LUISS Guido Carli University
Einaudi Institute for Economics and Finance

Department of Economics and Finance

The impact of Artificial Intelligence on the Italian Labor Market

Chair of Applied Microeconomics

Supervisor:

Matteo Paradisi

Co-supervisor:

Fabiano Schivardi

Student:

Ottavia Pesenti

Student ID: 732861

Academic Year 2021/2022

The impact of Artificial Intelligence on the Italian Labor Market

Ottavia Pesenti*

RoME - Master in Economics

EIEF and LUISS

Abstract

What is the effect of the application of AI-based technologies to tasks currently performed by humans? We exploit Webb (2019) measure of exposure to AI and ISTAT LFS data to study the relationship between artificial intelligence and labor market outcomes. AI capabilities, as measured by patenting intensity, are most common in high-wage occupations, while being unbiased in educational attainment. The present analysis estimates no significant impact of occupational exposure to artificial intelligence on wages and employment between 2014 and 2019 in Italy: either adjustments in the labor market have not occurred yet, or occupation-industry level reinstatement and displacement effects counterbalance, or even Webb (2019)'s AI exposure, as a measure of potential exposure, does not reflect actual AI applications.

*I express my deep and sincere gratitude to the supervisor of this thesis, Professor Matteo Paradisi, for his guidance and support and for inspiring my interest in Labor Economics. I also acknowledge the insightful comments that I have received by EIEF faculty during the Thesis Control Sessions.

Contents

1 Introduction

As past waves of automation have reshaped our economies, artificial intelligence is promising far-reaching innovations in every field, from consumer to industrial applications. It is at the core of the fourth industrial revolution. Companies are increasingly investing and deploying artificial intelligence systems.

What is the effect of the application of technologies to tasks that were previously performed by humans? The answer is not univocal. In the big picture of creative destruction, some tasks are being replaced by robots or software, from contactless cashiers to welding drones, and some occupations are being created, such as social media managers or cloud specialists. The impact of technologies on occupations is not confined to labor displacement or labor reinstatement: some tasks are being transformed by the technological advances, e.g. nurses, exempted from recording patient history, will fully devote themselves to monitoring patients. Furthermore, new technologies increase the productivity of labor in the tasks they perform, which in turn boosts labor demand (Acemoglu & Restrepo, 2019). Thus, the impact of the application of technologies on the labor market is not obvious.

U.S. evidence suggests that advances in artificial intelligence have so far led to a small increase in wages, with no change in employment (Felten et al., 2019), associated with greater job stability (Fossen & Sorgner, 2019) and a shift in the composition of job postings away from occupations most exposed to AI to occupations least exposed (Acemoglu et al., 2020).

In this thesis, we aim at furthering our understanding of the impact of the advent of artificial intelligence on the Italian labor market. Past waves of automation enable us to use historical cases to discipline our modeling of AI. The nature of AI anticipates that it is likely to penetrate tasks that were previously thought to be out of reach from automation. Which are the occupations, and consequently the demographic groups, that are suffering more in the aftermath of AI advent, and whether the Italian labor market has assimilated the AI takeover are the insights offered in this thesis.

What are the capabilities of AI? Turing (1950)'s question "Can machines think?" provides intuition in the classification of technologies as artificial intelligence. The Imitation game, subsequently renamed Turing Test after its author, aims to assess whether a machine can exhibit intelligent behavior equivalent to that of humans. At the time Turing was writing, machine learning was in an embryonic phase of development, as computers still could only apply commands, not store them.

In 1956 the Dartmouth Summer Research Project on Artificial Intelligence was carried out "to proceed on the basis of the conjecture that every aspect of learning or any other feature of artificial intelligence can in principle be so precisely described that a machine can be made to simulate it". In the 1990s early advances in the "intelligent machines" became available, with computers playing chess and recognizing speech.

Over the past twenty years, breakthroughs in AI development have been made, calling for the need of defining, regulating, and analyzing the impact of this technology. The nature and the boundaries of artificial intelligence are continually expanding, exacerbating the difficulty in defining AI.

In November 2018, the OECD Expert Group on AI (AIGO) developed a description of the AI system, to outline the scope of applicability of policies and regulations. According to AIGO (OECD, 2022),
"An AI system is a machine-based system that is capable of influencing the Environment by making recommendations, predictions or decisions for a given set of Objectives. It does so by utilising machine and/or human-based inputs/data to:

1. *perceive real and/or virtual environments;*

2. *abstract such perceptions into models manually or automatically;*

3. *use Model Interpretations to formulate options for outcomes."*

To estimate the capabilities of AI and identify the aggregate effects of AI on the Italian labor market, we exploit the AI occupations' exposure proposed by Webb (2019). Getting rid of expert judgments about occupations' susceptibility to AI applications, Webb (2019) relies on patents as the source of information to understand which tasks AI can perform. His methodology contributes to the debate on the construction of an indicator of occupations' exposure. The AI Occupational Exposure (AIOE) by Felten et al. (2021), the Suitability for machine learning (SML) by Brynjolfsson et al. (2018), and the Fourth Industrial Revolution exposure by Meindl et al. (2021) will provide a benchmark to test against Webb (2019)'s measure.

Webb's technology exposure measures depict heterogeneous impact on the workforce, depending on the technology employed. Where the robotics and software result more pervasive in routine and repetitive tasks (Acemoglu & Autor, 2011) performed by young and low-skilled workers, AI's capabilities will mostly encroach on better-paid occupations, requiring emotional intelligence (Brynjolfsson & Mitchell, 2017).

In examining the changes in wage and employment structure at the occupation-industry level in the Italian market between 2014 and 2019, no significant relationship with the exposure to artificial intelligence was detected. Either labor reinstatement and labor displacement induced by technology advent counteracted each other, or the effect of the technology deployment has not been absorbed yet. A further possible explanation concerns the suitability of the exposure measure for our analysis: as a forward-looking metric, it accounts for all the *potential* areas of application of the so-far developed innovations, not only for the *effective* ones.

The thesis proceeds as follows. Section 2 describes the data we will employ throughout our analysis. In Section 3 we illustrate the methodology of the exposure measures. In Section 4 we use Webb (2019)'s exposure measure to link technologies to occupations, demographic groups, wages and employment. In Section 5 we describe the steps taken to test the robustness of Webb (2019)'s exposure measure and his methodology. Section 6 examines possible limitations to our analysis. Section 7 concludes.

2 Data

The analysis we are proposing requires data on the occupations' exposure to artificial intelligence and the labor market outcomes, such as wages and employment, at the occupational level.

2.1 Exposure Data

To quantify the susceptibility of occupations to see their work complemented, substituted, or transformed by the advent of technology, we leverage on four measures of exposure to artificial intelligence - Webb (2019) Exposure to AI, Felten et al. (2021) AI Occupational Exposure (AIOE), Brynjolfsson et al. (2018) Suitability for machine learning (SML), and Meindl et al. (2021) 4IR exposure score -, one measure of exposure to software - Webb (2019) Exposure to software -, and two measures of exposure to robots - Webb (2019) Exposure to robots, and Paolillo et al. (2022) measure of automation risk. We provide details about the metrics' construction in Section 3. The authors' rubrics are publicly accessible at the 6-digit O*NET SOC occupation level[1].

The exposure score proposed by Webb (2019) will be the reference indicator throughout our analysis. Webb (2019) provides occupational exposure scores for the three technologies of robots, software, and AI, which allows us to examine the differential occupations' exposure to the three technologies while keeping the measure construction methodology constant. Moreover, by exploiting the other measures of AI and robot exposure cited above, we can test against Webb (2019)'s methodology, by comparing first Webb (2019) exposure to AI to Felten et al. (2021) AI Occupational Exposure (AIOE), Brynjolfsson et al. (2018) Suitability for machine learning (SML), and Meindl et al. (2021) 4IR exposure score, and then Webb (2019) robot exposure to Paolillo et al. (2022) measure of automation risk.

2.2 Labour Data

To measure the relationship between the exposure metrics at hand and the labor market outcomes, we use ISTAT Labor Force Survey cross-sectional quarterly data for the 1977-1992 and 2008-2020 periods. ISTAT Labor Force Survey collects data on demographic characteristics, such as age, gender, educational attainment and residence, and labor market outcomes, such as labor status, type of work, work experiences, type of work contract, and salary. From 1977 to 2020, the survey has undergone substantial changes in the data collection technique and methodology. The 1984-1992 time window data are aggregated at the 2-digit ATECO industry level (12 classes). A higher level of disaggregation has been introduced for the 2008-2020 survey: from 2008 to 2013, individuals were asked about the economic activity they were performing among 99 classes, while, from 2014 to 2020, individuals reported their occupation

[1] While Felten et al. (2021) arrange the AIOE metric at the 6-digit O*NET SOC occupation level, Webb (2019), Brynjolfsson et al. (2018), and Meindl et al. (2021) produced an exposure measure at the 8-digit O*NET SOC level. Note that in converting exposure scores to the ISCO level, when using the SOCPC crosswalk an aggregation of the 8-digit O*NET SOC level occupations to the 6-digit O*NET SOC level was required. The aggregation has been accomplished by simple means.

according to the 3-digit ISCO classification (126 classes). This inconsistency represents a limit in our analysis because it prevents us from quantifying the change in labor market outcome variables for the entire period, 1977-2020, at the occupation level, but only at the industry level.

Online job advertisement data, that serve as a proxy for job demand, are gathered from Skills Online Vacancy Analysis Tool for Europe (Skills OVATE). From more than 100 million online job advertisements in 28 European countries, Cedefop has collected information on occupations, skills, competencies, and qualifications required. In particular, we employ data on the 2021 Italian online job advertisements at the 3-digit ISCO level.

We relied on the U.S. Bureau of Labor Statistics database to collect employment projections data. The Employment Projections (EP) program develops information about the labor market for ten years in the future, providing 2020 employment, 2030 projected employment, and 2020-2030 occupational openings statistics at the 6-digit SOC level for the U.S. labor market.

2.3 SOC - ISCO Crosswalk

Given the different occupations taxonomy used for the exposure measures and the labor market data, a conversion system between the 6-digit O*NET SOC codes and the 3-digit ISCO08 codes was needed. We employed the crosswalk prepared by Wojciech Hardy from the University of Warsaw and provided by the Institute for Structural Research (IBS). As a robustness check, we also adopted the crosswalk designed by the Bureau of Labor Statistics on behalf of the Standard Occupational Classification Policy Committee (SOCPC).

The conversion between the two occupations' classifications reduced the variability that we could exploit: out of the 964 6-digit O*NET SOC codes for which Webb (2019) provides exposure estimates, after the conversion, we ended up with 126 3-digit ISCO codes.
First of all, some variation is lost when converting a 6-digit SOC code to a 4-digit ISCO code. Secondly, the 4-digit ISCO classification has been mapped to the 3-digit ISCO classification, to ensure correspondence with the ISTAT Labor Force Survey data. This step leads to further inaccuracy in the resulting exposure measure. See Appendix A for an explanatory example.

3 AI Exposure Measures

Investigating the impact of artificial intelligence on the labor market requires identifying the capabilities of AI and discerning which tasks AI technologies cover.

In the present analysis, we will employ four measures of exposure to artificial intelligence. Table 1 provides an overview.

3.1 AI Exposure score by Webb (2019)

Webb (2019) develops an innovative method to measure the extent to which occupations are exposed to technologies. The underlying idea is to quantitatively assess to what extent AI-based technologies and their functionalities can replace human abilities required in jobs. Operationally, the core of his strategy is to quantify the overlap between the tasks performed in each occupation and the tasks performed by technologies nowadays.

Figure 1: Illustration of process for constructing technology exposure measures, by Webb (2019)

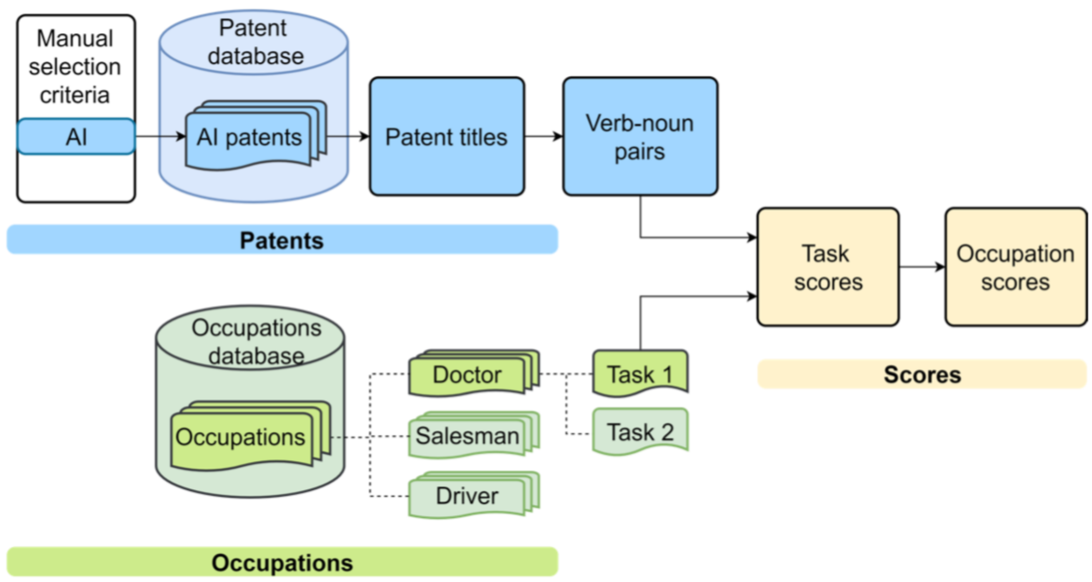

Figure 1 illustrates the process of construction of the exposure measure. On the patents' side, he extracts from the Google Patents Public Data, provided by IFI CLAIMS Patent Services, the set of patents that fall under a certain technology category: for example, for artificial intelligence he selects the patents that in their title use the keywords "supervised learning", "reinforcement learning", "deep learning", and "neural network". From the list of titles of all the patents collected, he extracts all the verb-noun pairs. On the occupations' side, from the O*NET Database of occupations, he extracts all the verb-noun pairs of every task performed in each occupation. Operationally, extracting verb-noun pairs was carried out by a dependency parsing algorithm.

