U.PORTO

MASTER
ECONOMICS

Artificial Intelligence and Work: Evidence from the Portuguese Labor Market

João Ricardo Oliveira Reis

M

2026



U.PORTO
FEP FACULDADE DE ECONOMIA
UNIVERSIDADE DO PORTO

ARTIFICIAL INTELLIGENCE AND WORK: EVIDENCE FROM THE PORTUGUESE LABOR MARKET

João Ricardo Oliveira Reis

Dissertation
Master in Economics

Supervised by
Hugo de Almeida Vilares

2026

Acknowledgements

What had begun as curiosity has today become this dissertation.

Sometimes the path seemed more difficult, but I was never alone.

To the people who accompanied me, I leave my most sincere thanks.

To my supervisor, Prof. Hugo Vilares, for his tireless help, without which this work would not have been possible.

To my friends and colleagues who made this process lighter.

To my parents, for their unconditional support and for always believing in me. This dissertation is dedicated to them.

Abstract

Artificial intelligence technologies have been a topic that has been considerably debated in recent years. While some are optimistic about their potential effects on productivity and efficiency, others are apprehensive about their effects on job replacement. In fact, AI differs from previous automation technologies – such as computerization, industrial robots and software – mainly by its ability to perform complex and considerably less routine tasks that would normally require greater cognitive effort. It is important, however, to understand how these capabilities effectively interact with human skills and influence the labor market. While this debate remains open, AI continues to develop new capabilities. This dissertation examines the extent to which workers and companies most exposed to AI in Portugal see changes in their wage levels, as well as productivity and employment, respectively. To this end, an econometric model was estimated by applying the C-AIOE index of Pizzinelli et al. (2023) to administrative data at the level of Portuguese workers and companies, in order to quantify the impact of exposure to AI on labor market outcomes. The results revealed the existence of a wage premium as well as productivity and employment losses, for employees and firms, respectively, most exposed to AI. In addition, occupations with higher levels of exposure – professionals, managers, administrative staff and technicians – are both the best paid and most qualified.

JEL codes: J00, J24, J31

Keywords: Artificial Intelligence, Automation, Wages, Productivity, Employment

Resumo

A tecnologias de inteligência artificial têm sido tema sido um tema consideravelmente debatida nos últimos anos. Enquanto alguns revelam-se otimistas quanto aos seus potenciais efeitos sobre a produtividade e eficiência, outros mostram-se apreensivos quanto os seus efeitos sobre a substituição do emprego. Efetivamente, a AI distingue-se das tecnologias de automação anteriores – como computorização, robôs industriais e software – sobretudo pela sua capacidade em desempenhar tarefas complexas e consideravelmente menos rotineiras que normalmente exigiriam maior esforço cognitivo. Importa, no entanto, perceber de que forma essas capacidades efetivamente interagem com as competências humanas e influenciam o mercado de trabalho. Enquanto este debate que continua em aberto a AI continua a desenvolver novas capacidades. Esta dissertação examina até que ponto trabalhadores e empresas mais expostas à AI em Portugal assistem a alterações dos seus níveis de salários, bem como produtividade e emprego, respetivamente. Para tal, foi estimado um modelo econométrico aplicando o índice C-AIOE de Pizzinelli et al. (2023) a dados administrativos ao nível dos trabalhadores e empresas portuguesas, de forma a quantificar o impacto da exposição à AI sobre os outcomes do mercado de trabalho. Os resultados revelaram a existência de um prémio salarial, bem como perdas de produtividade e emprego, sobre colaboradores e as empresas, respetivamente, mais expostas à IA. Além disso, as profissões com níveis mais elevados de exposição – profissionais, gestores, funcionários administrativos e técnicos – são tanto as mais bem pagas como as mais qualificadas.

JEL codes: J00, J24, J31

Palavras-chave: Inteligência Artificial, Automação, Salários, Produtividade, Emprego

Index of Contents

1 Introduction **1**

2 Literature Review 3

2.1 Conceptual Framework: Technology, Tasks and Labor Demand 3

2.1.1 Skill Biased Technological Change 3

2.1.2 From Skill-Based to Task-Based Approaches 4

2.1.3 Routine-Biased Technological Change and Job Polarization 5

2.1.4 Modern Framework: AI & Automation 7

2.2 Automation vs Augmentation: Competing Mechanisms 8

2.2.1 Automation 8

2.2.2 Augmentation 10

2.3 New Tasks, Structural Change and Employment 12

2.3.1 New Tasks and Job Creation 12

2.3.2 Structural Change and Labor Reallocation 13

2.4 Artificial Intelligence and Work 14

3 Data 22

3.1 Sources 22

3.1.1 Worker level data 22

3.1.2 Firm level data 22

3.1.3 C-AIOE 22

3.2 Overview 23

4 Method 29

4.1 Exposure Measurement 29

4.1.1 Conceptual explanation 29

4.1.2 Measure replication 31

4.2 Estimation strategy 32

4.2.1 Baseline OLS model 32

4.2.2 IV estimation 34

5 Results 37

6 Discussion & Limitations **43**

6.1 Discussion 43

6.2 Limitations 44

7 Conclusion 46

References 48

List of Graphics

Graphic 1. Average wage by occupational groups 23

Graphic 2. Average exposure by year (2016-2023) 24

Graphic 3. Average exposure by education levels 25

Graphic 4. Average exposure by qualification levels 25

Graphic 5. Average exposure by occupational groups 26

Graphic 6. Average exposure by regional groups (NUTSII) 27

Graphic 7. Average exposure by activity sectors (CAE) 28

List of Tables

Table 1. OLS estimates of the impact of C-AIOE on Wages 37

Table 2. OLS estimates of the impact of C-AIOE on GVA 38

Table 3. OLS estimates of the impact of C-AIOE on Employment 39

Table 4. IV estimates of the impact of C-AIOE on Wages 40

Table 5. IV estimates of the impact of C-AIOE on GVA 41

Table 6. IV estimates of the impact of C-AIOE on Employment 42

1 Introduction

Technological evolution is one of the fundamental determinants of the evolution of the structure of the labor market. Over time, the emergence of successive technologies such as computing, industrial robots, software and machine learning has been the focus of several contributions in the literature on their consequences for labor dynamics. Models based on skill-biased technological change, for example, were one of the hypotheses developed by several authors. Subsequently, the limitations of this approach motivated the emergence of task-based models, in which the focus of impact happens at the level of tasks instead of occupations as a whole. Within the scope of this approach, the concept of routine-biased technological change eventually emerged, according to which automation technologies would tend to replace human labor in repetitive and routine tasks. From this framework, however, artificial intelligence technologies diverge, which, due to their general-purpose capabilities, can perform cognitive, complex and non-routine tasks. This divergence motivated several authors to develop different measures of exposure to AI in order to study its impact on labor market outcomes. However, most of these studies are focused on the context of large developed economies, leaving the investigation of the effects on smaller economies relatively less explored. In the context of this gap, the motivation for the work developed in this dissertation arises.

The main objective of this work was to estimate the impact of AI exposure, measured by the C-AIOE index of Pizzinelli et al. (2023), on workers' wages and the level of productivity and employment of companies in the context of the Portuguese economy. To this end, the C-AIOE measure was replicated based on the AIOE indicator by Felten et al. (2021) and on work context and job zones data taken from O*NET. Subsequently, their occupational classification from SOC to CPP2010 was converted and the exposure measure was applied to data at the worker and firm level, taken from Quadros de Pessoal and SCIE databases. In this sense, the research question of this dissertation is: how do wages, productivity and employment vary in the Portuguese economy when workers and companies are more exposed to artificial intelligence technologies?

In the context of this branch of literature, the relevance of this research lies mainly in the relatively little explored use of the C-AIOE index to quantify the effects of IA on the labor market. In addition, it contributes to the understanding of the potential impact of this

technology on the Portuguese economy – an application still with few contributions in the literature.

To perform the desired analysis, a methodological approach was applied divided into two phases. First, a base OLS model was used to estimate the impact of C-AIOE on wages, gross value added and employment. For each labor market outcome, different specifications of the model were estimated, combining control variables, cluster robust standard-errors and fixed effects at the worker and company level. Subsequently, to deal with suspected endogeneity at the C-AIOE level, a model IV was estimated using a bartik Shift-share instrument based on the exposure measure.

Finally, the rest of this work is organized as follows. Section 2 presents the literature review, theoretically framing from the Skill biased technological change to the AI exposure models. Section 3 presents the data used, their sources and an overview of the descriptive statistics. Section 4 introduces the C-AIOE replication process and the OLS and IV estimation strategy. Section 5 reports the empirical results obtained. Section 6 discusses the estimation results and recognizes the main limitations of the study. Finally, Section 7 provides the conclusion of the research, while Appendix has graphs on descriptive statistics and tables with estimation outputs.

2 Literature Review

2.1 Conceptual Framework: Technology, Tasks and Labor Demand

2.1.1 Skill Biased Technological Change

In recent years many economies worldwide have experienced a wage inequality phenomenon. Such disproportional distribution of earnings seems to point out a wage premium towards more educated workers. On another words, while labor demand for low educated workers tend to fall, high educated workers demand has been growing. Therefore, understanding such shift on labor demand has been one of the central questions that many economists try to answer. One explanation emphasized in the literature is the role of technological change as a determinant of the demand structural bias, by the Skill-Biased Technological Change (SBTC) hypothesis.

SBTC literature claims that technological waves are an endogenous process guided by economic incentives, rather than a neutral process that impacts all workers the same way. In this context, Acemoglu (2002) documents that while in the nineteenth-century technological innovations were mainly skilled workers replacing, in the twentieth century they became largely skill biased. Such phenomenon, according to the author, aligns with the skill biases technological change hypothesis. As such, he argues that once literature registers a great increase of skills supply and yet skill premium did not fall, it implies that skill demand must have shifted outward and so skill demand is defined as not exogenous - firms innovate towards more profitable technologies, which depend on market size. Therefore, more skilled workers lead to more incentives to invest in high skill technologies and rising skills supply, what means that skills supply can shape future demand of skills. Based on this model the author argued that in the 19th century there was an innovation bias toward unskilled labor, due to the abundance of unskilled workers, that according to the model led to unskilled technological innovations. In the same way, the 20th century technologies were skill biased because of the higher abundance of skilled workers.

Acemoglu (2002) provided a coherent theoretical foundation for SBTC based on the relative abundance of skilled workers. However, its empirical implications require validation on long-run data. A central contribution in this sense is presented by Goldin and Katz (2007), that aimed to explain the long-run evolution of wage inequality in the US economy over the 20th century. More specifically, the authors deduced a theoretical skill-based framework in which wage inequality is a function of demand (e.g. technologies) and supply (e.g. education) of

skills and applied it to empirical data. Once skill demand is not directly observable it is inferred as a residual. So, if supply rose considerably and inequality decreased, it means that supply prevailed demand. On the other hand, if supply rose a bit and inequality increased, demand dominated. Based on this dynamic, Goldin and Katz (2007) argue that while from 1915 to 1980 wage inequality fell due to a rapid expansion of education (supply > demand), from 1980 to 2005 wage inequality increased due to an educational slowdown and a stable level of technology (supply < demand).

Despite its success in explaining the rise in wage inequality, SBTC framework has critical limitations. First, SBTC decomposes workers only by two main groups – skilled and unskilled or high skilled and low skilled. The problem with such level of aggregation is that it abstracts from heterogeneity within the same educational group, overlooking that workers with the same level of education may perform different tasks with different exposure to technologies. Another limitation is that SBTC does not explain task reallocation within the same occupations and so assumes that technology substitutes the whole skill group, instead of specific tasks.

2.1.2 From Skill-Based to Task-Based Approaches

These limitations of the SBTC approach have been motivating the development of task-based models in literature. While in SBTC frameworks technologies impact workers by their skill level, task-based models state that technologies affect specific tasks within an occupation, reshaping the task content of occupations rather than automate the whole occupations. Different tasks may have different suitability for technological automation.

A seminal contribution to the literature shift towards tasks is provided by Autor et al. (2003), who aimed to explain how computerization changes the tasks performed by humans and alters the job skill demand. The authors created a theoretical task-based model in which workers perform routine tasks, non-routine cognitive tasks and non-routine manual tasks. While routine tasks can be automated by computers, non-routine cognitive tasks are complemented. Based on this framework they built an economic model to predict the task demand response to an economywide decline in computer capital prices. This model revealed that as computer capital prices decline, industries with occupations intensive in routine tasks invest more in computer's capital, decreasing routine labor tasks inputs and increasing nonroutine labor task inputs. Hence, the demand of highly educated workers will rise, once they have a comparative advantage in non-routine tasks. Therefore, computerization affects

labor demand through the task content rather than through skill-biased logic, benefiting highly educated workers not because of their skills, but because they have a comparative advantage performing nonroutine tasks.

Another groundbreaking contribution in this context was developed by Acemoglu and Autor (2011) that identified the limitations of the canonical skill-based model on explaining recent labor market trends and suggested an alternative task-based model. The authors argue that the canonical model fails to explain the low skilled workers decrease in real earnings, significant shifts in wages and employment across workers and occupations, polarization of educational returns and the reallocation of certain tasks within occupations from labor to new technologies. The authors relate these limitations to strict assumptions and hence suggest the relaxation of: (1) The fixed allocation of skills to tasks and (2) the technological change that only roles an augmentation effect of high skill workers. Therefore, they built an alternative task-based model in which workers are divided into high, medium and low skilled. Each skill group may perform different tasks depending on wages, productivity and comparative advantages. Theres also three groups of tasks: routine tasks (performed by medium skilled workers), non-routine abstract tasks (performed by high skilled workers) and non-routine manual tasks (performed by low skilled workers). While routine tasks are automated, non-routine abstract tasks are complemented by technologies and, therefore, medium skilled workers became displaced and high skilled workers complemented. According to the authors this task-based model provides an empirically consistent explanation of recent trends, more accurate than the canonical skill-based framework.

