NETWORK LISER

Discussion Paper Series

IZA DP No. 18881

August 2026

Labor-Market Adjustment to Technological Change: The Role of Occupational Accessibility

Arnaud Dupuy
University of Luxembourg
and IZA@LISER

Morgan Raux
Aix-Marseille School of Economics
and IZA@LISER

Sara Signorelli
CREST - École Polytechnique
and IZA@LISER

IZA DP No. 18881 August 2026

Labor-Market Adjustment to Technological Change: The Role of Occupational Accessibility*

Abstract

Technological change requires workers to reallocate across occupations, but it may also reshape how easily they can do so by changing occupational skill requirements. We exploit digitalization during the 2010s to study this mechanism. We measure occupational accessibility by comparing occupational skill requirements in job-posting data. Combining these measures with French matched employer–employee data, we find that occupations becoming more similar in their digital skill profiles experience greater worker mobility. Counterfactual simulations based on a structural two-sided matching model indicate that changes in occupational accessibility generate worker reallocation amounting to 21 percent of that generated by observed labor-demand shifts.

JEL classification

J23, J24, J62

Keywords

occupation mobility, technological change, matching

Corresponding author

Sara Signorelli

sara.signorelli@polytechnique.edu

* This paper previously circulated under the name "Digitalization, Change in Skill Distance Between Occupations and Occupational Mobility". We would like to thank Philippe Kircher, Jeremias Klaeui, Paul Muller, Emilie Rademakers, Roland Rathelot, and Alexandra Spitz-Oener for their helpful comments, as well as to the participants of the Workshop on Labor Market Transformations organized by ESMT Berlin, the 8th KOF-ETH-IZA Workshop on Matching Workers and Jobs Online and the TASKS VII Conference on The Economic Impacts of AI on Work and Labor Markets. Access to some confidential data, on which is based this work, has been made possible within a secure environment offered by CASD – Centre d'accès sécurisé aux données (Ref. 10.34724/CASD). All remaining mistakes are our own.

1 Introduction

Technological change continuously reshapes the tasks performed by workers. As new technologies diffuse throughout the economy, labor demand decreases in occupations whose tasks are substituted by these technologies and increases in occupations whose tasks complement them (Acemoglu and Restrepo, 2018, Acemoglu and Restrepo, 2019). These changes in the distribution of jobs requires workers to reallocate across the occupational structure. Understanding what makes these adjustments easier or more difficult is therefore central to understanding the labor-market consequences of technological change.

Existing research has largely studied how labor markets adjust to technological change through the lens of changing labor demand, examining how technological change affects the demand for different occupations, the resulting patterns of worker displacement, and the costs of reallocation. Yet, workers do not respond to changing labor demand in a vacuum. Their ability to adjust also depends on the set of alternative occupations they can realistically access. Because workers move more easily between occupations requiring similar skills than between occupations requiring very different ones (Cortes and Gallipoli, 2018), changes in occupational skill requirements alter the opportunities available for worker reallocation. By making occupations more or less similar in the skills they require, technological change reshapes the occupational opportunity set available to workers and, ultimately, their ability to adjust to changing labor demand. Our main finding reveals that changes in occupational accessibility generate worker reallocation amounting to 21% of that induced by labor demand shifts, with a much stronger contribution among high-skilled workers. Importantly, these gains facilitate structural adjustment by disproportionately directing worker flows toward expanding occupations.

To study this mechanism, we exploit the digitalization of the economy during the 2010s, a technological transformation characterized less by the creation of new digital occupations than by the diffusion of foundational digital skills throughout the occupational structure (Tong et al., 2021; ILO, 2022; OECD, 2022). This setting enables us to identify the skills associated with the digital transition and trace how they spread across occupations. We measure these changes in skill requirements by combining two complementary sources of information: the near universe of online job vacancies posted in the United States between 2011 and 2019 and collected by Lightcast, and the evolving taxonomy of work activities assembled by O*NET. Using these data, we develop a time-varying measure of dyadic occupational accessibility that captures how technological change expands or contracts workers' access to alternative occupations by altering the skill distance between them. We start by describing the extent to which digital skills diffused across most occupations in the labor market. We then combine this measure with French matched employer–employee data to estimate workers' mobility responses using the gravity representation of a structural two-sided matching model à la Galichon and Salanié (2024). Measuring occupational accessibility outside the French labor market allows us to relate worker mobility in France to technological changes in occupational skill requirements that are determined independently of French worker reallocation. Finally, we use the estimated model to quantify the extent to which changes in occupational accessibility contribute to labor market adjustment following the digitalization shock, relative to changes in demand for workers' initial occupations.

We first show that digitalization induced widespread convergence in skill requirements across occupations. Tracking occupations relative to the digital skill set of reference defined by IT occupations at the beginning of the decade, we find that digitalization is characterized not simply by the diffusion of technical skills, but by the diffusion of the broader bundle of technical and complementary skills embodied in digital occupations. This process extends well beyond traditionally digital occupations and causes 65 percent of occupation pairs to become more similar in their skill requirements, while only 35 percent become more dissimilar. By bringing occupations closer together in the skill space, digitalization substantially expands the occupational opportunity set available to most workers.

Next, we assess whether these heterogeneous changes in occupational accessibility translate into meaningful differences in workers' occupational mobility. Using a gravity-style framework, we regress pairwise occupational mobility flows on the initial level and subsequent changes in accessibility, controlling for origin and destination fixed effects and bilateral controls. We find that mobility increased between pairs of occupations that became more digitally similar and declined between those that grew more distant, consistent with meaningful changes in transition frictions. This relationship remains economically significant after controlling for overall (non-digital) skill distance across occupations and is robust to alternative measures of accessibility and to alternative specifications.

Finally, our counterfactual analysis shows that the changes in occupational accessibility induced by digitalization substantially increased worker reallocation across all skill groups. The additional mobility generated amounts to 21% of that induced by changes in occupational demand. This average effect, however, masks considerable heterogeneity. The accessibility channel is substantially more important for high-skill workers, where it amounts to 50% of the mobility generated by demand shifts, compared with only around 10% for low-skill workers. Interestingly, for both skill groups, slightly less than half of the increase in mobility induced by improved accessibility is explained by higher match productivity. The remainder operates through residual channels, likely including lower informational and psychological barriers to occupational mobility. Lastly, we show that these mobility gains facilitate structural adjustment by disproportionally directing worker flows toward expanding occupations rather than slowing the reallocation required by demand shifts.

Our paper contributes to the literature on the labor-market consequences of technological change. Existing work has primarily examined how technological change reshapes labor markets by changing the demand for different occupations and has documented the resulting patterns of worker displacement, occupational upgrading, and reallocation (Acemoglu and Restrepo, 2018, 2019). More recent work has focused on the digitalization shock and its implications for workers and firms (Webb, 2019; Deming and Noray, 2020; Acemoglu et al., 2022; Gathmann et al., 2024), as well as on how technological change affects worker flows across occupations (Cortes, 2016; Adão et al., 2024; Battisti et al., 2023; Bessen et al., 2023; Edin et al., 2023; Bocquet, 2024). Closest to our paper, Adão et al. (2024) and Bocquet (2024) show that occupational transition costs slow workers' adjustment to technological shocks, particularly when affected occupations are more skill-distant from alternative employment opportunities. We complement this literature by showing that technological change not only interacts with occupational similarity but also reshapes it. By endogenizing the occupational opportunity set, we identify changes in occupational accessibility as a distinct channel through which technological change affects labor-market adjustment, quantify its contribution relative to changes in labor demand, and characterize how this channel facilitates worker adjustment across occupations.

Our paper also contributes to the literature on occupational mobility and the role of skills in shaping worker transitions (Gathmann and Schönberg, 2010; Lazear, 2009; Yamaguchi, 2012; Poletaev and Robinson, 2008; Lise and Postel-Vinay, 2020; Cortes and Gallipoli, 2018; Azmat et al., 2024; Dabed et al., 2025; Klaeui et al., 2026). This literature models occupations as multidimensional bundles of skills and shows that worker mobility declines with skill distance between occupations. Building on Cortes and Gallipoli (2018), who quantify the effects of fixed occupational skill distances using a gravity framework, we show that occupational similarity itself evolves during periods of rapid technological change. Methodologically, we develop a time-varying measure of occupational accessibility that tracks these changes over time and combine it with a structural two-sided matching model following Galichon and Salanié (2024) to quantify the contribution of occupational accessibility to labor-market adjustment while accounting for congestion effects arising from workers' and firms' joint mobility decisions.

Finally, our model-based distinction between the productivity and non-productivity components of changes in occupational accessibility relates to a growing literature studying the determinants of worker mobility beyond wages. One strand interprets worker mobility toward lower-paying jobs as revealed preference for non-wage amenities.[1] Another emphasizes the role of information frictions, showing that workers may fail to consider suitable opportunities in alternative occupations because they have incomplete information about them. For example, Belot et al. (2026) show that providing job seekers with information about alternative occupations requiring similar skills expands their consideration sets and increases job search toward these occupations.

The remainder of the paper is organized as follows. Section 2 introduces the conceptual framework. Section 3 documents how digitalization reshaped occupational accessibility. Section 4 examines how these changes affected worker mobility. Section 5 quantifies the contribution of occupational accessibility to labor-market adjustment. Section 6 concludes.

2 A Conceptual Framework for Occupational Accessibility and Labor-Market Adjustment

Technological change affects labor-market adjustment through two distinct channels. The first, emphasized in the existing literature, operates through changes in labor demand. As new technologies diffuse throughout the economy, occupations whose tasks complement these technologies expand while occupations whose tasks are substituted by them contract (Acemoglu and Restrepo, 2018; Acemoglu and Restrepo, 2019). These changes create reallocation needs, requiring workers to move from declining occupations

[1] See, for example, Bonhomme and Jolivet (2009); Mas and Pallais (2017); Sorkin (2018); Lavetti (2023).

toward expanding ones.

The second channel, which is the main focus of this paper, operates through changes in occupational accessibility. Workers move more easily between occupations requiring similar bundles of skills than between occupations requiring very different ones (Cortes and Gallipoli, 2018). Occupational accessibility therefore determines the set of occupations that workers can realistically access. We refer to these mobility opportunities as reallocation possibilities.

Figure 1 illustrates how these two channels interact. Panel (a) represents the economy before technological change. Occupations are characterized by their skill requirements and workers move primarily between occupations that are close in skill space, generating a natural rate of occupational mobility. Panel (b) illustrates how technological change reshapes occupational skill requirements. As new technologies diffuse, some occupations become more similar while others become more dissimilar. This reorganizes the occupational skill space and changes workers' mobility opportunities, even if labor demand remains unchanged. Panel (c) illustrates the traditional labor-demand channel. Changes in labor demand alter the relative size of occupations and generate incentives for workers to reallocate, while occupational accessibility remains unchanged. Finally, Panel (d) combines both mechanisms. Technological change simultaneously changes where workers need to move by altering labor demand and whether they can move by reshaping occupational accessibility. Distinguishing these two channels is the central objective of the empirical analysis that follows.

3 Evidence on the Evolution of Occupational Accessibility

This section documents how the diffusion of digital technologies during the 2010s reshaped occupational accessibility. It first explains why digitalization provides a particularly informative empirical setting for studying the mechanism introduced in the previous section and describes how we measure changes in occupational skill requirements over time. It then documents two empirical findings. First, occupations progressively converged toward the IT reference skill profile as digital technologies diffused throughout the economy. Second, this widespread diffusion fundamentally reorganized occupational accessibility, bringing some occupations closer together while moving others further apart.

3.1 Digitalization as an Empirical Setting

The conceptual framework developed in the previous section applies to any episode of technological change that reshapes occupational skill requirements. Empirically, however, studying this mechanism requires overcoming two important challenges. First, it requires detailed and comparable measures of occupational skill requirements over a sufficiently long period to observe how occupations evolve. Second, it requires identifying the skills associated with the technological transformation under study in order to distinguish its effects from broader changes in occupational requirements. These two conditions are rarely met simultaneously, making it difficult to document how technological change reshapes occupational accessibility.

The digitalization of the economy during the 2010s provides a particularly informative empirical setting because it satisfies both conditions. At the beginning of the decade, digital skills were disproportionately concentrated in a relatively small set of digital-intensive occupations. Over the following years, these skills became increasingly important across a much broader range of occupations as digital technologies diffused throughout the economy (Tong et al., 2021; ILO, 2022; OECD, 2022). This initial concentration provides a natural way to identify the bundle of skills associated with the technological shock. Combined with newly available longitudinal data on occupational skill requirements derived from job postings, it allows us to trace the diffusion of these skills across occupations over time and quantify how technological change reshapes occupational accessibility throughout the labor market.

3.2 Measuring Occupational Accessibility

To measure how digitalization reshaped occupational accessibility, we combine two complementary sources of information on occupational skill requirements.

Our primary source is the near universe of online job vacancies posted in the United States, which are collected by Lightcast (formerly Burning Glass Technologies) between 2011 and 2019. These data provide detailed information on the skills employers require when recruiting workers, covering approximately 13,000 distinct skill descriptors mapped to occupations and grouped into 620 skill clusters. Because these skills are extracted directly from job postings, they provide a direct measure of the skills employers seek in the labor market at each point in time and allow us to precisely track the evolution of occupational skill requirements throughout the decade.

We complement these data with the evolving taxonomy of work activities assembled by O*NET. Whereas Lightcast records highly granular employer-reported skill requirements, O*NET organizes occupational content into a consistent set of broader work activities that are assessed by experts. The two sources therefore provide complementary perspectives on occupational requirements. Throughout the paper, our baseline measures build on the Lightcast data, while we assess the robustness of our results by constructing the same measures using O*NET.

In both datasets, each occupation j in year t is represented by a vector of skill intensities z_j^t,

$$z_j^t = \left(z_{j,1}^t, \ldots, z_{j,K}^t \right), \tag{1}$$

where each element measures the importance of a particular skill or work activity within the occupation. In the Lightcast data, skill intensities correspond to the share of job postings for occupation j mentioning each skill in year t. In O*NET, they correspond to the average importance assigned by occupational experts to each work activity. Following Gathmann and Schönberg (2010) and Cortes and Gallipoli (2018), we measure the similarity between occupations using the angular separation between their skill profiles,

$$d\left(z_i^t, z_j^t\right) = \arccos\left(\frac{z_i^t \cdot z_j^t}{\|z_i^t\| \, \|z_j^t\|} \right). \tag{2}$$

Angular separation measures the similarity in the composition of occupational skill requirements independently of their overall intensity and provides a natural measure of occupational similarity.

To isolate the component of occupational similarity associated with digitalization, we exploit the fact that, at the beginning of the decade, the skills associated with digital technologies were disproportionately concentrated in a relatively small set of IT occupations. We define the reference digital skill profile as the average skill profile of the eight detailed PCS occupations classified as IT engineers or IT technicians. This reference profile provides the benchmark against which we measure the diffusion of digital skills across occupations over time. The complete list of occupations used to construct the reference digital skill profile is reported in Table 1.

For every occupation j, we measure its proximity to the IT reference occupations as the negative distance between its skill profile and the corresponding reference profile:

$$q_j^t = -d\left(z_j^t, z_d\right), \tag{3}$$

where z_d denotes the skill profile of the reference IT occupations. Higher values of q_j^t indicate occupations whose skill requirements more closely resemble those of the reference digital occupations. Importantly, the digital reference profile z_d is held fixed over time. To measure the diffusion of digital skills from the IT reference occupations to the rest of the economy, we define the digital reference using the beginning of the period. This avoids a moving benchmark that would otherwise confound skill diffusion with the simultaneous upgrading of skills within IT occupations themselves.

Finally, we characterize occupational accessibility using pairwise distances in digital proximity,

$$d_{ij,t}^{Digital} = \left|q_i^t - q_j^t\right|, \tag{4}$$

so that occupations with similar levels of digital proximity are considered closer along the digital dimension. Appendix A reports additional descriptive evidence on the digital accessibility measure.

3.3 Convergence Toward the Digital Skill Set

We first show that the skills initially concentrated in IT occupations at the beginning of the period progressively diffused throughout the occupational structure during the 2010s. Figure 2 reports the evolution of occupations' distance to the digital skill profile between 2011 and 2019. Because the digital skill set is defined using the skill profile of IT occupations at the beginning of the decade, a decline in this distance indicates that an occupation's skill requirements became more similar to those characterizing the IT reference occupations.

The evidence reveals broad-based convergence toward the digital skill profile. Using the Lightcast measure, 60 percent of occupations experienced a decline in their distance to the IT reference between 2011 and 2019. The corresponding figure reaches 70 percent when occupational skill requirements are measured using O*NET. Together, these results indicate that the skills initially concentrated in IT occupations progressively spread throughout the occupational structure. Further descriptive evidence is provided in Appendix A.

Figure 3 sheds further light on the nature of this convergence by decomposing the reduction in distance to the digital skill set into two margins: the increasing importance of digital skills ("*skill adoption*") and the declining importance of non-digital skills ("*skill abandonment*"). Convergence may arise because occupations adopt digital skills, because they rely less on non-digital skills, or through a combination of both.[2] To illustrate these mechanisms, we report the five skills contributing most to convergence through each margin. The figure shows that convergence reflects both processes. Remarkably, the skills contributing most to convergence through the adoption margin correspond to what could be considered core digital skills in both the Lightcast and O*NET measures, despite the two datasets being constructed from entirely different sources of occupational information. The top two skills contributing to convergence through adoption in the Lightcast data are "SQL" and "Microsoft Office", while they are "Interacting with Computers" and "Processing Information" in the O*NET data. This provides direct evidence that the observed convergence reflects the diffusion of digital technologies rather than unrelated changes in occupational content.

While this evidence establishes that convergence reflects both the adoption of new skills and the abandonment of existing ones, it does not allow us to assess their relative importance in shaping worker mobility. We address this question in the next section using a gravity framework that separately identifies the effects of skill adoption and skill abandonment on occupational mobility.

3.4 Digitalization Reshaped Occupational Accessibility

While the previous subsection examined convergence toward the reference digital skill profile, the question of interest for worker mobility is whether occupations became more similar to one another. Convergence toward the reference digital skill profile does not necessarily imply that occupations become more similar to each other. As occupations adopt digital skills, some occupational pairs may become more similar, whereas others may diverge. We therefore turn to the evolution of pairwise occupational distances, which determine how easily workers can move between occupations.

We find that digitalization substantially reorganized occupational accessibility. Figure

[2]We interpret both margins as manifestations of digitalization. Occupations may become closer to the IT reference profile either by adopting skills characteristic of the IT reference occupations or by reducing the importance of skills that are comparatively less prevalent among them.

4 reports the distribution of changes in pairwise digital distances between occupations between 2011 and 2019. 65 percent of occupation pairs experienced a reduction in their digital distance, while 35 percent moved further apart.

Digitalization therefore did not simply make occupations more alike. Instead, it reorganized the occupational opportunity set by simultaneously creating new mobility opportunities between some occupations while reducing them between others. Technological change thus reshaped labor-market adjustment not only by changing the skill requirements of individual occupations, but also by reorganizing the network of occupations through which workers can reallocate.

4 Occupational Accessibility and Worker Mobility

This section examines whether changes in occupational accessibility translate into changes in worker mobility. We address this question by combining our measures of occupational accessibility with French administrative data on worker transitions. We first describe the data used to measure occupational mobility before introducing a gravity framework to estimate the effect of changing occupational accessibility on worker mobility.

