the occupation level. Formally, we estimate:

$$y_o = \beta X_o + \Gamma' Z_o + \varepsilon_o, \tag{2}$$

where y_o is the change in the (standardized, importance-weighted) level of the worker requirements at the occupation level, X_o is the occupation-level AI exposure index , $L_{o,2011}$ captures the initial level of the worker requirements , and Z_o is a vector of controls including the initial level of the worker requirements, education shares, age composition, demographic characteristics, and sector fixed effects

In this specification, β again measures enrichment versus impoverishment, but now at the occupation level: a positive (negative) β indicates that occupations with higher exposure to AI tend to become more (less) intensive in the underlying worker requirements on average. This provides a more "macro–micro" perspective, as it no longer focuses on the distribution of skills within occupations, but rather on the evolution of occupations themselves.

3.5.1 Descriptive Evidence and Stylized Facts

This section provides descriptive statistics for the AI exposure measures and documents the evolution of Abilities, Knowledge, and Skills over time. Table 3.5.1 reports summary statistics for the two measures of Artificial Intelligence exposure (the variable x), at the worker requirements-by-occupation level.

Table 3.5.1: **Summary Statistics for AI Exposure Measures (Worker requirement–occupation level)**

Variable	**Obs.**	**Mean**	**Std. Dev.**	**Min**	**Max**
AI LLM	78,117	0.230	0.120	0	1
AI Topics	74,147	0.214	0.090	0	0.82

Notes: The table reports summary statistics for the two AI exposure measures at the worker requirement–occupation level.

The two measures are available for 78,177 and 74,147 cells in our dataset. The AI LLM index displays a slightly higher average value than the alternative measure, consistent with the fact that the latter has a lower upper bound, with a maximum of 0.82.

Table 3.5.2 reports the same set of summary statistics, but in this case computed at the occupation level.

Table 3.5.2: **Summary Statistics for AI Exposure Measures at the Occupation Level**

Variable	Obs.	Mean	Std. Dev.	Min	Max
AI LLM	797	0.350	0.180	0	1
AI Topics	797	0.390	0.150	0	1

Notes: The table reports summary statistics for the two AI exposure measures at the occupation level.

The descriptive statistics show clear differences between the AI exposure measures computed at the worker-requirement–occupation level and those obtained after aggregating the dataset to the occupation level.

The first set of values (78,117 and 74,147 observations) corresponds to the exposure indices calculated for each worker requirement within each occupation. These measures reflect how strongly individual Abilities, Skills, and Knowledge elements are exposed to AI. At this granular level, both indices exhibit relatively low average exposure (0.23 for AI_LLM and 0.21 for AI Topics), which is consistent with the fact that most worker requirements have limited direct association with AI-related benchmarks. After aggregating the dataset to the occupation level, we compute an occupation-level AI score by averaging the worker-requirement exposure values within each occupation. The distribution of the indices changes as a mechanical result of this aggregation and because the two measures rely on different numerical scales: the mean AI_LLM exposure at the occupation level is 0.35, while the AI Topics measure shows a substantially higher mean of 0.39

Thus, differences between the two sets of statistics do not reflect differences in the underlying concept of exposure, but rather the level of aggregation.

Table 3.5.3: **Correlation Between AI Exposure Measures**

	AI LLM	AI Topics
AI LLM	1.000	0.243
AI Topics	0.243	1.000

Notes: The table reports the correlation between the two AI exposure measures at the worker requirement–occupation level.

Table 3.5.4: **Correlation Between Occupation-Level AI Exposure Measures**

	AI LLM	**AI Topics**
AI LLM	1.000	0.786
AI Topics	0.786	1.000

Notes: The table reports the correlation between the two AI exposure measures at the occupation level.

Looking at Tables 3.5.3 and 3.5.4, which report the correlations between the two AI exposure measures at each level of aggregation, we observe a clear pattern. At the worker-requirement–occupation level—where both indices are originally constructed—the two measures display a relatively low correlation, reflecting the conceptual difference between the LLM-based (narrow AI) and the AI-Topics (general AI) indicators. Once we aggregate the data to the occupation level, however, the correlation between the two measures increases substantially. This suggests that part of the difference between the indices dissipates when moving from fine-grained worker requirements to broader occupational aggregates.

Tables 3.5.5 and 3.5.6 report the five most and least exposed occupations for both AI exposure measures at the worker–requirement–occupation level. These rankings already indicate that the two indices capture different dimensions of AI exposure. For the LLM-based measure, the most exposed worker requirement is *Deductive Reasoning*, a highly cognitive ability typically associated with high-skill occupations, consistent with recent evidence on the types of tasks most affected by AI (Felten et al., 2021; Lassébie and Quintini, 2022). Conversely, the least exposed cases correspond to knowledge domains that are either peripheral or not relevant for the occupations in which they appear.

By contrast, the AI Topics measure identifies *Stamina* as one of the most exposed worker requirements. According to O*NET, stamina refers to "the ability to exert yourself physically over long periods of time without getting winded or out of breath," a feature more commonly associated with exposure to broader technological families such as automation or robotics. This pattern confirms that the AI Topics index captures a more general notion of artificial intelligence exposure, whereas the LLM-based index reflects a narrower focus on cognitive and reasoning-intensive tasks.

We also construct the corresponding tables after aggregating the data at the occupation level. The resulting rankings, reported in Tables 3.5.7 and 3.5.8, show the most and least exposed occupations according to the two AI exposure measures once worker requirements are collapsed to the occupational dimension.

Table 3.5.5: **Occupations with the Highest and Lowest AI LLM Exposure**

Panel A: Highest AI LLM Exposure		
Occupation	**Worker Requirement**	**AI LLM**
Judges, Magistrate Judges, and Magistrates	Deductive Reasoning	1.000
Clinical Neuropsychologists	Deductive Reasoning	0.934
Preventive Medicine Physicians	Deductive Reasoning	0.934
Radiologists	Deductive Reasoning	0.934
Education Administrators, K–12	Deductive Reasoning	0.898
Panel B: Lowest AI LLM Exposure		
Occupation	**Worker Requirement**	**AI LLM**
Fundraisers	Physics	0.000
Human Resources Assistants, Except Payroll and Timekeeping	Physics	0.000
Bicycle Repairers	Medicine and Dentistry	0.000
Patternmakers, Wood	Food Production	0.000
Bicycle Repairers	Biology	0.000

Notes: The table reports the five occupations with the highest and lowest AI LLM exposure for the corresponding worker requirement. Deductive Reasoning is classified as an Ability, whereas Physics, Medicine and Dentistry, Food Production, and Biology are classified as Knowledge domains.