Once all the verb-noun pairs from the patent titles and the occupations' tasks have been collected, for each occupation he considers which extracted pairs appear in the set of pairs extracted from the patents. The frequency of the matched verb-noun pairs in the technology patents, together with the importance of the tasks in the occupation, are employed in the exposure measure. The exposure of occupation i to technology t is thus:

$$Exposure_{i,t} = \frac{\sum_{k \in K_i} [w_{k,i} \cdot \sum_{c \in S_k} rf_c^t]}{\sum_{k \in K_i} [w_{k,i} \cdot |\{c : c \in S_k\}|]}$$

where K_i is the set of tasks in occupation i, S_k the set of verb-noun pairs extracted from task $k \in K_i$, $w_{k,i}$ the weight of task k in occupation i, and $rf_c^t = \frac{f_c^t}{\sum_{c \in C^t} f_c^t}$ the relative frequency of aggregated[2] verb-noun pair c in technology t patent title.

Thus "an occupation's exposure score for technology t expresses the intensity of patenting activity in technology t directed towards the tasks in that occupation". The three technologies for which Webb (2019) constructs this metric are robots, software, and artificial intelligence.

Webb (2019)'s measures are employed for the first empirical evidence provided in this research.

3.2 AI Occupational Exposure by Felten et al. (2021)

An alternative measure has been created by Felten et al. (2021): the AI Occupational Exposure (AIOE). Instead of exploiting the overlapping between AI-performed and occupations-performed tasks, the authors consider the relatedness of technology applications to the occupational abilities.

AI applications information is collected from the Electronic Frontier Foundation (EFF) AI Progress Measurement Project, which provides statistics about the progress of AI across different applications. The authors first select the "10 AI applications[3] for which the EFF has recorded measured scientific activity and progress in the technology from 2010 onward", and which, according to experts, cover the most common uses of AI and are the most likely to impact the labor market. 200 Amazon's Mechanical Turkers are then surveyed to get their opinion on whether each of the ten AI applications considered from EFF "is related to or could be used for each" of the 52 abilities in the O*NET database. The survey responses provide an estimate of how related the AI applications are to each ability. Having obtained a matrix that connects the ten EFF AI applications to the 52 O*NET occupational abilities, the ability-level exposure is constructed as the sum of the 10 application-ability relatedness scores:

$$A_j = \sum_{i=10}^{10} x_{ij}$$

where i is the AI application and j the occupational ability. Then, the Artificial Intelligence Occupational Exposure (AIOE) for occupation k is constructed as:

$$AIOE_k = \frac{\sum_{j=1}^{52} A_{ij} \cdot L_{jk} \cdot I_{jk}}{\sum_{j=1}^{52} L_{jk} \cdot I_{jk}}$$

where L_{jk} is the ability's prevalence and I_{jk} is the ability's importance. The occupation level metric is thus obtained by multiplying the ability level measure by the prevalence and importance scores O*NET assigns to each ability within each occupation. The scaling by

[2] Verb-noun pairs are aggregated using WordNet, a database that groups nouns into a hierarchy of concepts.

[3] The 10 selected applications are: abstract strategy games, real-time video games, image recognition, visual question answering, image recognition, reading comprehension, language modeling, translation, speech recognition, and instrumental track recognition.

the sum of the prevalence and importance of all the abilities used in an occupation allows for accounting for the number of abilities' requirements by occupations (i.e. to avoid that occupations are assigned higher exposure scores just because they involve more abilities).

Their "methodology produces a score that measures how the most prevalent applications of AI are related to occupations at the six-digit SOC level based on their ability composition as defined by O*NET."

3.3 Suitability for Machine Learning by Brynjolfsson et al. (2018)

Brynjolfsson et al. (2018) construct a measure of "suitability for machine learning" (SML), based on the technical feasibility of the O*NET detailed work activities (DWA). In particular, experts in the area are asked 23 questions to evaluate each DWA based on eight criteria of technical feasibility[4]. DWA level scores are then aggregated at the task level and the occupation level, weighting by O*NET tasks' importance in each occupation.

Addressed questions, e.g. "task information is *recorded or recordable* by computer" are formulated in such a way to grasp not only current but also potential applications of ML. The nature of the questions used to assess the work activity suitability for automation implies that Brynjolfsson et al. (2018) metric is indicative of the "general automation potential once technologies are fully implemented" (Meindl et al., 2021).

3.4 Fourth Industrial Revolution by Meindl et al. (2021)

A further empirical indicator of technological progress has been developed by Meindl et al. (2021). It is similar in nature to the one constructed by Webb: both are based on a mapping between patents and occupations at the task level. Instead of focusing on one technology at time, Meindl et al. (2021) consider all the patents that shape the fourth industrial revolution, and which include, e.g., AI, machine vision, and autonomous robots. Their approach is as follows:

1. from the PATSTAT patent database, they extract patent abstracts for one patent per family, preferring in order USPTO, WIPO and EPO patents in order;

2. from O*NET database they extract task descriptions;

3. a preprocessing step is carried out to clean patents and tasks' text from stop-words, non-alphabetical characters, etc., identify noun-chunks, tokenize text, weight tokens based on their inverse document frequency;

4. transform text into word vectors (using fastText word embedding or Risch and Krestel (2019)'s word embedding) and compare word vectors with task descriptions' word vectors;

[4]The eight criteria considered are: "Learning a function that maps well-defined inputs to well-defined outputs, large (digital) data sets exist or can be created containing input-output pairs, the task provides clear feedback with clearly definable goals and metrics, no long chains of logic or reasoning that depend on diverse background knowledge or common sense, no need for detailed explanation of how the decision was made a tolerance for error and no need for provably correct or optimal solutions, the phenomenon or function being learned should not change rapidly over time, no specialized dexterity, physical skills, or mobility required".

5. select the patents that exceed a certain similarity threshold and those that fall within one of the 4IR classifications defined by Ménière et al. (2019);
6. 4IR exposure score is obtained by calculating the mean of log 4IR patents per task. For the aggregation by occupation, they first weight each task by the importance score of the tasks in the occupation.

Their methodology delivers a "patent-based measure of exposure to 4IR technologies" at 8-digit O*NET SOC level.

Table 1: Overview of AI exposure indicators

	AI Exposure	**AIOE**	**SML**	**4IR**
Authors	Webb (2020)	Felten et al. (2021)	Brynjolfsson et al. (2018)	Meindl et al. (2021)
Description	Exposure to AI patents	AI applications relatedness	Suitability for machine learning	Exposure to technologies of the fourth industrial revolution
Level of Analysis	Occupation-task (O*NET)	Occupation-ability (O*NET)	Occupation-work activities (O*NET)	Occupation-task (O*NET)
Index basis	Relevant patents per task	Scientific activity and progress, as described by EFF, per ability	Survey of SML characteristics per task	Relevant patents per task
Method	Dependency parsing algorithm to measure overlapping between patents and occupations' description	mTurkers assessing application-ability relatedness	Experts assessing technical feasibility of DWA	Word embedding and text similarity
Measure Interpretation	Existing AI capabilities based on patents, applied or not	Existing technology AI progress based on scientific articles	Potential future capabilities of ML	Existing 4IR capabilities based on patents, applied or not

4 Results

Previous-generation automation has affected routine manual and routine cognitive tasks. Applications were limited to areas where, once instructions are specified in advance for each contingency, human knowledge or expertise is not required. By contrast, the nature of artificial intelligence opens a new set of possibilities for automation. We expect the labor market impact of robots and software to differ from that induced by artificial intelligence.

Section 4.1 provides preliminary evidence on the heterogeneity of exposure to robots, software, and artificial intelligence, across occupations and demographic groups in the Italian market. Section 4.2 proceeds with an analysis of the effects of the exposure on the Italian labor market.

4.1 Preliminary evidence on exposure

Given that robots' diffusion and adoption have preceded the AI revolution, it could represent a benchmark of likely results to expect from AI. However, the likely labor market outcomes of technologies depend on the occupations' exposure: if occupations that adopt robots to perform manual routine tasks do not embrace AI-based systems, we expect the effects of AI to diverge from the effects of robots. Therefore the preliminary evidence we are providing in this section involves a comparison of differential exposure of occupations, industries, and demographic groups to the three technologies of robots, software, and artificial intelligence. The exposure measure is drawn from Webb (2019). As a consequence, we are dealing with *potential* exposure of occupations to technology, predicting areas in which adoption may occur, silent about the actual technology employment or the direction of the impact.

4.1.1 Most and least exposed occupations

Table 2 displays the five occupations most and least exposed to robots, according to Webb's metric[5]. Jobs that involve routine tasks, and in particular manual routine tasks, such as mobile plant operators or cleaners, are among the most exposed to robots. Vice-versa, occupations that require social skills and creativity, such as artists or clergy, are said to be safe from automation.

Similarly, Table 3 reports the five occupations with the highest and lowest Webb's index of software exposure. The most exposed occupations include manufacturing laborers, telecommunications and broadcasting technicians, or locomotive engine drivers, jobs that require processing information according to pre-specified rules, without substantial case-by-case human judgment. For example, in the last years, content companies have increasingly started using software for content programming and traffic management. Conversely, the least exposed occupations, such as university teachers, journalists, and hairdressers, involve either manual tasks or interpersonal skills.

[5]The statistics in the present section are derived from the matching between the 2019 ISTAT Labor Force Survey and Webb (2019) occupation level exposure indicators, after occupation codes' conversion from SOC to ISCO.

Table 2: Occupations with the highest and lowest exposure to Robots

Most exposed occupations	Least exposed occupations
Mobile plant operators	Sales workers
Domestic, hotel, office cleaners and helpers	Creative artists
Manufacturing laborers	Government regulatory associate professionals
Food preparation assistants	Clergy
Elementary workers	University and higher education teachers

Table 3: Occupations with the highest and lowest exposure to Software

Most exposed occupations	Least exposed occupations
Manufacturing labourers in Mining, Costruction, Transport	University and Higher Education Teachers
Telecommunications and Broadcasting Technicians	Information and Communications Technology Services Managers
Mixed Crop and Animal Producers	Authors, Journalists and Linguists
Mobile Plant Operators	Hairdressers, Beauticians
Locomotive Engine Drivers	Food Preparation Assistants

Boundaries between software and artificial intelligence are not always clear to the audience. Conceptually, the main difference is that every action performed by software has been programmed in advance by human programmers anticipating any possible contingency the software could face; AI, instead, is programmed to *learn* how to perform tasks. Operationally, patents selected in the software index construction had keywords ”software”, ”computer”, or ”program”. The search terms employed in the construction of the AI exposure score were ”supervised learning”, ”reinforcement learning”, ”neural network”, and ”deep learning”.

Given AI performs different tasks compared to robots and software, we expect exposed occupations to differ substantially. Table 4 reports the five occupations most and least exposed to artificial intelligence.

Table 4: Occupations with the highest and lowest exposure to Artificial Intelligence

Most exposed occupations	Least exposed occupations
Physical and engineering science team	University and higher education teachers
Process control technicians	Cooks, waiters and bartenders
Legal professionals	Child care workers
Skilled agricultural workers	Street vendors
Optometrists	Creative artists

The occupations most exposed to AI according to Webb’s measure are indeed those that in recent years have experienced wide application of AI-based technologies, such as physical and engineering science teams, process control technicians, and legal professionals. In the field of chemical engineering, AI-based systems have been employed in modeling, process control, classification, fault detection, and diagnosis and have proven to be superior to human performance. Moreover, process control technicians use AI to monitor equipment, detect failures, determine when to replace or repair it, and optimize energy consumption and storage. Legal professionals are already embracing AI: legal research, document management, due diligence, and litigation analysis are examples of fields in which AI has been employed. The greatest success of AI in the area of medicine is optometrists, moving toward the development of diagnostic recommendations. Even skilled agricultural workers have seen their work complemented or substituted in the tasks of climate modeling and weather forecasting, farm finance, and administration, "intelligent" irrigation systems, or recognition of plants and animals' disease (Sparrow et al., 2021).

As for robots and software, subjective inputs and interpersonal skills secure occupations from technology substitution, as well as reasoning about situations never seen before, even in the case of manual work, e.g. for waiters and bartenders. Note that AI exposure so far does not seem to be a skill-biased technological change: the least exposed ones are both high-skill workers, such as university teachers, and medium or low-skill ones, such as sales workers.

The above comparison of occupations' exposure to robots, software, and artificial intelligence is aimed at highlighting that the different exposure is going to determine distinct effects on the labor market.

This occupations' ranking has been made by exploiting Webb's measure of robots, software, and AI exposure. However, the results only roughly coincide with the ranking delivered by Webb (2019). The reason for this inconsistency lies in the matching between the 6-digit O*NET SOC codes and 3-digit ISCO-08 codes: there is not a one-to-one correspondence between the two classifications and occupations often end up being aggregated as a result of the conversion[6].

4.1.2 Exposure heterogeneity across industries

Occupations' exposure and industries' composition determine industries' susceptibility to the employment and deployment of technologies.

Figure 2 represents the standardized exposure score[7] of each industry, according to the ATECO classification, to the three technologies under exam, i.e. robots, software, and artifi-

[6]For the sake of clarity consider an example: Chief Executives (O*NET SOC code 11-1011.00, AI exposure score of 0.20) and Emergency Management Directors (O*NET SOC code 11-9161.00, AI exposure score of 0.75) are both associated by the SOC-ISCO crosswalk to Legislator category (ISCO08 code 111). Consequently, our measure of exposure for ISCO08 code 111 will be the mean of the exposure of the different SOC categories of occupations associated with the ISCO code 111 by the crosswalk. Legislator occupation (ISCO08 code 111) is indeed assigned AI exposure of 0.37, missing the information about the greater exposure of Emergency Management Directors compared to Chief Executives. See Section 2.3 for further details on the codes' conversion.

[7]Standardization has been applied to the 2019 ISTAT Labor Force Survey data, after merging them with exposure scores.

cial intelligence. As expected from Section 4.1.1, Hotels and Restaurants and Education and Health sectors are among those that will experience the lowest adoption of AI-based systems.

Figure 2: Standardized exposure score, by technology and industry

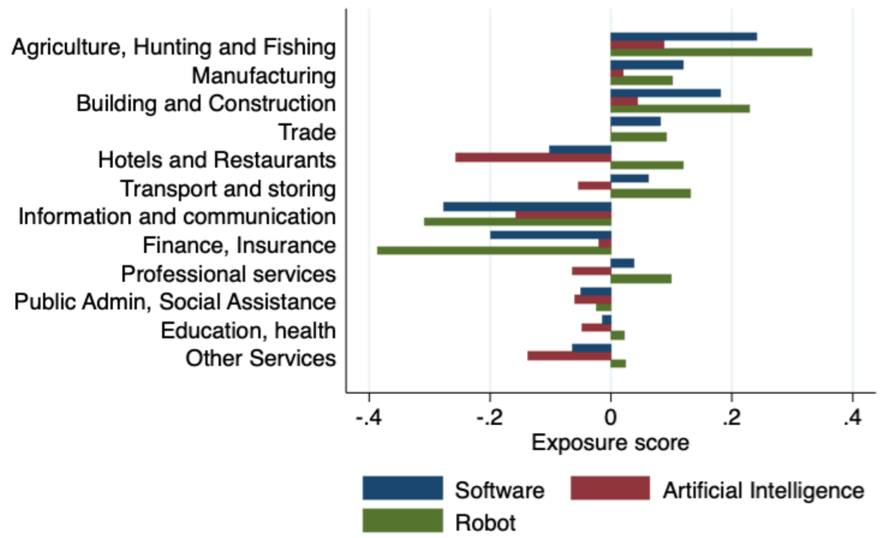

Figure 3 represents the standardized exposure score of broad occupations, according to the 1-digit ISCO classification, to the three technologies. Service and Sales workers is the category least exposed to AI and software, while elementary occupations are the most exposed to robots.

