

The Role of Firms in Occupational Mobility

Jeremias Klaeui^{a,b}, Daniel Kopp^b, Rafael Lalive^{c,*}, Michael Siegenthaler^b

^a CREST, ENSAE, Institut Polytechnique de Paris, 91120 Palaiseau, France

^b KOF Swiss Economic Institute, ETH Zurich, 8092 Zurich, Switzerland

^c Department of Economics, University of Lausanne, 1015 Lausanne, Switzerland

* Corresponding author. Tel.: +41 21 692 34 31; rafael.lalive@unil.ch

October 8, 2026

Abstract

Using recruiter clickstream data and application diaries, we analyze when and why firms hire occupational movers. Movers are contacted 21% and hired 26% less often than incumbents. Holding job-seeker quality fixed, movers' contact gap falls to 10.4%, implying movers are negatively selected. Four-fifths of this gap reflects movers' lower expected overlap with the job's requirements. The gap is 3.3 times larger where an occupation's jobs share most requirements than where they differ widely, and it narrows when incumbents are scarce. Firms' greater openness accounts for 30% of the shift in contacts and hires from incumbents to movers as occupations tighten.

Keywords: occupations; occupational mobility; job requirements overlap; labor demand; labor supply; labor market tightness

JEL Codes: J24, J62, J64

Acknowledgments: We thank Maria Balgova, Michael Böhm, Marius Brülhart, Pierre Cahuc, Nikolaj Harmon, Simon Jäger, Attila Lindner, Alan Manning, Oskar Nordström Skans, Roland Rathelot, Ayşegül Şahin, Sara Signorelli, Daphné Skandalis, Jason Sockin, Lennart Ziegler, and Josef Zweimüller for comments. This project benefited from financial support through the Swiss National Science Foundation (SNSF) grant No. 407740, entitled “What Workers Want”, and from SNSF grant No. 204575. Jeremias Klaeui acknowledges support from the London School of Economics during his stay from summer 2023 to spring 2024. We thank Tobias Lehmann and Fabrizio Colella for support in analyzing and interpreting the job vacancy data.

1 Introduction

Occupational transitions are essential for workers’ ability to adapt to shocks and structural changes in labor demand. They allow workers to move out of declining industries and occupations and into those in demand, thereby fostering structural change in the economy (Acemoglu & Autor, 2011). High occupational mobility can reduce mismatch unemployment (Herz & Van Rens, 2020; Guvenen et al., 2020; Şahin et al., 2014) and dampen the adverse effects of layoffs, automation, and employer concentration on workers’ employment and wages (Bloesch et al., 2022; del Rio-Chanona et al., 2021; Gathmann & Schönberg, 2010; Macaluso, 2025; Robinson, 2018; Schubert et al., 2024).

A substantial body of literature conceptualizes occupational mobility from the perspective of workers.¹ In these accounts, firms typically do not play an active role: they consider candidates from other occupations whenever those candidates possess the required skills and credentials. But recent empirical evidence suggests that firms’ hiring decisions substantially shape realized occupational transitions (Behaghel et al., 2024; Azmat et al., 2025). In the Swiss job application data used in this article, 70% of applications come from “movers,” job seekers who last worked in a different four-digit ISCO occupation than the advertised position. Yet their hiring chances are a quarter lower than those of occupational “incumbents.” These results suggest that firms’ hiring decisions offset a substantial share of the occupational moves that workers attempt. Despite their apparent relevance, little research examines when and why firms prevent or facilitate occupational switches, likely in part because isolating firms’ role imposes steep data demands.

This paper uses a unique combination of matched Swiss administrative and online datasets to study firms’ role in occupational transitions. We analyze why firms hesitate to hire occupational movers, when they overcome that hesitation, and how much this shapes realized occupational transitions. We link unemployment insurance (UI) register data with two further sources: (i) clickstream data from “Job-Room,” the recruitment platform of the Swiss Public Employment Service, where recruiters screen the profiles of registered job seekers and decide whom to contact; and (ii) monthly application diaries, which record job seekers’ applications to vacancies and hiring outcomes. Together, these data address two core empirical challenges in studying firms’ role in occupational mobility. First, they allow us to separate workers’ and firms’ decisions, which realized occupational transitions conflate: both data sources show whom firms could hire and whom they choose. Second,

¹For example, Yamaguchi (2012), Autor & Handel (2013), Guvenen et al. (2020), and Lise & Postel-Vinay (2020) conceptualize occupational mobility through the lens of Roy-type choice models. In these models, heterogeneous job seekers sort endogenously into occupations based on the (perceived) fit between their skills and the occupation’s task requirements. Similarly, in the models of Gathmann & Schönberg (2010), Cortes & Gallipoli (2018), and Robinson (2018), workers choose the occupation that offers the highest reward for their skills. Cortes & Gallipoli (2018) also account for the costs of mobility as a function of the task distance between occupation pairs and occupation-specific entry costs, such as professional training, licensing, and union membership requirements.

the platform data allow us to account for job seekers’ potential selection into occupational moves. We recover the causal effect of job seekers’ occupational background on recruiters’ contact decisions using a within-person design that exploits the fact that the same person appears as an incumbent in some recruiter searches and as a mover in others.

Our data document the importance of firms as gatekeepers of occupational mobility. In the application diaries, movers’ applications are 26% less likely to result in a hire than applications of incumbents. On the recruitment platform, recruiters are 21% less likely to contact movers in their result list. However, these gaps, in part, reflect movers’ negative selection: job seekers who attempt to switch occupations have lower chances of being contacted in any occupation. When we compare the same candidate as a mover and as an incumbent, among job seekers seen by the same recruiter in the same search, the “mover gap” falls to 0.77 percentage points, or 7% of the baseline contact rate of 11%. The estimated mover gap is similar when controlling for interactions between the searched occupation and all 64 profile characteristics, which flexibly control for a candidate’s fit in a particular occupation. Since the platform features fewer movers attempting distant switches than the application diaries, our preferred estimate reweights the platform data to the diaries’ mix of movers and vacancy occupations. Reweighting leads to a mover gap of 10.4% of the contact rate.

We rationalize firms’ reluctance to hire movers in a framework of screening under uncertainty. Firms prefer incumbents because they do not observe a candidate’s complete skill set. They therefore infer expected fit from the requirements of the candidate’s previous occupation, her experience in the vacancy occupation, and her occupational credentials. Because movers’ last occupation overlaps less with the vacancy than that of incumbents, recruiters expect movers to meet fewer of the position’s requirements. Testing this notion requires measuring overlap for incumbents as well as movers, since even incumbents’ last jobs may have required different skills than the vacancy. But existing indices of occupational similarity do not quantify overlap between jobs *within* the same occupation. We thus construct a novel measure using data on job requirements extracted from the near-universe of online job postings in Switzerland. We randomly select pairs of vacancies from two occupations and estimate the proportion of job requirements present in both—a proportion that we can also compute for ads within the same occupation. We validate this measure by showing that it is highly predictive of whether two occupations are classified as “related” according to O*NET. It also strongly predicts the occupations to which job seekers apply.

Using the new measure, we estimate the contact gap for movers who, in expectation, meet the same share of the position’s requirements as incumbents. The gap drops by 80% in this comparison, suggesting that movers’ lower overlap is the main reason why recruiters contact them less. The same pattern appears in realized hires in the application diaries, albeit less precisely. The screening framework further predicts that within-

occupation homogeneity matters: being an incumbent signals greater overlap if jobs in an occupation share many requirements. The platform and application data support this prediction. On the platform, the mover gap is 3.3 times larger in the most homogeneous than in the most heterogeneous occupations. Together, these results suggest that recruiters read a candidate’s last occupation as information about her expected productivity. Their behavior is consistent with skill-based hiring: recruiters treat a candidate’s occupation as a signal of the skills she likely brings, not as a credential in itself. That recruiter behavior largely matches the predictions of our overlap measure also suggests that recruiters have accurate beliefs about the skills that experience in an occupation typically brings.

Next, we exploit the large number of recruiter decisions on the platform to systematically characterize the influence of occupational experience and occupational credentials on the mover gap. Because experience and a credential in the vacancy occupation signal fit directly, regardless of what the candidate’s last occupation required, the framework predicts that the gap is largest for job seekers from distant occupations with no prior experience in the searched occupation. Indeed, this group, which accounts for slightly more than a quarter of applications from other occupations in the diary data, faces a mover gap of 17.5%. While overlap has the largest effect on the mover gap across mover profiles, experience and a credential in the searched occupation each lower the gap independently. Return movers who hold a credential in the searched occupation and come from a very close occupation face no contact gap at all, but this profile accounts for just 2% of mover applications.