Thus, the task-based models shift the target of technological change from skills groups to a variety of tasks. Instead of automating the whole occupation performed by a skill group, technologies may either automate or complement different tasks within the same occupations. Such framework leads to a task reallocations process in which the allocation of tasks across workers and technologies is adjusted due to technological change.

2.1.3 Routine-Biased Technological Change and Job Polarization

The reallocation of tasks within occupations and their impact on the labor market is a key branch in task-based literature. Different tasks may have different exposure levels to technologies. For example, the occupation of an accountant involves a different set of tasks, some which may be automated – such as data entry, bookkeeping or tax form standardization – while others are more difficult to automate. This process leads to a reallocation of tasks, in

which the automatable ones are performed by technology (e.g. software) rather than by workers. This is a process that skill-based models do not consider and because of that, they also fail to explain job polarization. According to those models, technology is skill biased what would lead to increase high skill workers employment and decrease in low skill workers levels. However, as identified by Acemoglu and Autor (2011), recent trends show employment increase in both high and low skill workers and a decrease in middle skill – job polarization. Task-based models, on the other hand, account for the reallocation of tasks, hence it is capable of explaining job polarization through Routine Biased Technological Change for technologies such computers, software and robots, that tend to be more likely to automate routine tasks.

An important contribution to the understanding of task reallocation and job polarization was presented by Goos and Manning (2007). In this study, to explain the concentration of employment in high and low-skill occupations in the UK, the authors presented a task-based framework based on Routine-Biased Technological Change (RBTC) in which there are routine tasks (easily automated), non-routine cognitive tasks and non-routine manual tasks. Since routine tasks are primarily performed by middle-skilled workers, once they are automated it leads to job polarization – rise in employment among high and low-skilled workers, alongside a decline in middle-skilled workers. Then, motivated by this framework they retrieve data on occupations classified by wage level and task content from 1975 to 1999, compute routine intensity by task and regress employment growth on wage rank and the routine intensity measure. The results obtained shed light on a U-shaped pattern on employment levels, with high and low-wage jobs experiencing employment growth while middle-wage jobs (in routine-intensive occupations) a decline. As such, Goos and Manning (2007) founded evidence that the job polarization documented in the UK was due to Routine-Biased Technological Changes.

Another relevant contribution was that of Autor and Dorn (2013) in which they tried to explain the US employment and wages polarization based on the decline of routine jobs and expansion of low-skill service jobs. The authors also started with a task-based framework with an output function depending on highly skilled workers (abstract tasks), medium skilled workers (routine tasks), low skilled workers (manual tasks) and technological capital (mainly computers). While routine tasks are fully automated and abstract tasks are complemented by technology, manual tasks are hard to automate. Motivated by the theoretical framework, they measured the routine task intensity from 722 US local labor markets on commuting zones

and used it to regress labor market outcomes, such employment and wages[1]. The results revealed that: (1) regions with higher routine task intensity experienced a decrease in employment – illustrating the automation effect of routine tasks presented in the previous theoretical framework – and (2) while high and low skill jobs employment and wages levels increased, medium skill jobs experienced an employment decrease and wage stagnation.

Despite the success of Routine-Biased Technological Change to present a reasonable explanation of job polarization, it remains mainly focused on the automation of routine tasks suitable for computerization type technologies. However, recent technological developments in fields such as machine learning and artificial intelligence are characterized by broader automation capabilities. While computers and robots are more suitable for performing routine tasks, the recent technological wave is capable of performing a wider range of cognitive and complex tasks requiring fewer human instructions. Hence, it takes a broader task reallocation framework capable to account for the automation capabilities of new technologies such as artificial intelligence.

2.1.4 Modern Framework: AI & Automation

Artificial intelligence can be defined as a general-purpose technology as it goes beyond routine tasks. Hence, both skill and routine biased technological change frameworks fail to explain AI's impact on task reallocation. Unlike computers, artificial intelligence can perform complex cognitive tasks. For example, a radiologist may perform tasks such as clinical image analysis (e.g., interpreting X-rays), which are susceptible to being automated by AI, while also engaging in medical decision-making and patient interaction, considerably less likely to be fully automated. As such, developments in new theoretical frameworks plausible to explain AI's impact on the labor market have been encouraged.

In this context, Acemoglu and Restrepo (2018) displacement-reinstatement effect model gained significant relevance in literature. In this paper, the authors aimed to present a task-based model of technological change, coherent with earlier contributions to explain phenomena such as job polarization, wage inequality and employment shifts, based on moder automation technologies. According to this model technology impacts work by two main channels: displacement and reinstatement effect. The displacement effect refers to the substitution of workers by technology on routine cognitive and routine manual tasks.

[1] Once Routine Tasks Intensity (RTI) was likely to be endogenous, it was used an instrumental variable approach based on the US industry structure in 1950, using routine intensive industries as IV.

Reinstatement effect, on the other hand, results either by the creation of new occupations and task categories (in which labor has a comparative advantage) and by the complementary effect of technology on abstract tasks, which are performed by high skill workers, increasing their productivity and wage levels. In this way, reinstatement acts as a countervailing effect to displacement (mainly due to the new task creation) and so the actual impact of technology is given by the net effect between displacement and reinstatement.

In a later work Acemoglu and Restrepo (2019a) adjusted this task framework to AI. They argued that since artificial intelligence is capable of automating, more than just routine tasks, cognitive and non-routine tasks, its displacement effect may be stronger affecting a wider set of jobs. But rather than just replace the authors claim that AI also creates new tasks in which labor has a comparative advantage (reinstatement effect). However, these new tasks also imply the need for new skills. According to the authors, if the new skill supply does not keep up with the new task creation, the adjustment process to AI may be slowed down, compromising both the task creation and the reallocation of labor from existing tasks, and so compromising the reinstatement effect.

Within this framework the actual impact of technological change on the labor market is not pre-determined and so is going to depend on the net effect between displacement and reinstatement effect. In the case of artificial intelligence, its wider automation capabilities contribute to a stronger displacement. Hence, the countervailing forces on such effect is going to depend greatly on the keep up of skill supply to the creation of new tasks. In this way, skill reinforcement policies such as education and qualification of workers are crucial to empowering the reinstatement effect of AI. As such, the overall impact of this technology on employment and wages is still ambiguous once it is going to depend on the strength of these two opposing mechanisms. In this way, empirical literature on whether technological change is going to substitute or complement human work has been encouraged.

2.2 Automation vs Augmentation: Competing Mechanisms

2.2.1 Automation

As mentioned above, one central channel through which technological change affects the labor market is by the displacement effect, whereby automating a set of tasks they substitute human labor. In the context of task-based framework the automation of tasks previously performed by workers lead to labor demand decrease, shrinking employment and wages in occupations exposed to automation. As discussed previously, different types of technologies

may have different suitability of automation. While computers and robots mainly affect routine manual tasks, software and AI are capable of performing more complex and cognitive tasks. As such, many empirical contributions in literature focus on deepening the understanding of these effects.

One of the central contributions in this sense was provided by Acemoglu and Restrepo (2020) where they studied the empirical effects of industrial robots on the US labor market from 1990 to 2007. More specifically, they developed created an exposure metric of occupations to robots and applied in an econometric approach, regressing employment and wages outcomes on the exposure. The results documented present a negative and robust impact of industrial robot adoption on employment and wages. Such negative effect appears to be stronger at local level (commuting zones) than at the aggregated level. Furthermore, the largest job losses were in routine intensive occupations, affecting mostly low and middle skill workers. Besides that, the authors did not find strong evidence on job creation due to robot adoption, what according to the previous model explains the high displacement levels.

With a similar approach, Graetz and Michaels (2018) also tried to explain the impact of industrial robots, but for a set of seventeen developed countries from 1993 to 2007. Based on data of fourteen industries of those developed countries, the authors created a robot intensity measure and used to regress productivity, employment and wages[2]. The results revealed a general productivity increase alongside an output price decrease, suggesting that robots improved efficiency and lowered production costs. Besides that, robots appear to be skill biased towards low-skilled workers, despite there being no significant effect on total employment, the low-skill workers share decreased. Furthermore, there were no relevant wage differences across skill groups, despite an average wage increase.

Dauth et al. (2017) also aimed to study how industrial robots affect employment, wages and worker careers, but applied to the German labor market. To do so, they first estimated the effects of robots at the regional level on local labor markets and then studied how that impact translates to individual workers' level. At the regional level they found that their manufacturing employment decreased at the pace that service sector employment grew, meaning that there was reallocation of workers from manufacturing sector to the service sector, resulting in a null net employment change. Besides that, wage polarization was also

[2] Once robot density could be potentially endogenous the authors used an instrumental variable approach based on two instruments: replaceable tasks (industries with more replaceable tasks are more likely to adopt robots) and reaching and handling tasks, that measures physical tasks suitable for robotic arms.

documented, once high skill workers' wages increased, while for medium skill workers remained steady and on low skill workers had mixed and small effects. At the individual level, however, robots appear to decrease the entry of new workers into more exposed occupations, instead of replacing the current workers. Furthermore, high skill workers' wages increased, medium skilled workers decreased and, once more, there were mixed and small effects on low skill workers earnings, confirming the idea that robots automate routine tasks (medium skilled) and complement abstract tasks (high skilled).

Hence, there are different results on automation effect. While Acemoglu and Restrepo (2020) concluded that robots tend to shrink employment and wages spatially at the local level, Graetz and Michaels (2018) found no net effect on average employment, but an increase in average net wages. Finally, Dauth et al. (2017) also found no net effect on average employment (spatially at local level) alongside with increase in high skilled wages, decrease or null effect in medium skilled and mixed changes on low skilled workers. Hence, automation may substitute workers, but its effects vary across countries, time, sectors and type of technology.

2.2.2 Augmentation

On the other hand, besides job substitution technological change may also have a complementarity effect on tasks that are not fully automated, as mentioned above by many contributions such: Autor et al. (2003), Acemoglu and Autor (2011), Autor and Dorn (2013), Acemoglu and Restrepo (2018) and Dauth et al. (2017). Aside from this, there are another relevant contribution to this sense.

A central contribution in this regard was presented by Bessen et al. (2022) that focused on understanding what happens to the tasks performed by workers that are not replaced by technologies. To do so, they created a task-based model in the spirit of Acemoglu and Restrepo (2018) and yielded a "remainder effect" in which the returns of skills of tasks not automated by technologies are increased. The main conclusion retrieved from such framework was that for a given set of occupations that needs tasks, technological change displaces labor in some of them. Hence, those occupations now rely on capital (technologies) and on tasks that were not automated. These tasks became more skill intensive and higher quality, leading to higher wage levels. In this way, technological automation may increase wages even with low productivity gains, if automation requires higher-quality performance on remaining tasks.

Another contribution was developed by Michaels at al. (2014) that aimed to explain whether ICT in a set of eleven OECD countries led to labor market polarization by increasing high-educated workers' demand and decreasing middle educated workers. Based in an econometric approach motivated by a task-based model, the authors found evidence on job polarization (on both employment and wage levels) in which routine tasks are fully substituted by ICT, while manual non-routine tasks are little affected and, most importantly, abstract tasks became complemented.

There appears to be a consensus in literature that technological change provides heterogenous effects across tasks and consequently skill levels. According to task-based frameworks, different tasks have different suitability to be performed by technologies. For instance, routine manual tasks may be easily automated by robots, while routine cognitive tasks are more likely to be automated by AI technologies. As such, once routine manual tasks tend to be traditionally performed by medium-low skill workers and routine cognitive tasks by some relatively higher skilled labor, the formers tend to be automated by robots while the latter by software and AI technologies.

However, technological change effects do not end on automation – there is also an augmentation/complementarity effect that often leads to higher wages and labor demand. This effect tends to arise mainly on tasks that are not automated and that are complementary to technology. Complementary tasks lean towards complex non-routine tasks that require a great deal of judgment. For example, AI-assisted medical diagnosis, investment and management decisions may use artificial intelligence to process and analyze relevant data, but the decision-making still needs to be made by humans. Furthermore, despite technological complementary tasks usually performed by higher skilled labor, this connection is not strict forward, since some lower skilled tasks are also complementary (e.g. machine operators assisted by digital systems).

Taking together the evidence on literature on both sides, these findings suggest that the impact of technological changes on labor market and how it is going to change the reallocation of tasks cannot be understood through a single mechanism. It is necessary to account for both displacement and complementarity effects.

2.3 New Tasks, Structural Change and Employment

2.3.1 New Tasks and Job Creation

However, if technological change displaces labor, why does employment still exist? There are many contributions in literature regarding this question. One of the most relevant was presented by Acemoglu and Restrepo (2019b), that scoped the study on explaining why technological breakthroughs did not destroy jobs. The authors empirical decomposition the displacement-reinstatement model, previously presented in Acemoglu and Restrepo (2018), on US labor market data from 1947 to 2017. The results show that since 1980 the labor share of earnings fell significantly, due to a considerable increase in automation alongside a slowdown in new task creation. On another words, they point that the sizeable job destruction documented in the US was due to the reinstatement effect played by new task creation that could not keep up with the displacement effect played by automation, leading to an unbalance technological change.

Autor et al. (2022) also contributed to the understanding of how the creation of new tasks may offset the displacement of jobs in the US economy. The novelty of this paper is that they measured the creation of new tasks over eighty years a systematic measurement of this effect. More specifically, they start by analyzing new job titles across the sample from 1940 to 2018. Then, link the occupational text data to patent text descriptions (as a proxy of innovations). Next, they classified patents by automation innovations (like tasks performed by workers) and augmentation innovations (like the output of occupations). Finally, they construct exposure measures of occupations to automations and augmentation innovations and use them to regress outcomes such as employment by sector. The results show that new jobs represent about 60% of the total employment of the US economy in 2018, which is a massive proportion. Such new jobs were mainly created by augmentation effect (technology complements tasks performed by humans) and labor demand shifts – new tasks countervail displacement, but not all task creation comes from technology. Furthermore, the evidence suggests a shift of new tasks from middle skill jobs (1940-1980) to high and low skill jobs, showing job polarization.