Table 3.5.6: **Occupations with the Highest and Lowest AI Topics Exposure**

Panel A: Highest AI Topics Exposure		
Occupation	**Worker Requirement**	**AI Topics**
Farmworkers and Laborers, Crop, Nursery, and Greenhouse	Stamina	0.815
Stonemasons	Stamina	0.722
Structural Iron and Steel Workers	Stamina	0.690
Reinforcing Iron and Rebar Workers	Stamina	0.657
Heat Treating Equipment Setters, Operators, and Tenders, Metal and Plastic	Stamina	0.627
Panel B: Lowest AI Topics Exposure		
Occupation	**Worker Requirement**	**AI Topics**
Model Makers, Metal and Plastic	Medicine and Dentistry	0.000
Bookkeeping, Accounting, and Auditing Clerks	History and Archeology	0.000
Septic Tank Servicers and Sewer Pipe Cleaners	Fine Arts	0.000
Industrial Production Managers	Fine Arts	0.000
Speech-Language Pathologists	Building and Construction	0.000

Notes: The table reports the five worker requirement–occupation cells with the highest and lowest AI Topics exposure. Stamina is classified as an Ability, whereas Medicine and Dentistry, History and Archeology, Fine Arts, and Building and Construction are classified as Knowledge domains.

Table 3.5.7: **Highest and Lowest Occupation-Level Exposure According to AI LLM**

Panel A: Highest AI LLM Exposure	
Occupation	**AI LLM**
Mathematicians	1.000
Statisticians	0.890
Mathematical Science Teachers, Postsecondary	0.865
Biostatisticians	0.858
Actuaries	0.842
Panel B: Lowest AI LLM Exposure	
Occupation	**AI LLM**
Dishwashers	0.017
Fence Erectors	0.016
Roof Bolters, Mining	0.007
Foundry Mold and Coremakers	0.006
Pressers, Textile, Garment, and Related Materials	0.000

Notes: The table reports the five occupations with the highest and lowest values of AI LLM, the occupation-level AI exposure index obtained by aggregating worker requirement exposures.

Table 3.5.8: **Highest and Lowest Occupation-Level Exposure According to AI Topics**

Panel A: Highest AI Topics Exposure	
Occupation	**AI Topics**
Mathematicians	1.000
Financial Quantitative Analysts	0.922
Statisticians	0.874
Astronomers	0.870
Mathematical Science Teachers, Postsecondary	0.868
Panel B: Lowest AI Topics Exposure	
Occupation	**AI Topics**
Dishwashers	0.071
Meat, Poultry, and Fish Cutters and Trimmers	0.069
Cleaners of Vehicles and Equipment	0.025
Fast Food and Counter Workers	0.016
Pressers, Textile, Garment, and Related Materials	0.000

Notes: The table reports the five occupations with the highest and lowest values of the AI Topics index. Higher values indicate greater exposure based on topic-model similarity to AI-related domains.

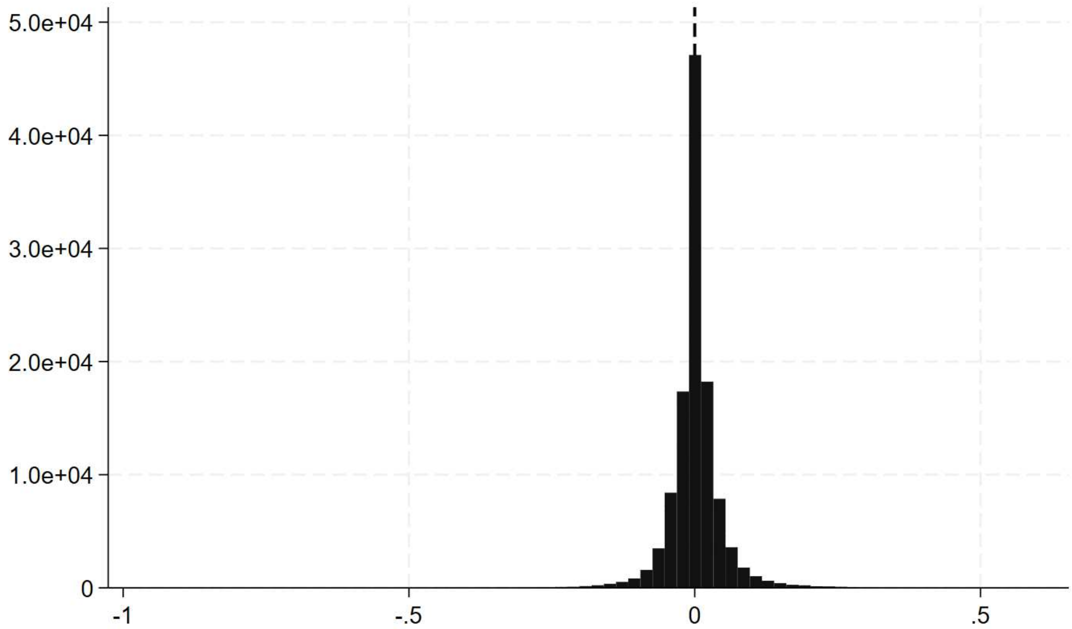

Figure 3.5.1: **Distribution of Changes in Workers Requirements , 2011–2025**

Notes: Positive values indicate increases in required worker-requirement levels, while negative values indicate decreases.

Although the distribution of workers requirements is centered close to zero, Figure 3.5.1 highlights a substantial dispersion in the evolution of worker requirements between 2011 and 2025. The presence of pronounced tails indicates that several skills experienced meaningful increases or decreases over time. This variability confirms that, despite an overall stable core, occupations underwent significant internal adjustments in their abilities, knowledge and skills composition.

To better capture how worker requirements have evolved across occupations, we aggregate SOC major groups into four macro-occupational categories (cognitive/high-skilled; social and service; sales and office; manual and physical). The pronounced blue-to-red gradient across workers requirements confirms that both the direction and magnitude of changes differ substantially between occupational domains. Across all groups, the heatmap reveals a clear predominance of red shading, indicating widespread declines in required skills, abilities, and knowledge between 2011 and 2025.

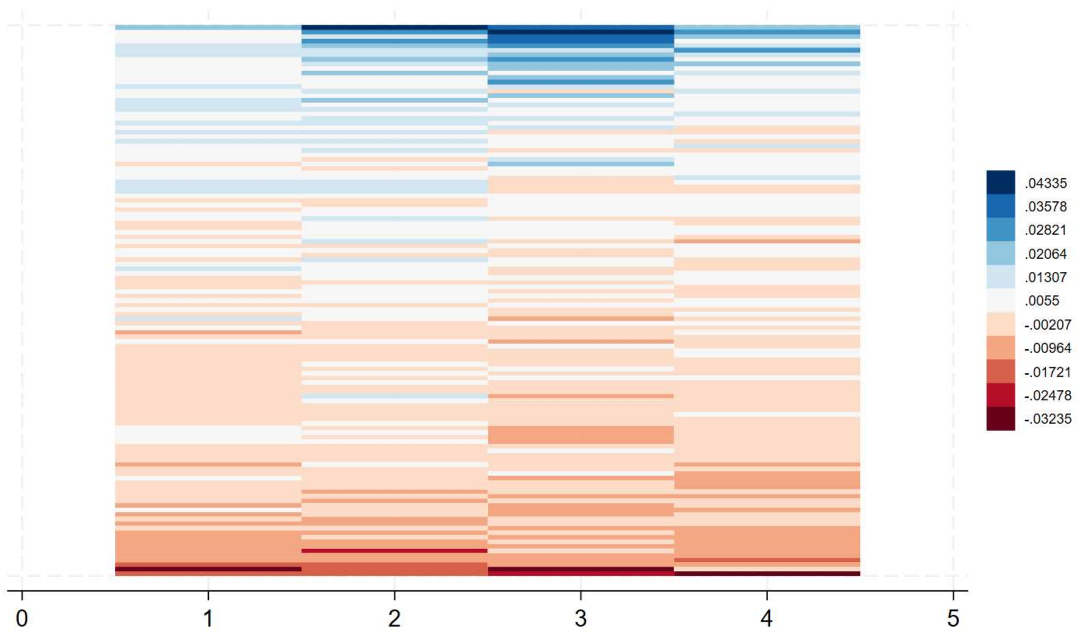

Figure 3.5.2: **Heatmap of Changes in Worker Requirements by Macro-Occupational Group (2011–2025)**