Figure 3: Standardized exposure score, by technology and broad occupation

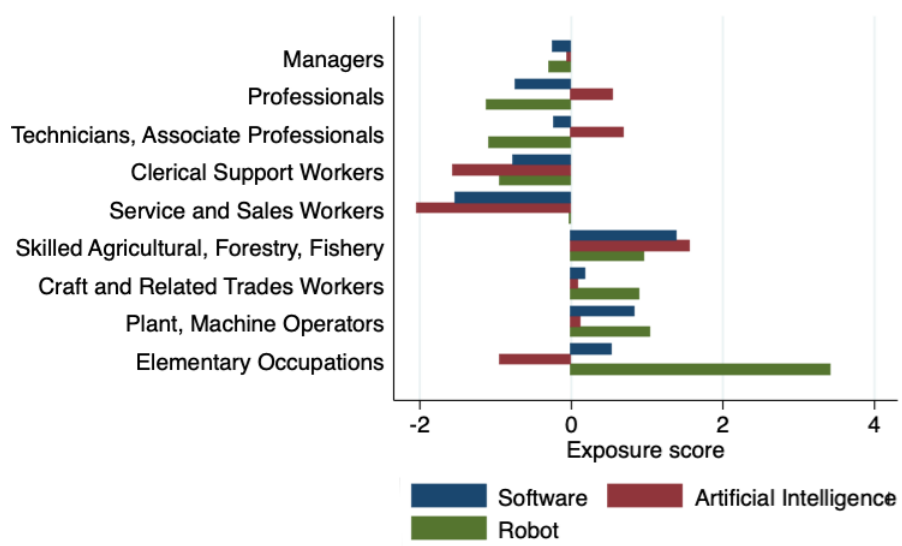

It is noteworthy that Agriculture, Hunting, and Fishing is the industry that, according to Webb's measures of exposure and the ISTAT composition of the industry workforce, has the highest propensity to see their tasks complemented or substituted by all the three types of technologies. Looking at the exposure scores by broad occupation helps explain such high exposure. Workers employed in Agriculture, Hunting, and Fishing industry are either skilled agricultural, forestry, and fishery workers, such as field crop and vegetable growers or apiarists and sericulturists, or laborers performing elementary tasks, such as crop farm laborers or livestock farm laborers. Thereby the high exposure of the sector to artificial intelligence and robots arises respectively from the high exposure of skilled and elementary workers.

4.1.3 Exposure heterogeneity by age

Figure 4 shows how the average standardized exposure score varies with age. The occupational mix of individual and occupation exposure determines the exposure life-cycle. For each age, the average exposure index, from the 2019 ISTAT Labor Force Survey[8], is plotted for each technology, smoothed using a locally weighted smoothing regression.

Figure 4: Standardized exposure score, by technology and age

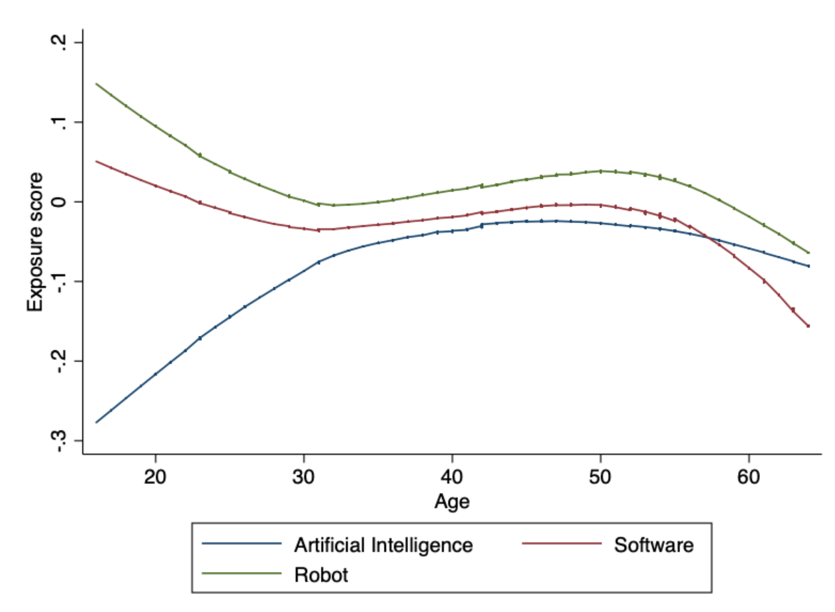

The propensity to work in occupations exposed to robots' applications starts declining from the start of working life, until the age of 35. It probably reflects the prevalence of routine tasks in the very first years of workers' careers. The decreasing path between the ages of 18 and 25 could also be explained by the composition of the labor force. The average exposure score for eighteen-year-old individuals is made up of exposure scores of the occupations in which eighteen-year-old individuals are working: if working and not undertaking a degree at the age of eighteen, it is likely they commit to jobs with lower skill-requirement.

For artificial intelligence the opposite is true: individuals in their forties and fifties are the most exposed. This trend could mirror the higher skill or degree requirements in the occupations that happen to be the most exposed to AI. However, as seen in Section 4.1.1, among the most exposed, also low or medium skill level occupations appeared. There is no necessary inconsistency between the two pieces of evidence, as in Section 4.1.1 only five occupations are reported. How exposure varies by educational attainment and wage, proxies for skills, will be investigated further in Section 6.

Figure 5 replicates Figure 4, additionally splitting by gender.

[8]The locally weighted regression of standardized exposure scores on age has been carried out after having excluded individuals aged above 65. It holds for all locally weighted regressions provided in Section 4.1.3.

Figure 5: Standardized exposure score, by technology, sex and age

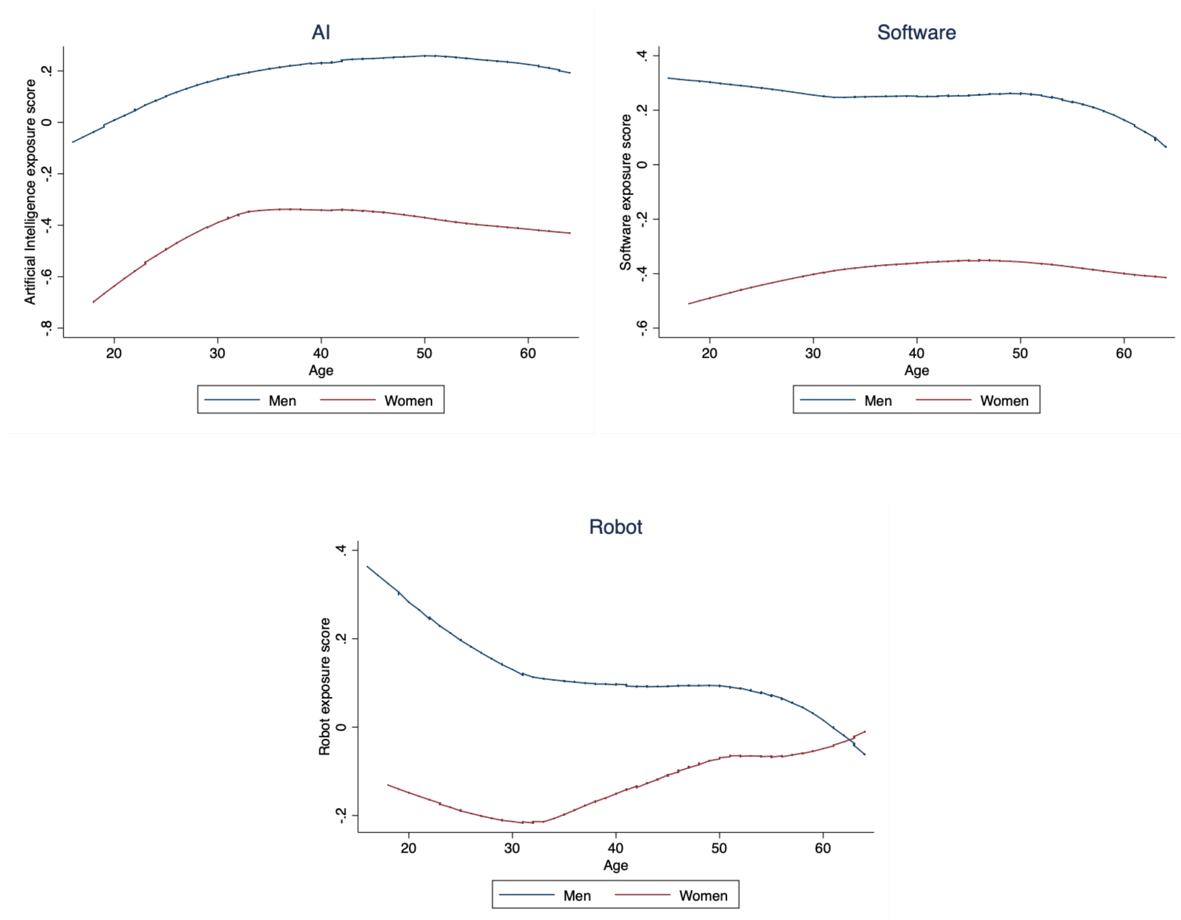

What emerges from Figure 5 is that for each technology, men over their entire lives tend to work in occupations with above-average exposure scores, while women in occupations with below-average exposure scores. It is consistent with female propensity to work in occupations requiring interpersonal skills and male propensity to work in occupations requiring technical skills (Webb, 2019). Different gender occupational placement is a leading factor of the differential exposure to technologies. According to ISTAT data, in 2019 74% of workers employed in the Agriculture, Hunting, and Fishing sector, the most exposed industry (see Section 4.1.2), were men. Occupational segregation places men and women at different risks of seeing their work substituted or complemented by technologies.

4.1.4 Exposure heterogeneity by skills

Whether the artificial intelligence revolution will exacerbate labor market polarization has been under debate in the last years. There is little evidence so far. Felten et al. (2021) document a positive relationship between his measure of AI exposure with wage and employment growth for high-income occupations, while no significant relationship for low and middle-income ones. Investigating the heterogeneity in exposure by skill level anticipates whether AI will have a skill-biased impact on the labor market, whatever the direction of the impact. Educational attainment and wage (Autor & Dorn, 2013) act as proxies for skills. The following paragraphs will provide evidence of the heterogeneity of exposure across educational and income groups.

Figure 6 depicts the average exposure score by education group, according to the ISCED classification[9], for each technology.

Figure 6: Standardized exposure score, by technology and education level

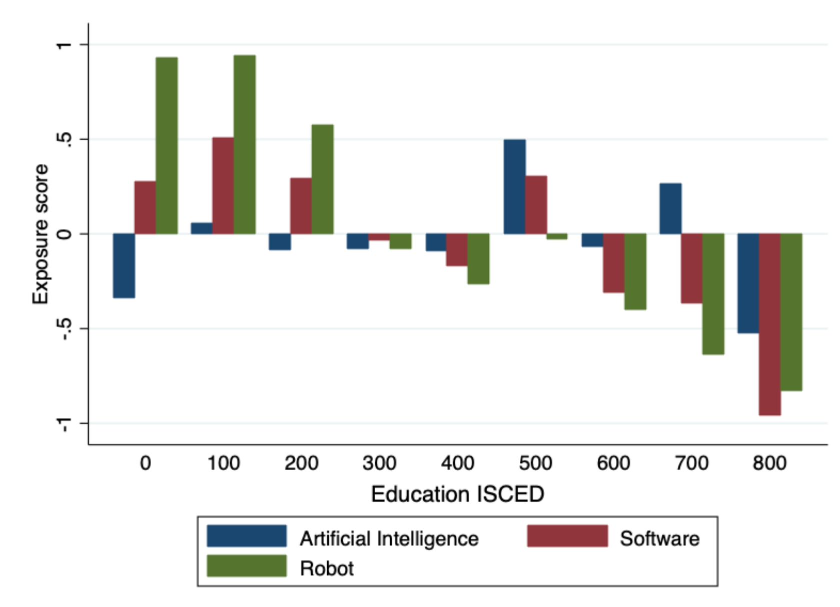

Exposure to software and robots is declining on average in the educational attainment: occupations performed by workers with at most primary or lower secondary education tend to be more exposed than those carried out by individuals with a degree. Software and robotics capabilities are more common in low-education occupations. It reflects the prevalence of manual tasks and low skill requirements in high-robot-exposure occupations.

Artificial intelligence exposure instead shows a flat pattern in education level: no significant difference in average exposure for an individual with secondary school or a bachelor's degree. As anticipated by the evidence about the most and least exposed occupations in Section 4.1.1, AI-based technologies' application seems to be unbiased in skills.

Figure 7 plots the average exposure score for each technology by occupational wage percentile. Peaks in exposure to robots, software, and artificial intelligence are reached respectively in low-income, middle-income, and high-income occupations. It reflects the fact that manual work, highly substitutable by robots, tends to be lower-paid, and tasks that involve reasoning, accomplishable by AI, tend to be higher-paid. If wage proxies for skills, Figures 6 and 7 provide conflicting evidence about skill-biased exposure.

4.1.5 Exposure heterogeneity by Italian region

The identification of the demographic groups more likely to be affected by the aftermath of AI possibly informs policy design. Policymakers may accordingly outline unemployment benefits or retraining programs for affected workers. Whether technology advent will exacerbate the Italian North-South divide is an issue to be addressed.

[9] In classifying the education levels, ISCED levels correspond to: Level 0 early childhood education, Level 1 primary education, Level 2 lower secondary education, Level 3 upper secondary education, Level 4 post-secondary education, Level 5 short-cycle tertiary education, Level 6 bachelor or equivalent level, Level 7 master's or equivalent level, Level 8 doctoral or equivalent level. Note that the ISCED level individuals are assigned in the ISTAT Labor Force Surve individuals are assigned the highest education level achieved.