Moving on to another contribution, Lin (2011) also studied the creation of new tasks, however at the local level. More specifically, the author aimed to measure technological adaptation through the creation of new occupations and explain why some locations are better at generating new work. The author also started with a task-based model in which

adaptation the technological change leads to the creation of new tasks. According to the framework, the adaptation process is faster in cities intensive in human capital and more diverse industries, due to knowledge spillovers and wider skill offer. In this sense, they document data on new work creation based on new job titles (as done by Autor et al. (2022)) and regresses the probability of a worker to be in a "new job", based on individual and city characteristics. The results reveal that workers are more likely to be in a new job in cities with higher human capital and variety of industries. Besides, new jobs tend to pay higher wages and workers with higher educational levels are also more likely to be in a new job – evidence that new job creation is skill biased.

2.3.2 Structural Change and Labor Reallocation

Instead of studying what technologies do to jobs, as previous papers, Acemoglu and Restrepo (2022) focus on what determines technological adoption. In this work the authors test the hypothesis that technological adoption is an endogenous process that responds to economic incentives. More specifically, they test the hypothesis that demographic changes such as population aging leads to labor scarcity (especially of younger workers) triggering wage growth. Hence, investment in automation becomes relatively cheaper, increasing incentives to adoption of automation technologies. To test this hypothesis, they regressed robot adoption (as proxy of automation) on demographic aging, at the regional level. The results appear to show empirical validation on the hypothesis tested: countries with higher population aging tend to adopt more automation technologies. Besides that, they also find that automation increases productivity and output production. However, technological adoption does not seem to shrink aggregate employment, what makes sense because such technological adoption responds to labor scarcity created initially. In this way, the most important contribution of this paper is to the understanding that structural changes in economics (such as demographics) may influence the adoption of automation technologies. Hence, the relationship between economic factor and automation operates in both directions, rather than in a unidirectional way.

Another relevant contribution in this regard is the work conducted by Acemoglu and Loebbing (2026) where the authors developed a new theoretical approach to explain why automation has typically been focused on middle-skill tasks. Acemoglu and Loebbing (2026) developed an assignment model in which tasks may be assigned to either labor or capital. Furthermore, task complexity and workers skill level are ranked on a scale from 0 to 1 and

there are four levels of automation: no automation, interior automation (intermediate tasks are automated); low-skill automation (both intermediate and simple tasks are automated) and high-skill automation (complex tasks are automated). According to the authors, capital automates mostly intermediate tasks, since simple tasks are cheaper when performed by humans and complex ones are too expensive to be automated. In this way, middle-skill workers became displaced, shrinking their wage and employment levels – leading to polarization. Hence, interior automation is the key mechanism that drives polarization. However, Acemoglu and Loebbing (2026) also argue that AI may increase inequality – once it has a wider automation potential and smaller deepening automation effect – but not necessarily polarization. They state that AI may negatively impact on more low-skill workers, since they do not benefit from AI complementarity effects, like high-skill workers.

2.4 Artificial Intelligence and Work

According to Webb (2020), artificial intelligence differs from the other types of automation technologies (e.g. software, computers and robots) for its ability to perform cognitive complex tasks while requiring much less human instructions. As mentioned before, different technologies may have different suitability to automate certain tasks. As such, to understand how AI impacts labor and study its effects across workers, it is necessary to measure AI adoption. There are many contributions in literature in this regard.

Brynjolfsson et al. (2018) on the study of what occupations are more exposed to machine learning (ML) technologies, argued that the measurement must be based on the capabilities of the technology, rather than by the whole occupations – as mentioned previously in the sections 1.1.2 and 1.1.3. To do so, the authors started by linking 964 different US economic occupations from O*NET database to 18.156 tasks and about 2.069 Direct Work Activities (DWA). Then, applying an adapted version of the Brynjolfsson and Mitchell (2017) task evaluation rubric, each DWA was ranked by 7 to 10 respondents that scored 23 statements (regarding ML suitability to DWA) on a 5-point scale by a crowdsourcing platform. Next, the evaluations were aggregated by DWA and occupations, creating a Suitability for Machine Leaning (SML) score. Finally, Brynjolfsson et al. (2018) linked the SML measure to BLS 2016 data on employment and wage for the US economy and correlated the exposure measure with wage percentiles and total bill by jobs. The results revealed high variability on SML scores within occupations – with a SML standard deviation of 0.596, representing around 17.2% of the mean SML score – suggesting that rather full automate jobs, ML leads to task

reallocation within the same occupations. Moreover, there was also found a low correlation to occupational wages (-0.14) and wage bill percentiles (0.10). The top 3 SML scores were the occupations of concierges, mechanical drafters and morticians, undertakers, and funeral directors, while the lowest 3 SML scores were found on massage therapists, animal scientists and archeologists.

In the work of Babina et al. (2024) the authors aimed to measure firm-level AI adoption and estimate its impact on firm growth, employment, and innovation. According to their conceptual model, AI impact labor by a product innovations channel (that leads to higher demand and firm growth) and process efficiency channel (that leads to cost reduction and higher productivity). To identify the dominant channel, authors regressed product innovation and process efficiency proxies to AI workers share as main explanatory variable. Regarding product innovation, the results reveal a positive and significant impact of AI investment on both trademarks and product patent variables, unlike product mix that became insignificant when added control variables. On the other hand, concerning process efficiency, instead of cost reduction, authors found a significant operating cost increase relation driven by AI investment and no significant impact on productivity and process patents. Hence, product innovation seems to be the main driver to which artificial intelligence impacts firm growth. However, such AI investment benefits appear to be biased to initially larger firms, leading to higher industry concentration. Besides, Babina et al. (2024) found no statistically significant evidence of AI on workers displacement, even down does not rule out possibility of task reallocation.

In order to fill in a data gap regarding diffusion of modern technologies (such as AI and ML), Zolas et al. (2020) introduced a new technology module to the 2018 Annual Business Survey (ABS). Such module addresses three main questions. The first one relates to digital share on six key business activities: personnel, financial, customer feedback, marketing, supply chain, and production. The second one regards purchases of cloud services for eight IT functions: security of firewall, servers, data storage and management, collaboration and file synchronization, data analysis, billing and account management, customer relationship management and all IT functions. The final question directly surveys the adoption level of the following modern technologies: augmented reality, automated guided vehicles (AGV) or AGV systems, automated storage and retrieval systems, machine learning, machine vision software, natural language processing, radio-frequency identification (RFID) inventory system, robotics and touchscreens/kiosks for customer interface. After conducting the

survey to a sample of 850.000 businesses, the authors had to face a non-response and "don't know" answers of about 17 to 23% of the total module responses. To deal with this problem, Zolas et al. (2020) applied an Ordered Probit Model to input this firm's adoption levels. More specifically, they estimated a set of probabilities of technology adoption (e.g. probability of use augmented reality at up to 5%, between 5% and 25% or more than 25%), based on observable firms' characteristics and responses to other questions regarding simpler technologies, such as digitalization and cloud computing. To validate the imputation, authors applied a prediction/recall analysis, where the non-imputed answers data were randomly divided into a training set (70%) and a test set (30%) and used the same probit model to estimate the test set based on training data. The validation suggested a better prediction performance of the probit model on extensive margins (whether firms use the technology or not) rather than intensive margin (intensity of firm technology adoption). Regarding the survey results, Zolas et al. (2020) found that while digitalization is widely used by the majority of firms (90%), cloud computing is moderately purchased by firms (78.4%) and only a few firms (10.3%) applied advanced technologies to their business. However, despite the low adoption rata of advanced technologies, the employment-weighted adoption rate is substantially greater – in fact, as it is for the three types of technologies. Besides, advanced technologies (where artificial intelligence is included) seem to be mainly adopted by larger and older firms. While adoption and firm age are positively related to large-scale businesses, for many cases the scenario is inverted when considering smaller sizes. Moreover, firms that adopt advanced technologies also tend to be more process and product innovative – a result that supports Babina et al. (2024) finding of product innovation as the main channel through which AI impacts businesses.

In the aim of quantifying the potential consequences of artificial intelligence on occupations, Tolan et al. (2021) developed an innovative framework where AI exposure is studied based on their actual capabilities. They claim that AI performs specific tasks, rather than automate full jobs and that AI and tasks are linked by their ability to perform cognitive abilities – as presented in the following hierarchy: artificial intelligence performs abilities; abilities are needed to perform tasks; tasks perform jobs. In this way, based on O*NET data on task and ability taxonomy, authors linked a set of 59 general tasks to a set of 14 cognitive abilities through a Delphi Method, creating matrix *U*. Next, they also linked 328 artificial intelligence benchmarks – which scores AI performance on each ability – to the same set of abilities, producing matrix *R*. Then, authors counted and normalized the number of references to

each benchmark related topic founded in the AI Topics Archive, creating vector *a*. Finally, multiplying matrix W (task weight on occupations) with vector *a*, they obtained the occupational exposure level across abilities (matrix V) and the AI Exposure Score, which corresponds to the sum of vector V for all abilities. By normalizing the exposure measure into percentiles (high, medium or low skilled), Tolan et al. (2021) found that AI seems to be targeting mainly high skilled and high wage workers (e.g. electrotechnology engineers) and less the low skilled ones (e.g. cleaners and helpers), suggesting a weaker polarization potential, unlike older technologies, such as robots. Besides, AI's relation to medium skill workers appears to be inconclusive. Moreover, authors underlined that this exposure does not mean job destruction. On the contrary, they defend that AI is more likely to complement human work, rather than fully replacing it. Furthermore, Tolan et al. (2021) also found evidence that supports the 'Modern Productivity Paradox' (MPP). Introduced by Brynjolfsson et al. (2019), the MPP concept argues that despite groundbreaking progress on AI capabilities, their impact on productivity is far below expectations. As a result of their study, the authors propose four possible explanations for such paradox: Mismeasurement errors, concentrated distribution of gains, overestimated hopes and lags on implementation of AI and complementary technologies.

Another widely cited contribution is the technological exposure measures developed by Webb (2020). In his work, the author created a general-purpose exposure measure that could predict the impact of AI, software and robots on occupations, based on a task-based framework inspired by Acemoglu and Restrepo (2018). To do so, he used a Natural Language Processing (NLP) approach to quantify the match level between jobs and patent text descriptions. While job descriptions, from O*NET database, present detailed textual information on human tasks performed by almost 1.000 occupations, patent text description, from Google Patents Public Data, presents the actual capabilities of each technology. By a dependency parsing algorithm the author extracted verb-noun pairs (e.g. [diagnose, disease]) from both datasets and grouped them by a hierarchy of concepts, through WordNet lexical database. Then, calculated the relative frequency of verb-noun-pairs for each set of patents by technology, assigned it to O*NET extracted data and scored each task based on the relative frequency of their verb-noun-pairs by patents for each technology. Next, obtained the final occupational exposure levels, by the weighted average of the task scores, based on the task weight by occupations. Since this is a prediction measure, Webb (2020) first applied it to well known cases such as industrial robots and software, in order to validate and fine-

tune the exposure score. Then, for both technologies, regressed outcome variables (such as wage and employment share variations) with the respective exposure measures for the period of 1980 to 2010. This process revealed a linear and negative relation between robots and software exposure and labor market outcomes. Since AI is a recent technology, their actual adoption levels still to be developed, and its effects are not yet reflected in the data. Hence, Webb (2020) assumed that the relation between AI exposure and labor outcome follows the same negative and linear pattern founded for robots and software. As results showed that high-skill workers at the 90th percentile of wage distribution are more likely to be exposed to artificial intelligence, it is estimated that this technology will reduce the 90:10 inequality levels. On the other hand, is expected to increase 99:90 inequality, once the top 1% (99th percentile) of the distribution is less exposed – for instances, the author exemplifies with the CEO as one of the top earning occupations that is, however, low exposed to AI. Furthermore, it is also documented that while artificial intelligence targets mainly high skill workers, robots and software exposure are higher for low and medium skill workers, respectively.

Another central breakthrough to the understanding of how AI impacts labor was presented by Felten et al. (2021). In the aim to fill the gap of relevant data to measure the effects of artificial intelligence, the authors created the Artificial Intelligence Occupational Exposure (AIOE) – a metric that indicates how likely is a given occupation to be affected by an AI specific technology, without specifying weather it is going to either substitute or complement human work. To construct this measure, first the authors linked 10 AI applications – retrieved from Electronic Frontier Foundation (EFF) – to a set of 52 O*NET abilities by an application-ability relatedness matrix, constructed through a crowed-sourcing survey[3]. The respondents were asked to give a binary answer ("Yes" or "No") to the relatedness between each ability and AI function[4]. Next, they calculated the ability-level AI exposure (A) by the sum of average responses by ability for all 10 AI applications and weighed the exposure measure by the ability importance (I) and requirement level (L) for each occupation. Then, to obtain the AIOE measure, the authors divided the sum of weighed ability-level AI exposure by the sum of abilities weight by occupation (sum of the product between I and L)[5], in order to prevent biases towards occupations that require more abilities. Furthermore,

[3] According to Felten et al. (2021) the crowed-sourcing survey had approximately 2.000 responses provided by Amazon's Mechanical Turk "Gig Workers".

[4] The authors argue that the responses were limited to a binary format in order prevent answers biased by respondents believes on AI complementarity and substitution potential.