Notes: Macro-occupational groups are constructed by aggregating SOC major groups as follows: (1) cognitive/high-skilled (code 11, 13, 15, 17, 19, 23, 25, 27); (2) social and service (code 21, 29, 31, 33, 39); (3) sales and office (code 41, 43); (4) manual and physical occupations (code 35, 37, 45, 47, 49, 51, 53). Each row represents an ONET abilities, knowledge or skills, ordered from the largest to the smallest average change in worker requirements between 2011 and 2025. Colors indicate the mean change in worker requirements for each group combination

3.6 Results

3.6.1 Results by occupation–worker requirements level

Table 3.6.1 reports the estimates from regression (1). We find a positive association between the change in worker requirements over time and the exposure of abilities, skills, and knowledge within each occupation. This result holds for both the exposure index constructed from LLM statistics—which captures exposure to machine-learning models—and the more general AI exposure index derived from AI Topics. According to the LLM-based AI exposure indicator, in the United States moving from the 25th to the 50th percentile of the exposure distribution is associated with an increase in worker requirements of about 0.5%, while using the measure provided by AI Topics the estimated increase in worker requirements is about 3.5%, indicating a substantial rise. The finding of a positive association supports the view that, within each occupation, artificial intelligence has contributed to an enrichment of worker requirements over this period.

Whereas other studies have documented an increase in tasks compatible with AI's current capabilities (Acemoglu et al., 2020), our results suggest that worker requirements themselves have also grown. This result can be interpreted consistently with other evidence on the effects of AI (McElheran et al., 2023), which, for example, finds that AI use is highest among more educated, more experienced, and younger owners, including those motivated by bringing new ideas to market. This may imply that artificial intelligence has favored certain types of abilities, skills, and knowledge.

However, a broader issue remains. The dependent variable is constructed as the difference between the level of each specific worker requirement in 2025 and 2011, weighted by importance. The level itself comes from survey questions asking respondents to report the extent to which a given skill is used—for example, the level of inductive reasoning—without distinguishing the strictly human component from the performance of the job as a whole. Several studies have shown the complementarity between tasks and artificial intelligence (Eloundou et al., 2024; Pizzinelli et al., 2023; Walkowiak, 2025), suggesting that task execution increasingly reflects a true combination of human and AI inputs, making it difficult to isolate a purely human skill level. For instance, regarding inductive reasoning, a worker may report a higher level today because such reasoning is now performed with assistance from AI systems, rather than solely by the human worker. What remains unclear, therefore, is the evolution of the strictly human capability: in other words, whether that skill level would have improved or deteriorated in the absence of tools such as ChatGPT. Table 3.6.2 provides further confirmation of these regression results when we split the sample into Abilities, Skills, and Knowledge. In all three categories the estimated coefficients remain positive, although their magnitudes vary. This suggests that the positive association between AI exposure and the evolution of worker requirements is robust across different types of occupational attributes.

Table 3.6.1: **2011–2025 Variation in Abilities, Skills, and Knowledge**

	(1)	**(2)**
AI Topics	0.070***	0.070***
	(0.001)	(0.0018)
Observations	74,147	74,147
AI LLM	0.020***	0.019***
	(0.0055)	(0.0052)
Observations	78,117	78,117

Source: Authors' calculations.
Notes: The dependent variable measures the change in worker-requirement levels (Abilities, Skills, and Knowledge) between 2011 and 2025. Standard errors are clustered at the worker-requirement level and reported in parentheses. Regression (1) includes occupation fixed effects (O*NET-SOC 6-digit), while regression (2) excludes occupation fixed effects.
Significance levels: $*p < 0.05$; $**p < 0.01$; $***p < 0.001$.

Table 3.6.2: **2011–2025 Variation in Abilities, Skills, and Knowledge**

	Abilities (A)	**Knowledge (K)**	**Skills (S)**
AI Topics	0.0004***	0.0031***	0.0006***
	(0.0001)	(0.0010)	(0.0002)
Observations	32,674	18,066	23,407
AI LLM	0.019***	0.061**	0.037***
	(0.003)	(0.019)	(0.008)
Observations	32,674	21,928	25,728

Source: Authors' calculations.
Notes: The dependent variables measure changes in worker-requirement levels between 2011 and 2025, separately for Abilities, Skills, and Knowledge. Standard errors are clustered at the worker-requirement level and reported in parentheses.
Significance levels: $*p < 0.05$; $**p < 0.01$; $***p < 0.001$.

In Figures 3.6.1 and 3.6.2, we further disaggregate the analysis by estimating separate regressions for each worker requirement with a sufficient number of observations, and we display the results using forest plots. In these figures, the red dots represent the estimated coefficients, while the horizontal whiskers denote the confidence intervals. The plots confirm that the positive impact of artificial intelligence (according to both exposure indices) holds for many individual worker requirements.

The only negative or near-zero effects emerge for *physical abilities*—which is unsurprising, as firms adopting AI rely more heavily on digitized information and cloud computing, suggesting complementarities with other "enabling" technologies (Kapoor and Teece, 2021; McElheran et al., 2023)—and, more interestingly, for *system skills* (which include, for example, Judgment and Decision Making) and for some *process skills* (such as Critical Thinking, Active Learning, and Learning Strategies). This pattern appears consistent with the perspective advanced by Gerlich (2025), although this interpretation is not fully supported in Figure 5, which reports the estimates based on the AI–LLM exposure index.

