ROCKWOOL Foundation Berlin
Institute for the Economy and the Future of Work (RFBerlin)

DISCUSSION PAPER SERIES 062/26

Digitalization, Change in Skill Distance Between Occupations and Occupational Mobility

Arnaud Dupuy, Morgan Raux, Sara Signorelli

www.rfberlin.com MARCH 2026

Digitalization, Change in Skill Distance Between Occupations and Occupational Mobility

Authors

Arnaud Dupuy, Morgan Raux, Sara Signorelli

Reference

JEL Codes: J23, J24, J62

Keywords: Occupation mobility, Technological change, Matching

Recommended Citation: Arnaud Dupuy, Morgan Raux, Sara Signorelli (2026): Digitalization, Change in Skill Distance Between Occupations and Occupational Mobility. RFBerlin Discussion Paper No. 062/26

Access

Papers can be downloaded free of charge from the RFBerlin website: https://www.rfberlin.com/discussion-papers

Discussion Papers of RFBerlin are indexed on RePEc: https://ideas.repec.org/s/crm/wpaper.html

Disclaimer

Imprint

RFBerlin
ROCKWOOL Foundation Berlin –
Institute for the Economy
and the Future of Work

Gormannstrasse 22, 10119 Berlin
Tel: +49 (0) 151 143 444 67
E-mail: info@rfberlin.com
Web: www.rfberlin.com

Digitalization, Change in Skill Distance Between Occupations and Occupational Mobility*

Arnaud Dupuy (University of Luxembourg and IZA),
Morgan Raux (Aix-Marseille School of Economics and IZA),
Sara Signorelli (CREST - Ecole Polytechnique and IZA)

February 2026

Abstract

Technological change affects labor markets not only by shifting labor demand across occupations, but also by reshaping the skill distances that govern workers' ability to move between jobs. This paper studies the digitalization wave of the 2010s using task data from online job postings, matched employer–employee data, and a gravity framework of occupational mobility. We show that while most occupations became more digital, skill distances converged for some occupation pairs and diverged for others, increasing mobility along some pathways and reducing it along others. Counterfactual simulations show that these frictions are meaningful and slow reallocation out of shrinking occupations.

JEL Classification: J23, J24, J62.

Keywords: Occupation mobility, Technological change, Matching.

*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. Contact the corresponding author, Sara Signorelli, at sara.signorelli@polytechnique.edu.

1 Introduction

Large waves of technological change are major sources of disruption in labor markets, reshaping how work is performed and creating substantial challenges for workers' ability to adapt. As new technologies diffuse, some jobs expand while others decline, forcing workers to reallocate across occupations in order to remain employed and productive. The costs workers face when transitioning between occupations play a central role in determining who benefits from technological progress and who is left behind. Yet, while a large literature has documented how technology shifts labor demand across occupations, much less is known about how technological change reshapes the pathways through which workers adjust. In particular, by altering the skill requirements of jobs, new technologies may change how close or distant occupations are from one another in the skill space, lowering mobility costs for some transitions while raising them for others. This paper investigates this margin of adjustment and quantifies its role in shaping occupational mobility and its distributional consequences in the wake of the digitalization shock that characterized the 2010s.

A large literature studies the labor-market effects of technological change by focusing on shifts in labor demand across occupations, emphasizing complementarities with new technologies and substitution effects.[1] More recent work emphasizes that occupations are bundles of multidimensional tasks, highlighting how technological change reshapes the skill content of jobs in more nuanced ways.[2] While this body of work has substantially improved our understanding of which occupations expand or contract following technological shocks, it typically abstracts from how technological change reshapes the costs workers face when moving between occupations. As a result, much less is known about how changes in skill similarity across jobs affect workers' ability to adjust to technological change through occupational mobility.

[1] This literature comprises the theory of skill-biased technological change (e.g.Katz and Murphy, 1992; Autor et al., 1998; Berman et al., 1998; Card and DiNardo, 2002) and its subsequent extension to routine-biased technological change (e.g.Autor et al., 2003; Goos and Manning, 2007; Acemoglu and Autor, 2011; Goos et al., 2014), which together constitute the main analytical frameworks through which economists have examined the labor-market effects of new technologies. Recent work has applied this framework to AI (Webb, 2019; Eloundou et al., 2024; Gathmann et al., 2024; Bloom et al., 2025; Hampole et al., 2025).

[2] See for instance Deming and Noray (2020); Braxton and Taska (2023); Lipowski et al. (2024); Althoff and Reichardt (2025); Autor and Thompson (2025). This development is also made possible by the increased availability of detailed task data coming from job adds (e.g. Deming and Kahn, 2018; Atalay et al., 2020; Modestino et al., 2020).

In this paper, we examine how technology-driven changes in the skill content of occupations also affect workers' labor market prospects by altering the skill distances between jobs. We propose a framework in which both workers and occupations are represented as multidimensional bundles of skills and tasks. Within this framework, technological change has two distinct effects: it shifts employment across occupations by changing occupational demand, and it alters the costs of moving between occupations by reshaping skill distances. We quantify the relative importance of these two channels for occupational mobility and assess their distributional and welfare implications.

We take this framework to the data by studying the digitalization wave of the 2010s and its impact on occupational mobility in France. This period is characterized by the widespread diffusion of digital technologies—such as cloud computing, data analytics, and machine learning—across most sectors of the economy, together with the growing prevalence of basic IT skill requirements across a wide range of occupations, including low-skilled ones. We measure the multi-dimensional task content of occupations and their evolution using the near-universe of online job postings collected by Lightcast from U.S. job boards between 2010 and 2019. Constructing skill measures from U.S. data provides an important identification advantage, as these measures are unlikely to be influenced by occupational mobility or institutional features specific to the French labor market and instead primarily reflect technology-driven changes. We combine these measures with matched employer–employee data covering the universe of the French labor force over the same period, providing a well-suited setting to study how changes in digital skill distance affect occupational mobility.

We begin by examining how technological change reshaped the digital skill content of occupations during the 2010s. We measure digitalization by tracking changes in the Euclidean distance between each occupation and a set of reference IT occupations, chosen to capture the bundle of skills required to work with digital technologies. We compute changes in digital skill distance for all pairs of occupations based on their relative proximity to these IT reference jobs. We show that 80 percent of occupations became more similar to IT jobs over the period, reflecting the pervasive diffusion of digital tasks across all segments of the labor market. Yet, this widespread digitalization did not translate into a uniform convergence across occupations. About half of the occupation pairs experienced a decline in digital skill distance—implying lower mobility costs—while the remaining pairs became more distant, facing higher transitional frictions. This heterogeneity indicates that digitalization simultaneously opened new mobility pathways for some workers while creating new barriers for others.

Next, we assess whether these heterogeneous changes in digital skill distance 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 distance, 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 digital skill distance and to alternative specifications.

Gravity models provide a natural framework to study workers' flows across occupations. Building on this approach, we micro-found our gravity equation using a discrete version of the two-sided matching model developed by Galichon and Salanié (2022a). This framework provides a structured interpretation of occupational mobility and its underlying determinants, and delivers three main advantages. First, it supports the interpretation of the reduced-form gravity estimates by explicitly modeling occupational mobility. Second, it allows us to conduct counterfactual simulations and quantify how occupational mobility would have evolved in the absence of changes in digital skill distance, using mobility shifts driven by changes in own-occupation demand-the primary focus of the technological change literature-as a benchmark for magnitude. Third, the model enables us to uncover the mechanisms behind distance-driven mobility by decomposing observed transitions into productivity-related and non-productivity-related components, leveraging detailed wage information from French administrative data.

Using counterfactual simulations of occupational mobility, we quantify the contribution of technological change and disentangle the roles of demand shifts and changes in digital skill distance. As expected, during the 2010s, demand shocks pushed workers out of shrinking, routine-intensive occupations such as clerical jobs and attracted workers into high-skill occupations. Changes in digital skill distance, however, operated differently: rising distance significantly reduced mobility out of clerical and blue-collar occupations, slowing reallocation along these margins. These distance-induced frictions are economically meaningful, comparable in magnitude to own-occupation demand shocks, and act to dampen adjustment out of declining occupations. In the short run, this does not necessarily imply worse outcomes for low-skill workers, as many demand-driven outflows are toward lower-paying occupations, and increased digital distance disproportionately lowers such downward moves. Over longer horizons, however, higher relocation frictions may increase the risk of unemployment by limiting workers' ability to adjust to adverse demand shocks. For high-skill workers, distance effects are unambiguously positive: they reduce transitions to lower-paying jobs while increasing moves to higher-paying ones. More broadly, we show that digital distance shapes mobility through both productivity and non-productivity channels. While mobility from high-skill occupations is primarily driven by wage gains, a substantial share of mobility from lower-skill occupations reflects non-wage amenities.

This paper contributes to the literature on the labor market effects of technological change, with particular emphasis on work studying the recent digitalization shock (Webb, 2019; Deming and Noray, 2020; Acemoglu et al., 2022; Gathmann et al., 2024), and on studies analyzing how technological change shapes workers' flows across occupations (Cortes, 2016; Adão et al., 2024; Battisti et al., 2023; Bessen et al., 2023; Edin et al., 2023; Bocquet, 2024). Most closely related to our work, Adão et al. (2024) and Bocquet (2024) show that adjustments to technological shocks are slowed by transition costs between occupations, particularly when the affected occupations are more skill-distant from others. Our contribution differs in both focus and scope. Rather than quantifying aggregate output losses from slow reallocation, we isolate and measure the role of changes in skill distance as an indirect channel through which technological change affects occupational mobility, benchmarking its magnitude against the more widely studied own-occupation demand effect, and we document how its implications vary across different groups of workers and firms.

Second, our work contributes to the literature on the skill content of jobs and the mobility costs associated with job-to-job transitions (Cortes and Gallipoli, 2018; Gathmann and Schönberg, 2010; Lazear, 2009; Lise and Postel-Vinay, 2020; Yamaguchi, 2012) and with transitions involving unemployment spells (Poletaev and Robinson, 2008; Azmat et al., 2024; Dabed et al., 2025; Klaeui et al., 2026). Following this literature, we conceptualize occupations as multidimensional bundles of tasks and reject the notion of distinct labor markets. Instead, we emphasize that occupations are interconnected through worker flows, which are larger between occupations with more similar skill requirements. Building on Cortes and Gallipoli (2018), who quantify the costs of fixed skill distances using a gravity model, we show that changes in skill distance significantly affect worker mobility, particularly during periods of rapid technological change. We focus specifically on shifts driven by the digitalization shock, allowing us to compare these effects with those driven by changes in skill demand. Finally, unlike their random utility framework, we introduce a micro-founded matching model that accounts for congestion effects from supply and demand, capturing the competitive constraints workers face when moving between occupations.

Finally, our model-based distinction between productivity-driven and non-productivity-driven effects of changes in skill distance relates closely to a growing literature that interprets worker mobility toward lower-paying jobs as revealed preference for non-wage amenities.[3]

The remainder of the paper is organized as follows. Section 2 outlines our conceptual framework for analyzing the effects of technological change. Section 3 describes the data. Section 4 introduces the gravity specification. Section 5 reports the empirical results. Section 6 develops the matching model that micro-founds the gravity equation. Section 7 conducts counterfactual exercises and disentangles the underlying mechanisms. Section 8 concludes.

2 Conceptual framework

In this section, we formalize our conceptual framework for thinking about the effects of technological change on the labor market. At the center of our framework is the notion of skills distance and, in particular, the distance between the skills a worker possesses and the skills required for a job.

Let z^p be a vector of Z skills a worker possesses, grouping workers into discrete types so that all workers of type i have skills $z^p = z_i^p$. Similarly, let jobs be defined by their vector z^r of Z required skills, grouping jobs into discrete types of occupations so that all jobs of type j have required skills $z^r = z_j^r$.[4] Although the skills possessed and the skills required are two distinct concepts, there is a clear mapping between the two. The first relates to the knowledge and know-how possessed by individuals. The second

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

[4]Implicit in this setting is the idea that required skills are associated to tasks. For instance, a required skill might be to program in Python. The task "programming in Python" is associated with the required skill of the same name. For all practical matters, we use required skills and tasks interchangeably.

relates to the activities to be performed on the job. In particular, the bundle of skills possessed by each individual determines its relative productivity in each occupation. Optimal productivity in an occupation j that requires skills z_j^r is assumed to be obtained with workers whose skills perfectly match those requirements. Consider, for example, workers of type i with skills z_i^p. If $z_i^p = z_j^r$, then the skills of workers of type i match perfectly those required in the occupation j. In contrast, if $z_i^p \neq z_j^r$, the skills of workers of type i do not perfectly match those required by occupation j, i.e. the skills distance between the skills possessed by the worker and those required by the occupation is greater than 0. We denote the distance between the skills possessed by the worker and those required by the occupation by $d\left(z_i^p, z_j^r\right)$.[5]

Figure 1: Visualization of the conceptual framework

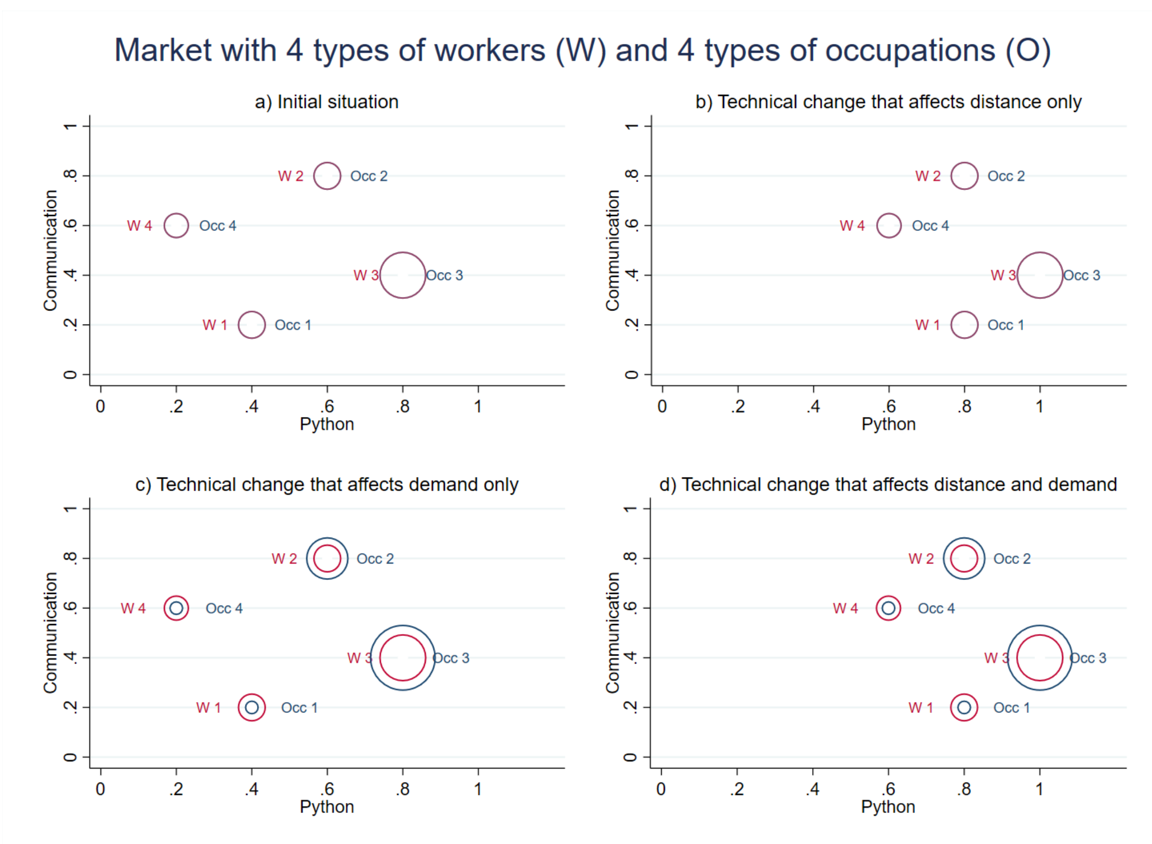

Notes: Circle's size reflects mass of jobs/workers. The figure sketches the two distinct effects of technological change in the context of our conceptual framework.