Are recruiters more open to movers when incumbents are scarce? In our framework, a scarcity of incumbents lowers the threshold that candidates must clear to be contacted. This benefits movers disproportionately if more movers than incumbents sit just below the initial threshold. To test how tightness affects the mover gap, we measure tightness at the occupation and region level as the ratio of the near-universe of online job openings to the number of registered unemployed job seekers who last worked in that occupation. We find that the mover gap is 8.8% of the baseline contact rate in the slackest occupations, while it is only 3.4% in the tightest occupations. For movers from closely related occupations, the gap in tight occupations becomes close to zero. These results suggest that the mover gap partly reflects how selective firms can afford to be, not only how well movers fit the position.

Finally, we take a step toward quantifying firms’ importance for occupational transitions by asking how much of the shift toward movers in tight occupations is due to firms. The object of interest is the difference between contacted (or hired) movers and contacted (or hired) incumbents per candidate in the pool. This difference rises in both data sources as incumbents become scarce, reflecting both changes in how firms treat movers—the demand side—and changes in how many movers are in the candidate pool—

the supply side. We use a Oaxaca-Blinder decomposition that attributes to firms the part of the rise stemming from the estimated change in the mover gap, and to the supply side the part stemming from the changing composition of the pool. This attribution assumes that the change in the mover gap reflects firms’ behavior rather than a change in the quality of the movers they see, which is plausible in the platform data because we control for candidate quality. We find that firms’ screening response accounts for about 30% of the shift from incumbents to movers as occupations tighten, in both the platform and the diary data.

Our paper makes four contributions. First, it extends the literature on the determinants of occupational mobility by evaluating the role of firms. Recent studies use online job-search data to study workers’ occupational scope (Banfi et al., 2019; Azar et al., 2022; Altmann et al., 2023; Klaeui, 2026). Much less is known about the demand side. The study closest to ours is the contemporaneous work by Azmat et al. (2025), who conduct a large-scale correspondence study across six occupations to analyze how retraining programs affect employer callbacks to fictitious movers. We extend this literature by conceptualizing how firms screen incumbents and movers under uncertainty about candidates’ skills, by explaining why firms are reluctant to hire movers, and by identifying factors that make firms more willing to consider them. Because we observe the candidate pool facing a vacancy, we can provide a first quantification of the roles of both market sides in the occupational transitions we observe.

Second, our analyses contribute to the nascent literature on the impacts of labor market tightness on firms’ hiring behavior. Several studies show that firms facing hiring difficulties lower their hiring standards in dimensions such as average candidate acceptance probability, skill requirements, or experience requirements (Modestino et al., 2016, 2020; Lochner et al., 2021; Le Barbanchon, Ronchi, & Sauvagnat, 2023; Ziegler, 2024; Cullen et al., 2025). While prior work has examined how job seekers’ occupational search scopes respond to tightness (e.g., Altmann et al., 2023; Ma & Samaniego de la Parra, 2021), we show that firms’ hesitation to hire movers is markedly weaker in tight occupations, which opens opportunities for job seekers to transition into them.

Third, our work provides an empirical foundation for targeted job search advice and informs a growing literature testing interventions that offer job seekers tailored occupational recommendations (Belot et al., 2019; Belot, Kircher, & Muller, 2025; Altmann et al., 2022; Dhia et al., 2022; van der Klaauw & Vethaak, 2022; Kircher, 2022; Le Barbanchon, Hensvik, & Rathelot, 2023; Bächli et al., 2025; Harmon et al., 2025). Our results suggest that movers face smaller gaps when they target occupations that share many job requirements with their last occupation and, less obviously, when they favor occupations in which jobs differ widely in their requirements. Movers also face smaller gaps in occupations that are currently tight.

Finally, our work contributes to the literature on measures of occupational similarity,

which are relevant for understanding labor market boundaries and labor reallocation. Previous studies measure occupational distance using observed occupational transitions (Schubert et al., 2024; del Rio-Chanona et al., 2021; Schmutte, 2014; Belot et al., 2019; Belot, Kircher, & Muller, 2025), task overlap between occupations², surveys among workers (Gathmann & Schönberg, 2010), or vocational education and training curricula (Eggenberger et al., 2018). Our measure instead builds on a few existing studies and computes overlap using job requirements extracted from online postings.³ This approach quantifies similarity based on up-to-date job requirements and could be used to estimate the similarity of jobs between regions, industries, or companies. A novelty is that our measure quantifies job requirements overlap within occupations, a feature that reveals substantial differences in preferences for incumbents depending on the homogeneity of jobs within occupations.

This paper is organized as follows. Section 2 presents a framework for screening under uncertainty that serves as the conceptual lens for our analyses. Section 3 describes the data and our new measure of job requirements overlap. Section 4 estimates the mover gap, examines the selection of movers, and tests the framework’s core predictions. Section 5 shows how the candidate pool and firms’ screening change as occupations tighten and quantifies firms’ contribution to occupational transitions. Section 6 concludes.

2 Screening Job Seekers from Different Occupations: A Framework

This section develops a simple framework that formalizes recruiters’ decisions when screening applicants from different occupations under uncertainty. The framework also motivates our empirical strategy.

The framework follows the logic of employer screening under uncertainty (Phelps, 1972). Recruiters cannot observe a candidate’s skills directly and therefore rely in part on signals they observe in the hiring process. A candidate’s previous occupation is one such

²Most of these studies extract the tasks or skills that are particularly prevalent in an occupation using occupational skill classification systems such as the Dictionary of Occupational Titles (DOT) (Poletaev & Robinson, 2008; Yamaguchi, 2012; Cortes & Gallipoli, 2018; Robinson, 2018; Bocquet, 2024), its successor O*NET (Alabdulkareem et al., 2018; Belot et al., 2019; Belot, Kircher, & Muller, 2025; Macaluso, 2025; Lyshol, 2022), or the Operational Directory of Trades and Jobs (ROME) (Goos et al., 2019).

³Leping (2009) measures skill overlap between firms based on job requirements in vacancies from an online job-search platform in Estonia. An advantage of our measure is that our skill and task classification considers over a thousand categories, whereas Leping (2009) considers only three broad skill groups. Bloesch et al. (2022) construct a measure of within-firm, across-position task differentiation based on US job posting data from Burning Glass Technologies (now Lightcast). Djumalieva et al. (2018) apply semi-supervised machine learning techniques to classify occupations based on skill requirements provided in online job ads in the UK. Finally, Dupuy et al. (2024) measure skill overlap and changes over time using job requirements in vacancies advertised in the US. An important difference to these papers is that we also compute the overlap between different jobs *within* an occupation.

signal. Occupations produce output by bundling tasks in specific ways, and these tasks require specific skills (Robinson, 2018).⁴ A candidate’s previous occupation therefore contains information about the skills she has likely acquired. This logic may explain why firms prefer incumbents. An average candidate from another occupation meets fewer of the job’s requirements than an average incumbent.

Let a vacancy be described by a binary requirement vector $x \in \{0, 1\}^K$, where $x_k = 1$ if the job requires skill k , and let the candidate’s skill vector be $a \in \{0, 1\}^K$, where $a_k = 1$ if she possesses skill k . The recruiter does not observe a , only signals z about the candidate. From these, she forms a posterior $p_k(z) = \Pr(a_k = 1 \mid z)$ that the candidate possesses skill k , collected in the vector $p(z)$. Assuming that all job requirements are weighted equally,⁵ the expected number of requirements the candidate meets is

$$m(x, z) = \mathbb{E}[x'a \mid z] = \sum_{k=1}^K x_k p_k(z) = x'p(z). \quad (1)$$

The recruiter contacts the candidate if this expected match clears a threshold, $x'p(z) \geq \tau_s$.⁶ The threshold is specific to the search s because it depends on the recruiter’s selectivity and on the candidates she can choose from. A pool with many strong applicants might let her set a higher bar. Throughout, we assume that recruiters hold correct beliefs: the posterior $p(z)$ coincides with the true conditional probability that the candidate possesses each skill.⁷

Employers can learn about a candidate’s skills from the requirements of her previous job. Suppose that a candidate possesses a given skill with probability π_1 if her previous job required it and with probability $\pi_0 < \pi_1$ otherwise. The baseline probability π_0 reflects the candidate’s general quality, since a more skilled job seeker is more likely to possess any given skill. The difference $\pi_1 - \pi_0$ is the additional probability of possessing a skill that the previous job required.