4.1 Measuring occupational mobility

To examine whether changes in occupational accessibility translate into changes in worker mobility, we combine our U.S. measures of occupational accessibility with French matched employer–employee data. Measuring occupational accessibility outside the French labor market substantially reduces concerns that the estimated relationship reflects endogenous changes in worker mobility. This helps us to isolate the role of technological change in shaping occupational accessibility. We map U.S. SOC occupations to the French PCS classification of occupations using the crosswalk described in Appendix B.1.

To measure occupational mobility we rely on the *Déclarations Annuelles de Données Sociales* (DADS), an administrative registry covering the near universe of employees in France used to compute payroll taxes. The DADS records provide detailed information on workers' occupations, earnings, hours worked, employers, and demographic characteristics. The version made available by the French statistical office (INSEE) follows workers over time only for a one-twelfth sample of employees. Babet et al. (2022) show how to reconstruct the quasi-universe of individual employment histories by matching overlapping information across consecutive annual waves. We follow their procedure and observe occupational transitions for nearly the entire French workforce between 2011 and 2019. In addition, for every transition, we observe the origin and destination wage, which we use in the final section of the paper to decompose the occupational accessibility effect into a productivity component and a residual non-productivity component.

Because our objective is to characterize worker reallocation across the entire occupational structure, we impose only minimal sample restrictions. For workers holding multiple jobs within a year, we retain the main employment relationship, defined as the job generating the highest annual earnings. We exclude observations with incomplete occupation codes and hourly wages below the statutory minimum wage, and restrict the sample to workers aged between 20 and 60 years. Importantly, we do not impose any sectoral or occupational restrictions, allowing us to capture worker reallocation across the full set of occupations potentially affected by digitalization.

Using the reconstructed worker panel, we aggregate individual transitions into bilateral mobility flows between the full matrix of occupation pairs. Our analysis includes the 385 occupations of the French PCS classification together with a non-employment state, yielding 148,995 origin–destination occupation pairs.[3] Including both stayers (i.e. workers who stay in the same occupation) and transitions to and from non-employment allows us to characterize worker mobility over the entire labor market, an important feature given the equilibrium interpretation of our gravity framework.

Our benchmark mobility outcome records all occupational transitions observed between 2018 and 2019. By that time, the changes in occupational accessibility documented in Section 3 had largely materialized. We additionally construct a baseline measure using transitions observed between 2011 and 2012, before the widespread diffusion of digital skills. Comparing mobility at the beginning and the end of the decade allows us to relate changes in occupational accessibility to changes in worker mobility and provides an important robustness check for our main results.

Table 2 reports summary statistics for the distribution of worker flows across occupation pairs. Pairwise mobility flows are highly skewed, with many occupation pairs record-

[3]Transitions from non-employment to non-employment are not observed in employer–employee data and are therefore excluded from the analysis.

ing no transitions and median flows substantially below the mean. Total occupational mobility also declined slightly over the period, with the average number of bilateral transitions falling from 217 between 2011 and 2012 to 189 between 2018 and 2019.

4.2 Gravity Framework

Our reduced-form approach relies on estimating a gravity model relating bilateral worker flows between occupations to changes in occupational accessibility. Gravity models provide a natural framework for identifying the determinants of bilateral worker mobility. By controlling for all origin- and destination-specific factors affecting occupational transitions, they isolate the role of pair-specific characteristics, such as occupational dyadic distance, in shaping worker flows. Following Cortes and Gallipoli (2018), we interpret occupational distance as increasing the frictions workers face when moving across occupations.

A conventional gravity specification relates worker mobility observed at the end of the period to occupational distance measured at the same date. In our setting, however, end-of-period digital distance can be naturally decomposed into its initial level and the change that occurred over the decade,

$$d_{ij,2019}^{Digital} = d_{ij,2011}^{Digital} + \Delta d_{ij}^{Digital}, \tag{5}$$

where Δ denotes the change between 2011 and 2019. Estimating these two components separately is central to our empirical strategy. The initial distance, $d_{ij,2011}^{Digital}$, captures persistent differences in digital accessibility across occupation pairs, whereas $\Delta d_{ij}^{Digital}$ isolates the variation arising from the diffusion of digital technologies over the decade. Our main parameter of interest is therefore the effect of changes in digital accessibility on worker mobility, conditional on initial occupational distance.

Our empirical specification is therefore given by:

$$Mobility_{ij,2019} = \exp\Big(\beta_1\, d_{ij,2011}^{Digital} + \beta_2\, \Delta d_{ij}^{Digital} + X_{ij}'\delta + \gamma_i + \gamma_j + \varepsilon_{ij,2019}\Big), \tag{6}$$

where $Mobility_{ij,2019}$ denotes the number of workers moving from occupation i to occupation j between 2018 and 2019. Our main coefficient of interest is β_2, which measures whether occupation pairs that experienced larger improvements in digital accessibility over the decade exhibit greater worker mobility at the end of the period, conditional on their initial level of digital accessibility. Because the model controls for initial digital accessibility, β_2 is identified exclusively from cross-sectional differences in cumulative changes in digital accessibility generated by the diffusion of digital technologies between 2011 and 2019. Origin and destination fixed effects absorb time-invariant occupation-specific determinants of worker mobility, including differences in labor demand, labor supply, occupational size, and wage levels. We estimate the model using Poisson pseudo-maximum likelihood (PPML), the standard estimator for gravity models. PPML naturally accommodates the large number of zero bilateral flows documented in Table 2. Standard errors are double-clustered by origin and destination occupations.

Our specification allows the initial level of occupational distance and its subsequent change to have different associations with worker mobility. Rather than imposing that the initial level of occupational distance and its subsequent change have the same association with worker mobility, we estimate the two coefficients separately. This flexibility is central to our empirical strategy because it allows us to isolate the contribution of technological changes in occupational accessibility while controlling for persistent differences in occupational similarity.

As a robustness check, we exploit the panel dimension of the data to examine whether changes in occupational accessibility are associated with changes in worker mobility. Specifically, we estimate a bilateral fixed-effects model using observations from the beginning (2011–2012) and the end (2018–2019) of the decade:

$$Mobility_{ij,t} = \exp\left(\alpha_1 \, d_{ij,t}^{Digital} + X'_{ij,t}\omega + \gamma_{ij} + \varepsilon_{ij,t}\right), \tag{7}$$

where bilateral fixed effects γ_{ij} absorb all time-invariant determinants of worker mobility between occupation pairs. Consequently, α_1 is identified from within-pair changes in both digital occupational accessibility and worker mobility between 2011 and 2019. It therefore measures whether occupation pairs experiencing larger changes in occupational accessibility also experienced larger increases in worker mobility over the decade.

The identifying assumptions underlying the cross-sectional and panel specifications differ, but both require that the remaining time-varying bilateral determinants of worker mobility be adequately controlled for. The composition of the bilateral controls, $X_{ij,t}$, is therefore central to our empirical strategy and deserves careful discussion.

4.3 Identification strategy

Identifying the effect of digitalization on worker mobility requires distinguishing changes in occupational similarity induced by digitalization from the many other forces simultaneously reshaping occupations over the decade. Our empirical strategy addresses this challenge through two complementary elements. First, we exploit variation in digital occupational distance measured independently of the French labor market. Second, we control for changes in occupational similarity unrelated to digitalization.

Our first source of identification exploits the fact that digital occupational distances are constructed exclusively from U.S. job-posting data and O*NET. Their evolution is therefore measured independently of worker mobility in France and plausibly captures technological shocks common across developed countries. By exploiting variation in occupational skill requirements outside the French labor market, this approach mitigates concerns that digital occupational accessibility and worker mobility are jointly determined.

The second component of our identification strategy recognizes that digitalization is only one of several forces capable of reshaping occupational similarity over time. Occupations may also become more similar because of changes in recruitment practices, management methods, regulation, or broader transformations in the organization of work. Ignoring these alternative sources of convergence could therefore lead us to attribute some of their effects on worker mobility to digitalization.

A natural solution would be to control for changes in overall occupational distance. This approach is problematic, however, because overall skill convergence mechanically incorporates the digital convergence that constitutes our object of interest. Occupation pairs becoming more similar in terms of digital skills also become more similar in their overall skill profiles. Conditioning on the raw change in overall occupational distance would therefore absorb part of the variation generated by digitalization itself, creating a bad-control problem.

Our objective is consequently not simply to control for changing occupational similarity, but to isolate the component of occupational convergence unrelated to digitalization. We construct a non-digital occupational distance measure by combining two approaches, each designed to remove a different component of digital convergence.

The first approach restricts the skill space to skills that are not disproportionately used by IT occupations. This removes convergence generated directly by the adoption of digital skills. Digitalization, however, diffuses together with a broader set of complementary skills that are also disproportionately represented in IT occupations. Because these complementary skills remain in the restricted skill space, this approach alone would continue to absorb part of the variation generated by digitalization.

The second approach is geometric. Rather than removing individual skills on the basis of their classification, we project the average skill profile of IT occupations in 2011 out of each occupational skill vector before computing bilateral distances. This removes convergence along the average direction associated with digitalization. Because IT occupations are themselves heterogeneous, however, a single projection does not eliminate every dimension of digital convergence. Differences across IT occupations—for example, between software developers and network engineers—may continue to generate residual variation correlated with digitalization.

Our preferred non-digital distance measure combines the two approaches. We first remove all skills disproportionately used by IT occupations, thereby eliminating convergence driven directly by digital skill adoption. Within the remaining skill space, we then project out the average non-digital skill profile of IT occupations in 2011. This second operation removes the complementary skills that systematically accompany digitalization while preserving other dimensions of occupational convergence. The resulting measure captures changes in occupational similarity unrelated to digitalization without absorbing the variation through which digitalization reshapes occupational accessibility. Appendix B.2 provides a detailed description of the construction of this non-digital accessibility measure.

In all specifications, the vector X_{ij} includes both the initial level and the change in this non-digital distance measure. It also includes bilateral controls for remaining in the same occupation and for transitions involving changes in socioeconomic status, which capture mobility frictions beyond occupational similarity. Our identifying variation therefore comes from comparing occupation pairs that experienced similar changes in non-digital occupational similarity but different changes in digital occupational distance. This comparison isolates the contribution of digitalization-induced changes in occupational accessibility to worker mobility.

4.4 Main Results

Table 3 presents the main estimates of the relationship between occupational accessibility and worker mobility. Because the model is estimated by Poisson pseudo-maximum likelihood (PPML), the reported coefficients are incidence-rate ratios obtained by exponentiating the estimated coefficients. Values below one therefore correspond to a negative effect and values above one to a positive one. Since the regression relates bilateral worker flows to occupational distance, coefficients smaller than one imply that occupations becoming closer in the digital skill space (i.e. becoming more accessible) exchange more workers. Finally, to ease interpretation, we standardize all distance measures to have mean zero and a standard deviation of one, such that the coefficients can be interpreted as the effect of a one standard deviation increase in distance on percentage mobility flows.

Column 1 reports the baseline gravity specification including the initial level and the change in digital distance. The estimated coefficient on the change in digital distance is significantly below one, indicating that reductions in digital distance between occupations are associated with higher worker mobility. This suggests that the changes in occupational accessibility documented in the previous section translated into changes in workers' occupational opportunity set and, ultimately, into actual worker reallocation across occupations.

Columns 2 and 3 examine whether changes in digital occupational accessibility predict worker mobility beyond broader changes in occupational similarity. This is a considerably more demanding test because it asks whether digital convergence has explanatory power over and above the general evolution of occupational skill similarity. Column 2 additionally controls for the level of overall occupational distance in 2019, computed using the full set of skills, including both digital and non-digital ones. Column 3 replaces the overall occupational distance in 2019 with its initial level in 2011 and its change between 2011 and 2019, mirroring the decomposition used for digital occupational distance. The estimated coefficient associated with the change in digital distance decreases in magnitude but remains statistically significant across these specifications. In column 3, a one standard deviation reduction in digital occupational distance over the period is associated with approximately 10% higher worker mobility in 2019. Notably, the estimated coefficient on the initial level of digital distance in 2011 is very similar in magnitude to that on the change in digital distance.

As discussed in the previous subsection, however, controlling for the raw change in overall occupational distance raises an important identification concern. Because overall skill convergence mechanically incorporates convergence generated by digitalization itself, this specification conditions on part of the treatment and therefore risks understating the true effect of digitalization.

Columns 4 through 6 address this concern by progressively replacing the raw overall distance with the non-digital control measures introduced in Section 4.3. Column 4 excludes skills disproportionately used by IT occupations, thereby removing convergence driven directly by digital skill adoption. Column 5 instead projects out the average IT direction in skill space before computing occupational distance. Column 6 combines both approaches, removing both IT-intensive skills and the complementary dimensions of convergence associated with digitalization. The estimated coefficient on digital distance change remains stable across all these specifications, at a level of about 10-12% additional mobility per standard deviation reduction in digital skill distance. This stability indicates that the relationship between digital convergence and worker mobility is robust to controlling for broader changes in occupational similarity unrelated to digitalization.

Finally, Columns 7 and 8 demonstrate that the results are robust to alternative measures of occupational skill requirements. Column 7 reconstructs both digital and non-digital distances using O*NET work activities instead of Lightcast job postings. Column 8 builds the measures using the 620 aggregated Lightcast skill categories rather than the full vocabulary of roughly 13,000 individual skills. Both specifications produce estimates that are quantitatively and statistically similar to our favorite specification reported in column 6. This confirms that the results are not sensitive to the measurement of occupational skills. Appendix figure A.8 further summarizes the stability of the estimated coefficients across a broader set of alternative definitions of both the digital and non-digital accessibility measures.

The gravity estimates establish the existence of the mechanism proposed in this paper: occupations becoming more accessible because of digitalization systematically exchange more workers. However, the economic magnitude of these estimates is not directly interpretable, as the reported incidence-rate ratios quantify changes in bilateral worker flows rather than the contribution of occupational accessibility to aggregate labor-market adjustment. We therefore turn in the next section to a structural matching framework, which uses the gravity estimates to quantify how much of the observed labor-market reallocation can be attributed to changes in occupational accessibility induced by digitalization.

4.5 Robustness

Table 4 examines the robustness of our main findings, taking the specification reported in Column 6 of Table 3 as the benchmark. Across all specifications, the estimated effect of digital distance remains remarkably stable, confirming that the relationship between occupational accessibility and worker mobility is robust to sample restrictions, alternative measures of occupational distance, and the empirical specification.

The first set of robustness checks examines whether the results are driven by specific types of occupational transitions. Column 2 excludes stayers, that is, workers remaining in the same occupation, while Column 3 excludes transitions to non-employment.These observations are retained in the benchmark specification to capture all labor market transitions and to remain consistent with the structural matching model. However, they provide no identifying variation for the effect of changes in occupational accessibility because digital distance is mechanically equal to zero in all periods. We therefore verify that our results are robust to excluding them. In both cases, excluding these observations leaves the estimated coefficient on digital distance virtually unchanged, indicating that our results are not driven by occupational persistence or transitions into and out of employment.

Column 4 considers an alternative measure of digital distance motivated by a different assumption about workers' skill accumulation. Our benchmark specification implicitly assumes that workers continuously update their skills as the requirements of their occupation evolve over time. Accordingly, changes in occupational accessibility are measured using the contemporaneous skill profiles of both the origin and destination occupations, while the digital reference remains fixed at its 2011 definition. An alternative view is that workers retain the skill bundle they acquired when entering their occupation and experience technological change as the obsolescence of these initial skills. To capture this alternative mechanism, we construct an obsolescence measure that holds the origin occupation fixed at its 2011 skill profile while allowing only the destination occupation to evolve. This specification also implies that stayers experience changes in digital distance as the skill requirements of their occupation evolve over time. The estimated effect remains statistically and economically similar to the baseline, indicating that our findings are robust to alternative assumptions about how workers' skills evolve over their careers.

Finally, Columns 5 and 6 examine the robustness of the results to alternative empirical specifications. Column 5 controls for baseline worker mobility observed between 2011 and 2012, ensuring that the estimated relationship is not driven by persistent mobility patterns across occupation pairs. Column 6 instead estimates the dyadic panel fixed-effects specification in equation 7, which relates within-pair changes in worker mobility to within-pair changes in occupational accessibility while absorbing all time-invariant pair-specific determinants of worker mobility. The estimated coefficient remains statistically significant and becomes even larger in magnitude, confirming that the baseline relationship is not driven by persistent differences across occupation pairs but instead reflects a strong association between changes in occupational accessibility and changes in worker mobility.

These robustness exercises reinforce the central finding of the paper: occupations becoming more accessible because of digitalization consistently exchange more workers, regardless of the sample considered, the assumptions made about workers' skill accumulation, or the empirical specification. Further robustness analyses are reported in Appendix A.

4.6 Channel

As discussed in Section 3, digitalization reshapes occupational accessibility through two distinct margins: the adoption of new skills and the abandonment of existing ones. Understanding the relative importance of these two channels for worker mobility provides insight into the mechanisms through which technological change reshapes occupational accessibility. We examine their contributions separately in Table 5.

To do so, we decompose changes in digital distance into two components. Our measure of skill adoption is constructed exactly as our baseline measure of digital distance, except that it is computed using only the skills whose usage increased between 2011 and 2019 in both the origin and destination occupations. It therefore captures the extent to which two occupations became closer because they jointly adopted new digital skills. Conversely, our measure of skill abandonment is computed using only the skills whose usage declined in both occupations over the same period, after excluding skills disproportionately used by the IT reference occupations in 2011. It therefore captures convergence arising from the joint disappearance of previously required non-digital skills. Finally, to verify that the adoption channel is genuinely driven by digitalization rather than by the adoption of new skills more generally, we construct a placebo measure of skill adoption by applying the same procedure after excluding all skills disproportionately used by the IT reference occupations in 2011.

Column 1 replaces the overall change in digital distance in our benchmark specification with the separate measures of skill adoption and skill abandonment. Only the coefficient associated with skill adoption is significantly below one, while the coefficient on skill abandonment is small and statistically insignificant. Columns 2 and 3 confirm this result by introducing each channel separately. The effect of skill adoption remains virtually unchanged when estimated on its own, whereas skill abandonment continues to exhibit no statistically significant relationship with worker mobility.

Finally, Column 4 replaces the digital skill-adoption measure with its placebo counterpart. The estimated coefficient becomes statistically insignificant, indicating that the adoption channel documented in the previous columns is specific to the diffusion of digital skills rather than the adoption of new skills more generally.

These results suggest that the effect of digitalization on occupational accessibility operates primarily through the adoption of digital skills rather than through the abandonment of obsolete skills. This finding provides direct evidence for the mechanism proposed in the paper: workers reallocate because occupations acquire common digital skills that reduce mobility frictions, not because they jointly abandon declining tasks.

5 Quantifying the Accessibility Channel

The gravity estimates establish that digitalization-induced improvements in occupational accessibility increase worker mobility. However, the estimated coefficients alone do not quantify the aggregate importance of this mechanism. Improvements in occupational accessibility affect not only bilateral worker flows but also the equilibrium allocation of workers and wages throughout the labor market. To account for these general-equilibrium interactions, we embed the estimated gravity relationship into a structural two-sided matching model. We then use the model to answer two questions. First, how important was the occupational-accessibility channel for worker reallocation during the 2010s? Second, how did improvements in occupational accessibility reshape the labor-market adjustment generated by changes in occupational demand?[4]

5.1 Structural Two-Sided Matching Model

To quantify the contribution of occupational accessibility to labor-market adjustment, we model the labor market as a competitive two-sided matching market with transferable utility following Choo and Siow (2006) and Galichon and Salanié (2024). Workers are characterized by their occupation of origin, which determines the bundle of skills they possess, while firms are characterized by the occupation they seek to fill, which determines the bundle of skills they require. Occupational transitions arise from decentralized matching between workers and firms, where the surplus generated by each match depends on occupational accessibility, that is, the ease with which workers' skills can be transferred across occupations.