Figure 7: Standardized exposure score, by technology and wage quantiles

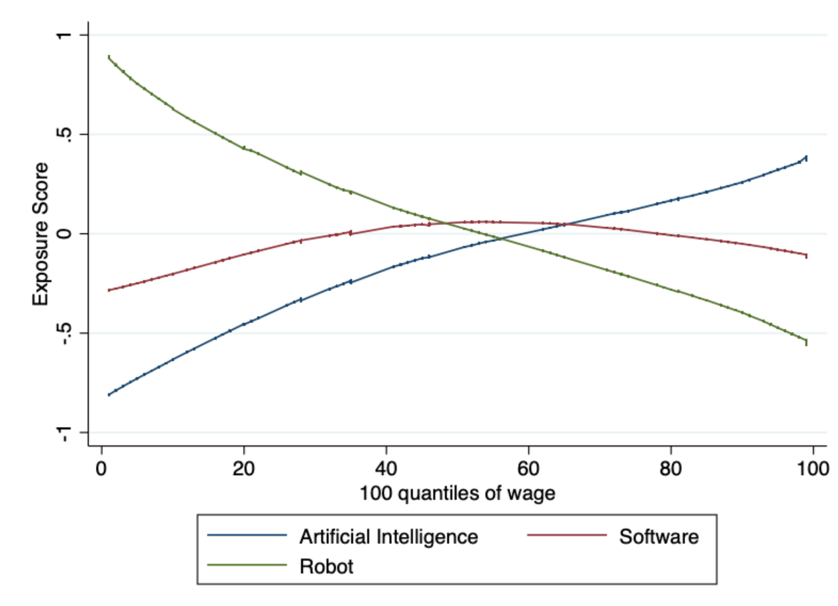

Figures 8, 9, and 10 plot the average standardized exposure score for each technology and Italian region.

Figure 8: AI standardized exposure score by region

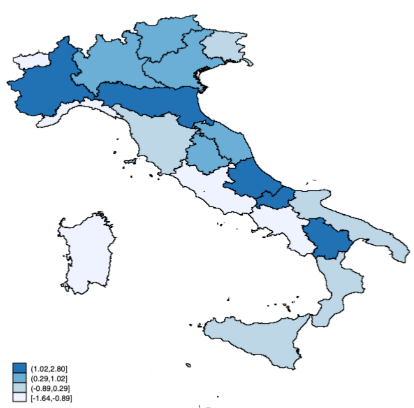

Figure 9: Software standardized exposure score by region

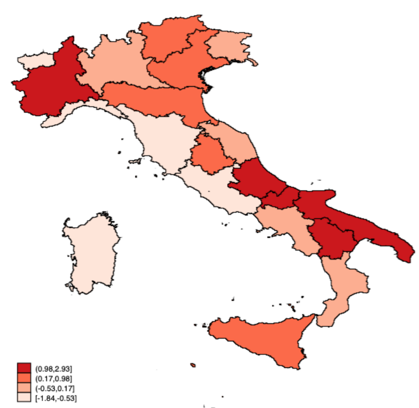

Figure 10: Robot standardized exposure score by region

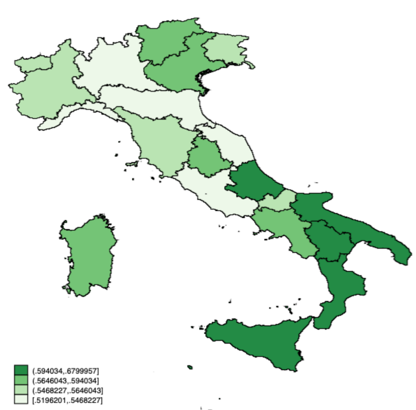

The greater exposure to artificial intelligence in the North and the greater exposure to robots and software in the South can be traced back to the differential wage and educational structure. The higher educational achievements (Argentin & Triventi, 2015) and work remuneration (Daniele & Malanima, 2017) documented for individuals in the northern area, together with evidence about high-skill workers in Section 4.1.4, can explain the pattern in the Figures above.

The Italian North-South divide places workers at different risks of job displacement or reinstatement.

4.2 Labor market outcomes

Section 4.1 investigates the *potential* exposure of occupations to artificial intelligence. Webb (2019) measure of exposure to technologies is *forward-looking*: it takes into account all the possible applications of AI, not only areas in which adoption has already occurred, but also areas in which it might occur in the future, given the nature of the patent. Moreover, Webb's measure is agnostic about whether exposure takes the form of labor displacement, labor reinstatement, or within-occupation change in performed tasks.

So far our analysis does not comment on the direction and magnitude of the impact. To study empirically the relationship between exposure and changes in employment and wages, we exploit ISTAT Labor Force Survey data to estimate the following regression:

$$\Delta y_{o,i,t} = \alpha_i + \beta Exp_o + \gamma Z_{i,o} + \epsilon_{o,i,t}$$

where $\Delta y_{o,i,t}$ is the difference in labor market outcome from 2014 to 2019[10] for occupation o in industry i, α_i the industry fixed effects, Exp_o the technology exposure measure at occupation level, $Z_{i,o}$ the occupation-industry level controls. The labor market outcomes considered are wage and employment: the change in wage is measured as log change in weekly wage, while the change in employment as change of an occupation-industry share of total employment. Controls include the percentage of workers in an occupation-industry cell having achieved different levels of education, the percentage of men in an occupation-industry cell, the occupation-industry average number of years of workers' tenure, the occupation-industry average workers' age, and the percentage of self-employed workers in occupation-industry cell. As shown in Section 4.1.2, there is substantial heterogeneity in exposure by industry. To avoid results being driven by industry constant unobserved factors, we include industry fixed effects.

Identification assumption $E[\epsilon|Exp, Z] = 0$ may not hold: progress in the development and the application of AI is not exogenous and may be influenced by characteristics that impact also wages and employment growth. Consider for example changes in the labor supply side: if low-skill workers are employed in occupations highly exposed to AI, and an educational upgrade reduces the relative supply of low-skill workers compared to high-skill ones, this decline will be attributed to reduced demand for low-skill labor induced by substitution by

[10]The time window considered, from 2014 to 2019, is relatively short. From ISTAT we have access to Labor Force Survey for the periods 1977-1992 and 2008-2020. Studying the change in wage and employment between 1990 to 2020 would have been preferable, considering the timing of the adoption of artificial intelligence technologies. However, the structure of the survey has changed from 1977 to 2020, and individuals' occupation information at the ISCO level of aggregation is available only since 2014 (see Section 2.2 for further details). Hence, the choice of focusing on 2014-2019 changes. Expanding the time window would have come at the expense of losing variation in exposure measure: industry level exposure delivers an imprecise and insufficient variability. The year 2020 has been excluded to avoid results being driven by COVID-19 shocks.

AI-based systems. Or even if task offshorability is positively correlated with exposure, a decrease in employment will be mistakenly attributed to AI.

Figure 11 is the residual binscatter of the relationship between employment-weighted vingtiles of AI exposure, measured via Webb's scores, and the change in wage and employment between 2014 and 2019, at the occupation-industry level, after controlling for industry-level, occupational and demographic variables. Figure 17 presents the corresponding regression results.

Figure 11: Change in wages and employment by exposure to AI, 2014-2019

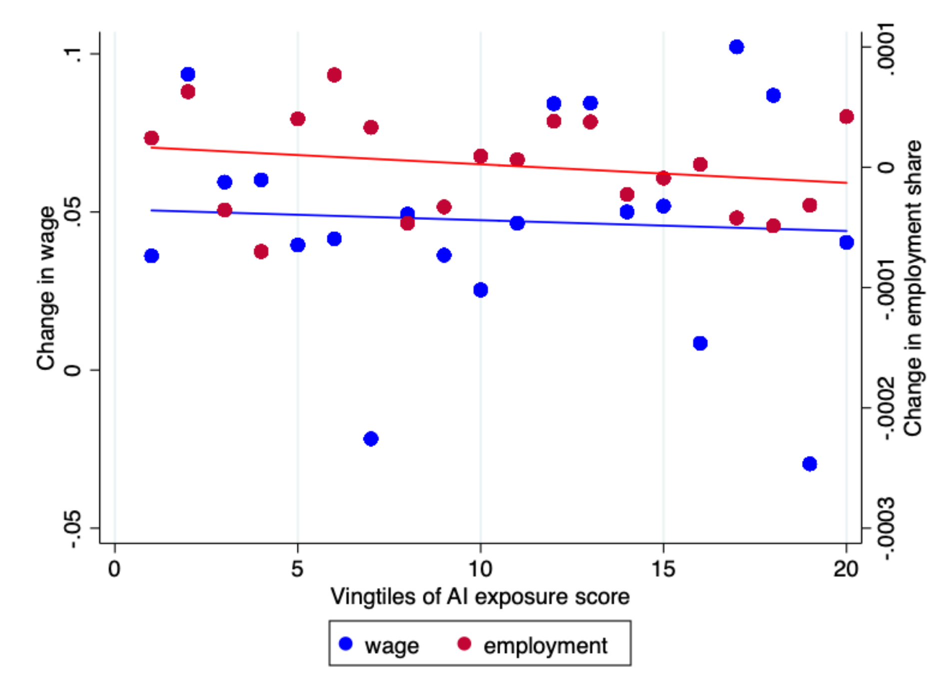

Note: Plot is a binscatter. Change in wage is measured as log difference in occupation-industry weekly wage. Change in employment is measured as change of an occupation-industry share of employment. Changes are winsorized at the top and bottom 1%. Controls are added at the occupation-industry level. Standard errors are clustered at the industry level.

The results are very small and statistically insignificant: in Italy between 2014 and 2019 occupations with high exposure to artificial intelligence, as estimated by Webb (2019), have experienced no additional change in wage and employment, compared to occupations characterized by low levels of exposure. Felten et al. (2021) provided evidence of positive but not statistically significant relationship between employment growth and his measure of AI Occupational Exposure in the U.S. market between 2010 and 2015, while small positive and statistically significant relationship with wage growth. The different impact of AI on the Italian and U.S. market may be the result of the greater rigidity of the Italian labor market.

Although we cannot attribute causality to our results, one possibility is that the advent of AI has not altered yet the labor market. As the diffusion of AI-based systems began to be felt in the second half of the 2000s, it could be that employers' adjustments, following AI deployment, have not occurred yet. Another possibility, that our data do not allow to control for, is that occupation-industry level reinstatement and displacement effects counterbalance, generating negligible effect overall. A further possible explanation targets our exposure measure. The occupational exposure score constructed by Webb (2019), by measuring the overlapping between tasks performed by patents and humans, takes into consideration all the possible,

Figure 12: Change in wage and employment vs. exposure to artificial intelligence

	Dependent variable:			
	Change in wage		Change in employment	
	(1)	(2)	(3)	(4)
Webb (2019) AI exposure	-0.000842	-0.00115	-1.32e-05	-1.33e-05
	(0.0155)	(0.0156)	(1.50e-05)	(1.47e-05)
Observations	834	834	881	881
Demographic controls	Y	Y	Y	Y
Occupational controls	Y	Y	Y	Y
Industry FE		Y		Y

Note: Each observation is an occupation-industry cell. Dependent variables are change in log wage and change of a cell's share of overall employment between 2014 and 2019, winsorized at the top and bottom 1%. Controls are added at the occupation-industry level. Standard errors are clustered by industry. *p<0.1; **p<0.05; ***p<0.01

actual or future, conceivable or not, applications of the existing patents. It provides an estimate of the occupations' exposure once the potential of the so-far developed patents has been fully realized. As such, if, until 2019, there are still areas of work that could deploy AI, our measure is going to display little relationship with the labor market outcomes.

4.2.1 Labor demand

As anticipated, artificial intelligence is in the early stages of its development and our measure of exposure to technologies is a forward-looking index. However so far we relied on historical data, in which the effects of AI deployment have hardly been realized. We speculate however that consequences of technology application materializes first in job openings, rather than in wage adjustments or firings.

Online job advertisement, as a proxy for labor demand, possibly provides us a useful insight on the impact of advances in technology adoption. Figure 13 shows how the average number of online job advertisements in Italy in 2019, collected from Skills Online Vacancy Analysis Tool for Europe (Skills OVATE), varies by exposure score for each technology. We take the average number of advertisements across all occupations that share the same standardized exposure score, and plot these averages, smoothed using a locally weighted smoothing regression. Thus it represents how the occupational mix of job openings changes by exposure.

There seems to be no clear trend in labor demand, if related to exposure. Technology-induced change in employment, even at the occupation-industry level, is not going to be monotone. Individuals are seeing their work substituted, complemented or even transformed by the advent of technologies. Exposure score is not able to predict the direction of the impact on labor demand. Artificial intelligence exposure, which shows a sharper relationship

Figure 13: Online job advertisement, in thousands, by standardized exposure score and by technology

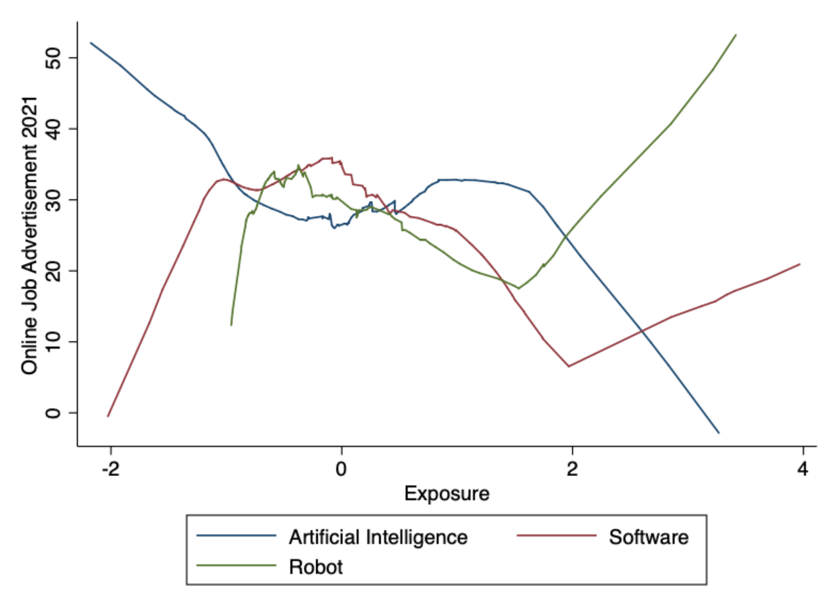

with online job advertisements, compared to robots and software, could be indicative of tasks substitution patterns induced by AI applications. If occupations involve tasks that have the potential to be automated by AI-based patents, the number of job openings declines.