[5] For an extended explanation of the Artificial Intelligence Occupational Exposure, consult section 4.

they created another two measures – the Artificial Intelligence Industry Exposure (AIIE) and the Artificial Intelligence Geographic Exposure (AIGE). AIIE – that represents the exposure levels by industry – was obtained through the average of weighted AIOE by industry employment data. Then, the AIGE was calculated by aggregating the AIIE by different industries within each geographic region, based on country employment data. The analysis of the measures results by the authors revealed that white-collar and high-skill occupations are systematically more exposed to the AIOE, while blue-collar jobs seem to be less exposed once they rely mainly on physical skills. In this extent, industries intensive in high education demanding occupations[6] are in average the more exposed to the AIIE, while the low skill intensive industries[6] appear as the less exposed to AI. Finally, the authors found evidence that, despite some exceptions, the US urban areas concentrate the highest levels of exposure, while rural ones are associated with low exposure levels[7].

In a later contribution, Felten et al. (2023) adapted Felten et al. (2021) Artificial Intelligence Occupational Exposure to account only for AI language modeling (LM) technologies, so they could study their impact on occupations and industries. The adjusted measure was obtained as the original one but weighing the ability-level AI exposure (A) by a binary parameter α that assumes 1 for AI LM applications and 0 for others.

Based on descriptive analysis the authors concluded that language intensive jobs (e.g. telemarketers and history teachers) are the most exposed to LM (as would be expected) and that legal services and securities, commodities, and investments are the top exposed industries. Besides that, exposure to AI LM seems to be stronger in high wage jobs.

Also relating to AI language models, there is the contribution of Eloundou et al. (2023) where authors aimed to quantify how broadly large language models (LLM) could affect the labor market. To do so, they constructed a new metric of exposure to LLM, using an interesting approach – based on task-level data, the authors used Chat GPT-4 by API prompting to classify tasks direct exposed (tasks suitable to LLMs perform) and indirect exposed (tasks that can be done by LLMs assisted by additional tools). Then, the tasks were aggregated by

[6] The top three exposed industries are "Securities, commodity contracts, and other financial investments and related activities", "Accounting, tax preparation, bookkeeping, and payroll services" and "Insurance and employee benefit funds" (Felten et al, 2021), while the bottom three exposed industries are "Support activities for crop production", "Services to buildings and dwellings" and "Foundation, structure, and building exterior contractors" (Felten et al, 2021).

[7] According to Felten et al. (2021) urban areas in the US economy are more exposed to AI once they tend to employ systematically more educated and more skilled worker, and rural areas are less exposed to AI because they rely mainly on low skill and educational demanding occupations.

occupations, and it was constructed another three measures: lower bound (number of tasks that LLMs can do by their own); preferred measure (number of tasks that LLMs can do either directly or with tools); upper bound (number of tasks that LLMs could do with a fully integrated AI system). Eloundou et al. (2023) analyzed the exposure distribution of LLMs across employment, wages and educational levels. The authors found out an economy-wide impact potential of LLMs (about 80% of workers have at least 10% of tasks exposed) and a significant share of tasks affected (around 15% of tasks are directly exposed, 30% exposed including tools and 50% exposed including a full ecosystem). Furthermore, exposure levels seem to be widespread across sectors, supporting the idea that LLMs are a general-purpose technology. Besides, as pointed out by Felten et al. (2023), high-skill jobs appear to be more exposed, once LLMs affect a wide range of higher-income and more educated occupations. Gmyrek et al. (2023) also focused on the effects of generative AI technologies and studied its possible effects on jobs around the world through the creation of a task-based exposure measure. Similarly to Eloundou et al. (2023) work, in this paper the authors also use Chat GPT-4 with API prompting. First, they ask GPT to classify tasks with a score within a continuous range from 0 to 1, based on task-level data (e.g. ISCO code and task description). Next, based on the GPT output score and the weight of task on occupations, they built the weighted average of task automation potential – the exposure measure to AI by occupation. Then, derived from the mean score (average automation potential) and standard deviation (heterogeneity of scores across tasks) the authors grouped occupations on task-level scores: automation potential (low standard deviation and high mean), augmentation potential (high standard deviation and low mean), not affected (low standard deviation and low mean), and the big unknown (high standard deviation and high mean). By analyzing the measure distribution across task and work-level data, Gmyrek et al. (2023) found evidence that most jobs are not fully automatable and there is a strong and significant complementation effect. Based on Felten et al. (2021) work, Pizzinelli et al. (2023) argues that the exposure measure AIOE accounts only for a potential interaction of AI to work and so do not tell if it either replaces or complements workers. The authors claim that exposure by itself is not enough, once for a given level of exposure it may either mean a substitution or a complementarity effect. Moved by this limitation, Pizzinelli et al. (2023) introduce the concept of complementarity-adjusted exposure and propose an AIOE adjusted measure – the C-AIOE – that indicates if artificial intelligence is either more likely to automate or complement labor in each occupation. To do so, they retrieved 11 work contexts and a single job zone value by

occupations from O*NET and distribute then into 6 variable groups (communication, responsibility, physical conditions, criticality, routine and skills). Then, the author measured the average work context and job zone value of occupations by each respective variable group(s) to construct the adjustment parameter θ – while values closer to zero indicate a higher potential for substitution, values closer to one indicate a higher potential for complementarity. Since both AIOE and θ range from 0 to 1, the parameter is subtracted from AIOE to obtain C-AIOE. Hence, the closer θ is to 1, the smaller C-AIOE became and so the lower the substitutions potential is. Their results reveal that occupations highly exposed to substitution effect tend to experience lower employment growth. On the other hand, occupations more likely to be complemented by AI appear to experience higher employment growth – these empirical findings support Acemoglu and Restrepo (2018) concept that the net effect between displacement and reinstatement will determine the actual impact of AI.

3 Data

To conduct the aimed analyses, it was necessary to collect data at both worker and enterprise level: 'Quadros de Pessoal' (QP) and 'Sistema de Contas Integradas das Empresas' (SCIE), respectively. Furthermore, the C-AIOE measure was also reproduced based on Pizzinelli et al. (2023) and Felten et al. (2021) work.

3.1 Sources

3.1.1 Worker level data

To obtain data at workers level such as wages, tenure, education, sex and age, we used Quadros de Pessoal (QP). This database is supported by INE, how gathers legally mandatory firm reporting data on workers and creates a harmonized dataset based on CPP2010 for the whole Portuguese economy through several years. From this source we retrieved data on workers number, wage, education, tenure, age, sex and work contract type, plus the fictitious firm identification number and CPP codes.

3.1.2 Firm level data

Regarding firm level data, we used as source the Sistema de Contas Integradas das Empresas (SCIE), also created by INE that gathers legally mandatory firm level data reporting, such as IES[8]. From this source we retrieved data on firms' GVA at market prices, working people by gender, sales, turnover, net income, investment in software and GFCF, plus the fictitious firm identification number and NUTSII.

3.1.3 C-AIOE

To serve as an AI exposure proxy we used the C-AIOE measure. As mentioned before, it results from an adjustment made by Pizzinelli et al. (2023) to Felten et al. (2021) AIOE original measure. As such, we retrieved the original measure from Felten et al. (2021) replication package and then calculated the adjustment parameter θ from Pizzinelli et al. (2023) raw data and instructions. Next, we applied the adjustment parameter to AIOE and used an SOC10-ISCO08 crosswalk to convert and organize the final C-AIOE measure by ISCO codes.

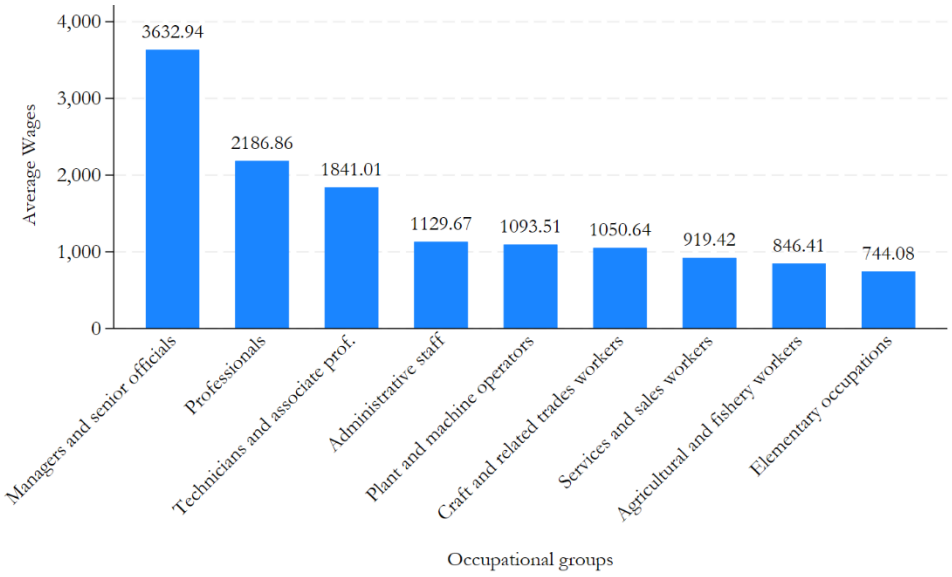
[8] IES: 'Informação Empresarial Simplificada'

3.2 Overview

This study used a total of 3.586.096 workers[9] with an average age of 37 years and composed by 55.7% by males and about 44.3% female workers. These workers are distributed across 417 different occupations, composing a set of 9 main occupational groups: Managers and senior officials, professionals, technicians and associate professionals, administrative staff, services and sales workers, agricultural and fishery workers, craft and related trades workers, plant and machine operators and elementary occupations.

As presented on Graphic 1, when comparing the average wages across the nine CPP occupational groups, it becomes clear that one group stands out from the others: 'Managers and senior officials', with an outstanding average wage of 3633 euros, approximately – the only group above the 3000 euros threshold. Below this stage 'Professionals' and 'Technicians and associate professions' also stand out with an average wage of 2187 and 1841 euros, approximately and respectively. On the other hand, at the tail of the distribution the three lowest groups are 'Services and Sales workers', 'Agriculture and fishery workers' and 'Elementary occupations' – all below the 1000 euros threshold.

Graphic 1. Average wage by occupational groups

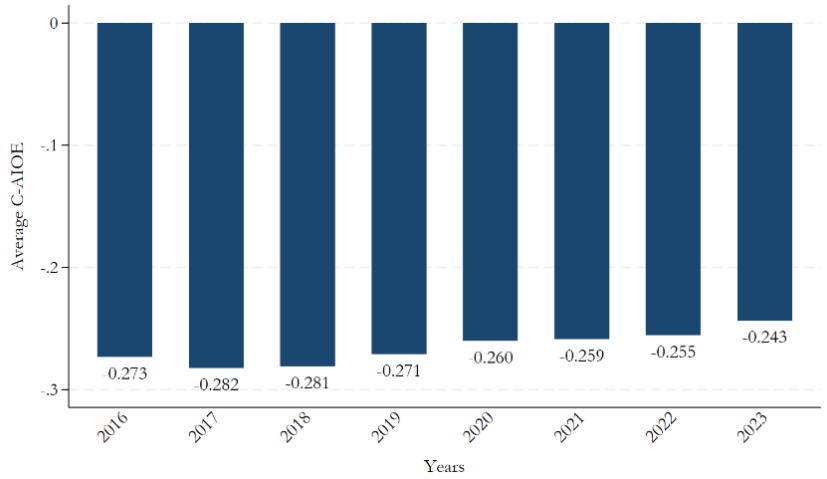


As explained previously, C-AIOE measures the relative exposure of an occupation to Artificial Intelligence technologies. Once this metric results from the complementarity adjustment of AIOE, higher scores translate higher exposure to substitution effect, while lower scores suggest higher complementarity potential. Based on the sample of 3.586.096

[9] Note that the sample of 3.586.096 workers only covers employees.

workers from 2016 to 2023, C-AIOE is comprehended between a minimum of -1.84 and maximum of 1.20, with an average score of -0.27, which draws the image of a moderately low exposed sample. When looking at Graphic 2 it is interesting to notice that across time the average C-AIOE score started to increase consistently since 2017. Once this measure is indexed at the occupational level, the increase of average score in time suggests the movement of the work composition towards occupations more likely to be performed by AI technologies.