Let $Mobility_{ij}$ denote the number of workers moving from occupation i to occupation j. In equilibrium, worker flows satisfy the market-clearing conditions:

$$\sum_j Mobility_{ij} = L_i, \tag{8}$$

$$\sum_i Mobility_{ij} = E_j, \tag{9}$$

where L_i denotes the supply of workers initially employed in occupation i and E_j denotes labor demand in destination occupation j.[5] These conditions ensure that every worker is matched to a destination while every occupation receives the required number of workers.

Under the assumptions of the model, equilibrium worker flows satisfy the gravity equation:

$$Mobility_{ij} = \exp\left(\frac{S_{ij} - u_i - v_j}{2}\right), \tag{10}$$

[4]The section below presents the main equations of the model in their simplest form. See Appendix C for a more detailed presentation of the model.

[5]Recall that one occupational category corresponds to non-employment. The model therefore does not impose full employment, but instead allows worker flows between employment and non-employment to vary endogenously.

where S_{ij} denotes the systematic surplus generated by matching workers from occupation i with firms recruiting in occupation j, while u_i and v_j are equilibrium terms that adjust to satisfy the market-clearing conditions above. The bilateral surplus summarizes the gains generated by each worker-firm match, reflecting workers' preferences over occupations, match productivity, and occupational accessibility. Equation (10) provides the microfoundation for the gravity specification estimated in Section 4. In particular, the origin and destination fixed effects capture the equilibrium forces that balance worker supply and labor demand across occupations.

Occupational accessibility affects worker mobility by increasing the surplus generated by matches between workers and occupations requiring more similar skills. To connect the structural model to the reduced-form evidence, we parameterize bilateral surplus using the same determinants of worker mobility estimated in the gravity model:

$$S_{ij} = \beta_1 d_{ij,2011}^{Digital} + \beta_2 \Delta d_{ij}^{Digital} + \beta_3 d_{ij,2011}^{Non-digital} + \beta_4 \Delta d_{ij}^{Non-digital} + X_{ij}'\delta, \tag{11}$$

where the coefficients are taken directly from the preferred specification estimated in Section 4. Improvements in occupational accessibility therefore increase the surplus generated by worker–firm matches and, in equilibrium, raise worker mobility between the corresponding occupations. The remaining bilateral controls capture other determinants of occupational mobility unrelated to digitalization.

Before turning to the counterfactual analysis, we assess the model's ability to reproduce the observed pattern of worker mobility. Appendix Figure C.1 compares the bilateral worker flows predicted by the model with those observed in the data. By construction, the model matches the observed occupational employment shares but does not impose the size of individual bilateral mobility flows. Its ability to reproduce these bilateral flows therefore provides a meaningful validation of the model. The fitted relationship closely tracks the data, with a correlation coefficient of nearly 0.5, although the model tends to slightly over-predict smaller flows and under-predict larger ones.

The estimated surplus function, together with the market-clearing conditions, allows us to compute counterfactual labor-market equilibria. In these counterfactuals, occupational accessibility is held fixed at its 2011 level, while the observed evolution of labor demand is left unchanged. Comparing the resulting equilibria with the observed equilibrium isolates the contribution of occupational accessibility to worker reallocation during the 2010s while fully accounting for the general-equilibrium interactions generated by the matching market.

Following the decomposition proposed by Dupuy and Galichon (2022), an additional implication of the transferable-utility framework is that the effect of occupational accessibility on worker mobility can be decomposed into a productivity component and a residual non-productivity component.[6] Using data on wages for workers moving from occupation i to occupation j, we recover the component of bilateral surplus associated with match productivity as:

$$Productivity_{ij} = Wages_{ij} + 1/2 \times (S_{ij} + v_j - u_i). \tag{12}$$

The remaining component $S_{ij} - Productivity_{ij}$ captures the increase in worker mobility associated with improvements in occupational accessibility that is not explained by higher match productivity. We interpret this residual as reflecting non-pecuniary mechanisms that facilitate occupational mobility, including lower informational frictions, reduced psychological barriers to occupational transitions, and other factors that encourage workers to move across occupations without increasing the productivity of worker-job matches.

5.2 How Important Is the Occupational-Accessibility Channel?

We assess the importance of the occupational-accessibility channel by computing a series of counterfactual equilibria that isolate the contribution of changes in occupational accessibility and labor demand to worker reallocation. The first counterfactual isolates the accessibility effect by holding digital skill distance at its 2011 level while allowing occupational employment shares to evolve as observed between 2011 and 2019. Following the literature on technological change that has primarily emphasized shifts in labor demand as the main driver of worker reallocation, we benchmark the accessibility effect against the worker reallocation generated by changes in occupational employment shares. We refer to this contribution as the demand effect. This second counterfactual isolates the demand effect by holding occupational employment shares fixed at their 2011 level while allowing occupational accessibility to evolve as observed. We compare both counterfactuals to a common baseline equilibrium in which neither occupational accessibility nor occupational employment shares change over the period.

[6] See Appendix C for more details.

Because the model is nonlinear, the worker reallocation generated by changes in occupational accessibility depends on whether it is evaluated before or after changes in labor demand. As a result, the contribution attributed to each factor depends on the order in which the counterfactuals are implemented. We address this path dependence using the Shapley decomposition, which averages each factor's contribution across all possible elimination orders.

Our simulations reveal that changes in occupational accessibility generated an additional 3.3% of occupational mobility relative to the baseline equilibrium. By comparison, the demand effect increased mobility by 15%. Occupational accessibility therefore amounts to approximately 21% of the worker reallocation generated by the demand effect, indicating that it constitutes a quantitatively important channel of labor-market adjustment.

This comparison should nevertheless be interpreted as conservative. Ideally, we would benchmark the accessibility effect generated by technological change against the worker reallocation induced by the corresponding demand shifts attributable to technological change. Because this component cannot be separately identified, we instead compare it with the overall worker reallocation associated with changes in occupational employment shares between 2011 and 2019. Because changes in occupational employment shares reflect not only the labor-demand consequences of technological change but also a wide range of other structural forces, this benchmark is likely larger than the labor-demand effect attributable to technological change alone. As a result, our estimate likely understates the quantitative importance of the occupational-accessibility channel relative to the demand effect attributable to technological change alone.

The quantitative importance of occupational accessibility varies substantially across skill groups, as illustrated in Figure 5. Improvements in occupational accessibility generate substantially more worker reallocation among high-skilled occupations, increasing outflows by slightly more than 5%, compared with about 2% among low-skilled occupations (middle panel). By contrast, changes in occupational employment shares induce considerably larger worker reallocation among low-skilled occupations (left panel). As a result, occupational accessibility amounts to roughly half of the demand effect among high-skilled occupations but only around one tenth among low-skilled occupations (right panel). These findings indicate that occupational accessibility is a particularly important driver of worker reallocation among high-skilled workers.

The middle panel of the figure further decomposes the accessibility effect into a productivity component and a residual non-productivity component. For both skill groups, slightly less than half of the increase in mobility is explained by improvements in match productivity. The remaining effect is captured by the residual non-productivity component, indicating that occupational accessibility facilitates worker reallocation through mechanisms beyond productivity improvements. Among the mechanisms potentially captured by this residual is a reduction in informational frictions about alternative careers. For example, Belot et al. (2026) show that providing job seekers with information about alternative occupations requiring similar skills expands their consideration sets and increases job search toward these occupations.

We next examine whether improvements in occupational accessibility facilitate labor-market adjustment to changing labor demand. Figure 6 classifies occupations into three groups according to changes in their employment shares over the period: growing occupations (employment share increases by more than 10%), stable occupations (employment share changes by less than 10%), and shrinking occupations (employment share decreases by more than 10%). The results show that improvements in occupational accessibility primarily increase worker mobility toward growing occupations. Accessibility raises transitions from shrinking to growing occupations by 4.1% and from stable to growing occupations by 5.2%, compared with increases of only 2.3% for transitions from growing to shrinking occupations and 2.6% from stable to shrinking occupations. These findings indicate that improvements in occupational accessibility facilitate labor-market adjustment by easing the reallocation of workers from shrinking to growing occupations.

6 Conclusion

Technological change is viewed as reshaping labor markets by changing the demand for workers across occupations. This paper argues that another important, but largely overlooked, channel operates through occupational accessibility. As occupations require new skills, they become more or less accessible to workers from other occupations, altering the ease with which labor can reallocate in response to changing employment opportunities.

Using the diffusion of digital technologies during the 2010s as an empirical setting, we introduce a measure of occupational accessibility based on the similarity of occupational skill requirements. We first document that digital skills spread far beyond traditionally

IT-intensive occupations, substantially reorganizing the occupational landscape. We then show that occupations becoming more accessible because of digitalization exchanged more workers, even after controlling for broader changes in occupational similarity. Finally, embedding these estimates into a structural two-sided matching model, we show that the worker reallocation induced by changes in occupational accessibility amounts to 21% of that induced by observed changes in labor demand over the decade. This contribution is substantially larger for high-skilled occupations, where the accessibility channel reaches roughly one-half the size of the demand channel. Moreover, improvements in occupational accessibility facilitate worker reallocation toward growing occupations, particularly by increasing transitions from shrinking to growing occupations.

More broadly, our findings suggest a different way of thinking about the labor-market consequences of technological change. Technological change reshapes labor markets along two distinct dimensions: it changes where workers are needed, but it also changes how easily workers can move to where they are needed. The speed and magnitude of labor-market adjustment therefore depend jointly on changes in labor demand and changes in occupational accessibility. Focusing exclusively on the former risks overlooking an important determinant of workers' ability to adapt to structural transformation.

References

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

Acemoglu, D. and P. Restrepo (2018). Artificial intelligence, automation, and work. In *The economics of artificial intelligence: An agenda*, pp. 197–236. University of Chicago Press.

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

Adão, R., M. Beraja, and N. Pandalai-Nayar (2024). Fast and slow technological transitions. *Journal of Political Economy Macroeconomics 2*(2), 183–227.

Azmat, G., L. Behaghel, Y. Hazard, R. Rathelot, and J. Sultan (2024). Occupational mobility and retraining: Experimental evidence on firms' hiring preferences.

Babet, D., O. Godechot, and M. G. Palladino (2022). In the land of akm: Explaining the dynamics of wage inequality in france.

Battisti, M., C. Dustmann, and U. Schönberg (2023). Technological and organizational change and the careers of workers. *Journal of the European Economic Association*, jvad014.

Belot, M., P. Kircher, and P. Muller (2026). Do the long-term unemployed benefit from automated occupational advice during online job search? *The Economic Journal 136*(673), 184–206.

Bessen, J., M. Goos, A. Salomons, and W. Van den Berge (2023). What happens to workers at firms that automate? *The Review of Economics and Statistics*, 1–45.

Bocquet, L. (2024). The network origin of slow labor reallocation.

Bonhomme, S. and G. Jolivet (2009). The pervasive absence of compensating differentials. *Journal of Applied Econometrics 24*(5), 763–795.

Chen, L., E. Choo, A. Galichon, and S. Weber (2021). Matching function equilibria with partial assignment: Existence, uniqueness and estimation. *arXiv preprint arXiv:2102.02071*.

Choo, E. and A. Siow (2006). Who marries whom and why. *Journal of Political Economy 114*(1), 175–201.

Cortes, G. M. (2016). Where have the middle-wage workers gone? a study of polarization using panel data. *Journal of Labor Economics 34*(1), 63–105.

Cortes, G. M. and G. Gallipoli (2018). The costs of occupational mobility: An aggregate analysis. *Journal of the European Economic Association 16*(2), 275–315.

Dabed, D., S. Genz, and E. Rademakers (2025). Equalising the effects of automation? the role of task overlap for job finding. *Labour Economics*, 102766.

Deming, D. J. and K. Noray (2020). Earnings dynamics, changing job skills, and stem careers. *The Quarterly Journal of Economics 135*(4), 1965–2005.

Dupuy, A. and A. Galichon (2022). A note on the estimation of job amenities and labor productivity. *Quantitative Economics 13*(1), 153–177.

Edin, P.-A., T. Evans, G. Graetz, S. Hernnäs, and G. Michaels (2023). Individual consequences of occupational decline. *The Economic Journal 133*(654), 2178–2209.

Galichon, A. and B. Salanié (2022). Cupid's invisible hand: Social surplus and identification in matching models. *The Review of Economic Studies 89*(5), 2600–2629.

Galichon, A. and B. Salanié (2024). Estimating separable matching models. *Journal of Applied Econometrics 39*(6), 1021–1044.

Gathmann, C., F. Grimm, and E. Winkler (2024). Ai, task changes in jobs, and worker reallocation. Technical report, CESifo Working Paper.

Gathmann, C. and U. Schönberg (2010). How general is human capital? a task-based approach. *Journal of Labor Economics 28*(1), 1–49.

ILO (2022). Digitalization and employment. Technical report, International Labour Organization, Geneva.

Klaeui, J., D. Kopp, R. Lalive, and M. Siegenthaler (2026). Adapting to scarcity: The role of firms in occupational transitions. CEPR Discussion Paper 21027, CEPR Press, Paris and London.

Lavetti, K. (2023). Compensating wage differentials in labor markets: Empirical challenges and applications. *Journal of Economic Perspectives 37*(3), 189–212.

Lazear, E. P. (2009). Firm-specific human capital: A skill-weights approach. *Journal of political economy 117*(5), 914–940.

Lise, J. and F. Postel-Vinay (2020). Multidimensional skills, sorting, and human capital accumulation. *American Economic Review 110*(8), 2328–2376.

Mas, A. and A. Pallais (2017). Valuing alternative work arrangements. *American Economic Review 107*(12), 3722–3759.

OECD (2022). Skills for the digital transition. Technical report, OECD Publishing, Paris.

Poletaev, M. and C. Robinson (2008). Human capital specificity: evidence from the dictionary of occupational titles and displaced worker surveys, 1984–2000. *Journal of Labor Economics 26*(3), 387–420.

Sorkin, I. (2018). Ranking firms using revealed preference. *The quarterly journal of economics 133*(3), 1331–1393.

Tong, D., L. Wu, and J. A. Evans (2021). Low-skilled occupations face the highest upskilling pressure. *arXiv preprint arXiv:2101.11505*.

Webb, M. (2019). The impact of artificial intelligence on the labor market. *Available at SSRN 3482150*.

Yamaguchi, S. (2012). Tasks and heterogeneous human capital. *Journal of Labor Economics 30*(1), 1–53.

7 Figures

Figure 1: Visualization of the conceptual framework

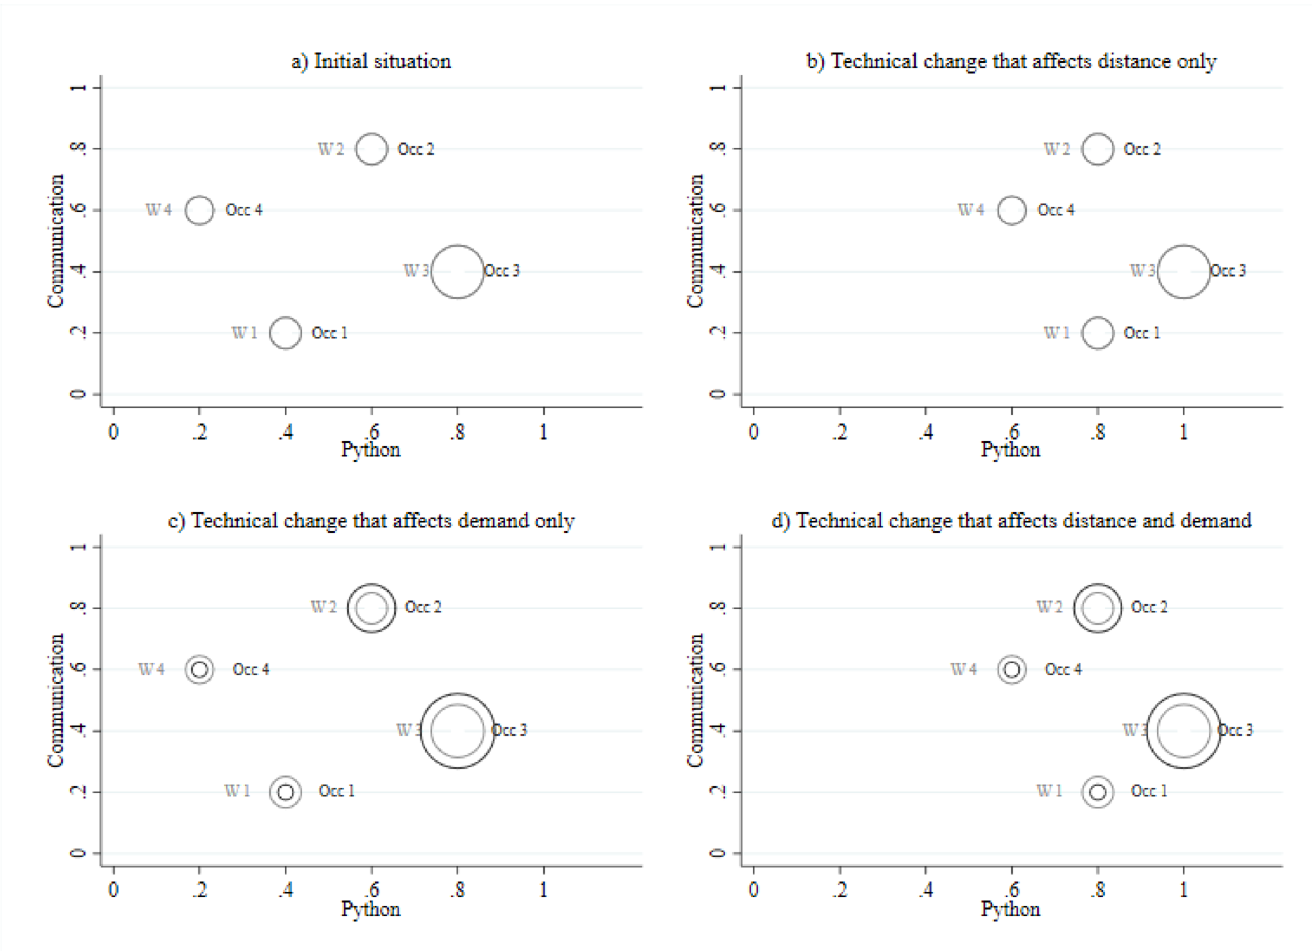

Notes: The figure illustrates a labor market with four worker types (W) and four occupation types (O). The size of each circle represents the mass of workers or jobs, while its position reflects the intensity of the two skills shown: "communication" and "Python". The figure illustrates the two distinct effects of technological change in the conceptual framework developed in the paper.

