Pillars – Pathways to Inclusive Labour Markets

Report

Skill Demand, Skill Supply, and the Prevalence of Skill Mismatch in the European Union

October 2022

This project receives funding from the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement No. 101004703.

WP 2.1 : Skill Demand, Skill Supply, and the Prevalence of Skill Mismatch in the European Union

Yuchen Guo[a,b], Christina Langer[a,b,c], Fabio Mercorio[d,e], Mario Mezzanzanica[d,e], Filippo Pallucchini[d,e], Francesco Trentini[d,e], Simon Wiederhold[a,b]

[a] *KU Eichstätt-Ingolstadt* [b] *ifo Institute Munich* [c] *Harvard Kennedy School* [d] *CRISP Research Centre - University of Milano-Bicocca* [e] *University of Milano-Bicocca*

Abstract

Labor market developments such as globalization, structural transformation, and accelerating technological change can lead to mismatches between firms' skill demand and employees' skill supply. While skill mismatch is heavily discussed in research and policy, empirical evidence on the existence and determinants of skill mismatch in Europe is very scarce. In this paper, we develop novel measures of skill mismatch in Europe to address various questions of high relevance for labor market policies in the European Union: (1) How prevalent is skill mismatch in Europe? (2) What are the drivers of skill mismatch and how can workers better prepare for the skill demand of employers? We draw on innovative online job ad data and skill survey data for 17 European countries to measure skill demand and supply, respectively. Applying modern machine learning techniques to link the demand and supply data, we develop new measures of skill mismatch that allow us to investigate mismatch across occupations, skill domains, gender, industries, and regions. We document that skill mismatch is a widespread phenomenon in Europe, while the extent and direction of mismatch varies across occupations and regions. We also show that skill mismatch across European regions exists at two margins: First, the skills current workers possess do not match the skills demanded by employers (intensive margin). Second, the number of early-career workers who are qualified to work in certain occupations is not sufficient to fill the vacancies in these occupations (extensive margin). We further study skill mismatch at the regional level and its relationship to economic, industrial, and structural characteristics. We show that more prosperous and more dynamic regions systematically face less skill shortage, while regions more exposed to technological change (in particular, automation) are more severely affected by skill shortage. At the same time, the prevalence of on-the-job training reduces skill shortage, pointing to the role of education and training systems to ensure the employability of the workforce in Europe.

Keywords: skill mismatch, online job ads, labor market, word embeddings

*This project receives funding from the European Union's Horizon 2020 Resesarch and Innovation Programme under Grant Agreement No 101004703. See Guo et al. (2022) for a policy brief summarizing some of the main results and policy conclusions from this report.

1. Introduction

Labor markets face challenges due to globalization, structural transformation, and advancing technological change. These can lead to a mismatch between firms' skill demand and employees' skill supply, which can go in two directions: Workers having a skill surplus, where skill supply exceeds demand, or workers experiencing skill shortage, where firms' skill demand is larger than the skills workers possess. In light of this, the European Union's Agenda for New Skills and Jobs (European Commission, 2011, p. 2) states that creating a more skilled workforce "is a considerable challenge given the rapidly-changing skills needed and the persistent skill mismatches in the EU labor market". In the same report, the European Commission also "established the anticipation and matching of labor market and skill needs as a top priority for the EU" (European Commission, 2011, p. 8). However, despite the enormous policy interest, empirical evidence on the prevalence of skill mismatch in Europe is scarce.

We contribute to the understanding of skill mismatch in the European Union by presenting novel evidence on skill mismatch across countries, occupations, skill domains, gender, industries, and regions for 17 European countries. Our skill mismatch measures are based on several innovative data sources. In particular, we leverage online job advertisement (OJA) data to study skills requested by employers (skills demand). OJA data provide real-time labor market data on employers' skill demand at an unprecedented level of granularity. This allows us to observe variation in skill demand between occupations, industries, and regions. Further, in contrast to previous measures of skill mismatch, we do not rely on survey data to measure skill demand, which suffer from significant time lags and capture past labor market demands. Rather, OJA data provide information on current skill demand in the labor market in real time.

To capture the skills provided by workers (skill supply), we draw on representative survey data from the Programme for the International Assessment of Adult Competencies (PIAAC). These data offer comprehensive information about skills used at the workplace for a wide range of skill domains. We link OJA data and PIAAC survey data using modern machine learning methods of word embeddings. The unique level of granularity of both OJA and PIAAC data allows us to quantify *which* and by how much skills are lacking at the level of occupations, industries, and regions. We coin mismatch between the skills supplied by workers employed in a given occupation and the skills that employers demand from workers in this occupation as skill mismatch at the *intensive margin*.

Using our novel measure of skill mismatch at the intensive margin, we document several key findings: First, skill gaps in European countries exist, but the extent and direction vary across occupation types. Workers in cognitive-intensive occupations provide more skills than are demanded (skill surplus), whereas workers in manual-intensive occupations face higher skill demand compared to the skills they have (skill shortage). The described pattern is consistent across all 17 European countries that are included in our analysis. This suggests that overall patterns of skill mismatch do not reflect country-specific factors but are rather a European-wide phenomenon. Second, we document a similar pattern of skill mismatch for different skill domains (i.e., digital, numeracy, literacy, and social skills). Thus, the observed skill gaps are not driven by a lack of specific skills such as digital or social skills. However, repeating the analysis separately for males and females shows a slightly different picture. While the skill gaps point in the same direction for male and female cognitive non-routine and manual routine workers, this is not always the case for the cognitive routine and manual non-routine occupation types: For instance, in cognitive routine occupations, females exhibit skill surpluses in numeracy, literacy, and social skills, while males face skill shortages in these skill domains. However, the overall patterns of skill mismatch do not differ much by gender.[1]

Our novel measure of skill mismatch also allows us depict the degree of skill mismatch across European regions. Our data suggests that, on average across regions, the skills possessed by currently employed workers do not fall short of employers' skill requirements. However, we find considerable heterogeneity in skill mismatch across regions and countries. While countries such as Greece, Italy, Spain, and Poland generally suffer from skill shortages, the other European countries are characterized by skill surpluses. However, there is also considerable variation in the degree of skill mismatch within countries. We thus investigate potential determinants of differences in skill shortage across regions, such as a region's industry specialization and economic performance. We find that economically stronger and more dynamic regions systematically face less skill shortage, possibly due to higher incentives to invest in skills in these regions. We also assess the role of trade competition and exposure to automation technology in explaining variation in skill shortage across regions. While regions more exposed to trade and import competition do not suffer from higher skill shortage, we do observe a positive association between skill shortage with exposure to automation risk, as automation technologies shift existing skill demands more rapidly. At the same time, we document a negative association of skill shortage with on-the-job training, suggesting that training policies could be a poten-

[1] Note that this result could partly be driven by the fact that OJAs are always targeted at both males and females, so we could not derive a gender-specific measure of skill demand.

tial means to help workers to cope with changing skill requirements and keep their skills up-to-date.

Lastly, we complement our measure of skill mismatch at the *intensive margin* with a quantification of skill mismatch at the *extensive margin* across European regions, i.e., the gap between the number of workers demanded in different occupations and the number of individuals that are suited to work in these occupations. To construct our measure of extensive-margin skill mismatch, we combine several data sets in a novel way to represent the skill supply side. First, we leverage data from the European Tertiary Education Roster (ETER), providing the number of graduates by field of study and region from the near universe of higher education institutions in Europe. Together with information on the occupational distribution of graduates by field of study from PIAAC, we infer the potential supply for selected occupations at the regional level. This allows us to compare the number of online job advertisements for a given occupation with the number of graduates potentially suited to work in this occupation across European regions.

Our data suggest that the skill shortage at the extensive margin varies considerably across European regions for occupations such as business administration professionals or legal, cultural, and social professionals. For instance, the potential supply of business administration professionals exceeds the demand by 38 percent in the region of Eastern Macedonia in Greece (i.e., extensive-margin skill surplus), while demand exceeds potential supply by as much as 94 percent in the Brussels region (i.e., extensive-margin skill shortage). We further document that skill shortages of qualified graduates for, e.g., health professionals exist especially in Germany and the United Kingdom. However, at the same time, we find pronounced skill shortages for information and communication technology (ICT) and science, technology, engineering, and mathematics (STEM) workers in almost all European regions.

Our measure of skill shortage at the extensive margin shows *how many workers* are lacking to satisfy employers' occupational demand, complementing our analysis at the intensive margin of which skills are lacking in an occupation. Both measures depict skill mismatch from different angles and can potentially lead to different conclusions. For instance, at the intensive margin, the skills provided by individuals actually working as ICT professionals closely match the skills that are required by firms. However, at the extensive margin, there exists a pronounced shortage of ICT professionals as the number of vacancies greatly exceeds the potential regional supply of individuals suited to work in this occupation.

Overall, our analyses of intensive-margin and extensive-margin skill shortage suggest that skill mismatch is a challenge in the European Union for at least to reasons: First, education and training systems need to provide a fitting set of skills allowing individuals to match the skill demands on the labor market. Our investigation of intensive-margin skill mismatch shows *which* and by how much skills are lacking. Our results suggest that the extent of skill shortages varies widely across occupations and regions, suggesting different needs for policy action. Thus, our findings offer insights into which types of workers and regions might particularly benefit from targeted education and (re-)training policies. One such case are occupations or regions that were exposed to technological change at a more rapid pace, and thus have been particularly affected by the automation of routine tasks.

Second, there may also be shortages in the number of people that are qualified to work in a given occupation, which we coin extensive-margin skill mismatch. This is a measure of *how many workers* would be needed to fulfill the occupational demand by employers across European regions. Our results suggest that almost all European regions suffer shortages when it comes to ICT and STEM workers. This is consistent with two much-discussed developments. First, that the proliferation of personal computers caused a shift away from routine tasks–that is, those more amenable to automation–toward problem-solving and complex communication tasks (typically called “nonroutine abstract tasks”). This argument was first made by Autor et al. (2003) when developing their task-based approach to skill-biased technological change.[2] Our results are also consistent with the concern frequently expressed by researchers and policymakers that the supply of STEM workers is not sufficient to meet the demands of the labor market (see e.g., Xue and Larson, 2015). Our findings, based on the unique combination of OJA and worker survey data, suggest that these concerns are not only valid, but that the shortage of ICT and STEM professionals is substantial across Europe. Thus, almost all regions in Europe could potentially benefit from targeted education and occupational re-training policies to increase the number of ICT and STEM workers.

Our analyses of skill demand, skill supply, and skill mismatch contributes to the literature in several ways: First, we use novel OJA data for a large number of European countries to depict employers’ demand for skills. OJA data provide real-time skill demand data on a very granular level, which accounts for heterogeneity between occupations, industries, and regions. Unlike previous studies, we do not rely on survey data to measure skill demand, which are subject to significant time lags and may reflect a labor market

[2]See Acemoglu and Autor (2011) as well as Autor (2015) for reviews of this literature.

situation that has already been superseded by more recent developments (for an overview, see Hartog, 2000). OJA data are also more objective than survey data and not prone to reporting bias. Second, on the supply side, we use large international worker skill data, complemented with information from the near universe of higher education institutions in European countries. Linking demand and supply, the resulting skill mismatch measure is more granular than aggregate measures such as the number of vacancies requiring a specific college major relative to the number of graduates holding this degree (Berkes et al., 2018) or the fraction of potential hiring that does not take place because of an occupational mismatch of unemployed workers relative to the occupational distribution of online job ads (Burke et al., 2020). Our paper is the first that linked real-time labor market data from OJAs to depict employers' skill demand with ETER/PIAAC data to measure skill supply. Doing so, we achieve a better depiction of the mismatch of skills actually required by firms and those supplied by (prospective) workers than ever before in a cross-country setting. Furthermore, we advance the literature by investigating the role of technological change and industrial transformation, as well as trade and the reshuffling of global value chains, as potential determinants of skill mismatch.

This report proceeds as follows. In Section 2, we provide an overview of the data used in our analyses. Section 3 develops two novel measures of skill mismatch by linking skill demand and supply data. Section 4 presents out results on skill mismatch across occupations, skill domains, gender, industry, and regions. We also investigate potential determinants of skill mismatch and the economic cost of mismatch. Furthermore, we provide a measure of extensive-margin skill mismatch at the regional level in Section 5. Section 6 concludes.

2. Data

Our analysis leverages a variety of data sources. We use the European Skills, Competences, Qualifications and Occupations (ESCO) skill classification, online job ads (OJA) data, data from the Programme for International Assessment of Adult Competencies (PIAAC), the European Tertiary Education Register (ETER), and the European Labour Force Survey (EU-LFS), all of which are described below.

2.1. ESCO

ESCO (European classification of Skills, Competencies, and Occupations) provides a multilingual dictionary of occupations and skill requirements organised along two main pillars. The first is the "occupation" pillar, which is referenced to the ISCO08 standard.

The second is the “skills” pillar, which lists and describes skills linked to occupations. ESCO, therefore, provides a list of occupations and related skills organised as a network; however, it gives no information on the importance of skills in the considered occupation. We use the occurrences of skills in online job advertisements to complement ESCO with information on skill relevance.

2.2. OJA

We exploit online job advertisements (OJA) to capture skill demand. These real-time data provide information on the skills that firms request at an unprecedented level of resolution. Our OJA data are obtained from Eurostat and CEDEFOP as part of the Web Intelligence Hub-Online Job Advertisements (WIH-OJA) project. The system has been collecting online job advertisements to analyse vacancies and emerging skill requirements across all EU countries since the last quarter of 2018. We use the UK data to train the word embedding, but use data from a total of 17 European countries to analyse skill mismatch. We use data from 2019, the first full year available, and extract the universe of OJAs collected by the data production system. In total, our skill demand analyses rely on 17,966,812 observations for 17 European countries in the year 2019.

2.3. PIAAC

On the skill supply side, we use survey data from the Programme for International Assessment of Adult Competencies (PIAAC). PIAAC is a survey administered by the Organisation for Economic Co-operation and Development (OECD) and comprises representative samples of working-age individuals (16-65 years) from 33 OECD countries, including 17 European countries: Belgium, Cyprus, Czech Republic, Denmark, France, Germany, Greece, Ireland, Italy, Lithuania, Netherlands, Poland, Slovak Republic, Slovenia, Spain, Sweden, and the United Kingdom (UK). Data were collected in two rounds between 2011 and 2015.[3] In total, the PIAAC sample comprises 250,000 observations, with sample sizes typically ranging from 4000-8000 observations in each country.

PIAAC contains information on key cognitive and workplace skills of adults. PIAAC assesses cognitive skills in literacy, numeracy, and problem-solving in technology-rich environments using tests.[4] For the purpose of our analysis, we focus on skill use at the

[3] A third round of data collection took place in 2017, which included Hungary as the only European country. However, to ensure that the skill data are not too far apart in time, we refrain from using round 3 data in most of our analysis. We use Hungary in our investigation of skill mismatch at the extensive margin in Section 5.

[4] These data have been used to estimate returns to cognitive skills across countries (e.g., Hanushek et al., 2015; Falck et al., 2021).

workplace and personal attitudes elicited in sections F through I of the PIAAC background questionnaire. For instance, subjects are asked to which extent they, e.g., teach people, use calculators, or read financial statements.[5] In total, the PIAAC survey contains 73 items eliciting skill use at work. Moreover, the PIAAC data contain an array of background characteristics, such as respondents' gender, level of education, field of study, occupation at the ISCO 2-digit level, and industry of employment at the ISIC 1-digit level.

One limitation of the PIAAC survey for the purpose of our analysis is its timing. The first PIAAC cycle took place in the period 2012-2017 in three rounds (2012, 2015, 2017), while (with the exception of the United States) each country participated only in one round.[6] The countries used in our analysis mostly come from round 1 (2012), with the exception of Greece, Lithuania and Slovenia (round 2, 2014).We use Hungary, which participated in round 3 of PIAAC in 2017, in our investigation of skill mismatch at the extensive margin in Section 5. However, the earliest job advertisement data from CEDEFOP stem from 2019. To remedy this issue of temporal misalignment between PIAAC and OJA, we use the skill change in the United States, for which the PIAAC survey has been conducted twice, in 2012 and 2017. We use the U.S. data to project skill changes for all countries in our sample. Assuming that changes in the occupational skill content in the United States represent changes at the technological frontier (Caunedo et al. (2021)), these changes can be used to project an upper bound of how skills have evolved in other PIAAC countries.[7]

2.4. European Tertiary Education Roster (ETER)

The European Tertiary Education Roster (ETER) collects information from the near universe of higher education institutions (HEI) in European countries from 2011-2019. Inter alia, ETER provides information on the number of graduates by field of study and qualification level in each year at the HEI level. Moreover, the data contain information on the region (NUTS2) of the respective higher education institution, which allows us to

[5]The response scale differs by the type of variable that is considered. In the cases relevant to us, there are three relevant scales. The first, which is used to measure frequency in time units, is a 5-point Likert scale (1 - Never, 2 - Less than once a month, 3 - Less than once a week but at least once a month, 4 - At least once a week but not every day, 5 - Every day). The second scale is used to measure the extent to which the respondents think the statement represents them, again measured on a 5-point Likert scale (1 - Not at all, 2 - Very little, 3 - To some extent, 4 - To a high extent, 5 - To a very high extent), and the third scale is simply binary (Yes/No).

[6]Data of cycle 2 of PIAAC will become available only in 2024, and thus after the end of the PILLARS project.

[7]Our results are qualitatively similar without this U.S.-based skill adjustment.

calculate the number of graduates by field of study and NUTS2 region in Europe. We will exploit this information to construct a measure of skill shortage at the extensive margin in European regions by comparing the number of individuals needed in each occupation from OJA data to the number of graduates that are potentially suited to work in the respective occupation from ETER (see Section 3.3). For the year 2019, ETER covers 2,620 higher education institutions in 305 different NUTS2 regions in Europe.

2.5. European Labor Force Survey (EU-LFS)

The European Labor Force Survey (EU-LFS) is a representative household survey conducted in all member states of the European Union and the United Kingdom covering the years 1983-2020. The EU-LFS provides data on labor force participation of individuals aged 15 and above, as well as information on the occupation (ISCO2) and industry (NACE2) of employment, and the geographic region of the workplace (NUTS2). We use EU-LFS to obtain the industrial and occupational composition in European regions, and to obtain labor market measures at the regional level.

We use the data described in this section to construct novel measures of skill mismatch. The first step of our analysis, presented in Section 3.1, requires the whole set of skills provided in the ESCO skill classification and a list of selected items of the PIAAC questionnaire as input taxonomies, as well as the OJA data to train embedding algorithms to provide the context for the evaluation of similarity among the two taxonomies.[8] The second step uses the PIAAC to ESCO link to compare PIAAC questions and OJA skills to calculate measures of skill mismatch (Section 3.2 and Section 3.3).

3. Methodology

We propose a novel measure of the gap between the skills demanded by employers and the skills provided by workers. On the skill demand side, we rely on OJA data from CEDEFOP, collected in 2019, to capture skills required by employers. On the skill supply side, we use PIAAC data. However, these two data sources are not directly linked at the skill level. To link the data, we develop an Artificial Intelligence (AI)-driven tool using word embeddings that maps skills mentioned in OJA to the skills elicited in PIAAC. We present the method, its implementation, and validation in Section 3.1, while relegating the details and technical descriptions to Appendix A. In Section 3.2, we apply our linkage

[8] See Table A.1 for an overview of PIAAC items with at least one corresponding ESCO skill.