Figure 1 shows a simplified version of our conceptual framework with only two skills ($Z = 2$), communication (vertical axis) and the Python programming language (hori-

[5]In the next section, we introduce the metric used to compute these distances.

zontal axis). There are four types of workers (W1-W4) and four types of occupations (Occ1-Occ4) represented by circles, in red for workers and blue for occupations. The coordinates of the center of each circle correspond to the possessed / required skills of the associated group. For example, workers of type 1 have skills $z_1^p = (0.4, 0.2)$, while occupation 2 requires skills $z_2^r = (0.6, 0.8)$. The distance between the skills possessed by a worker of type 1 and those required in the occupation of type 2 is then $d\left(z_1^p, z_2^r\right)$ which is best understood using the Euclidean distance in this example. The size of each circle indicates the employment weight of the associated type of worker/occupation in the economy.

Panel a) of Figure 1 represents our baseline scenario and corresponds to a situation before the advent of technological change. As drawn, it is assumed that there is a perfect match between the skills possessed by workers of type i and the skills required by the occupations of type i and there are as many jobs of type i as workers of type i (circles of the same size). Hence, at baseline, we assume that workers of type i are matched with occupations of type i and the distance between workers' skills and those required in their occupations is 0. It is important to notice that, even in the absence of technical change, over time, some workers might change occupation because of idiosyncratic shocks. We call this "the natural rate of mobility",[6] and expect it to occur predominantly between occupations close to each other in terms of required skills, i.e. small distance $d\left(z_j^r, z_k^r\right)$ for two occupations $j \neq k$.

With this framework in mind, we can now conceptualize technological changes as bringing about two major effects. First, technological change can trigger changes in the skills required in each occupation. In panel b) of Figure 1 we see that while communication requirements remain unchanged in all occupations, all occupations face a growing requirement for python, and more so in occupations that had low python requirements at baseline. It is also important to note that, as depicted, we make the assumption that workers' skills evolve with the skills requirements of their matching occupation at baseline. Hence, workers of type i, who are employed in the occupation of type i at baseline, see their skills evolve as the skills requirements of the occupation i. This assumption is valid when workers learn on-the-job and firms invest in constant training to keep their workers up to date.[7] Importantly, as depicted, technological change results in occupa-

[6]We herewith make a reference to the "natural rate of unemployment" which occurs because of idiosyncratic shocks (search frictions for instance).

[7]Our framework can accommodate the opposite assumption : that workers possess the skills required in the occupation at the moment they are hired, but that their skills become obsolete when the required

tions getting closer to each other, i.e. the distance $d\left(z_j^r, z_k^r\right)$ for each pair of occupations is smaller in panel b) than in panel a). Since our framework predicts that workers move more towards occupations whose required skills are close to their own skills, we can expect some additional mobility of workers of type 2 (W2) towards occupations of type 4 (Occ 4), and some additional mobility of workers of type 3 (W3) towards occupations of type 1 (Occ 1) compared to the baseline situation (without technological change). Finally, note that the size of the circles has remained constant, so that the demand in each occupation and the supply of each type of worker have not changed. This means that the mobility observed in this scenario would be merely the result of 1) "natural mobility" (idiosyncratic shocks) and 2) changes in the skills requirements between occupations.

The second major effect of technological change is that it varies the demand for different occupations and thus their relative size in the labor market. This effect is depicted in panel c) of Figure 1. In this example, while the supply of workers is the same as in the baseline situation (same size of red circles), the demand in occupations of types 2 and 3 has increased (larger blue circles) while the demand in occupations of types 1 and 4 has declined (smaller blue circles). However, note that the skills requirements are the same as in the baseline situation. Everything else equal, we expect increased mobility of workers away from shrinking occupations towards growing occupations, and more so towards growing occupations that are closer. The mobility observed in this scenario would be simply the result of 1) "natural mobility" and 2) changes in demand.

Finally, the total effect of technological change is shown in panel d) of Figure 1, where both the distribution of occupations and the skills requirements have changed. The observed mobility would then combine the three components: 1) "natural mobility" 2) changes in the required skills, and 3) changes in demand. The remainder of the paper focuses on disentangling the respective roles of these two effects of technological change by expressing them as shares of the underlying "natural rate of mobility".

---

skills of their occupation evolves (Deming and Noray, 2020). In our empirical analysis, we present robustness tests where the distance between occupations, including once own, is defined using the initial bundle of required skills in the origin occupation and the final bundle of required skills in the destination occupation.

3 Digitalization and change in skill distance

3.1 Measuring skill distances

To document how technological change during the 2010s reshaped the skill requirements of jobs, we exploit detailed information on the skills listed in online job postings. We characterize the evolution of occupational skill content using three related measures. First, we capture digitalization at the occupational level by tracking changes in the Euclidean distance between each occupation and a set of reference IT occupations, which proxy for the bundle of skills associated with digital technologies. This measure forms the basis of our pairwise analysis. Second, we construct a measure of digital skill distance between every pair of occupations, defined as the difference in their respective distances to the IT reference occupations. This measure captures changes in relative digital proximity across occupations. Third, we compute an overall skill-distance measure for each occupation pair that reflects differences in their full skill profiles, independent of the digital dimension. We use this broader measure as a control to distinguish the effects of digital skill distance from changes in overall skill similarity.

We construct our measures of occupational skill content using data from Lightcast (formerly Burning Glass Technologies), which cover online job postings in the United States between 2011 and 2019. A key advantage of this data source is its extensive coverage: Lightcast collects the near-universe of online vacancies by daily web-scraping approximately 40,000 job boards and firm websites, identifying roughly 3.4 million active postings at any point in time. This coverage is widely viewed as capturing the vast majority of vacancies posted online in the United States during this period. Lightcast further applies text-analysis algorithms to remove duplicate postings and to classify job characteristics into standardized formats. Of particular relevance for our analysis, the data extract detailed information on required skills—covering roughly 13,000 distinct skill descriptors—which are mapped to SOC occupation codes. For more details on the data, see Carnevale et al. (2014). Owing to these features, Lightcast data have been widely used to study job characteristics and the evolution of skill demand.[8]

Because our objective is to measure changes in mobility costs between occupations, we

[8] See for instance Bloom et al. (2021); Dillender and Forsythe (2022); Acemoglu et al. (2022); Braxton and Taska (2023).

refine the skill set by excluding skills that can be considered as generic, since present across the majority of job ads, regardless of the occupation. We define generic skills as those appearing in more than 90% of job postings and exclude them from the analysis, leaving us with a set of skills that we believe are informative about occupational specialization.[9]

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 exogenous to French occupational mobility, affecting it only through shared technological change. To apply these measures to French administrative data, we map U.S. SOC occupation codes to French PCS occupation codes.

We consider each occupation j as a vector of Z required skills in year t denoted $z_j^{r,t} = \left(z_{j,1}^{r,t}, z_{j,2}^{r,t}, ..., z_{j,Z}^{r,t}\right)$ where $z_{j,k}^{r,t}$ corresponds to the share of ads for occupation j requiring skill k in year t. To avoid well-known issues when comparing vectors, for each occupation j, we normalize vector $z_j^{r,t}$ by its Euclidean norm $\left\|z_j^{r,t}\right\|$, so that $z_j^{r,t}$ is of unit length for all j and t. We measure the skill distance between occupations i and j using the Euclidean distance between the vectors $z_i^{r,t}$ and $z_j^{r,t}$ as follows:

$$d\left(z_i^{r,t}, z_j^{r,t}\right) = \left(\sum_k \left(z_{i,k}^t - z_{j,k}^t\right)^2\right)^{1/2}. \quad (1)$$

Identifying which skills are complementary to the technologies introduced by recent waves of technological change is not straightforward. Rather than relying on an a-priori classification of digital skills, our baseline approach uses the skill requirements listed in online job postings for a list of reference IT occupations, which we view as capturing the relevant bundle of skills needed to work with digital technologies.[10] Changes in the distance to these reference IT occupations therefore provide a flexible measure of how closely each occupation aligns with the digital skill frontier over time. We thus

[9] In total, out of the roughly 13,000 skills extracted by Lightcast, we classify 134 as generic, including items such as "Analytical Skills," "Speaking English," "Communication Skills," and "Time Management."

[10] The appendix Table A1 presents the list of such reference occupations, which in short includes all jobs in the category of IT engineers and IT technicians.

measure the digitalization of occupations by tracking changes in the Euclidean distance between each occupation's skill profile and that of our set of reference IT occupations. We consider this list of IT occupations as a composite reference category and denote this digital occupation $j = d$, where d stands for digital. In order to quantify the degree of digitalization of an occupation j we construct the scalar q_j^t as minus the Euclidean distance in required skills between that occupation j and the required skills in the typical digital occupation d:

$$\begin{aligned} q_j^t &= -d\left(z_j^{r,t}, z_d^{r,t}\right) \\ &= -\left(\sum_k \left(z_{j,k}^{r,t} - z_{d,k}^{r,t}\right)^2\right)^{1/2}. \end{aligned} \tag{2}$$

Finally, using the obtained occupation-level measures of digitalization q_j^t, we derive pairwise indicators of skill distance that capture how technological change affects mobility costs between occupations. For each occupation pair, we thus compute changes over time in their relative distance to the set of reference IT occupations, which defines our measure of digital skill distance :

$$DigiDist_{ij}^t = d\left(q_i^{r,t}, q_j^{r,t}\right) = |q_i^{r,t} - q_j^{r,t}| \tag{3}$$

This measure reflects how digitalization alters the similarity of skill requirements across occupations and, in turn, the ease of transitioning between them. In parallel, we calculate changes in overall skill distance between occupation pairs using the Euclidean metric described in Equation (1), which accounts for differences in their complete skill profiles and serves as a control in our analysis. Finally, we assess the robustness of these measures by constructing alternative distance metrics based on Mahalanobis and angular separation, and by using a more aggregated skill taxonomy provided by Lightcast, which reduces the skill space from roughly 13,000 to about 656 descriptors.

3.2 Digitalization of jobs

We first document the digitalization of occupations by computing changes in the Euclidean distance between each occupation and the reference IT category observed over the period 2011 to 2019, as defined in Equation 2. Figure 2 ranks occupations according to these changes and shows that the skill content of approximately 80 percent of occupations moved closer to that of IT jobs over the period. This pattern provides clear evidence of the broad diffusion of digital skills across the labor market during the 2010s. For the majority of occupations, the reduction in Euclidean skill distance relative to the reference IT category is modest, amounting to less than 10 percent of their initial 2011 distance. However, occupations in the top decile of the distribution experienced substantially larger changes, with reductions in skill distance reaching up to 30 percent of their initial 2011 level.

Figure 2: Percentage change in skill distance relative to reference IT jobs

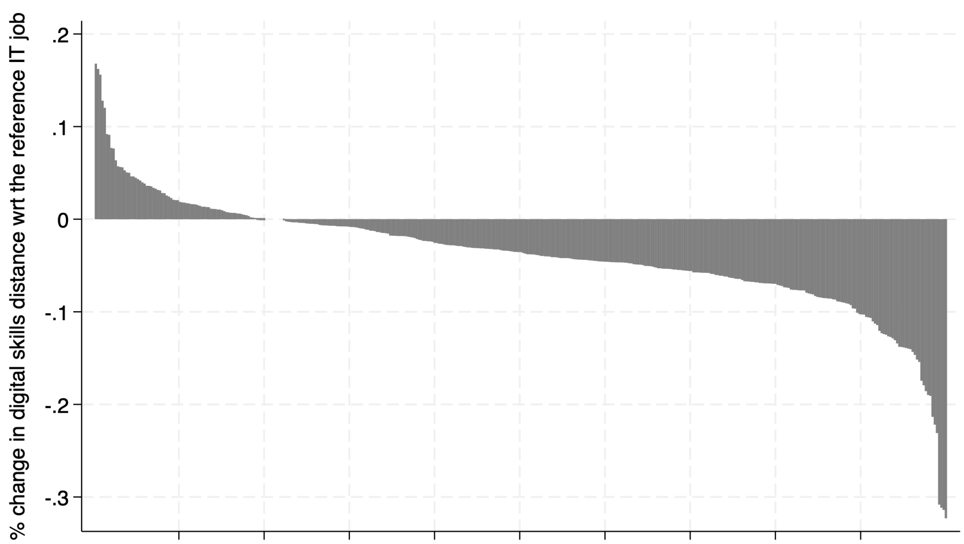

Notes: The figure ranks occupations according to the percentage change in distance relative to the reference category of IT jobs observed between 2011 and 2019. The skill distance measure is computed following equation 2.

The Appendix Figure A1 shows that this pattern is broadly similar across the four main socio-professional categories in the French occupational classification: executives and professionals (including managers and engineers), intermediate professions (including mid-level managers, technicians, and foremen), clerical workers, and blue-collar workers. With the exception of clerical jobs, roughly 80 percent of occupations in each category got closer to IT jobs over the period. For clerks, this share is slightly lower at around 70 percent, still indicating substantial digitalization. These findings confirm the permeation of digital skills across most sectors of activity, including among low-skill jobs. However, the magnitude of these changes varies across groups. Occupations in intermediate professions and clerical categories experienced the smallest changes along the distribution, while the top decile of blue-collar occupations exhibits the largest reductions in distance to IT jobs. For executives, the few occupations that became more distant from IT jobs experienced relatively small changes compared to other categories.