However, the present assessment of technology-induced impact on job advertisements has two limitations for such a conclusion to be drawn. First of all, it covers only the year 2021, which does not provide enough time coverage. For this reason, we performed the same exercise using U.S. Bureau of Labor Statistics (BLS) employment projections data at occupation-level in the United States between 2020 and 2030. BLS allows to expand the time span considered and examine medium-term employment outcomes, at the expenses of relying on U.S. market data, not Italian one. Under the assumption that employment trend in the United States for the 2020 decade is similar to the Italian one, and given that Figure 14 does not depart markedly from Figure 13, we can conclude that year 2021 is representative of the decade. Moreover, the second limitation stems from Cedefop data which collects statistics about existing occupations, and it does not speak to whether *new* types of job will be created.

Figure 14: Job openings in 2020-2030, in thousands, by standardized exposure score and by technology

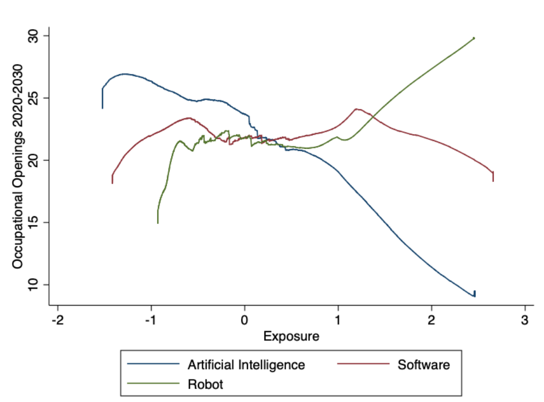

Figure 15: Percent employment change 2020-2030, by standardized exposure score and technology

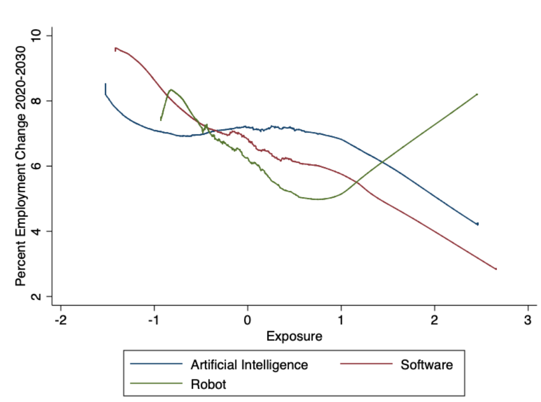

Figures 13 and 14 make use of the absolute number of online job advertisement or occupational openings, notwithstanding the occupational group size. If exposure correlates negatively with number of employees for reasons other than task substitution, our analysis would mistakenly attribute the negative trend of job openings in AI exposure to the technology impact. Therefore, Figure 15, plotting the percent employment change, instead of the absolute one, should isolate the present assessment from such an issue. Again, as data is drawn from U.S. Bureau of Labor Statistics (BLS) employment projections, they provide greater evidence, at the expenses of country coverage.

Percent employment change from 2020 to 2030 in the U.S. exhibit a sharper pattern as occupational exposure increases. Occupations whose tasks, according to Webb (2020), are likely to be performed by software or AI based technologies, will experience lower demand until 2030. Instead, the relationship between employment change and robot exposure is not monotone: job recruitment is concentrated at the head and tail of occupations' exposure to robots. Note however that the determinants of the employment projections are not necessarily exogenous to the exposure measure.

5 Exposure measure's robustness check

Given artificial intelligence is still a nascent phenomenon and little is known so far about its impact on the labor market, there is no clear benchmark to test against our measure. As a robustness check, we propose two exercises. First of all, we compare the results obtained using Webb (2019) measure of exposure to AI to alternative measures of occupational risk to AI, i.e. Felten et al. (2021) AI Occupational Exposure (AIOE), Brynjolfsson et al. (2018) Suitability for Machine Learning (SML), and Meindl et al. (2021) Fourth Industrial Revolution exposure (4IR). Secondly, to test Webb (2019) methodology itself, rather than the constructed measure, we perform an equivalent analysis of robots' impact on the Italian labor market, exploiting Paolillo et al. (2022) measure of automation risk as a test against Webb (2019) measure of robot exposure.

5.1 Alternative AI exposure metrics

In this section we compare the AI exposure score by Webb (2019) to other indicators of technological progress: AIOE by Felten et al. (2021), SML by Brynjolfsson et al. (2018), and 4IR by Meindl et al. (2021). Section 3 and Table 1 provide an overview of the indicators.

As shown in Figure 16, we cannot find a strong correlation between Webb (2019)'s AI exposure and the other three metrics considered. Modest correlation is found among AIOE, SML, and 4IR measures, as it appears from Table 5. Appendix B investigates the heterogeneity in exposure across occupations and demographic groups, using as a reference index the four measures under exam. There is a significant discrepancy in the potential exposure predicted by the four indicators.

Table 5: Pearson correlations between AI exposure indicators

	AI	AIOE	SML	4IR
AI	1	-0.10	-0.07	-0.17
AIOE	-0.10	1	0.53	0.68
SML	-0.07	0.53	1	0.63
4IR	-0.17	0.68	0.63	1

Figure 16: Correlation between AI exposure measures, employment weighted

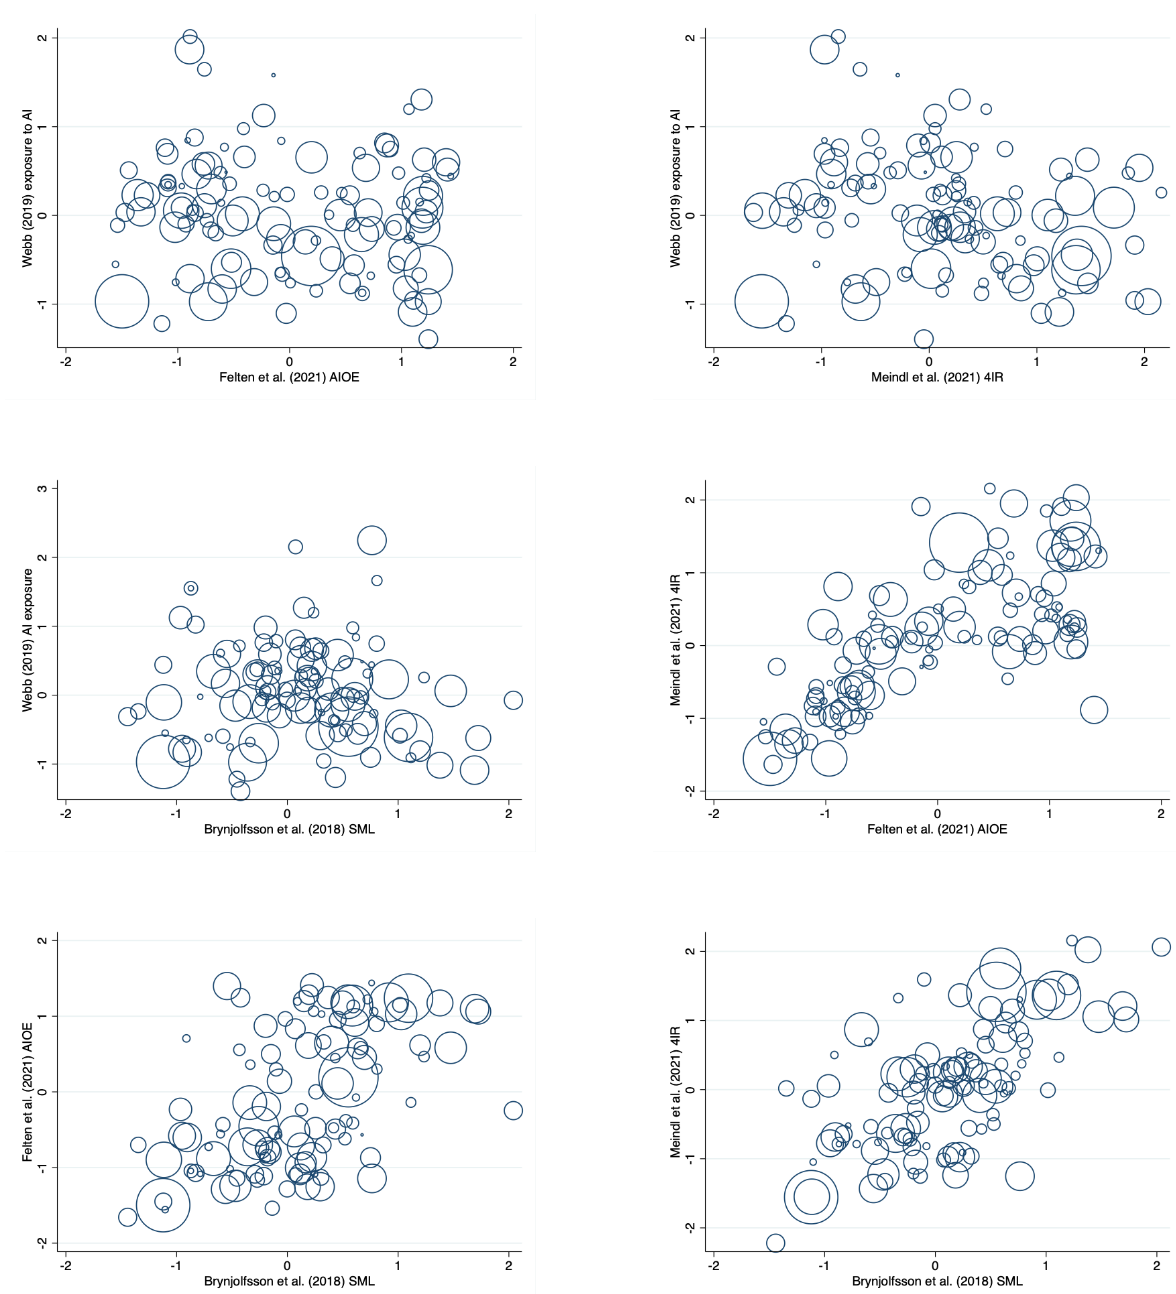

Note: All exposure scores are z-scores. Standardization is not employment-weighted. The corresponding Pearson correlation values are presented in Table 5.

While Meindl et al. (2021) expand to all the technologies of the Fourth Industrial Revolution and Brynjolfsson et al. (2018) target a narrower area of AI, both Webb (2019) and Felten et al. (2021) focus on artificial intelligence. However, the occupations' potential exposure predicted by the two indicators of Webb (2019) and Felten et al. (2021) differ. We envisage several reasons for such discrepancy. First, the AI data sources can be a reason why exposure scores differ. Webb (2019) collects all patents in Google Patents Public Data, provided by IFI CLAIMS Patent Services, that use as keywords, in titles or abstracts, "supervised learning", “reinforcement learning”, “neural network”, and “deep learning”. Conversely, Felten et al. (2021) employ only the 10 AI applications for which the Electronic Frontier Foundation (EFF) has recorded ”measured scientific activity and progress in the technology from 2010 onward” (Felten et al., 2021). In principle, every patent selected by Webb (2019) should be classifiable in one or more of the 10 AI applications considered by Felten et al. (2021). However, it could be the case that by limiting to 10 the number of applications under exam, they’re excluding some sets of applications for which patents exist and are contemplated by Webb (2019)’s measure. The difference in AI data source is reflected in the inconsistency between the two measures.

Second, the fundamentally different algorithms for mapping AI and occupations can explain why scores per occupation differ. Webb (2019) links AI capabilities to occupations by quantifying the overlapping between patents’ text and occupational tasks’ description. In Felten et al. (2021) instead, the link between AI applications and occupations’ abilities is established by survey responses of Amazon’s Mechanical Turkers about the application-ability relatedness.

Patents are increasingly being used as indicators of technological progress or innovative activity (Meindl et al., 2021), as they provide a rigorous and detailed source of information. ”The purpose of the patent system is to encourage invention and technical progress both by providing a temporary monopoly for the inventor and by forcing the early disclosure of the information necessary for the production of this item or the operation of the new process” (Griliches 1998). Given a patent is granted only if its novelty and potential utility are recognized, using patents as a source of available technology helps focus on up-to-date and likely-to-be-employed innovations. Further, patent documents contain technical, legal, and business information in a standardized format, which makes patents well suited to text extraction. Survey responses, vice-versa, provide a potentially biased source of information on the level of relatedness of abilities and technology: mTurkers may be influenced by their perceptions or understanding of the progress and feasibility of the technology.

The debate about patents and surveys does not necessarily lean toward patents. Indeed Webb’s methodology is limited in the extent to which human and technology tasks’ overlapping erroneously predicts areas of AI application, which instead is prevented in Felten et al. (2021) if mTurkers have correct knowledge about AI diffusion. For the sake of clarity, consider the following example. Among the tasks performed by Cardiologists (O*NET SOC code 29-1212.00) there is ”Answer questions that patients have about their health and well-being”. If in Google Patents Public Data there is a patent whose title contains the verb-noun pair ”answer question”, then Webb (2019) will count it as an area of AI application, independently of whether the patent has been employed by the health system. If mTurkers have expertise in the ”Reading comprehension” and ”Language modeling” applications, when asked about the relatedness of the cited O*NET ability to these AI applications, they will correctly report the relatedness.

A further issue about Felten et al. (2021) index is that mTurkers are probably not prone to identify *unconventional* ways of employing AI-based technologies in the workplace, as Webb

(2019) can do. In this sense, areas in which patents have not been applied so far and that generate matching in Webb's metric, can represent future creative diffusion of AI technologies, but will not be recognized by Felten et al. (2021).

Deployment of patent system is common to Webb (2019) and Meindl et al. (2021) metrics. Still, the two measures provide different exposure results. The underlying reasons are the following: (1) *technology type*, Webb (2019) create an AI-specific score, while Meindl et al. (2021) widen to all technologies that shape the Fourth Industrial Revolution; (2) *patent selection*, Webb (2019) selects patents through text search, while Meindl et al. (2021) through CPC classification; (3) *mapping algorithm*, Webb (2019) employs a dependency parsing algorithm, while Meindl et al. (2021) relies on text similarity.

Brynjolfsson et al. (2018)'s rubric is similar to Felten et al. (2021)'s one, as far as both are constructed on experts' judgments. However, Brynjolfsson et al. (2018) restricts attention to machine learning, which is only one of the branches of AI. Given that SML presents the occupations' potential to be automated by machine learning, once technologies are fully implemented, and given AI exposure and AIOE capture the availability and applicability of AI technologies, SML can be regarded as complementary to AI exposure or AIOE. Occupations in the first and fourth quadrants of the third correlation plot display high suitability for machine learning, but differ in the extent to which they are exposed to AI patents, meaning that occupations in the first quadrants are expected to be subject to greater pressure to adapt to new technologies, either in terms of job displacement or task reorganization.