⁴Lazear (2009) makes this argument at the firm level. Individual skills may be general, but their combination is often specific, which makes switching easier between firms with a similar skill mix. Applied to occupations, several studies show that the transferability of skills shapes occupational transitions and their consequences for wages and re-employment (Poletaev & Robinson, 2008; Gathmann & Schönberg, 2010; Yamaguchi, 2012; Cortes & Gallipoli, 2018; Robinson, 2018; Goos et al., 2019; Macaluso, 2025; Böhm et al., 2024, among others).

⁵In a robustness check, we instead attach skill-specific weights and use $\sum_k w_k x_k a_k$. In this alternative approach, rare distinctive requirements matter more than common ones (Section 3.3). The framework’s insights carry over to this case, with the weighted overlap $\sum_k w_k x_k \bar{x}_k$ replacing $x'\bar{x}$ throughout.

⁶The additive form of the expected match assumes that a candidate can compensate for a missing requirement by meeting others. In practice, some requirements may not be perfectly substitutable. Recruiters may then screen out any candidate who lacks them regardless of her other skills. The most prominent example is an occupational license, which could be a hard requirement. Indeed, we find that the contact gap is roughly 45% larger in licensed than in unlicensed occupations (Section 4.3).

⁷The assumption of accurate beliefs follows the classic theory of statistical discrimination in hiring (Phelps, 1972; Aigner & Cain, 1977). It is consistent with our results below, where measurable signals of candidates’ skills explain recruiters’ contact decisions well.

Recruiters rarely know the requirements of an applicant's previous job, at least not at the pre-selection stage of the hiring process. What they observe is her previous occupation, and with it the requirement profile of that occupation, $\bar{x} \in [0, 1]^K$, whose k -th entry is the share of jobs in the occupation that require skill k . From the recruiter's point of view, $\bar{x}_k$ is thus the probability that the candidate's previous job required skill k , and her posterior that the candidate possesses it is $p_k(\bar{x}) = \pi_1 \bar{x}_k + \pi_0 (1 - \bar{x}_k)$. The expected match then decomposes into

$$x'p(\bar{x}) = \pi_0 |x| + (\pi_1 - \pi_0) x'\bar{x}, \quad (2)$$

where $|x|$ is the number of requirements in the vacancy and $x'\bar{x} = \sum_k x_k \bar{x}_k$ is the expected overlap between the vacancy and a job in the candidate's previous occupation. The baseline term $\pi_0 |x|$ is the candidate's expected fit absent any overlap. Recruiters are more likely to contact job seekers with higher overlap.

Some of the candidates the recruiter sees are incumbents, that is, candidates who last worked in the vacancy occupation. Recruiters form expectations about their fit in the same way. Let $\bar{x}^I$ denote the requirement profile of the vacancy occupation and $\bar{x}^M$ the profile of a mover's previous occupation. Because jobs *within* an occupation also differ in their requirements, the entries of $\bar{x}^I$ lie below one. Even an incumbent's expected overlap thus falls short of the vacancy's requirements, $x'\bar{x}^I < |x|$. The size of the shortfall depends on how homogeneous the occupation is. In occupations where jobs share many requirements, an incumbent's previous job very likely required the skills the vacancy asks for.

Let π_0^I and π_0^M denote the general quality of the incumbent and the mover. With expected matches $m^I = x'p(\bar{x}^I)$ and $m^M = x'p(\bar{x}^M)$, and assuming an identical return to overlap for movers and incumbents ($\pi_1^I - \pi_0^I = \pi_1^M - \pi_0^M = \pi_1 - \pi_0$), the mover gap in expected match is

$$m^M - m^I = \underbrace{(\pi_0^M - \pi_0^I) |x|}_{\text{selection}} + \underbrace{(\pi_1 - \pi_0) (x'\bar{x}^M - x'\bar{x}^I)}_{\text{occupation}}. \quad (3)$$

The occupation term in (3) reflects the difference in overlap between movers and incumbents. It is negative on average, since movers have less overlap with the vacancy than incumbents. The selection term reflects differences in general quality between the two groups. For instance, if lower-ability candidates are more likely to attempt to switch occupation, the selection term would be negative and widen the mover gap.

This framework makes several predictions about the factors that affect the gap in expected match. First, the gap depends on the difference in overlap and should therefore shrink as a mover's overlap with the vacancy increases. Second, it should be larger when the vacancy occupation is homogeneous, since an incumbent's previous job is then more

likely to have required the vacancy’s skills. Third, it may be smaller when incumbents are scarce, because scarcity may induce recruiters to lower τ_s . Lowering the bar raises the probability of contact for movers and incumbents alike, but possibly not by the same amount. If movers are more likely to sit just below the original threshold, they may disproportionately benefit from a lower threshold.

A final prediction concerns signals of occupation-specific fit beyond the last occupation. Recruiters also see credentials and work history that speak directly to a candidate’s fit with the advertised position. Signals such as prior occupational experience raise the probability that the candidate possesses the vacancy’s required skills even when her previous occupation did not require them. A richer framework that incorporates such signals thus predicts that these credentials raise contact rates beyond what overlap alone predicts.

We recover the contact gap between movers and incumbents in (3) by regressing a contact indicator on a mover indicator. Our main specification includes candidate fixed effects, which hold constant each job seeker’s baseline probability of possessing any required skill (π_0). Together with occupation-specific valuation controls introduced in some specifications, the individual fixed effects also absorb the selection term in Equation (3). In addition, the specification contains search fixed effects, which absorb everything common to the candidates in a single result list, including the threshold τ_s .

3 Data and Measurement

Our analyses draw on three Swiss data sources: clickstream data from a recruitment platform, job seeker application diaries, and online vacancy postings. We link the first two to unemployment register data. This section describes each in turn and introduces our measure of job requirements overlap.

3.1 Recruiter Clickstream Data from Job-Room

Our core dataset is recruiter clickstream data from “Job-Room”, the recruitment platform of the Swiss Public Employment Service (PES), where recruiters search a database of standardized CVs of registered job seekers. These data reveal every step of the screening process: the search terms recruiters enter, which candidates appear in their results, and whom they choose to contact. The data cover March to December 2017. Hangartner et al. (2021) use them to study ethnic and gender discrimination in the Swiss labor market, and Kopp (2025) uses them to document recruiters’ hiring penalties against male job seekers who wish to work part-time.

The candidate profiles on the platform belong to job seekers registered with the PES. We can therefore link the profiles to job seekers’ entries in the Swiss unemployment

register, which records their age, education, nationality, occupational history, and unemployment history, as well as the date they leave unemployment. Job-Room covers about 80% of all registered job seekers in Switzerland. Registration is open to all job seekers, but mandatory for UI recipients. Of the registered job seekers, 61% receive UI benefits.

The recruiters who use the platform typically seek suitable candidates for one or several open positions. We identify recruiters using a cookie installed in users' browsers when they first visit the platform after December 2016.⁸ Recruiters usually start by entering the occupation they are hiring for, and most restrict the search to a canton or region. The platform then returns a list of up to 100 candidates who match the search criteria exactly. The platform matches the searched occupation against the occupations listed on candidates' profiles. The list order is unrelated to a candidate's last occupation. Thus, incumbents—job seekers whose last occupation is the searched occupation—are not ranked above movers. The result list displays the candidates' gender, canton of residence, and desired work volume (in full-time equivalents), and whether job seekers are “immediately available” or can start to work “by agreement.” It also shows the content of an unstructured text field with additional “knowledge, abilities, skills” of job seekers. This field is specific to the recruiters' searched occupation and is visible only if the job seeker filled it out with her caseworker at the PES. The result list does not show a candidate's last occupation. To see it, the recruiter has to open the candidate's full profile.

Figure 1 shows a screenshot of the full candidate profile.⁹ Similar to a standard CV, the profile shows detailed information on candidates' skills and credentials, including their language proficiency, work experience, and educational attainment, as well as personal characteristics such as gender, name, and nationality. At the top of the profile, recruiters can see the occupation a candidate most recently worked in. Below that, the profile lists all other occupations the job seeker is willing to work in. For each occupation, the profile reveals a candidate's years of occupational experience in categorical format (none, less than 1 year, 1–3 years, and more than 3 years). It also shows any professional degree or training in that occupation, and whether that degree is Swiss, foreign and recognized in Switzerland, or foreign and not recognized. The profile also repeats the text field with job seekers' “knowledge, abilities, skills.”