Appendix Table A2 reports the ten occupations that experienced the largest percentage convergence toward IT reference jobs. Most of these occupations are concentrated in electricity and electronics–related activities, suggesting that this industry was among those most exposed to the wave of digitalization during the 2010s. Finally, Appendix Figure A2 validates our digitalization measure by correlating the percentage change in skill distance relative to reference IT jobs with the change in the average number of IT skills mentioned in job ads. Here, IT skills are defined according to our manual classification of the 13,000 Lightcast skills into software-related versus non-software-related categories.[11] We see that, as expected, the closer an occupation became relative to IT, the higher the average number of software-related skills mentioned in the ads as required.

While we document the widespread digitalization of occupations during the 2010s, this process does not necessarily imply uniform convergence in skills across occupations. Figure 3 ranks all occupation pairs by their percentage change in digital skill distance, measured following equation 3, and show that these changes are highly heterogeneous. Approximately half of occupation pairs experienced a decline in digital skill distance—implying lower mobility costs—while the remaining pairs became more distant, facing higher transitional frictions. Moreover, the magnitude of these changes is markedly asymmetric: increases in digital skill distance among diverging pairs are substantially larger than the reductions observed among converging pairs. Appendix Figure A3 shows that this pattern of heterogeneous convergence and divergence is remarkably similar across all major occupational categories, where categories are defined by the origin occupation in the pair. This heterogeneity in the evolution of digital skill distance suggests that digitalization may have uneven effects on workers' occupational mobility depending on their occupation of origin.

[11]Some examples of IT skills include "MS PowerPoint", "MS Excel", "Python", and "GitHub".

Figure 3: Percentage change in digital skill distance between all occupation pairs

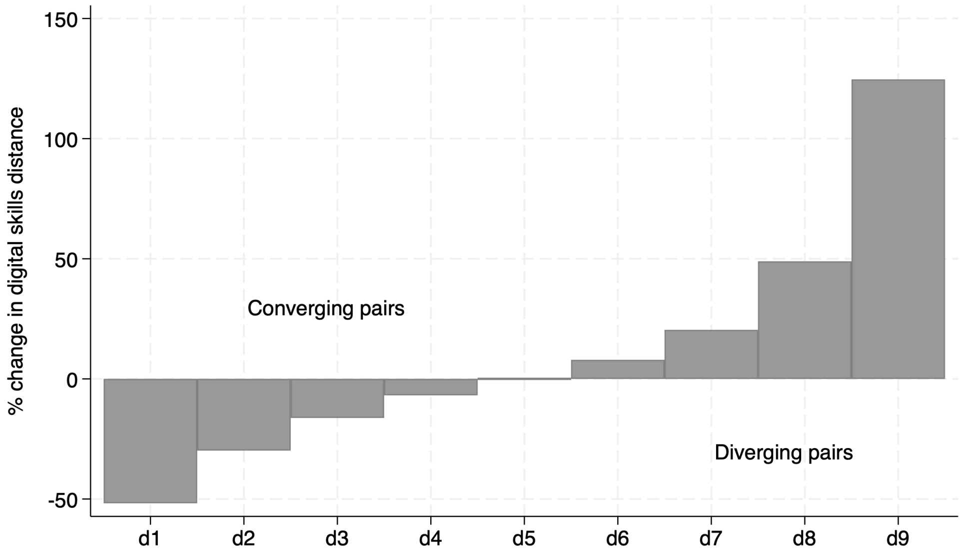

Notes: The figure ranks occupation pairs according to the percentage change in digital skill distance observed between 2011 and 2019. The digital skill distance measure is computed following equation 3.

4 Change in skill distance and occupational mobility

4.1 Measuring occupational mobility

We study occupational mobility in France using French registry data. In particular, we rely on the Payroll Tax records called *DADS poste*, which collect information on all employees active in the French labor market, including their annual salary, the total number of hours worked in the year and the occupation in which they are employed. Although the version made available by the French statistics office INSEE does not include individual worker identifiers that can be followed over time, except for a sub-sample of 1/12th of the employees, a recent contribution by Babet et al. (2022) explains how to re-construct the quasi-entirety of the worker panel using the information available in the registry.[12]

[12]In short, each dataset reports individual level information relative to the activity in the current year t as well as information relative to the year that preceded it ($t-1$). As such, the same information is reported twice for two consecutive years, once for year t and once for year $t-1$ for the wave afterward. Babet et al. (2022) show that the match on the available overlapping information is unique for 98% of the individuals, thus allowing to reconstruct a worker-level panel for the quasi-universe of the observations. The only caveat that exists is for individuals who remain out of employment for more than one year and

Given that we want to portray the entirety of the French labor market to take into account all general equilibrium effects of changes in distance between occupations, our data cleaning is restricted to the bare minimum. In particular, for workers with multiple jobs, we only consider the main one, defined as the one that paid the highest total salary over the year, and we further drop workers with incomplete occupation codes and with hourly wages below the statutory minimum wage. Finally, we only keep workers in prime age, defined between 20 years old and 60 years old. In contrast, we do not apply any sectoral or occupational restriction.

Using this reconstruction of the worker panel, we build three distinct datasets: i) baseline mobility, measuring all occupation flows observed between 2011 and 2012, ii) end-line mobility, measuring all occupation flows observed between 2018 and 2019, and iii) long-run mobility, measuring all occupation flows observed between 2011 and 2019. For all of the three datasets we further compute the average origin and destination wage relative to all occupation pairs, and we include all pairs of stayers, defined as the flows of workers that remained in the same occupation during the period, and all flows towards and from non-employment to obtain a complete picture of employment in the labor market. In total in our final sample we have 385 occupation codes plus one non-employment category, which give rise to 148,995 pairs of occupations.[13]

Appendix Table A3 reports the distribution of the mobility outcomes. In roughly 35% of the occupation pairs, we observe zero flows, which is not surprising given the wide variety of jobs included. On average, we observe 23 moves between an occupation pair in 2018-2019 and 33 moves over the entire decade. However, the averages mask substantial heterogeneity, where the pair that has the highest number of moves reaches 20 thousands at endline and 45 thousands in the long run. Figure A4 in the appendix shows a matrix of mobility patterns with all occupation pairs ranked according to the French PCS classification. Broadly speaking, the rank follows socio-professional categories, going from executives and engineers, to intermediate professionals (including mid-level management, technicians and foremen), to clerical office workers, to skilled, and finally unskilled blue collar workers. To improve readability, diagonal pairs, corresponding to stayers who remain in the same occupation over the period, have been dropped. What we can observe is that mobility patterns cluster around similar occupations, and that this

who are therefore considered a new individual once they return.

[13] We do not observe flows from non-employment to non-employment, as such transitions are not recorded in the employer–employee data. These flows are thus excluded, as they are not central given our focus on occupation-to-occupation mobility.

is especially true among higher ranked professions.

Finally, Appendix Figure A5 portrays our correlation of interest between endline and long-run mobility and the change in digital skill distance, controlling for deciles of initial digital distance. Even if this figure is purely descriptive, as it does not include any fixed effects nor bilateral control, it does suggests that changes in digital skill distance over the period did affect mobility patterns in the expected direction.

4.2 Estimation strategy

Gravity models are widely used in the trade and migration literature to analyze how distance between geographic units—typically countries—shapes bilateral flows of goods, people, or capital, net of origin- and destination-specific factors common to all pairs. Cortes and Gallipoli (2018) are the first to adapt this framework to occupational mobility within labor markets, using it to quantify the frictions associated with skill distance.

Our analysis builds on their approach but departs from it along several dimensions. First, we relax the assumption that occupational skill requirements are fixed, and instead examine how changes in skill distance over time affect mobility flows. Second, we focus specifically on digital skills, linking mobility patterns to the technological transformations that characterized the 2010s. This allows us to compare the role of evolving skill distances with that of changes in demand for workers' initial occupations.

We consider the following gravity model :

$$Mob_{ij}^{T-19} = \beta_0 + \beta_1 DigiDist_{ij}^{11} + \beta_2 \Delta DigiDist_{ij}^{11-19} + \alpha X_{ij} + \gamma_i + \gamma_j + \epsilon_{ij} \tag{4}$$

Where Mob_{ij}^{T-19} captures the total flows between origin i and destination j, including pairs where $i = j$ (stayers) and pairs where $i = 0$ or $j = 0$ (movers to and from non-employment). We estimate the model on two mobility outcomes: end-line mobility ($T = 2018$) and long-run mobility ($T = 2011$). The first has the advantage of capturing the total effect of the changing distances since all moves occur at the end of the period. The second also includes some changes that happened early in the 2010s, and thus not yet affected by the full skill distance change, but has the advantage of capturing longer career trajectories.

$DigiDist_{ij}^{11}$ represents the level of pairwise digital skill distance in 2011, while $\Delta DigiDist_{ij}^{11-19}$ is our main regressor of interest, allowing us to capture the effect of changes in digital skill distance on mobility flows. In the empirical section, both measures are standardized to have mean zero and standard deviation one to facilitate the interpretation of regression coefficients.[14] The multilateral resistance parameters γ_i and γ_j absorb all the determinants of mobility flows that are driven by origin and destination factors, thus absorbing all the effects driven by changes in labor demand and supply.[15]

X_{ij} controls for additional bilateral factors affecting occupation flows. In particular, we control for an indicator for stayers, which takes into account the fact that there are some additional fix costs associated with switching occupation that go beyond the simple effect of skill distance. In addition, we include a dummy equal to one for all occupation switches that involve a change in socioeconomic status, which are expected to be more costly than changing jobs within a given rank.[16] Most importantly, we progressively introduce controls for overall skill distance between occupations—measured as in equation 1—to verify that our effect of interest persists beyond non–technology-driven changes in occupational distances. Such changes may reflect a variety of factors, including (but not limited to) shifts in how recruiters write job advertisements, changes in the relative importance of soft skills, or evolutions in regulatory frameworks common to France and the United States. Because relying on the same Euclidean notion of distance induces a strong correlation between our measures of overall distance and digital distance, we also explore an alternative specification in which this control is constructed using a slightly modified formula.

Finally, standard errors are double-clustered at the level of origin occupation i and destination occupation j, and the model is estimated using pseudo-poisson maximum likelihood (PPML), as is standard in gravity models to avoid biases coming from the large portion of zeros in bilateral flows (Silva and Tenreyro, 2006).

[14]In the counterfactual exercises presented in section 6, we include both the level of digital skill distance and its change without standardizing them, in order to achieve a precise decomposition of the effect.

[15]In the counterfactual exercise presented in section 6 we disentangle which portion of total changes in mobility are driven by direct changes in own occupation demand versus indirect changes through evolving skill distances to other occupations.

[16]These switches might involve additional institutional constraints such as the need for a higher education diploma. In practice, we define changes of socioeconomic status using the French occupational classification, as previously described.

4.3 Estimation results

Table 1 reports the main results obtained from the estimation of equation 4. All regressions are estimated using PPML and coefficients are reported in their exponentiated form, which means that coefficients larger than 1 signal a positive effect, and coefficients smaller than 1 signal a negative effect. Panel A reports results on endline mobility (2018-19), while panel B reports results on long-run mobility (2011-19). Column (1) shows the effects of digital distance level in 2011 and digital distance change between 2011 and 2019 without controlling for overall distance. As already mentioned in the previous section, we are interested in estimating the effect of digital distance above and beyond the one of overall distance, which might capture a variety of factors. However, an issue that arises is that the two measures, when both constructed using the same Euclidean distance formula, are highly correlated (correlation coefficient of 0.2), making it difficult to disentangle the two effects. As such, we show different solutions : in column (2), we control for overall distance level and change using a weighted Euclidean formula;[17] in Column (3) we control for overall distance level using the standard euclidean formula but not its change; in Column (4) we control for overall distance change using the standard euclidean formula but not its level; in Column (5) we control for both using the standard Euclidean formula; and in Column (6) we residualize the level and change in digital distance with respect to the corresponding measures of overall distance (using the standard Euclidean formula), to capture the residual effect of digital skill distance.

When we do not control for overall skill distance, a one–standard deviation higher digital distance in 2011 is associated with a 30% reduction in endline mobility and a 27% reduction in long-run mobility. More importantly for our research question, changes in digital skill distance over time exhibit a very large effect, indicating that the 2011 level alone is insufficient to account for mobility patterns at the end of the period. Specifically, a one-standard deviation increase in digital distance between 2011 and 2019 is associated with 20% lower endline mobility and 21% lower long-run mobility. This result strongly suggests that mobility frictions across occupations are not time-invariant, particularly

[17] The weighted Euclidean formula $dw\left(z_i^{r,t}, z_j^{r,t}\right)$ applies skill weights w_s to the standard formula reported in equation 1: $dw\left(z_i^{r,t}, z_j^{r,t}\right) = \left(\sum_k w_s \left(z_{i,k}^t - z_{j,k}^t\right)^2\right)^{1/2}$, where w_s are constructed as the inverse of the skill prevalence in all other occupations. Using this measure, rare skills gain importance, thus emphasizing differences in distance driven by rare (specialized) skills rather than ubiquitous ones. The correlation coefficient between digital distance measures computed using the standard Euclidean formula and overall distance measures computed using the weighted Euclidean formula are much closer to 0.

Table 1: Effect of digital distance changes on mobility flows

	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Endline Mobility (2018-2019)					
	PPML	PPML	PPML	PPML	PPML	PPML
Digi distance 2011	0.703***	0.815***	0.921*	0.704***	0.931	
	(0.0385)	(0.0427)	(0.0433)	(0.0391)	(0.0462)	
Change digi distance 2011-19	0.818***	0.910***	0.917**	0.822***	0.954	
	(0.0281)	(0.0305)	(0.0318)	(0.0288)	(0.0369)	
Digi distance 2011 (residualized)						0.929**
						(0.0282)
Δ Digi distance 2011-19 (residualized)						0.878***
						(0.0281)
Controls :						
Distance level		✓	✓		✓	
Distance change		✓		✓	✓	
Stayers dummy	✓	✓	✓	✓	✓	✓
Level switch dummy	✓	✓	✓	✓	✓	✓
Observations	148,995	148,995	148,995	148,995	148,995	148,995
	Panel B: Longrun Mobility (2011-2019)					
	PPML	PPML	PPML	PPML	PPML	PPML
Digi distance 2011	0.731***	0.864***	0.997	0.743***	1.033	
	(0.0347)	(0.0292)	(0.0294)	(0.0367)	(0.0330)	
Δ Digi distance 2011-19	0.789***	0.882***	0.910**	0.844***	1.017	
	(0.0346)	(0.0388)	(0.0389)	(0.0361)	(0.0446)	
Digi distance 2011 (residualized)						0.986
						(0.0389)
Change digi distance 2011-19 (residualized)						0.891***
						(0.0344)
Controls :						
Distance level		✓	✓		✓	
Distance change		✓		✓	✓	
Stayers dummy	✓	✓	✓	✓	✓	✓
Level switch dummy	✓	✓	✓	✓	✓	✓
Observations	148,995	148,995	148,995	148,995	148,995	148,995

Notes: Robust standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1. Coefficients obtained from the estimation of equation (4) using PPML, reported in exponentiated form. Panel A shows the results on endline mobility (2018-2019), while panel B shows results for long-run mobility (2011-2019). Column (1) includes no controls for overall distance. Column (2) controls jointly for the level and change in overall distance using the weighted Euclidean measure. Columns (3) and (4) control, respectively, for the level of overall distance in 2011 and for changes in overall distance, both measured using the standard Euclidean formula. Column (5) controls jointly for the level and change in overall distance using the standard Euclidean formula. Finally, Column (6) residualizes the level and change in digital distance with respect to the corresponding measures of overall distance.

for skills that are most exposed to technological change.