The above illustration provides sufficient reasons why the occupational exposure predicted by Webb (2019)'s metric and the alternative indicators differ.

In spite of the cited differences, the labor market impact estimated under the exposure metric by Webb (2019) is consistent with the results obtained via the other metrics. Figure 18 reports the results of the regression equivalent to the one in Section 4.2, employing Webb (2019) exposure score in Columns (1)-(5), Felten et al. (2021) metric in Columns (2)-(6), Brynjolfsson et al. (2018) index in Columns (3)-(7), and Meindl et al. (2021) rubric in Columns (4)-(8). All the metrics suggest that greater exposure predicts negative effect on wage change at the occupation-industry level, still statistically insignificant, and negligible effect on employment share change.

Figure 17: Change in wages and employment vs. exposure to artificial intelligence, 2014-2019

	Dependent variable:							
	Change in wage				Change in employment share			
AI exposure	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Webb (2019) AI exposure	-0.00115				-1.33e-05			
	(0.0156)				(1.47e-05)			
Felten et al. (2021) AIOE		-0.0153				6.84e-06		
		(0.0143)				(1.25e-05)		
Brynjolfsson et al. (2018) SML			-0.000890				-1.78e-05	
			(0.0128)				(1.48e-05)	
Meindl et al. (2021) 4IR				-0.00746				5.71e-06
				(0.0112)				(1.47e-05)
Observations	834	834	798	834	881	881	840	881
Demographic controls	Y	Y	Y	Y	Y	Y	Y	Y
Occupational controls	Y	Y	Y	Y	Y	Y	Y	Y
Industry FE	Y	Y	Y	Y	Y	Y	Y	Y

Note: Each observation is an occupation-industry cell. Dependent variables are change in log wage and change of a cell's share of overall employment between 2014 and 2019, winsorized at the top and bottom 1%. Controls are added at the occupation-industry level. Standard errors are clustered by industry. *p<0.1; **p<0.05; ***p<0.01

5.2 Alternative methodologies

Artificial intelligence represents the upcoming wave in a process of automation that has been ongoing for at least 200 years. The different timing of the AI automation and robotization implies robotic appliances have already encroached on human tasks, with subsequent adjustments on the side of the workers, in terms of occupation position or retraining programs, while AI effects have hardly been realized. Consequently, from the standpoint of the researchers, the robotics case study provides a benchmark in the study of other technologies' applications. Therefore, comparing Webb (2019)'s measure of robots' exposure to equivalent measures serves as a robustness check of the consistency of Webb (2019)'s methodology of construction of the exposure measure. For this exercise, we will employ Paolillo et al. (2022) Automation Risk Index (ARI).

Paolillo et al. (2022) Automation Risk Index (ARI) can be interpreted as the relative risk of a job being automated by a robot. To construct it, they first match the 87 human abilities, as defined by the O*NET database, to the 26 robotic abilities, derived from the European H2020 Robotics Multi-Annual Roadmap (MAR) released by SPARC, a partnership between the European Commission and the European Association of Robotics. Human and robotic abilities are matched when their definitions are found to be conceptually similar. A technological readiness level (TRL) is then assigned to each of the matched abilities, according to the scale defined by the European Union.

The index is thus built as a function of the importance and the level of the required human abilities (HA), and the TRL of the corresponding robotic abilities:

$$ARI = \frac{1}{2}(f(HA_{importance}, HA_{level}, TRL_{ha}) + f(HA_{importance}, HA_{level}, TRL_{ha}))$$

For the matched abilities, $TRL_{ha} = TRL_{la}$. For the unmatched abilities, $TRL_{ha} = 0$ and TRL_{ha}, which is equivalent to assuming that there is an equal probability of a low-automation and a high-automation scenario, with, respectively, the corresponding technology not developed at all or fully developed.

The index can be interpreted as the proportion of abilities, required by a job, that robots can perform, accounting for the importance of the abilities in the occupation.

Figure 18 reveals a high correlation between the two measures. In Appendix C we perform the exercises proposed in Sections 4.1.3 and 4.1.4 for the two indices, which seems to provide support in favor of the consistency of Webb (2020)'s measures.

As a final robustness check, we estimate the labor market impact of robots, using both Webb (2020) and Paolillo et al. (2020) metrics. To study the link between occupations' exposure to robots and changes in wages and employment, we follow the same empirical strategy as for artificial intelligence, described in Section 4.2.

Figures 19 and 20 are the residual binscatters of the change in wage and change in employment for each vingtile of robot exposure at the occupation-industry level. Figure 21 shows the corresponding regression results. Greater exposure to automation, either measured by Webb's metric or Paolillo et al.'s index, is associated with a statistically significant increase in wages between 2014 and 2019, and a small and not significant change in employment share. Given

Figure 18: Correlation between Webb (2020) exposure to robots and Paolillo et al. (2022) automation risk index, employment weighted

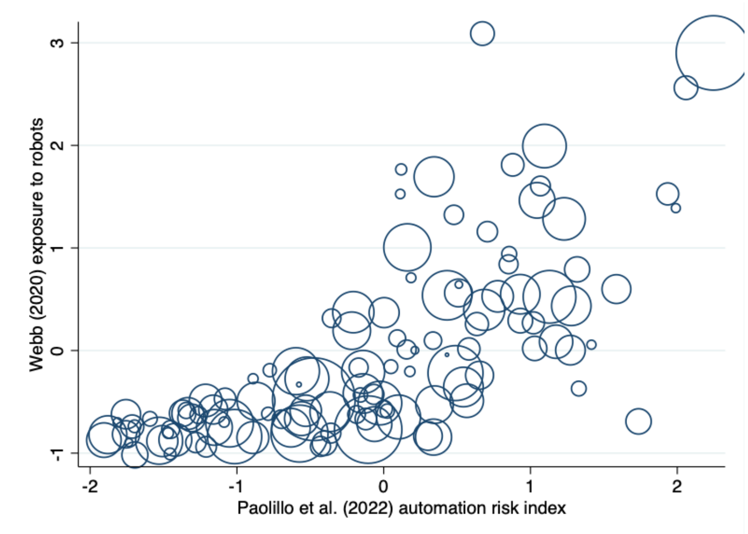

that the estimates of the exposure impact do not differ significantly across the two metrics, the comparison provides support in favor of Webb (2019)'s methodology.

Webb (2019) and Acemoglu and Restrepo (2020) provide evidence from the U.S. labor market, i.e. a strong negative relationship between exposure to robots and wage and employment changes, which is apparently contrasting with ours. However, periods of analysis differ: while Webb (2019) examines 1980-2010 and Acemoglu and Restrepo (2020) 1990-2007, our period of interest is 2014-2019. Consequently, given robotics technology advances can be traced back to the 1990s and 2000s, it is likely that the employment of robots provoked job losses in the first years while inducing some labor market adjustments in the following years, in view of possible additional automation waves, securing jobs in the 2010s.

Figure 19: Change in wages by exposure to robots, 2014-2019

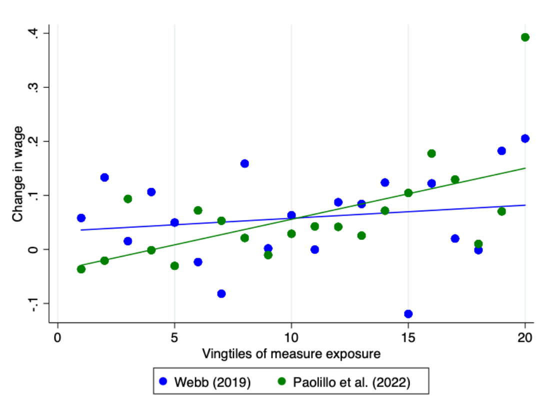

Figure 20: Change in employment by exposure to robots, 2014-2019

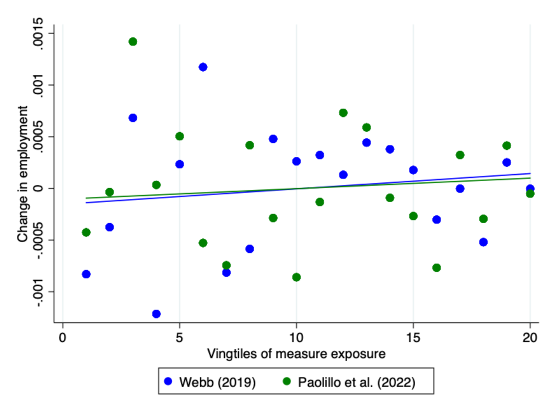

6 Limitations

The analysis conducted provides insights into the labor market impact induced by new technologies. Empirical indicators of technological progress linked to occupations lie at the heart of our research. However, there are several limitations to the employment of these indicators.

Figure 21: Change in wage and employment vs. exposure to robots, 2014-2019

	Dependent variable:			
	Change in wage		Change in employment share	
Robot exposure	(1)	(2)	(3)	(4)
Webb (2019) robots exposure	0.0537***		5.32e-05	
	(0.0138)		(0.000144)	
Paolillo et al. (2022) Automation Risk		0.0532***		-8.12e-06
		(0.0121)		(0.000167)
Observations	151	151	159	159
Demographic controls	Y	Y	Y	Y
Occupational controls	Y	Y	Y	Y

Note: Each observation is an occupation-industry cell. Dependent variables are change in log wage and change of a cell's share of overall employment between 2014 and 2019, winsorized at the top and bottom 1%. Controls are added at the occupation-industry level. Sample is restricted to occupations in the manufacturing sector. *p<0.1; **p<0.05; ***p<0.01

First, the time covered. Webb (2019) metrics are constructed using O*NET task descriptions from 2017, and the ISTAT Labor Force Survey data we employ in Section 4.2 cover the period 2014-2019. As anticipated in Section 4.2, expanding to wider periods could have provided more robust results. However, Webb (2019) scores accurately reflect occupations in 2017 and adjacent years, while may not reflect the occupations' tasks conducted decades ago. Thus employing 2014-2019 wages and employment data reassures potential occupations' tasks inconsistency between labor and exposure data due to the evolution of the jobs. As anticipated in Section 4.2, the choice of the period under exam has been constrained by the absence of workers' occupations ISCO classification in years preceding 2014. The shortness of the period we are left with impedes to observing substantial adjustments in the labor market, following AI applications.

Second, all the technology exposure indicators investigated are based on the U.S. occupational structure, as outlined by the O*NET database. The U.S. occupations and their task descriptions do not necessarily reflect the social, cultural, and economic establishment that shapes the Italian labor market.

Moreover, Webb (2019), Felten et al. (2021), Brynjolfsson et al. (2018), and Meindl et al. (2021) measures provide indicators of *theoretical* exposure, rather than *actual* exposure. Data about AI-based technologies' adoption per occupation would be a valuable resource and would provide an overview of the current versus future technology adoption.

A further limitation generated by the exposure metrics at hand is that they focus on technical feasibility, "silent on the economic, organizational, legal, cultural, and societal factors" (Brynjolfsson et al., 2018) that can affect the adoption of AI systems.

7 Conclusions

The new automation patterns are creating fears of machines making human work obsolete. While robotic appliances have already encroached on human tasks, with subsequent adjustments on the side of the workers, in terms of occupation position or retraining programs, AI effects have hardly been realized. The present analysis, exploiting Webb (2019)'s AI exposure indicator, focuses on the identification of the occupations, and consequently the demographic groups, more likely to be affected by the AI advent, and on whether the Italian labor market has assimilated the AI takeover.

Given the nature of AI and its capabilities in terms of perception, reasoning, and learning, it is likely to affect occupations that were previously thought to be safe from automation. High-skill occupations, out of reach from robotics and software automation, are found to be the most susceptible to artificial intelligence.

No significant impact of artificial intelligence exposure has been estimated on wages and employment. Either employers' adjustments, following AI deployment, have not occurred yet, or occupation-industry level reinstatement and displacement effects counterbalance. Or even Webb (2019)'s AI exposure, being a *forward-looking* measure of the *potential* exposure, is not the appropriate metric to examine the 2014-2019 labor market outcomes.

Several empirical indicators of technological progress, based on tasks' feasibility to be performed by AI-based systems, were leveraged and compared as robustness check: Webb (2019), Felten et al. (2021), Brynjolfsson et al. (2018), and Meindl et al. (2021). Each of the metrics provides different and complementary insights into the labor market effect of AI advances. However, having at disposal occupation level data on the penetration of AI systems would qualify as the most accurate indicator of technology deployment.

The results presented in this thesis provide preliminary insights into the impacts of artificial intelligence on the Italian labor market. Whether AI will lead to within-occupation tasks' transformation or the creation of new jobs, or it will induce human capital investments are further areas of research to investigate.

Appendices

A SOC-ISCO occupations' classification conversion

As anticipated in Section 2.3, out of the 964 6-digit SOC codes for which Webb (2020) provides exposure estimate, after the conversion, we end up with 126 3-digit ISCO08 codes.

First of all, some variation is lost when converting the 6-digit SOC codes to the 4-digit ISCO codes. Consider the example proposed in Table 6: Astronomers (19-2011 SOC code) and Physicists (19-2012 SOC code) are both assigned to the category of Physicists and Astronomers in the ISCO taxonomy (2111 ISCO code). The resulting exposure for ISCO 2111 will be an average of the two, and will not reflect the greater exposure of physicists.

Secondly, the 4-digit ISCO classification has been mapped to the 3-digit ISCO level, to ensure correspondence with the ISTAT Labor Force Survey data. This step leads to further inaccuracy in the resulting exposure measure. All occupations listed in Table 6, with an AI exposure ranging from 0.50 to 1.14, are assigned to the 3-digit ISCO code 211. The associated AI exposure is consequently an average of the 8 occupations listed. The different exposure of the SOC occupation of Physicists (19-2012 SOC code), the 24^{th} most exposed to AI, out of 964, according to Webb (2020), and the SOC occupation of Material Scientists (19-2032 SOC code), the 334^{th} most exposed, fade out. The result is an AI exposure score of 0.77 for the ISCO minor group 211 of Physicists, Chemists and Related Professionals.