Figure 1: Mapping Process between PIAAC and ESCO

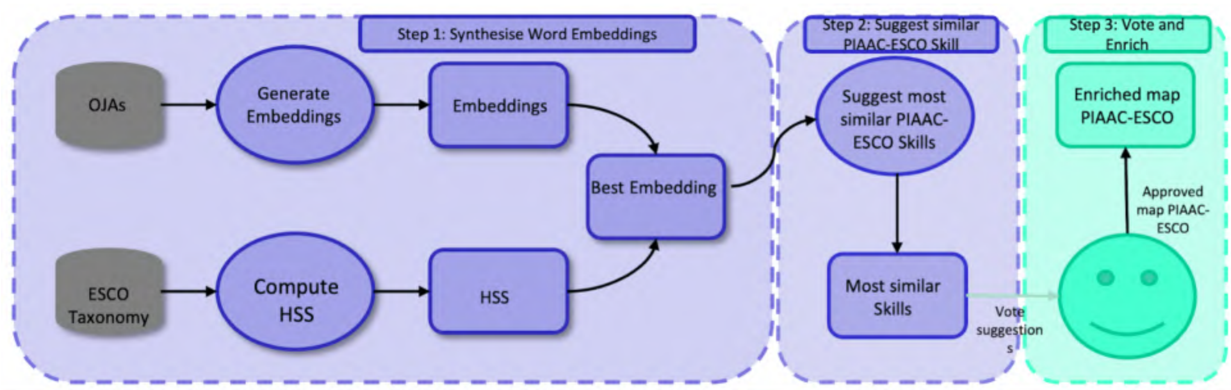

Notes: Figure provides a graphical overview of the process that generates the mapping between PIAAC and ESCO. HSS is short for Hierarchical Semantic Similarity.

to develop a quantification of the skill mismatch (i.e., skill demand minus skill supply) at the *intensive margin*. To complement our intensive-margin measure, Section 3.3 develops a measure of skill mismatch at the *extensive margin*.

3.1. Linking PIAAC and ESCO Data

This section describes the global approach to construct the crosswalk between the PIAAC background questionnaire and the ESCO skill taxonomy. The process is depicted in Figure 1.

The first step in creating our skill mismatch measure is to train a word embedding model and suggest possible alignments between each PIAAC skill item to skills at the third level of the ESCO Skill Pillars, the leaf concepts of our hierarchy. This set of paired PIAAC question - ESCO skill provides the domain experts with a number of possible suggestions for the skill taxonomy alignment, which would otherwise have been done manually in the second step. The procedure follows the work done by Giabelli et al. (2022b) and applies it to the context of PIAAC and ESCO.

The main goal of the first step is to introduce a vector representation of taxonomic terms that represents the similarity of words within the taxonomy in the best possible way. To accomplish this, we perform three distinct tasks. First, we generate word embeddings through the state-of-the-art method FastText (Bojanowski et al., 2017). This word embedding method considers sub-word information and can deal with out-of-vocabulary words (Giabelli et al., 2022b). Following Baroni et al. (2014), we perform an intrinsic evaluation to select the best embedding model. The authors select the word vectors model with the maximum correlation between its cosine similarity and a benchmark semantic similarity value and use a handcrafted data set of pairwise semantic similarity between common words as the gold benchmark. However, those resources usually have low coverage, especially when used in specific domains such as those related to the labor market. For this reason, we resort to a measure of semantic similarity in taxonomies developed by Malandri et al. (2021), which measures semantic similarity in a taxonomy based on the structure of the hierarchy itself without using any external resource; thus, in a sense, preserving the semantic similarity intrinsic to the taxonomy. The latter method, Hierarchical Semantic Similarity (HSS), has proven to help select embeddings in several contexts, e.g., taxonomy enrichment (Giabelli et al., 2021a) and job-skill mismatch analysis in the labor market Giabelli et al. (2021b). The output of this method is a constrained number of suggestions to the domain experts to simplify their work of skill taxonomy alignment that otherwise would be all manual. A manual alignment would imply an expert to match each of the almost 13'500 possible ESCO skills to each PIAAC question of interest, a task with high complexity and high cost. Moreover, in order to have a valid mapping, validation would be required, having more than one person perform the task and setting a rule to resolve conflicts in the matched skills. Our procedure, therefore, supports the experts by narrowing down the number of possible alternative ESCO skills matches per PIAAC question. The last step of our method requires the involved experts to vote if and to what extent the ESCO skills suggestions were relevant and consistent with the PIAAC questions, using a 5-point Likert scale (ranging from 1 (low agreement) to 5 (high agreement)).

Figure 2 reports the results of the validation exercise, which involved three experts of the PILLARS consortium. Vote scores are concentrated in the upper part of the graph for most skills, suggesting the experts' high level of agreement with the results of the word embedding model.The final subset of PIAAC question - ESCO skill pairs is made of pairs with a score of at least 4 for all experts.

3.2. Measuring Skill Mismatch at the Intensive Margin

The mapping of skills in OJA to skill items in PIAAC allows us to develop a quantification of *which and by how much* skills are lacking in an occupation, i.e., a measure of skill shortage at the *intensive margin.* To quantify the gap between skills on the demand side from OJA data and the skills on the supply side from PIAAC, we develop a novel measure of skill shortage based on the importance of each skill in each occupation. Our measure is based on the concept of revealed comparative advantage (RCA), originally developed in international trade economics to describe countries' export specialization (e.g., Balassa, 1964). The measure applied in the context of occupations and skills can be understood

Figure 2: Expert Validation of the PIAAC to ESCO Mapping

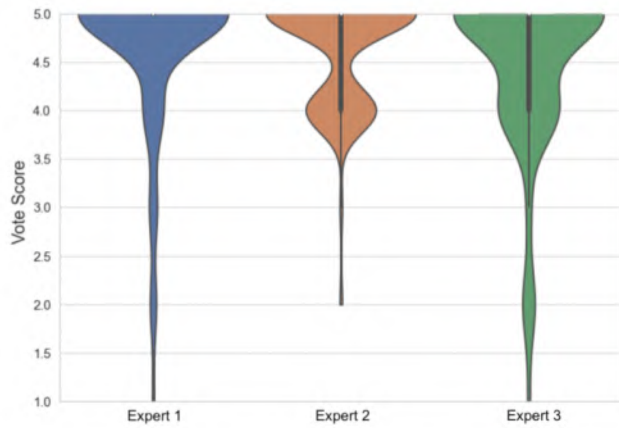

Notes: Figure shows the distribution of votes for all validated skills separately for each expert rating. Votes are expressed on a Likert scale with 5 indicating highest agreement with the mapping and 1 indicating lowest agreement.

as the relevance of a skill in an occupation relative to all other occupations and to all other skills in the same occupation. We follow the approach developed by Alabdulkareem et al. (2018), who calculate the RCA using the O*NET dictionary of occupations and skills, which surveys a sample of workers in the Unites States to assess the relevance of a skill for each occupation.[9] The ESCO skill classification does not provide any indicator of relevance of the skill in each occupation. Following Giabelli et al. (2022a), we measure the skill relevance by referring to the occurrence of skills in OJAs. In particular, the skill relevance is computed as the frequency of a given skill in the job ads of a specific occupation relative to the skill's frequency in job ads in all other occupations. Analogously, the RCA of the skill in PIAAC is computed as the frequency of the use of this skill among survey respondents in a given occupation relative to this skill's use in all other occupations.[10]

Analytically, we define the skill frequency, *sf*, as the number of observations in an occupation that use the skill of interest divided by the total number of observations in

[9] See `https://www.onetcenter.org/` for details on O*NET.

[10] As explained in Section 2, PIAAC answers are measured with different scales, depending on the context. When it comes to skill use, both Likert scales and dichotomous scales are used. To dichotomise the Likert scale in PIAAC, we set cutoffs at which we define that a skill is used by a worker (i.e., the respective PIAAC item takes a value of 2 or more on the 5-point Likert scale). This way, we encode each PIAAC question as an indicator of the skill being used on the job by workers at all. This step allows having an analogous structure for PIAAC and OJA data as in both cases we can identify at the occupation level the cases in which the skill is used or demanded. However, our results are robust to other choices of the cutoff value.

that occupation. Given a set of occupations $\bar{O} = \{o_k, k = 1, \ldots, m\}$ and a set of skills $\bar{S} = \{s_j, j = 1, \ldots, p\}$, we define the skill frequency as:

$$sf(o_k, s_j) = \frac{\sum_{i=1}^{n} I(o_i = o_k) \cdot I(s_i = s_j)}{\sum_{i=1}^{n} I(o_i = o_k)} \tag{1}$$

where I denotes the indicator function and $\sum_{i=1}^{n} I(o_i = o_k) \cdot I(s_i = s_j)$ the count of the occurrences of skill s_j within the occupation o_k. The term $\sum_{i=1}^{n} I(o_i = o_k)$ refers to the total number of observations in occupation o_k in either PIAAC or OJA.

Iterating over skills and occupations, we obtain a matrix $\underset{m \times p}{M}$ of the skill frequency for each pair of occupations $o_k \in \bar{O}$ and skills $s_j \in \bar{S}$. The revealed comparative advantage, *rca*, for o_i and s_l is defined as:

$$rca(o_i, s_l) = \frac{sf(o_i, s_l) / \sum_{j=1}^{p} sf(o_i, s_j)}{\sum_{k=1}^{m} sf(o_k, s_l) / \sum_{k=1}^{m} \sum_{j=1}^{p} sf(o_k, s_j)} \tag{2}$$

which ranges between $[0, +\infty)$.

To measure the gap between skills on the demand side from OJA data and the skills provided on the supply side in PIAAC, we make use of the RCA measures described above in the following way. First, we calculate the RCA of each skill for each occupation in both OJA and PIAAC. Thus, for OJA and PIAAC, we calculate the RCAs according to equation 2. Next, for each skill, we compute the rank of the RCA across all occupations on both the demand and supply side. For instance, if the RCA of "Teaching people" is ranked at the top five percent for teaching professionals on the supply side, this means that teachers use this skill very intensively as compared to individuals in other occupations and to all other skills in the teaching professionals occupation.

Thus, differences in the RCA ranks of skills between demand and supply reflect potential mismatches in skill relevance.[11] We define the gap in skill ranks as the percentile rank on the demand side minus the percentile rank on the supply side. Thus, a positive value indicates that the RCA in skill demand is larger than the RCA of skill supply in a particular occupation. Since positive values of our skill mismatch measure indicate that a specific skill is more relevant on the demand than on the supply side, we frequently refer to this case as *skill shortage*. For instance, the importance of "Reading financial statements" is at the 61st percentile on the demand side for teaching professionals, while it is at the 43rd percentile on the supply side, resulting in a positive skill gap of 18 percentile

[11] See, e.g., Wyness and Murphy (2020) for a similar measure of mismatch between students and colleges.

ranks. Thus, teaching professionals have a skill shortage in the skill of “Reading financial statements”. In contrast, a negative skill gap indicates a *skill surplus*. For instance, for teaching professionals, the skill “Reading e-mails” ranks at the 68th percentile on the demand side and at the 82nd percentile on the supply side, resulting in a negative skill gap of −14 percentile ranks (i.e., skill surplus).

3.3. Measuring Skill Mismatch at the Extensive Margin

Our measure of skill shortage at the intensive margin provides a measure of the gap between the skills workers in an occupation possess and the skills employers request. At the same time, there might be a mismatch between the number of individuals who are potentially qualified to be employed in these occupations and the number of workers in these occupations demanded by employers. Thus, we refer to skill shortage (skill surplus) at the extensive margin if the number of vacancies for a given occupation exceeds (falls short of) the number of suited candidates in a region.

To quantify skill shortages at the extensive margin in European regions, we draw on data from the European Tertiary Education Roster (ETER) providing the number of university graduates by field of study and region. We exploit this information to construct a measure of skill mismatch at the extensive margin by comparing the number of individuals requested by firms in a given occupation to the number of graduates that are potentially suited to work in the respective occupation.

To infer the potential supply of workers in an occupation, we obtain country-specific occupational distributions of individuals with a given field of study in tertiary education from PIAAC.[12] Thus, we obtain the share of individuals working as, say, business administration professionals or teaching professionals given that they have completed a business degree during university. For instance, assume that 90 percent of individuals with a business degree are employed as business administration professionals, while the remaining 10 percent work as teachers. Accordingly, if there are 100 business graduates in 2019 in a region and given these occupational distributions from PIAAC, there are 90 individuals that will potentially find employment in business administration professionals occupations, while 10 individuals will potentially work as teaching professionals. Thus, variation in the potential supply of workers in a given occupation between regions is driven by the regional variation in the number and field of study of university graduates.

[12] PIAAC provides information about field of study in seven categories: General programmes; teacher training and education science; humanities, languages and arts; social sciences, business and law; science, mathematics and computing, health and welfare; services

More technically, we calculate the supply of individuals suited to work in occupation i in NUTS2 region r, $Supply_{ir}$, as:

$$Supply_{ir} = \sum_{j=1}^{J} Graduates_{jr} \times Pr(Occ_i|Field_j), \tag{3}$$

where $Graduates_{jr}$ is the number of graduates in field of study j in region r obtained from ETER, and $Pr(Occ_i|Field_j)$ is the country-specific probability of individuals with a degree in field j to work in occupation i obtained from PIAAC. Employers' demand for the respective occupation in a given region is simply the number of vacancies in the occupation-region pair from OJA data. Skill mismatch at the extensive margin for an occupation is the difference between demand and supply, divided by demand, at the regional level – i.e., the share of vacancies in an occupation that cannot be filled by the regional graduate pool.[13]

While it would be possible in principle to construct skill mismatch at the extensive margin for entire regions (i.e., aggregating over all occupations in a region), we refrain from doing so because the ETER data only provide the number of *university* graduates. Thus, our measure only accurately depicts skill mismatch in occupations with a higher share of workers with tertiary education. Below, we will show skill mismatch at the extensive margin for selected high-skilled occupations across Europe (Section 5).

4. Skill Mismatch at the Intensive Margin

In the following, we provide a quantification of skill demand, skill supply, and skill shortage in Europe. First, we consider the intensive margin. We start by investigating skill demand, skill supply, and skill shortage by occupation.[14] We also provide an analysis by gender to see whether gender differences in education or occupational choices lead to differences in skill mismatch. Further, we investigate how European regions differ in the degree of skill shortage and explore potential determinants and outcomes of skill shortage at the regional level. Second, we complement our regional analysis of intensive-margin skill

[13]Note that not all vacancies are suited for early-career workers as they require certain occupational experience or labor market tenure. Our measure of skill mismatch at the extensive margin implicitly assumes that there are no systematic differences across European regions in the propensity that university graduates are regarded as suited by the employers to fill an OJA.

[14]In Appendix B, we aggregate our measures of mismatch at the industry level to obtain a picture of skill mismatch in different industries in European countries.

shortage with a depiction of skill shortage at the extensive margin for selected occupations, showing which regions are most affected by excess demand in these occupations.

4.1. Skill Shortage by Occupation

Figure 3 shows the overall demand for skills (upper panel) and the supply of skills (lower panel) for each ISCO2 occupation. For instance, the average skill demand in the occupation "Business Associates" amounts to the 32nd percentile. This indicates that on average, the RCA or relative importance of the skills demanded in this occupation ranks at the 32nd percentile across all occupations. Skill supply in the business associates occupation is higher at the 74th percentile. This suggest that on average across all skills, the RCA or relative specialisation of the skills supplied by business associates ranks at the 74th percentile. Figure 4 provides our measure of skill shortage, defined as the difference between skill demand and skill supply in an occupation. Thus, business associates are characterized by a skill surplus (negative skill shortage) of 42 percentiles.

The literature often distinguishes between four occupation types (Autor et al., 2003): manual routine, manual non-routine, cognitive non-routine, and cognitive routine occupations. These occupations differ in the tasks workers need to perform on the job. For instance, food preparation assistants perform predominantly manual, routine-intensive tasks, such as manual assembling and quality checks. On the other hand, teaching professionals perform predominantly cognitive and non-routine tasks, such as using advanced mathematics and teaching people. At the same time, structural transformation and technological change have different impacts on different types of tasks. Automation technologies have particularly rendered codifiable routine and manual tasks susceptible to substitution by automation. As the task composition and thus the skill requirements of different occupations are affected differently by technological change, this also renders occupations more or less susceptible to changing skills demands and skill shortages.

Figure 4 shows that there is an intriguing difference between manual and cognitive workers in the average skill shortage. While cognitive non-routine and cognitive routine workers have a skill surplus (negative shortage) on average, manual non-routine and manual routine workers typically exhibit a (positive) skill shortage. For instance, business associates exhibit the highest skill surplus: they provide more skills than are required in respective job ads. For this occupation, the relevance of the skills demanded ranks below the skills provided by business administration professionals by 42 percentiles. Food preparation assistants, on the other hand, show the most pronounced skill shortage: The skill requirements in this occupation exceed the skill supply by 31 percentiles. Health professionals, electrical workers, ICT professionals, and ICT technicians feature the narrowest

Figure 3: Skill Demand and Skill Supply by Occupation

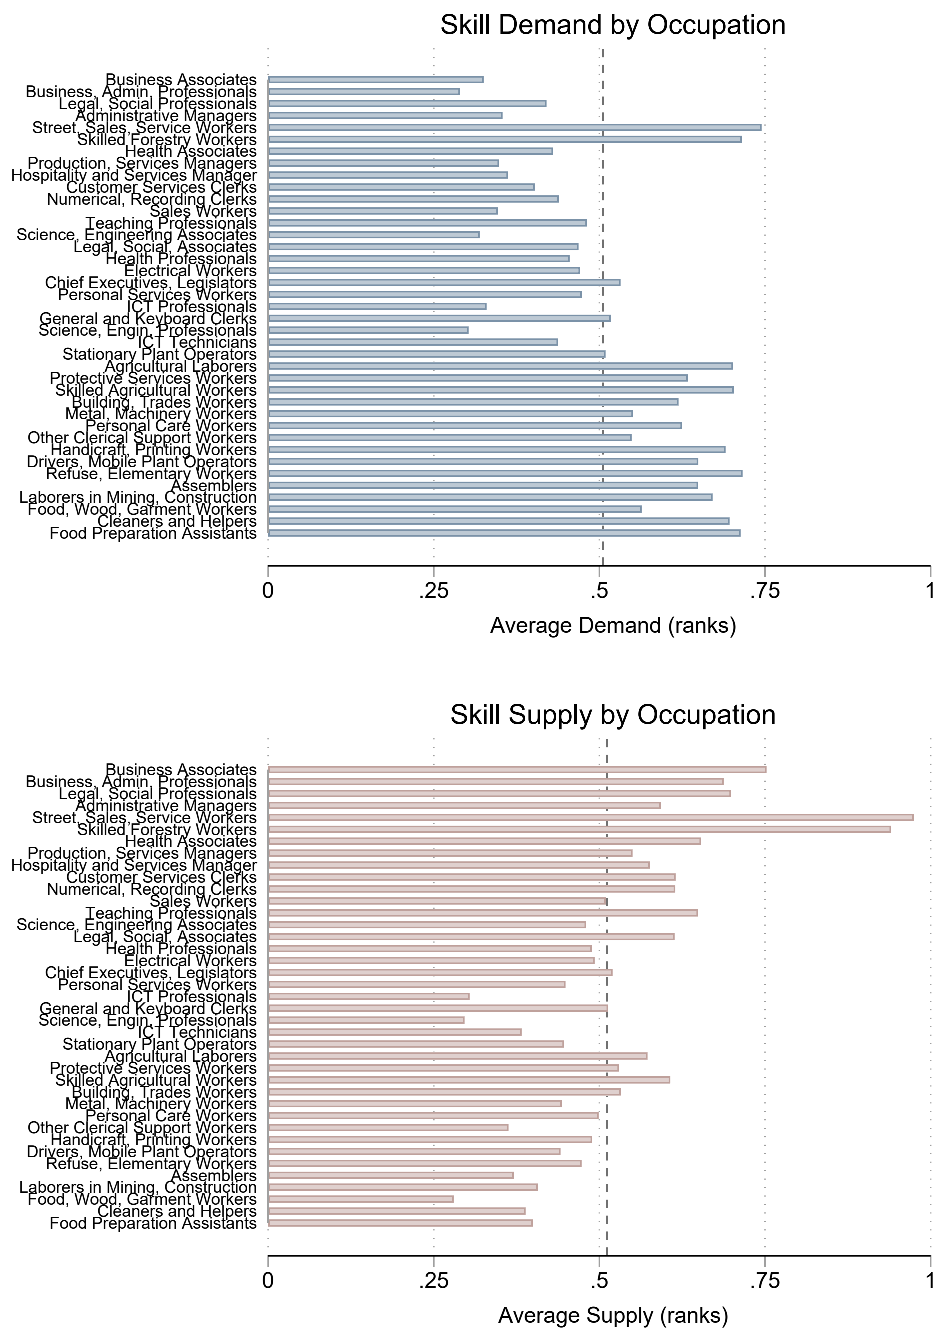

Notes: The figure shows the skill demand and skill supply by occupation and pooled for 17 European countries. Data are from CEDEFOP and PIAAC.