Figure 2: Convergence Toward the Digital Skill Set, 2011–2019

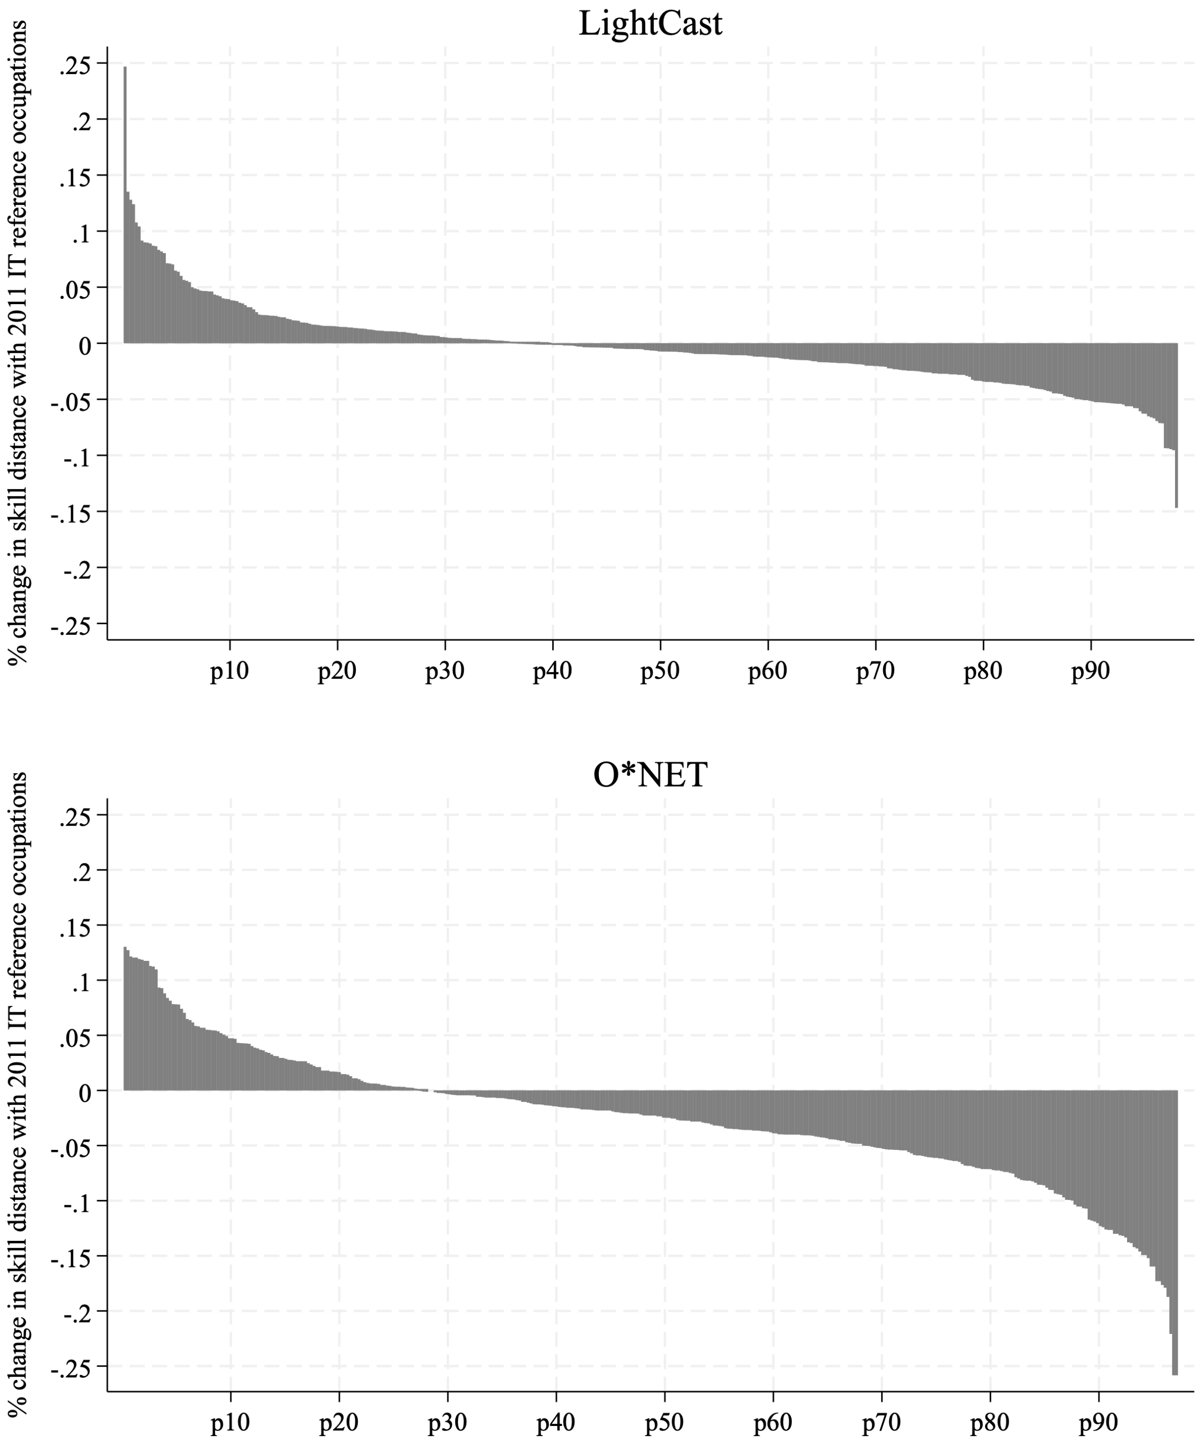

Notes: The figure reports the change in each occupation's distance to the digital skill reference between 2011 and 2019. The digital skill reference is defined by the set of IT occupations. A negative change indicates that an occupation moved closer to the digital reference by becoming more similar to these occupations in terms of its skill requirements. Distances are computed using the occupational digital skill vectors described in Section 3. The top panel uses digital skills identified from LightCast job postings, while the bottom panel relies on O*NET work activities.

Figure 3: Digital Skill Adoption and Skill Abandonment as Sources of Convergence

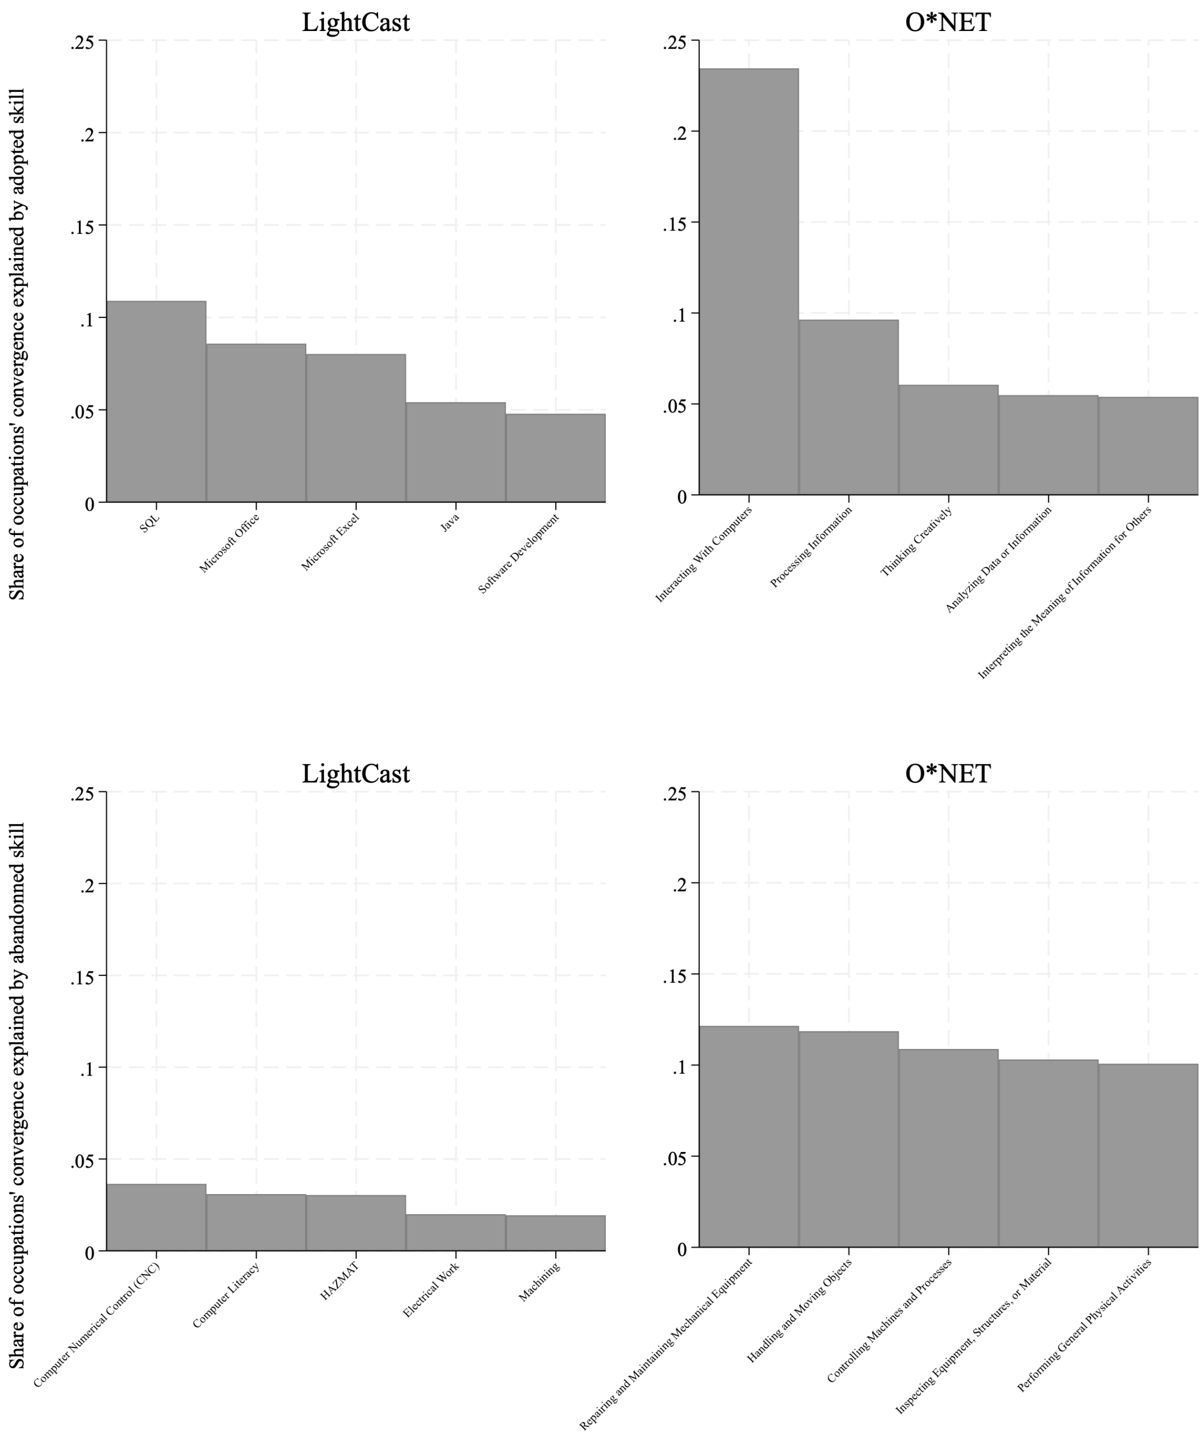

Notes: The figure reports the five individual skills accounting for the largest share of the reduction in distance to the digital skill reference through the adoption margin (left panel) and through the abandonment margin (right panel). The contribution of each skill is measured as its share of the total convergence explained by skill adoption or skill abandonment, respectively. Results are shown separately for the LightCast and O*NET measures.

Figure 4: Changes in Occupational Accessibility, 2011–2019

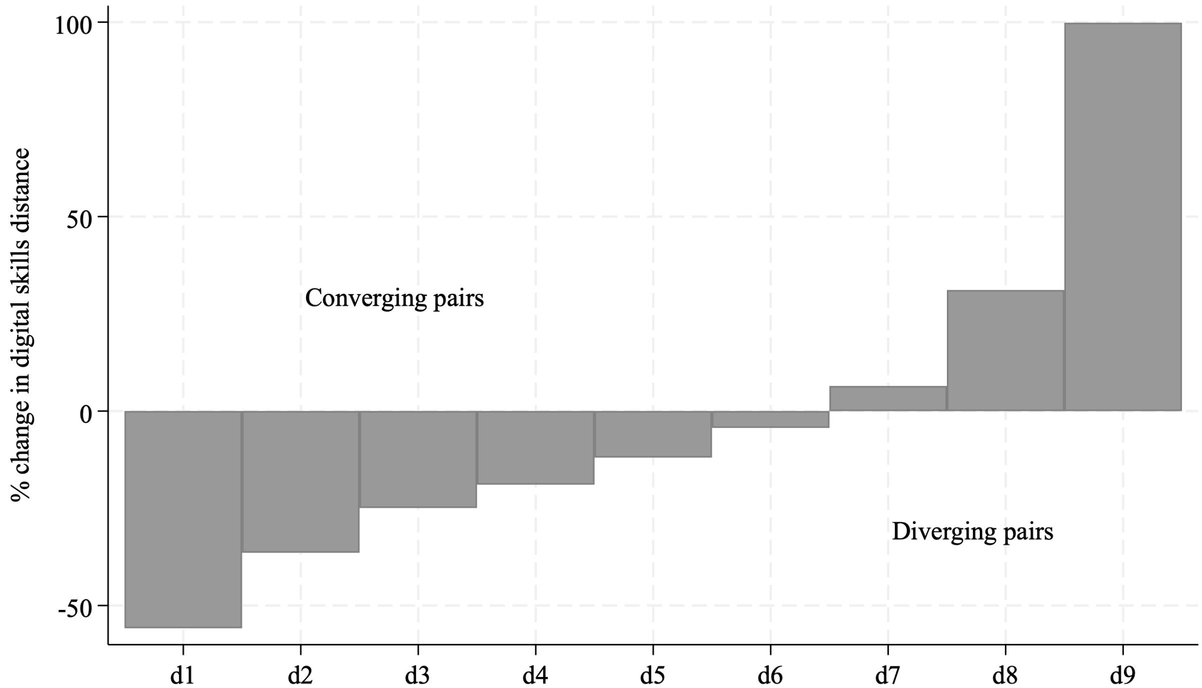

Notes: The figure reports the distribution of changes in pairwise digital skill distance between occupations from 2011 to 2019. Negative values indicate that two occupations became more similar in their digital skill requirements and therefore more accessible to one another, while positive values indicate that occupations became less similar and less accessible. Pairwise distances are computed using the digital skill measures described in Section 3. Results are shown separately for the LightCast and O*NET measures. These pairwise changes in occupational accessibility constitute the key explanatory variable used in the gravity analysis of Section 4

Figure 5: Contribution of Occupational Accessibility to Worker Reallocation

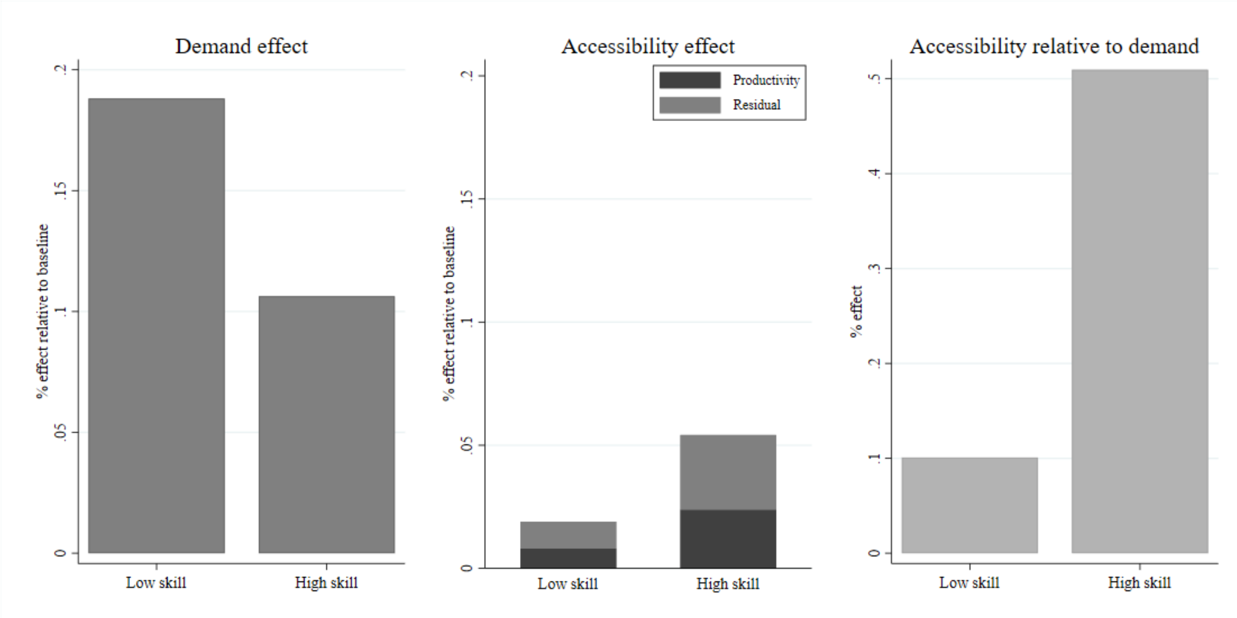

Notes: The left panel reports the additional worker mobility generated by changes in occupational demand relative to the counterfactual without technological change, separately for low- and high-skill occupations. The middle panel reports the additional mobility generated by changes in occupational accessibility. The colored bars decompose the accessibility effect into a productivity component and a residual component, capturing amenities and imperfect information. Finally, the right panel expresses the accessibility effect as a share of the demand effect.

Figure 6: Occupational Accessibility Facilitates Labor-Market Adjustment

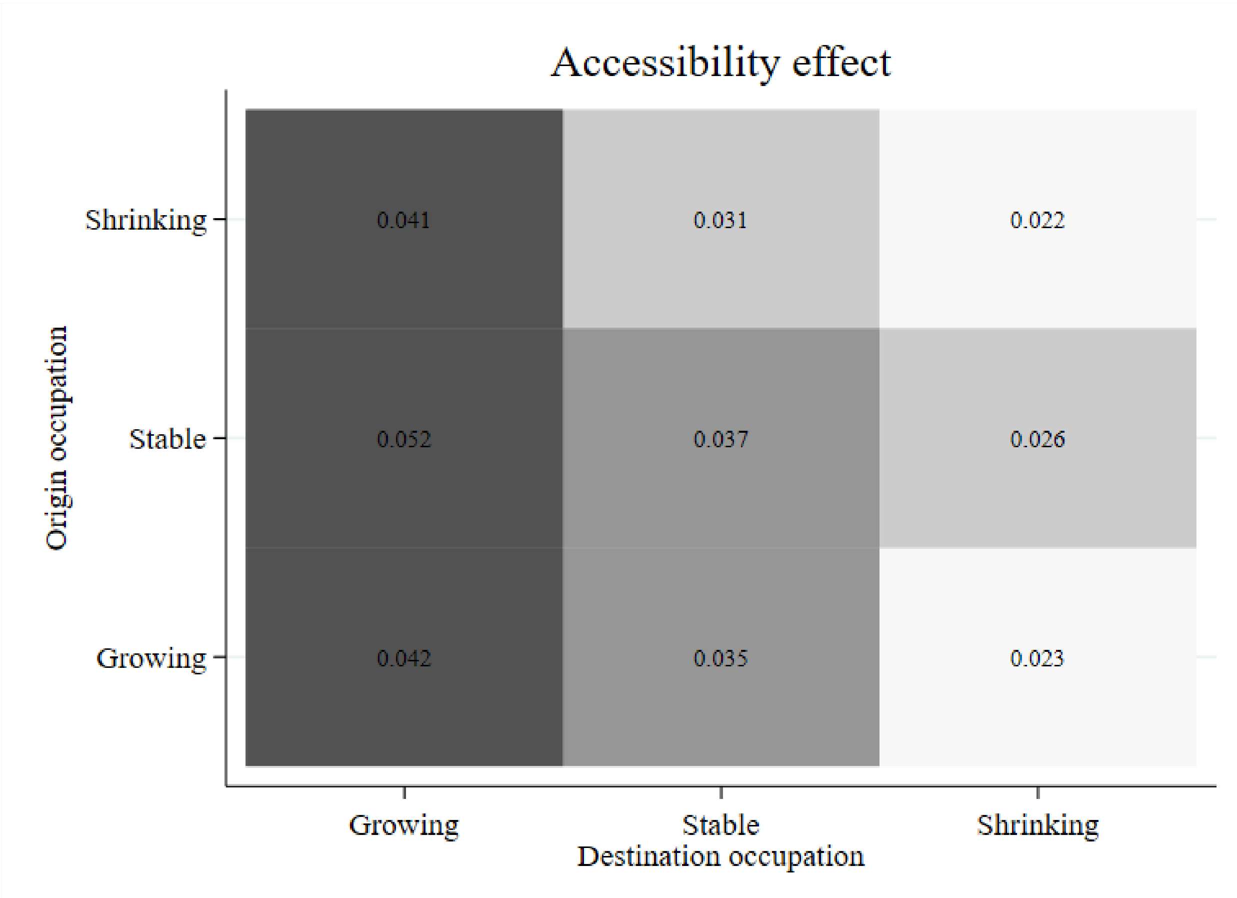

Notes: Occupations are classified into three groups according to changes in their relative employment share over the period: growing (employment share increases by more than 10%), shrinking (employment share decreases by more than 10%), and stable (employment share changes by less than 10%). Each group contains approximately one-third of all occupations. For each origin–destination combination, the figure reports the accessibility effect, measured as the percentage change in worker flows relative to the counterfactual without technological change. The four color categories correspond to the quartiles of the distribution of estimated effects.