Column (2) shows that adding controls for the level and change in overall distance—using a measure that is less prone to collinearity—reduces the coefficients by roughly half. Specifically, a one–standard deviation higher digital distance in 2011 is associated with a 19% (14%) lower endline (long-run) mobility, while a one–standard deviation increase in digital distance over the period is associated with a 9% (12%) reduction in endline (long-run) mobility. These estimates remain stable when we instead control for the level of overall distance in 2011 measured using the standard Euclidean formula, and they become larger when we control only for changes in overall distance. When both the level and change in overall distance are included using the standard Euclidean measure—which is more susceptible to multicollinearity—the magnitude of the coefficient on digital distance change on endline mobility remains similar, but it loses statistical significance. Finally, Column (6) shows that residualizing digital distance with respect to overall distance restores a statistically significant effect of digital distance change, with a magnitude comparable to that in Column (2): a one–standard deviation increase in digital distance over the period reduces both endline and long-run mobility by 11%. Overall, these results indicate that the effect of changes in digital distance is robust across alternative ways of controlling for overall distance. We therefore select the specification in Column (2) as our preferred model for the remaining of the paper.

While these results suggest that digital skill distance between occupations constitutes an important source of mobility frictions, and especially, that these frictions evolve over time in ways that shape workers' outside opportunities across occupations, it is difficult to assess their economic significance based on these regressions alone. To fully grasp the magnitude of these effects, we present counterfactual exercises that estimate how mobility from different origin occupations would have differed had digital distance remained constant over the 2010s (see Section 6). Before turning to these counterfactuals, we report additional robustness checks supporting the validity of our results and examine whether the effects of changes in digital distance are heterogeneous across demographic groups and/or asymmetric depending on the relative characteristics of occupations i and j.

4.3.1 Robustness

First, our regression model assumes linearity in the effect of digital distance level and change. We test the validity of this assumption in two ways : first, we re-estimate our unconditional model (Table 1 Col 1) and our preferred model (Table 1 Col 2) including quadratic terms of the main variables of interest, and secondly we re-estimate our preferred model including quantile dummies for the main variables of interest. Table A4 shows that the level of digital distance in 2011 exhibits some degree of nonlinearity: the marginal effect of distance declines at higher levels, consistent with a model in which mobility becomes increasingly difficult beyond a certain threshold, so that additional distance has little incremental impact. By contrast, the effect of changes in digital distance appears approximately linear, as the quadratic term is not statistically significant and close to zero in magnitude. Since changes in digital distance are our primary object of interest, this result supports the validity of our main specification. Figure A6 provides visual confirmation: although the relationship is not perfectly linear, the decline in mobility across quantiles of digital distance change can be well approximated by a linear trend.

Second, in this paper we measure skill distances using the Euclidean norm, which is both intuitive and well suited to our setting. Because we normalize skill vectors to have unit length, this choice also avoids issues arising from differences in scale across skill dimensions. However, alternative distance measures are possible, such as angular separation or the Mahalanobis distance. In addition, our baseline analysis relies on the most disaggregated skill taxonomy available in the Lightcast data (approximately 13,000 skill categories), which provides a highly detailed characterization of occupational skill profiles. However, the presence of many skills with very small weights may introduce measurement error. To assess the robustness of our results, we therefore recompute all three distance measures using a more aggregated Lightcast skill taxonomy comprising 656 skill categories.

Appendix Figure A7 reports the estimated effect of changes in digital distance on endline and long-run mobility using alternative distance measures. The coefficients for endline mobility are remarkably stable across all six specifications. While their magnitude is somewhat smaller when Euclidean and angular separation measures are computed using the aggregated skill taxonomy, declining from roughly a 10% reduction in mobility per one–standard deviation increase in digital distance to about 5%, the estimates remain highly statistically significant, and the confidence intervals largely overlap across specifications. For long-run mobility, the lower effect of aggregate skill measures is confirmed with less confidence interval overlap. Taken together, these patterns suggest that the effect of changes in digital distance are broadly robust to different measures, and that the more granular skill taxonomy captures additional meaningful variation rather than introducing noise.

Third, we assess the robustness of our results to alternative model specifications. Figure 4 displays the baseline coefficients for endline and long-run mobility on the left for reference. The second set of coefficients introduces fixed effects for all interactions between socio-professional categories, rather than a single indicator for level switches, effectively identifying the effect of digital distance from variation within a given pair of socio-professional levels. The third set adds a control for baseline mobility (2011–12), which captures pre-existing determinants of mobility flows above and beyond our distance level in 2011 and other controls, in the spirit of a panel fixed-effects specification. The fifth set of coefficients includes an additional control for relative tightness—a dummy equal to one if occupation i is tighter than occupation j—as well as its interaction with changes in digital distance. This specification addresses the concern that the estimated effect of digital distance change may be confounded by correlated changes in relative labor market tightness. Finally, the rightmost specification recomputes all distance measures under the assumption that workers do not update their skills with those of their origin occupation over time, implying positive distance even for stayers as occupations evolve. Overall, the effect of changes in digital distance is highly robust across specifications, with the exception of this obsolescence test: although the magnitude of the coefficients remain similar, we lose statistical significance. We interpret this result as evidence that workers acquire skills on the job, rendering the full obsolescence assumption unrealistic and introducing additional measurement noise.

4.3.2 Heterogeneity and asymmetry

Finally, we examine whether the impact of changes in digital distance is heterogeneous across demographic groups and/or asymmetric depending on the relative characteristics of occupations i and j. Appendix Figure A8 presents heterogeneity estimates by age (below versus above 35) and gender. To obtain these results, we construct separate mobility matrices for each mutually exclusive group and estimate the model independently

Figure 4: Robustness to different specifications

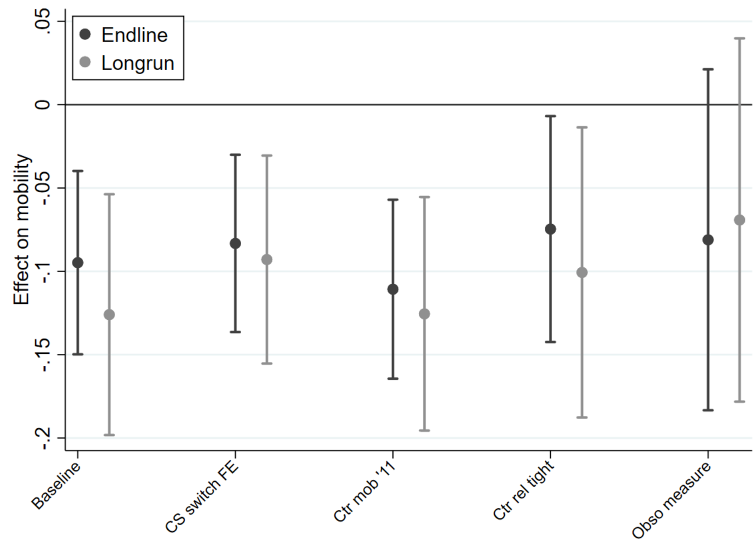

Notes: The figure reports robustness checks for the effect of changes in digital distance on endline and long-run mobility. From left to right, coefficients correspond to the baseline specification; specifications with socio-professional category interaction fixed effects; a specification controlling for baseline mobility (2011–12); a specification controlling for relative labor market tightness and its interaction with digital distance change; and a specification assuming no skill updating within occupations over time.

for each subsample.

We find little evidence of age-related heterogeneity: the effect of changes in digital distance on occupational mobility is remarkably similar across younger and older workers. In contrast, some differences emerge along the gender dimension. Women appear more sensitive than men to increases in digital distance, particularly in the long run. Importantly, this pattern cannot be explained by women's higher propensity to remain in the same occupation, as all specifications control for stayers.[18]

Table 2 examines whether the effect of changes in digital distance is asymmetric with respect to the relative characteristics of occupations i and j. For example, one might hypothesize that when two occupations become closer in digital terms, mobility increases disproportionately from the occupation that was initially more digitally intensive toward the less intensive one. Analogous asymmetries could arise when the origin occupation is initially more intensive in mathematical skills or ranked as more skilled in the French socio-professional classification.

[18] Appendix Table A5 reports summary statistics on mobility patterns by gender and age.

Table 2: Asymmetry of the effect of digital distance change

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Endline Mobility (2018-2019)				Longrun Mobility (2011-2019)			
VARIABLES	PPML	PPML	PPML	PPML	PPML	PPML	PPML	PPML
Digital distance 2011	0.815***	0.814***	0.805***	0.812***	0.864***	0.861***	0.838***	0.854***
	(0.0427)	(0.0424)	(0.0474)	(0.0425)	(0.0292)	(0.0300)	(0.0365)	(0.0329)
Change digital distance 2011-19	0.910***	0.907**	0.895**	0.923**	0.882***	0.871***	0.891**	0.893**
	(0.0305)	(0.0400)	(0.0412)	(0.0339)	(0.0388)	(0.0381)	(0.0475)	(0.0468)
Change digital distance 2011-19 x higher digi		1.006				1.028		
		(0.0694)				(0.0593)		
higher digi		0.893				0.838		
		(0.210)				(0.0952)		
Change digital distance 2011-19 x higher math			1.030				1.054	
			(0.0471)				(0.0526)	
higher math			1.220*				0.971	
			(0.141)				(0.0738)	
Change digital distance 2011-19 x higher skill				0.948				0.970
				(0.0574)				(0.0597)
higher skill				1.760				2.334***
				(1.010)				(0.593)
Observations	148,995	148,995	148,224	148,995	148,995	148,995	148,224	148,995

Notes: Robust standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1. Coefficients obtained from the estimation of equation (4) using PPML, reported in exponentiated form. Columns (1) to (4) show the results on endline mobility (2018-2019), whileColumns (5) to (8) show the results for long-run mobility (2011-2019). All the columns are based on our preferred specification (column (2) of table 1), but include a variable for the characteristics of occupation *i* relative to *i* alone and interacted with changes in digital distance.

To test these hypotheses, we include a dummy equal to one if occupation *i* is initially more digital-, math-, or skill-intensive than occupation *j*, and interact this indicator with our measure of changes in digital distance.[19] The results show no evidence of asymmetry: none of the interaction terms is statistically significant, and all estimates are close to zero. This indicates that reductions in digital distance lower transition frictions symmetrically, irrespective of the relative digital, cognitive, or skill intensity of the two occupations. In the next section, we introduce a model that allows us to generate counterfactuals and assess the economic significance of the changes in digital distance that took place over the decade of the 2010s.

[19]Initial digital intensity is constructed using our manual classification of 13,000 Lightcast skills into software-related and non-software-related categories, as already used in Appendix Figure A2. Initial mathematical intensity is measured using the O*NET occupation score for "mathematical knowledge." Skill intensity is defined using the four socio-professional ranks captured by the first digit of the French PCS classification.

5 Theoretical micro-foundation

To provide a structural interpretation of the gravity equation estimated in the previous section and to conduct counterfactual analyses, we embed occupational mobility in a two-sided one-to-one matching model with transferable utility à la Choo and Siow (2006) and Galichon and Salanié (2022b).[20] Workers and jobs are characterized by discrete types, with a large mass of agents of each type. Workers choose occupations to maximize utility, firms choose workers to maximize profits, and wages act as transfers that clear the market under competitive, price-taking behavior. In equilibrium, bilateral mobility flows reflect both match-specific productivity and amenity components, together with market-clearing constraints on both sides. This structure allows us to interpret the reduced-form gravity coefficients, separate productivity from non-productivity forces, and compute counterfactual equilibria under alternative demand and digital skill distance scenarios.

In what follows, we present the main features of the model and the key equations for identification, estimation, and computation for our counterfactual exercises succinctly. For a detailed presentation of the model and an extension, we refer the interested reader to Appendix (B.8).

Let O be the list of occupations. We say that a job in the occupation $j \in O$ is of type j. Types of workers are also defined using occupations. In particular, a worker is said to be of type $i \in O$ if the occupation of her previous job was 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 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. 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 that $\sum_{i \in O_0} Y_i \geq \sum_{j \in O} X_j$.[21]

[20] See also Galichon and Salanié (2022a) for a generalized version and Dupuy and Galichon (2022) for a continuous-type application to risky jobs.

[21] In appendix B.8 we show an extension of the model where jobs may also remain vacant. Because

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. $w_{i\emptyset} = 0$.[22] 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 a unit scaling factor.

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

$$\max_{j \in O_{\emptyset}} \left(\alpha_{ij} + w_{ij} + \varepsilon_j \right). \tag{5}$$

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 a unit scaling factor.

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

$$\max_{i \in O_0} \left(\gamma_{ij} - w_{ij} + \eta_i \right). \tag{6}$$

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

$$\log X_{ij}^S = \alpha_{ij} + w_{ij} - s_i \quad \forall (i,j) \in O_0 \times O_{\emptyset} / (0, \emptyset), \tag{7}$$

$$\log X_{ij}^D = \gamma_{ij} - w_{ij} - m_j \quad \forall (i,j) \in O_0 \times O. \tag{8}$$

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.

of the logit structure of the model, i.e., the Independence of Irrelevant Alternatives applies, excluding the possibility of vacant jobs does not affect the remaining log-odds, and the main analysis remains unchanged.

[22]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 $j \in O$.

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,

$$
\begin{aligned}
X_{ij} &= \exp\left(\frac{\varphi_{ij} - s_i - m_j}{2}\right), \forall\,(i,j) \in O_0 \times O && (9)\\
X_{i\varnothing} &= \exp\left(\alpha_{i\varnothing} - s_i\right), \forall i \in O && (10)\\
w_{ij} &= \frac{1}{2}\left(\gamma_{ij} - \alpha_{ij} + s_i - m_j\right), \forall i \in O_0, j \in O. && (11)
\end{aligned}
$$

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

$$
\begin{aligned}
\sum_{j \in O_\varnothing} X_{ij} &= Y_i,\ \forall i \in O_0, && (12)\\
\text{and} \sum_{i \in O_0} X_{ij} &= X_j, \forall j \in O, && (13)
\end{aligned}
$$

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

Three important remarks are in order. First, clearly, this equilibrium 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. Note that while both productivity (γ_{ij}) and amenities (α_{ij}) increase the equilibrium mass of workers of type i matched to jobs of type j, equilibrium transfers increase with the former but decrease with the latter. This is the source of the separate identification of the productivity and amenities components in this matching model. We exploit this identification strategy in our counterfactual exercises.