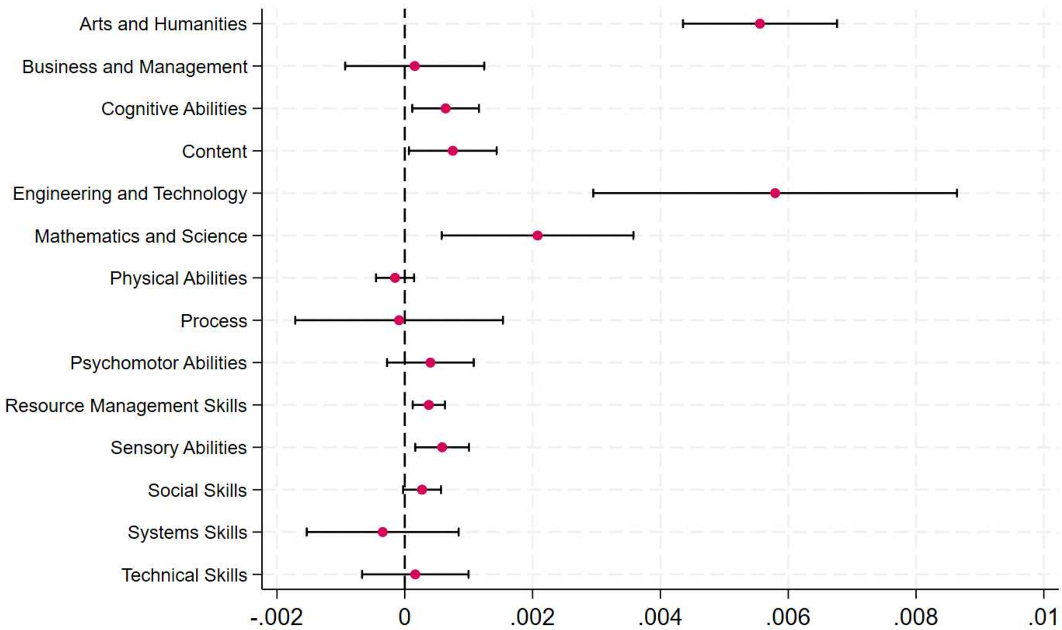

Figure 3.6.1: **Forest Plot of AI–Topics Exposure Effects on Individual Worker Requirements**

Note: Each red dot represents the estimated coefficient from a separate regression of the change in a given worker requirement on the AI–Topics exposure index, controlling for occupation fixed effects and clustering standard errors at the worker-requirement level. Horizontal whiskers show the 95% confidence intervals. Only worker requirements with a sufficient number of observations are included. Positive estimates indicate stronger increases in required worker requirements levels between 2011 and 2025 for worker requirements more exposed to AI Topics.

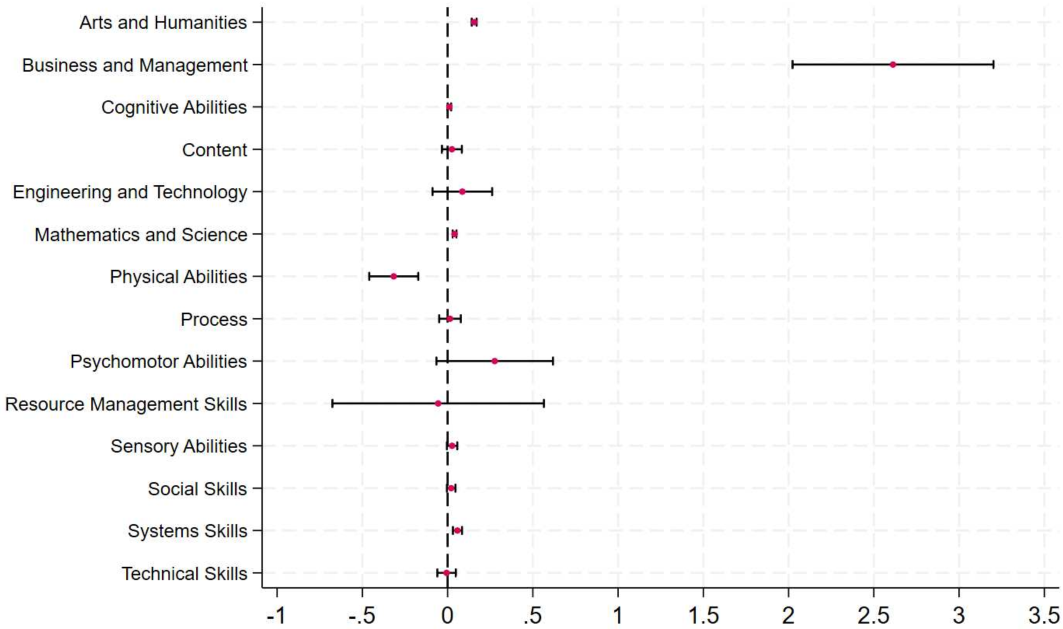

Figure 3.6.2: **Forest Plot of LLM-Based AI Exposure Effects on Individual Worker Requirements**

Note: Each red dot represents the estimated coefficient from a separate regression of the change in a given worker requirement on the AI–LLM exposure index, controlling for occupation fixed effects and clustering standard errors at the worker-requirement level. Horizontal whiskers show the 95% confidence intervals. Only worker requirements with a sufficient number of observations are included. Positive estimates indicate stronger increases in required worker requirements levels between 2011 and 2025 for worker requirements more exposed to AI Topics.

3.6.2 Results at the occupational level

Finally, we aggregate our dataset at the occupational level to examine whether the occupations most exposed to AI are also those whose worker requirements experienced enrichment or erosion. In contrast to the previous subsection—where the analysis was conducted at a micro–micro level—this section adopts what can be described as a macro–micro perspective. As shown in Table 3.6.3, which reports in the first column the estimates based on the full dataset and in the subsequent columns the results for Abilities, Knowledge, and Skills separately, the picture changes substantially. In this case the coefficients become negative. Moving from the 25th to the 50th percentile of the AI exposure distribution is associated with a decrease in worker requirements of about 6.5% when using the AI Topics indicator and about 3.5% when using the LLM-based exposure index.

The negative occupational-level effect of AI exposure is largely driven by Abilities, while Skills and Knowledge show no statistically significant adjustment. This heterogeneity across requirement types is consistent with the notion that AI primarily reshapes cognitive and reasoning-intensive (Felten et al., 2021) dimensions of work rather than broader skill or knowledge domains. The positive coefficients at the worker-requirement level capture how AI exposure is associated with increases in some specific requirements within occupations. However, when aggregating to the occupational level, these heterogeneous within-occupation adjustments are averaged, and the dominant pattern becomes a net reduction in overall worker-requirement levels for more AI-exposed occupations. This divergence is therefore consistent with AI simultaneously raising certain detailed requirements while reducing the aggregate intensity of requirements at the occupation level.

Table 3.6.3: **2025 Cross-Sectional Association Between AI Exposure and Tasks, Abilities, Skills, and Knowledge**

	Tasks (T)	**Abilities (A)**	**Knowledge (K)**	**Skills (S)**
AI Topics	-0.260***	-0.130**	-0.090	-0.039
	(0.065)	(0.059)	(0.059)	(0.102)
Observations	372	372	372	372
AI LLM	-0.140**	-0.120**	-0.018	-0.077
	(0.065)	(0.057)	(0.058)	(0.143)
Observations	372	372	372	372

Source: Authors' calculations.
Notes: The table reports cross-sectional estimates for 2025. All specifications control for baseline worker-requirement levels, occupational composition by education, age, gender, and race, and include controls for fifteen macro-industry shares. Robust standard errors are reported in parentheses.
Definitions: T = Tasks; A = Abilities; S = Skills; K = Knowledge.
Significance levels: *$p < 0.05$; **$p < 0.01$; ***$p < 0.001$.