8 Tables

Table 1: List of occupations classified as fully digital

PCS code	Occupation
388a	IT engineers in R&D
388b	IT engineers in charge of maintenance, support and user services
388c	IT project managers, IT managers
388d	Engineers and technical sales executives in IT and telecommunications
478a	IT design and development technicians
478b	IT production and operations technicians
478c	IT installation, maintenance, support and user services technicians
478d	Telecommunications and network computing technicians

Notes: The table includes the list of occupations classified as fully digital, which serve as comparison group to define digital distance between any two pairs of occupations. In practice, they represent all the occupation codes reported for the job of IT engineers and IT technicians.

Table 2: Summary statistics of pairwise occupation mobility

	Mobility flows		Mobility flows restricted to occupation switchers	
	Baseline (2011-12)	Endline (2018-19)	Baseline (2011-12)	Endline (2018-19)
Mean	153.9	172.5	26.0	23.3
SD	3701.9	4072.3	217.1	189.4
Min	0	0	0	0
p25	0	0	0	0
p50	2	1	2	1
p75	10	7	9	7
Max	515860	538008	37778	21515
N. of obs.	148’995		147’840	

Notes: The table summarizes the distribution of pairwise occupation mobility at the beginning and at the end of the period. The left two columns present the entire sample used for the analysis, including stayers in the same occupation and movers to and from non-employment. The right two columns restrict the sample to pairs of different occupations, and thus report the distribution of mobility across occupations.

Table 3: Occupational Accessibility and Worker Mobility

Dependent variable:	Endline mobility (2018-2019)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Baseline	Overall (2019)	Overall (2011 & Δ)	Excl. IT	IT Projection	Combined control	O*NET	Skill Clusters
Change in digital distance	0.6867***	0.9200*	0.9130*	0.9020**	0.9044*	0.8839**	0.8454***	0.9252**
	(0.0399)	(0.0464)	(0.0488)	(0.0417)	(0.0483)	(0.0473)	(0.0352)	(0.0311)
Digital distance	0.5983***	0.8844*	0.8782**	0.7656***	0.9399	0.9061*	0.6590***	0.8962**
	(0.0358)	(0.0579)	(0.0561)	(0.0427)	(0.0525)	(0.0460)	(0.0422)	(0.0422)
Change in overall/non-digital distance			0.7931***	0.7638***	0.7830***	0.8157***	0.7766***	0.8559***
			(0.0301)	(0.0295)	(0.0272)	(0.0268)	(0.0583)	(0.0235)
Overall/non-digital distance		0.4848***	0.4790***	0.4687***	0.5149***	0.5285***	0.5508***	0.5808***
		(0.0148)	(0.0151)	(0.0144)	(0.0123)	(0.0132)	(0.0653)	(0.0136)
Observations	148,995	148,995	148,995	148,995	148,995	148,995	145,935	148,995
Stayer dummy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Level Switch dummy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Origin FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Destination FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes:The dependent variable is the number of workers moving from origin occupation i to destination occupation j between 2018 and 2019. All specifications are estimated by Poisson pseudo-maximum likelihood (PPML) and include origin and destination occupation fixed effects. Standard errors are double-clustered at the origin and destination occupation levels. Bilateral controls include an indicator for remaining in the same occupation and an indicator for transitions involving a change in socioeconomic status. Column (1): Baseline specification including only the initial level and the change in digital distance. Column (2): Adds the 2019 level of overall occupational distance. Column (3): Replaces the previous control with the initial level and the change in overall occupational distance. Column (4): Controls for non-digital distance constructed after excluding skills disproportionately used by the occupations defining the digital frontier. Column (5): Controls for non-digital distance after projecting out the direction of the IT centroid from the occupational skill space. Column (6): Preferred specification combining both adjustments: the non-digital distance is constructed after excluding IT-intensive skills and projecting out the IT centroid from the remaining non-digital skill space. Column (7): Replicates Column (6) using O*NET to construct both digital and non-digital distance measures. Column (8): Replicates Column (6) using the 700-category LightCast skill clusters taxonomy instead of the full skill classification.

Table 4: Alternative Specifications

Dependent variable:			Endline mobility (2018-2019)			Δ mobility
	(1)	(2)	(3)	(4)	(5)	(6)
	Preferred	No Stayers	No Non-Employment	Skill Obsolescence	Baseline Mobility	Mobility Changes
Change in digital distance	0.8839**	0.8504**	0.8794**	0.8102***	0.9014*	0.7962**
	(0.0473)	(0.0575)	(0.0454)	(0.0320)	(0.0494)	(0.0730)
Digital distance	0.9061*	0.9219	0.8937**	0.8505***	0.9408	
	(0.0460)	(0.0630)	(0.0474)	(0.0402)	(0.0428)	
Change in non-digital distance	0.8157***	0.8055***	0.8095***	0.9117*	0.7864***	1.3461***
	(0.0268)	(0.0326)	(0.0260)	(0.0507)	(0.0276)	(0.0408)
Non-digital Distance	0.5285***	0.5446***	0.5254***	0.5243***	0.5298***	
	(0.0132)	(0.0169)	(0.0130)	(0.0214)	(0.0135)	
Observations	148,995	148,610	148,225	148,995	148,995	222,148
Stayer dummy	Yes	Yes	Yes	Yes	Yes	Yes
Level Switch dummy	Yes	Yes	Yes	Yes	Yes	Yes
Origin FE	Yes	Yes	Yes	Yes	Yes	No
Destination FE	Yes	Yes	Yes	Yes	Yes	No
Origin × Destination FE	No	No	No	No	No	Yes
Baseline mobility	No	No	No	No	Yes	No

Notes: The dependent variable is the number of workers moving from origin occupation i to destination occupation j. All specifications are estimated by Poisson pseudo-maximum likelihood (PPML) and include the preferred measure of digital accessibility and the preferred control measure for non-digital occupational distance introduced in Table 3. Origin and destination occupation fixed effects, together with the bilateral controls for remaining in the same occupation and for transitions involving changes in socioeconomic status, are included in all specifications unless otherwise noted. Standard errors are double-clustered at the origin and destination occupation levels.Column (1): Preferred specification from Column (6) of Table 3 Column (2): Excludes stayers (occupation pairs with i=j). Column (3): Excludes transitions to and from non-employment. Column (4): Replaces the baseline measure of occupational accessibility with the skill-obsolescence measure, holding the origin occupation fixed at its 2011 skill profile. Column (5): Controls for baseline bilateral worker mobility observed between 2011 and 2012. Column (6): Estimates the model using panel data over two periods (2011-12 and 2018-19) and including occupation pair fixed effects.

Table 5: Skill Adoption, Skill Abandonment, and Worker Mobility

Dependent variable:	Endline mobility (2018-2019)			
	(1)	(2)	(3)	(4)
	Horse Race	Adoption Only	Abandonment Only	Adoption Placebo
Change in Adoption	0.9391*	0.9385*		
	(0.0335)	(0.0334)		
Change in Abandonment	0.9777		0.9764	
	(0.0240)		(0.0239)	
Digital distance	0.9722	0.9726	1.0052	1.0361
	(0.0437)	(0.0437)	(0.0503)	(0.0640)
Change in non-digital distance	0.8102***	0.8093***	0.8100***	0.8090***
	(0.0256)	(0.0254)	(0.0257)	(0.0255)
Non-digital distance	0.5242***	0.5242***	0.5238***	0.5237***
	(0.0129)	(0.0129)	(0.0130)	(0.0130)
Change in placebo Adoption				1.0653
				(0.0412)
Observations	148,995	148,995	148,995	148,995
Stayer dummy	Yes	Yes	Yes	Yes
Level Switch dummy	Yes	Yes	Yes	Yes
Origin FE	Yes	Yes	Yes	Yes
Destination FE	Yes	Yes	Yes	Yes

Notes: The dependent variable is the number of workers moving from origin occupation i to destination occupation j between 2018 and 2019. All specifications build on the preferred gravity specification reported in Column (6) of Table 1. Skill adoption and skill abandonment replace the change in digital distance while all other controls remain unchanged. All specifications are estimated by Poisson pseudo-maximum likelihood (PPML) and include origin and destination occupation fixed effects. Standard errors are double-clustered at the origin and destination occupation levels.Column (1): Includes both skill adoption and skill abandonment. Column (2): Includes only skill adoption. Column (3): Includes only skill abandonment. Column (4): Replaces skill adoption with the placebo measure.

A Additional Empirical Evidence

This appendix presents additional empirical evidence supporting the main results of the paper. We first provide descriptive evidence on how occupational accessibility evolved during the period of digitalization. We then present a series of robustness exercises and supplementary analyses examining the stability of the main estimates across alternative specifications, worker groups, and patterns of occupational mobility.

A.1 Digital Skill Diffusion

This subsection documents how occupational accessibility evolved between 2011 and 2019. We first examine changes in the distribution of digital accessibility across occupation pairs, before comparing these changes with those observed in other dimensions of occupational similarity. Together, these figures provide descriptive evidence that the digitalization of the 2010s primarily reduced digital barriers between occupations rather than generating a broader convergence in occupational skill requirements.

Figure A.1 reports the distribution of digital distances in 2011 and 2019 for all occupation pairs (left panel) and, separately, for pairs in which the destination occupation j belongs to the set of IT reference occupations (right panel). The dashed vertical lines indicate the mean distance in each year. In both samples, the entire distribution shifts to the left over the decade, with both the average and the maximum distance in 2019 lying below their 2011 counterparts. The shift is considerably more pronounced for pairs involving an IT reference occupation, further confirming the evidence presented in the main text that digitalization brought occupations closer to the IT skill set.

Figure A.2 further explores where in the initial distribution of digital distance occupation pairs experienced the greatest convergence. Interestingly, convergence increases monotonically with the initial level of digital distance: occupation pairs that were initially the furthest apart became closer together the most over the decade. This suggests that the digitalization of the 2010s did not create a distinct group of already-digital occupations that pulled further away from the rest. Instead, it fostered convergence in digital skill requirements across initially distant occupations, compressing the overall distribution of digital distances.

Figure A.3 reports the relationship between digital and non-digital distance, both in levels (2011) and in changes over time. Non-digital distance is measured using our preferred specification, excluding both IT skills and the IT projection, and constitutes the main control variable in our benchmark specification. It is therefore important to assess the extent to which the two measures are correlated. As expected, occupation pairs that are closer in the digital skill space in 2011 also tend to be closer in the non-digital skill space, although the correlation is relatively modest (0.20). More interestingly, the correlation falls to just 0.06 when considering changes over time, indicating that changes in digital distance are largely orthogonal to changes in other skill dimensions. This suggests that the evolution of digital distance captures a distinct source of variation beyond general changes in occupational skill similarity.

Finally, Figure A.4 compares the average change in digital distance for all occupation pairs with that for pairs involving an IT reference occupation (similar in spirit to Figure A.1), and contrasts these patterns with those observed for non-digital distance. Consistent with the evidence presented earlier, the decline in digital distance is substantially larger for pairs involving an IT reference occupation, indicating that occupations became digitally closer primarily by acquiring the skills characteristic of the IT reference occupations. Importantly, this result is not mechanical: changes in digital distance are computed using the fixed set of IT reference skills identified in 2011, so the stronger convergence toward IT occupations is not driven by changes in the benchmark itself. By contrast, changes in non-digital distance are much smaller and exhibit virtually no difference between all occupation pairs and those involving an IT reference occupation. This further suggests that the transformation observed over the decade is specific to the digital skill dimension, rather than reflecting a broader convergence in occupational skill requirements.

Figure A.1: Distribution of digital distance measure

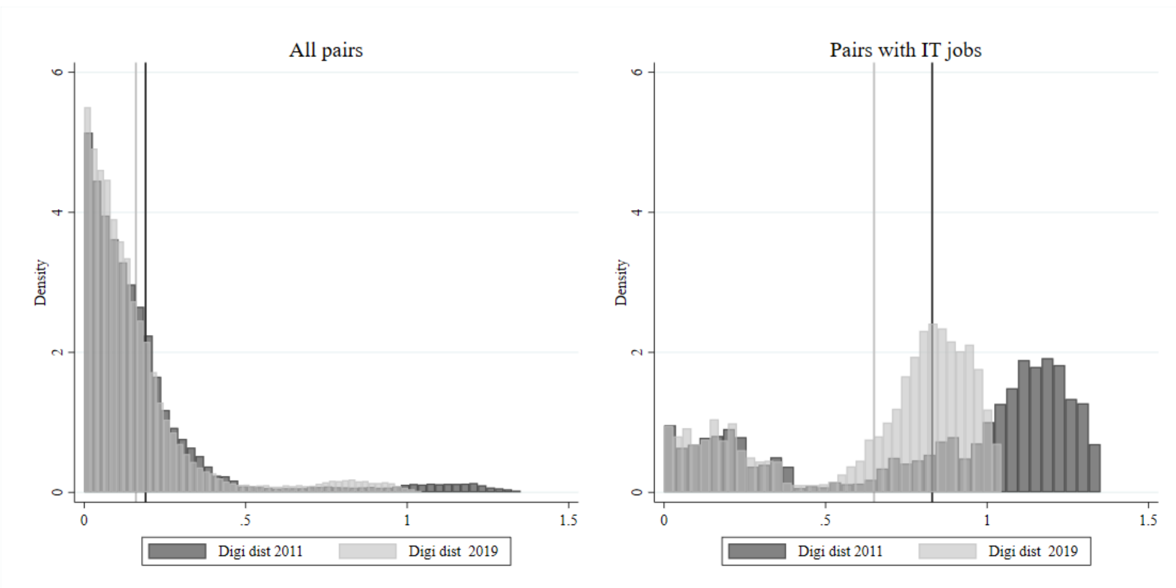

Notes: The left panel shows the distribution of dyadic digital distances across all occupation pairs in 2011 and 2019. The vertical lines indicate the corresponding mean distances. The right panel presents the same distributions, restricting the sample to occupation pairs in which the destination occupation (j) belongs to the set of IT reference occupations.

Figure A.2: Changes in digital distance by initial level of distance

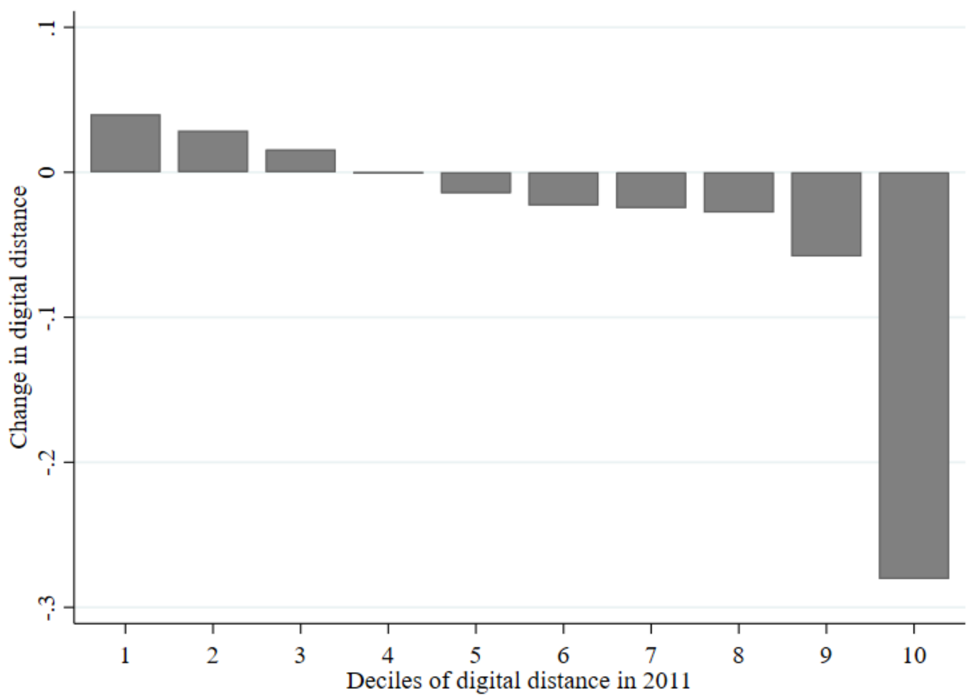

Notes: The figure plots the average change in digital distance across deciles of digital distance measured in 2011.

Figure A.3: Correlations between digital and non-digital distance

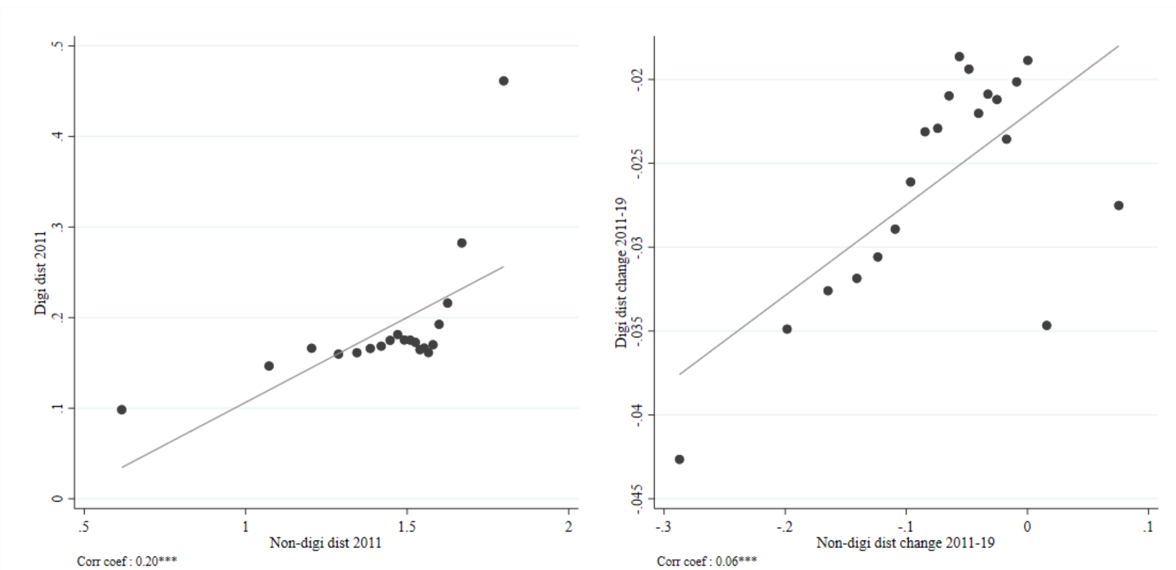

Notes: The left panel plots the binned relationship between the initial levels of digital and non-digital distance, where non-digital distance is measured using our preferred specification excluding both IT skills and the IT projection. The right panel plots the corresponding relationship for changes in the two measures over the sample period. The correlation coefficient is reported at the bottom of each panel.