Figure 4

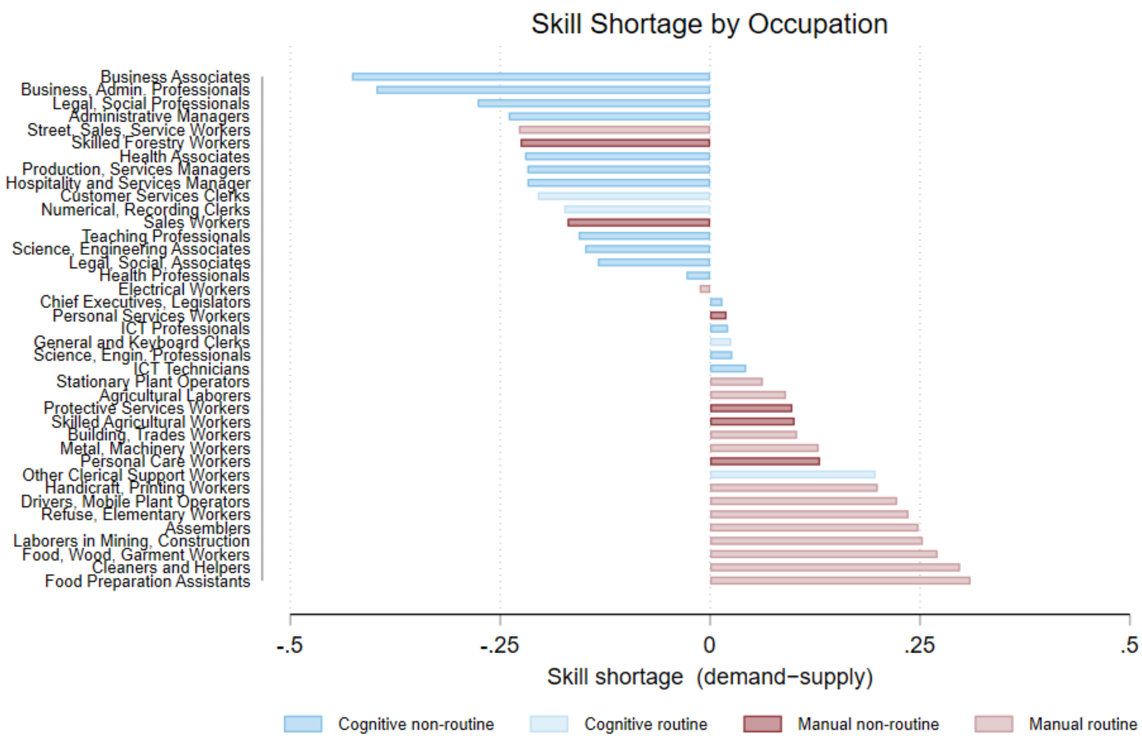

Notes: The figure shows the skill shortage by occupation, separately for each occupation type and pooled for 17 European countries. Data are from CEDEFOP and PIAAC.

skill gaps; i.e., the skills required in these occupations closely match the skills workers actually possess. In Section 4.4, we discuss potential determinants and mechanisms underlying these patterns across occupations, such as the risk of automation and employees' participation in on-the-job training.

4.2. Skill Shortage by Occupation Type, Skill Domain, and Gender

One potential driver of the skill shortage of manual workers could be a lack of specific skills that have gained importance for employers in recent years, such as social skills or digital skills. Indeed, several recent studies suggest substantial labor market returns to certain types of skills, such as numeracy skills (e.g., Hanushek et al., 2015, 2017) and digital skills (e.g. Falck et al., 2021). Other research indicates the economic importance of social skills. For instance, Deming (2017) shows that the fastest-growing occupations in the U.S. require a high level of communication and teamwork. Weidmann and Deming (2021) provide evidence for the value of social skills in team production. Piopiunik et al. (2020) show experimentally that HR managers value social skills (e.g., signaled by social volunteering) when making hiring decisions. In an RCT of a vocational training program in Colombia, Barrera-Osorio et al. (2020) find that technical skills and social skills increase formal employment.

In Figure 5, we investigate skill demand, supply, and shortage in each occupation group for the following skill domains: digital, numeracy, literacy, and social skills.[15] Intriguingly, the pattern of skill mismatch for the four occupation types is very similar across all skill domains: in each domain, cognitive workers show a skill surplus on average, while manual workers exhibit a shortage on average. Further, the skill shortage is largest for manual routine workers in each skill domain, while the surplus is largest for cognitive non-routine workers in each domain.

Skill mismatch could also differ by gender, for instance, due to gender-specific education trajectories, occupational choice, or heterogeneous mobility potential. Figure 6 thus provides the skill shortage analysis shown in Figure 5 separately for men and women. As PIAAC contains information on gender, we split the sample by gender and construct gender-specific RCA measures and RCA ranks on the supply side for each skill and occupation by calculating these quantities separately by gender. While we can observe the gender of PIAAC respondents, we have no way of knowing whether a job ad is targeted at men or women (officially, they will always be targeted at both). Thus, we subtract

[15]Note that in the PILLARS application, we refer to these skills as digital skills, hard (non-digital) skills (e.g., numeracy and literacy skills), and soft skills (i.e., social skills).

Figure 5: Skill Demand, Supply, and Shortage by Skill Domain and Occupation Type

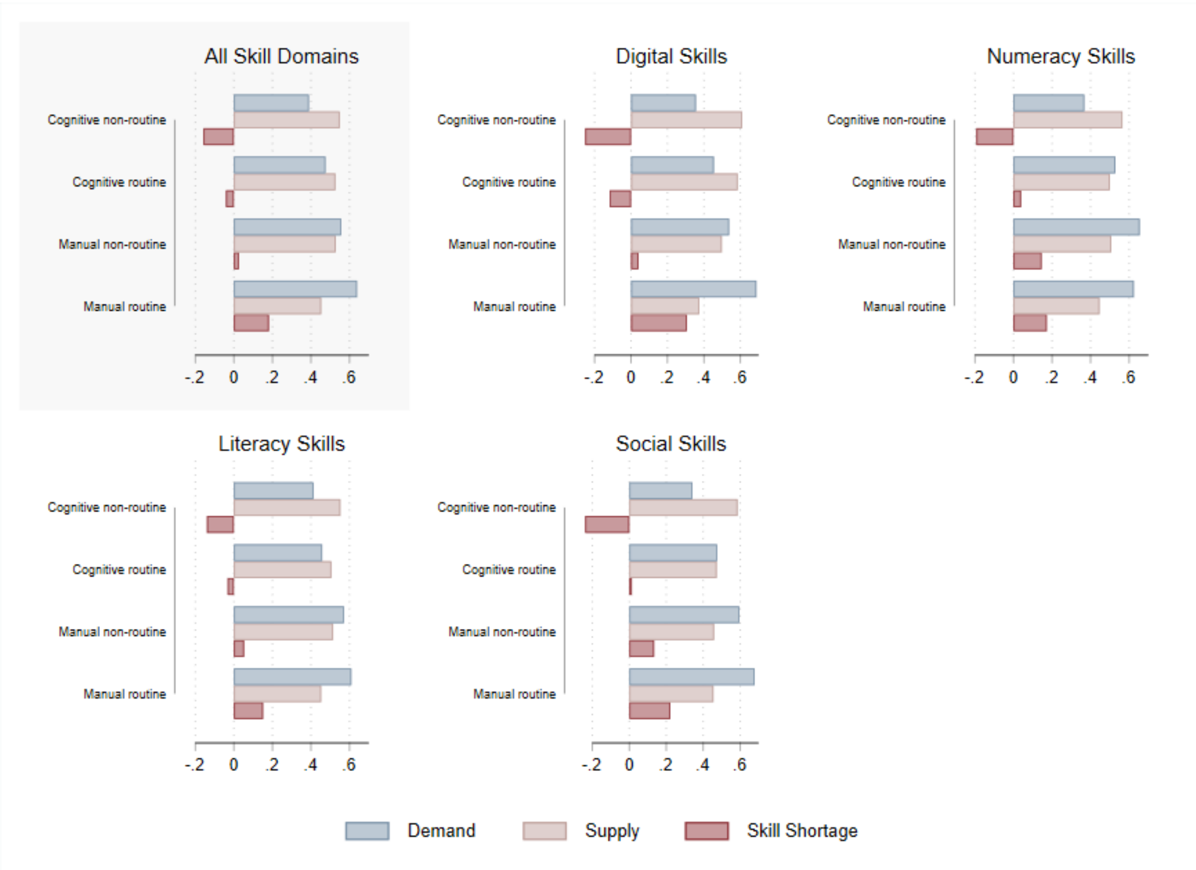

Notes: Pooled skill demand, supply, and shortage by skill domain (digital, numeracy, literacy, and social) and occupation type (cognitive non-routine, cognitive routine, manual non-routine, and manual routine) for 17 European countries. Data are from CEDEFOP and PIAAC.

Figure 6: Gender-Specific Skill Shortage

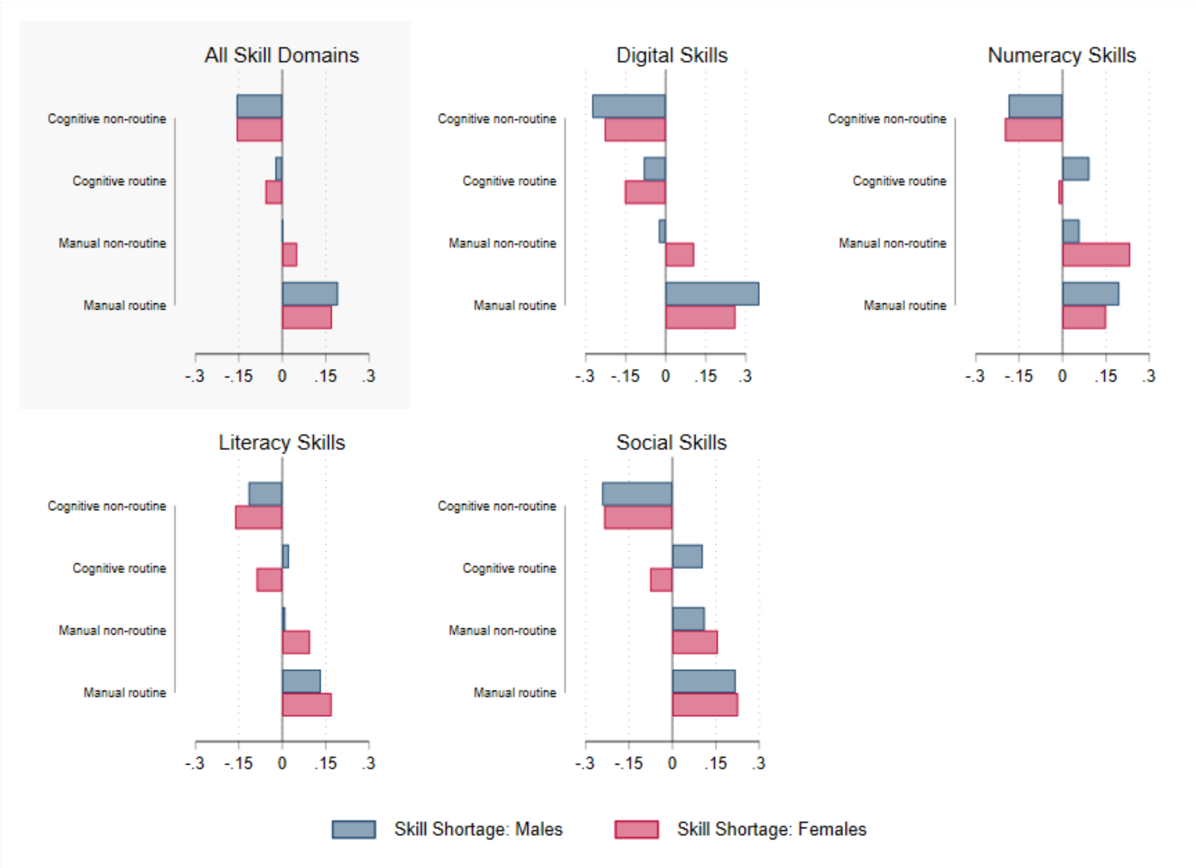

Notes: The figure shows skill shortage separately for men and women. Skill shortage is further depicted by skill domain (digital, numeracy, literacy, and social) and occupation type (cognitive non-routine, cognitive routine, manual non-routine, and manual routine) for 17 European countries. Data are from CEDEFOP and PIAAC.

the gender-specific skill supply from the pooled skill demand in all of our job ads. Skill shortage in Figure 6 generally follows the same patterns for males and females across all skill domains: Cognitive non-routine jobs show skill surpluses, while manual routine jobs are characterized by skill shortages for both men and women. However, there are some noteworthy deviations from the overall patterns in certain skill domains: For instance, female cognitive routine workers exhibit skill surpluses in numeracy, literacy, and social skills, whereas male cognitive routine workers face skill shortages in these domains. The magnitude of mismatch also differs between men and women in certain skill domains (e.g., for manual non-routine workers in the domain of numeracy skills).

Figure 7: Skill Shortage across European Regions

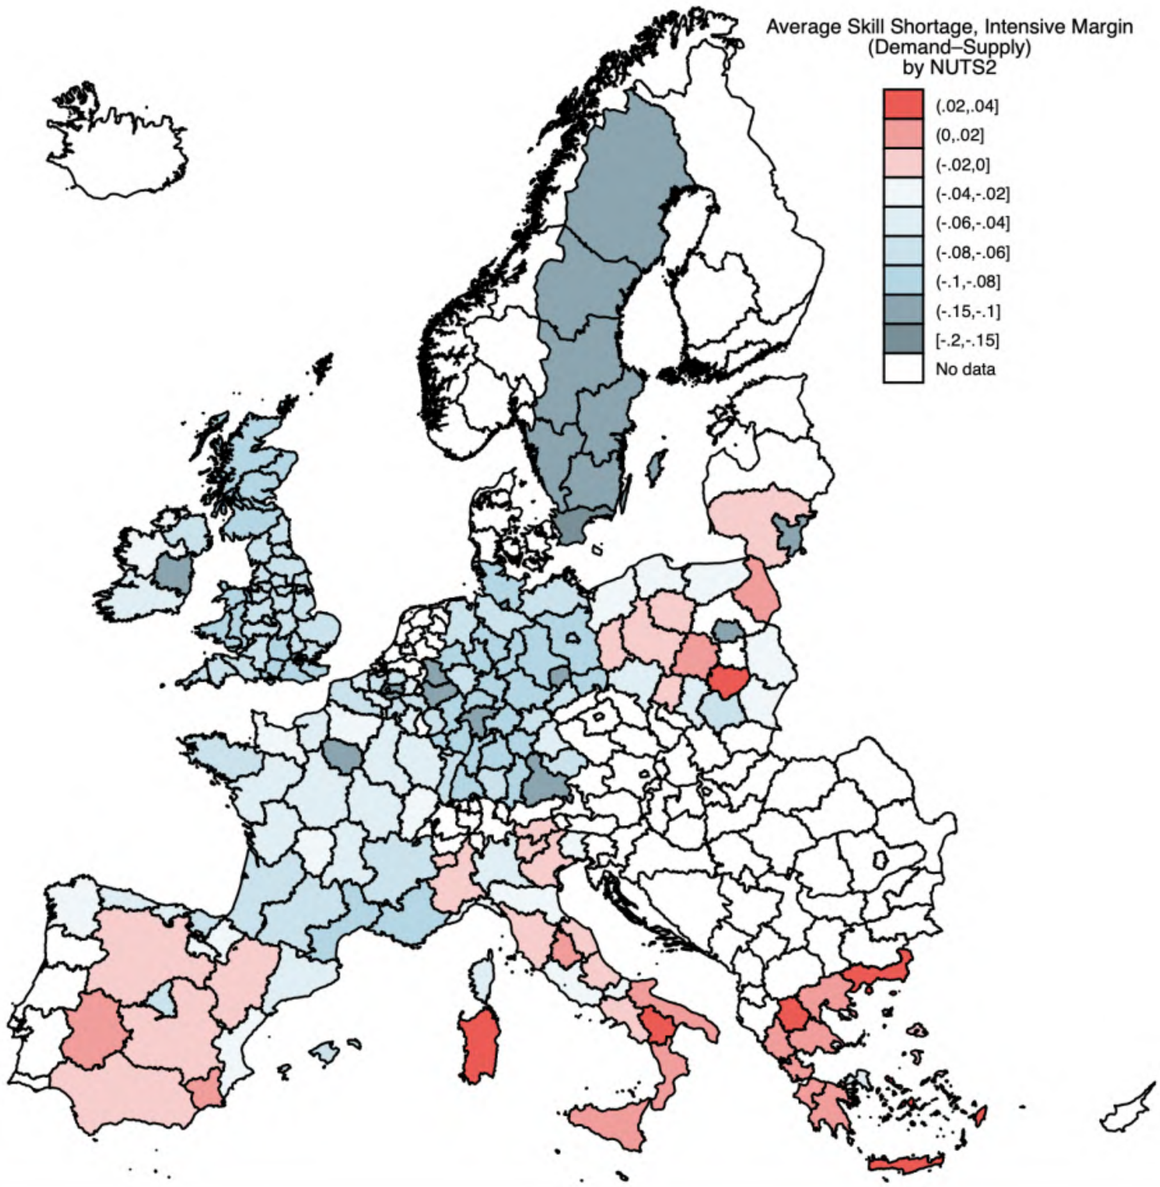

Notes: Figure shows the average intensive-margin skill shortage, weighted by occupation shares, for NUTS2 regions in Europe. Data are from CEDEFOP and PIAAC.

4.3. Skill Shortage by Region

Skill shortage may not only differ by workers' characteristics, but also geographically. To analyze which regions are most affected by skill shortages, we now provide a depiction of skill shortage at the intensive margin across European regions.

Figure 7 provides the occupation-weighted average skill shortage (intensive margin) for NUTS2 regions in Europe. Specifically, we obtain employment shares for each ISCO 2-digit occupation at the NUTS2 level from the EU-LFS in 2019. We obtain a region-level skill shortage measure by aggregating the occupation-country-specific shortage measure over all occupations in a region, using the occupation's employment share in the region as weight.