3.7 Conclusion

In this paper, we explore how exposure to artificial intelligence shapes variation in worker requirements—abilities, skills, and knowledge—both within occupations and across occupations. Rather than examining the impact of AI on employment or the redefinition of tasks, as emphasized in much of the existing literature, our focus is on the evolution of the subjective determinants of work, a dimension that has received comparatively little attention. To do so, we construct two AI exposure measures at the worker-requirement level: one capturing exposure to general-purpose AI using AI Topics, and the other capturing exposure to narrow, model-based AI using LLM statistics.

Our results show that, at the worker requirement–occupation level, AI exposure has a positive effect: within occupations, the worker requirements most exposed to AI tend to increase over time. This pattern is broadly consistent with task-based evidence suggesting complementarity between AI and certain cognitive dimensions of work Acemoglu et al. (2020); Pizzinelli et al. (2023). The result is remarkably stable across nearly all individual Abilities, Skills, and Knowledge items in O*NET.

However, when we aggregate the analysis to the occupational level, the pattern reverses. The positive within-occupation effect disappears, and more AI-exposed occupations exhibit a markedly stronger decline in overall worker requirements. This divergence suggests the following interpretation: while AI elevates the requirements associated with the subset of abilities, skills, and knowledge that are directly complemental to it, the net effect across the entire occupation is a progressive reduction in the average requirement level. In other words, AI appears to *reshape the distribution of worker requirements from within*, enriching specific dimensions while contributing to a general erosion of occupational content. This paper opens the door to future research aimed at investigating the subjective content of work. For this reason, we did not adopt a task-based approach but instead relied on a skill-based framework; the challenge ahead is to develop methodologies capable of integrating these two perspectives without erasing the subjective dimension of worker requirements. Moreover, it remains crucial to understand how LLM-based systems will further reshape the contours of worker requirements as their adoption expands, given that we are only at the beginning of their implementation. Additional research is also needed to trace the long-run evolution of abilities, skills, and knowledge, for instance by identifying and analyzing specific clusters within each of these domains, and to examine how these components of work may continue to transform in response to advances in AI.

Bibliography

Acemoglu, D. (2002). Technical change, inequality, and the labor market. *Journal of Economic Literature*, 40(1):7–72.

Acemoglu, D. (2024). The simple macroeconomics of AI. NBER Working Paper 32487, National Bureau of Economic Research (NBER).

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*, volume 4, chapter 12, pages 1043–1171. Elsevier, 1 edition.

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2020). AI and jobs: Evidence from online vacancies. NBER Working Paper 28257, National Bureau of Economic Research, Cambridge, MA.

Acemoglu, D. and Johnson, S. (2023). *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity.* PublicAffAIrs.

Acemoglu, D. and Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30.

Acemoglu, D. and Zilibotti, F. (2001). Productivity differences. *The Quarterly Journal of Economics*, 116(2):563–606.

Agrawal, A., Gans, J. S., and Goldfarb, A. (2019). Artificial Intelligence: The ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives*, 33(2):31–50.

Auer, R., Köpfer, D., and Sveda, J. (2024). The rise of generative AI: Modelling exposure, substitution and inequality effects on the us labour market. BIS Working Paper 1207, Bank for International Settlements (BIS).

Autor, D. (2024). Applying AI to rebuild middle class jobs. NBER Working Paper 32140, National Bureau of Economic Research (NBER).

Autor, D. H. (2013). The "Task Approach" to labor markets: An overview. *Journal for Labour Market Research*, 46(3):185–199.

Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review*, 103(5):1553–1597.

Autor, D. H. and Handel, M. J. (2013). Putting tasks to the test: Human capital, job tasks, and wages. *Journal of Labor Economics*, 31(2):S59–S96.

Autor, D. H., Katz, L. F., and Kearney, M. S. (2006). The polarization of the u.s. labor market. *American Economic Review*, 96(2):189–194.

Autor, D. H., Katz, L. F., and Kearney, M. S. (2008). Trends in u.s. wage inequality: Revising the revisionists. *Review of Economics and Statistics*, 90(2):300–323.

Autor, D. H., Katz, L. F., and Krueger, A. B. (1998). Computing inequality: Have computers changed the labor market? *The Quarterly Journal of Economics*, 113(4):1169–1213.

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

Beaudry, P., Green, D. A., and Sand, B. M. (2016). The great reversal in the demand for skill and cognitive tasks. *Journal of Labor Economics*, 34(S1):199–247.

Berman, E., Bound, J., and Machin, S. (1998). Implications of skill-biased technological change: International evidence. *The Quarterly Journal of Economics*, 113(4):1245–1279.

Bessen, J. (2018). AI and jobs: The role of demand. NBER Working Paper 24235, National Bureau of Economic Research, Cambridge, MA.

Braverman, H. (1974). Labor and monopoly capital. *Monthly Review*, 26(3):1–.

Brynjolfsson, E., Mitchell, T., and Rock, D. (2018). What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings*, 108:43–47.

Buchanan, B. G., Eckroth, J., and Smith, R. (2013). A virtual archive for the history of AI. *AI Magazine*, 34(2):86–86.

Calvino, F., Haerle, D., and Liu, S. (2025). Is generative AI a general purpose technology? implications for productivity and policy. OECD Artificial Intelligence Papers 40, OECD Publishing, Paris.

Eisfeldt, A. L., Schubert, G., Zhang, M. B., and Taska, B. (2023). Generative AI and firm values. Ssrn working paper, SSRN.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). Gpts are gpts: Labor market impact potential of llms. *Science*, 384:1306–1308.

Esposito, P. and Stehrer, R. (2009). The sector bias of skill-biased technical change and the rising skill premium in transition economies. *Empirica*, 36(3):351–364.

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.

Felten, E. W., Raj, M., and Seamans, R. (2019). The occupational impact of Artificial Intelligence: Labour, skills, and polarization. Ssrn scholarly paper, SSRN, Rochester, NY.

Fernandez, R. M. (2001). Skill-biased technological change and wage inequality: Evidence from a plant retooling. *American Journal of Sociology*, 107(2):273–320.

Fernández-Macías, E. and Hurley, J. (2017). Routine-biased technical change and job polarization in europe. *Socio-Economic Review*, 15(3):563–585.

Galli, C. (2025). *Tecnica*. Il Mulino, Bologna.

Georgieff, A. and Hyee, R. (2021). Artificial Intelligence and employment: New cross-country evidence. OECD Social, Employment and Migration Working Papers 265, OECD Publishing.

Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies*, 15(1):6.

Gmyrek, P., Berg, J., and Bescond, D. (2023). Generative AI and jobs: A global analysis of potential effects on job quantity and quality. ILO Working Paper 995324892702676, International Labour Organization (ILO).

Goldin, C. and Katz, L. F. (2007). The race between education and technology: The evolution of u.s. educational wage differentials, 1890 to 2005. NBER Working Paper 12984, National Bureau of Economic Research (NBER).

Goos, M. and Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in britain. *The Review of Economics and Statistics*, 89(1):118–133.