A candidate appears in a recruiter's results whenever she lists the recruiter's searched occupation on her profile as an occupation she is willing to work in. The job seeker and her PES caseworker define the list of occupations during the initial meeting. It is rarely updated during the unemployment spell and is binding: a job seeker can be required to accept an offer in any of her listed occupations. Job seekers therefore tend to list

⁸What we term ‘recruiters’ are thus typically the users of the same browser on the same workstation. The presumably few recruiters who deleted their browser's cache during the sample period may be counted as several recruiters.

⁹Appendix Figure A.1 shows a screenshot of the original, German version of the profile.

Figure 1. Recruiter clicks: screenshot of the candidate profile on Job-Room

Candidate's Occupations, Qualifications and Experience	
Last occupation held:	Factory Worker
Years of experience in occupation:	more than 1 year
Occupation title:	Auto Mechanic
Qualification for occupation:	foreign, not recognized in Switzerland
Years of experience in occupation:	more than 3 years
Knowledge, Abilities, Skills:	Training in Macedonia
Occupation title:	Construction Laborer
Years of experience in occupation:	more than 1 year
Candidate's Language Skills	
German:	oral: good, written: basic knowledge
Albanian:	oral: very good, written: very good
Additional Candidate Information	
Availability:	by arrangement
Work type:	no information
Gender:	male
Maximum work percentage:	100%
Possible employment duration:	permanent
Highest education level:	Secondary Level II - ISCED 3
Education:	Sec. II - Vocational Education and Training (VET) or equiv. Sec. I - Compulsory School
Desired work region(s):	Major Region 4 (ZH, SH, TG, SG, AI, AR, GL, GR)
Driver's license categories:	B
Further information provided by:	
Address:	RAV Frauenfeld, Thundorferstrasse 37, 8510 Frauenfeld Cantonal Administration
Show contact details	
Back Send this candidate as link Print view ◀◀ ▶▶	

→ Jobseeker's last occupation

→ Contact button

Notes: The figure shows a screenshot of an example candidate profile on Job-Room.ch (own translation).

occupations in which they have work experience. The platform uses a detailed internal Swiss classification (AVAM), which we aggregate to four-digit ISCO-19 occupations to harmonize it with our other data sources.¹⁰

Our main dependent variable is whether the recruiter clicks the “Show contact details” button on a candidate’s profile. In 2017, recruiters could not contact and hire a candidate through Job-Room without clicking this button. Hangartner et al. (2021) validate the outcome by showing that each contact click raises a candidate’s exit rate from unemployment within three months by 2.1%.

We impose the following sample restrictions on the original sample of 469,779 recruiter searches. Most importantly, we drop 62,777 searches without a specified occupation, since we cannot classify candidates as movers or incumbents in these cases. We also drop candidate profiles without a last occupation, searches with missing search-identifier tokens, and the few searches in military occupations. The analysis sample comprises 15,266,233 result-list items from 406,873 recruiter searches of 37,665 distinct recruiters. The recruiters saw 325,545 distinct job seekers.

Table 1 reports descriptive statistics of our analysis sample, distinguishing all job seekers (Column (1)), incumbents (Column (2)), and movers (Column (3)). The top panel shows that the average job seeker lists 2.3 occupations, with 95% listing their last occupation. 69% of job seekers list at least one occupation beyond their last. Incumbents and movers look quite similar. Incumbents are slightly younger and more experienced in the searched occupation, but many movers also have experience in that occupation, seeking to return to a position they previously held (“return movers”).

The middle panel of Table 1 shows that recruiters find, on average, 37.5 candidates per search, of which 19.6 are incumbents. Recruiters’ contact rates differ by occupational background: they contact 12.2% of incumbents but only 9.6% of movers, an unconditional mover gap of 21%. The table also shows where this gap arises. Recruiters open the full profile at relatively similar rates for incumbents and movers (25% and 23%), consistent with the fact that the result list does not reveal the last occupation. Among viewed profiles, however, recruiters contact 48% of incumbents but only 41% of movers. Recruiters search disproportionately in craft and trade occupations, but we observe several thousand searches in all major occupational groups (see Appendix Table A.1).

¹⁰We convert AVAM codes to four-digit ISCO-19 using a crosswalk built from unemployment records, which report each job seeker’s last occupation in both classifications. The crosswalk assigns each AVAM code the ISCO four-digit occupation most frequent among registered job seekers with that AVAM code; for 4,995 of 5,801 AVAM occupations, the mapping is unique. We aggregate occupational experience accordingly.

Table 1. Recruiter clicks: descriptive statistics

	(1) All job seekers	(2) Incumbent job seekers	(3) Job seekers from a different occupation
<i>Job seeker characteristics:</i>			
Number of occupations listed by jobseeker	2.32	—	—
Probability to list last occupation (%)	94.58	—	—
Probability to list an occupation other than the last (%)	69.11	—	—
Share female (%)	41.52	42.97	40.59
Average age	39.32	38.71	39.71
<i>Share by highest education (%):</i>			
Tertiary	20.45	20.54	20.38
Upper secondary	51.69	51.50	51.81
Lower secondary	22.06	21.32	22.53
<i>Experience in searched occupation (%):</i>			
More than 3 years	70.84	76.72	64.03
1 to 3 years	16.72	15.26	18.41
Less than 1 year	7.94	8.02	7.84
<i>Experience in last occupation (%):</i>			
More than 3 years	69.38	76.72	61.37
1 to 3 years	19.04	15.26	23.16
Less than 1 year	11.58	8.02	15.47
<i>Recruiter search results:</i>			
Candidates in result list per search	37.52	19.58	17.94
Candidate views per search	9.14	4.95	4.19
Candidate contacts per search	4.13	2.40	1.73
Share contacted (%) (cond. on in result)	11.00	12.24	9.64
Share contacted (%) (cond. on viewed)	45.13	48.40	41.27
<i>Count data:</i>			
Recruiters	37 665	—	—
Searches	406 873	—	—
Result list entries	15 266 233	—	—
Distinct job seekers in result lists	325 545	—	—

Notes: The table reports descriptive statistics of the candidate pool and recruiters' click behavior on Job-Room. Column (1) covers all candidate profiles appearing in recruiters' result lists; Columns (2) and (3) split them into incumbents, whose last occupation is the searched occupation, and movers, whose last occupation differs. The top panel reports statistics on registered job seekers visible on the platform. The middle panel reports averages over recruiter searches. The bottom panel reports the underlying counts.

3.2 Online Application Diaries

The platform data enable us to analyze the causal effect of job seekers’ most recent occupation on recruiters’ contact rates. But they have two drawbacks. First, they reveal only the initial contact attempt, not whether a candidate is ultimately hired. If the last occupation influences success at the interview stage, we may mismeasure its importance. Second, it is unclear whether the behavior of recruiters who actively approach potential employees generalizes to hiring in the wider labor market. Recruiters who actively search for candidates may, for instance, be more open to considering job seekers with a lower fit to the position.

To address these drawbacks, we supplement the platform data with the application diaries of unemployed job seekers. These data capture actual applications and their hiring outcomes. As in the Job-Room data, we can link the application diary data to job seekers’ UI register entries. We therefore observe their occupational history, their experience in each occupation, and their occupational credentials. The full data description and cleaning steps are in Appendix B.

Switzerland’s UI system requires benefit claimants to document their application activities each month in application protocols. Since 2022, they can do so digitally through Job-Room by importing a vacancy’s details from the posting, which links the application to the official vacancy identifier. The identifier links the application to the record of the vacancy in the UI register, which reports the vacancy’s four-digit ISCO occupation. Imported applications make up 34.7% of all applications recorded by the job seekers who use the import function.