Figure 7 shows considerable variation in the average extent of skill shortage across European NUTS2 regions: The skill shortage ranges from –23 percentile ranks in the Stockholm Metropolitan region in Sweden (i.e., abundant skill supply) to a skill shortage of 4 percentile ranks in the Voivodeship of Mazovia in Poland. While the average skill shortage across European regions is –5.4 percentile ranks (i.e., skill surplus on average), several regions in Greece, Italy, Spain, and Poland are characterized by skill shortages. This suggests that workers in countries such as Greece, Italy, Spain, and Poland are often not equipped with the right skills do meet the skill demand from the labor market.

4.4. Determinants of Skill Mismatch

4.4.1. Regional Level

After having investigated skill shortage at the intensive margin from different angles and in different dimensions, we now explore potential mechanisms that might drive the shortage.

Table 1 shows the association of skill shortage at the regional level with regional characteristics potentially associated with skill shortage, such as a region's industry structure and economic performance. In particular, we investigate the role of regional characteristics measured in 2010, and thus almost 10 years before skill shortage (measured in 2019). Table 1 reports coefficients from OLS regressions with country fixed effects, i.e., relationships that are purged of unobserved country-specific factors that affect both regional characteristics and skill shortage. Standard errors are clustered at the country level.

Column 1 of Table 1 suggests a positive and significant association between the employment share in the manufacturing sector in 2010 (obtained from EU-LFS) and skill shortage in 2019. In particular, a one percentage point increase in the manufacturing employment share is associated with an increase in shortage of 0.17 percentile ranks. This is approximately the difference in skill shortage between the regions Nordjylland in Denmark and Limousin in France. As the manufacturing sector has experienced the largest decline in the labor share over the last decades (e.g., Alvarez-Cuadrado et al., 2018), the labor and skill inputs have also likely changed rapidly over time. Thus, regions with a

Table 1: Determinants of Skill Shortage in Europe

	(1)	(2)	(3)	(4)	(5)
Manufacturing share 2010	0.165**				
	(0.0671)				
Automation exposure		2.212***			
		(0.264)			
KIBS share 2010			-0.455**		
			(0.155)		
log GDP per capita 2010				-6.398***	
				(1.597)	
Import share 2010					0.161
					(0.204)
Country Fixed Effects	Yes	Yes	Yes	Yes	Yes
Observations (Regions)	148	145	148	141	148
Countries	17	17	17	17	17

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Dependent variable: skill shortage (intensive margin) in 2019. Skill shortages is calculated as an occupation-weighted average at the NUTS2-level. Manufacturing share and KIBS share are obtained from EU-LFS. GDP per capita at the NUTS2 level is obtained from EUROSTAT. Import share is obtained from UN-Comtrade and EU-LFS; it is calculated as industry employment weighted import share over the total trade volume. Automation exposure is obtained from Nedelkoska and Quintini (2018) based on PIAAC; it is weighted by occupational employment (obtained from EU-LFS data) to arrive at a region-level measure. Standard errors clustered at the country level in parentheses.

higher manufacturing share might have not been able to keep up with this change in manufacturing production in terms of the skills of the workforce.

In a similar vein, the manufacturing sector has traditionally been employing high shares of manual and routine workers (e.g., OECD, 2019). At the same time, these occupational groups are most exposed task displacement (Acemoglu and Restrepo, 2019), and skill requirements have changed over time. Thus, workers might have not been able to keep up with changing skill demands. Indeed, we observe that regions more affected by automation exhibit a higher shortages. In Column 2 of Table 1, we construct the average exposure to automation risk at the regional level by weighing country- and occupation-specific measures of automation exposure by (Nedelkoska and Quintini, 2018) with the occupational composition at the NUTS2 level from EU-LFS. We find a significantly positive association between the average skill shortage in a region and the automation exposure of its workforce. We will explore this aspect of skill shortage and automation risk in more detail below.

Column 3 of Table 1 suggests that more dynamic, innovative regions suffer less from skill shortage. A one percentage point increase in the employment share in knowledge intensive business service (KIBS) industries in a region in 2010 is related to a 0.46 percentile decrease in the average skill shortage. This approximately corresponds to the difference in skill shortage between the regions Emilia-Romagna in Northern Italy and La Rioja in central Spain. Relative to stagnant manufacturing wages, returns to skilled labor and wages of high-skilled knowledge workers have increased over the last decades (e.g., Acemoglu and Autor, 2011). Thus workers in regions with a high KIBS share might have had incentives to invest in skills, resulting in a lower skill shortage today.

Similarly, Column 4 of Table 1 suggests that economically stronger regions are less affected by skill shortage. There is a significantly negative relationship between skill shortage and past economic performance, as measured by the purchasing power adjusted log GDP per capita in 2010 (obtained from EUROSTAT). A one percent higher GDP per capita in 2010 is associated with a 0.064 percentile decrease in skill shortage. For instance, the difference in past economic performance between, e.g., the Stuttgart region in Germany and Andalusia in Spain of 52 percent accounts for more than half of the difference in skill shortage between both regions (5.9 percentiles). Hanushek et al. (2017) provide a potential explanation why regions that have performed well in the past show less skill shortage today. They show that returns to skills are systematically larger in countries that have grown more rapidly in the past. Thus, regions that performed well economically in the past might have had higher returns to skills and thus higher incentives for its workforce to invest in skills.

Similarly to a region's manufacturing share, import competition might change the labor and skill inputs required, as local industries are forced to change their position in the global value chain and the type of intermediary goods produced (also see PILLARS Work Package 1.3). Thus, regions that are more exposed to import competition in the past might show higher skill shortages today if labor inputs could not keep up with the transformation of global value chains. In Column 5 of Table 1, we investigate the role of import competition for skill shortages at the regional level. To this end, we obtain data on import and export volumes at the industry level for each country from UN-Comtrade and calculate the import share as imports over the total absorption in each industry for each country in our sample. We follow the literature (e.g., Autor et al., 2013; Dauth et al., 2017) and construct measures of import exposure at the regional level by weighting industry specific import shares by the industry composition at the regional level. However, while the point estimate does suggest that regions with higher import competition suffer from higher skill shortages, the coefficient is not statistically significant.

We interpret our results regarding the determinants of skill shortage as suggesting that exposure to technology might be a more relevant driver of skill shortages than trade and import exposure. This corroborates the notion that as automation capital substitutes existing tasks, skill demand changes more rapidly and workers face larger challenges to meet these new skill demands. In the following, we thus investigate the relationship between automation risk and skill shortage at the occupational level in more detail.

4.4.2. Occupational Level

A large body of literature suggests that routine occupations are particularly exposed to automation risks (e.g., Frey and Osborne, 2017; Arntz et al., 2016; Nedelkoska and Quintini, 2018). Thus, occupations with higher susceptibility to automation are at larger risk that their tasks are replaced by robots and automation technologies. Accordingly, the skill requirements for these occupations change more rapidly (e.g., Acemoglu and Restrepo, 2019; Deming and Noray, 2020), and occupations with a higher risk of automation should face larger skill gaps. Our data provide suggestive evidence for this assertion. Figure 8 depicts the relationship between skill shortage and automation risk across occupations. Our measure of automation risk stems from Nedelkoska and Quintini (2018), who constructed an occupation-specific measure of automation risk for all our sample countries in the PIAAC data. For instance, teaching professionals have a probability of 31 percent that their occupation is being substituted by automation, while it is 68 percent

Figure 8: Automation Risk and Skill Shortage

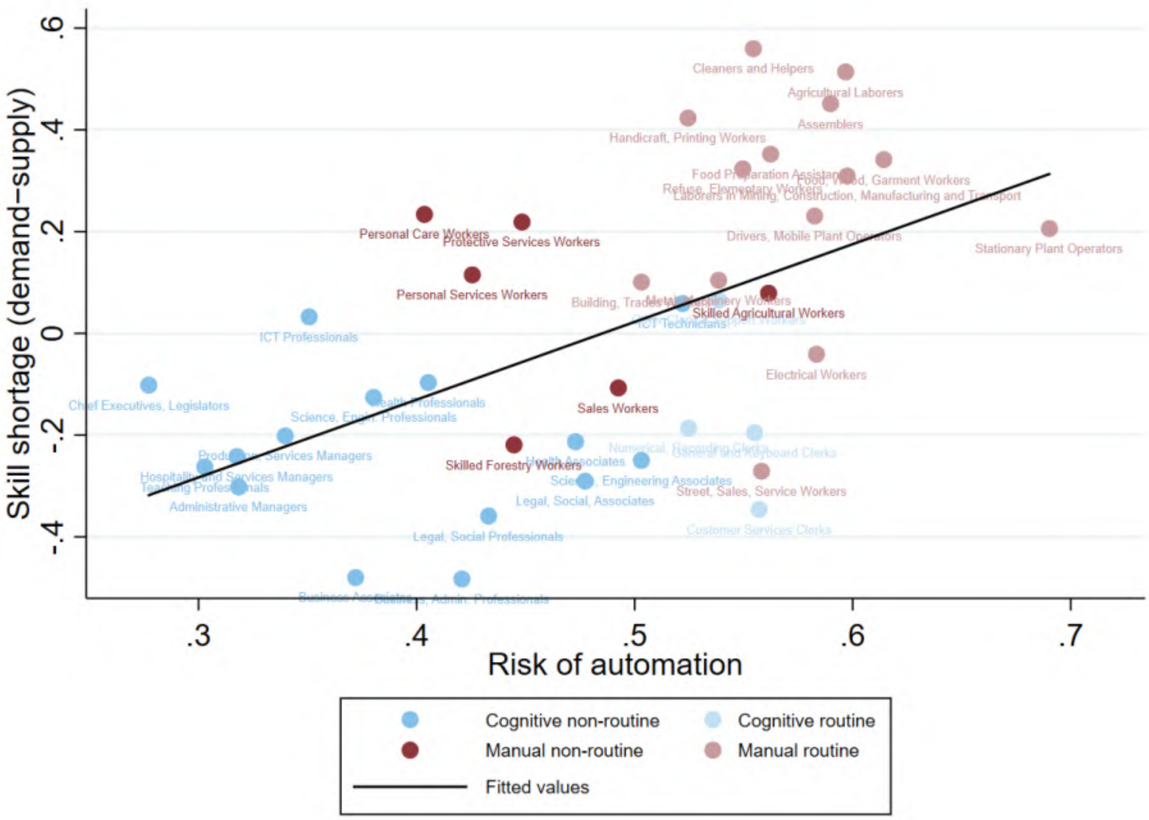

Notes: Correlation between automation risk and skill shortage, pooled for 17 European countries. Our measure of automation risk stems from Nedelkoska and Quintini (2018), who constructed the automation probability for all occupations and countries in our sample using PIAAC data. Data are from CEDEFOP and PIAAC.

for food preparation assistants. Figure 8 shows a positive relationship between the risk of automation and the average skill shortage of occupations: a higher risk of automation is associated with higher skill shortage. Thus, occupations more exposed to the risk of automation, such as food preparation assistants and plant operators, exhibit larger skill shortages. This is consistent with the notion that as automation technologies become able to perform existing tasks, skill demand for these occupations changes more rapidly and workers face larger challenges to meet these new skill demands. At the same time, occupations at lower risk of automation have skill surpluses, indicating that they provide more skills than are currently required from employers.

We have shown that workers in manual occupations are more exposed to automation risk and face larger skill shortages, which potentially stem from more rapid changes in skill

Figure 9: Job Training and Skill Shortage

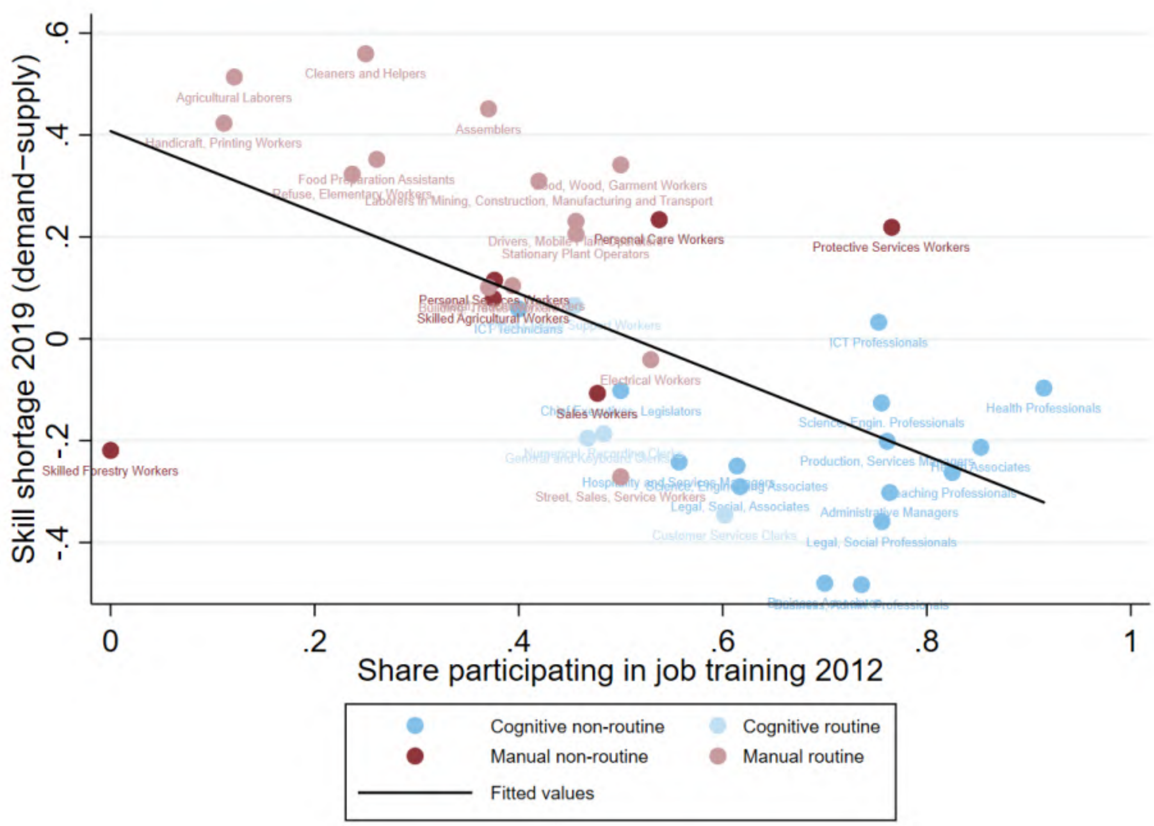

Notes: Correlation between on-the-job training (measured in 2012) and skill shortage (measured in 2019), pooled for 17 European countries. Data are from CEDEFOP and PIAAC.

requirements due to the automation of existing tasks. One way to mitigate the adverse ramifications of technological change on skill gaps is training to re-educate employees so as to prepare them for changing skill demands. Indeed, Figure 9 shows that occupations with a higher share of workers participating in on-the-job training in the PIAAC data (measured in 2012 or 2015, depending on the PIAAC country) are less likely to exhibit skill shortages in 2019. We use the training intensity in the year 2012/2015, as it is likely that the degree of skill mismatch of employees in an occupation in 2019 depends on the participation in, or missing out on, training in the past. Thus, a potential explanation for this could be that workers in occupations with more training were better at anticipating skills demand changes and investing in on-the-job training to stay on the frontier of what is demanded in their respective job. Conversely, there is a skill shortage for occupations that showed low rates of training in the past. At the same time, occupations with lower training intensities in the past are also those more exposed to automation risks, such as manual occupations like agricultural laborers and food preparation assistants. This suggests that workers who did not invest in job training were less prepared for changing skill requirements and thus show more pronounced skill shortages.[16]

4.5. The Cost of Skill Shortage

Table 2 shows how (intensive margin) skill shortage is related to economic outcomes at the regional level, thus depicting the cost of skill shortages. The table reports coefficients from OLS regressions with country fixed effects and standard errors clustered at the country level. We find a significantly negative relationship between the skill shortage and GDP per capita, both measured in 2019: a one percentile increase in skill shortage in a region is associated with a decrease in GDP per capita of 5.1 percent (Column 1). This is approximately the difference in GDP per capita between the Madrid metropolitan region and the Basque Country in northern Spain.

Looking at labor market outcomes, we observe that skill shortage is positively and statistically significantly associated with the unemployment rate in a region. Specifically, a one percentile increase in the skill shortage is associated with a 0.38 percentage point increase in the unemployment rate (Table 2, Column 2). Further, skill shortages might be particularly detrimental at the beginning of individuals' careers. For instance, when employers have limited information about young workers' actual skills (e.g., Altonji and Pierret, 2001), firms might be particularly cautious to hire young applicants. Indeed,

[16]See also (Falck et al., 2022), who show a positive association between training participation and digital skills for elderly workers.

Table 2: Economic Cost of Skill Shortage in Europe

Outcome:	(1) log GDP per capita	(2) Unemployment rate	(3) Youth unemployment rate
Skill shortage	-0.0510***	0.384*	0.443**
	(0.00845)	(0.201)	(0.160)
Country Fixed Effects	Yes	Yes	Yes
Observations (Regions)	141	149	149
Countries	17	17	17

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Dependent variable is indicated in the column header. Skill shortage (intensive margin) is calculated as an occupation-weighted average at the NUTS2-level. (Log) GDP per capita at the NUTS2 level is obtained from EUROSTAT. Unemployment rate is obtained from LFS. Youth unemployment rate is defined as the share of unemployed of the labor force aged 17-24. All variables are measured in 2019. Standard errors clustered at the country level in parentheses.

Column 3 of Table 2 shows that a one percentile increase the skill shortage is associated with a 0.44 percentage point increase in the youth unemployment rate (for individuals aged 17-24). The negative relation between skill shortage and unemployment rate, both for the working age population in general and for the youth in particular, suggests that higher skill shortages reduce the workers' employability at various points during their careers. The large literature showing that entering the labor market during bad macroeconomic times leads to sizable and persistent earnings losses (Kahn, 2010; Oreopoulos et al., 2012; Arellano-Bover, 2022) suggests that skill shortages at the beginning of workers' careers, which impair their initial employment prospects, might also entail long-term costs by putting workers on worse career tracks.

5. Skill Mismatch at the Extensive Margin

Our previous analyses provided a picture of skill shortages at the intensive margin across European regions and investigated potential determinants and cost of intensive margin skill shortages. However, this intensive margin measure of skill shortage does not provide the complete picture of skill mismatch. To see that, consider that our measure only shows the degree of mismatch between skill supply and demand for workers already employed in certain occupations, e.g., Information and Communication Technology (ICT) professionals and for Science and Engineering professionals (see Figure 4). This means we only observe the skills of workers in jobs whose skill set closely matches the skills firms require from workers in these occupations. However, at the same time, there might be a mismatch between the number of individuals who are potentially suited to work in these occupations and the number of workers in these occupations demanded by employers. Thus, if the number of vacancies for a given occupation exceeds the number of suited candidates in a region, this would constitute skill shortage at the *extensive margin*.

Since our measure of extensive margin skill supply is based on the number of university graduates by field from ETER, and because university graduates tend to select into high-earning occupations to reap the returns to their human capital investments (e.g., Wiswall and Zafar, 2018), we will focus on occupations that have an above-median probability of being chosen by university graduates. These are predominantly white-collar occupations. Below, we depict the regional variation in skill shortage at the extensive margin for selected occupations across European regions. The figures show the share of vacancies (from OJAs) in a region that cannot be filled by the regional pool of university graduates (from ETER), as described in Section 3.3.