Goos, M., Manning, A., and Salomons, A. (2009). Job polarization in europe. *American Economic Review*, 99(2):58–63.

Goos, M., Manning, A., and Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, 104(8):2509–2526.

Humlum, A. and Vestergaard, E. (2024). The adoption of ChatGPT. BFI Working Paper 2024-50, Becker Friedman Institute for Economics, University of Chicago.

Hutter, C. and Weber, E. (2017). The effects of skill-biased technical change on productivity flattening and hours worked. IAB Discussion Paper 32/2017, Institute for Employment Research (IAB), Nürnberg.

Kapoor, R. and Teece, D. J. (2021). Three faces of technology's value creation: Emerging, enabling, embedding. *Strategy Science*, 6(1):1–4.

Katz, L. F. and Murphy, K. M. (1992). Changes in relative wages, 1963–1987: Supply and demand factors. *The Quarterly Journal of Economics*, 107(1):35–78.

Lassébie, J. and Quintini, G. (2022). What skills and abilities can automation technologies replicate and what does it mean for workers? new evidence. OECD Social, Employment and Migration Working Papers 282, OECD Publishing.

Marglin, S. A. (1974). What do bosses do? the origins and functions of hierarchy in capitalist production. *Review of Radical Political Economics*, 6(2):60–112.

Matthes, B., Christoph, B., Janik, F., and Ruland, M. (2014). Collecting information on job tasks: An instrument to measure tasks required at the workplace in a multi-topic survey. *Journal for Labour Market Research*, 47(4):273–297.

McElheran, K., Li, J. F., Brynjolfsson, E., Kroff, Z., Dinlersoz, E., Foster, L., and Zolas, N. J. (2023). A.I. adoption in america: Who, what, and where. *Journal of Economics and Management Strategy*. Forthcoming.

Milanez, A. (2023). The impact of AI on the workplace: Evidence from oecd case studies of AI implementation. OECD Social, Employment and Migration Working Papers 289, OECD Publishing, Paris.

Mokyr, J. (2002). *The Gifts of Athena: Historical Origins of the Knowledge Economy*. Princeton University Press.

Montobbio, F., Staccioli, J., Virgillito, M. E., and Vivarelli, M. (2021). Labour-saving automation and occupational exposure: A text-similarity measure. LEM Papers Series 2021/43, Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies, Pisa, Italy.

Montobbio, F., Staccioli, J., Virgillito, M. E., and Vivarelli, M. (2024). Labour-saving automation: A direct measure of occupational exposure. *The World Economy*, 47:332–361.

National Center for O*NET Development. O*net online help: Scales, ratings, and standardized scores. O*NET OnLine. Accessed: March 4, 2026.

National Center for O*NET Development. The o*net® content model. O*NET Resource Center. Accessed: March 4, 2026.

OECD (2023). *OECD Employment Outlook 2023: Artificial Intelligence and the Labour Market*. OECD Publishing, Paris.

Pizzinelli, C., Panton, A., Tavares, M. M., Cazzaniga, M., and Li, L. (2023). Labor market exposure to AI: Cross-country differences and distributional implications. IMF Working Paper 23/216, International Monetary Fund (IMF).

Reimers, N. and Gurevych, I. (2019). Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.

Rodrigues, M., Fernández-Macías, E., and Sostero, M. (2021). A unified conceptual framework of tasks, skills and competences. JRC Working Papers on Labour, Education and Technology 2021-02, Joint Research Centre (JRC).

Spitz-Oener, A. (2006). Technical change, job tasks, and rising educational demands: Looking outside the wage structure. *Journal of Labor Economics*, 24(2):235–270.

Taylor, F. W. (1911). *The Principles of Scientific Management*. Number taylor1911 in History of Economic Thought Books. McMaster University Archive for the History of Economic Thought.

Thompson, P. (1983). Braverman and the re-discovery of the labour process. In *The Nature of Work*. Palgrave, London.

Tolan, S., Pesole, A., Martínez-Plumed, F., Fernández-Macías, E., Hernández-Orallo, J., and Gómez, E. (2021). Measuring the occupational impact of AI: Tasks, cognitive abilities and AI benchmarks. *Journal of Artificial Intelligence Research*, 71:191–236.

Walkowiak, E. (2025). Substitution and complementarity between generative AI and human creativity: Evolutionary rules governing job transformation. AvAIlable at SSRN.

Webb, M. (2019). The impact of Artificial Intelligence on the labor market. Ssrn working paper, SSRN.

Zhai, C., Wibowo, S., and Li, L. D. (2024). The effects of over-reliance on AI dialogue systems on students' cognitive abilities: A systematic review. *Smart Learning Environments*, 11:28.

Appendix

3.A Worker Requirements Described by O*NET

Table 3.A.1: Overview of O*NET Worker Requirements

Category	**General Definition**	**Main Subdomains (Examples)**
Abilities	Enduring attributes of the individual that influence performance.	*Cognitive Abilities* (verbal abilities, idea generation and reasoning, quantitative abilities, memory, perceptual and spatial abilities, attentiveness); *Psychomotor Abilities* (fine manipulation, control of movement, reaction time and speed); *Physical Abilities* (strength, endurance, flexibility, balance, coordination); *Sensory Abilities* (visual, auditory, and speech perception, including near/far vision, color discrimination, night vision, peripheral vision, depth perception, and speech recognition).

Continued on next page

Table 3.A.1 – continued from previous page

Category	General Definition	Main Subdomains (Examples)
Skills	Developed capacities that facilitate learning or the rapid acquisition and application of knowledge.	*Content Skills* (reading comprehension, active listening, writing, speaking, mathematics, science); *Process Skills* (critical thinking, active learning, learning strategies, monitoring); *Social Skills* (social perceptiveness, coordination, persuasion, negotiation, instructing, service orientation); *Complex Problem-Solving Skills*; *Technical Skills* (operations analysis, technology design, programming, installation, operation and control, maintenance, troubleshooting, repairing, quality control analysis); *Systems Skills* (judgment and decision making, systems analysis, systems evaluation); *Resource Management Skills* (time, financial, material, and personnel resource management).

Continued on next page

Table 3.A.1 – continued from previous page

Category	**General Definition**	**Main Subdomains (Examples)**
Knowledge	Organized sets of principles and facts applying to broad content domains.	*Business and Management* (administration, economics, sales and marketing, human resources); *Manufacturing and Production* (production and processing, food production); *Engineering and Technology* (computers and electronics, engineering, design, construction, mechanical systems); *Mathematics and Science* (mathematics, physics, chemistry, biology, psychology, sociology, anthropology, geography); *Health Services* (medicine and dentistry, therapy and counseling, education and training); *Arts and Humanities* (languages, fine arts, history, archaeology, philosophy, theology); *Law and Public Safety* (public safety, security, law and government); *Communications and Transportation* (telecommunications, media, transportation); *Education* (required level of education, certifications, and subject-specific educational background).

Notes: Definitions are based on the O*NET Content Model (U.S. Department of Labor).