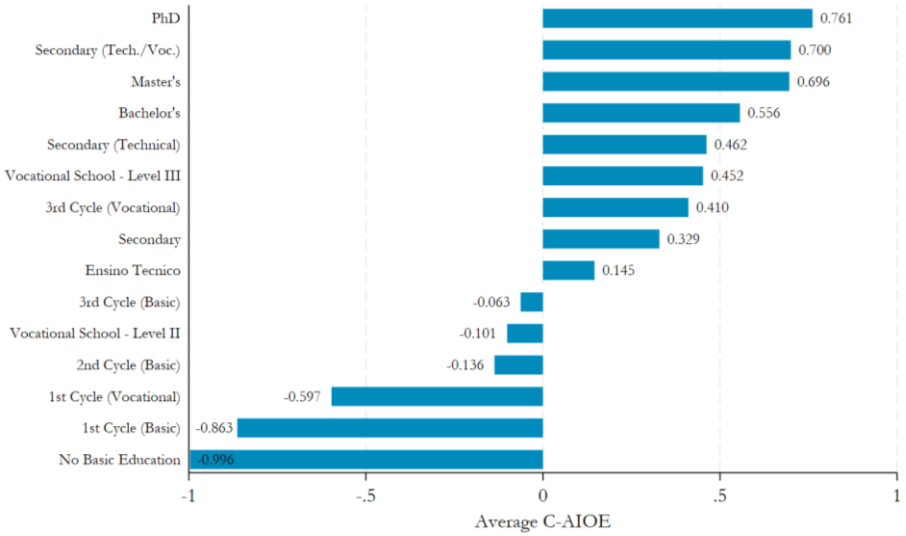
Graphic 2. Average exposure by year (2016-2023)



Graphic 3 represents the relation between exposure and educational levels. As we can see, as the educational level increases the highest is the exposure to AI. On another words, it means that highly educated workers are substantially more exposed to the substitution effect of AI, than less instructed employees. It is curious, however, to notice that, besides high education degrees such as PhD, Master or bachelor, there are relatively lower educational levels also highly exposed. More than that, some of them are even more exposed than a few high educational degrees – for example, workers with a technical-vocational secondary diploma have a higher C-AIOE score than those who hold a bachelor's or master's degree. Another curious pattern is that across all education levels, the only ones that also score a positive exposure to AI, besides higher education degrees, are workers with professional secondary education. This pattern may be a sign that Ai technologies are capable of performing cognitive demanding tasks (usually performed by highly educated workers) but also tasks that require a more technical skill set. When looking at the other side of the coin, we can notice that employees that only hold a first cycle of basic education or less, are extremely low exposed – especially the ones with no basic education that register a score remarkably close to -1. A possible explanation for these low scores may be that workers who are substantially

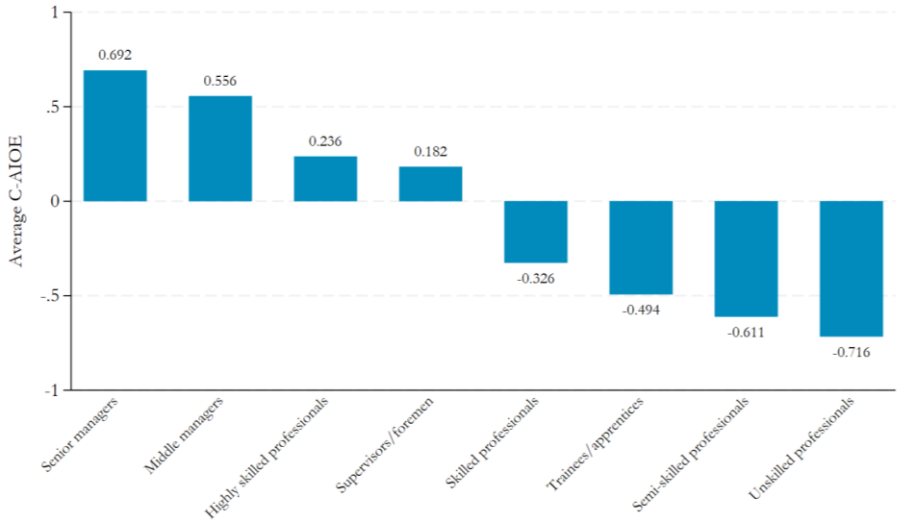
less educated usually perform manual tasks that require a high component of physical work which AI is less likely to perform.

Graphic 3. Average exposure by education levels



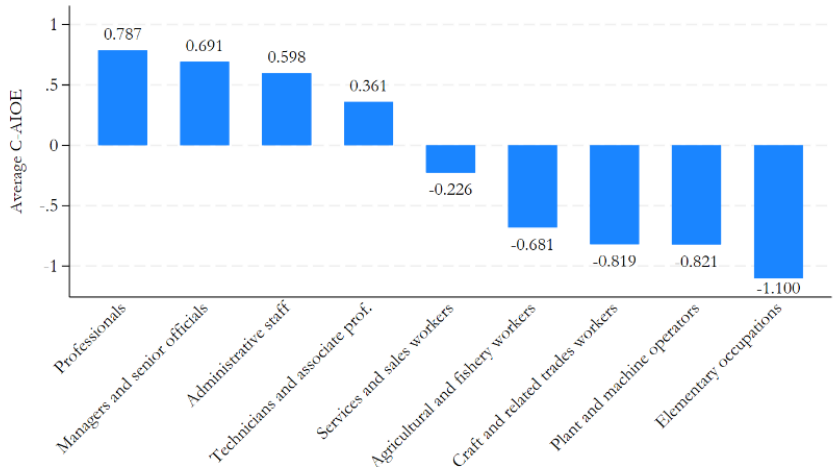
Graphic 4 represents the exposure by levels of qualification. Similarly to education levels, the exposure scores appear to increase as the qualification level reaches higher thresholds. On another words, workers in high qualified positions – such as managers that deal with complex and demanding tasks – are more exposed to AI's substitution effect than low-skilled or unskilled employees that perform less demanding tasks. This evidence is consistent with those found on Graphic 3 – it is plausible to assume that highly educated workers are more likely to perform high qualified tasks, since they are more prepared to deal with complex and demanding work.

Graphic 4. Average exposure by qualification levels



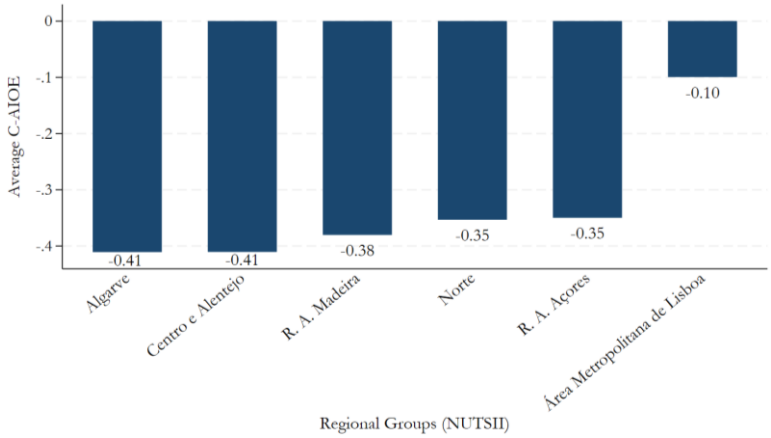
Graphic 5, however, is where the most interesting hint comes up. When looking at the AI exposure levels given by C-AIOE across occupational groups, we can see that the highest paid and qualified occupations also seem to be more exposed to AI substitution effect, while the lowest paid and qualified occupations appear to be less exposed. Such relation is consistent with most findings on literature regarding occupational AI exposure studies – as reported in section 1.4. Nevertheless, the exposure relations suggested by these graphics are merely descriptive statistics and as such they must be taken only as leads to a further econometric approach.

Graphic 5. Average exposure by occupational groups



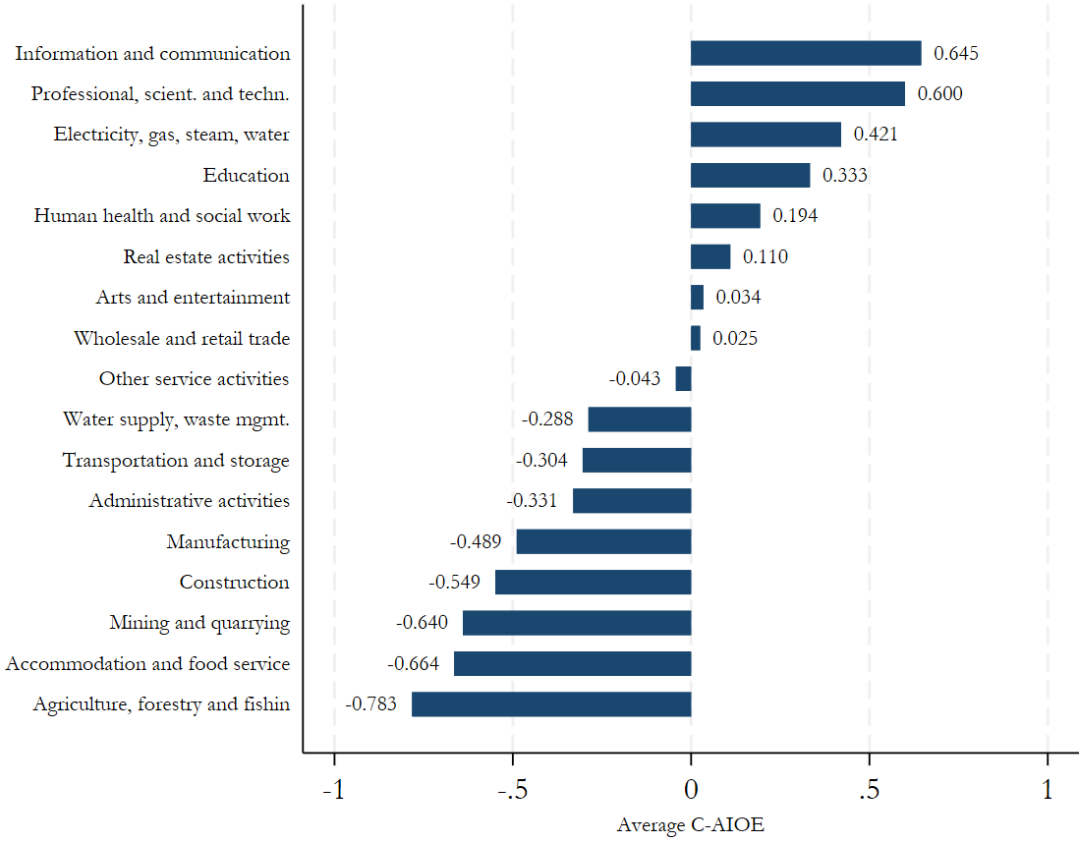
Furthermore, when looking at the average C-AIOE levels across NUTSII (Graphic 6), it is also curious to notice that 'Área Metropolitana de Lisboa' is the region with the highest average exposure level (-0.05) and very distant from other region's scores that go from -0.18 to -0.22. Hence, employees of firms from 'Área Metropolitana de Lisboa' seem to be more exposed to AI than other regions, what suggest that exposure is heterogeneously distributed across regions due to different work occupational compositions. In this case, firms from 'Área Metropolitana de Lisboa' appear to be substantially more intensive in occupations more suitable for AI.

Graphic 6. Average exposure by regional groups (NUTSII)



This database used a set of 662 different activity sectors based on CAE classification codes. On Graphic 7 is displayed the average C-AIOE across 17 groups of activity sectors, based on CAE Rev. 3 table, retrieved from INE. Such graphic illustrates two clearly distinct groups: those with positive values and those with negative values. As would be expected 'Information and communication' and 'Professional, scientific and technical activities' activity sectors are the most exposed to AI and the most likely to be substituted. This result is intuitive once related tasks such as computer program editing and data processing can be easily performed by AI with little human action, as we well know. Other sectors such as 'Electricity, gas, steam and air conditioning supply', 'Education' and 'Human health and social work activities', 'Real estate activities', 'Arts, entertainment and recreation' and 'Wholesale and retail trade, repair of motor vehicles and motorcycles' also registered a positive exposure score, even down progressively weaker. On the other hand, the less exposed sectors are 'agriculture, forestry and fishing', 'accommodation and food service activities' and 'mining and quarrying', followed by 'manufacturing', 'construction', 'administrative and support service activities', 'transportation and storage', 'water supply, sewerage, waste management and remediation activities' and 'other service activities'. It is interesting to notice that, as we go down from top to the bottom exposed scores, the respective sectors became progressively more intensive in human work. Data processing tasks can be performed by AI almost on its one. It is also curios that 'Human health and social work activities' is one of the most exposed sectors, as it is a sector that requires a significant amount of human work. However, this sector groups two main components: human health and social work. While the later registers low exposure levels, the former is substantially more exposed, due to hospital and laboratorial occupations that require cognitive effort on analysis and diagnose tasks that suit AI's capabilities.

Graphic 7. Average exposure by activity sectors (CAE)



4 Method

This section starts by describing how the C-AIOE measure from Pizzinelli et al. (2023) was replicated and applied to the Portuguese economy labor structure. Then, we present the econometric approach used to estimate the impact of adoption of AI technologies on wages, productivity and employment levels.

4.1 Exposure Measurement

4.1.1 Conceptual explanation

The original AIOE was created by Felten et al. (2021) with the purpose of measuring the relative exposure of occupations to artificial intelligence, through the overlap between a set of 10 AI applications and a sample of 52 human capabilities required to perform each occupation. Analytically speaking, this measure is given by:

$$AIOE_i = \frac{\Sigma_{j=1}^{52} A_{kj} \cdot L_{ji} \cdot I_{ji}}{\Sigma_{j=1}^{52} L_{ji} \cdot I_{ji}}$$

Let k represent the AI application, L_{ji} how often the human ability j is required to perform the occupation i, I_{ji} how crucial is that ability to perform the occupation and A_{kj} how exposed to AI is that same ability. Formally:

$$A_{kj} = \sum_{k=1}^{10} x_{kj}$$

Regarding Felten et al. (2021) work, Pizzinelli et al. (2023) claimed AIOE to be agnostic about exposure, once it only indicates how occupations are exposed to AI but does not inform whether that exposure is towards substitution or complementarity effects. Built on this limitation, Pizzinelli et al. (2023) adjusted the AIOE measure by applying a complementarity correction factor, creating the so-called C-AIOE.

To create the adjustment parameter θ, the authors started by selecting a set of 11 work contexts from O*NET database that they considered to be most likely to have human work displaced by AI. Then, based on the selected work contexts and the O*NET job zones, they created six groups as cornerstones of θ:

1. **Communication:** Face-to-Face and Public speaking.
2. **Responsibility:** Responsibility for outcomes and Responsibility for others' health.
3. **Physical Conditions:** Exposure to outdoors environments and Physical proximity to others.

4. **Criticality:** Consequence of errors, Freedom of decisions and Frequency of decisions.
5. **Routine:** Degree of automation and Unstructured vs structured work.
6. **Skills:** O*NET job zones.