Note that the number of vacancies in an occupation generally exceeds the number of graduates suited to work in the occupation, for at least three reasons: First, not all vacancies are targeted at early-career workers but require a certain amount of occupational or work experience. Second, vacancies cannot only be filled by the regional pool of university graduates, but also by university graduates from other regions. Finally, we compare the number of OJAs posted and individuals graduating in the same year (i.e., 2019). However, if individuals graduated in the year before did not find employment instantaneously after graduation, open vacancies in 2019 could also be filled with graduates from preceding cohorts. However, differences across European regions in the share of vacancies that cannot be filled by current graduates from the local universities still provide an interesting picture of European-wide differences in the ability of firms to fill their vacancies by means of the regional (early-career) workforce and show in which regions certain occupations are in particularly high demand.[17]

Figure 10 shows the extensive margin skill shortage for business professionals. There is considerable variation across regions in the degree of shortage of business administration professionals. For instance, the potential supply of business administration professionals exceeds the demand by 38 percent in the Eastern Macedonia region in Greece (i.e., extensive margin skill surplus), while demand for business administration professionals exceeds potential supply in the Brussels region by as much as 94 percent (i.e., extensive mar-

[17] For this exercise, we use the occupational distributions by field of study for each country from PIAAC which is arguably stable over time as they reflect major and workplace preferences (e.g., Arcidiacono, 2004) that change only slowly over time. We thus included Hungary, which participated in PIAAC round 3 (in 2017) as the only European country, in our extensive margin analysis for completeness.

gin skill shortage). Further, metropolitan areas such as Athens, Dublin, and Stockholm show pronounced extensive margin skill shortages of business administration professionals, although neighbouring regions show excess supply. This might reflect that knowledge intensive service firms employing, e.g., business administration occupations, tend to cluster in the central business districts (CBDs) of urban agglomerations (e.g., Duranton and Puga, 2004; Content et al., 2022).

Similarly, Figures 11 and 12 show considerable heterogeneity in the skill shortage at the extensive margin for health professionals and legal, social, and cultural professions across European countries and regions. For instance, while there is excess supply of health professionals in countries such as Denmark, Italy, Greece, and Lithuania, as well as in Eastern Europe, Germany faces severe shortages of health professionals in almost all NUTS2 regions. In fact, in most German regions more than 90 percent of the vacancies for health professionals cannot be filled with the regional pool of university graduates.

However, there are other occupations for which there is a high degree of excess demand in almost all European regions and countries. One such example is information and communication technology (ICT) professionals (Figure 13). Here, the average extensive margin skill shortage across European regions is 74 percent, with many regions in which more than 90 percent of vacancies for ICT professionals cannot be filled by people graduating from university in the same region. This is consistent with the expanding diffusion of digital technologies across all economic sectors in European countries (e.g., European Central Bank, 2020), which increases the demand for ICT professionals and information-processing tasks (e.g., Brynjolfsson et al., 2002). This pronounced skill shortage for ICT professionals at the extensive margin contrasts the small extent of skill shortage at the intensive margin (see Figure 4). Thus, there exists a shortage of ICT professionals in the sense that the number of vacancies exceeds the potential supply of individuals suited to work in this occupation, while the skills provided by individuals actually working as ICT professionals closely match the skills that are required by firms. One potential interpretation of this finding is that on-the-job training offered to ICT professionals is effective in providing the skills that are needed in this occupation (small intensive-margin mismatch). At the same time, the higher education systems do not seem to be able to "produce" a sufficient amount of graduates being able to work as ICT professionals to cope with the high demand by firms (high extensive-margin mismatch).

Similarly, Figure 14 shows substantial extensive margin skill shortages for science and engineering professionals in almost all European regions. This is consistent with high demand for STEM occupations on the one hand, and generally low shares of STEM graduates on the other.[18] These results fuel the concern of policymakers and industry leaders worldwide regarding a shortage in the STEM workforce believed necessary to sustain innovative performance and global competitiveness. For instance, according to the President's Council of Advisors on Science and Technology, the United States would need to increase its annual production of undergraduate STEM degrees by 34 percent over current rates to match the demand forecast for STEM professionals (President's Council of Advisors on Science and Technology, 2012). Our results point to a need for more STEM professionals in Europe that is at least in a comparable order of magnitude as in the United States.[19]

Our novel data and measure of skill shortage allow a depiction and quantification of extensive margin shortages at the occupational and regional level. For instance, our results suggest that skill shortages for business administration and health professionals vary considerably between European regions while skill shortages are very high for ICT and STEM occupations across all European regions. These results help identify the regions and occupations that could potentially benefit from targeted education and occupational re-training policies to counteract skill mismatch.

[18]Research and policy have considered not only overall issues of attracting more people into STEM fields but also the large gender disparities in these choices (Altonji et al., 2016; Stoet and Geary, 2018; UNESCO, 2017; Deming and Noray, 2020). For instance, it has been shown that males are more likely to take mathematically oriented courses in school and to obtain bachelor's degrees in STEM disciplines compared to females (National Science Foundation, 2016, 2017). However, while the under-performance of females in mathematics and physics tests has narrowed or even reversed in many countries (OECD, 2016), the gender gaps in STEM university enrollment still remain. See Goulas et al. (2022) for a recent discussion of the role of peer effects in explaining gender differences in students' decisions to select and specialize in a STEM field.

[19]In our data, 5.5 percent of OJAs are for STEM occupations, compared to a share of graduates who are potentially qualified to work in STEM occupations of 3.1 percent. Thus, graduates in STEM would need to increase by 75 percent to close this gap. One of the reasons for why the STEM shortage is lower in the United States than in Europe many be that the high costs of attending college in the United States increases individuals' probability to choose high-paying degrees, such as STEM majors.

Figure 10: Skill Shortage (Extensive Margin) for Business Administration Professionals across European Regions

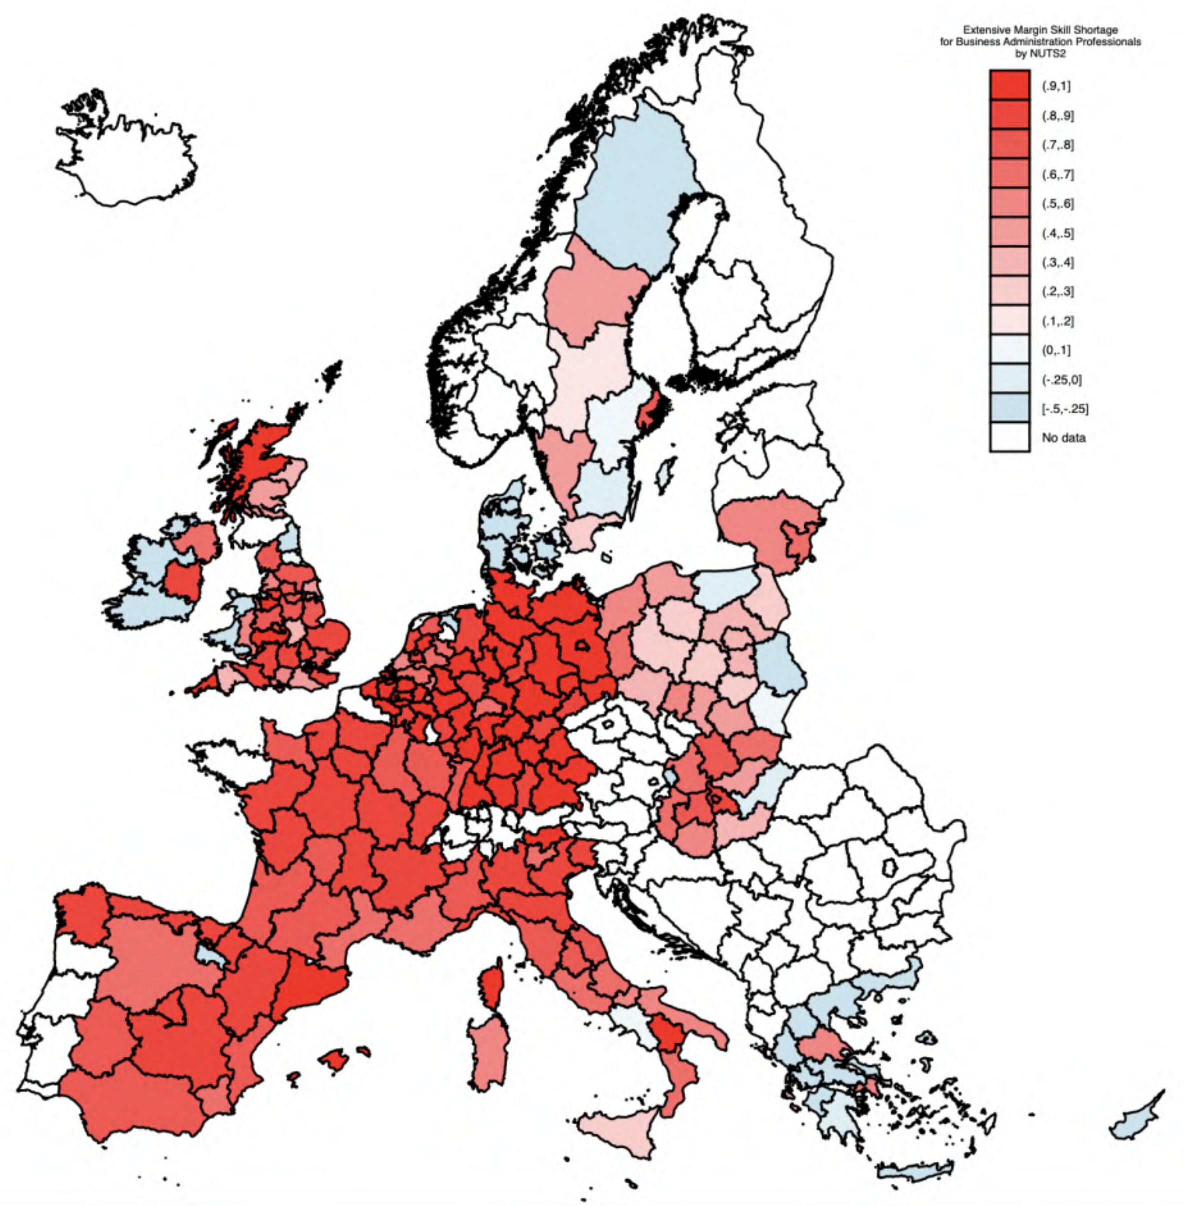

Notes: Figure shows skill shortage at the extensive margin for business administration professionals by NUTS2 region in Europe. Data are from CEDEFOP (demand side) as well as from PIAAC and ETER (supply side). The measure of extensive margin skill shortage is calculated as the difference between the number of OJAs for business administration professionals (demand) and the number of university graduates suited to work as business administration professionals (supply), divided by demand. Thus, the measure depicts the share of vacancies in a region that cannot be filled by the regional pool of university graduates. The average skill shortage for business administration professionals across European regions is 57 percent.

Figure 11: Skill Shortage (Extensive Margin) for Health Professionals across European Regions

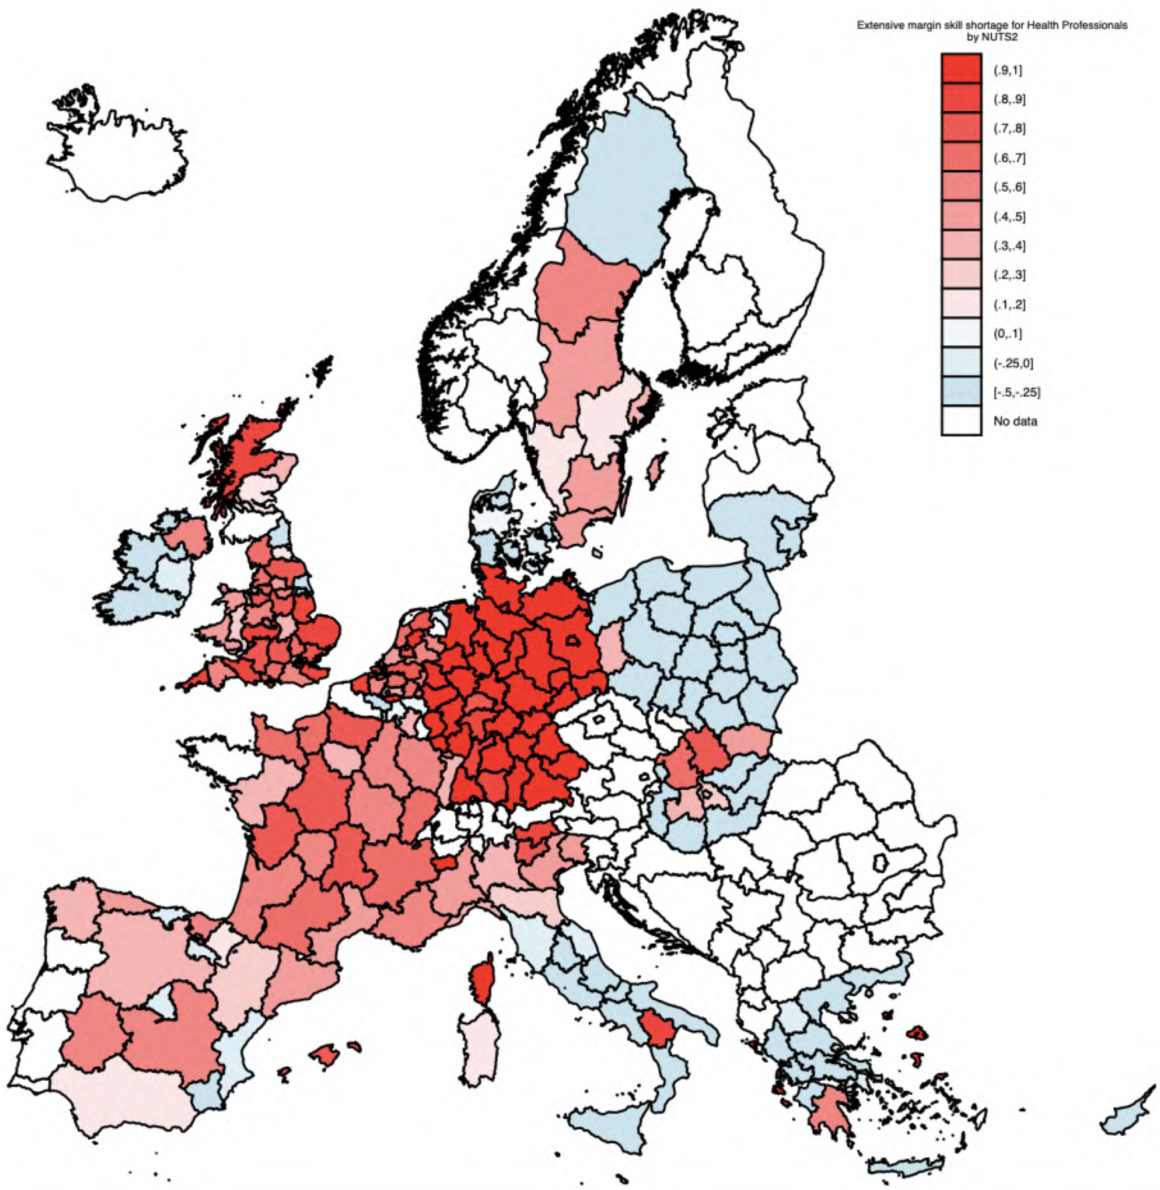

Notes: Figure shows skill shortage at the extensive margin for health professionals by NUTS2 region in Europe. Data are from CEDEFOP (demand side) as well as from PIAAC and ETER (supply side). The measure of extensive margin skill shortage is calculated as the difference between the number of OJAs for health professionals (demand) and the number of university graduates suited to work as health professionals (supply), divided by demand. Thus, the measure depicts the share of vacancies in a region that cannot be filled by the regional pool of university graduates. The average skill shortage for health professionals across European regions is 33 percent.

Figure 12: Skill Shortage (Extensive Margin) for Legal, Cultural, and Social professionals across European Regions

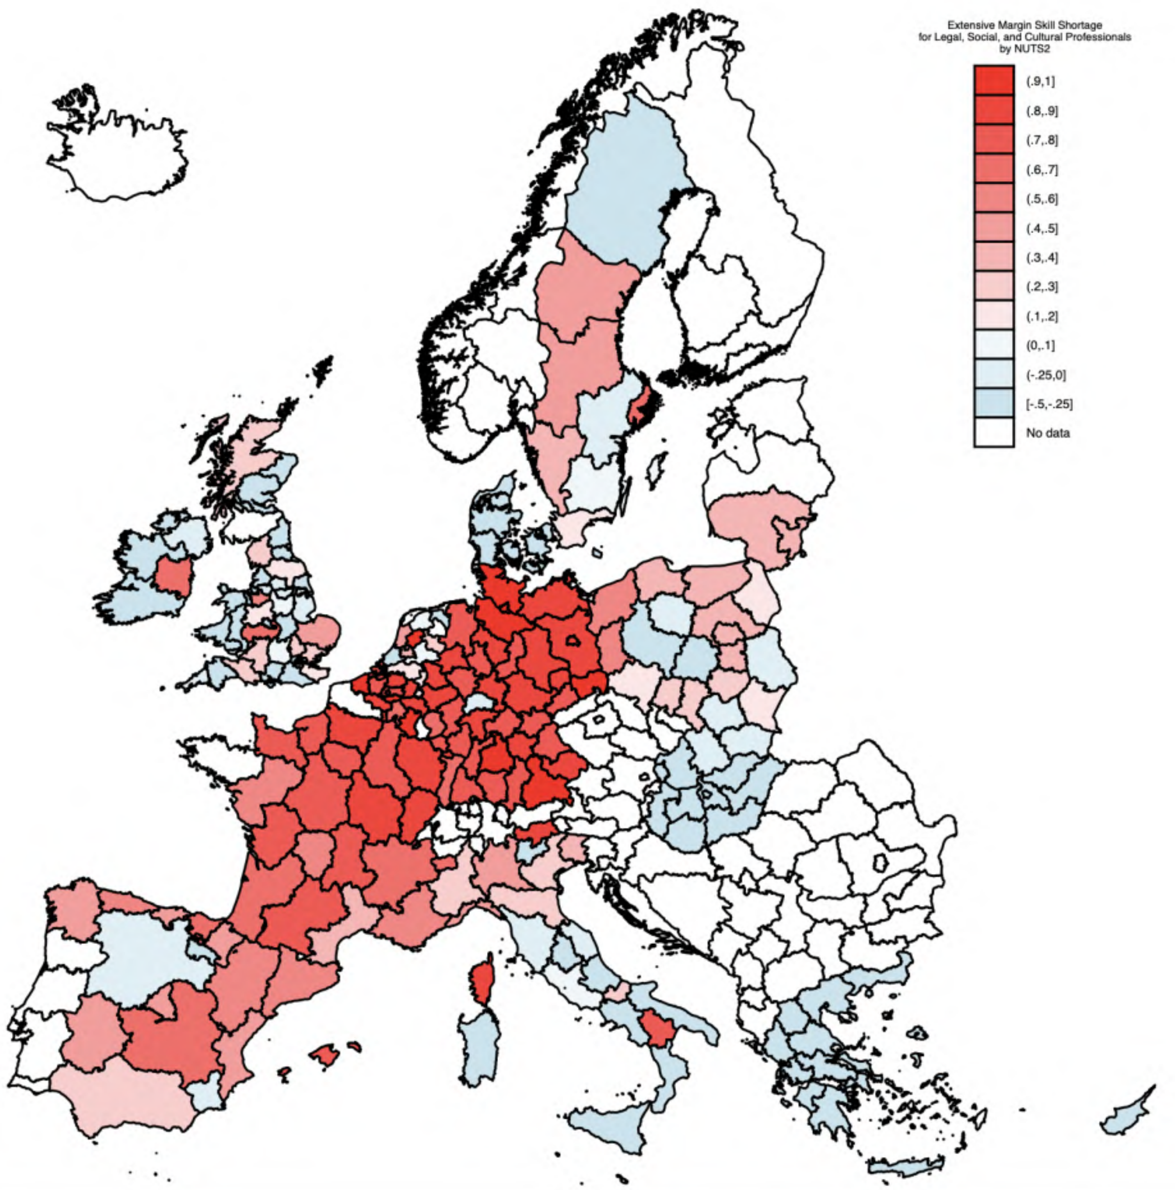

Notes: Figure shows skill shortage at the extensive margin for legal, cultural, and social professionals by NUTS2 region in Europe. Data are from CEDEFOP (demand side) as well as from PIAAC and ETER (supply side). The measure of extensive margin skill shortage is calculated as the difference between the number of OJAs for legal, cultural, and social professionals (demand) and the number of university graduates suited to work as legal, cultural, and social professionals (supply), divided by demand. Thus, the measure depicts the share of vacancies in a region that cannot be filled by the regional pool of university graduates. The average skill shortage for legal, cultural, and social professionals across European regions is 21 percent.