3.B Benchmarks Used in the AI Exposure Measure

Table 3.B.1: Core Benchmarks Used to Construct Ability-, Skill-, and Knowledge-Based AI Exposure Measures

Benchmark	**Description**
GSM8K	Grade-school reasoning problems requiring multi-step arithmetic reasoning.
MATH	12,500 competition-level mathematics problems with step-by-step solutions.
MATH-500	Hard subset of MATH, covering seven mathematical subjects.
AIME 2024	Olympiad-level reasoning benchmark requiring multi-step logical inference.
BBH	Big-Bench Hard: 23 challenging reasoning tasks beyond standard model capabilities.
HellaSwag	Adversarial commonsense reasoning dataset requiring physical and causal inference.
TruthfulQA	Measures whether models avoid common human misconceptions.
ARC-C	Challenge set of grade-school science questions requiring multi-hop reasoning.
ARC-E	Easier subset of ARC, solvable via factual recall with conceptual understanding.
MMLU	57-subject multitask exam spanning humanities, STEM, and social sciences.
MMLU-Pro	Harder variant with expanded options and filtered trivial questions.
HumanEval	Code synthesis benchmark measuring functional correctness of generated programs.
HumanEval+	Enhanced version with substantially expanded test coverage for robustness.
MBPP	974 Python programming tasks solvable by entry-level programmers.
MBPP+ / EvalPlus	MBPP enriched with substantially more test cases for rigorous evaluation.
SWE-bench Verified	Human-validated GitHub issues requiring multi-file code edits.
RepoBench	Repository-level code completion and retrieval tasks.
BigCodeBench	1,140 real-world programming tasks requiring tool and library usage.
MathVista	Multimodal mathematical reasoning benchmark combining figures, charts, and equations.

Continued on next page

Table 3.B.1 – continued from previous page

Benchmark	**Description**
ScienceQA	Multimodal science questions with diagrams, explanations, and reasoning chains.
VQA-v2	Visual question answering benchmark requiring recognition and reasoning.
Video-MME	Large-scale multimodal video understanding benchmark across 30 subfields.
LVBench	Long-video understanding benchmark with videos up to two hours.
RULER	Synthetic benchmark testing long-context capabilities with multi-hop tracing.
MRCR v2	Multi-round coreference resolution requiring tracking of multiple entities in long contexts.
LongBench v2	Long-context problem-solving benchmark across single- and multi-document reasoning.

Notes: The listed benchmarks are used to proxy AI capabilities across abilities, skills, and knowledge domains, forming the basis of the LLM-based AI exposure measure.

3.C Examples of Boolean Queries Used to Search AI-Related Publications

This appendix provides representative examples of the Boolean queries used to link O*NET worker requirements (Knowledge, Skills, and Abilities) to AI-related research publications. The queries combine semantic descriptors derived from O*NET definitions with AI-specific filters applied separately in the search process.

The examples below are illustrative and not exhaustive.

Knowledge

Engineering and Technology

```
("engineering principles" OR "engineering rules" OR "engineering knowledge" OR
 "engineering information" OR "engineering data" OR "engineering facts" OR
 "engineering concepts" OR "engineering theories" OR "engineering literature"
    OR
 "engineering book" OR "engineering books" OR "engineering publications" OR
 "engineering archive" OR "engineering archives" OR "engineering documentation"
    OR
 "engineering records" OR "engineering datasets" OR "engineering methods" OR
 "engineering techniques" OR "technology principles" OR "technology rules" OR
 "technology knowledge" OR "technology methods" OR "technology techniques")
```

Economics and Accounting

```
("economics principles" OR "economics rules" OR "economics knowledge" OR
 "economics information" OR "economics data" OR "economics facts" OR
 "economics concepts" OR "economics theories" OR "economics literature" OR
 "economics book" OR "economics books" OR "economics publications" OR
 "economics archive" OR "economics archives" OR "economics documentation" OR
 "economics records" OR "economics datasets" OR "economics methods" OR
 "economics techniques" OR "accounting principles" OR "accounting rules" OR
 "accounting knowledge" OR "accounting information" OR "accounting data" OR
 "accounting facts" OR "accounting concepts" OR "accounting theories" OR
 "accounting literature" OR "accounting book" OR "accounting books" OR
 "accounting publications" OR "accounting archive" OR "accounting archives" OR
 "accounting documentation" OR "accounting records" OR "accounting datasets" OR
 "accounting methods" OR "accounting techniques" OR "banking principles" OR
 "banking rules" OR "banking knowledge" OR "finance principles" OR
 "finance rules" OR "finance knowledge")
```

Philosophy and Theology

```
("philosophy and theology principles" OR "philosophy and theology rules" OR
 "philosophy and theology knowledge" OR "philosophy and theology information"
   OR
 "philosophy and theology data" OR "philosophy and theology facts" OR
 "philosophy and theology concepts" OR "philosophy and theology theories" OR
 "philosophy and theology literature" OR "philosophy and theology book" OR
 "philosophy and theology books" OR "philosophy and theology publications" OR
 "philosophy and theology archive" OR "philosophy and theology archives" OR
 "philosophy and theology documentation" OR "philosophy and theology records"
   OR
 "philosophy and theology datasets" OR "religious texts" OR
 "philosophical texts" OR "religious literature")
```

Arts and Humanities

```
("arts and humanities principles" OR "arts and humanities rules" OR
 "arts and humanities knowledge" OR "arts and humanities information" OR
 "arts and humanities data" OR "arts and humanities facts" OR
 "arts and humanities concepts" OR "arts and humanities theories" OR
 "arts and humanities literature" OR "arts and humanities book" OR
 "arts and humanities books" OR "arts and humanities publications" OR
 "arts and humanities archive" OR "arts and humanities archives" OR
 "arts and humanities documentation" OR "arts and humanities records" OR
 "arts and humanities datasets" OR "arts and humanities methods" OR
 "arts and humanities techniques")
```

Skills

Reading Comprehension

(Semantic component only; AI-specific filters are applied separately.)

```
("reading comprehension" OR "text understanding" OR
 "natural language understanding" OR "document summarization" OR
 "question answering" OR "legal document analysis" OR
 "reading assistance" OR "understand sentences" OR
 "understand paragraphs")
```

Active Listening

```
("active listening" OR "conversational agents" OR
 "speech analysis" OR "listening comprehension" OR
 "dialogue systems" OR "customer service bots" OR
```

```
 "meeting transcription" OR "sentiment detection in speech" OR
 "give attention" OR "understand points" OR
 "ask questions" OR "avoid interruptions")
```

Critical Thinking

```
("decision support" OR "critical thinking models" OR
 "strategy evaluation" OR "policy analysis" OR
 "reasoning engines" OR "identify strengths" OR
 "identify weaknesses" OR "evaluate solutions")
```

Abilities

Manual Dexterity

```
("manual dexterity" OR
 ("move" AND "hands quickly") OR
 ("manipulate" AND "objects") OR
 "hand coordination")
```

Stamina

```
("stamina" OR
 ("exert" AND "physically") OR
 ("maintain" AND "physical activity") OR
 "sustained performance")
```

Memorization

```
("memorization" OR
 ("remember" AND "information") OR
 ("recall" AND "data") OR
 "memory encoding" OR
 "working memory")
```

Notes: These Boolean queries are illustrative examples. The full set of queries is generated programmatically from O*NET definitions and combined with AI-specific filters when querying publication databases.