Table 6: Example SOC-ISCO conversion

6-digit SOC code			4-digit ISCO code			3-digit ISCO code		
Code	Label	AI exposure	Code	Label	AI exposure	Code	Label	AI exposure
19-2011	Astronomers	0.95	2111	Physicists, astronomers	1.05	211	Physicists, Chemists and Related Professionals	0.77
19-2012	Physicists	1.14	2111	Physicists, astronomers	1.05	211	Physicists, Chemists and Related Professionals	0.77
19-2021	Atmospheric, Space Scientists	0.91	2112	Meteorologists	0.91	211	Physicists, Chemists and Related Professionals	0.77
19-2031	Chemists	0.76	2113	Chemists	0.63	211	Physicists, Chemists and Related Professionals	0.77
19-2032	Material Scientists	0.50	2113	Chemists	0.63	211	Physicists, Chemists and Related Professionals	0.77
19-2042	Geoscientists, Except Hydrologists, Geographers	0.54	2114	Geologists, geophysicists	0.58	211	Physicists, Chemists and Related Professionals	0.77
19-2043	Hydrologists	0.66	2114	Geologists, geophysicists	0.58	211	Physicists, Chemists and Related Professionals	0.77
19-2099	Physical Scientists, All other	0.68	211	Geologists, geophysicists	0.68	211	Physicists, Chemists and Related Professionals	0.77

B Artificial Intelligence exposure measures' comparison

Figure 22 plots the average standardized metrics of exposure to artificial intelligence for each age, smoothed using a locally weighted smoothing regression, separately for the four measures developed by Webb (2019), Felten et al. (2021), Brynjolfsson et al. (2018), and Meindl et al. (2021). Figures 23 and 24 replicate Figure 22 splitting by gender.

Figure 22: Standardized exposure to AI, by measure and age

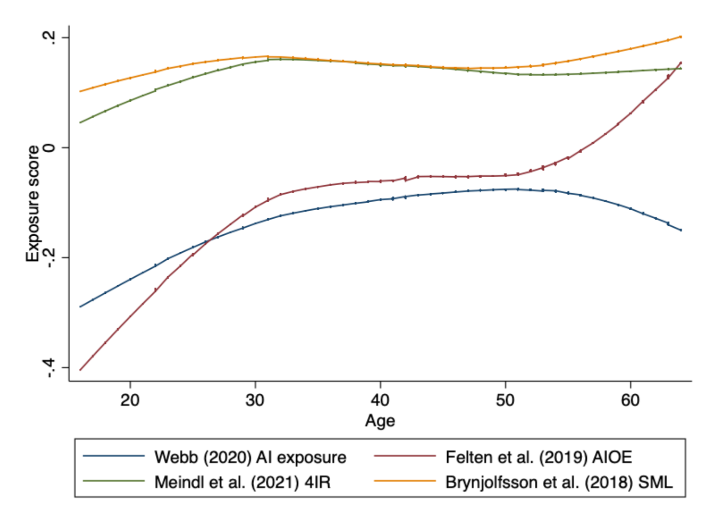

Figure 23: Male standardized exposure to AI, by measure

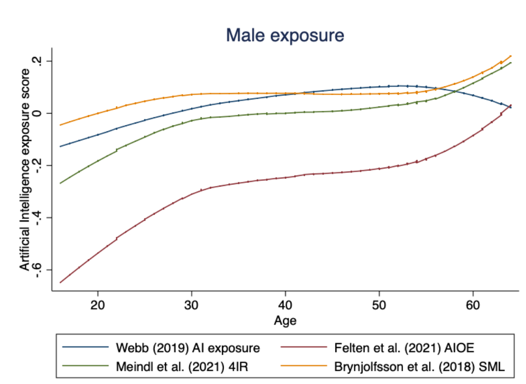

Figure 24: Female standardized exposure to AI, by measure

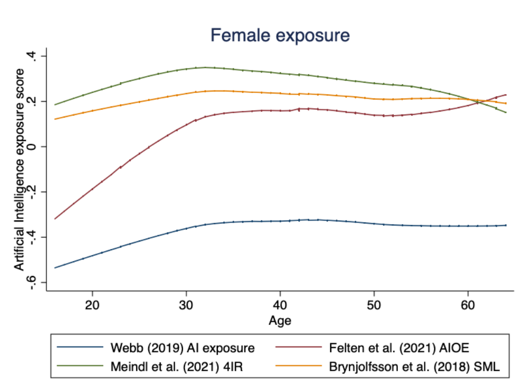

Figure 25 shows the average AI exposure score by Webb (2019), the AIOE index by Felten et al. (2021), the SML metric by Brynjolfsson et al. (2018), and the 4IR measure by Meindl et al. (2021) for each ISCED education level.

Figure 25: Standardized exposure to AI, by measure and education level

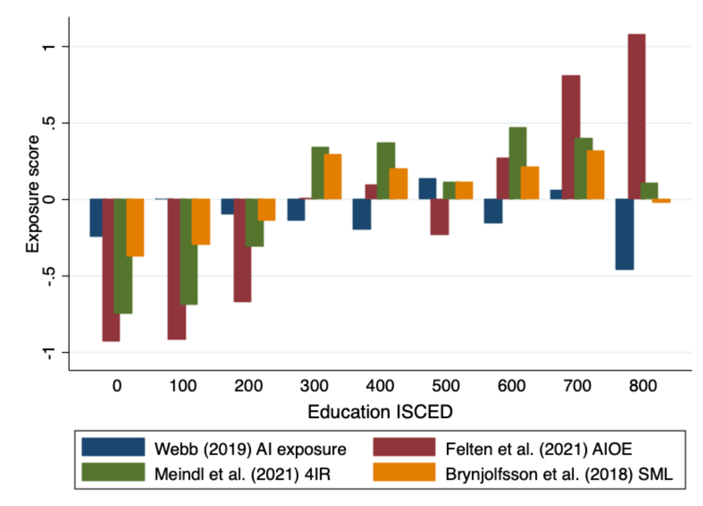

Figure 26 depicts for each wage quantiles the average standardized Webb (2019)'s AI exposure, the average standardized Felten et al. (2021)'s AI Occupational Exposure, the average standardized Brynjolfsson et al. (2018)'s Suitability for Machine Learning index, and the average standardized Meindl et al. (2021)'s Fourth Industrial Revolution exposure, smoothed using a locally weighted smoothing regression.

Figure 26: Standardized exposure to AI, by measure and wage quantiles

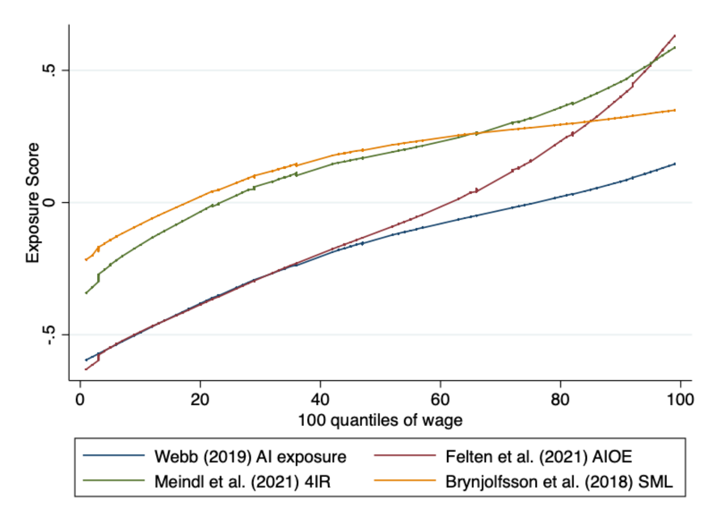

Maps in Figures 30 and 28 plot the regional average exposure to AI, using respectively Webb (2019), Felten et al. (2021), Brynjolfsson et al. (2018), and Meindl et al. (2021) metrics.

Figure 27: Webb (2020) standardized exposure to AI, by region

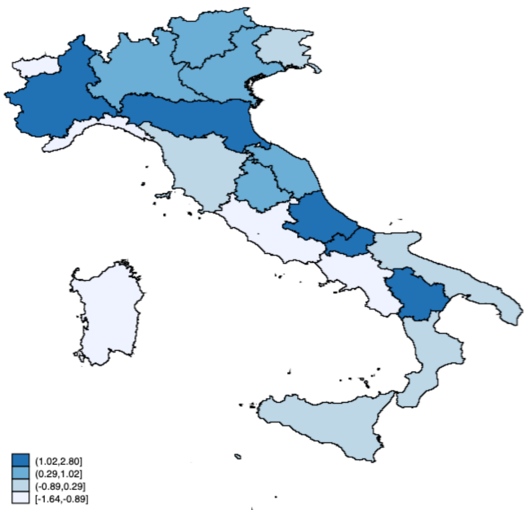

Figure 28: Felten et al. (2021) standardized AIOE, by region

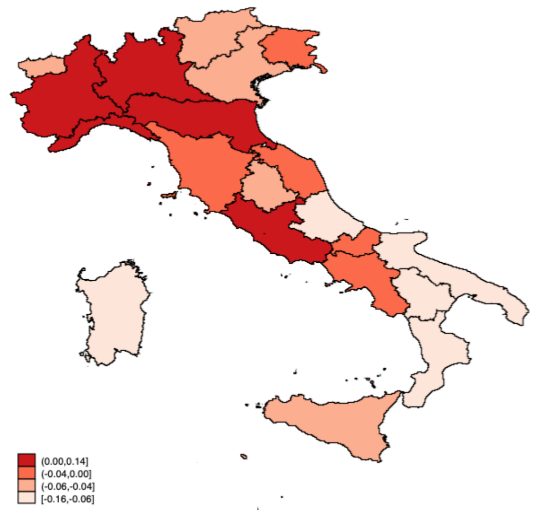

Figure 29: Brynjolfsson et al. (2018) standardized SML, by region

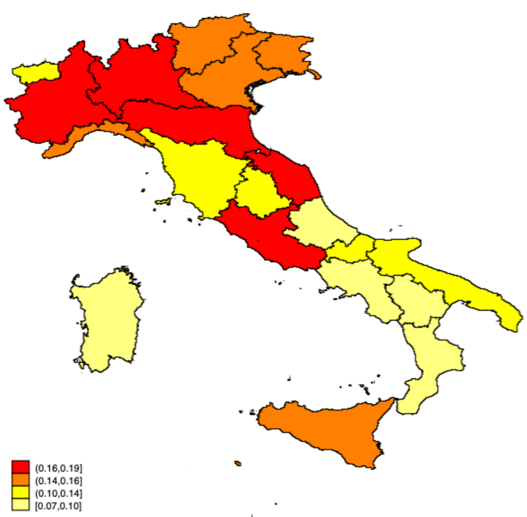

Figure 30: Meindl et al. (2021) standardized 4IR, by region

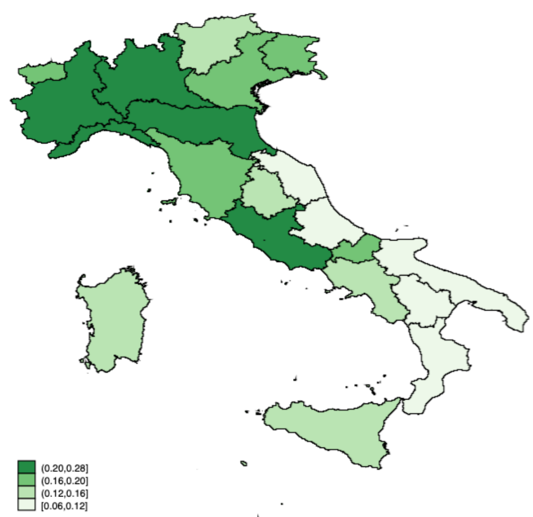

C Robot exposure measures' comparison

Figure 31 plots the average standardized metrics of exposure to robots for each age, smoothed using a locally weighted smoothing regression, separately for the two measures developed by Webb (2020) and Paolillo et al. (2022). Figures 32 and 33 replicate Figure 31 splitting by gender.

Figure 31: Standardized exposure to robots, by measure and age

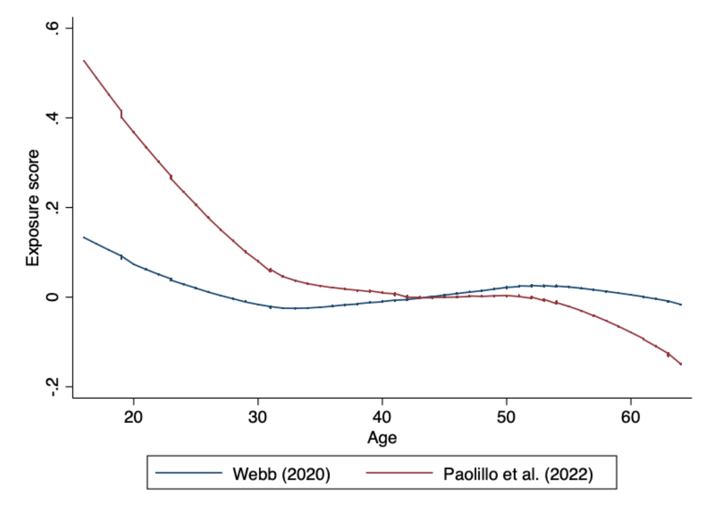

Figure 32: Male standardized exposure to robots, by measure

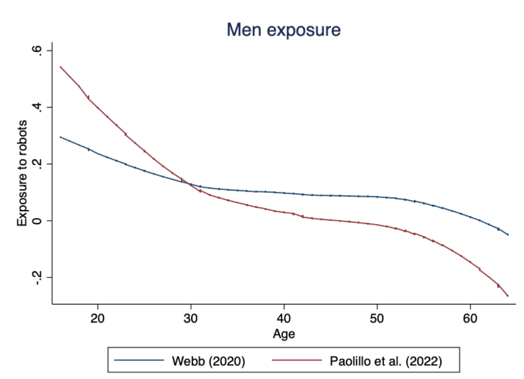

Figure 33: Female standardized exposure to robots, by measure

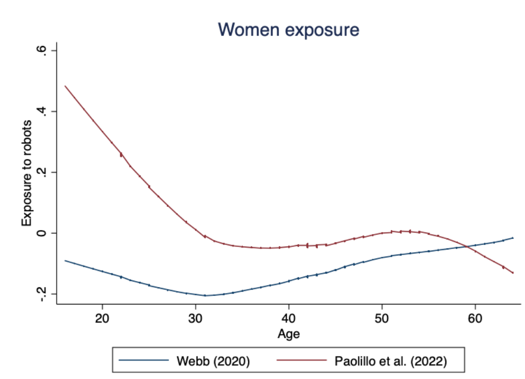

Figure 34 shows the average robot exposure score by Webb (2020) and the average automation risk index by Paolillo et al. (2022) for each ISCED education level.