Second, an important advantage of the micro-foundation of the gravity equation through our matching model is that it provides us with an algorithm to compute the equilibrium associated with counterfactuals of interest. Indeed, note that given parameters γ_{ij} and α_{ij}, the constraints (12-13) form a system of equations that can be solved for s_i and m_j using equations (9-10), see Appendix (B.8).

Finally, estimation of the parameters of interest can be done using similar tools as the ones advocated for the estimation of a gravity equation. To see this, 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. As recently shown in Galichon and Salanié (2022b), 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 of 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 parameterization, now read as

$$w_{ij} = \gamma_{ij}^b - \frac{1}{2}\varphi_{ij}^\beta + \frac{1}{2}\left(s_i - m_j\right), \forall i \in O_0, j \in O. \tag{14}$$

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

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

where w_{ij}^o 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

$\hat{a}_k = \hat{\beta}_k - \hat{b}_k$.[23]

Appendix Figure A9 evaluates the fit of the model. By construction, the model reproduces the total number of flows in the economy, but not necessarily the correct bilateral flows. We see that we tend to slightly over-predict mobility at lower levels and under-predict mobility at higher levels, but that the linear fit is close to the 45-degree line. The correlation coefficient is about 0.5.

6 Counterfactual Analysis

Section 4 shows that changes in digital skill distance are systematically associated with occupational mobility, with more mobility between occupation pairs that became more similar and less mobility between pairs that became more distant. The key question is how large these distance-induced mobility effects are in economic terms. We address this question by benchmarking the mobility generated by changes in digital skill distance against the mobility driven by shifts in occupational labor demand, which has been the dominant focus of the technological change literature. In this context, the demand channel captures the reallocation of labor away from occupations adversely affected by technological change—and therefore shrinking—toward occupations that are complementary to new technologies and consequently expanding.

To quantify the contribution of each channel, we conduct a set of counterfactual simulations within the model. First, we simulate occupational mobility holding digital skill distances fixed at their 2011 level while allowing occupational demand to evolve as observed. This counterfactual captures how mobility would have changed if technological change had affected labor demand but not the digital distance between occupations. Second, we simulate mobility holding the relative size of all occupations constant over the period, while allowing digital skill distances to evolve.[24] This scenario isolates the

[23] Appendix B.9 presents an extension in which we also incorporate information on wages at t.

[24] In this counterfactual, we impose that the distribution of employment across occupations remains constant between 2011 and 2019, while allowing the total number of jobs to evolve as observed. This choice reflects the fact that changes in overall employment are driven by many factors orthogonal to technological change (e.g., demographics). By contrast, we attribute all changes in relative occupational sizes to technological change, which may overstate its role. Under this assumption, the contribution of changes in digital skill distance to the overall effect of technological change should therefore be interpreted as a lower bound.

mobility effects driven by changes in digital skill distance alone. Finally, we simulate mobility holding both occupational demand and digital skill distances fixed at their initial levels. This counterfactual recovers the level of occupational mobility that would have prevailed in the absence of technological change and corresponds to the "natural rate of mobility" defined in the framework section.

Comparing these simulations to observed mobility flows allows us to quantify the share of mobility attributable to technological change and to disentangle the respective roles of demand shifts and changes in digital skill distance.[25] Because the ordering of the simulations matters—shutting down demand before distance yields slightly different contributions than reversing the order—we apply a Shapley decomposition, which averages the effects obtained under both sequences. Finally, as detailed in Section 5, we exploit wage information across occupations to further decompose the effect of changes in digital distance into a productivity and a non-productivity channel. While the productivity channel raises both mobility and wages, the non-productivity channel (often interpreted as amenities) increases mobility while lowering wages.

Figure 5 shows that, as expected, demand shifts over the 2010s reduced outflows from executive and professional occupations while increasing outflows from lower-skilled occupations—most notably clerks, which are commonly identified as the most routine-intensive group. More strikingly, changes in digital distance operate in the opposite direction, substantially reducing outflows from clerical and blue-collar occupations. This pattern is potentially concerning, as it suggests that increases in digital skill distance raise additional frictions precisely in the occupations facing adverse demand shocks that require relocation. Appendix Figure A10 confirms this intuition by documenting a generally negative correlation between the demand-driven and distance-driven effects on outflows.

Appendix Table A6 reports the magnitudes underlying Figure 5. Although both demand and distance effects are small relative to the natural rate of mobility, reflecting substantial labor market churning and the fact that many occupational transitions are unrelated to technological change, their relative magnitudes are nevertheless economically meaning-

[25]In all counterfactuals, we allow overall skill distance to evolve as observed by continuing to include it as a control in the model rather than setting it to zero. We do so because changes in overall skill distance are likely driven by multiple factors that may be orthogonal to technological change. In this section, all counterfactuals are based on endline mobility, results obtained on long-run mobility are very similar in magnitude and available upon request.

Figure 5: Demand and distance effect on outflows

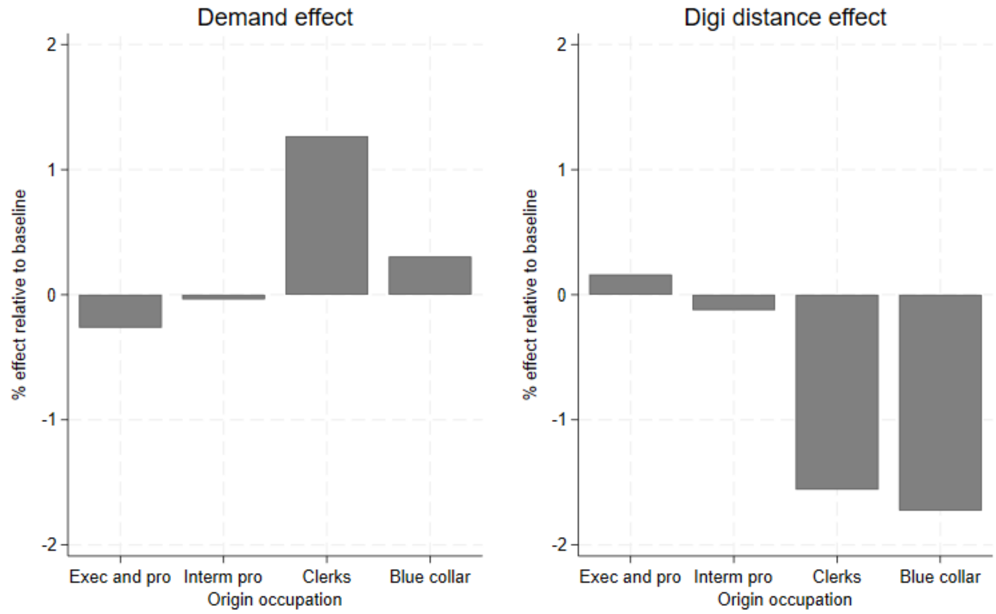

The figure shows the predicted effects of demand and distance changes relative to the baseline "natural rate of mobility".

ful. Changes in labor demand reduce outflows from executive and professional positions by 0.3%, increase outflows from blue-collar occupations by 0.3%, and raise outflows from clerical jobs by 1.3%, with little effect on intermediate professions. By contrast, changes in digital skill distance increase outflows from executive and professional occupations by 0.2%, reduce outflows from intermediate professions by 0.1%, and substantially decrease outflows from clerical and blue-collar jobs, by 1.6% and 1.7%, respectively. Overall, the indirect effects associated with changes in digital distance are comparable in magnitude to own-occupation demand shocks and, for lower-skilled occupations, they reduce mobility by considerably more than the increase in mobility induced by declining demand, highlighting the importance of taking these effects into account.

Appendix Figure A11 plots the digital distance effect, distinguishing between converging and diverging occupation pairs and across occupational categories. We estimate that, in the absence of changes in digital skill distance, mobility between converging occupation pairs would have been 4.5 percent lower for executive and professional occupations, 3 percent lower for intermediate professions, 2.5 percent lower for clerical occupations, and 2 percent lower for blue-collar occupations. Conversely, mobility between diverging occupation pairs would have been 4 percent higher for executive and professional occupations, 4.5 percent higher for intermediate professions, 4.5 percent higher for clerical occupations, and 6 percent higher for blue-collar occupations. Further highlighting the non-negligible magnitude of the digital distance change effect.

Table 3: Counterfactual results by growing versus shrinking origin occupations

		Demand effect		Distance effect	
	baseline outflows	absolute	relative to baseline	absolute	relative to baseline
Panel A : Outflows from shrinking occupations					
Executives	224149	9563	4.3%	-1363	-0.6%
Professionals	393260	11636	3.0%	-729	-0.2%
Clerks	710127	25286	3.6%	-10405	-1.5%
Blue collar	593108	14387	2.4%	-9729	-1.6%
Total	1920644	60871	3.2%	-22226	-1.2%
Panel B : Outflows from growing occupations					
Executives	490759	-11442	-2.3%	2518	0.5%
Professionals	394175	-11922	-3.0%	-236	-0.1%
Clerks	281356	-12714	-4.5%	-5040	-1.8%
Blue collar	337269	-11545	-3.4%	-6331	-1.9%
Total	1503559	-47623	-3.2%	-9089	-0.6%

Notes: The table summarizes estimated "natural rate of out-mobility", which is referred to as baseline. It then summarizes the effect on outflows generated by changes in demand and by changes in digital distance, separately for each socio-demographic category and for occupations that are growing in relative size between 2011 and 2019, and occupations that are shrinking in relative size between 2011 and 2019.

We further distinguish occupations according to whether they expand or contract in relative size over the 2010s. Table 3 reports simulation results for these two groups across socio-professional categories. As expected, demand shocks increase outflows from shrinking occupations and reduce outflows from growing occupations. By contrast, changes in digital skill distance decrease outflows in nearly all cases. As a result, distance amplifies the reduction in outflows for growing occupations and dampens the increase in outflows for shrinking ones. This pattern indicates that workers in occupations experiencing declining demand due to technological change face tighter reallocation constraints, as increased frictions limit their ability to transition to alternative occupations.

In the short run, this does not imply that low-skilled workers are unambiguously worse off. Appendix Table A7 shows that most of the demand-induced increase in outflows for low-skill workers is directed toward lower-paying jobs. Increased distance thus disproportionately reduces these downward transitions. In the longer run, however, increased frictions following negative demand shocks may translate into a higher risk of unemployment, as relocation to alternative occupations becomes more difficult. We do not show simulated flows into and out of non-employment, as digital skill distance is not meaningfully defined for these transitions in our framework.[26] Appendix Table A7 also shows that, for executives and professionals, the distance effect is clearly favorable, as it raises transitions to higher-paying jobs while reducing moves to lower-paying ones.

Figure 6: Distance effect by channel

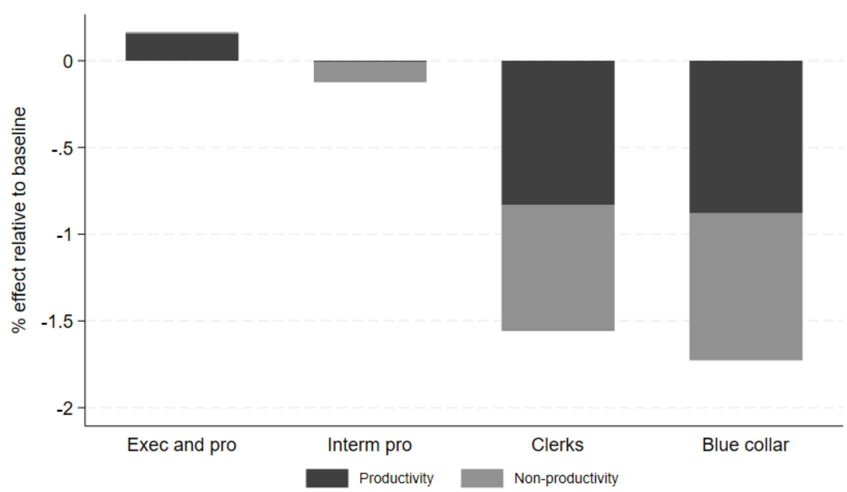

Notes: The figure decomposes the effect of digital distance changes between outflows driven by productivity gains (i.e. wage increases) and outflows driven by non-productivity gains (i.e. wage decreases).

Finally, Figure 6 reports the decomposition of the digital distance effect on outflows between productivity driven flows and non-productivity (or amenity) driven flows. We

[26]In our data we have occupation pairs involving cells of non-employment, because we want to correctly estimate the multilateral resistance parameters for the entire economy. However, we have to assign an arbitrary distance level to these pairs (we set it to zero), which does not vary over time. As such, it does not make sense to look at simulations of what would have happened to these transitions had digital skill distance not changed over the period.

can see that, while virtually all additional outflows from executives and professionals are driven by higher wage opportunities, roughly half of outflows observed among clerks and blue collar are driven by better non-wage conditions.

7 Conclusion

This paper shows that technological change affects workers' career trajectories not only by reshaping occupational demand, but also by altering the skill distances that structure mobility across jobs. Focusing on the digitalization wave of the 2010s, we document substantial heterogeneity in how occupations converged toward or diverged from one another in digital skill space, despite the widespread diffusion of digital tasks. These changes in digital skill distance translate into meaningful differences in occupational mobility, with workers moving more easily between occupations that became more similar and facing higher frictions when distances increased.

By embedding these empirical patterns in a micro-founded matching framework, we show that distance-induced mobility frictions are quantitatively important and comparable in magnitude to demand-driven forces emphasized in the existing literature. Crucially, changes in digital skill distance often operate in the opposite direction of demand shifts: while demand shocks push workers out of routine-intensive and declining occupations, rising skill distances reduce their ability to reallocate. This mechanism slows adjustment and disproportionately affects lower-skilled occupations, while facilitating productivity-enhancing mobility for higher-skilled workers.

More broadly, our findings suggest that studies of the labor market impacts of technological change may understate adjustment costs if they abstract from how technology reshapes occupational mobility frictions. By treating skill distances as fixed, much of the existing literature implicitly assumes that workers' adjustment options remain unchanged, even as technologies transform the task content of jobs. Our results highlight that technological change alters not only where labor demand shifts, but also how easily workers can respond to these shifts. Accounting for the endogenous evolution of mobility frictions may therefore be essential for understanding the full distributional and welfare consequences of technological change.

Taken together, our findings suggest that policies aimed at easing labor market adjustment to technological change should not only address displaced labor demand but also target the evolving skill distances between occupations, for instance through training and credentialing systems that reduce mobility frictions in the skill space.

References

Acemoglu, D. and D. Autor (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics*, Volume 4, pp. 1043–1171. Elsevier.

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.

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

Althoff, L. and H. Reichardt (2025). Task-specific technical change and comparative advantage.

Atalay, E., P. Phongthiengtham, S. Sotelo, and D. Tannenbaum (2020). The evolution of work in the united states. *American Economic Journal: Applied Economics 12*(2), 1–34.