Then, once job zones values range from 0 to 100 but work contexts from 1 to 5, the authors rescaled job zones by multiplying them by 20, so they range from 20 to 100:

$$WC_i \in [0, 100]$$

$$JZ_i \in \{1,2,3,4,5\}$$

$$JZ_i^{rescaled} = JZ_i \times 20 = \{20, 40, 60, 80, 100\}$$

Besides that, before aggregating the work context values by group, it is mandatory to ensure that all components are scaled to the same extent: the higher the value, the greater its complementarity potential is. In this way, the work context 'Degree of Automation' had to be invested, by subtracting the original value from 100, so that higher values of this work context mean higher complementarity potential:

$$Degree\ of\ Automation_i^{inverted} = 100 - Degree\ of\ Automation_i^{original}$$

With all components scaled accordantly, the authors then calculated the average value of the six groups – let them be denoted by C_{gi}, where g represents each group and i each occupation. For groups 1 to 5 C_{gi} corresponds to the simple arithmetic average of work contexts, while for group 6, naturally, C_{gi} has directly correspondence to the job zones values. Then, they applied the simple arithmetic average of C_{gi} for all groups g and divided everything by 100, so that θ ranged from 0 to 1. Formally:

$$\theta_i = \frac{\frac{1}{6}\cdot\sum_{g=1}^{6} C_{gi}}{100}, \text{ where } \theta_i \in [0,1]$$

The closer θ values are to 1, the higher the complementarity potential of occupation i with AI technologies. Conversely, the closer they are to 0, the less likely that occupation is to be complemented. However, the original AIOE from Felten et al. (2021) is a relative exposure measure, what means that it measures how exposed is a given occupation comparatively to the others. Hence, instead of applying the adjustment factor directly, Pizzinelli et al. (2023)

adjust θ_i relatively top its lower values, so that it represents the complementarity potential each occupation comparatively to the lowest scored one[10]. Analytically:

$$CAIOE_i = AIOE_i \cdot [1 - (\theta_i - \theta_{min})] \quad (4.1)$$

Based on the argument that AIOE measures exposure in a broad sense (absorbing both substitution and complementarity effects), by subtracting the relative complementarity score, we obtain the remaining substitution effect, denoted by C-AIOE. Occupations with a relatively higher complementarity potential will be less likely to be substituted by AI and vice versa. In this way, higher C-AIOE levels translate stronger exposure to substitution, while lower C-AIOE values reveal a prevalent complementarity potential. Notice, however, that a low C-AIOE does not mean a less exposed occupation to AI, but rather that the exposure is towards complementarity instead of substitution. It is also important to take note that the trend of exposure towards substitution or complementarity is not binary. C-AIOE is a continuous variable, what means that there is a whole spectrum between substitution and complementarity, with different possible combinations. As an analogy, we may think about it as a needle that ranges from one side to the other. For instance, let's say that the left side is the substitution likelihood and on the right side the complementarity likelihood. If the needle is closer to the left side than it is to the right side, then that occupation is more likely to be substituted by AI. However, it does not mean that some tasks within that occupation may not be complemented. The same logic applies when it is closer to the right side.

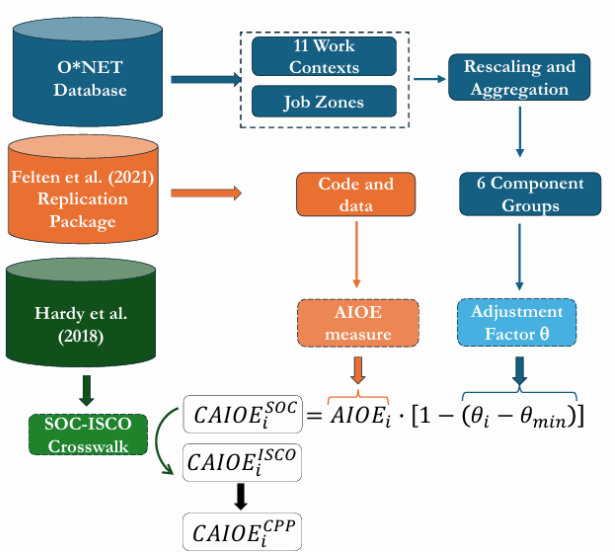
4.1.2 Measure replication

To replicate the pretended measure, as illustrated on Figure 1, we started by extracting directly the AIOE metric from Felten et al. (2021) replication package. Then, based on Pizzinelli et al. (2023) description on how they created the C-AIOE measure, we replicated the same process step by step. As illustrated below on Figure 1, we started by extracting the work context, job zones and SOC codes from O*NET database. Made the mentioned rescale adjustments, aggregated the variables on the same six groups, calculated θ (equation X), applied it to AIOE and obtained the C-AIOE (equation X). Finally, to convert the SOC classification into ISCO08 we used the occupation classifications SOC provided by Hardy et al. (2018). As the CPP10 were created based on the structure of ISCO08, there is direct correspondence between both classification codes up to 4 digits for all classification groups

[10] As mentioned by Pizzinelli et al. (2023), notice that for the occupation with the lowest θ, the AIOE and C-AIOE measures will have the same value.

(INE, 2011, pp. 14-15), except for 'Forças Armadas'. Since military occupations go outside the scope of this study, the use of this assumption does not affect the validity of the correspondence approach adopted. On total, our sample comprises a set of 417 distinct CPP10 unit group codes – a value consistent with the expected universe of Portuguese civilian occupations.

Figure 1. Illustration of the construction of C-AIOE by CPP10



Source: Own elaboration.

4.2 Estimation strategy

4.2.1 Baseline OLS model

To estimate the impact of AI on wages, productivity and employment we started by dividing the approach into two parts: the workers level estimation (wages) and firm level estimation (productivity and employment).

First, to establish a baseline benchmark of the relation between exposure (given by C-AIOE) and the targeted dependent variables, we started by running OLS regressions with pertinent control variables and fixed effects, depending on the dependent variable.

Regarding the worker level estimation, presented in equation 4.2, we used the logarithm of wage for interpretation proposals and to smooth out the variable oscillations. The wage level of an individual may be affected by multiple determinants. Individual characteristics such as tenure, age, level of education or sex can influence how much that worker earns. For instance, on average, a person with higher education has access to job opportunities that pay higher wages. To account for such external determinants and isolate the influence of exposure on wages, we used tenure, age, education and their quadratic forms as control variables (denoted

by $\mathbf{X'_{ijt}}$), plus fixed effects (represented by **FE**) by sex, year, firm and worker. Besides, as wages vary across workers and workers within the same firm are exposed to common shocks, we used cluster-robust standard errors at firm level.

To test different scenarios, we runed four different regressions, with different combinations of fixed effects. On regressions 1 and 2 we runed the estimation with fixed by year and by firm, accordantly, and kept all controls. Then, regression 3 used fixed effects by workers and omitted the control dummy sex, once the fixed effects used already capture workers individual characteristics. Equation 4.2 may be represented as follows, where subscripts i, j, o and t represent workers, firms, occupations and years:

$$ln(wage)_{ijot} = \alpha + \beta \cdot C\text{-}AIOE_{ot} + b \cdot \boldsymbol{X'_{ijt}} + \boldsymbol{FE} + \varepsilon_{ijot} \quad (4.2)$$

Moving on to firm level estimation, we used GVA at market prices as proxy of productivity and people employed by the firm as proxy of employment. As done for wage regressions, here we also used the logarithm of the dependent variables and applied fixed effects as well as cluster-robust standard errors at firm level.

To isolate AI exposure effects on GVA we used equity, gross fixed capital formation, people employed and investment in software as control variables ($\mathbf{X'_{jt}}$), plus fixed effects **(FE)** by year, firm, region and activity sector. To explore different scenarios, three different OLS combinations were run. On regression 1 we estimated the baseline regression of the impact of C-AIOE on GVA, controlling by the mentioned variables and fixed effects at year and firm level. However, as documented in section 2.2, the exposure levels don't appear to be randomly distributed by sectors and regions. As such, to make sure that there is no spurious correlation between exposure and GVA induced by geographic presence and sector of activity of firms, on regression 2 we added fixed effects for both region and activity sector. Furthermore, as suggested by some contributions in literature, AI adoption by itself may not be enough, requiring investment in complementary factors, such as support infrastructures, to boost AI's impact. In this way, on regression 3 we added the interaction between C-AIOE and investment in software, to explore whether firms that invest more in software have hire impact of AI on GVA.

Regarding employment, regression 1 estimated the baseline effects using equity, gross fixed capital formation and investment in software as control variables ($\mathbf{X'_{jt}}$) and fixed effects (**FE**) by year and firm. Then, as done for GVA estimation, we also applied fixed effects by region and activity sector on regression 2 and the interaction between C-AIOE and investment in

software on regression 3. Regressions 1 and 2 on GVA and employment are represented by equation 4.3, while both regressions 3 are represented on equation 4.4.

$$(4.3) \quad ln(Y)_{jt} = \alpha + \beta \cdot C\text{-}AIOE_{jt} + b \cdot \boldsymbol{X'_{jt}} + FE + \varepsilon_{jt}$$

$$(4.4) \quad ln(Y)_{jt} = \alpha + \beta_1 \cdot C\text{-}AIOE_{jt} + \beta_2 \cdot (C\text{-}AIOE \times Invest.\ Soft.)_{jt} + b \cdot \boldsymbol{X'_{jt}} + FE + \varepsilon_{jt}$$

4.2.2 IV estimation

4.2.2.1 Model specification

Despite the use of relevant fixed effects and control variables, the OLS estimations do not exclude the possibility of endogeneity at C-AIOE level. More specifically, there may be worker and firm characteristics not included in the regressions that affect simultaneously the exposure measure and the dependent variable – s.a. organizational culture, management quality and innovation capability, which are difficult to measure. Besides that, we also cannot exclude the possibility of reverse causality, since already highly productive and high-employing firms may attract highly exposed occupations, which is the opposite of what we are trying to estimate. Such problems do not ensure Gauss-Markov assumption of zero conditional mean between error term and C-AIOE, which makes OLS coefficients biased. To address this potential endogeneity concern, we employed an instrumental variable (IV) approach, creating a Bartik shift-share instrument for C-AIOE, based on the exposure measure itself. Formally Bartik instrument is composed of a shift-share logic. In this scenario, "shift" represents the average aggregated growth of exposure by sector (s) between 2016 and year t, while "share" denotes the average exposure of qualifications by sector in 2016.

$$bartik_{q,s,t} = S_{q,s}^{2016} \times \left(1 + g_{s,t}\right)$$

For the "share" component ($S_{q,s}^{2016}$) we applied the following, where $N_{q,s}$ is the number of workers by qualifications (q) and sector (s) and $C\text{-}AIOE_{i,2016}$ the exposure by sector in 2016:

$$S_{q,s} = \overline{C\text{-}AIOE}_{q,s,2016} = \frac{1}{N_{q,s}} \sum_{i \in (q,s)} C\text{-}AIOE_{i,2016}$$

Then, for the component "shift" $(1 + g_{s,t})$ we calculated $g_{s,t}$ as presented on below , where $\overline{C\text{-}AIOE}_{s,t}$ is the average exposure level by sector *s* and year *t* and $\overline{C\text{-}AIOE}_{s,2016}$ the average exposure of sector *s* for the baseline year 2016.

$$g_{s,t} = \frac{\overline{C\text{-}AIOE}_{s,t} - \overline{C\text{-}AIOE}_{s,2016}}{\overline{C\text{-}AIOE}_{s,2016}}$$

$$\overline{C\text{-}AIOE}_{s,t} = \frac{1}{N_{s,t}} \sum_{i \in s,t} C\text{-}AIOE_{i,t}$$

$$\overline{C\text{-}AIOE}_{s,2016} = \overline{C\text{-}AIOE}_{s,t}|_{t=2016}$$

Finally, moving to the IV approach, $bartik_{q,s,t}$ was used to instrument C-AIOE on a 2SLS estimation sequence. The first stage (equation 4.5) regresses C-AIOE by the instrument $bartik_{q,s,t}$ and the control variables $X'_{\cdot t}$, plus the specified fixed effects, obtaining the exogenous $\widehat{C\text{-}AIOE}_{\cdot t}$. On second stage (equation 4.6), the previous OLS regressions are estimated but replacing the C-AIOE with their exogenous version $\widehat{C\text{-}AIOE}_{\cdot t}$.

$$C\text{-}\widehat{AIOE}_{\cdot t} = \pi_0 + \pi_1 bartik_{q,s,t} + X'_{\cdot t}\pi + FE + \upsilon_{\cdot t} \tag{4.5}$$

$$Y_{\cdot t} = \alpha + \beta_1 C\text{-}\widehat{AIOE}_{\cdot t} + X'_{\cdot t}\gamma + FE + \varepsilon_{\cdot t} \tag{4.6}$$

Although, in theory, the panel data structure would suggest the use of clustered robust standard errors by firm level, we decided to estimate IV regressions with robust standard errors instead, because the constructed instrument has low within-cluster variation and varies across sector instead of firms.

4.2.2.2 Validity of bartik instrument

The presented IV approach is, however, only acceptable to solve C-AIOE endogeneity problem if following conditions are held:

Condition 1: under the condition of relevance, the bartik instrument must be correlated with the endogenous variable so that the IV estimator is sufficiently strong to isolate the exogenous effects from C-AIOE. Formally:

$$\text{Cov}\left(bartik_{q,s,t}, C\text{-}AIOE_{i,t}\right) \neq 0$$

Condition 2: under the condition of exogeneity, the bartik instrument must be independent from non-observable factor that influence the outcomes. In another words, the instrument

must not be correlated with possible omitted variables, represented by the error term. Formally:

$$\text{Cov}\left(bartik_{q,s,t}, \varepsilon_{\text{jt}}\right) = 0$$

For the first condition, we analyzed the underidentification and weak instrument identification test. Both tests check the correlation between the instrument and the instrumented variable. The underidentification test is the first stage of analyses and tests the hypotheses of null correlation. As for weak instrument identification test, it reports how strong is that correlation by comparing the Kleibergen-Paap rk Wald F statistic with Stock-Yogo weak ID test critical values. For regressions 1 to 3, underidentification test rejected the null hypotheses, indicating identified models – as would be expected, since the instrument was constructed based on the endogenous variable. Then, regarding weak identification test, the Wald F statistic reported values significantly above the 10% maximum IV size threshold from Stock-Yogo weak ID test critical values, meaning that there is no evidence of weak instrument and, therefore, condition 1 is verified.