Figure 13: Skill Shortage (Extensive Margin) for ICT Professionals across European Regions

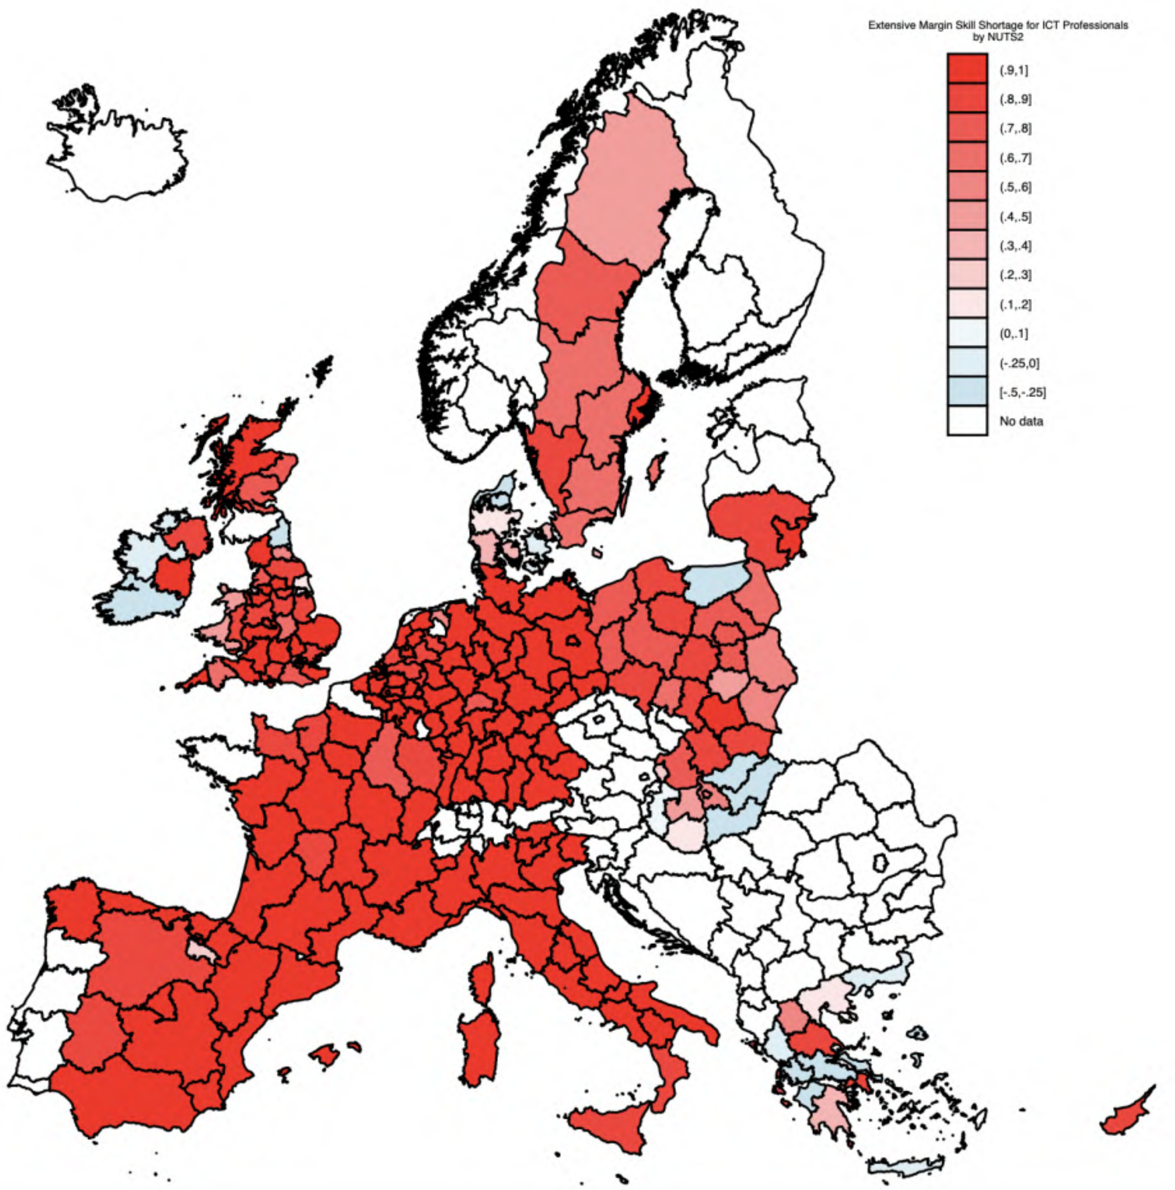

Notes: Figure shows skill shortage at the extensive margin for information and communication technology (ICT) professionals by NUTS2 region in Europe. Data are from CEDEFOP (demand side) as well as from PIAAC and ETER (supply side). The measure of extensive margin skill shortage is calculated as the difference between the number of OJAs for ICT professionals (demand) and the number of university graduates suited to work as ICT professionals (supply), divided by demand. Thus, the measure depicts the share of vacancies in a region that cannot be filled by the regional pool of university graduates. The average skill shortage for ICT professionals across European regions is 74 percent.

Figure 14: Skill Shortage (Extensive Margin) for Science and Engineering Professionals across European Regions

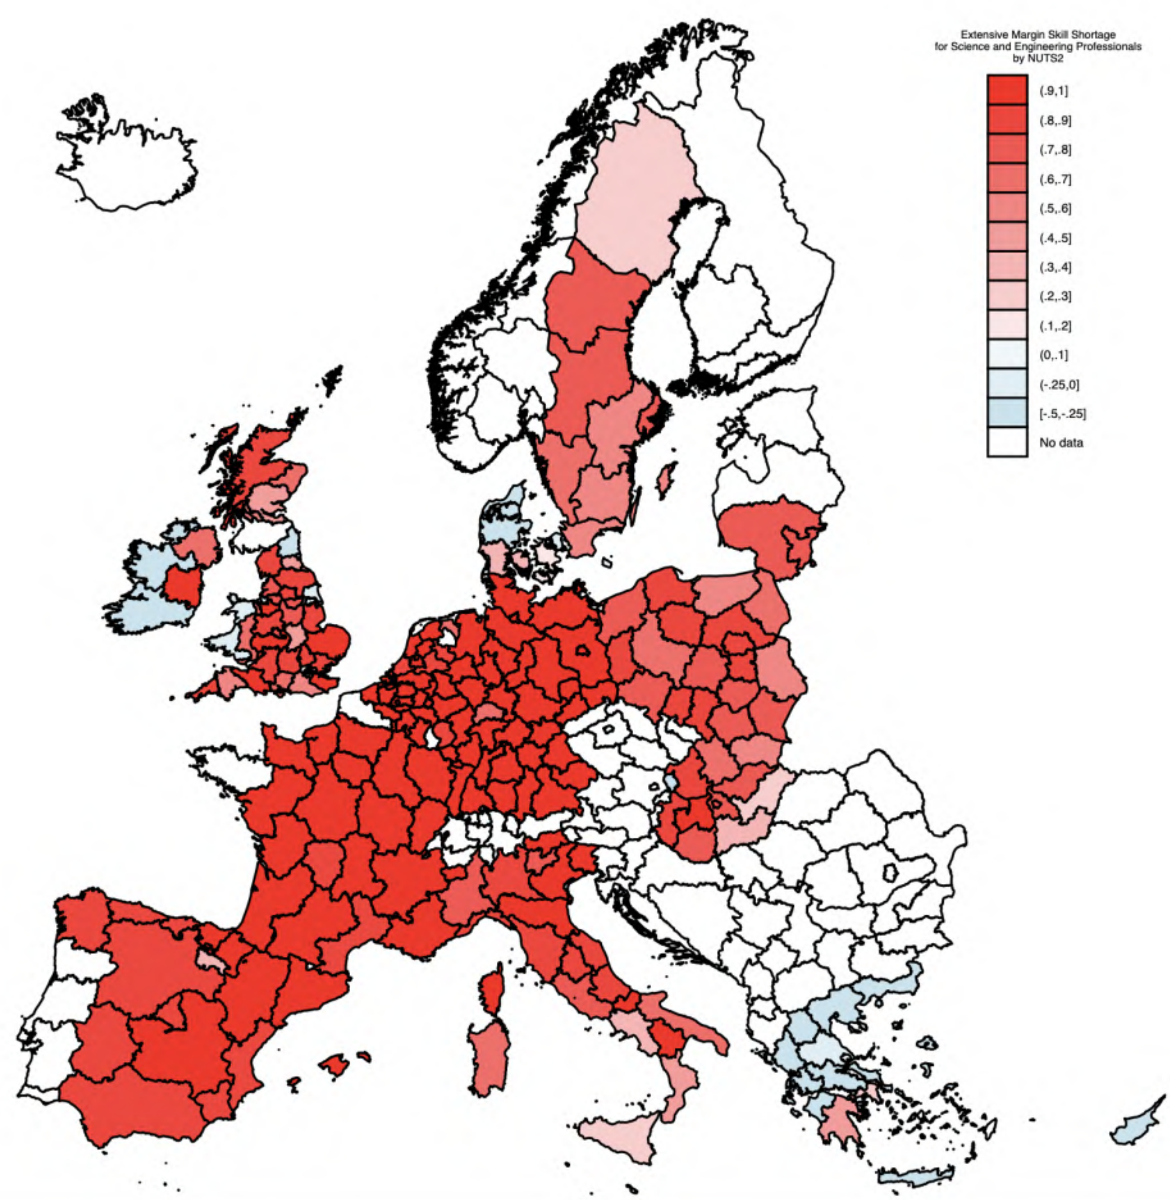

Notes: Figure shows skill shortage at the extensive margin for science and engineering professionals by NUTS2 region in Europe. Data are from CEDEFOP (demand side) as well as from PIAAC and ETER (supply side). The measure of extensive margin skill shortage is calculated as the difference between the number of OJAs for science and engineering professionals (demand) and the number of university graduates suited to work as science and engineering professionals (supply), divided by demand. Thus, the measure depicts the share of vacancies in a region that cannot be filled by the regional pool of university graduates. The average skill shortage for science and engineering professionals across European regions is 68 percent.

6. Conclusion

In this paper, we address multiple questions of high relevance for labor market policies in the European Union: (1) How prevalent is skill mismatch in Europe, and (2) what are the drivers of these skill gaps and how can workers better prepare for the skill demand of employers. Drawing on innovative online job ad data and skill survey data for 17 European countries, we develop novel measures of skill mismatch at two margins: At the intensive margin, we ask how well the skills that workers currently use at the workplace fit the skills employers require when adverting new jobs. Complementing this, skill mismatch at the extensive margin shows to what extent the regional pool of university graduates is actually able to fill the vacancies firms have.

We first look at skill mismatch at the intensive margin. We document that skill shortages in Europe exist, but the extent and direction of these shortages vary considerably by type of occupation and across regions. In particular, manual workers generally face skill shortages, while cognitive workers exhibit skill surpluses. We further find that the observed differences in skill shortage are not driven by specific skill domains, such as digital or social skills, but are consistent across all occupational types and skill domains. Intriguingly, the overall patterns of skill shortage do not differ much between women and men.

We also present a quantification of skill shortage across European regions. Our data suggest that, on average, skill supply exceeds skill demand at the intensive margin. That is, the amount of skills workers report to need in their jobs are on average higher than the skills required by employers in their job ads. However, this does not necessarily imply that employers demand lower skills than their employees exhibit: First, employers are potentially more specific about the skills they seek in highly technical occupations and thus list a lower, but more explicit set of skills in their job ads. Second, employers could infer skills from a degree requirement or occupational title and thus not specifically state them in the requirement section as they deem these skills to be obvious (e.g., a welder needs to have welding skills). However, we find substantial heterogeneity in skill mismatch across regions and countries: While skill shortage is prevailing in countries such as Greece, Italy, Spain, and Poland, the other European countries are rather characterized by skill surpluses.

We also investigate potential drivers of this variation in skill shortage. First, we find that economically stronger and more dynamic regions systematically face less skill shortage. These regions may yield higher returns to skills, thus incentivizing workers to invest in their skills and meet the skill demands employers have. Second, there is a positive association between skill shortage and exposure to automation risk. This underpins the notion that automation technologies shift existing skill demands more rapidly. Third, regions more exposed to trade and import competition do not suffer from higher skill shortage. Finally, we also document that skill shortages entail economic costs, such as lower economic output and higher unemployment rates in a region.

Furthermore, our measure of skill mismatch at the extensive margin provides a novel, European-wide quantification of the gap in the number of workers demanded in an occupation and the number of individuals who are suited to work in this occupation. Using this measure, we find especially pronounced skill shortages for information and communication (ICT) and STEM workers across almost all European regions.

Our results thus suggest that skill shortage poses a challenge for at least two reasons: First, education and training systems need to ensure that the skills individuals provide match the skills demanded by employers. Our analyses show that there is considerable heterogeneity in this dimension of skill shortage across occupations, skill domains, and regions. At the same time, we provide evidence that skill shortage is systematically lower in occupations that provide more on-the-job training. Thus, our results emphasise further training as a potential measure to meet changing skill requirements and mitigate skill shortages in the European Union. Beyond, our analyses offer insights into which types of workers and regions might particularly benefit from targeted training policies, namely, occupations or regions that underwent technological change at a more rapid pace, and thus have been particularly affected by the automation of routine tasks.

Second, beyond potential skill mismatch of workers already employed in an occupation (intensive margin), there may also be shortages in the number of people who are qualified to work in certain occupations that are currently in high demand. We see evidence for such extensive-margin skill mismatch particularly in ICT and STEM occupations. One way to counteract such shortages are education and occupational re-training policies; these, however, would surely take some time to show effects. A possible means of expanding employers' talent pool already in the short-run could be a reevaluation of existing hiring practices. For instance, employers could consider alternatives like skills-based hiring, placing higher emphasis on skills instead of degrees. Especially in the context of ICT occupations, this could be a fruitful pathway to decrease mismatch at the extensive margin (Fuller et al., 2022).

In sum, our descriptive evidence points to the fact that skill gaps are prevalent in the European Union and are potentially accompanied by skill depreciation and lower adaptability to technological change. This has potential adverse impacts on workers in terms of earnings and job satisfaction, but also on the productivity of firms and, as we show, the economic development of European regions. Anticipation of future skill needs and providing the opportunity to meet these needs are thus of fundamental importance for European countries to increase productivity, job satisfaction, and competitiveness of both employers and employees.

References

Acemoglu, D. and D. Autor (2011). Chapter 12 - Skills, Tasks and Technologies: Implications for Employment and Earnings. In D. Card and O. Ashenfelter (Eds.), *Handbook of Labor Economics*, Volume 4, pp. 1043–1171. Elsevier.

Acemoglu, D. and P. Restrepo (2019). Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives 33*(2), 3–30.

Alabdulkareem, A., M. R. Frank, L. Sun, B. AlShebli, C. Hidalgo, and I. Rahwan (2018). Unpacking the Polarization of Workplace Skills. *Science Advances 4*(7).

Altonji, J. G., P. Arcidiacono, and A. Maurel (2016). Chapter 7 - The Analysis of Field Choice in College and Graduate School: Determinants and Wage Effects. In E. A. Hanushek, L. Woessmann, and S. Machin (Eds.), *Handbook of the Economics of Education* (1st ed.), Volume 5, pp. 305–396. Amsterdam: North Holland: Elsevier.

Altonji, J. G. and C. R. Pierret (2001). Employer Learning and Statistical Discrimination. *The Quarterly Journal of Economics 116*(1), 313–350. Publisher: Oxford University Press.

Alvarez-Cuadrado, F., N. V. Long, and M. Poschke (2018). Capital-Labor Substitution, Structural Change and the Labor Income Share. *Journal of Economic Dynamics and Control 87*, 206–231.

Arcidiacono, P. (2004). Ability Sorting and the Returns to College Major. *Journal of Econometrics 121*(1), 343–375.

Arellano-Bover, J. (2022). The Effect of Labor Market Conditions at Entry on Workers' Long-Term Skills. *The Review of Economics and Statistics 104*(5), 1028–1045.

Arntz, M., T. Gregory, and U. Zierahn (2016). The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis. Technical report, OECD, Paris.

Autor, D. H. (2015). Why Are There Still So Many Jobs? The History and Future of Workplace Automation. *Journal of Economic Perspectives 29*(3), 3–30.

Autor, D. H., D. Dorn, and G. H. Hanson (2013). The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American Economic Review 103*(6), 2121–2168.

Autor, D. H., F. Levy, and R. J. Murnane (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics 118*(4), 1279–1333.

Avesani, P., F. Giunchiglia, and M. Yatskevich (2005). A Large Scale Taxonomy Mapping Evaluation. In *International Semantic Web Conference*, pp. 67–81. Springer.

Balassa, B. (1964). The Purchasing-Power Parity Doctrine: A Reappraisal. *Journal of Political Economy 72*(6), 584–596.

Baroni, M., G. Dinu, and G. Kruszewski (2014). Don't Count, Predict! a Systematic Comparison of Context-Counting vs. Context-Predicting Semantic Vectors. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 238–247.

Barrera-Osorio, F., A. D. Kugler, and M. I. Silliman (2020). Hard and Soft Skills in Vocational Training: Experimental Evidence from Colombia. NBER Working Paper Nr. 27548.

Berkes, E., P. Mohnen, and B. Taska (2018). The Consequences of Initial Skill Mismatch for College Graduates: Evidence from Online Job Postings. mimeo.

Bojanowski, P., E. Grave, A. Joulin, and T. Mikolov (2017). Enriching Word Vectors with Subword Information. *Transactions of the Association for Computational Linguistics 5*, 135–146. Publisher: MIT Press.

Brynjolfsson, E., L. M. Hitt, and S. Yang (2002). Intangible Assets: Computers and Organizational Capital. *Brookings Papers on Economic Activity 2002*(1), 137–198.

Burke, M. A., A. Sasser Modestino, S. Sadighi, R. B. Sederberg, and B. Taska (2020). No Longer Qualified? Changes in the Supply and Demand for Skills within Occupations. Working Paper Federal Reserve Bank of Boston.

Camacho-Collados, J. and M. T. Pilehvar (2018). From Word to Sense Embeddings: A Survey on Vector Representations of Meaning. *JAIR 63*.

Caunedo, J., E. Keller, and Y. Shin (2021). Technology and the Task Content of Jobs Across the Development Spectrum. NBER Working Paper Nr. 28681.

Collobert, R. and J. Weston (2008). A Unified Architecture for Natural Language Processing: Deep Neural Networks with Multitask Learning. In *The 25th International Conference on Machine Learning*, pp. 160–167.

Content, J., N. Cortinovis, K. Frenken, and J. Jordaan (2022). The Roles of KIBS and R&D in the Industrial Diversification of Regions. *The Annals of Regional Science 68*(1), 29–64.