Figure A.4: Comparing Changes in Digital and Non-Digital Distance Across Occupation Pairs

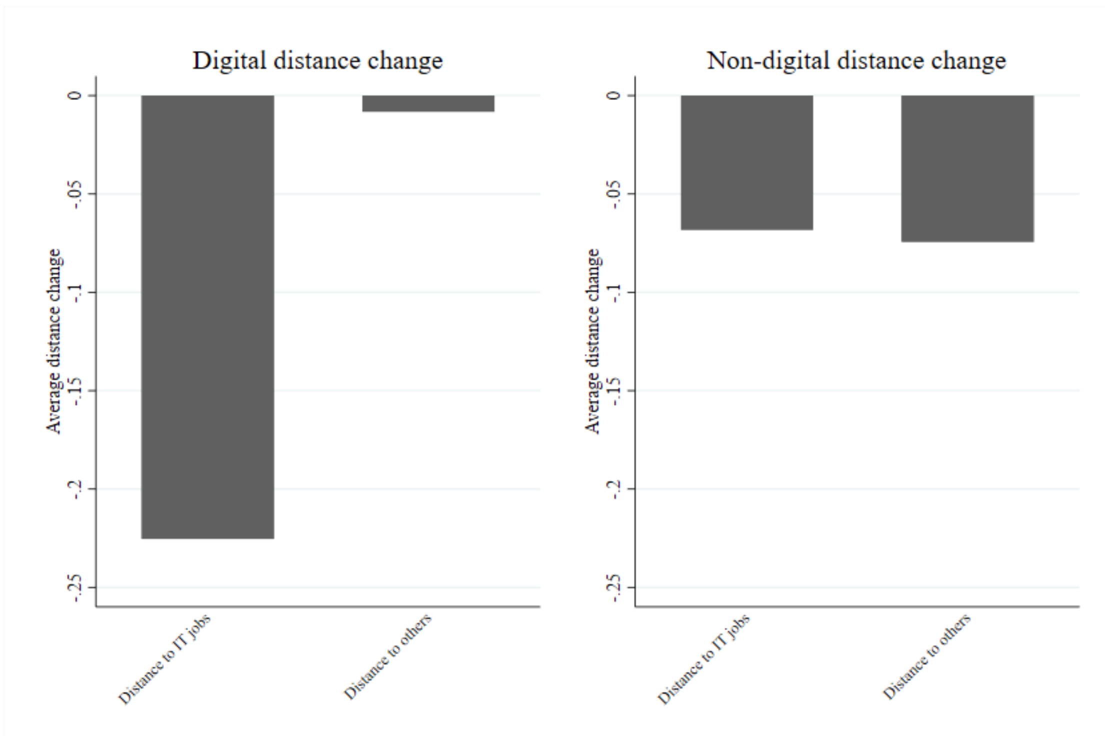

Notes: The left panel compares the average change in digital distance across all occupation pairs and pairs involving an IT reference occupation. The right panel presents the corresponding comparison for non-digital distance.

A.2 Robustness of the Gravity Estimates

This subsection examines the robustness of the estimated relationship between changes in digital accessibility and worker mobility. We first assess the sensitivity of the estimated coefficient to alternative specifications and measures of occupational similarity. We then examine whether the results are disproportionately driven by occupation pairs involving the IT reference occupations used to construct the digital accessibility measure.

Figures A.5 and A.6 summarize the estimated coefficient on changes in digital distance across the specifications reported in the main text. Figure A.5 displays the estimates and 90% confidence intervals for the benchmark specifications presented in Table 3, while Figure A.6 reports the corresponding results for the robustness exercises in Table 4. In both figures, the estimate from our preferred specification is highlighted in dark bold.

The figures show that the baseline specification, which does not control for other dimensions of occupational distance, yields a somewhat larger estimate of the effect of changes in digital distance. Once controls for overall or non-digital distance are introduced, however, the estimated coefficient remains remarkably stable across specifications. This includes specifications using alternative measures of occupational similarity, such as those constructed from O*NET or from the more aggregated Lightcast skill clusters, as well as specifications that explicitly account for skill obsolescence. Importantly, the estimated effect is even larger in the occupation-pair fixed-effects specification, where identification relies on within-pair changes in digital distance and bilateral mobility over time while controlling for changes in non-digital distance.

An additional concern is whether our results are driven primarily by occupation pairs involving one of the IT reference occupations. Since Figures A.1 and A.4 show substantially larger convergence for these pairs, despite their representing only a small fraction of all occupation pairs, it is natural to ask whether they account for the estimated relationship. To address this concern, Figure A.6 reports the estimated coefficient on changes in digital distance for the main specifications after excluding all occupation pairs involving an IT reference occupation.

Excluding these pairs has little effect on the estimated magnitude of the coefficient. In our preferred specification, the estimate becomes marginally insignificant (p-value = 0.13), while remaining very similar in size to the benchmark estimate. Statistical significance is restored once occupation pairs involving non-employment are excluded. These pairs add substantial noise to the estimation because we cannot compute a meaningful measure of occupational distance for non-employment and therefore assign them a distance of zero by construction. Reassuringly, even without excluding non-employment pairs, the coefficient remains statistically significant and of comparable magnitude across the other key specifications, including those based on O*NET, Lightcast skill clusters, skill obsolescence, and occupation-pair fixed effects.

Finally, figure A.7 summarizes the robustness of the estimated effect of changes in digital occupational accessibility to alternative implementations of both the treatment and control variables. The first estimate corresponds to the preferred specification reported in Column 6 of Table 3. The following ten estimates modify the construction of the digital accessibility measure by considering alternative definitions of the IT reference occupations. Specifically, we construct the digital reference list using only IT technician occupations, only IT engineering occupations, and eight leave-one-out definitions that sequentially exclude one of the eight reference IT occupations. The final estimate instead modifies the construction of the non-digital accessibility control by replacing our preferred two-step procedure with an alternative implementation. The estimated coefficient on changes in digital occupational accessibility remains stable across all specifications. The figure reports point estimates together with 90% confidence intervals.

Figure A.5: Coefficient of interest in main results table

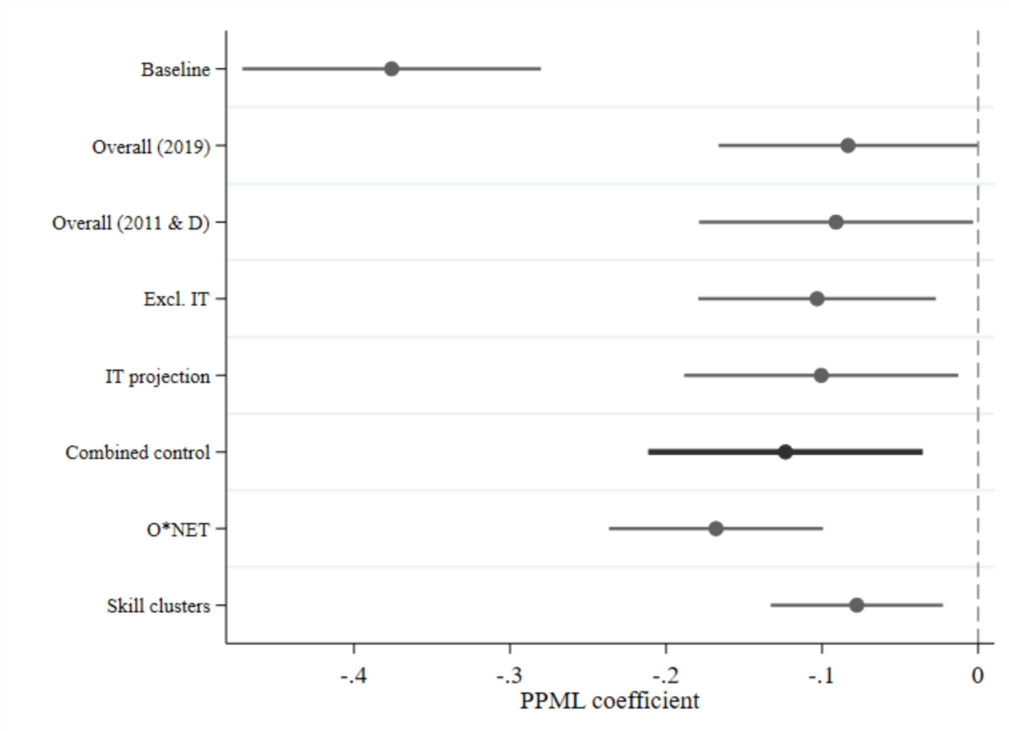

Notes: The figure reports the coefficient of interest associated with the change in digital distance and its 90% confidence interval for all specifications reported in the main results table (Table 3). The preferred model is reported in dark bold.

Figure A.6: Coefficient of interest in alternative specifications table

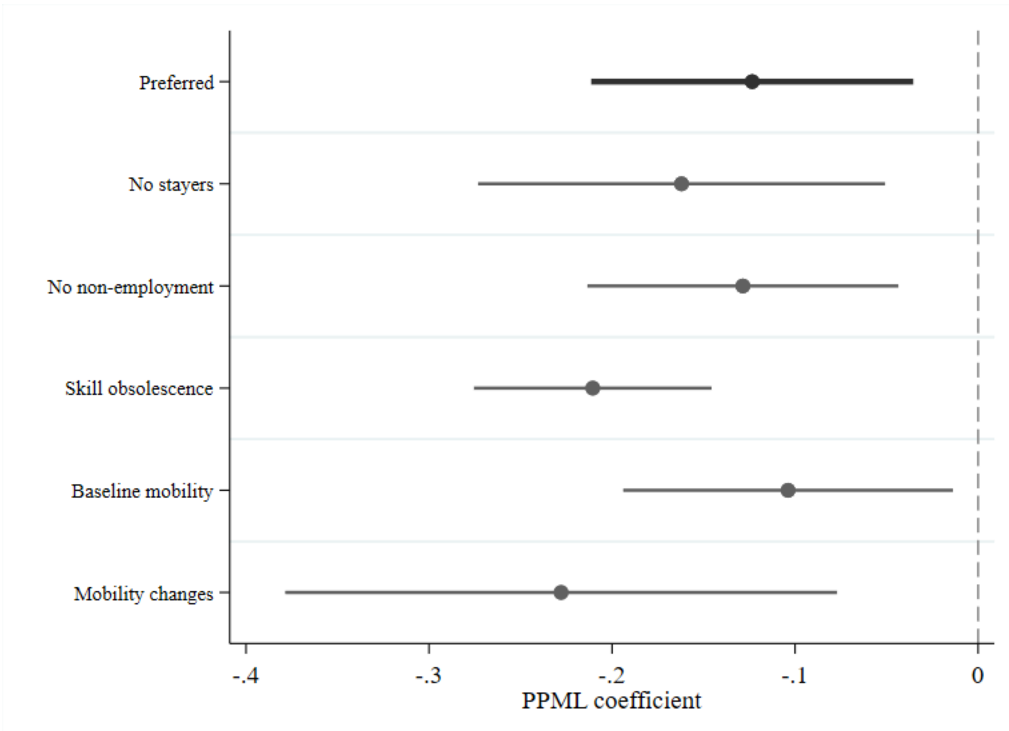

Notes: The figure reports the coefficient of interest associated with the change in digital distance and its 90% confidence interval for all specifications reported in the alternative specifications table (Table 4). The preferred model is reported in dark bold.

Figure A.7: Coefficient of interest and robustness to excluding the IT references

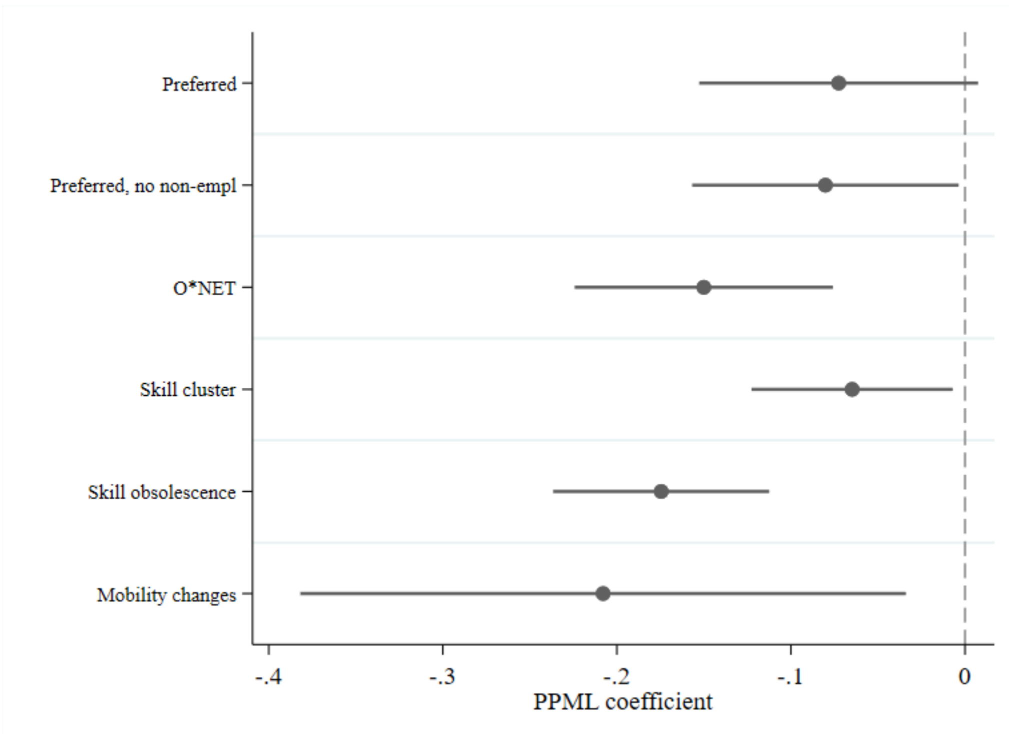

Notes: The figure reports the coefficient of interest associated with the change in digital distance and its 90% confidence interval for the most important specifications, obtained when excluding from the sample all occupation pairs involving at least one IT reference occupation.

Figure A.8: Coefficient of interest with alternative definitions of the IT reference list

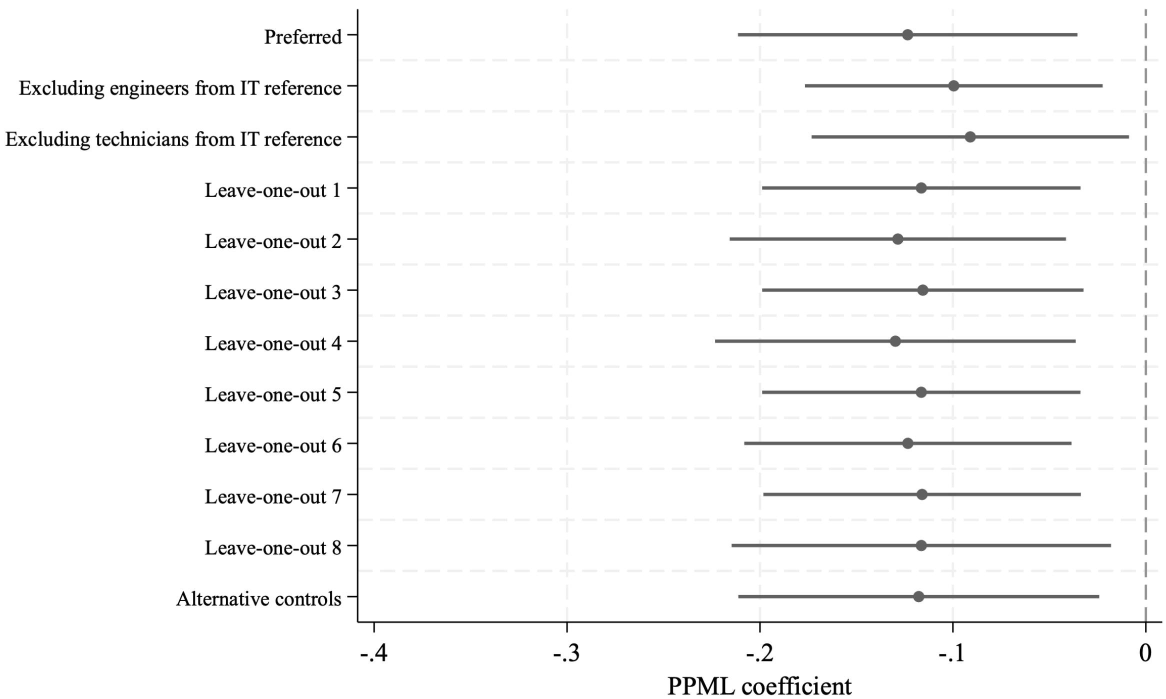

Notes: The figure reports the coefficient of interest associated with the change in digital distance and its 90% confidence interval.

A.3 Heterogeneity

One may also wonder whether our effect is heterogeneous across demographic groups, or whether it is asymmetric across relative characteristics of origin and destination occupations.

Table A.1 estimates our preferred specification separately across demographic groups. Columns (1) and (2) report estimates for workers aged 20–35 and 36–60, respectively, while columns (3) and (4) report estimates for female and male workers.

Results are remarkably similar across age groups. Although younger workers change occupations more frequently than older workers (the average bilateral occupation flow is 12.1, compared with 11.0 for workers aged 36–60), both the effect of the initial level of digital distance and that of changes in digital distance are nearly identical across the two samples.

The results are similarly consistent across gender. While the coefficient on changes in digital distance is slightly larger and statistically significant for men but only marginally insignificant for women, the pattern is reversed for the initial level of digital distance, which is estimated more precisely and is slightly larger for women. Men also exhibit higher occupational mobility on average, with an average bilateral occupation flow of 13.3 compared with 10.0 for women. Overall, these differences are modest and provide little evidence that the relationship between digital accessibility and occupational mobility differs systematically across demographic groups.

A.4 Symmetry of Occupational Accessibility

Finally, table A.2 examines whether changes in bilateral digital accessibility—a symmetric measure by construction—have asymmetric effects depending on the relative characteristics of the origin and destination occupations. Each column interacts the change in digital distance with an indicator capturing a relative characteristic of the two occupations. In column (1), the indicator equals one if the origin occupation i was initially more digitally intensive than the destination occupation j, as measured by the average initial distance to the IT reference occupations. In column (2), it equals one if occupation i was initially more math-intensive than occupation j, as measured by the O*NET "Mathematical Knowledge" score. In column (3), it equals one if occupation i was initially more

Table A.1: Heterogeneity of the effect of digital accessibility across demographics

Dependent variable	Endline mobility (2018-2019)			
	(1)	(2)	(3)	(4)
	Younger workers (<35)	Older workers (>35)	Female workers	Male workers
Change in digital distance	0.895*	0.872***	0.908	0.866***
	(0.0536)	(0.0459)	(0.0583)	(0.0443)
Digital distance	0.937	0.881**	0.857**	0.924*
	(0.0450)	(0.0505)	(0.0610)	(0.0399)
Change in non-digital distance	0.797***	0.838***	0.796***	0.837***
	(0.0267)	(0.0291)	(0.0303)	(0.0288)
Non-digital distance	0.0541***	0.520***	0.548***	0.536***
	(0.0157)	(0.0127)	(0.0148)	(0.0144)
Observations	148,995	148,995	148,995	148,995
Stayer dummy	yes	yes	yes	yes
Level switch dummy	yes	yes	yes	yes
Origin FE	yes	yes	yes	yes
Destination FE	yes	yes	yes	yes

Notes: The dependent variable is the number of workers moving from origin occupation i to destination occupation j between 2018 and 2019. All specifications are estimated by Poisson pseudo-maximum likelihood (PPML) and include origin and destination occupation fixed effects. Standard errors are double-clustered at the origin and destination occupation levels. Bilateral controls include an indicator for remaining in the same occupation and an indicator for transitions involving a change in socioeconomic status. Columns (1) and (2) estimate the preferred specification separately for workers aged 20–35 and 36–60, respectively. Columns (3) and (4) estimate the preferred specification separately for female and male workers, respectively.

skilled than occupation j, as measured by the first digit of the French PCS classification. In column (4), it equals one if occupation i was initially tighter than occupation j, as measured by labor market tightness statistics from the French Public Employment Service.