Autor, D. and N. Thompson (2025). Expertise. *Journal of the European Economic Association*, jvaf023.

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

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

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.

Berman, E., J. Bound, and S. Machin (1998). Implications of skill-biased technological change: international evidence. *The quarterly journal of economics 113*(4), 1245–1279.

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.

Bloom, D. E., K. Prettner, J. Saadaoui, and M. Veruete (2025). Artificial intelligence and the skill premium. *Finance Research Letters*, 107401.

Bloom, N., T. A. Hassan, A. Kalyani, J. Lerner, and A. Tahoun (2021). The diffusion of disruptive technologies. Technical report, National Bureau of Economic Research.

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.

Braxton, J. C. and B. Taska (2023). Technological change and the consequences of job loss. *American Economic Review 113*(2), 279–316.

Card, D. and J. E. DiNardo (2002). Skill-biased technological change and rising wage inequality: Some problems and puzzles. *Journal of labor economics 20*(4), 733–783.

Carnevale, A. P., T. Jayasundera, and D. Repnikov (2014). Understanding online job ads data. Technical report, A technical report. MS o. PP Center on Education and the Workforce.

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. and L. B. Kahn (2018). Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics 36*(S1), S337–S369.

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

Dillender, M. and E. Forsythe (2022). Computerization of white collar jobs. Technical report, National Bureau of Economic Research.

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.

Eloundou, T., S. Manning, P. Mishkin, and D. Rock (2024). Gpts are gpts: Labor market impact potential of llms. *Science 384*(6702), 1306–1308.

Galichon, A. and B. Salanié (2022a). 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 (2022b). Estimating separable matching models. *arXiv preprint arXiv:2204.00362*.

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.

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

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

Hampole, M., D. Papanikolaou, L. D. Schmidt, and B. Seegmiller (2025). Artificial intelligence and the labor market. Technical report, National Bureau of Economic Research.

Katz, L. F. and K. M. Murphy (1992). Changes in relative wages, 1963–1987: supply and demand factors. *The quarterly journal of economics 107*(1), 35–78.

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.

Lipowski, C., A. Salomons, and U. Zierahn (2024). Expertise at work: New technologies, new skills, and worker impacts. *ZEW-Centre for European Economic Research Discussion Paper* (24-044).

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.

Modestino, A. S., D. Shoag, and J. Ballance (2020). Upskilling: Do employers demand greater skill when workers are plentiful? *Review of Economics and Statistics 102*(4), 793–805.

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.

Silva, J. S. and S. Tenreyro (2006). The log of gravity. *The Review of Economics and statistics 88*(4), 641–658.

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

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.

Appendix

A Additional Tables and Figures

A.1 figures

Figure A1: Percentage change in skill distance relative to reference IT jobs by broad occupation category

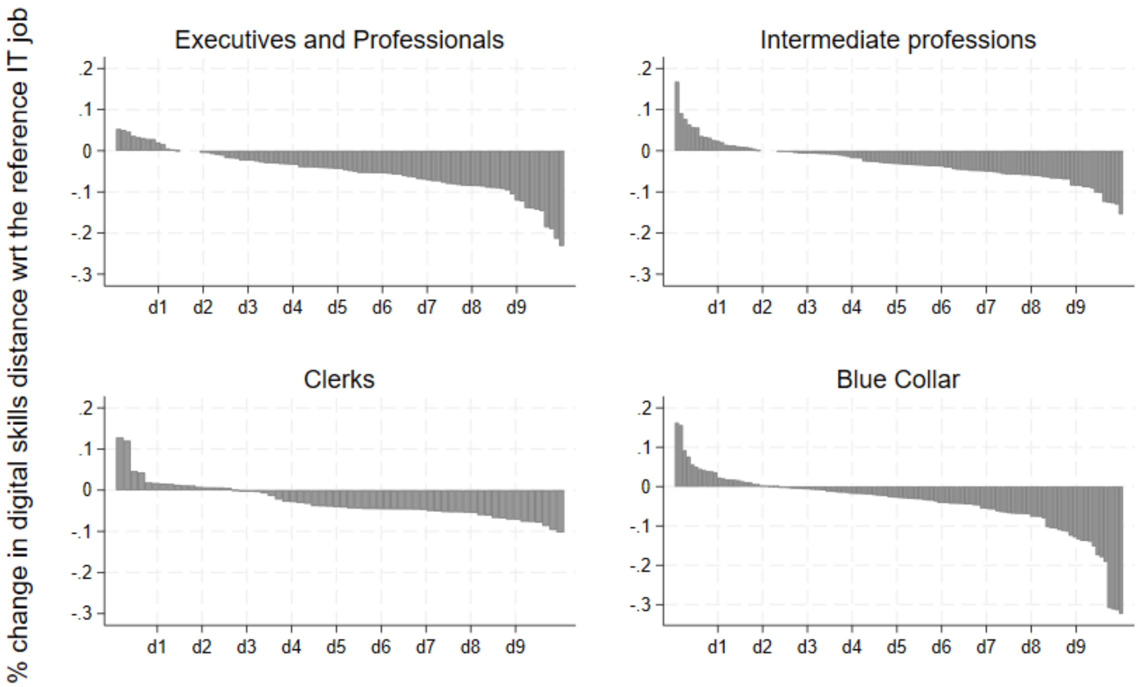

Notes: The figure ranks occupations according to the percentage change in distance relative to the reference category of IT jobs observed between 2011 and 2019, separately for each broad socio-professional category of the French occupational classification. The skill distance measure is computed following equation 2.

Figure A2: Validation of our digitalization measure

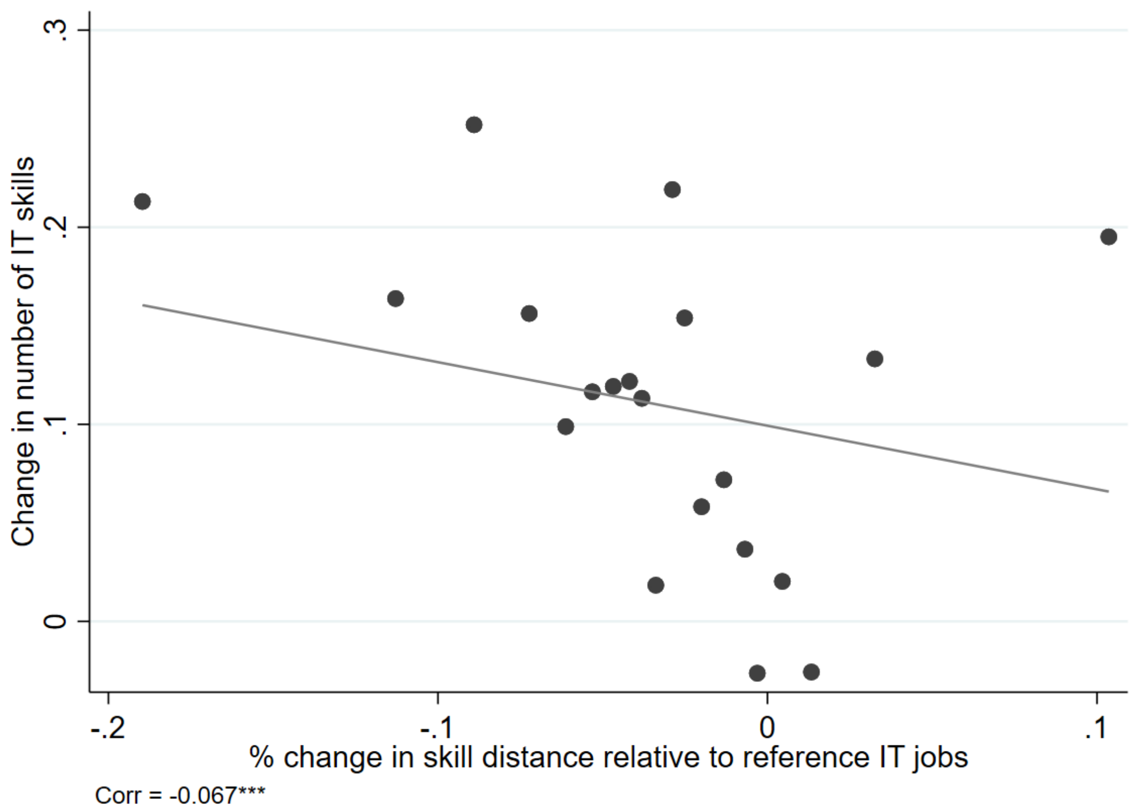

Notes: Binned scatter plot correlating the change in skill distance relative to reference IT jobs to the change in the average number of required IT skills mentioned in the ads of the occupation, where IT skills correspond to our manual classification of the 13,000 skills collected by Lightcast into Software related vs not. Some examples of IT skills include "MS PowerPoint", "MS Excel", "Python", and "GitHub".

Figure A3: Percentage change in digital skill distance between all occupation pairs by broad occupation category

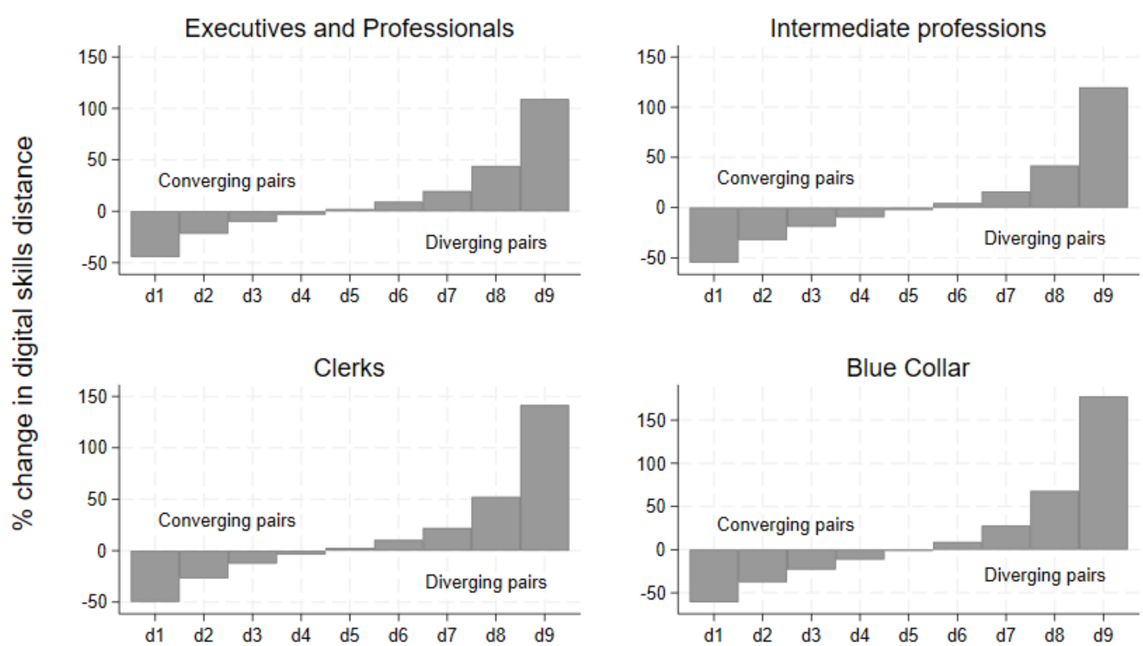

Notes: The figure ranks occupation pairs according to the percentage change in digital skill distance observed between 2011 and 2019, separately for each broad socio-professional category of the French occupational classification. The digital skill distance measure is computed following equation 3.

Figure A4: Matrix of bilateral flows from 2018 to 2019

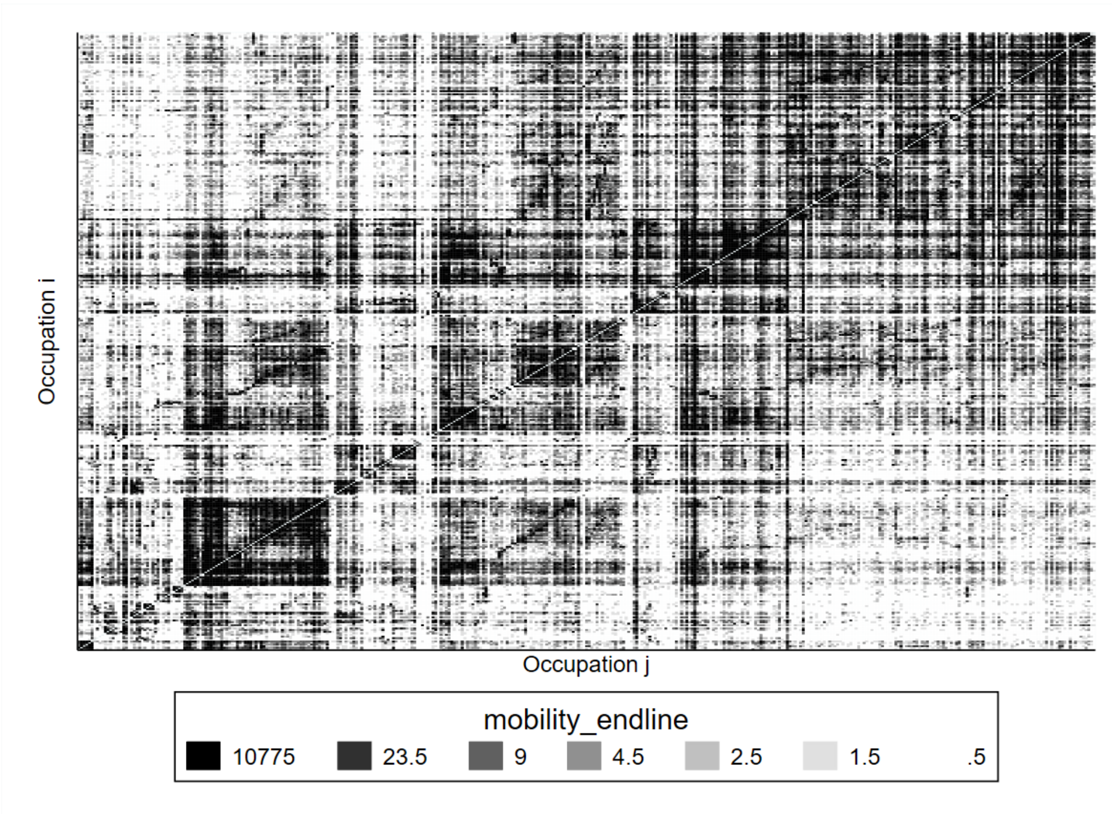

Notes: The figure shows the matrix of mobility flows between 2018 and 2019, where darker areas indicate higher flows. To increase readability the stayers have been dropped (flows within the same occupation along the diagonal) as well as flows to and out of non-employment. Occupations are ranked following the French PCS classification, which goes from executives and engineers, to technicians and intermediate professionals, to clerical workers, to skilled and finally unskilled blue collar workers. The ordering thus corresponds coarsely to socio-economic status.