Our primary outcome variable in the diary data is an indicator of a successful application. This indicator is one if the firm the job seeker applied to is the firm at which the UI register records her re-employment. Constructing this outcome requires matching employer names in the UI register with those in the diaries. We thus implement a string-matching procedure that achieves 91.9% balanced accuracy on a hand-labeled test set (see Appendix B.1). Job seekers in our analysis sample have a 1.2% probability of being hired conditional on applying.¹¹

Our diary sample, drawn from the second to the fourth quarter of 2022, contains 23,477 job seekers, of whom 12,230 have a completed spell with a known re-employment firm. We restrict the sample to these job seekers because we cannot observe the hiring outcome for those still unemployed at the end of our observation window, and it is unobservable when the register does not record the re-employment firm. In total, we observe 167,024 applications to 96,611 distinct vacancies. Our analysis sample is 166,029 applications

¹¹This number exceeds the 0.9% job finding rate in Zuchuat et al. (2023), who use manually coded Swiss application protocols from 2012 and 2013 with direct information on whether an application was successful. Part of the difference reflects that our sample conditions on job seekers who eventually found a job.

after dropping applications with an unmatched employer, a missing overlap measure, or missing controls or fixed-effect identifiers. The analyses that condition on labor market tightness additionally focus on the subset of 159,061 applications for which the vacancy occupation and canton identify a tightness cell (Appendix B).

Despite its size, the application sample is not a random sample of job seekers. Appendix Table B.1 shows that job seekers who use the online form are somewhat more likely to change occupation upon re-employment than registered job seekers overall (58% versus 53%). Otherwise, the sample resembles the registered job seeker population closely.

The linked diary-UI register data also provide a first descriptive look at the mover gap in hires. Movers represent the majority in the applicant pool: 70% of applications come from job seekers whose last job was in a different four-digit ISCO occupation than the vacancy occupation. But their applications are less likely to succeed. A mover's application results in a hire 1.12% of the time, against 1.51% for an incumbent. These numbers suggest that firms are 26% less likely to hire an applicant who comes from another occupation than one from the vacancy's occupation.

3.3 Job Requirements Overlap: A New Index

Testing our theoretical framework requires measuring the overlap in requirements between two jobs in two distinct occupations and comparing it with the overlap between two jobs within the same occupation. We therefore cannot resort to standard measures of occupational similarity based on the overlap of tasks according to occupational skill classifications (used, e.g., by Gathmann & Schönberg, 2010; Cortes & Gallipoli, 2018; Belot, Kircher, & Muller, 2025; Eggenberger & Backes-Gellner, 2023) because they do not allow measuring the similarity of jobs within an occupation. Since transitions across occupations are the outcome of interest, measuring similarity based on observed transitions between occupations, as in del Rio-Chanona et al. (2021) or Schubert et al. (2024), is not an option either. Building on a few previous studies (Leping, 2009; Djumalieva et al., 2018; Bloesch et al., 2022; Dupuy et al., 2024), we thus construct a measure of job requirements overlap from the requirements listed in job postings. The vacancy data cover the near-universe of online job postings in Switzerland between 2016 and 2022 (see Appendix C.1 for details). Each of the more than 1 million job postings is assigned to a four-digit ISCO occupation based on the job title and the ad's context. The data also include a detailed list of job requirements extracted from the ad text and the job title.

A concern when using job vacancy data to measure overlap is that a job ad may not capture all the skills a position demands. In the framework of Woessmann (2025), vacancy data primarily capture applied skills, the skills required to perform concrete job tasks. They are less likely to record basic skills, the foundational skills that facilitate skill acquisition and are likely relevant for hiring. To capture basic skills as far as possible,

we include the educational degrees and professional certificates an ad requires, which are more likely to reflect them. We then compute overlap across all requirements an ad lists: tasks, skill requirements, and degrees. We include the listed tasks both because a successful candidate must likely possess the skills required to perform them and because distinguishing skill requirements from job tasks in postings is technically challenging.¹²

We compute the index in two steps (see Appendix C.2 for details and Figure C.1 for a stylized example). For a given pair of occupations, we first draw 5,000 random pairs of postings, one from each occupation. We focus on occupations with at least 200 postings.¹³ We then compute the Jaccard similarity for each occupation pair, defined as the number of requirements appearing in both ads divided by the number of distinct requirements across both ads. We then average over the 5,000 pairs. We apply the same procedure to two postings in the same occupation, which places within- and between-occupation overlap on a common scale. We compute the index from postings in the German-speaking part of Switzerland, where the extraction of requirements is more reliable, pooling all postings from 2016 to the second quarter of 2022. Since not every shared requirement is equally informative about how interchangeable two jobs are, we also compute an alternative version that weights requirements by their inverse document frequency (see Appendix F for details).

The resulting overlap index is the empirical counterpart of the expected overlap in the framework. Averaging over random pairs of postings corresponds to the expectation over jobs in an occupation. Hence, the between-occupation index corresponds to $x'\bar{x}^M$ and the within-occupation index to $x'\bar{x}^I$. In a slight deviation from our framework, which focuses on the raw count of shared requirements, the Jaccard index normalizes by the number of distinct requirements in a pair. This normalization is empirically attractive because it accounts for the sizable cross-occupational differences in the number of job requirements posted per vacancy.

Figure 2 shows the estimated job requirements overlap for the ten occupations with the most registered job seekers. As expected, overlap is higher within occupations than between them. Across all occupation pairs, it averages 5.5% between two distinct occupations and 9.8% within the same occupation.¹⁴ To put this difference in perspective, the median ad lists 9 distinct requirements, and the median pair of ads shares about one of them (a 5.5% Jaccard). A difference of roughly 4 percentage points thus corresponds to about 0.7 additional shared requirements. Within-occupation overlap also varies substantially. Among the ten occupations in the figure, it ranges from 6.1% for manufacturing

¹²Job postings often detail the duties and tasks of the open position, as well as the skill requirements for suitable candidates. For example, a job posting could list the word “programming” as both a skill requirement and a job task.

¹³This covers 80% of all pairs of four-digit occupations. For the remaining pairs, we impute overlap from the average among pairs at the three-digit level or, where necessary, at the two- or one-digit level (Appendix C.3). On the platform, 2.7% of result-list entries rely on an imputed value.

¹⁴The distribution of overlap across all occupation pairs is shown in Appendix Figure C.2.

laborers to 11.7% for cooks, and across all occupations from 4.8% for university teachers to 35.6% for typists. Jobs within some occupations thus differ as much from each other as jobs across occupations.


Figure 2. Job requirements overlap: between and within the ten most common occupations

	Office clerks	Manufacturing labourers	Shop sales assistants	Cleaners and helpers	Waiters	Building construction labourers	Kitchen helpers	Freight handlers	Cooks	Material-recording and transport clerks
Office clerks	7.8	5.3	6.9	6	6.1	4	5.4	4.8	5.7	6
Manufacturing labourers	5.3	6.1	6.3	6.2	5.7	6.1	5.5	6.1	5.3	5.7
Shop sales assistants	6.9	6.3	10.9	7.8	8.1	4.9	6.8	5.9	6.9	6.2
Cleaners and helpers	6	6.2	7.8	8.6	8	4.7	7.3	5.4	6.9	5.7
Waiters	6.1	5.7	8.1	8	9.3	4.4	6.8	5	6.8	5.2
Building construction labourers	4	6.1	4.9	4.7	4.4	10.5	4.5	7	4.1	4.9
Kitchen helpers	5.4	5.5	6.8	7.3	6.8	4.5	6.7	5.1	7	5.2
Freight handlers	4.8	6.1	5.9	5.4	5	7	5.1	7.7	4.6	6.9
Cooks	5.7	5.3	6.9	6.9	6.8	4.1	7	4.6	11.7	5.2
Material-recording and transport clerks	6	5.7	6.2	5.7	5.2	4.9	5.2	6.9	5.2	9.3

Notes: The figure shows the average job requirements overlap, in percent, between and within the ten occupations with the most registered job seekers in the Swiss unemployment register. Each number is the average Jaccard index over 5,000 randomly drawn pairs of job postings from the two occupations.

We validate the measure in three ways. First, Appendix Figure C.3 shows that our measure of job requirements overlap lines up with an external classification of related occupations from O*NET. Occupation pairs with less than 4% overlap are practically never classified as related by O*NET. In contrast, pairs with more than 10% overlap have about a 12% chance of being classified as related. The latter share is low because O*NET lists at most ten related occupations per occupation, which keeps the base rate of related pairs low. Second, the measure strongly predicts job seekers' applications. Figure 3 uses the diary data to relate how often a job seeker applies to an occupation over an unemployment spell to the overlap between that occupation and her most recent occupation. The figure shows that job seekers rarely apply to occupations sharing less than 5% of their last occupation's requirements. In contrast, they apply more than three times per spell to occupations with more than 11% overlap (Panel a). The same pattern holds when we restrict attention to applications outside the job seekers' most recent occupation (Panel b). Third, Section 4 shows that the index also strongly predicts recruiters' contact decisions.