Across all specifications, we find no evidence of asymmetric effects. A reduction in digital distance between two occupations increases mobility to a similar extent in both directions, regardless of which occupation is initially more digitally intensive, more math-intensive, more skilled, or characterized by a tighter labor market. This does not imply that digitalization generates no aggregate reallocation across occupations. Rather, it indicates that such reallocation arises because digital accessibility improves unevenly across occupation pairs, not because a given bilateral improvement has intrinsically asymmetric effects.

Table A.2: Asymmetry of the effect of digital accessibility

Dependent variable	Endline mobility (2018-2019)			
	(1)	(2)	(3)	(4)
	By digital intensity	By math intensity	By skill	By tightness
Change in digital distance	0.890	0.853***	0.885**	0.908*
	(0.0631)	(0.0507)	(0.0518)	(0.0525)
Change in digital distance * higher digi i	0.986			
	(0.0770)			
Higher digi i	1.164			
	(0.184)			
Change in digital distance * higher math i		1.042		
		(0.0646)		
Higher math i		1.244**		
		(0.120)		
Change in digital distance * higher skill i			0.992	
			(0.0546)	
Higher skill i			1.766	
			(0.984)	
Change in digital distance * higher tightness i				0.905
				(0.0660)
Higher tightness i				1.018
				(0.238)
Observations	148,995	148,995	148,995	148,995
Stayer dummy	yes	yes	yes	yes
Level switch dummy	yes	yes	yes	yes
Digital distance	yes	yes	yes	yes
Change in non-digital distance	yes	yes	yes	yes
Non-digital distance	yes	yes	yes	yes
Origin FE	yes	yes	yes	yes
Destination FE	yes	yes	yes	yes

Notes:The dependent variable is the number of workers moving from origin occupation i to destination occupation j between 2018 and 2019. All specifications are estimated by Poisson pseudo-maximum likelihood (PPML) and include origin and destination occupation fixed effects. Standard errors are double-clustered at the origin and destination occupation levels. Bilateral controls include the level and change in non-digital distance, an indicator for remaining in the same occupation and an indicator for transitions involving a change in socioeconomic status. Each column interacts the change in digital distance with an indicator capturing a relative characteristic of the origin and destination occupations. In column (1), the indicator equals one if the origin occupation *i* was initially more digitally intensive than the destination occupation *j*, as measured by the average initial distance to the IT reference occupations. In column (2), it equals one if occupation *i* was initially more math-intensive than occupation *j*, as measured by the O*NET "Mathematical Knowledge" score. In column (3), it equals one if occupation *i* was initially more skilled than occupation *j*, as measured by the first digit of the French PCS classification. In column (4), it equals one if occupation *i* was initially tighter than occupation *j*, as measured by labor market tightness statistics from the French Public Employment Service.

B Data and Measurement

This section provides additional details on the measurement procedures underlying the empirical analysis. We first describe the crosswalk between U.S. SOC occupations and French PCS occupations, before detailing the construction of the non-digital accessibility measure used as a control in the empirical specifications.

B.1 Occupation Crosswalk

Using U.S. job-posting data to construct measures of skill distance offers an important identification advantage in our French setting. Changes in skill requirements measured in the United States are unlikely to be driven by French institutional features or labor market conditions. Instead, they plausibly reflect technology-driven shifts common to advanced economies. We therefore interpret these measures as predetermined with respect to French occupational mobility. To combine these measures with French administrative data, however, U.S. occupations must first be mapped to their French counterparts.

Lightcast job postings are classified according to the 2010 U.S. Standard Occupational Classification (SOC), whereas the mobility data and all occupation-level measures used elsewhere in the paper are defined using the French Professions et Catégories Socio-professionnelles (PCS 2003) classification. Because no official correspondence between the two occupational classifications exists, we construct a SOC–PCS crosswalk using the International Standard Classification of Occupations (ISCO-08) as a common intermediate classification. Specifically, we combine the official correspondence table between SOC 2010 and ISCO-08 published by the U.S. Bureau of Labor Statistics with the official correspondence table between PCS 2003 and ISCO-08 published by INSEE.

A SOC occupation and a PCS occupation are linked whenever they are both associated with at least one common ISCO-08 occupation. Operationally, this amounts to joining the two official correspondence tables on the ISCO-08 code and retaining every resulting SOC–PCS pair. The resulting crosswalk contains 4,181 SOC–PCS pairs linking 858 distinct SOC occupations to 485 distinct PCS occupations.

Because both official correspondence tables allow one-to-many mappings, the resulting crosswalk is itself many-to-many: on average, each SOC occupation is linked to 4.9 PCS occupations, while each PCS occupation is associated with 8.6 SOC occupations. Rather than selecting a single correspondence for each occupation, we retain every SOC–PCS pair implied by the two official crosswalks. This avoids imposing arbitrary restrictions when occupations span multiple ISCO categories and preserves all occupation matches supported by the official classifications.

B.2 Construction of the Non-Digital Accessibility Measure

This subsection describes the construction of the non-digital occupational accessibility measure used as a control throughout the empirical analysis. Starting from occupation-level skill vectors, the measure removes the influence of digital technologies in two stages. First, we exclude skills that are either software-specific or disproportionately concentrated in IT occupations. Second, we remove the residual direction associated with IT occupations from the remaining skill vectors before computing bilateral occupational distances. All classification decisions are based on the 2011 skill distribution and are held fixed throughout the sample period. We conclude by describing the alternative implementations reported in Table 3 and the construction of the cross-year distance measures used in the robustness analysis.

B.2.1 Inputs and Timing Convention

Let $z_j^t = (z_{j,1}^t, \ldots, z_{j,\bar{K}}^t)$ denote the Lightcast skill vector of occupation j in year $t \in \{2011, 2019\}$, where $z_{j,k}^t$ is the share of job postings mentioning skill k. The data contain $\bar{K} = 15{,}377$ distinct skills and 476 detailed PCS occupations observed in both years. Let $\mathcal{J}_{IT}$ denote the eight IT occupations used to define the digital frontier in Section 3, and let $Soft_k$ denote the Lightcast indicator identifying software-specific skills.

All classification decisions described below are based exclusively on the 2011 skill distribution and are subsequently held fixed throughout the sample period. In particular, the identification of IT-associated skills and the construction of the IT reference direction are predetermined with respect to the subsequent diffusion of digital technologies.

B.2.2 Removing Digital Skills

The first stage removes skills that are directly associated with digital technologies. We exclude both software-specific skills identified by the Lightcast taxonomy and skills that are disproportionately concentrated in IT occupations. The objective is to retain only skills whose diffusion is unlikely to mechanically reflect the spread of digital technologies.

For each skill, we compare its average 2011 intensity among the eight IT occupations with its average intensity among all remaining occupations. A skill is classified as IT-associated whenever its average intensity among IT occupations exceeds 0.001 and is at least 1.5 times larger than among non-IT occupations. Together with the software classification, this defines the set of non-digital skills,

$$\mathcal{K}^{ND} = \left\{ k : Soft_k = 0 \text{ and } \frac{\bar{z}_k^{IT}}{\bar{z}_k^{N}} < 1.5 \text{ or } \bar{z}_k^{IT} < 0.001 \right\}, \tag{13}$$

where $\bar{z}_k^{IT}$ and $\bar{z}_k^{N}$ denote the corresponding average skill intensities computed from the 2011 data.

Restricting each occupation's skill vector to the coordinates in $\mathcal{K}^{ND}$ removes skills whose diffusion is primarily driven by digital technologies.

B.2.3 Removing the Residual IT Direction

Even after excluding IT-associated skills, occupations may still differ along broader dimensions that characterize IT occupations, such as analytical reasoning, project management, or complementary technical competencies. The second stage therefore removes the residual direction associated with IT occupations before occupational distances are computed.

Within the restricted skill space $\mathcal{K}^{ND}$, we compute the average 2011 skill vector of the IT occupations,

$$R = \frac{1}{|\mathcal{J}_{IT}|} \sum_{j \in \mathcal{J}_{IT}} z_j^{2011}, \qquad \hat{R} = \frac{R}{\|R\|}, \tag{14}$$

where all vectors are understood to be restricted to the coordinates in $\mathcal{K}^{ND}$.

Each occupation's skill vector is then orthogonally projected onto the subspace perpendicular to $\hat{R}$,

$$\tilde{z}_j^t = z_j^t - (\hat{R} \cdot z_j^t)\hat{R}, \tag{15}$$

again restricting attention to the coordinates in $\mathcal{K}^{ND}$.

The resulting non-digital occupational distance is defined as

$$d_{ij,t}^{ND} = \delta\left(\tilde{z}_i^t, \tilde{z}_j^t\right), \tag{16}$$

where $\delta(\cdot, \cdot)$ denotes the angular distance measure introduced in Section 3. This is the preferred measure of non-digital occupational accessibility used throughout the paper.

B.2.4 Alternative Implementations

Table 3 reports three nested implementations of this procedure. The specification labelled *ndig_xw* (Column 4) performs only the first stage, removing software-specific and IT-associated skills before computing occupational distances. The specification labelled *itperp* (Column 5) performs only the second stage, projecting out the IT reference direction while retaining all non-software skills. Our preferred specification, labelled *ndigxwp* (Column 6), combines both stages and corresponds to the measure $d_{ij,t}^{ND}$ used throughout the paper.

B.2.5 Alternative Assumptions on Skill Accumulation

The construction of cross-year changes in occupational distance depends on the assumed process of workers' skill accumulation. The baseline specification assumes that workers continuously update their skills as occupational requirements evolve, so that both occupations in a pair are evaluated using their current skill profiles. As discussed in Section 4, we also consider an alternative skill-obsolescence assumption, under which workers retain the skill bundle associated with their occupation when they entered it. These alternative assumptions give rise to different measures of changes in occupational distance.

Define

$$d_{ij}^{cross} = \delta\left(\tilde{z}_i^{2019}, \tilde{z}_j^{2011}\right), \tag{17}$$

that is, the occupational distance obtained by comparing occupation i's 2019 skill profile with occupation j's 2011 skill profile.

Under the baseline assumption, changes in occupational accessibility are measured by

$$\Delta d_{ij}^{ND} = d_{ij,2019}^{ND} - d_{ij,2011}^{ND}, \tag{18}$$

which allows the skill requirements of both occupations to evolve between 2011 and 2019.

Under the skill-obsolescence assumption, the worker's origin occupation is instead held fixed at its 2011 skill profile while the destination occupation is allowed to evolve. The corresponding change measure is

$$\Delta d_{ij}^{ND,obs} = d_{ji}^{cross} - d_{ij,2011}^{ND}. \tag{19}$$

This measure captures how the skill bundle associated with the worker's original occupation compares with the current requirements of potential destination occupations. It is used in the robustness specification reported under "Obsolescence" in Table 4, together with the corresponding digital accessibility measure constructed under the same fixed-origin assumption.

C Structural Model

The labor market we consider is a market where on both sides, workers and jobs are grouped into discrete types (occupations), and types are defined by a vector of (required) skills.

We use the two-sided one-to-one matching model with transferable utility a la Choo and Siow (2006) and Galichon and Salanié (2024) to model this market.[7] In this model, there is a large number of workers of each type and a large number of jobs of each type. Both workers and employers aim to match up with one agent on the other side to maximize their utility. Transfers, in the form of wages, are possible, but the market being competitive and workers and employers being price-takers, transfers are determined in equilibrium.

C.1 Workers' and jobs' types

We denote by O the list of occupations. We say that a job is of type $j \in O$ when this job's occupation is j. Let X_j be the mass of jobs of type j in the market so that there are $\sum_{j \in O} X_j$ jobs in the market.

Types of workers are also defined using occupations. A worker is of type $i \in O$ if the occupation of her previous job is i. Workers who were previously not employed are taken into account by extending the set of occupations to include a category "0", i.e. $O_0 = O \cup \{0\}$, so that a worker who was previously not employed is said to be of type $i = 0$. Let Y_i for all $i \in O_0$, be the mass of workers of type i on the market, so that, for example, Y_0 indicates the mass of workers who were not previously employed. There are $\sum_{i \in O_0} Y_i$ workers on the market.

Note that we do not restrict the mass of workers $\sum_{i \in O_0} Y_i$ to be equal to the mass of jobs $\sum_{j \in O} X_j$, as workers can remain not employed. However, since our empirical analysis does not consider job vacancies, the model is such that all jobs must be matched to a worker. Therefore, it must be $\sum_{i \in O_0} Y_i \geq \sum_{j \in O} X_j$.[8]

[7]See also Galichon and Salanié (2022) for a generalized version and Dupuy and Galichon (2022) for a continuous type version of the model applied to the labor market for risky jobs.

[8]In appendix C.8 we show an extension of the model where jobs may also remain vacant. Because of the logit structure of the model, i.e. the Independence of Irrelevant Alternatives applies, and excluding the possibility of vacant jobs does not affect the remaining log-odds and the main analysis remains unchanged.

C.2 Matching

Let X_{ij} denote the mass of workers of type i matched to a job of type j and let $X_{i\emptyset}$ be the mass of workers of type i that remain not employed. A feasible matching is then a tuple $\left\{ (X_{i\emptyset})_{i\in O_0, j\in O}, (X_{ij})_{i\in O_0, j\in O} \right\}$ satisfying the accounting constraints[9]

$$\begin{aligned} X_{i\emptyset} + \sum_{j\in O} X_{ij} &= Y_i, \ \forall i \in O_0, & (20) \\ \text{and} \sum_{i\in O_0} X_{ij} &= X_j, \forall j \in O. & (21) \end{aligned}$$

The first accounting constraint indicates that the total mass of workers of type i matched to any type of jobs (not employed included) is exactly the mass of workers of type i available on the market. The second accounting constraint indicates that the total mass of jobs of type j filled by any type of workers is exactly the mass of jobs of that type available on the market.[10]

C.3 Match values and choices

Let α_{ij} be the systematic intrinsic utility derived by a worker of type $i \in O_0$ when working in a job of type $j \in O_\emptyset = O \cup \{\emptyset\}$ and w_{ij} be the monetary transfer, typically the wage, paid in jobs of type $j \in O_\emptyset$ for workers of type $i \in O_0$, with the convention that when workers do not work, i.e. $j = \emptyset$, they receive no transfers, i.e. $w_{i\emptyset} = 0$.[11] Further, let ε_j be a worker-specific, idiosyncratic taste for occupation $j \in O_\emptyset$, drawn from a (centered) Gumbel type I distribution with unit scaling factor.

[9]Underlying this notation and interpretation of the matching model is the assumption that there are no vacant jobs in the economy. The market clears with all jobs being filled. This is because we assume that $\sum_i X_{ij}$, the sum of all workers working in j is also the mass of jobs in j, i.e. $\sum_i X_{ij} = X_j$. As shown in Appendix C.8, this could be accommodated within the same framework. Estimation of this model, however, requires one to observe the mass of vacant jobs by occupation, which typically is not observed in matched-employer-employee data. Note that one could circumvent this issue by collecting data on vacancies by occupation and appending that information to the matched data.

[10]Using our notation one then has that X_{0j} is the mass of workers that previously were not employed, who are working in occupation j. Moreover, $\sum_j X_{0j} = Y_0$ is the total mass of workers that previously were not employed whereas X_{i0} is the mass of workers that were previously employed in occupation i and are currently not employed. Finally, $\sum_i X_{i\emptyset} = X_\emptyset$ is the mass of workers that are currently not employed. Note that since individuals who were previously not employed and are still not employed are typically not observed in our data, we simply add the restriction $X_{0\emptyset} = 0$.

[11]Note that w_{0j} $\forall j \in O$ needs not be 0 as it corresponds to the transfer paid to a worker of type 0, i.e. that was previously not employed, who is currently working in a job of type $j \in O$.

A worker of type $i \in O_0$ maximizes her utility by choosing the appropriate occupation, i.e. solves the problem

$$\max_{j \in O_\emptyset} \left(\alpha_{ij} + w_{ij} + \varepsilon_j \right). \tag{22}$$

Let γ_{ij} be the systematic productivity of a worker of type $i \in O_0$ in a job of type $j \in O$ and η_i be the idiosyncratic productivity of the job / employer when matched with a worker of type $i \in O_0$. It is assumed that the job-specific productivity η_i is drawn from a (centered) Gumbel type I distribution with unit scaling factor.

An employer with a vacant job of type $j \in O$ maximizes her profits by choosing the type of the worker to match with that solves the following problem

$$\max_{i \in O_0} \left(\gamma_{ij} - w_{ij} + \eta_i \right). \tag{23}$$

By an application of the Williams-Daly-Zachary theorem, each of these problems yields a solution of the form

$$\log X_{ij}^S = \alpha_{ij} + w_{ij} - s_i \quad \forall \, (i,j) \in O_0 \times O_\emptyset / \, (0, \emptyset) \,, \tag{24}$$
$$\log X_{ij}^D = \gamma_{ij} - w_{ij} - m_j \quad \forall \, (i,j) \in O_0 \times O. \tag{25}$$

The first equation can be thought of as the supply of workers of type i to jobs of type j, while the second equation can be thought of as the demand of employers with jobs of type j for workers of type i.

C.4 Equilibrium

In equilibrium, supply equates demand, i.e. $X_{ij}^S = X_{ij}^D = X_{ij}$ for all $(i,j) \in O_0 \times O$, and it follows that, by rescaling and adding equations (24) and (25),

$$X_{ij} = \exp\left(\frac{\varphi_{ij} - s_i - m_j}{2} \right), \forall \, (i,j) \in O_0 \times O \tag{26}$$
$$X_{i\emptyset} = \exp\left(\alpha_{i\emptyset} - s_i \right), \forall i \in O \tag{27}$$

where $\varphi_{ij} = \alpha_{ij} + \gamma_{ij}$ and with

$$\begin{aligned} \sum_{j \in O_\varnothing} X_{ij} &= Y_i, \ \forall i \in O_0, &(28)\\ \text{and} \sum_{i \in O_0} X_{ij} &= X_j, \forall j \in O, &(29) \end{aligned}$$

with $X_{0\varnothing} = 0$.