Figure 34: Standardized exposure to robots, by measure and education level

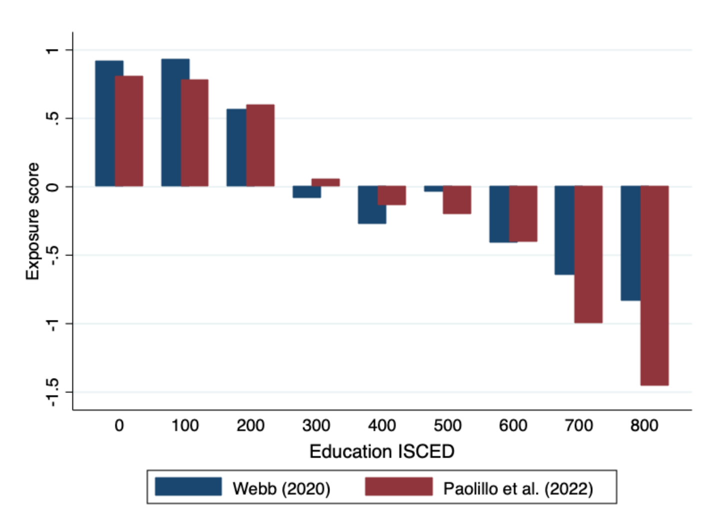

Figure 35 depicts for each wage quantiles the average standardized Webb (2020) robot exposure and the average standardized Paolillo et al. (2022) automation risk index, smoothed using a locally weighted smoothing regression.

Figure 35: Standardized exposure to robots, by measure and wage quantiles

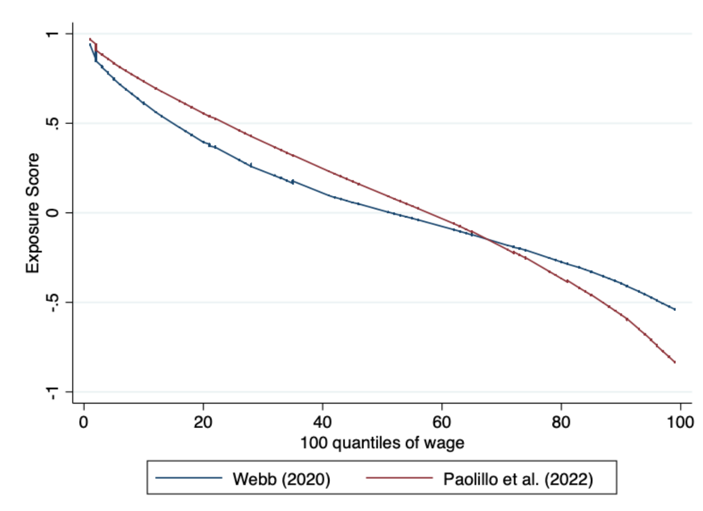

Maps in Figures 36 and 37 plot the regional average exposure to robots, using respectively Webb (2020) measure and Paolillo et al. (2022) index.

Figure 36: Webb (2020) standardized exposure to AI, by region

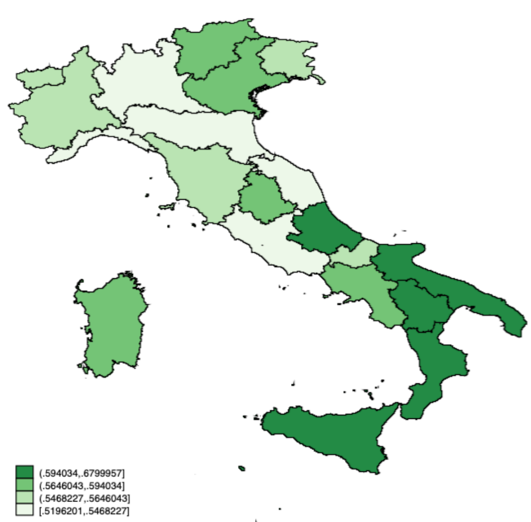

Figure 37: Paolillo et al. (2022) standardized ARI, by region

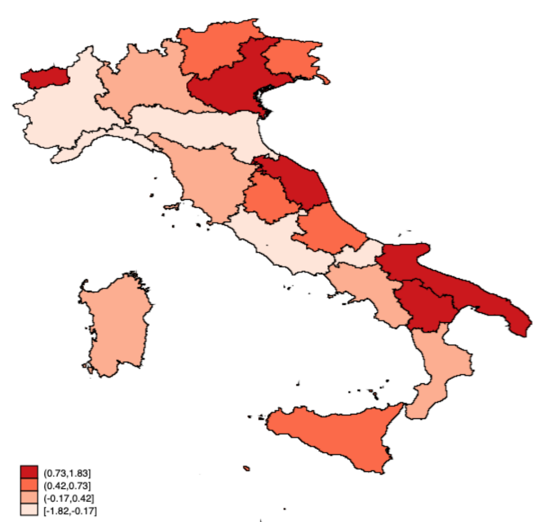

List of Figures

List of Tables

References

Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. *Handbook of Labor Economics*, *4*, 1043-1171.

Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2020). AI and Jobs: Evidence from Online Vacancies. *NBER Working Papers*, *28257*.

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

Acemoglu, D., & Restrepo, P. (2020). Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy*, *128*(6).

Argentin, G., & Triventi, M. (2015, June). The North-South Divide in School Grading Standards: New Evidence from National Assessment of the Italian Student Population. *Italian Journal of Sociology of Education*, *7(2)*, 157-185.

Autor, D., & Dorn, D. (2013). The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review*, *103*.

Brynjolfsson, E., & Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science*, *358*, 1530-1534.

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, American Economic Association*, *108*, 43-47.

Daniele, V., & Malanima, P. (2017). Regional Wages and the North-South Disparity in Italy after the Unification. *Rivista di Storia Economica*, *XXXIII*(2), 117-158.

Felten, E., Raj, M., & Seamans, R. (2019). The Occupational Impact of Artificial Intelligence: Labor, Skills, and Polarization. *NYU Stern School of Business*.

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.

Fossen, F., & Sorgner, A. (2019). New Digital Technologies and Heterogeneous Employment and Wage Dynamics in the United States: Evidence from Individual-Level Data. *IZA Discussion Paper Series*, *12242*.

Meindl, B., Frank, M., & Mendonça, J. (2021). Exposure of occupations to technologies of the fourth industrial revolution.

Ménière, Y., Philpott, J., Rodríguez, J. P., Rudyk, I., Wewege, S., & Wienold, N. (2019). Patents and the Fourth Industrial Revolution. *Munich: European Patent Office*.

OECD. (2022). Recommendation of the Council on Artificial Intelligence. *OECD/LEGAL/0449*.

Paolillo, A., Colella, F., Nosengo, N., Schiano, F., Stewart, W., Zambrano, D., ... Floreano, D. (2022). How to compete with robots by assessing job automation risks and resilient alternatives. *Science Robotics*, *7 (65)*.

Risch, J., & Krestel, R. (2019). Domain-specific word embeddings for patent classification. *Data Technologies and Applications*, *53*, 108-122.

Sparrow, R., Howard, M., & Degeling, C. (2021). Managing the risks of artificial intelligence in agriculture. *NJAS: Impact in Agricultural and Life Sciences*, *93:1*, 172-19.

Turing, A. M. (1950). I.—Computing machinery and intelligence. *Mind*, *LIX*(236), 433-460.

Webb, M. (2019). The Impact of Artificial Intelligence on the Labor Market. , Available at SSRN: https://ssrn.com/abstract=3482150.

Executive summary

As past waves of automation have reshaped our economies, artificial intelligence is offering far-reaching innovations in every field, from consumer to industrial applications, promising to generate productivity gains. It is at the core of the fourth industrial revolution. At the same time, AI also raises questions about its impact on the labor market.

Over the past twenty years, breakthroughs in AI development have been made, continually expanding the nature and the boundaries of AI, calling for the need of defining, regulating, and analyzing the impact of this technology.

What is the effect of the application of technologies to tasks that were previously performed by humans? The answer is not univocal. In the big picture of creative destruction, some tasks are being replaced by robots or software, from contactless cashiers to welding drones, and some occupations are being created, such as social media managers or cloud specialists. The impact of technologies on occupations is not confined to labor displacement or labor reinstatement: some tasks are being transformed by the technological advances, e.g. nurses, exempted from recording patient history, will fully devote themselves to monitoring patients. Furthermore, new technologies increase the productivity of labor in the tasks they perform, which in turn boosts labor demand (Acemoglu & Restrepo, 2019). Thus, the impact of the application of technologies on the labor market is not obvious.

There is little systematic evidence of the impact of AI on labor. Part of the reason for the lack of evidence is that the rapid advancement in AI is a nascent phenomenon. U.S. evidence suggests that advances in artificial intelligence have so far led to a small increase in wages, with no change in employment (Felten et al., 2019), associated with greater job stability (Fossen & Sorgner, 2019) and a shift in the composition of job postings away from occupations most exposed to AI to occupations least exposed (Acemoglu et al., 2020).

In this thesis, we aim at furthering our understanding of the impact of the advent of artificial intelligence on the Italian labor market. Which are the occupations, and consequently the demographic groups, that are suffering more in the aftermath of AI advent, and whether the Italian labor market has assimilated the AI takeover are the insights offered in this thesis.

To estimate the capabilities of AI and identify the aggregate effects of AI on the Italian labor market, we exploit the AI occupations' exposure proposed by Webb (2019). Getting rid of expert judgments about occupations' susceptibility to AI applications, Webb (2019) relies on patents as the source of information to understand which tasks AI can perform. He uses the text of patents to identify what the technology can do, then to quantify the extent to which each occupation in the economy involves similar tasks. Operationally, through a dependency parsing algorithm, for each O*NET occupation, he extracts all the verb-noun pairs from the set of patents of AI-based systems and the verb-noun pairs from the tasks of that occupations, and uses the relative frequency of the pairs that are common to the occupation and the patent to calculate the measure of occupation-level exposure. Webb (2019) provides similar measures for software and robots.

By linking Webb (2019) measure of exposure to robot, software, and AI to ISTAT Labor Force Survey data, we provide insights into the different likely impacts of the three technologies on the labor market, by looking at the heterogeneity in the exposure scores across demographic groups.

Previous-generation automation has affected routine manual and routine cognitive tasks. Applications were limited to areas where, once instructions are specified in advance for each contingency, human knowledge or expertise is not required. By contrast, the nature of artificial intelligence opens a new set of possibilities for automation. Routine tasks place workers at greater exposure to robots and software. Jobs degree requirements that could secure individuals from exposure to robots and software, do not secure them from AI. Still, for all the three technologies, subjective inputs and interpersonal skills protect occupations from technology substitution, as well as reasoning about situations never seen before.

Whether the artificial intelligence revolution will exacerbate labor market polarization has been under debate in the last years. Investigating the heterogeneity in exposure by skill level anticipates whether AI will have a skill-biased impact on the labor market, whatever the direction of the impact. Exposure to software and robots is declining on average in educational attainment: occupations performed by workers with at most primary or lower secondary education tend to be more exposed than those carried out by individuals with a degree. Conversely, no significant difference in average exposure is found for an individual with secondary school or a bachelor's degree. By looking at wage as a proxy for skills, robotics capabilities, as measured by patenting intensity, are most common in low-wage occupations, software capabilities in middle-wage occupations, and artificial intelligence capabilities in high-wage occupations.

So far our analysis does not comment on the direction and magnitude of the impact of AI on the Italian labor market, but only on the individuals more likely to be affected by the AI aftermath. Moreover, note that Webb (2019) measure of exposure to technologies is forward-looking: it takes into account all the possible applications of AI, not only areas in which adoption has already occurred, but also areas in which it might occur in the future, given the nature of the patent. Furthermore, Webb's measure is agnostic about whether exposure takes the form of labor displacement, labor reinstatement, or within-occupation change in performed tasks. By investigating the relationship between exposure and changes in employment and wages, we obtain very small and statistically insignificant results. In Italy between 2014 and 2019 occupations with high exposure to artificial intelligence, as estimated by Webb (2019), have experienced no additional change in wage and employment, compared to occupations characterized by low levels of exposure. Felten et al. (2021) provided evidence of a positive but not statistically significant relationship between employment growth and his measure of AI Occupational Exposure in the U.S. market between 2010 and 2015, while small positive and statistically significant relationship with wage growth. The different impact of AI on the Italian and U.S. market may be the result of the greater rigidity of the Italian labor market. Although we cannot attribute causality to our results, one possibility is that the advent of AI has not altered yet the labor market. As the diffusion of AI-based systems began to be felt in the second half of the 2000s, it could be that employers' adjustments, following AI deployment, have not occurred yet. Another possibility, that our data do not allow to control for, is that occupation-industry level reinstatement and displacement effects counterbalance, generating negligible effect overall. A further possible explanation targets our exposure measure. The occupational exposure score constructed by Webb (2019), by measuring the overlapping between tasks performed by patents and humans, takes into consideration all the possible, actual or future, conceivable or not, applications of the existing patents. It provides an estimate of the occupations' exposure once the potential of the so-far developed patents has been fully realized. As such, if, until 2019, there are still areas of work that could deploy AI, our measure is going to display little relationship with the labor market outcomes.

Given artificial intelligence is still a nascent phenomenon and little is known so far about its impact on the labor market, there is no clear benchmark to test against our measure. As a robustness check, we propose to compare the results obtained using Webb (2019) measure of exposure to AI to alternative measures of occupational risk to AI, i.e. Felten et al. (2021) AI Occupational Exposure (AIOE), Brynjolfsson et al. (2018) Suitability for Machine Learning (SML), and Meindl et al. (2021) Fourth Industrial Revolution exposure (4IR). Alternative measures of exposure to AI, although differing in the data sources and methodology, provide similar results. All the metrics suggest that greater exposure predicts a negative effect on wage change at the occupation-industry level, still statistically insignificant, and negligible effect on employment share change.

Additionally, to test Webb (2019) methodology itself, rather than the constructed measure, we perform an equivalent analysis of robots' impact on the Italian labor market, exploiting Paolillo et al. (2022) Automation Risk Index as a test against Webb (2019) measure of robot exposure. Greater exposure to automation, either measured by Webb's metric or Paolillo et al.'s index, is associated with an increase in wages and employment share between 2014 and 2019. Given that the estimates of the exposure impact do not differ significantly across the two metrics, the comparison provides support in favor of Webb (2019)'s methodology.

Our analysis however presents several limitations. First of all, the shortness of the period we are left with impedes to observing substantial adjustments in the labor market, following AI application. Moreover, all the technology exposure indexes investigated are based on the U.S. occupational structure, as outlined by the O*NET database, are indicators of theoretical exposure, rather than actual exposure, and they focus on technical feasibility, "silent on the economic, organizational, legal, cultural, and societal factors" (Brynjolfsson et al., 2018) that can affect the adoption of AI systems.

The results presented in this thesis provide preliminary insights into the impact of artificial intelligence on the Italian labor market. Whether AI will lead to within-occupation tasks' transformation or the creation of new jobs, or it will induce human capital investments are further areas of research to investigate.