Figure 3. Application diaries: job requirements overlap and job seekers' applications



Notes: Panels (a) and (b) are binned scatter plots. The y-axis shows the number of applications made to an occupation over an unemployment spell, the x-axis the average job requirements overlap between that occupation and the job seeker's last occupation. Panel (a) uses all occupations; Panel (b) restricts to occupations different from the last one. To fill in the occupations to which a job seeker has zero applications, we expand the data to job seeker-occupation combinations. For computational feasibility we restrict this grid in two ways: we draw a random fifth of the job seekers, and we only look at the 50 largest occupations in the data. Dot size is proportional to the number of job seeker-occupation combinations in the bin and the lines are local polynomial fits to these combinations. Overlap is winsorized at the 5th percentile for different-occupation pairs and at the 95th percentile for same-occupation pairs; application counts are winsorized at the 95th percentile.

4 Size and Determinants of the Mover Gap

In this section, we use recruiter clickstream data from Job-Room, application diary data from the UI register, and our measure of job requirements overlap to answer two questions. How large is the mover gap, the difference in contact rates between movers and incumbents when holding job-seeker quality constant? And what determines it? The framework predicts that the gap should shrink as the overlap between a candidate's last occupation and the vacancy rises and widen when the vacancy occupation is homogeneous. It should close when a mover provides direct evidence of fit, such as experience or a degree in the vacancy occupation.

We proceed as follows. Section 4.1 introduces the research design. Section 4.2 establishes the average mover gap and examines whether job seekers who attempt to switch occupations are positively or negatively selected. Sections 4.3 and 4.4 then test the framework's predictions on overlap, homogeneity, and occupation-specific signals of fit in turn.

4.1 Empirical Approach

In this section, we estimate the contact gap for job seekers who come from a different occupation relative to incumbents. In the framework, this is the mover gap in expected match (Equation (3)). The term of primary interest is the occupation term, which is the part of the mover gap that arises because a mover’s previous occupation overlaps less with the vacancy’s requirements than an incumbent’s. The selection term primarily reflects differences in general quality between movers and incumbents, which arise if job seekers who attempt to switch occupations are positively or negatively selected.

A major advantage of the Job-Room clickstream data is that they allow us to hold the selection term constant. To this end, we exploit the fact that we observe the same job seeker in the result lists of recruiters who search in different occupations. We can therefore compare how the same candidate is treated when she is an incumbent versus when she is a mover. Empirically, we estimate the mover gap in regression specifications with candidate fixed effects. These fixed effects absorb a candidate’s average contact rate across occupations, and thus control for a candidate’s overall quality (the baseline π_0 of the framework).

Candidate fixed effects alone do not remove the selection term entirely, however, as the same signal may influence job seekers’ expected fit more in some occupations than in others.¹⁵ For example, a recruiter hiring a sales assistant may place more value on strong German than a recruiter hiring a kitchen helper. In an extended specification, we therefore control for interactions between the searched occupation and the candidate’s characteristics, which allow each characteristic to be valued differently in each occupation.

More formally, we estimate the mover gap using variants of the following linear probability model:

$$y_{i,s} = \alpha \text{Mover}_{i,s} + \gamma_i + X'_{i,o(s)}\beta + \psi_{\text{rank}(i,s)} + \phi_s + \varepsilon_{i,s}. \quad (4)$$

The unit of observation is a job seeker i appearing in recruiter search s . The outcome variable, $y_{i,s}$, is an indicator equal to one if the recruiter clicks the button revealing candidate i ’s contact details. The main regressor, $\text{Mover}_{i,s} = \mathbb{1}\{o_l(i) \neq o(s)\}$, equals one if candidate i ’s last occupation, $o_l(i)$, differs from the occupation specified by the recruiter in search s , $o(s)$. The coefficient α captures the mover gap—the empirical counterpart of the occupation term in Equation (3). As discussed above, in our preferred specification, the coefficient reflects the gap net of the selection term. We first recover the contact gap for an average mover. Hence, we do not control for a candidate’s potential work experience or degrees in the searched occupation, as displayed to recruiters on Job-

¹⁵Formally, Equation (3) shows that the selection term, $(\pi_0^M - \pi_0^I)|x|$, multiplies the difference in general quality by the number of requirements of the vacancy, $|x|$, which varies across searches. Candidate fixed effects hold π_0 constant and search fixed effects hold $|x|$ constant, but neither absorbs their product.

Room. We treat occupation-specific skills and credentials as determinants of the mover gap and analyze them in Sections 4.3 and 4.4.

The main ingredients for identifying the mover gap in (4) are the search fixed effects, ϕ_s , and the candidate fixed effects, γ_i . The search fixed effects imply that we compare recruiter decisions across candidates who appear in the same search result list. Unless the recruiter uses the same search terms to fill several positions, we thus compare only job seekers evaluated for the same position. The search fixed effects also control for all unobserved recruiter characteristics, the searched occupation, other search criteria, and the labor market conditions at the time of the search. They further hold the candidate pool in search s constant, including its quality and the total number of candidates the recruiter contacts. As a result, the search fixed effects hold the search-specific hiring threshold τ_s constant, i.e., we compare only job seekers evaluated against the same decision criteria.

The candidate fixed effects control for all candidate characteristics that influence contact rates and are constant across recruiter searches. They hold the time-invariant content of the profile displayed to recruiters fixed, and with it the candidate’s average attractiveness across searches. They imply that the mover gap is identified only from job seekers who appear both as movers and as incumbents, which requires that they list at least one occupation beyond their last (69% of job seekers). Comparing the estimated mover gap with and without candidate fixed effects also allows us to test whether movers are positively or negatively selected.

Equation (4) contains two further elements. The control variables in the vector $X_{i,o(s)}$ are the occupation-specific valuation controls described above. The vector comprises first-order interactions between recruiters’ searched occupation (at the three-digit level) and the 64 candidate characteristics that do not vary by occupation (see the list in Appendix Table A.2). Finally, the rank fixed effects $\psi_{\text{rank}(i,s)}$ account flexibly for a candidate’s absolute and relative position in the result list. These controls are added primarily for precision and matter little for identification, as incumbents are not systematically better-ranked on Job-Room than movers.¹⁶

4.2 The Average Mover Gap and the Selection of Movers

Table 2 presents the baseline mover gap estimates based on the recruiter clickstream data. Standard errors are clustered two-way by recruiter and by candidate. Column (1) shows that candidates who last worked in an occupation different from the recruiter’s searched occupation face a highly significant contact gap relative to otherwise similar incumbents. The estimated mover gap amounts to -0.77 percentage points, or 7% of the average contact rate of 11%. Column (2) adds the occupation-specific valuation controls

¹⁶Within a search result list, the share of mover candidates rises just slightly with rank (see Appendix Figure F.1).

(see Section 4.1). The gap falls from -0.77 to -0.71 percentage points, suggesting that the valuation of job seekers' skills does not differ much across the occupations they are willing to work in.

The second main result of the table concerns the selection of movers. The mover gap is much smaller than the raw difference in contact rates of 2.6 percentage points (Table 1). This has two reasons. First, movers appear disproportionately in searches in which recruiters contact few candidates. Controlling for search fixed effects, which imply that we compare only candidates within the same search, reduces the gap to 1.5 percentage points (Column 3). Second, movers are negatively selected. Adding candidate fixed effects halves the gap again (Column 1). In the framework's terms, this implies that the selection term in Equation (3) is negative. The results suggest that job seekers who appear as movers have lower contact rates across all their searches than job seekers who appear as incumbents. This is in line with the conjecture in Section 2 that lower-ability candidates are more likely to attempt an occupational switch.

How large is a mover gap of 7%? Column (4) offers a benchmark. It drops the candidate fixed effects and controls for candidate characteristics instead, which identifies their effects. The mover gap then amounts to about two-thirds of the contact premium for good or very good German skills, relative to lower proficiency, in searches in German-speaking cantons.