Clearly, this solution is of the form of a typical gravity equation, since it contains a bilateral component (φ_{ij}) and two unilateral components associated with both sides of the market (s_i and m_j). Interestingly, the matching foundation offers an equilibrium transfer equation. Indeed, using $X_{ij}^S = X_{ij}^D = X_{ij}$ in equilibrium, solving equation (25) for w_{ij} and substituting the equilibrium expression of X_{ij} in equation (**??**) for X_{ij}^D, one obtains equilibrium outcome as[12]

$$\begin{aligned} X_{ij} &= \exp\left(\frac{\varphi_{ij} - s_i - m_j}{2}\right) \forall i \in O_0, j \in O, \\ X_{i\varnothing} &= \exp\left(\alpha_{i\varnothing} - s_i\right), \forall i \in O, \\ w_{ij} &= \gamma_{ij} - \frac{1}{2}\varphi_{ij} + \frac{1}{2}\left(s_i - m_j\right), \forall i \in O_0, j \in O. \end{aligned}$$

C.5 Identification

Assume that one has access to the data $\mathcal{D} = (X, W)$, that is, data on matches and transfers (i.e. wages). Then in what follows, we show that the productivity channel (γ_{ij}) is identified separately from the preference channel (α_{ij}). The intuition is that while productivity increases both the flow and the transfer, preferences increase the flow but decrease the transfer.

Formally, consider that the double difference operator Δ^2 applied to a variable Y_{ij} re-

[12]Note that this can also be written as

$$\begin{aligned} X_{ij} &= \exp\left(\frac{(\gamma_{ij} - m_j) + (\alpha_{ij} - s_i)}{2}\right) \forall i \in O_0, j \in O, \\ w_{ij} &= \frac{1}{2}\left((\gamma_{ij} - m_j) - (\alpha_{ij} - s_i)\right), \forall i \in O_0, j \in O. \end{aligned}$$

This notation makes the source of identification more transparent, as we show in the next section.

turns:

$$\Delta^2 Y_{ij} = \left[Y_{ij} - Y_{kj}\right] - \left[Y_{il} - Y_{kl}\right], \forall i \neq k, j \neq l.$$

Then, using the gravity equation note that

$$\begin{aligned}
\Delta^2 \log X_{ij} &= \frac{1}{2}\Delta^2 \left[\varphi_{ij} - s_i - m_j\right] \\
&= \frac{1}{2}\left(\Delta^2\varphi_{ij} - \Delta^2 s_i - \Delta^2 m_j\right) \\
&= \frac{1}{2}\Delta^2\varphi_{ij} \\
&= \frac{1}{2}\left(\Delta^2\gamma_{ij} + \Delta^2\alpha_{ij}\right).
\end{aligned}$$

However, note also that applying the operator on transfers one has

$$\begin{aligned}
\Delta^2 w_{ij} &= \Delta^2 \left[\frac{1}{2}\left(\left(\gamma_{ij} - m_j\right) - \left(\alpha_{ij} - s_i\right)\right)\right] \\
&= \frac{1}{2}\left(\Delta^2\gamma_{ij} - \Delta^2\alpha_{ij}\right).
\end{aligned}$$

It follows that rearranging these two results to express unknowns in terms of data $\mathcal{D}$, one has the following identification result:

$$\begin{aligned}
\Delta^2\alpha_{ij} &= \Delta^2 \log X_{ij} - \Delta^2 w_{ij}, \\
\Delta^2\gamma_{ij} &= \Delta^2 \log X_{ij} + \Delta^2 w_{ij}.
\end{aligned}$$

The amenities and productivity can be identified separately using the data on flows (X) and transfers (W).

C.6 Computation

A clear advantage of the micro-foundation of the gravity equation through our matching model is in providing us with an algorithm to compute the equilibrium associated with counterfactuals of interest. To see this, use the equilibrium expressions of X_{ij} and $X_{i\emptyset}$

(equations **??** and **??**) into the accounting constraints (equations 20 and 21), to obtain[13]

$$\exp\left(\alpha_{i\emptyset} - s_i\right) + \sum_{j \in O} \exp\left(\frac{\varphi_{ij} - s_i - m_j}{2}\right) = Y_i, \ \forall i \in O_0, \quad (30)$$

$$\sum_{i \in O_0} \exp\left(\frac{\varphi_{ij} - s_i - m_j}{2}\right) = X_j, \forall j \in O. \quad (31)$$

By simple factorization, this system can be written as

$$X_{i\emptyset} + X_{i\emptyset}^{1/2} \sum_{j \in O} K_{ij} M_j^{1/2} = Y_i, \ \forall i \in O_0, \quad (32)$$

$$M_j^{1/2} \sum_{i \in O_0} K_{ij} X_{i\emptyset}^{1/2} = X_j, \forall j \in O, \quad (33)$$

where $K_{ij} = \exp\left(\frac{\varphi_{ij} - \alpha_{i\emptyset}}{2}\right)$, $M_j = \exp\left(-m_j\right)$ and $X_{0\emptyset} = 0$.

The first equation of the system is a quadratic equation of the form

$$z^2 + 2Pz = Y$$

for $z = X_{i\emptyset}^{1/2}$, $Y = Y_i$ and $P = \frac{1}{2} \sum_{j \in O} K_{ij} M_j^{1/2}$, whose solution[14] is

$$z^2 = \left(\left(Y + P^2\right)^{1/2} - P\right)^2. \quad (34)$$

[13]Note that setting $\alpha_{0\emptyset} \to -\infty$, one has $X_{0\emptyset} = \exp\left(\alpha_{0\emptyset} - s_0\right) \to 0$.

[14]The solution is obtained by completing the square, i.e.

$$\begin{aligned} z^2 + 2Pz &: = (z+P)^2 - P^2 = Y \\ &\Leftrightarrow \\ (z+P)^2 &= Y + P^2 \\ &\Leftrightarrow \\ z &= -P + \left(Y + P^2\right)^{1/2} \\ &\Leftrightarrow \\ z^2 &= \left(\left(Y + P^2\right)^{1/2} - P\right)^2 \end{aligned}$$

It follows that the system can be expressed in terms of X_{i0} and M_j as follows:

$$X_{i\emptyset} = \left(\left(Y_i + \left(\frac{1}{2} \sum_{j \in O} K_{ij} M_j^{1/2} \right)^2 \right)^{1/2} - \frac{1}{2} \sum_{j \in O} K_{ij} M_j^{1/2} \right)^2, \tag{35}$$

$$M_j = \left(\frac{X_j}{\sum_{i \in O_0} K_{ij} X_{i\emptyset}^{1/2}} \right)^2. \tag{36}$$

This system actually provides an IPFP algorithm that admits a fixed point (see Chen et al. (2021)) which can be achieved by solving successively the first set of equations for $X_{i\emptyset}$ given all M_j 's and then the second set of equations for M_j given the solutions for $X_{i\emptyset}$'s obtained at the previous step.

This means that for known quantities for $\left(\varphi_{ij}, \alpha_{i\emptyset}\right)_{i,j}$ and $\left(Y_i, X_j\right)_{i,j}$, one can use the above algorithm to solve for an equilibrium $\left(X_{ij}, w_{ij}\right)_{i,j}$. We use this algorithm to compute the equilibrium associated with each of our counterfactuals once the parameters of the utilities $\left(\varphi_{ij}, \alpha_{i\emptyset}\right)_{i,j}$ have been estimated using the method outlined in the next section.

C.7 Estimation

Recall that jobs' types are defined by a vector of required skills, whereas workers' types are defined by a vector of possessed skills. For each occupation j and each worker i the distance between the required skills and the skills of the worker can be calculated using classical metrics (for example, Euclidean distance).

Let D_{ij}^k be a measure of the distance between the skills required for a job of type j and the skills of a worker of type i. For instance, one could define a measure of distance using the Euclidean norm

$$D_{ij}^1 = ||z_i - z_j||$$

where z_i is the vector of skills of a worker of type i and z_j is the vector of skills required for a job of type j.

Suppose that we parametrize $\alpha_{ij}^a = \sum_{k=1}^K a_k D_{ij}^k$ and $\gamma_{ij}^b = \sum_{k=1}^K b_k D_{ij}^k$ so that $\varphi_{ij}^\beta = \sum_{k=1}^K \beta_k D_{ij}^k$ where $\beta_k = a_k + b_k$. and D_{ij}^k are K basis functions of the "distance" between workers' types and jobs' types. With this parametrization of the model, the gravity equation now becomes

$$X_{ij} = \exp\left(\frac{\sum_{k=1}^{K}\beta_k D_{ij}^k - s_i - m_j}{2}\right) \forall i \in O_0, j \in O, \tag{37}$$
$$X_{i\emptyset} = \exp\left(-s_i\right), \forall i \in O,$$

assuming $\alpha_{i\emptyset} = 0$.

As recently shown in Galichon and Salanié (2024), this parametric version of the Choo and Siow (2006) equation can be estimated using GLM models, and in particular Pseudo-Poisson Maximum Likelihood as for the classical gravity equation. The main difference lies in the specification of appropriate weights (all terms in the exponential are divided by a factor 2 for pairs (i,j) unlike for transitions to not employed).

We therefore estimate the parameters (β, s, m), where s and m are workers' type fixed effects and jobs' type fixed-effects respectively, using the command ppmlhdfe in Stata. We herewith obtain estimates $\hat{\varphi}_{ij}^{\beta} = \sum_{k=1}^{K}\hat{\beta}_k D_{ij}^k$, $\hat{s}_i$ and $\hat{m}_j$ of the parameters of the model.

However, note that the model also provides a solution for the equilibrium transfers which given our parametrization now read as

$$w_{ij} = \gamma_{ij}^b - \frac{1}{2}\varphi_{ij}^{\beta} + \frac{1}{2}\left(s_i - m_j\right), \forall i \in O_0, j \in O. \tag{38}$$

Using the estimates from the gravity equation one can compute the variable

$$y_{ij} = w_{ij} - \left(-\frac{1}{2}\hat{\varphi}_{ij}^{\beta} + \frac{1}{2}\left(\hat{s}_i - \hat{m}_j\right)\right)$$

where w_{ij} are observed (log) wages. It follows that the parameters $(b_k)_k$ can be estimated applying a simple OLS regression of y_{ij} on the basis functions $\left(D_{ij}^k\right)_k$. This means that we recover estimates of the productivity parameters $\hat{b}_k$ and the amenity parameters $\hat{a}_k = \hat{\beta}_k - \hat{b}_k$. Appendix C.9 presents an extension in which we also incorporate information on wages at t.

C.8 Model Extension 1: including non-filled jobs by types of occupation

We still consider a highly differentiated labor market where workers and jobs are grouped into types.

The supply side is unchancged.

On the demand side, jobs are differentiated by their type denoted $j \in O$. There is a mass X_j of jobs of type j. However, we append the set of potential types of workers from which employers can choose from with $\{\emptyset\}$ which indicate the job is not filled. Hence employers can choose among $O_0^{\emptyset} = O_0 \cup \{\emptyset\}$.

It follows that Y_i is the mass of workers that were employed in occupation $i \in O$ at t whereas X_i is the mass of jobs available in occupation $i \in O$.

The worker's problem is unchanged but now the employer's problem reads as

$$\max_{i \in O_0^{\emptyset}} \gamma_{ij} - w_{ij} + \eta_i$$

where η_i is an idiosyncratic taste for workers of type i, γ_{ij} is the systematic productivity for a worker of type i in occupation j and w_{ij} is the wage paid by employers in occupation j to workers of type i. Note that w_{0j} needs not be 0 as it corresponds to the wage employers in occupation j have to pay to workers that were not employed previously, i.e. of type 0. In contrast, since $X_{\emptyset j}$ is the mass of vacant (unfilled) jobs in occupation j, one has $w_{\emptyset j} = 0$.

The problems on the two sides solve to yield

$$\begin{aligned} \log X_{ij} &= \alpha_{ij} + w_{ij} - s_i \forall i \in O_0, j \in O_0, \\ \log X_{ij} &= \gamma_{ij} - w_{ij} - m_j \forall i \in O_0^{\emptyset}, j \in O. \end{aligned}$$

Equilibrium is then characterized by

$$
\begin{aligned}
X_{ij} &= \exp\left(\frac{\varphi_{ij} - s_i - m_j}{2}\right), \forall i \in O_0, j \in O \\
X_{i0} &= \exp\left(\alpha_{i0} - s_i\right), \forall i \in O_0 \\
X_{\emptyset j} &= \exp\left(\gamma_{\emptyset j} - m_j\right), \forall j \in O
\end{aligned}
$$

where $\varphi_{ij} = \alpha_{ij} + \gamma_{ij}$ and with

$$
\begin{aligned}
\sum_{j \in O_0} X_{ij} &= Y_i,\ \forall i \in O_0 \\
\text{and} \sum_{i \in O_0^{\emptyset}} X_{ij} &= X_j, \forall j \in O.
\end{aligned}
$$

More specifically the accounting constraints are

$$
\begin{aligned}
\underset{\text{Employed in } i \text{ at } t \text{, not employed at } t+1}{X_{i0}} + \underset{\text{Employed in } i \text{ at } t \text{, employed at } t+1}{\sum_{j \in O} X_{ij}} &= \underset{\text{Workers in } i \text{ at } t \text{, from } i \text{ at } t+1}{Y_i} \\
, \forall i &\in O_0 \\
&\text{and} \\
\underset{\text{Vacant jobs in } j \text{ at } t+1}{X_{\emptyset j}} + \underset{\text{Filled jobs in } j \text{ at } t+1}{\sum_{i \in O_0} X_{ij}} &= \underset{\text{Total number of jobs in } j \text{ at } t+1}{X_j} \\
, \forall j &\in O.
\end{aligned}
$$

C.9 Model Extension 2: How to incorporate wages at t

Remember that the transfer w_{ij} was defined as the (log) wage at $t+1$ received by workers from i in j. These workers were in i at t and are in j at $t+1$.

Let us slightly change the notation in order to introduce the time dimension more explicitely.

Let w_{ij}^{t+1} be the (log) wage at $t+1$ recevied by workers having moved from i to j between t and $t+1$.

Similarily, let w_{ij}^t be the (log) wage at t received by workers having moved from i to j between t and $t+1$.

Let assume that the transfer w_{ij} is in fact the (log) wage differential

$$w_{ij} = w_{ij}^{t+1} - w_{ij}^t$$

it is therefore defined as the log wage differential necessary to attract a worker previously employed in occupation i at t to work in occupation j at $t+1$.

Note that we still have that

$$\begin{aligned} \log X_{ij} &= \frac{\alpha_{ij} + w_{ij} - s_i}{\sigma_1} \forall i \in O_0, j \in O_0, \\ \log X_{ij} &= \frac{\gamma_{ij} - w_{ij} - m_j}{\sigma_2} \forall i \in O_0, j \in O. \end{aligned}$$

It follows that in equilibrium

$$\begin{aligned} X_{ij} &= \exp\left(\frac{\varphi_{ij} - s_i - m_j}{2}\right), \forall i \in O_0, j \in O \\ X_{i0} &= \exp\left(\alpha_{i0} - s_i\right), \forall i \in O_0 \end{aligned}$$

so the equilibrium flows are not affected by our choice of definition for the transfer and they can still be estimated using the same technique (Poisson regression) and will yield the exact same results.

However, even though equilibrium transfers are still given as

$$w_{ij} = \frac{1}{2}\left(\left(\gamma_{ij} - m_j\right) - \left(\alpha_{ij} - s_i\right)\right), \forall i \in O_0, j \in O$$

the interpretation of the transfer as changed and is $w_{ij} = w_{ij}^{t+1} - w_{ij}^t$ instead of w_{ij}^{t+1}. It follows that the estimation of the transfer regression becomes a regression of the (log) difference in earnings instead of a (log) earnings regression. One indeed now has

$$\begin{aligned} \hat{w}_{ij} &\equiv w_{ij}^{t+1} - w_{ij}^t \\ &= \gamma_{ij} - \log X_{ij} - m_j + e_{ij} \end{aligned}$$

So this rewrites as

$$w_{ij}^{t+1} = \gamma_{ij} - \log X_{ij} - m_j + w_{ij}^t + e_{ij}$$

and we note that this is a (log) earnings regression as before except that we now additionally control for w_{ij}^t forcing a coefficient of 1.

C.10 Model's fit

Figure C.1 evaluates the model's ability to reproduce the observed pattern of occupational mobility. The figure plots binned averages of the bilateral worker flows predicted by the estimated model against the corresponding flows observed in the data. The fitted relationship is close to the 45-degree line, indicating that the model captures the average relationship between predicted and observed mobility flows. The correlation between predicted and observed bilateral flows is 0.48, suggesting that the model explains a substantial share of the cross-occupation variation in worker mobility despite the considerable idiosyncratic heterogeneity present at the bilateral level.

Figure C.1: Model fit

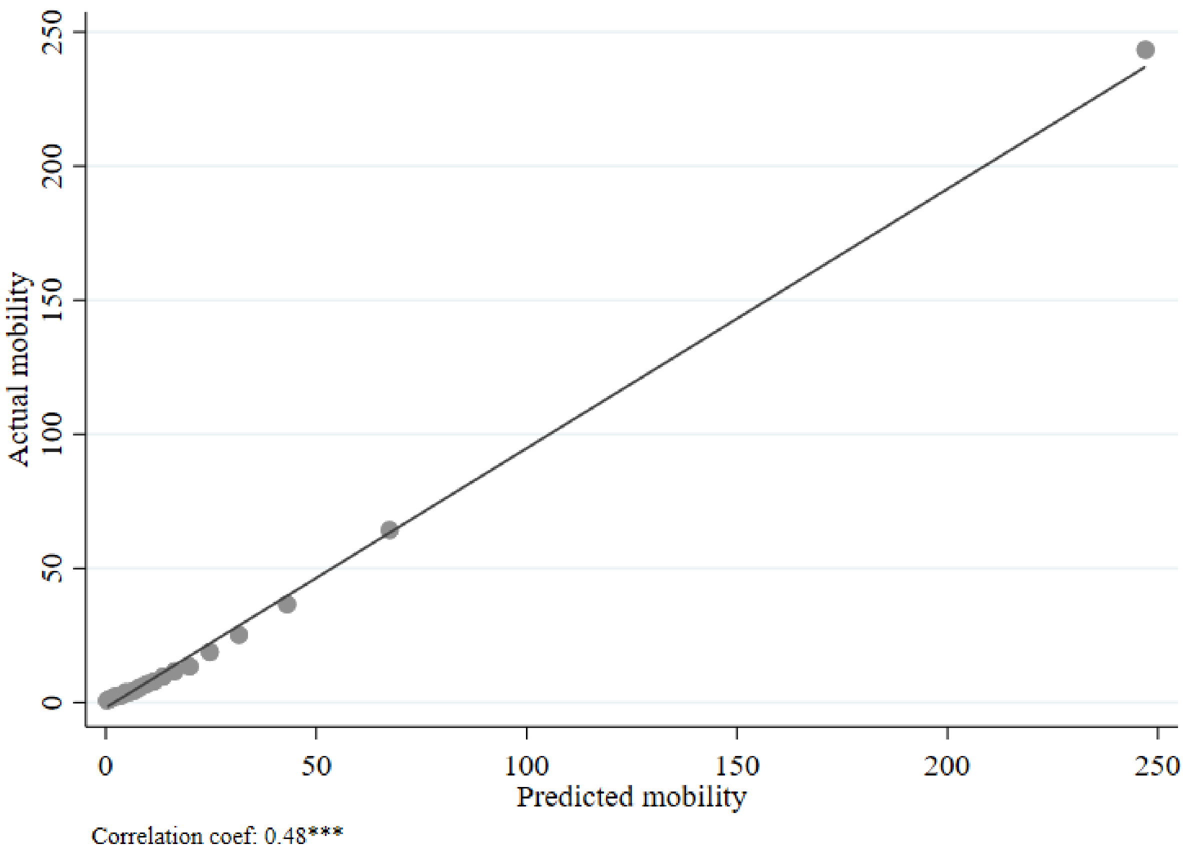

Notes: The figure plots the binned correlation between the bilateral flows predicted by the model and the corresponding bilateral flows observed in the data. The solid line shows the fitted linear relationship. The correlation coefficient between predicted and observed flows is reported at the bottom of the figure.