Moving to the second condition, as we are working with exactly identified models – exact and strict number of instruments and endogenous variables – it is not possible to formally test this preposition. As such, the validity of condition 2 rests entirely on theoretical argumentation. In this case, since the share component of the shift-share instrument is fixed at the baseline year (2016), it is impossible, by construction, for a reverse causality bias to arise from outcomes on subsequent years. Furthermore, the measurement of exposure is based on the occupational structure, who adapts itself to technological development by changing the allocation of productive factors (s.a. work and capital) across tasks. Besides, the shift component of the instrument represents the aggregated average exposure by sectors across time, so that firms individually don't have enough weight to influence significantly the sector average. However, this strategy is not bulletproof. We must acknowledge that, even down firm level shocks don't impact instrument's time-variant component, it may be affected by endogenous sector level shocks that impact simultaneously all firms. For example, if there is an external shock on aggregated demand caused by increase in exports of a specific sector, it could lead firms to rearrange labor into higher value-added products related to occupation naturally more exposed to AI. In this way, the shock impacted sector GVA and C-AIOE simultaneously, making model IV estimator biased.

5 Results

This section presents the empirical estimation results, organized in two parts. First, we report the OLS estimation for both worker and firm level regressions and then move to IV models. For each set of results, we interpret the estimated coefficients, report their significance and robustness and compare IV results with the base line OLS specifications.

Table 1 contains the OLS estimation results of C-AIOE on wages. As we can see, C-AIOE coefficients are positive and statistically significant at 1% across regressions 1 to 4. Since the outcome variable is in logarithmic form, these regressions are in a Log-Lin specification, what means that C-AIOE coefficient is a semi-elasticity, and so has to be interpreted as exact semi-elasticity, transforming log-points into percentage, applying:

$$\%\Delta Y = \left(e^{\beta_{C-AIOE}} - 1\right) \cdot 100$$

To this extent, for each additional unit of exposure score C-AIOE, workers' wages increase on average 17.27%, 14.98%, 7.77% and 6.57%, respectively for regressions 1, 2, 3 and 4. These results suggest that workers performing tasks on high exposed occupations to AI's substitution effect earn higher wages, on average. As the model becomes more saturated by fixed effects, the coefficient magnitude decreases progressively – from regression 1 (fixed effects by year) to regression 4 (fixed effects by year, firm and worker) the C-AIOE coefficient falls by less than half. This constatation may be a sign that there is non-observable heterogeneity being diluted, as we control for more fixed effects.

Table 1. OLS estimates of the impact of C-AIOE on Wages

OLS regressions	**(1)**	**(2)**	**(3)**	**(4)**
Dependent variable:	Log(wage)	Log(wage)	Log(wage)	Log(wage)
C-AIOE	0.1593*** *(0.0069)*	0.1396*** *(0.0047)*	0.0748*** *(0.0063)*	0.0636*** *(0.0056)*
Fixed effects				
Year	Yes	No	No	Yes
Firm	No	Yes	No	Yes
Worker	No	No	Yes	Yes
N	14,459,791	14,459,488	13,721,463	13,719,217
R^2	0.2938	0.4562	0.7171	0.7442
Within R^2	0.2818	0.1316	0.0830	0.0122

Note: Robust standard errors, clustered by firm, in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

Next, Table 2 presents the estimation of C-AIOE on firm GVA across regressions 1, 2 and 3. For the three columns, C-AIOE coefficients are negative and statistically significant at 1%. Here we also have a Log-Lin specification, what means that for each additional unit of exposure score C-AIOE, firm's GVA decreases on average 4.69% (regression 1) and 4.63% (regressions 2 and 3). On another words, the more workers a firm employs on high exposed occupations, the smallest GVA that firm will have. From regressions 1 to 2 and 3 the coefficient decreases slightly, what may be explained by the added fixed effect controls at sector and regional levels on regression 2 and 3. Furthermore, the coefficient of interaction between C-AIOE and firm's investment in software is non-significant, what implies that, for the data and specification used, there appears to be no evidence that investment in software influences C-AIOE impact on GVA.

Table 2. OLS estimates of the impact of C-AIOE on GVA

OLS regressions	**(1)**	**(2)**	**(3)**
Dependent variable:	Log(GVA)	Log(GVA)	Log(GVA)
C-AIOE	-0.0480*** *(0.0083)*	-0.0474*** *(0.0083)*	-0.0474*** *(0.0083)*
C-AIOE × Software investment	—	—	-1.87e-09 *(1.30e-08)*
Fixed effects			
Year	Yes	Yes	Yes
Firm	Yes	Yes	Yes
Sector (CAE)	No	Yes	Yes
Region (NUTS II)	No	Yes	Yes
N	223,452	223,447	223,447
R^2	0.9514	0.9518	0.9518
Within R^2	0.2727	0.2711	0.2711

Note: Robust standard errors, clustered by firm, in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

Table 3 contains the estimation results of the impact of C-AIOE on firm's employment level. Like previous OLS estimates, for regressions 1 to 3 the estimation coefficients of C-AIOE are highly significant at 1%. Besides, the coefficients are also interpreted as semi-elasticity and, as in GVA regressions, they reveal a negative impact of exposure on firm's employment levels – for each additional unit of C-AIOE, firm's will employ 14.27%, 14.17% and 14.16% less workers, respectively for equations 1, 2 and 3. However, unlike GVA, the interaction

between C-AIOE and investment in software is statistically significant, meaning that firms that invest more in software are more affected by the exposure to AI's substitution effect.

Table 3. OLS estimates of the impact of C-AIOE on Employment

OLS regressions	**(1)**	**(2)**	**(3)**
Dependent variable:	Log(employees)	Log(employees)	Log(employees)
C-AIOE	-0.1540*** *(0.0081)*	-0.1528*** *(0.0080)*	-0.1527*** *(0.0080)*
C-AIOE × Software investment	—	—	-1.77e-08** *(7.55e-09)*
Fixed effects			
Year	Yes	Yes	Yes
Firm	Yes	Yes	Yes
Sector (CAE)	No	Yes	Yes
Region (NUTS II)	No	Yes	Yes
N	198,572	198,567	198,567
R^2	0.9681	0.9685	0.9685
Within R^2	0.0507	0.0501	0.0502

Note: Robust standard errors, clustered by firm, in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

Now, moving on to the IV results, Table 4, 5 and 6 contain the 2SLS estimation of the relation between C-AIOE (instrumented by bartik instrument) and wages, GVA and employment, respectively, with robust standard errors.

First, on Table 4 we can see that the 2SLS estimation coefficient of C-AIOE is consistently positive and significative at 1% for regressions 1 to 4. For each additional point on C-AIOE, wages increase 37.14% (regression 1), 50.72% (2), 20.56% (3) and 25.92% (4).

The result on regression 4 allows us to predict the impact on wage differentials between occupations, using the following formula where a and b are different occupations:

$$\%\Delta Wage_{a,b} = \left(e^{\beta \times (C\text{-}AIOE_a - C\text{-}AIOE_b)} - 1\right) \times 100$$

Hence, for all controls and fixed effects used, the exposure to AI reflects an average wage difference between a financial analyst (the most exposed occupation) and a hand launderer (the less exposed occupation) of 101.50%. As for between a chief financial officer and a bank teller, for example, the predicted differential is about 29.14%, while between a civil engineer and a bricklayer is 62.05% and between a choreographer and a translator is 80.96%. Furthermore, while OLS coefficient 4 is lower than the other OLS, IV coefficient 4 is lower

than 1 and 2, but higher than 3. As regression 3 on OLS and IV models use fixed effects by worker, while 4 accounts effects on year, firm and worker level, the differences between coefficient magnitude is most likely due to the fact that IV specification absorbs only the exogenous relation between C-AIOE and outcomes – the instrumentation of C-AIOE in a 2SLS on IV models filters only the exogenous behavior of C-AIOE, to make sure that the coefficients are not biased. As such, it is plausible that the OLS estimates are being undervalued by endogeneity and eventual measurement errors.

Table 4. IV estimates of the impact of C-AIOE on Wages

IV regressions	**(1)**	**(2)**	**(3)**	**(4)**
Dependent variable:	Log(wages)	Log(wages)	Log(wages)	Log(wages)
C-AIOE (2SLS)	0.3158*** *(0.0008)*	0.4101*** *(0.0013)*	0.1870*** *(0.0042)*	0.2305*** *(0.0072)*
Fixed effects				
Year	Yes	No	No	Yes
Firm	No	Yes	No	Yes
Worker	No	No	Yes	Yes
N	12,462,920	12,462,645	11,721,968	11,719,545
R^2 (centered)	0.2568	0.0738	0.0698	0.0037
Instrument diagnostics				
Underidentification test	1.4e+06***	4.4e+05***	1.2e+05***	4.7e+04***
Weak identification test	7.4e+04	3.3e+04	2.4e+04	1.5e+04

Note: Robust standard errors, clustered by firm, in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

Next, on Table 5 is presented the estimation of IV regressions for GVA. Like in Table 2, the coefficients are negative and significant, even down here only at 5%. For each additional unit of firm C-AIOE (instrumented), their GVA decreases 19.33% (equation 1), 20.66% (2) and 20.63% (3). This means that firms that employ more workers in high exposed occupations will have higher losses of GVA than those how employ less workers on such occupations. Once more, the magnitude of IV estimates is far larger than in OLS models. On regression 1 we are controlling the fixed effects by year and firm. Then, regression 2 uses the same controls, but adds fixed effects by sector and region. Between these two regressions 2SLS coefficient decreased, while in OLS it increased slightly. Like on wage regressions, the endogeneity bias of OLS estimator is likely to explain such differences. Regarding regression 3, as also suggested in Table 2, there is no evidence that firms that invest more in software are higher or less exposed to AI.

Table 5. IV estimates of the impact of C-AIOE on GVA

IV regressions	**(1)**	**(2)**	**(3)**
Dependent variable:	Log(GVA)	Log(GVA)	Log(GVA)
C-AIOE (2SLS)	-0.2148** *(0.0878)*	-0.2312** *(0.0898)*	-0.2308** *(0.0899)*
C-AIOE × Software investment	—	—	-5.87e-09 *(2.62e-08)*
Fixed effects			
Year	Yes	Yes	Yes
Firm	Yes	Yes	Yes
Sector (CAE)	No	Yes	Yes
Region (NUTS II)	No	Yes	Yes
N	191,915	191,912	191,912
R^2 (centered)	0.2512	0.2491	0.2491
Instrument diagnostics			
Underidentification test	1,055.6***	1,014.6***	1,013.0***
Weak identification test	227.0	217.1	216.6

Note: Robust standard errors, clustered by firm, in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

Finally, Table 6 revealed that firms that employ more workers in high exposed occupations will have less employees have less than those with less workers in high exposed occupations. More specifically, the coefficients of C-AIOE appeared as negative and significant at 1%. For each addition unit of firm's C-AIOE, they will lose 31.81% employees (regression 1), 31.15% (2) and 31.10% (3). As also found in Table 3, the estimates do not change in a considerably way across the three columns. However, from OLS to 2SLS estimation the interaction between C-AIOE and investment in software lost its significance. A possible explanation for such results may be that regression 3 interacts with software investment with non-instrumented C-AIOE and, therefore, the estimate may be biased.

Table 6. IV estimates of the impact of C-AIOE on Employment

IV regressions	**(1)**	**(2)**	**(3)**
Dependent variable:	Log(employees)	Log(employees)	Log(employees)
C-AIOE (2SLS)	-0.3830*** *(0.0528)*	-0.3732*** *(0.0534)*	-0.3725*** *(0.0534)*
C-AIOE × Software investment	—	—	-1.44e-08 *(1.09e-08)*
Fixed effects			
Year	Yes	Yes	Yes
Firm	Yes	Yes	Yes
Sector (CAE)	No	Yes	Yes
Region (NUTS II)	No	Yes	Yes
N	194,583	194,580	194,580
R^2 (centered)	0.0298	0.0307	0.0309
Instrument diagnostics			
Underidentification test	1,126.4***	1,087.5***	1,086.4***
Weak identification test	235.2	225.0	224.5

Note: Robust standard errors, clustered by firm, in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

Overall, across wage, GVA and employment, the IV produced larger estimates in magnitude than OLS, especially for the more complete specifications. Based on these regressions, a one-unit increase in C-AIOE leads to a 25.92% increase in employees' wages, 20.63% decrease in firms' GVA and 31.10% decrease in firms' employment. These results support the idea of a wage premium for workers that perform tasks in occupations more exposed to AI and that firms how employ such workers obtain lower value added and employment levels, indicating that productivity and labor-displacing effects of AI are more noticeable at the firm level than the wage gains at worker level would suggest. Additionally, for GVA the interaction term between C-AIOE and firm's investment in software is statistically non-significant, reinforcing the suspects on OLS estimates of no evidence supporting that such investment would moderate the impact of AI's substitution effect on firms GVA. As for employment, this interaction term lost its significance found on OLS when applied the IV specification, most likely because such term is not instrumented and hence remains exposed to the endogeneity or eventual measurement error that motivated the AI approach. Gathering all these findings, it becomes clear that, if anything, the associations identified in the OLS estimations underestimates the true magnitude of C-AIOE's effect on wages, GVA and employment, reflecting the endogeneity and measurement error bias of the OLS estimators.

6 Discussion & Limitations

This section critically reflects on the empirical results presented above, comparing them with similar studies within literature on AI's impact on labor market outcomes and discussing the main limitations of this study's method.