3.D Top 10 Semantic Matches Between Benchmarks and Worker Requirements

As described in the Methodology section, we construct a semantic similarity measure linking textual descriptions of AI benchmarks to O*NET worker requirements (Abilities, Skills, and Knowledge). It is not a concern that certain worker requirements— such as Deductive Reasoning—appear multiple times among the top semantic matches. The semantic similarity procedure is used solely to identify conceptual links between benchmarks and O*NET requirements. As explained above, the AI exposure score for each worker requirement is ultimately computed as the average contribution of all benchmarks associated with it; therefore, repeated occurrences in the similarity ranking do not mechanically inflate its exposure measure.

Below we report the top semantic matches. For each pair, we show the benchmark name and description, together with the matched worker requirement and its definition.

Table 3.D.1: Top Semantic Similarity Links Between AI Benchmarks and O*NET Worker Requirements

Benchmark	**Benchmark Description**	**Matched Worker Requirement**
HMMT 2025	Harvard–MIT Mathematics Tournament 2025, a prestigious student-organized mathematics competition with individual tests, team rounds, and guts rounds held at MIT and Harvard.	**Deductive Reasoning** — The ability to apply general rules to specific problems to produce answers that make sense.
Natural2Code	NaturalCodeBench (NCB), a benchmark capturing complex real-world coding tasks across 402 problems in Python and Java derived from natural user queries.	**Idea Generation and Reasoning Abilities** — Abilities that influence the application and manipulation of information in problem solving.
Natural2Code	Same benchmark as above.	**Speech Recognition** — The ability to identify and understand the speech of another person.

Continued on next page

Table 3.D.1 – continued from previous page

Benchmark	Benchmark Description	Matched Worker Requirement
Natural2Code	Same benchmark as above.	**Speed of Closure** — The ability to quickly make sense of, combine, and organize information into meaningful patterns.
DS-FIM-Eval	DeepSeek's internal fill-in-the-middle evaluation dataset for assessing code completion performance in data science contexts.	**Deductive Reasoning** — The ability to apply general rules to specific problems to produce answers that make sense.
HumanEval+	Enhanced version of HumanEval with substantially expanded test coverage using EvalPlus.	**Idea Generation and Reasoning Abilities** — Abilities that influence the application and manipulation of information in problem solving.
WMT24++	Multilingual machine translation benchmark covering 55 languages and multiple content domains.	**Memorization** — The ability to remember information such as words, numbers, pictures, and procedures.
HMMT 2025	Same mathematics competition as above.	**Idea Generation and Reasoning Abilities** — Abilities that influence the application and manipulation of information in problem solving.
SWE-Bench Verified	Human-validated software engineering benchmark based on real GitHub issues.	**Speed of Closure** — The ability to quickly make sense of, combine, and organize information into meaningful patterns.

Continued on next page

Table 3.D.1 – continued from previous page

Benchmark	**Benchmark Description**	**Matched Worker Requirement**
CodeForces	Competitive programming benchmark using algorithmic problems from the CodeForces platform.	**Deductive Reasoning** — The ability to apply general rules to specific problems to produce answers that make sense.

Notes: These matches reflect semantic similarity only and do not mechanically affect the exposure index of a given worker requirement, since AI exposure scores are computed as averages across all benchmarks linked to that requirement.

3.E Top Five and Bottom Five Worker Requirements by Exposure to the AI Topics Index

This appendix reports the worker requirements with the highest and lowest exposure to the AI Topics index. The top five worker requirements all display an exposure value equal to one, while the bottom five exhibit substantially lower exposure levels.

Table 3.E.1: Top Five and Bottom Five Worker Requirements by Exposure to the AI Topics Index

Worker Requirement	Exposure Value	Description
Top Five		
Active Learning (S)	1.00	Understanding the implications of new information for both current and future problem-solving and decision-making.
Arm–Hand Steadiness (A)	1.00	The ability to keep hand and arm steady while moving or holding them in one position.
Deductive Reasoning (A)	1.00	The ability to apply general rules to specific problems to produce logically coherent answers.
Mathematical Reasoning (A)	1.00	The ability to choose appropriate mathematical methods or formulas to solve a problem.
Judgment and Decision Making (S)	1.00	Considering the relative costs and benefits of potential actions to choose the most appropriate one.
Bottom Five		
Installation (S)	0.00	Installing equipment, machines, wiring, or programs to meet specifications.
Customer and Personal Services (K)	0.05	Knowledge of principles and processes for providing customer and personal services, including customer assessment and service quality evaluation.

Continued on next page

Table 3.E.1 – continued from previous page

Worker Requirement	**Exposure Value**	**Description**
Personnel and Human Resources (K)	0.07	Knowledge of principles and procedures for personnel recruitment, selection, training, compensation, and labor relations.
Dynamic Strength (A)	0.31	The ability to exert muscle force repeatedly or continuously over time.
Equipment Selection (S)	0.33	Determining the kind of tools and equipment needed to perform a job.

Notes: Letters in parentheses indicate the O*NET category of each worker requirement: (S) = Skill; (A) = Ability; (K) = Knowledge.

3.F Top Five and Bottom Five Worker Requirements by Exposure to the AI LLM Index

This appendix reports the worker requirements with the highest and lowest exposure to the AI LLM index. The top five worker requirements display relatively high exposure values, while the bottom five exhibit minimal or no exposure.

Table 3.F.1: Top Five and Bottom Five Worker Requirements by Exposure to the AI LLM Index

Worker Requirement	Exposure Value	Description
Top Five		
Deductive Reasoning (A)	1.00	The ability to apply general rules to specific problems to produce logically coherent answers.
Speed of Closure (A)	0.70	The ability to quickly make sense of, combine, and organize information into meaningful patterns.
Oral Comprehension (A)	0.57	The ability to listen to and understand information and ideas presented through spoken words and sentences.
Mathematical Reasoning (A)	0.56	The ability to choose appropriate mathematical methods or formulas to solve a problem.
Speech Recognition (A)	0.55	The ability to identify and understand the speech of another person.
Bottom Five		
Education and Training (K)	0.00	Knowledge of principles and methods for curriculum design, teaching, instruction, and training effectiveness.
Dynamic Strength (A)	0.001	The ability to exert muscle force repeatedly or continuously over time.
Sociology and Anthropology (K)	0.002	Knowledge of group behavior, societal trends, cultural differences, and human origins.

Continued on next page

Table 3.F.1 – continued from previous page

Worker Requirement	Exposure Value	Description
Public Safety and Security (K)	0.003	Knowledge of equipment, policies, procedures, and strategies for effective protection of people, data, and property.
Chemistry (K)	0.03	Knowledge of chemical composition, structure, properties, processes, and transformations of substances.

Notes: Letters in parentheses indicate the O*NET category of each worker requirement: (A) = Ability; (K) = Knowledge.