Figure A5: Correlation between distance changes and mobility

(a) Endline mobility

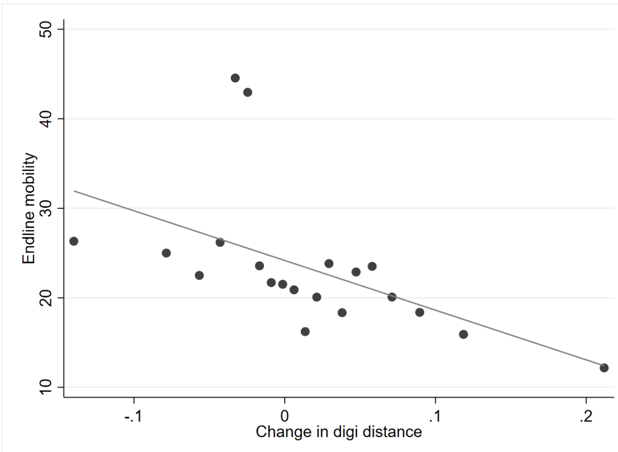

(b) Long-run mobility

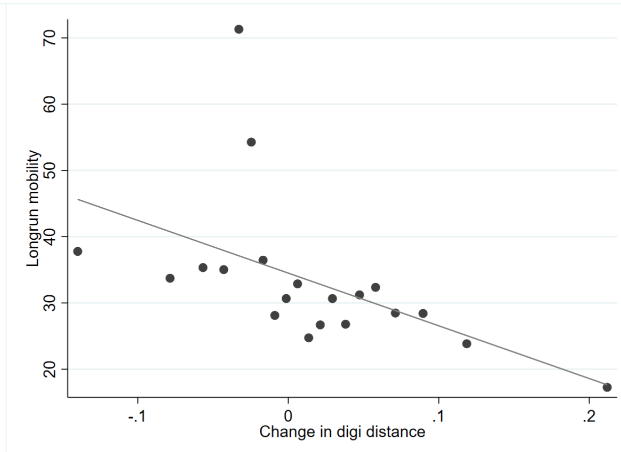

The figure shows the correlation between the change in digital distance observed between 2011 and 2019 and occupational endline mobility (panel a) and long-run mobility (panel b), controlling for deciles of initial digital distance in 2011.

Figure A6: Linearity of the effect of digital skill distance

(a) Digi distance level

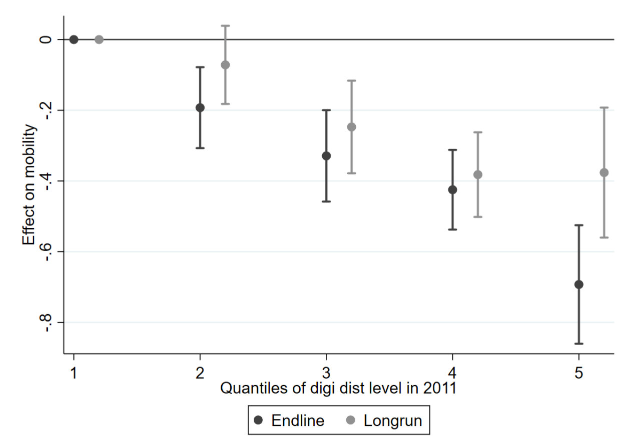

(b) Digi distance change

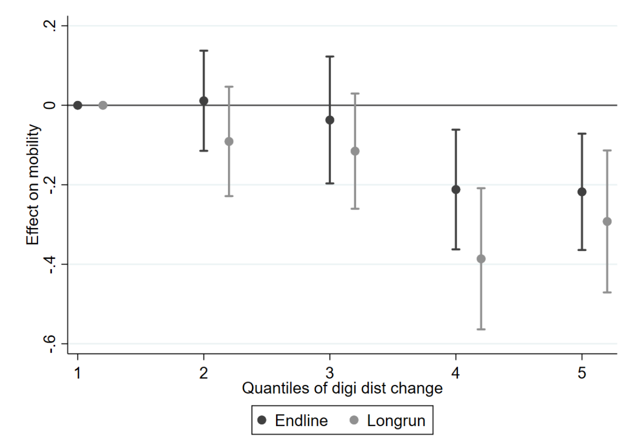

Notes: The figure shows the coefficients obtained from estimating equation 4 on 5 quantiles of digital skill distance level (panel a), and digital skill distance change (panel b). The first quantile serves as reference category. All the other controls are included in the regressions. The same figure reports the results for both endline and longrun mobility.

Figure A7: Robustness to different measures of digital distance

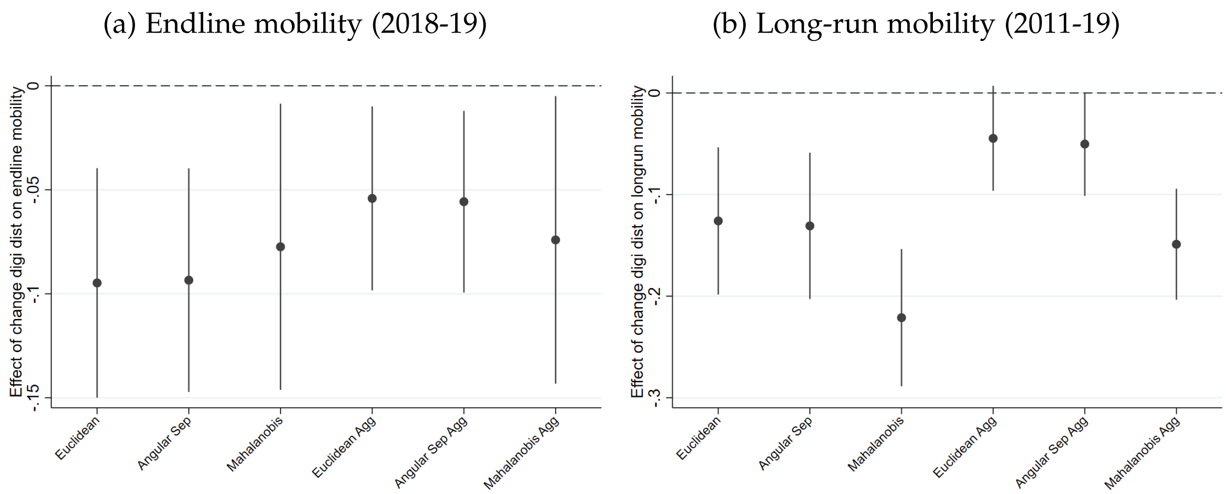

Notes: The figure shows the coefficients obtained from estimating equation 4 using different measures of digital skill distance. Panel A shows the effect of digital distance change on endline mobility, Panel B shows the effect of digital distance change on long-run mobility.

Figure A8: Heterogeneity of the effect across age and gender

(a) Distance changes by age

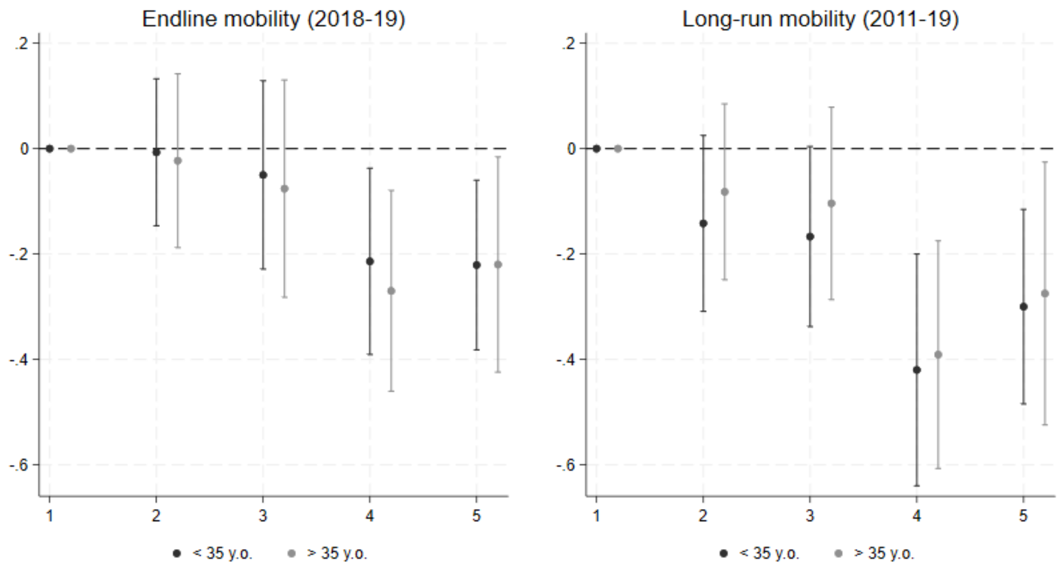

(b) Distance changes by gender

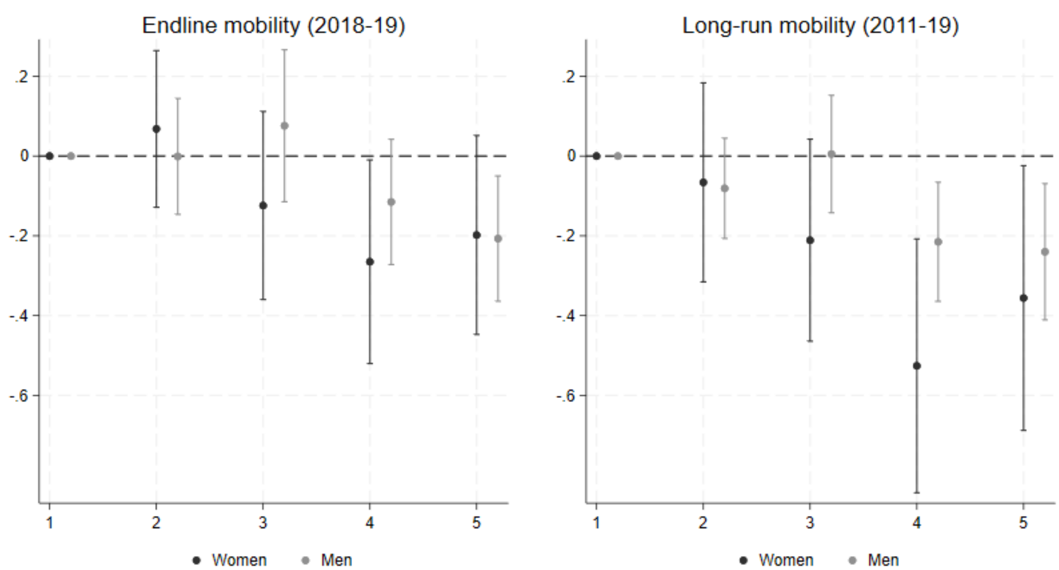

The figure shows the coefficients on quantiles of digital distance changes estimated on endline and long-run mobility obtained from two separate datasets. Panel a) constructs two separate mobility matrices by age, where the split is done at 35 years old. Panel b) constructs two separate mobility matrices by gender.

Figure A9: Evaluating the fit of the model

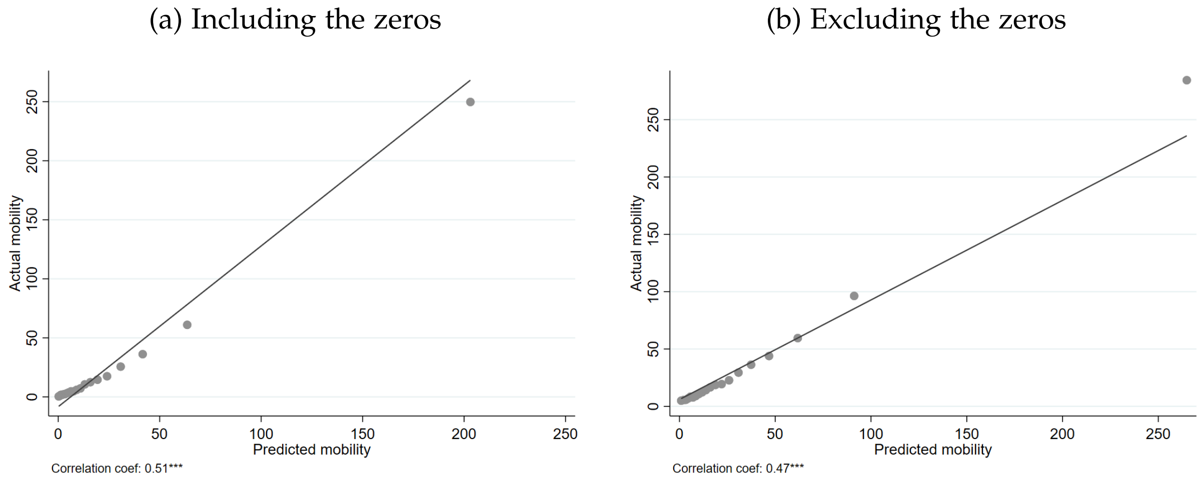

The figure shows the correlation between the observed bilateral flows and the total bilateral flows predicted by our model. Panel a) includes all apirs, including those with zero flows, panel b) excludes pairs with zero flows.

Figure A10: Correlation between the demand and distance effect on outflows for each origin occupation

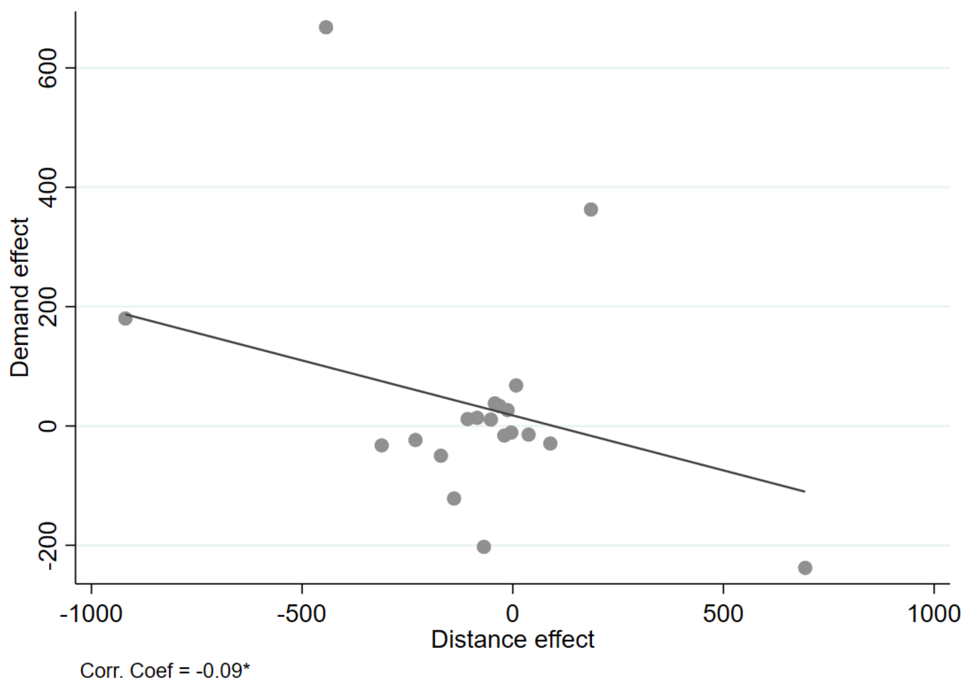

Notes: This figure shows the correlation between the demand and distance effect of technological change on occuational outflows. The correlation coefficient is -0.09, significant at the 10% level.