Dauth, W., S. Findeisen, J. Suedekum, and N. Woessner (2017). German Robots - The Impact of Industrial Robots on Workers. IAB-Discussion Paper 30/2017.

Deming, D. J. (2017). The Growing Importance of Social Skills in the Labor Market. *The Quarterly Journal of Economics 132*(4), 1593–1640.

Deming, D. J. and K. Noray (2020). Earnings Dynamics, Changing Job Skills, and STEM Careers. *The Quarterly Journal of Economics 135*(4), 1965–2005.

Devlin, J., M.-W. Chang, K. Lee, and K. Toutanova (2018). Bert: Pre-Training of Deep Bidirectional Transformers for Language Understanding. *arXiv preprint arXiv:1810.04805*.

Duranton, G. and D. Puga (2004). Chapter 48 - Micro-Foundations of Urban Agglomeration Economies. In J. V. Henderson and J.-F. Thisse (Eds.), *Handbook of Regional and Urban Economics*, Volume 4 of *Cities and Geography*, pp. 2063–2117. Elsevier.

European Central Bank (2020). Economic Bulletin. (Issue 8), 177.

European Commission (2011). An Agenda for New Skills and Jobs: A European Contribution Towards Full Employment. Technical report.

Euzenat, J., D. Loup, M. Touzani, and P. Valtchev (2004). Ontology Alignment with OLA. In *Proc. 3rd ISWC2004 workshop on Evaluation of Ontology-based tools (EON)*, pp. 59–68.

Falck, O., A. Heimisch-Roecker, and S. Wiederhold (2021). Returns to ICT Skills. *Research Policy 50*(7), 104064.

Falck, O., V. Lindlacher, and S. Wiederhold (2022). Elderly Left Behind? How Older Workers Can Participate in the Modern Labor Market. EconPol Forum 23 (5), 16-19.

Frey, C. B. and M. A. Osborne (2017). The Future of Employment: How Susceptible are Jobs to Computerisation? *Technological Forecasting and Social Change 114*, 254–280.

Fuller, J. B., C. Langer, J. Nitschke, L. O'Kane, M. Sigelman, and B. Taska (2022). The Emerging Degree Reset: How the Shift to Skills-Based Hiring Holds the Keys to Growing the U.S. Workforce at a Time of Talent Shortage. *Burning Glass Institute White Paper*.

Giabelli, A., L. Malandri, F. Mercorio, and M. Mezzanzanica (2022a). GraphLMI: A Data Driven System for Exploring Labor Market Information through Graph Databases. *Multimedia Tools and Applications 81*(3), 3061–3090.

Giabelli, A., L. Malandri, F. Mercorio, and M. Mezzanzanica (2022b). WETA: Automatic Taxonomy Alignment via Word Embeddings. *Computers in Industry 138*, 103626. Publisher: Elsevier.

Giabelli, A., L. Malandri, F. Mercorio, M. Mezzanzanica, and A. Seveso (2020). NEO: A Tool for Taxonomy Enrichment with New Emerging Occupations. In *The Semantic Web - ISWC 2020 - 19th International Semantic Web Conference*, Volume 12507 of *Lecture Notes in Computer Science*, pp. 568–584. Springer.

Giabelli, A., L. Malandri, F. Mercorio, M. Mezzanzanica, and A. Seveso (2021a). NEO: A System for Identifying New Emerging Occupation from Job Ads. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Volume 35, pp. 16035–16037. Issue: 18.

Giabelli, A., L. Malandri, F. Mercorio, M. Mezzanzanica, and A. Seveso (2021b). Skills2Job: A Recommender System that Encodes Job Offer Rmbeddings on Graph Databases. *Appl. Soft Comput. 101*, 107049.

Goulas, S., S. Griselda, and R. Megalokonomou (2022). Comparative Advantage and Gender Gap in STEM. *Journal of Human Resources*, 0320. Publisher: University of Wisconsin Press.

Guo, Y., C. Langer, F. Mercorio, and F. Trentini (2022). Skills Mismatch, Automation, and Training: Evidence from 17 European Countries Using Survey Data and Online Job Ads. EconPol Forum 23 (5), 11-15.

Hanushek, E. A., G. Schwerdt, S. Wiederhold, and L. Woessmann (2015). Returns to Skills around the World: Evidence from PIAAC. *European Economic Review 73*, 103–130.

Hanushek, E. A., G. Schwerdt, S. Wiederhold, and L. Woessmann (2017). Coping with Change: International Differences in the Returns to Skills. *Economics Letters 153*, 15–19.

Hartog, J. (2000). Over-Education and Earnings: Where are We, Where Should We Go? *Economics of Education Review 19*(2), 131–147.

Jung, J. J. (2008). Taxonomy Alignment for Interoperability between Heterogeneous Virtual Organizations. *Expert Systems with Applications 34*(4), 2721–2731. Publisher: Elsevier.

Kahn, L. B. (2010). The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy. *Labour Economics 17*(2), 303–316.

Levy, O. and Y. Goldberg (2014). Neural Word Embedding as Implicit Matrix Factorization. In *Advances in Neural Information Processing Systems*, Volume 27, pp. 2177–2185.

Maedche, A. and S. Staab (2001). Ontology Learning for the Semantic Web. *IEEE Intelligent systems 16*(2), 72–79. Publisher: IEEE.

Malandri, L., F. Mercorio, M. Mezzanzanica, and N. Nobani (2020). MEET: A Method for Embeddings Evaluation for Taxonomic Data. In *2020 International Conference on Data Mining Workshops (ICDMW)*, pp. 31–38. IEEE.

Malandri, L., F. Mercorio, M. Mezzanzanica, and N. Nobani (2021). MEET-LM: A Method for Embeddings Evaluation for Taxonomic Data in the Labour Market. *Computers in Industry 124*, 103341. Publisher: Elsevier.

Meng, Y., J. Huang, G. Wang, C. Zhang, H. Zhuang, L. Kaplan, and J. Han (2019). Spherical Text Embedding. In *Advances in Neural Information Processing Systems*, pp. 8208–8217.

Mikolov, T., I. Sutskever, K. Chen, G. S. Corrado, and J. Dean (2013). Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pp. 3111–3119.

National Science Foundation (2016). Science and Engineering Indicators. Technical Report 20161.

National Science Foundation (2017). Women, Minorities, and Persons with Disabilities in Science. Technical Report 17310.

Nedelkoska, L. and G. Quintini (2018). Automation, Skills Use and Training. Technical report, OECD, Paris.

Nguyen, K. A., M. Koeper, S. S. i. Walde, and N. T. Vu (2017). Hierarchical Embeddings for Hypernymy Detection and Directionality. *arXiv preprint arXiv:1707.07273*.

OECD (2016). PISA 2015 Results (Volume I): Excellence and Equity in Education.

OECD (2019). *OECD Employment Outlook 2019: The Future of Work*. Paris: Organisation for Economic Co-operation and Development.

Oreopoulos, P., T. von Wachter, and A. Heisz (2012). The Short- and Long-Term Career Effects of Graduating in a Recession. *American Economic Journal: Applied Economics 4*(1), 1–29.

Pennington, J., R. Socher, and C. Manning (2014). Glove: Global Vectors for Word Representation. In *EMNLP*, pp. 1532–1543.

Piopiunik, M., G. Schwerdt, L. Simon, and L. Woessmann (2020). Skills, Signals, and Employability: An Experimental Investigation. *European Economic Review 123*, 103374.

President's Council of Advisors on Science and Technology (2012). Report to the President. Engaged to Excel: Producing one Million Additional College Graduates with Degrees in Science, Technology, Engineering, and Mathematics. Technical report.

Stoet, G. and D. C. Geary (2018). The Gender-Equality Paradox in Science, Technology, Engineering, and Mathematics Education. *Psychological Science 29*(4), 581–593. Publisher: SAGE Publications Inc.

Turian, J., L. Ratinov, and Y. Bengio (2010). Word Representations: A Simple and General Method for Semi-Supervised Learning. In *ACL*.

UNESCO (2017). Cracking the code: girls' and women's education in science, technology, engineering and mathematics (STEM) - UNESCO Digital Library.

Wang, R. Y. and D. M. Strong (1996). Beyond Accuracy: What Data Quality Means to Data Consumers. *Journal of management information systems 12*(4), 5–33. Publisher: Taylor & Francis.

Weidmann, B. and D. J. Deming (2021). Team Players: How Social Skills Improve Team Performance. *Econometrica 89*(6), 2637–2657. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA18461.

Wiswall, M. and B. Zafar (2018). Preference for the Workplace, Investment in Human Capital, and Gender. *The Quarterly Journal of Economics 133*(1), 457–507.

Wu, T., G. Qi, H. Wang, K. Xu, and X. Cui (2016). Cross-Lingual Taxonomy Alignment with Bilingual Biterm Topic Model. In *AAAI*, pp. 287–293.

Wu, Z., H. Zhu, G. Li, Z. Cui, H. Huang, J. Li, E. Chen, and G. Xu (2017). An Efficient Wikipedia Semantic Matching Approach to Text Document Classification. *Information Sciences 393*, 15–28. Publisher: Elsevier.

Wyness, G. and R. Murphy (2020). What is the nature and extent of student-university mismatch? *IZA World of Labor, No 477*.

Xue, Y. and R. Larson (2015). STEM Crisis or STEM Surplus? Yes and Yes. *Monthly Labor Review*.

Appendix A. PIAAC to ESCO Mapping: Technical Appendix

In recent literature, taxonomy alignment has received considerable attention and different approaches have been proposed to create efficient mappings in a variety of context. In one of the first approaches, Euzenat et al. (2004) compute the similarity between entities through a system of quasi-linear equations. They start from lexical similarity derived by WordNet 2.0, a lexical database for English, and gradually include contributions from structure comparing functions. Avesani et al. (2005) uses both a syntactic and a semantic score of taxonomic similarity, called COMA and S-match, respectively. COMA exploits both element and structure-level syntactic similarity, while S-match uses Wordnet 2.0 to derive semantic similarity between words. In Wu et al. (2017), an approach based on Wikipedia-matching and keywords is considered to perform document classification without employing standard occurrence methods. Despite being relevant and widely used, those methods are built on specific lexical resources (WordNet, Wikipedia) and thus are unsuitable for different domains, such as job ads, which in turn have a specific jargon that our analysis must consider. In Jung (2008) the authors use Latent Semantic Analysis (LSA) to group sets of related entities based on their co-occurrence matrix and TF-IDF (term frequency-inverse document frequency), also considering the description of the taxonomic concepts. In Wu et al. (2016), authors train a bilingual topic model on contextual text extracted from the web to build semantic vectors of the topics of two multi-lingual taxonomies. The cosine similarity between those vectors represents the relevance of each concept in the source taxonomy and its candidate-matched categories. Each candidate entity is then evaluated through syntactic similarity. Those kinds of approaches make use of contextual information and learning algorithms. However, none considers the taxonomy's vertical structure to match entities nor employs distributional semantics, which has been beneficial in several Natural Language Processing (NLP) applications in the last years. More recently, Giabelli et al. (2022a) proposed the `WETA` (**W**eb **T**axonomy **E**mbedding **A**lignment) approach that exploits distributional semantic and context information to perform taxonomy alignment, blending a hierarchical approach based on cosine similarity and a machine learning classification task that uses the embeddings as input features. Moreover, it performs an intrinsic evaluation of the selected embedding model based on the structure of the taxonomy itself.

Introducing Word Embeddings. Evaluating the intrinsic quality of vector space models, as well as their impact when used as the input of specific tasks (*aka*, extrinsic quality), have a very practical significance (see, e.g. Turian et al. (2010); Camacho-Collados and Pilehvar (2018)), as this affects the believability[20] of the overall process or system in which they are used. In essence, we may argue that the well-known principle *"garbage-in, garbage-out"* that characterises the data quality research in many domains also applies to word embedding, that is, *the lower the quality of the word embeddings, the lower the performance of the tasks that are based on them.*

Word embeddings are vector representations of words based on the hypothesis that words occurring in a similar context tend to have a similar meaning. Words are represented by semantic vectors, which are usually derived from a large corpus using co-occurrence statistics, and their use improves learning algorithms in many NLP tasks. Two powerful methods to induce word embeddings are neural networks training (Collobert and Weston, 2008; Mikolov et al., 2013) and co-occurrence matrix factorisation (Pennington et al., 2014; Levy and Goldberg, 2014).

These techniques consider each word as a distinct vector and ignore the morphological similarity among them. More recently Bojanowski et al. (2017) developed a version of the continuous skip-gram model (Mikolov et al., 2013) which considers subword information. This architecture, called fastText, is an extension of word2vec for scalable word representation and classification. One of its major improvements to word2vec is to consider sub-word information by representing each word as the sum of its character n-gram vectors. Formally, given a word w, and a dictionary of size G, G_{w} is the set of n-grams of size G appearing in w. Denoted as z_g vector representation of the n-gram g, w will be represented as the sum of the vector representation of its n-grams and the score associated to the word w as:

$$\int(\mathrm{w},\mathrm{c}) = \sum_{g\in G_{\mathrm{w}}} z_g^{\intercal} v_{\mathrm{c}} \tag{A.1}$$

[20]Here the term believability is intended as "the extent to which data are accepted or regarded as true, real and credible"Wang and Strong (1996)

where v_c is the vector representing the context. This simple representation allows one to share information between words, making it useful to represent rare words, typos, and words with the same root. Other embedding models have been evaluated along with fastText. Nevertheless, none of them fit our conditions: Neither classical embedding models (Mikolov et al., 2013; Pennington et al., 2014) nor embeddings specifically designed to fit taxonomic data consider subword information. Moreover, they cannot be easily bonded with external sources in their generation phase, reducing the flexibility of . Regarding hyperbolic and spherical embeddings like HyperVec (Nguyen et al., 2017) or JoSe (Meng et al., 2019), we discarded them since (i) they also don't consider subword information, which is important for short text and many words with the same root (e.g. engineer-engineering, developer-developing) like OJAs, and (ii) HyperVec uses hypernym-hyponym relationships for training, while we train our models on a text corpus which has no such relationships. Finally, we considered context embeddings (see e.g. Devlin et al. (2018)). However, contextual embeddings represent words based on their context, thus capturing the uses of words across varied contexts. This is unsuitable for our case, where we aim to compare words in a corpus and their similarity with words of taxonomy in a given sense.

Formal definition of taxonomy. In this section, we introduce a formal definition of taxonomy and formulate the problem of taxonomy alignment, relying on the formalisation proposed by Maedche and Staab (2001). Then, we summarise how `WETA` (Giabelli et al., 2022a) bridges taxonomies through word embeddings.

Definition 1 (Taxonomy)**.** *A taxonomy $\mathcal{T}$ is a 4-tuple $\mathcal{T} = (\mathcal{C}, \mathcal{W}, \mathcal{H}^c, \mathcal{F})$.*

- *$\mathcal{C}$ is a set of concepts $c \in \mathcal{C}$ (aka, nodes) that can be classified in p different hierarchical levels: $\mathcal{C}_1, \ldots, \mathcal{C}_p$;*
- *$\mathcal{W}$ is a set of words (or entities, or leaf concepts) belonging to the domain of interest; each word $w \in \mathcal{W}$ can be assigned to none, one or multiple concepts $c \in \mathcal{C}$.*
- *$\mathcal{H}^c$ is a directed taxonomic binary relation between concepts, that is $\mathcal{H}^c \subseteq \{(c_i, c_j) \mid (c_i, c_j) \in \mathcal{C}^2) \wedge i \neq j\}$. $\mathcal{H}^c(c_1, c_2)$ means that c_1 is a sub-concept, or hyponym, of c_2, while c_2 is the hypernym of c_1, meaning c_2 has a broader meaning and constitutes a category into which c_1 falls. The relation $\mathcal{H}^c(c_1, c_2)$ is also known as $IS - A$ relation (i.e., c_1 $IS - A$ sub-concept of c_2).*
- *$\mathcal{F}$ is a directed binary relation mapping words into concepts, i.e. $\mathcal{F} \subseteq \{(c, w) \mid c \in \mathcal{C} \wedge w \in \mathcal{W}\}$. $\mathcal{F}(c, w)$ means that the word w is an entity of the concept c.*

$\mathcal{T}$ could be represented as a Directed Acyclic Graph (DAG). Therefore, the concepts at the most specific level have an in-degree of 0*, i.e. they don't have any incoming edge. We refer to those concepts as leaf concepts, which are concepts representing different entities or words. Note that in several taxonomies, the terms representing leaf concepts are also item words, while concepts at a higher level are not.*

Given an origin taxonomy $\mathcal{T}_o$ (i.e., PIAAC) and a destination taxonomy $\mathcal{T}_d$ (i.e., ESCO skill pillar), the goal of `WETA` is to suggest one or more concepts $c \in \mathcal{T}_d$ for each word $w \in \mathcal{T}_o$. More specifically, for each $w \in \mathcal{T}_o$, n possible $c \in \mathcal{T}_d$ are suggested based on the scoring function $\mathcal{S}$. More formally:

Definition 2 (Taxonomy Alignment Problem (TAP))**.** *Let $\mathcal{T}_o$ and $\mathcal{T}_d$ be respectively an origin and a destination taxonomy as in Def.1. A Taxonomy Alignment Problem (TAP) is a 3-tuple* $(\psi, h, \mathcal{S})$*, where:*

- $\psi : \mathcal{W}^o \times \mathcal{C}^d \rightarrow [0, 1]$ *is a scoring function that estimates the relevance of $c \in \mathcal{T}_d$ with respect to a word $w \in \mathcal{T}_o$ considering the prediction scores of a multi-class classification task;*
- $h : \mathcal{W}^o \times \mathcal{C}^d \rightarrow [0, 1]$ *is a scoring function that estimates the relevance of $c \in \mathcal{T}_d$ with respect to a word $w \in \mathcal{T}_o$ considering the semantic similarity of w with c and all its hypernyms;*
- $\mathcal{S}(\psi, h) \subseteq \{(w, c) \mid w \in \mathcal{W}^o \wedge c \in \mathcal{C}^d\}$ *is the score of an alignment relation existing between a word in $\mathcal{T}_o$ and a concept in $\mathcal{T}_d$, blending the results of the above mentioned scoring functions.*

A solution to TAP computed over $\mathcal{T}_o$ and $\mathcal{T}_d$ is a 3-tuple $\mathcal{T}_{o,d} = (\mathcal{W}^o, \mathcal{C}^d, \mathcal{S})$

Given a hierarchical level x, we define $\mathcal{C}^o_x$ and $\mathcal{C}^d_x$ as the set of all the $c \in \mathcal{T}_o$ and $c \in \mathcal{T}_d$ respectively, that are classified in level x.

Bridging PIAAC on ESCO. This section describes the global approach used to align the taxonomies $\mathcal{T}_o$ and $\mathcal{T}_d$. The first step allows us to train and select the best word embedding model, which is then used in the second step to suggest for each leaf concept $w_o \in \mathcal{W}^o$ n possible alignments $c_d \in \mathcal{C}^d_p$. The last step consists of the validation of the suggestions because the utility of `WETA` is the help it provides to the domain experts, narrowing the choices for the alignment that would otherwise be done by scratch.

Step 1: Generate and evaluate embeddings. The main goal of the first step of `WETA` is to induce a vector representation of taxonomic terms that represent as much as possible the similarity of words within the taxonomy. To accomplish this, we perform three distinct tasks. We (i) generate word embeddings through a state-of-the-art method; we (ii) compute the `HSS` of terms in $\mathcal{T}_o$ and $\mathcal{T}_d$, and finally, we (iii) select the embeddings for which the correlation between the cosine similarity between taxonomic terms and their `HSS` is maximised for both $\mathcal{T}_o$ and $\mathcal{T}_d$.