Three tests support the premise that recruiters act on the information displayed on the profile and not on information we do not observe. First, replacing candidate fixed effects with control variables for all candidate information visible on the Job-Room profile yields a similar estimate (Column 4), as expected if recruiters' contact decisions depend only on observable candidate characteristics. Second, Column (5) re-estimates Column (4) on the sample of candidates without an entry in the unstructured text field about their knowledge, abilities, and skills. The gap is essentially unchanged, suggesting that the hard-to-codify field does not drive the estimate. Third, Column (5) also controls for the candidate's age, which we observe in the UI register but recruiters do not see on Job-Room. We note that age has no effect conditional on the observables, consistent with recruiters acting only on the displayed profile.¹⁷

The estimate is robust along several further dimensions. First, the gap is virtually unchanged when we define a mover at the three-digit rather than the four-digit ISCO level (Appendix Table F.1, Panel A, Column 2). Second, the gap is smaller in relative terms but still highly significant when we restrict the sample to candidate profiles the recruiter either never opened or viewed for at most 20 seconds. In these decisions, recruiters are unlikely to have gathered outside information about candidates (Appendix

¹⁷This age placebo uses the sample without the unstructured text field, since the field sometimes contains information about candidates' experience that is correlated with age (e.g., "more than 20 years of work experience").

Table 2. Recruiter clicks: the contact gap for job seekers from a different occupation

	(1)	(2)	(3)	(4)	(5)
Last worked in different occupation	-0.0077*** (0.0004)	-0.0071*** (0.0004)	-0.0151*** (0.0006)	-0.0086*** (0.0004)	-0.0082*** (0.0006)
Good German skills				0.0126*** (0.0011)	0.0116*** (0.0018)
Age = 30-44					-0.0004 (0.0010)
Age = 45-59					0.0006 (0.0010)
Age = 60+					0.0005 (0.0018)
Candidate FE	X	X			
Recruiter search FE	X	X	X	X	X
Search rank (abs. and rel.) FE	X	X	X	X	X
Vacancy occ. FE x candidate observables		X			
Controls for candidate observables				X	X
-----	-----	-----	-----	-----	-----
Observations	15,266,233	15,266,233	15,266,233	15,266,233	4,624,254
R2	0.46442	0.46495	0.43919	0.44653	0.48927
N searches	406,873	406,873	406,873	406,873	406,873
N recruiters	37,665	37,665	37,665	37,665	37,665
Baseline prob.	0.1100	0.1100	0.1100	0.1100	0.1100
Sample	All	All	All	All	No text

Notes: The table reports estimates of the contact gap for movers relative to incumbents from Equation (4). Each observation is a candidate profile appearing in a recruiter’s result list. The outcome is whether the recruiter clicks the button revealing the candidate’s contact details. The main regressor is whether the candidate’s last occupation differs from the occupation the recruiter entered as a search term. Search rank fixed effects control for the absolute rank, with one dummy per rank, and the relative rank, one dummy per decile, of a candidate within the list. “Controls for candidate observables” are 64 dummies capturing all information displayed on the profile, listed in Appendix Table A.2. Column (2) interacts them with the three-digit searched occupation. “Good German skills” is an indicator for good or very good reported German, relative to lower proficiency, for candidates appearing in searches in German-speaking cantons. The regression also controls for good German skills in searches in non-German-speaking cantons. Standard errors, in parentheses, are clustered two-way on the recruiter and the candidate. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table F.1, Panel A, Column 3).¹⁸ Third, when we use a more concrete indicator of recruiter interest in a candidate, a click to print, email, or favorite the profile, the mover gap rises in relative terms (Appendix Table F.1, Panel A, Column 4). Fourth, the results are virtually identical when we control for the characteristics and qualifications of the job seekers ranked in the immediate vicinity of the focal candidate (Appendix Table F.1, Panel B). This limits the concern that recruiters mainly compare candidates listed next to each other.¹⁹ Finally, the data allow us to estimate the mover gap at a fine-grained occupational level. Appendix Figure F.2 reports the gap separately for each two-digit ISCO occupation. Although recruiters search disproportionately in some occupations, we find economically and often statistically significant mover gaps in almost all of them. The gaps are particularly large for health professionals, information and communication technicians, personal care workers, and handicraft and printing workers.

In sum, the findings show that movers face a significant contact gap of 7% relative to otherwise similar incumbents. The unconditional mover gap, which does not hold candidate quality constant, is twice as large, indicating that movers are negatively selected. In the next sections, we test the determinants of the mover gap suggested by our conceptual framework.

4.3 Job Requirements Overlap and the Mover Gap

Our conceptual framework suggests that the primary signal a recruiter reads from a candidate’s last occupation is how much that occupation’s job requirements overlap with those of the vacancy. It points to two determinants of the mover gap. The first is the within-occupation homogeneity of the vacancy occupation, which governs how much being an incumbent signals about meeting the vacancy’s requirements. The second is the expected overlap $x'\bar{x}$ between the candidate’s last occupation and the vacancy, which determines how many of the vacancy’s requirements a mover can be expected to meet.

Figure 4 tests both predictions by enriching the regression model in Equation (4) with the overlap measure developed in Section 3.3. Each panel is based on a separate regression. Panel (a) reports the mover gap by bins of the within-occupation overlap of the searched occupation, obtained by interacting the mover indicator with the bin indicators. The interaction compares the mover gap across searches in occupations of different homogeneity. Consistent with our framework, the mover gap is much larger in homogeneous occupations—occupations in which jobs share many requirements. Since incumbents’

¹⁸This behavior is unlikely to be common, because non-registered recruiters do not see a candidate’s name and because recruiters spend, on average, less than ten seconds viewing a job seeker’s profile before making a decision.

¹⁹We control for the leave-one-out mean of each job seeker characteristic and of the mover indicator across a centered window of one, three, or five candidates listed before and after the focal candidate. These controls test whether the immediate neighborhood matters beyond the average competitor quality that the search fixed effects hold constant.

previous jobs likely required most of the position’s skills, such occupations are hard to enter for movers. In contrast, the mover gap is smaller where jobs share few requirements and even incumbents’ previous jobs often required different skills. Focusing on the end-points of Panel (a), the mover gap is about 3.3 times larger in the most homogeneous occupations than in the most heterogeneous occupations. These results indicate that job requirements vary within occupations and that recruiters act on this variation. They are in line with other empirical evidence that jobs vary within occupations, including the high relevance of job titles for job seekers’ application behavior (Marinescu & Wolthoff, 2020) and the predictive power of task variation in explaining earnings variation within occupations (Autor & Handel, 2013).

Panel (b) of Figure 4 tests whether the mover gap reflects the shortfall in overlap between the candidate’s last occupation and the recruiters’ searched occupation. Motivated by the framework,²⁰ we measure this shortfall in terms of movers’ overlap gap to incumbents: the within-occupation overlap of the searched occupation minus the overlap of the candidate’s last occupation with the searched occupation, expressed in percentage points. We then add indicators for bins of the overlap gap to the baseline specification and compute the mover gap per bin as the sum of the mover coefficient and the bin coefficient. The overlap gap varies across the movers in the same search, because they come from different occupations.


The figure provides strong evidence that the overlap gap shapes firms’ decisions to contact movers. The contact gap shrinks steadily as a candidate’s overlap approaches the searched occupation’s within-occupation overlap. The gap is only 1.4% of the baseline contact rate in the rightmost bin, which contains movers who, in expectation, meet nearly the same shares of the vacancy’s job requirements as incumbents. This estimate is 80% smaller than the average mover gap, suggesting that job requirements overlap can account for four-fifths of the average gap.

The results are robust to weighting the overlap measure by inverse document frequency, which gives less weight to generic requirements and more weight to rare, distinctive ones. Appendix Figure F.4 and Appendix Table F.3 report the weighted results. Across the pairs of four-digit occupations in the estimation sample, the correlation between the weighted and the unweighted measure is 0.91. Under this weighted measure, a lack of overlap accounts for the entire mover gap: movers with no overlap gap face no detectable contact gap.

These results suggest that overlap matters for the mover gap. However, while the results show that overlap mediates the mover gap, they do not establish causality. A natural concern is that job requirements overlap proxies for other characteristics of the

²⁰In the framework, the occupation component of the mover gap is proportional to the difference between an incumbent’s and a mover’s expected overlap with the vacancy (Equation (3)). The framework predicts that the occupation component of the mover gap vanishes as this difference approaches zero.