Figure A11: Distance effects for converging and diverging pairs

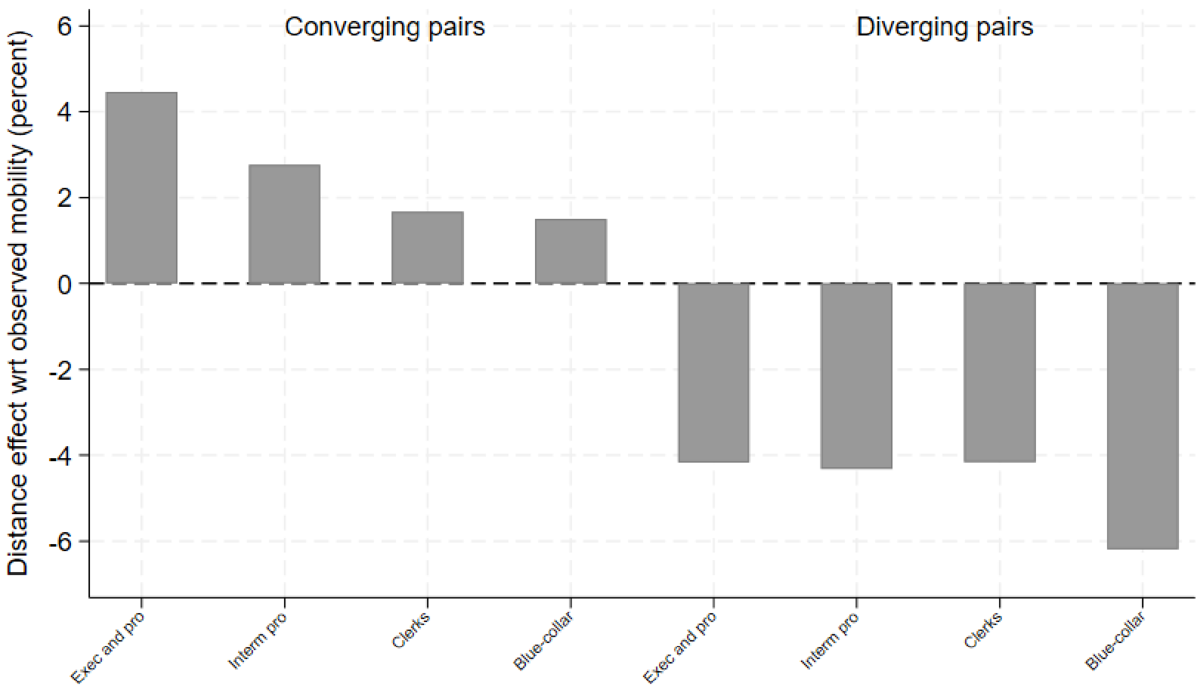

Notes: The figure shows the effect of digital distance changes separately for occupation pairs getting closer and occupation pairs getting further away from each other.

A.2 tables

Table A1: 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 A2: List of occupations that had the largest convergence in % terms towards reference IT occupations

PCS code	Occupation	% change in distance to IT jobs
633d	Maintenance electricians and electronics technicians (non-industrial equipment)	-31%
628b	Industrial maintenance electricians and electromechanics	-30%
633b	Consumer electronics repair technicians	-30%
633c	Automotive electrical and electronics maintenance technicians	-28%
332b	Local government and hospital engineers	-23%
463b	Technical sales representatives (capital and intermediate goods, excl. IT)	-21%
233c	Owners/managers of small commercial firms (10–49 employees)	-18%
386a	R&D engineers and managers in miscellaneous industries	-17%
377a	Hospitality and food service managers	-17%
386e	Manufacturing engineers and managers in miscellaneous industries	-17%

Notes: The table lists the fifteen occupations that had the largest decrease in distance relative to reference IT occupations.

Table A3: Summary statistics on occupation mobility

	Mobility flows			Mobility flows restricted to o-o switchers		
	Baseline (2011-12)	Endline (2018-19)	Longrun (2011-19)	Baseline (2011-12)	Endline (2018-19)	Longrun (2011-19)
Mean	153.9	172.5	222.8	26.0	23.3	33.3
SD	3701.9	4072.3	4293.6	217.1	189.4	306.2
Min	0	0	0	0	0	0
p25	0	0	0	0	0	0
p50	2	1	2	2	1	2
p75	10	7	12	9	7	11
Max	515860	538008	413237	37778	21515	44865
N. obs		148’995			147’840	

Notes: The table summarizes the main outcomes of interest for the analysis. On the left, statistics are obtained from all pairs of occupations, including pairs where the occupations of origin and destination are the same (stayers) and where one of the two pairs is non-employment. On the right, statistics are obtained from a subset of pairs that involve a switch in occupation, thus excluding stayers and movers from/to non-employment.

Table A4: Testing the linearity assumption

	(1)	(2)	(3)	(4)
	Endline Mobility (2018-2019)		**Longrun Mobility (2011-2019)**	
	PPML	PPML	PPML	PPML
Digi distance 2011	0.639***	0.639***	0.751***	0.751***
	(0.0404)	(0.0404)	(0.0475)	(0.0475)
Digi distance 2011 squared	1.141***	1.141***	1.073***	1.073***
	(0.0234)	(0.0234)	(0.0219)	(0.0219)
Δ Digi distance 2011-19	0.894***	0.894***	0.882***	0.882***
	(0.0287)	(0.0287)	(0.0356)	(0.0356)
Δ Digi distance 2011-19 squared	0.994	0.994	1.023	1.023
	(0.0236)	(0.0236)	(0.0205)	(0.0205)
Controls :				
Distance level		✓		✓
Distance change		✓		✓
Stayers dummy	✓	✓	✓	✓
Level switch dummy	✓	✓	✓	✓
Observations	148,995	148,995	148,995	148,995

Notes: Robust standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1. Coefficients obtained from the estimation of equation (4) using PPML, reported in exponentiated form. In addition, the model adds quadratic terms for the two main regressors of interest. Columns (1) and (2) show the results on endline mobility (2018-2019), while Columns (3) and (4) show the results for long-run mobility (2011-2019). Columns (1) and (3) includes no controls for overall distance. Columns (2) and (4) control jointly for the level and change in overall distance using the weighted Euclidean measure.

Table A5: Summary statistics on occupation mobility by demographics

	All pairs				Pairs occ i different occ j			
	mean	sd	min	max	mean	sd	min	max
Mobility endline young	68.6	1475	0	183937	12.3	101	0	11356
Mobility endline old	103.9	2771	0	440940	11.0	98	0	15377
Mobility long run young	98.4	2079	0	255477	16.8	135	0	14897
Mobility long run old	124.5	2452	0	243788	16.5	185	0	34949
Mobility endline female	81.3	2588	0	342826	10.0	119	0	17502
Mobility endline male	91.2	2084	0	332860	13.3	103	0	10509
Mobility long run female	105.0	2688	0	297724	14.6	223	0	35412
Mobility long run male	117.8	2191	0	191282	18.7	134	0	19007

Notes: The table summarizes the main outcomes of interest for the analysis separately for men and women and for younger workers (less than 35 years old) and older workers (more than 35 years old). On the left, statistics are obtained from all pairs of occupations, including pairs where the occupations of origin and destination are the same (stayers) and where one of the two pairs is non-employment. On the right, statistics are obtained from a subset of pairs that involve a switch in occupation, thus excluding stayers and movers from/to non-employment.

Table A6: Main counterfactual results

Mobility to other occupations					
starting occ.	Baseline	Demand		Distance	
		Abs	rel	Abs	rel
Executives	714908	-1879	-0.3%	1155	0.2%
Professionals	787435	-286	0.0%	-965	-0.1%
Clerks	991483	12572	1.3%	-15445	-1.6%
Blue collar	930377	2841	0.3%	-16060	-1.7%

Notes: The table summarizes estimated "natural rate of out-mobility", which is referred to as baseline. It then summarizes the effect on outflows generated by changes in demand and by changes in digital distance, separately for each socio-demographic category.

Table A7: Counterfactual results by movers towards higher paying versus lower paying occupations

	baseline outflows	Demand effect: absolute	Demand effect: relative to baseline	Distance effect: absolute	Distance effect: relative to baseline
Panel A : Outflows to worse paying jobs					
Executives	516609	-543	-0.1%	-860	-0.2%
Professionals	450946	632	0.1%	-1283	-0.3%
Clerks	354861	8590	2.4%	-5781	-1.6%
Blue collar	375429	7113	1.9%	-7770	-2.1%
Total	1697845	15793	0.9%	-15694	-0.9%
Panel B : Outflows to better paying jobs					
Executives	198299	-1336	-0.7%	2015	1.0%
Professionals	336489	-918	-0.3%	318	0.1%
Clerks	636621	3982	0.6%	-9664	-1.5%
Blue collar	554948	-4272	-0.8%	-8290	-1.5%
Total	1726358	-2544	-0.1%	-15620	-0.9%

Notes: The table summarizes estimated "natural rate of out-mobility", which is referred to as baseline. It then summarizes the effect on outflows generated by changes in demand and by changes in digital distance, separately for each socio-demographic category and for mobility flows towards better paying jobs and mobility flows towards worse paying jobs.

B Matching model details and extensions

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 (2022b) to model this market.[27] 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

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

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.

B.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 $\sum_{i \in O_0} Y_i \geq \sum_{j \in O} X_j$.[28]

B.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

[28] In appendix B.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.

$\left\{(X_{i\varnothing})_{i\in O_0, j\in O}, (X_{ij})_{i\in O_0, j\in O}\right\}$ satisfying the accounting constraints[29]

$$X_{i\varnothing} + \sum_{j\in O} X_{ij} = Y_i, \ \forall i \in O_0, \tag{15}$$

$$\text{and} \sum_{i\in O_0} X_{ij} = X_j, \forall j \in O. \tag{16}$$

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.[30]

B.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 $j \in O_\varnothing = O \cup \{\varnothing\}$ and w_{ij} be the monetary transfer, typically the wage, paid in jobs of type $j \in O_\varnothing$ for workers of type $i \in O_0$, with the convention that when workers do not work, i.e. $j = \varnothing$, they receive no transfers, i.e. $w_{i\varnothing} = 0$.[31] Further, let ε_j be a worker-specific, idiosyncratic taste for occupation $j \in O_\varnothing$, drawn from a (centered) Gumbel type I distribution with unit scaling factor.

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

[29]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 B.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.

[30]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, $\sum_i X_{i\varnothing} = X_\varnothing$ 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 $X_{0\varnothing} = 0$.

[31]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 $j \in O$.

i.e. solves the problem

$$\max_{j\in O_\varnothing}\left(\alpha_{ij}+w_{ij}+\varepsilon_j\right). \tag{17}$$

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

$$\max_{i\in O_0}\left(\gamma_{ij}-w_{ij}+\eta_i\right). \tag{18}$$

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

$$\log X_{ij}^S = \alpha_{ij}+w_{ij}-s_i \quad \forall\,(i,j)\in O_0\times O_\varnothing/\,(0,\varnothing)\,, \tag{19}$$
$$\log X_{ij}^D = \gamma_{ij}-w_{ij}-m_j \quad \forall\,(i,j)\in O_0\times O. \tag{20}$$

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.

B.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 (19) and (20),

$$X_{ij} = \exp\left(\frac{\varphi_{ij}-s_i-m_j}{2}\right),\forall\,(i,j)\in O_0\times O \tag{21}$$
$$X_{i\varnothing} = \exp\left(\alpha_{i\varnothing}-s_i\right),\forall i\in O \tag{22}$$

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

$$\begin{aligned} \sum_{j \in O_\varnothing} X_{ij} &= Y_i, \forall i \in O_0, & (23) \\ \text{and} \sum_{i \in O_0} X_{ij} &= X_j, \forall j \in O, & (24) \end{aligned}$$

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 (20) for w_{ij} and substituting the equilibrium expression of X_{ij} in equation (9) for X_{ij}^D, one obtains equilibrium outcome as[32]

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

B.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-

[32]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).

B.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 9 and 10) into the accounting constraints (equations 15 and 16), to obtain[33]

$$\begin{aligned} \exp\left(\alpha_{i\emptyset}-s_i\right)+\sum_{j\in O}\exp\left(\frac{\varphi_{ij}-s_i-m_j}{2}\right) &= Y_i,\ \forall i\in O_0, &(25)\\ \sum_{i\in O_0}\exp\left(\frac{\varphi_{ij}-s_i-m_j}{2}\right) &= X_j, \forall j\in O. &(26)\end{aligned}$$

By simple factorization, this system can be written as

$$\begin{aligned} X_{i\emptyset}+X_{i\emptyset}^{1/2}\sum_{j\in O}K_{ij}M_j^{1/2} &= Y_i,\ \forall i\in O_0, &(27)\\ M_j^{1/2}\sum_{i\in O_0}K_{ij}X_{i\emptyset}^{1/2} &= X_j, \forall j\in O, &(28)\end{aligned}$$

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[34] is

$$z^2=\left(\left(Y+P^2\right)^{1/2}-P\right)^2. \tag{29}$$

[33]Note that setting $\alpha_{0\emptyset}\to-\infty$, one has $X_{0\emptyset}=\exp\left(\alpha_{0\emptyset}-s_0\right)\to 0$.
[34]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:

$$X_{i\emptyset} = \left(\left(Y_i + \left(\frac{1}{2} \sum_{j \in O} K_{ij} M_j^{1/2} \right)^2 \right)^{1/2} - \frac{1}{2} \sum_{j \in O} K_{ij} M_j^{1/2} \right)^2, \quad (30)$$

$$M_j = \left(\frac{X_j}{\sum_{i \in O_0} K_{ij} X_{i\emptyset}^{1/2}} \right)^2. \quad (31)$$

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.

B.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

$$
\begin{aligned}
X_{ij} &= \exp\left(\frac{\sum_{k=1}^{K}\beta_k D_{ij}^k - s_i - m_j}{2}\right) \forall i \in O_0, j \in O, \\
X_{i\varnothing} &= \exp\left(-s_i\right), \forall i \in O,
\end{aligned} \tag{32}
$$

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

As recently shown in Galichon and Salanié (2022b), 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

$$
w_{ij} = \gamma_{ij}^{b} - \frac{1}{2}\varphi_{ij}^{\beta} + \frac{1}{2}\left(s_i - m_j\right), \forall i \in O_0, j \in O. \tag{33}
$$

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 B.9 presents an extension in which we also incorporate information on wages at t.

B.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

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

B.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.