6.1 Discussion

The founded wage premium on Table 4 is consistent with recent generative AI exposure literature that tests the possibility of complementarity effects. Many studies like Tolan et al. (2021), Webb (2020), Felten et al. (2021) and Eloundou et al. (2023) have found evidence that higher exposed occupations tend to demand more non-routine cognitive tasks which are usually performed by higher paid and qualified workers. This being said, we must interpret this wage premium very carefully. C-AIOE measure is constructed based on two factors: AIOE from Felten et al. (2021) and the complementarity adjustment factor θ from Pizzinelli et al. (2023). While AIOE captures the broad exposure to AI, the adjustment factor targets the complementarity potential of each occupation. As such, by construction, the C-AIOE represents the net exposure to AI's substitutions effects. However, in practice, a higher C-AIOE score could result either from a high general exposure to AI (AIOE) or a low complementarity potential of that occupation. So, to properly identify the source of the documented wage premium, we must bear in mind that AIOE is highly associated with workers' skills and wage levels (Felten et al., 2021) and therefore, it is very likely that such wage premium is mainly due to the occupational general relatedness to AI technologies.

As for the firm-level results, the finding of productivity losses associated with higher exposure to AI, documented in Table 5, also find support on evidence from literature. Studies like Babina et al. (2024), Bessen et al. (2022), Brynjolfsson et al. (2018) and Tolan et al. (2021) haven't also found evidence of productivity gains related to AI adoption and that, at most, there could be productivity losses. Brynjolfsson et al. (2018) and Tolan et al. (2021) report a Modern Productivity Paradox where even down AI's capabilities have been evolving, those breakthroughs don't generate productivity increases. As such, we can say that our study also found evidence that supports this paradox. However, it is important to notice that the estimate of 23.08% productivity decrease for each additional point of exposure seems to have a quite large impact, spatially when compared with literature's findings that are closer to zero than a considerable negative value. Such discrepancy may be due to the use of firm's GVA as productivity proxy, absorbing other effects different from productivity.

Finally, regarding the impact of AI on employment, the finding suggests the existence of a labor displacement effects leaded by firms exposure to AI. According to Acemoglu and Restrepo (2018) and Acemoglu and Restrepo (2019a), this could mean that the supply of employee's qualifications is not being enough to satisfy the creation of new tasks driven by AI. Yet, similar studies like Babina et al. (2024) and Webb (2020) have not found evidence of employment decrease effect. Such divergence, however, may be explained by differences in methodological approaches. While C-AIOE captures the exposure potential to AI, based on task composition of occupations, Babina et al. (2024) measured AI adoption through the actual share of workers with AI skills and Webb (2020) extrapolate AI's impact based on the assumption that it would follow the same linear pattern found on robots and software. Besides, adoption of this kind of technology is still low even in developed economies (Zolas et al., 2020) whereby it is possible that C-AIOE is not capturing the whole characteristics of exposed sectors, at this stage of AI diffusion.

6.2 Limitations

While previous results reveal meaningful insights about the relation between AI exposure and labor market outcomes, naturally there are some methodological limitations that must be disclosed.

The first limitation of this study lies in the use of the C-AIOE measure. Although this indicator allows distinguishing the replacement component with the complementarity of exposure to artificial intelligence, this measure only reports the theoretical potential exposure level. That is, based on the composition and nature of the tasks of each occupation, it measures how likely these occupations are to be affected by AI capabilities. In this way, cases of workers with high C-AIOE may be considered, as they perform tasks highly exposed to AI, but without necessarily implying that the companies where they work have adopted any AI technology. This limitation is relevant because Zolas et al. (2020), for example, reveal that the effective adoption of advanced technologies (in which AI fits) is still low, so C-AIOE measurements may be reporting exposure levels above what is actually verified. Thus, by using this measure as a proxy for the adoption of IA, the estimation results may be slightly inflated – a plausible hypothesis, since the estimates obtained have a considerably high magnitude: 23.05% in wages, 23.08% in GVA and 37.25% in employment.

In addition, C-AIOE is a measure built based on two factors according to equation X: AIOE (gross potential exposure of occupations to AI) and θ (correction factor for the potential

complementarity). This construction, although it allows a more informed reading of the effects of AI (substitution and complementarity) to which occupations are more exposed, in econometric terms, by regressing C-AIOE with labor market outcomes, the estimated coefficient cannot isolate which of the components is truly driving the results, losing the distinction between the substitution and complementarity effects. Still regarding the construction of this measure, in the process of converting the occupational classification of raw data from O*NET-SOC to ISCO08 and later from ISCO08 to CPP10, it was necessary to slightly aggregate the occupational levels for 4-digit codes, losing observational granularity and capturing heterogeneity within the occupations. In other words, workers within the same occupation, but who perform different tasks with different exposure to AI, are receiving the same C-AIOE.

Furthermore, there are also some limitations regarding the identification strategy of the used econometric approach. More specifically, the shift-share bartik instrument used to instrumentalize the C-AIOE variable is constructed from the endogenous variable itself, for both share (average C-AIOE in 2016) and shift (aggregate average variation of C-AIOE per firm) components. In this way, we are trying to absorb the heterogeneous effects of C-AIOE through the variation of the measure itself at firm level. As such, the exogeneity condition of the instrument, by construction, is not compromised by any endogenous shocks at the firm level, as these do not individually have sufficient weight to influence the average evolution of C-AIOE per sector. However, the same cannot be said for possible endogenous sectoral shocks. The fact that the variation component of the instrument is introduced by aggregating the endogenous variable itself does not guarantee that the instrument is not capturing any aggregated endogenous shocks at sectoral level, which simultaneously affect all undertakings in the sector, undermining the exogeneity of the Bartik instrument and thus its validity.

7 Conclusion

The present dissertation aimed to estimate the impact of exposure to artificial intelligence technologies on the wages, productivity and employment levels of the Portuguese labor market for the period from 2016 to 2023. To this end, a strategy IV was applied with an instrumental variable Shift-share in a 2SLS estimation, using the C-AIOE index – introduced by Pizzinelli et al. (2023) as a measure of AIOE gross exposure by Felten et al. (2021), adjusted for complementarity – as an explanatory variable of interest.

This research falls within the branches of empirical literature on the evaluation of the effects of IA on labor market outcomes and the task-based exposure measurement literature. In this sense, the results obtained contribute mainly as clues to understanding the effects of IA on labor market components. However, they also help to understand the level of potential exposure to which workers, companies and sectors of activity in Portugal are subject. In addition, this study also stands out for the application still little explored in the literature of the C-AIOE index by Pizzinelli et al. (2023) to national matched employer-employee administrative data in an econometric estimation process in the context of the national economy.

The results obtained at the employee level revealed the existence of a positive and statistically significant wage premium associated with C-AIOE, which aligns with previous literature such as Tolan et al. (2021), Webb (2020), Felten et al. (2021) and Eloundou et al. (2023). At the company level, negative and statistically significant effects were recorded regarding its GVA levels and employment. The negative effects on GVA suggest the existence of productivity losses resulting from exposure – a result that is part of the "Modern Productivity Paradox" proposed by Brynjolfsson et al. (2019) and reinforced by Tolan et al. (2021). On the other losses are in line with the displacement effect suggested by Acemoglu and Restrepo (2018, 2019a), but contrast with the results of Babina et al. (2024) and Webb (2020). In addition, no statistically significant empirical evidence was found regarding the effect that investment in complementary technologies (such as software) by companies could have on the impact of exposure to AI. Estimates IV revealed a magnitude systematically higher than OLS estimates, which is a result compatible with the suspicions of endogeneity and attenuation bias in OLS.

As in any study, there are some limitations that, for the sake of scientific transparency, should be clarified. First, as the C-AIOE index only reports potential exposure levels, it does not measure the actual adoption of AI technologies, so the interpretation of its results should

ensure that its effects are only approximate to what would actually be observed with the actual exposure level. On the other hand, the fact that it is an indicator constructed by the product of two components (substitution and complementarity) is difficult to econometrically isolate the effect of each component, not allowing us to know which of these components is being captured in the estimation results. In addition, in the conversion between occupational classifications there was a loss of granularity that may be hiding heterogeneous effects within the occupations. Finally, with regard to the identification strategy in IV, the use of an instrument based on the Shift-share aggregation of the endogenous variable itself may compromise the validity of the instrument in the face of possible endogenous sectoral shocks.

Regarding avenues for future research, it is suggested that further studies of the same scope use an identification strategy with alternative instrumental variables to the one used in this study, of a truly exogenous nature, in order to ensure greater robustness and instrumental validity. For example, the construction of an instrument based on the growth of patent design of AI technologies would introduce a more robust temporal variation component than the aggregate sectoral average of C-AIOE. At the same time, it would be interesting to use an impact variable that approaches the true levels of investment of companies in AI and that explores heterogeneity within occupations. To this end, an indicator of investment in AI could be constructed based on conducting a survey of a significant sample of resident firms. Alternatively, a proxy could be built through data on job openings for AI-related positions (such as deep learning, machine learning, neural networks or optical character recognition). In addition, as documented, the proliferation of AI in markets may not yet be sufficiently developed, so it will be relevant to extend the review period as more years become available.

In conclusion, this dissertation found robust evidence on how Portuguese workers exposed to artificial intelligence, measured by C-AIOE, benefit from wage premiums, while firms experience productivity and job losses. At a time when society is witnessing successive advances in artificial intelligence technologies, understanding their effects on economic activity is essential to allow policymakers to have access to relevant information to define policies that allow them to enjoy the benefits of AI without compromising the proper functioning of the labor market – an objective to which this research aimed to contribute.

References

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

Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In O. Ashenfelter & D. Card (Eds.), *Handbook of labor economics* (Vol. 4B, pp. 1043–1171). Elsevier.

Acemoglu, D., & Loebbing, J. (2026). Automation and polarization. *Journal of Political Economy*, *134*(3), 1017–1072. https://doi.org/10.1086/739330

Acemoglu, D., & Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, *108*(6), 1488–1542. https://doi.org/10.1257/aer.20160696

Acemoglu, D., & Restrepo, P. (2019a). Artificial intelligence, automation, and work. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 197–236). University of Chicago Press.

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

Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, *128*(6), 2188–2244. https://doi.org/10.1086/705716

Acemoglu, D., & Restrepo, P. (2022). Demographics and automation. *Review of Economic Studies*, *89*(1), 1–44. https://doi.org/10.1093/restud/rdab031

Autor, D., Chin, C., Salomons, A. M., & Seegmiller, B. (2022). *New frontiers: The origins and content of new work, 1940–2018* (Unpublished working paper). Massachusetts Institute of Technology.

Autor, D. H., & 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. https://doi.org/10.1257/aer.103.5.1553

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

Babina, T., Fedyk, A., He, A., & Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, *151*, Article 103745. https://doi.org/10.1016/j.jfineco.2023.103745

Bessen, J., Denk, E., & Meng, C. (2022). *The remainder effect: How automation complements labor quality* (Boston University School of Law, Law and Economics Research Paper No. 22-3). https://ssrn.com/abstract=4042317

Brynjolfsson, E., & Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science*, *358*(6370), 1530–1534. https://doi.org/10.1126/science.aap8062

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

Brynjolfsson, E., Rock, D., & Syverson, C. (2019). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 23–57). University of Chicago Press.

Dauth, W., Findeisen, S., Südekum, J., & Wößner, N. (2017). *German robots: The impact of industrial robots on workers* (IAB-Discussion Paper No. 30/2017). Institut für Arbeitsmarkt- und Berufsforschung.

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). *GPTs are GPTs: An early look at the labor market impact potential of large language models* (arXiv:2303.10130). arXiv. https://doi.org/10.48550/arXiv.2303.10130

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

Felten, E., Raj, M., & Seamans, R. (2023). *How will language modelers like ChatGPT affect occupations and industries?* (SSRN Working Paper No. 4375268). https://doi.org/10.2139/ssrn.4375268

Gmyrek, P., Berg, J., & Bescond, D. (2023). *Generative AI and jobs: A global analysis of potential effects on job quantity and quality* (ILO Working Paper No. 96). International Labor Organization.

Goldin, C., & Katz, L. F. (2007). *The race between education and technology: The evolution of U.S. educational wage differentials, 1890 to 2005* (NBER Working Paper No. 12984). National Bureau of Economic Research. https://doi.org/10.3386/w12984

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

Graetz, G., & Michaels, G. (2018). Robots at work. *The Review of Economics and Statistics*, *100*(5), 753–768.

Hardy, W., Keister, R., & Lewandowski, P. (2018). Educational upgrading, structural change and the task composition of jobs in Europe. *Economics of Transition*, *26*(2), 201–231. https://doi.org/10.1111/ecot.12145

Instituto Nacional de Estatística. (2011). *Classificação Portuguesa das Profissões 2010*. INE. https://www.ine.pt/xurl/pub/107961853

Lin, J. (2011). Technological adaptation, cities, and new work. *The Review of Economics and Statistics*, *93*(2), 554–574.

Michaels, G., Natraj, A., & Van Reenen, J. (2014). Has ICT polarized skill demand? Evidence from eleven countries over twenty-five years. *The Review of Economics and Statistics*, *96*(1), 60–77.

Pizzinelli, C., Panton, A. J., Tavares, M. M., Cazzaniga, M., & Li, L. (2023). *Labor market exposure to AI: Cross-country differences and distributional implications* (IMF Working Paper No. WP/23/216). International Monetary Fund. https://doi.org/10.5089/9798400254802.001

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

Webb, M. (2020). *The impact of artificial intelligence on the labor market* (SSRN Working Paper No. 3482150). https://doi.org/10.2139/ssrn.3482150

Zolas, N., Kroff, Z., Brynjolfsson, E., McElheran, K., Beede, D. N., Buffington, C., Goldschlag, N., Foster, L., & Dinlersoz, E. (2020). *Advanced technologies adoption and use by U.S. firms: Evidence from the Annual Business Survey* (NBER Working Paper No. 28290). National Bureau of Economic Research. https://doi.org/10.3386/w28290

FEP FACULDADE DE **ECONOMIA E GESTÃO**