Embeddings generation For the generation of the word embedding models, `WETA` employs the state-of-the-art method FastText (Bojanowski et al., 2017), a word embedding method that considers sub-word information and can deal with out-of-vocabulary words.

Selection of the best word embedding To select the best embedding model, we perform an intrinsic evaluation following Baroni et al. (2014). The authors select the word vectors model which has a maximum correlation between their cosine similarity and a benchmark value of semantic similarity. In Baroni et al. (2014), the authors use a handcrafted dataset of pairwise semantic similarity between common words as the gold benchmark. However, those resources usually have low coverage, especially in specific domains like the labor market. For this reason, we resort to a measure of semantic similarity in taxonomies developed in Malandri et al. (2021) and refined in Malandri et al. (2020), which measures semantic similarity in a taxonomy based on the structure of the hierarchy itself without using any external resource, thus, in a sense, preserving the semantic similarity intrinsic to the taxonomy. The `HSS` has proven to be useful in the selection of embeddings for several applications, like taxonomy enrichment (Giabelli et al., 2020, 2021a) and job-skill mismatch analysis in the field of labor market research (Giabelli et al., 2021b).

Since in Malandri et al. (2020) the authors want to encode semantic information from a semantic hierarchy built by human experts, they adopt those values as a proxy of human judgements. Therefore, they compute:

$$\hat{p}(c) = \frac{N_c}{N} \tag{A.2}$$

where N is the cardinality of the taxonomy, and N_c is the cardinality of the concept c and all its hyponyms. Note that $\hat{p}(c)$ is monotonic and increases with granularity.

Intuitively, given two words w_1, w_2 in the taxonomy, $c_1 \in s(w_1)$ and $c_2 \in s(w_2)$ are defined as all the concepts containing w_1 and w_2 respectively, i.e. the *senses* of w_1 and w_2. Therefore, there are $S_{w_1} \times S_{w_2}$ possible combinations of their word senses, where S_{w_1} and S_{w_2} are the cardinality of $s(w_1)$ and $s(w_2)$ respectively. $\mathcal{L}$ is the set of all the lowest common ancestors for all the combinations of $c_1 \in s(w_1)$ and $c_2 \in s(w_2)$. The hierarchical semantic similarity between w_1 and w_2 is defined by the authors of Malandri et al. (2020) as:

$$sim_{\texttt{HSS}}(w_1, w_2) = \sum_{\ell \in \mathcal{L}} \hat{p}\left(\ell = LCA \mid w_1, w_2\right) \times I(LCA) \tag{A.3}$$

where $I(c)$ is the self-information of the concept c and $\hat{p}(\ell = LCA \mid w_1, w_2)$ is the probability of LCA being the lowest common ancestor of w_1, w_2, and it can be computed as follows:

$$\hat{p}\left(\ell = LCA \mid w_1, w_2\right) = \frac{\frac{S_{<w_1,w_2>\in\ell}}{|descendants(\ell)|^2} \times \frac{N_\ell}{N}}{\sum_{k\in\mathcal{L}} \frac{S_{<w_1,w_2>\in k}}{|descendants(k)|^2} \times \frac{N_k}{N}} \tag{A.4}$$

where $S_{<w_1,w_2>\in\ell}$ is the number of pairs of senses of words w_1 and w_2 that have ℓ as lower common ancestor, the term $|descendants(\ell)|$ represents the number of sub-concepts of ℓ, and N_ℓ is the cardinality of ℓ and all its descendants. For more details see Malandri et al. (2020, 2021).

Step 2: Taxonomy alignment method. `WETA` proposes a methodology for taxonomy alignment that suggests, for each word, or leaf concept, $w_o \in \mathcal{T}_o$, a set of n possible destination concept in $\mathcal{T}_d$. The destination concepts are selected among most specialised concepts in $\mathcal{T}_d$, i.e. those which are at the lowest level p, that is $\{c_1, \ldots, c_n\} \in \mathcal{C}_p^d$.

To do this, we perform two different processes that lead to independent results, and then we blend their suggestions to obtain a robust mapping between taxonomies.

Hierarchical approach For each $w_o \in \mathcal{W}^o$, the set of words of the origin taxonomy, we create a list that contains the cosine similarity between w_o and each element in $w \in \mathcal{W}^d$:

$$L_H^{(w_o)} = \{(w, sim(w_o, w)) \mid w \in \mathcal{W}^d\} \; \forall w_o \in \mathcal{W}^o \tag{A.5}$$

where sim is the cosine similarity between the vector representations of the two inputs in the best word embedding model, see Sec. Appendix A.

Given the i-th pair $(w, sim)_i$, we can refer to its similarity as sim_i to define the function W as follows:

$$W(sim_i) = (sim_i)^2 \cdot (sim_i - sim_{i+1}) \tag{A.6}$$

where to transform each similarity score we consider the next similarity in the ordered list. Thanks to W, we can highlight the situations in which, for example, the similarity score between w_o and the first word in $L_H^{(w_o)}$ is significantly higher than the other scores in the ordered list, rather than a situation where all the elements have a high similarity with w_o. Now, we exploit the hierarchical concepts: for each $w \in L_H^{(w_o)}$ we extract its respective hypernym at the level p. We define $L_{Hp}^{(w_o)}$ as the list that contains all the level p hypernyms of every $(w, sim) \in L_H^{(w_o)}$, we order them, and we keep the n with associated the highest similarity. More formally:

$$\begin{gathered} L_{Hp}^{(w_o)} = \{(c_p, h_p) \mid \exists (w, sim) \in L_H^{(w_o)} : \mathcal{F}(c_p, w) \wedge c_p \in \mathcal{C}_p^d\} \; \forall w_o \in \mathcal{W}^o \\ \text{with} h_p = \max_{sim \in S_p} sim, \; S_p = \{sim | \; (w, sim) \in L_H^{(w_o)}\} \end{gathered} \tag{A.7}$$

We keep only the n pairs with the highest similarity score h_p.

Step 3: Evaluation of the suggestions. The usefulness of `WETA` is that it provides a limited number of suggestions to the domain experts to simplify their work of taxonomy alignment that otherwise would be all manual. The last step consists of the validation of the suggestions provided to complete the alignment procedure.

Let us consider the PIAAC question (G_Q05g) that asks workers to what extent they *"use a programming language to program or write computer code"*. Our approach automatically suggests a list of ESCO skills that are likely to be related to the PIAAC question. In this example, our approach proposes (limited to top-5 matches):

- computer programming
- Python (computer programming)
- Java (computer programming)
- software and applications development and analysis
- database and network design and administration

Figure 2 reports the results of the validation phase, made by involving experts of the PILLARS consortium. Each member was asked to vote using a Likert scale to what extent the ESCO suggestions were relevant and consistent with the PIAAC questions reported in Table A.1. We reported mean and standard deviation for each PIAAC skill for users; as you can see the mean value is over 4,6 and the mean of the standard deviation demonstrates that is a quite stable result. From this plot, it is clear that vote scores are concentrated in the upper part of the graph so for most skills there is at least one suggestion with a high level of agreement.

The final list of PIAAC questions that find one or more corresponding ESCO skills is presented in Table A.1.

Table A.1: Selected PIAAC questions and metadata

PIAAC item	Label	Set	Value scheme
F_Q02b	Teaching people	General	Frequency (time units)
F_Q02d	Selling	General	Frequency (time units)
F_Q04a	Influencing people	General	Frequency (time units)
F_Q05a	Simple problems	Problem solving	Frequency (time units)
G_Q01b	Read letters memos or mails	Literacy	Frequency (time units)
G_Q01g	Read financial statements	Literacy	Frequency (time units)
G_Q01h	Read diagrams maps or schematics	Literacy	Frequency (time units)
G_Q02a	Write letters memos or mails	Literacy	Frequency (time units)
G_Q03b	Calculating costs or budgets	Numeracy	Frequency (time units)
G_Q03c	Use or calculate fractions or percentages	Numeracy	Frequency (time units)
G_Q03d	Use a calculator	Numeracy	Frequency (time units)
G_Q03g	Use simple algebra or formulas	Numeracy	Frequency (time units)

Continued on next page

Table A.1 – continued from previous page

PIAAC item	Label	Set	Value scheme
G_Q03h	Use advanced math or statistics	Numeracy	Frequency (time units)
G_Q04	Experience with computer in job	ICT	Yes (1) / No (2)
G_Q05a	For mail	ICT - Internet	Frequency (time units)
G_Q05d	Conduct transactions	ICT - Internet	Frequency (time units)
G_Q05e	Spreadsheets	ICT - Computer	Frequency (time units)
G_Q05f	Word	ICT - Computer	Frequency (time units)
G_Q05g	Programming language	ICT - Computer	Frequency (time units)
I_Q04d	Like learning new things	Learning strategies	Extent of agreement
I_Q04l	Figure out how different ideas fit together	Learning strategies	Extent of agreement

Figure B.1: Skill Demand by Industry

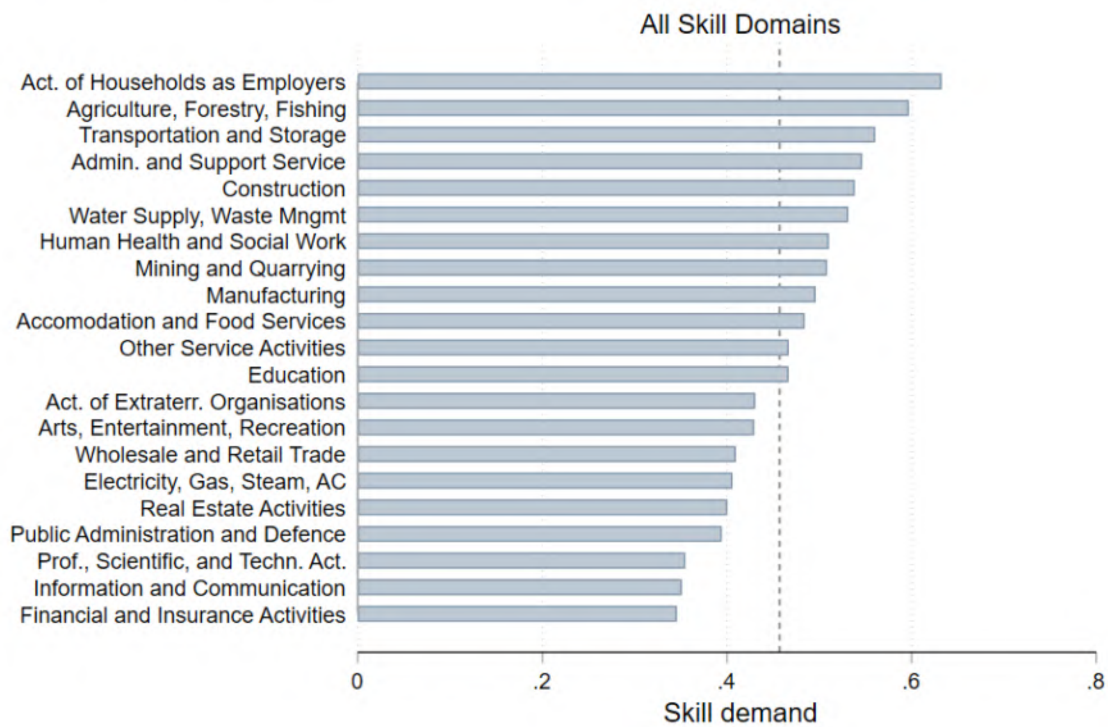

Notes: Skill demand by industry is shown pooled for 17 European countries. The dotted line shows the mean over all industries. Data are from CEDEFOP.

Appendix B. Skill Demand by Industry

Skill demand could not only vary substantially by occupations, but also by industry. To examine this, we first plot the pooled skill demand for all the countries in our sample by sector (NACE1) in Figure B.1. Activities of households as employers shows the highest skill demand, financial and insurance activities the lowest. This does not necessarily mean that industries like information and communication demand less skills. First, it might be that employers are more specific about the skills they seek in these highly technical occupations and thus list a lower variety of skill in their job ads. Second, employers might infer skills from a degree requirement and thus not specifically state them as they deem these skills to be obvious due to the degree requirement.

Second, we investigate potential differences in the demand for certain skills within a given industry. For this reason, we estimate the demand for digital, numeracy, literacy, and social skills separately by skill domain within each NACE1 industry. Figure B.2 shows the results. We find similar skill demand across skill domains within a given industry, again with manual occupations showing higher skill demand than technical occupations. In most of the industries, the pattern are strikingly consistent across skill domains. However, there are also industries with variation in skill demand across skill domains: For instance, in the education industry, demand for literacy skills is higher than, e.g., demand for numeracy skills.

Figure B.2: Skill Demand by Industry and Skill Domain

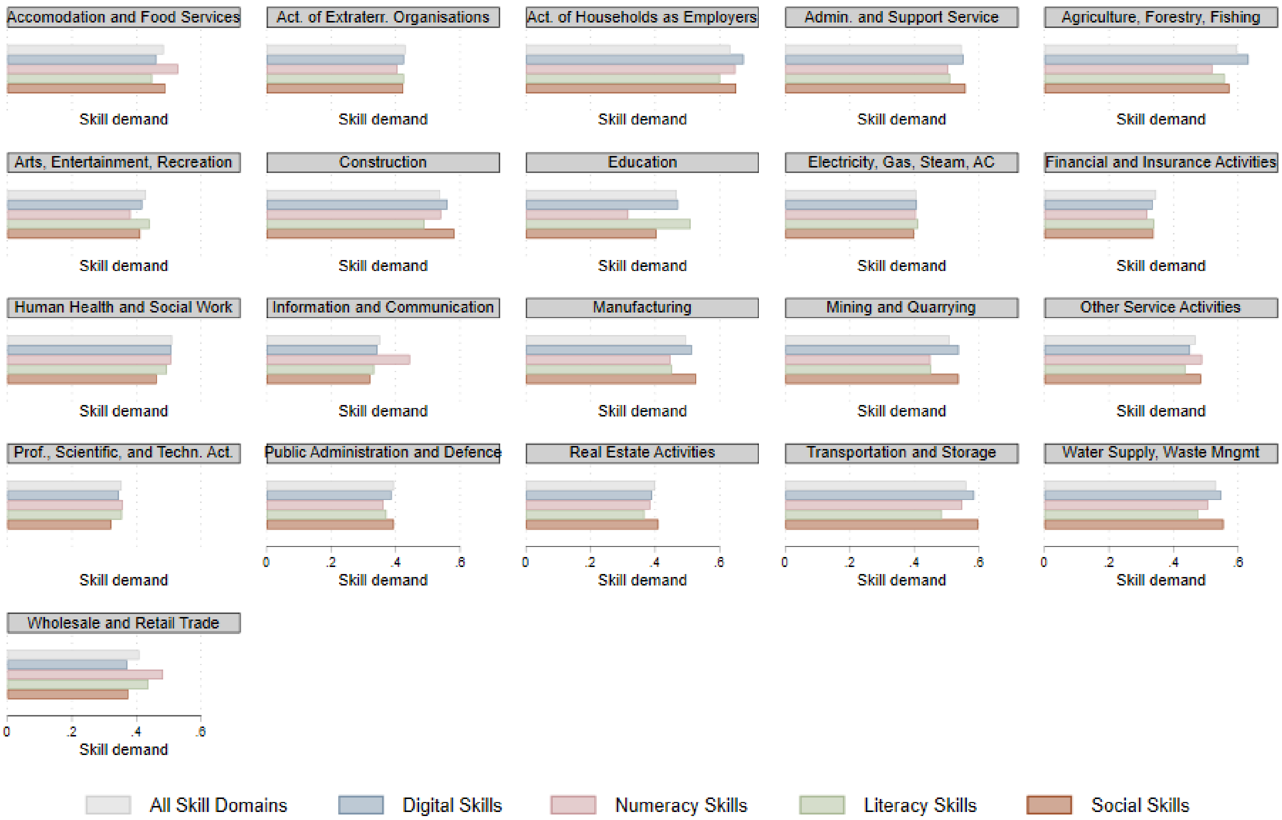

Notes: Skill demand by industry and skill domain is shown pooled for 17 European countries. Data are from CEDEFOP.

Appendix C. Skill Gaps by Country

The pattern of skill mismatch we describe in Section 4 is strikingly consistent across Europe. Figure C.3 plots the difference in skill gaps by occupation types for different European countries. Skill gaps by countries for cognitive non-routine, cognitive routine, manual non-routine, and manual routine occupational types are presented in panels 1, 2, 3, and 4, respectively. The dotted line represents the average skill gap over all countries within an occupation type. We can clearly see that the skill shortage for workers in manual-intensive occupations and the skill surplus for those in cognitive-intensive occupations is persistent across EU countries. For almost all countries, cognitive workers show a skill supply surplus, while manual workers have a skill supply shortage on average. The only exceptions are Sweden for manual non-routine workers and France for cognitive routine and manual non-routine workers.

Figure C.3: Skill Gaps by Country and Occupation Type

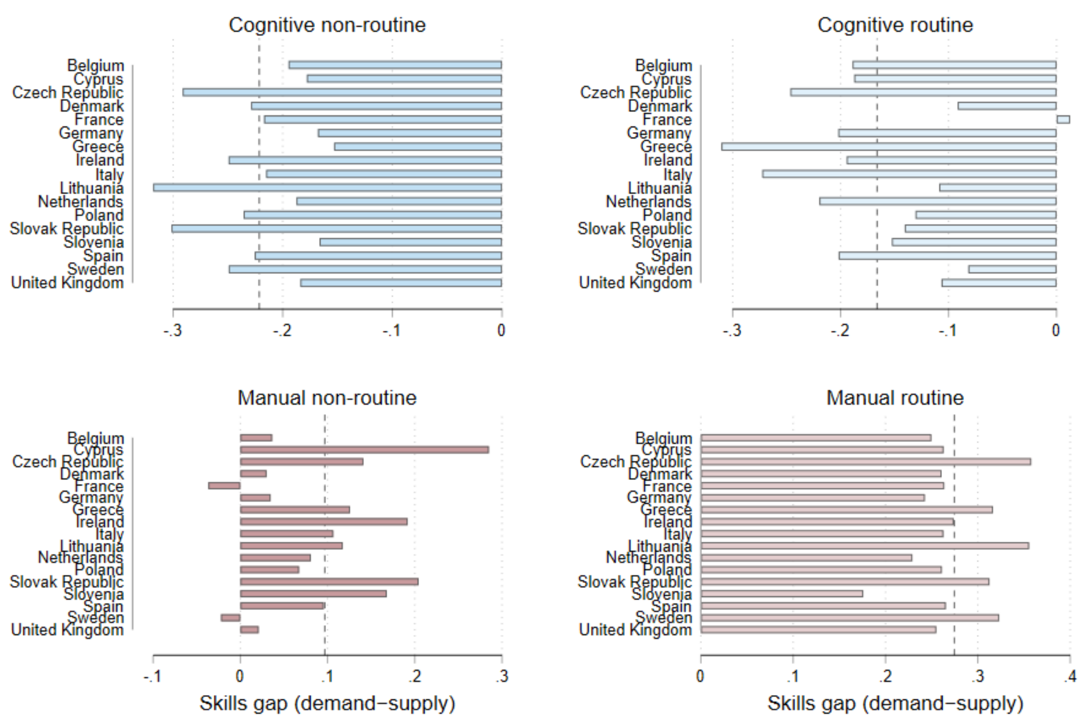

Notes: Skill gap by occupation type – cognitive non-routine, cognitive routine, manual non-routine, and manual routine – are shown for 17 European countries. The dotted line shows the occupation-type-specific mean over all countries. Data are from CEDEFOP and PIAAC.