Figure 4. Recruiter clicks: the contact gap for movers, by within-occupation homogeneity and by the overlap gap to the searched occupation



Notes: The figure plots the contact gap for movers, as a percent of the baseline contact rate, estimated from Equation (4) enriched with the job requirements overlap measure. Panel (a) interacts the mover indicator with bins of the within-occupation overlap of the searched occupation, in percentage points. Panel (b) adds bins of the overlap gap to incumbents, defined as the within-occupation overlap of the searched occupation minus the overlap of the candidate's last occupation with it, in percentage points. The dashed line marks the average mover gap from the baseline regression of Section 4.2. Both regressions include candidate, recruiter-search, absolute-rank, and relative-rank fixed effects. Error bars show 95% confidence intervals based on standard errors clustered two-way on the recruiter and the candidate. Appendix Table F.3, Columns (1) and (3), reports the underlying coefficients.

searched occupation that both influence the gap and are correlated with overlap. To limit such concerns, we estimate the gap jointly against four characteristics of the searched occupation that could independently shape firms’ willingness to contact movers: the tightness in recruiters’ searched occupation, whether the occupation requires a license, its wage level, and the typical education of incumbents.²¹

Appendix Figure F.3 reports how the mover gap varies with the overlap gap and the four vacancy characteristics. Appendix Table F.2, Panel A, decomposes each dimension’s effect into the part that survives conditioning on the other four and the part they explain, following Gelbach (2016). The results show that the overlap gap remains the most important dimension when the others are held fixed. A one standard deviation increase in the overlap gap widens the mover gap by 3.5% of the baseline contact rate on its own and by 3.9% once the other dimensions enter the regression. The only dimension with a comparable moderating effect is labor market tightness, which we analyze in detail in Section 5. Beyond these two, only occupational licensing displays a noticeable association: the mover gap is 45% larger in licensed occupations, in line with prior evidence that licensing matters for occupational transitions (Kleiner, 2000; Kleiner & Krueger, 2013; Blair & Chung, 2019; Kleiner & Xu, forthcoming; Johnson, 2026). Below, we also show that overlap remains a key mediator of the mover gap when only comparing movers with the same occupational experience and credentials.

In sum, these results limit the concern that the relationship between overlap and the mover gap reflects other characteristics of the searched occupation. To test whether these patterns extend to hiring outcomes beyond the platform, we next replicate our analyses using the application diary data, which record whether an application from a registered job seeker resulted in a hire (see Section 3.2).

The analyses using these data follow the empirical design used so far as closely as possible. The unit of observation is an application to a vacancy. The outcome is an indicator that is one if the job seeker is subsequently hired by the firm she applied to. We regress the hiring indicator on the mover indicator and the overlap measures. Instead of search fixed effects, we control for vacancy occupation, canton, and hiring-firm fixed effects. This implies that we compare applications received by the same firm. We also show results using vacancy fixed effects. Because we observe relatively few candidates with applications in different occupations, we replace the job seeker fixed effects from our Job-Room regressions with a rich set of job seeker characteristics.²²

²¹We measure tightness at the occupation level as explained in Appendix D. The measure is the stock of online vacancies in an occupation and region relative to the contemporaneous number of registered unemployed job seekers who last worked in that occupation and are willing to work in the region. Appendix E also describes how licensing, the wage level, and the typical education of incumbents are constructed.

²²The job seeker controls comprise fixed effects for gender, indicators for being at least 35 and for being at least 50 years old, having children and its interaction with gender, non-Swiss nationality, indicators for primary education and for university education, so that the omitted category is vocational or secondary

The diary analyses depart from the platform design in two respects. First, because a hire is a rare outcome—the mean hire rate is 1.2%—we estimate logit regressions rather than linear probability models. Since logit coefficients do not translate directly into differences in hire rates, we report average marginal effects, expressed as a percent of the mean hire rate. Second, because candidate controls replace the job seeker fixed effects, the estimates rely on a selection-on-observables assumption: they compare applications by movers and incumbents who apply to the same hiring firm for a vacancy in the same occupation and canton, holding constant the applicants’ observable characteristics. But selection on observables is less plausible in the diary data than on Job-Room because similar applicants can differ along dimensions unobserved to us but observed by recruiters, both at the application and later recruitment stages. Since our results suggest that movers are negatively selected, the mover gaps in hires likely combine the occupation component with part of a negative selection term, overstating the occupation component in absolute value. For the same reason, we abstain from interpreting the estimated coefficients causally.

The results based on the diary data are presented in Table 3. Columns (1)–(3) test the relevance of within-occupation overlap and Columns (4)–(6) assess the relevance of the overlap gap. We present specifications with three sets of fixed effects: Columns (1) and (4) compare applications to vacancies in the same occupation and canton. Columns (2) and (5) additionally compare applications received by the same firm. Columns (3) and (6) compare applicants to the same vacancy using vacancy fixed effects. Below, we focus on Columns (2) and (5) because the vacancy fixed effects in Columns (3) and (6) leave no identifying variation from vacancies with a single applicant or without a hire, which makes these estimates relatively imprecise. In any case, the estimated slopes are similar across specifications, suggesting that omitting the vacancy fixed effects does not materially bias the effects of interest.

The table first confirms that applications from movers are significantly less likely to result in a hire than applications from observationally similar incumbents. Although the standard errors are quite large, the estimated mover gaps in hires are substantially larger in magnitude than on the platform. This is consistent with the estimates reflecting part of the selection term. It may also reflect that they measure mover gaps in hiring rates and thus over more screening stages than the Job-Room estimates, or that firms are more reluctant toward movers when they receive applications than when they actively search for candidates on Job-Room.

Second, mover gaps in hiring vary with overlap in the same way as contact gaps on

education, an indicator for more than three years of experience in the last job, and the employment history, namely indicators for being employed one and two years earlier, log labor income one and two years earlier, and indicators for being registered as unemployed one and two years earlier. Log labor income is set to zero for job seekers who were not employed in the respective year, so that their level is absorbed by the corresponding employment indicator.

Job-Room. Column (2) shows that the mover gap decreases significantly with the degree of heterogeneity in the vacancy occupation. The hiring gap for movers in occupations at the 90th percentile of homogeneity is 61% of the mean hire rate, 1.7 times as large as the gap of 35% for movers in occupations at the 10th homogeneity percentile. Column (5) shows that the gap also decreases substantially with the expected overlap between the vacancy occupation and the job seeker’s last occupation: applicants at the 90th percentile of the overlap gap face a hiring gap of about 64%, almost twice the gap of about 34% for applicants at the 10th percentile.

Table 3. Application diaries: the hiring gap for movers by overlap

Dependent Var.:	Within-occupation overlap			Overlap gap to incumbents		
	(1)	(2)	(3)	(4)	(5)	(6)
Find job at firm	Find job at firm	Find job at firm	Find job at firm	Find job at firm	Find job at firm	Find job at firm
Mover (different occupation)	-0.4139*** (0.0672)	-0.5122*** (0.0840)	-0.5470*** (0.1498)	-0.2479*** (0.0870)	-0.3268*** (0.1116)	-0.3399** (0.1713)
Mover x within-occupation overlap (pp)	-0.0651** (0.0271)	-0.0662** (0.0331)	-0.0540 (0.0490)			
Overlap gap to incumbents (pp)				-0.0671*** (0.0225)	-0.0748*** (0.0284)	-0.0829** (0.0416)
Vacancy occupation x canton	X	X		X	X	
Firm FE		X			X	
Vacancy posting			X			X
Job seeker characteristics	X	X	X	X	X	X
Observations	163,629	163,629	166,029	163,629	163,629	166,029
Pseudo R2	0.14520	0.40263	0.81106	0.14546	0.40294	0.81125
Base hire rate	0.0124	0.0124	0.0124	0.0124	0.0124	0.0124
Predicted gap at 10th percentile (pp)	-0.333	-0.439	-0.506	-0.330	-0.419	-0.434
Predicted gap at 90th percentile (pp)	-0.676	-0.756	-0.756	-0.694	-0.787	-0.828

Notes: The table reports coefficients from logit regressions using the application diary data. Each observation is an application by a registered job seeker to a hiring firm. The outcome is whether the application resulted in the job seeker being hired by that firm. Columns (1) to (3) interact the mover indicator with the within-occupation overlap of the vacancy occupation, centered at its sample mean. Columns (4) to (6) add the overlap gap to incumbents to the mover indicator; the overlap gap is zero for incumbents. Every column includes the job seeker controls: gender, indicators for being at least 35 and at least 50 years old, having children and its interaction with gender, non-Swiss nationality, indicators for primary and for university education, an indicator for more than three years of experience in the last job, indicators for being employed one and two years earlier, log labor income one and two years earlier, and indicators for being registered as unemployed one and two years earlier. Columns (1) and (4) include a vacancy-occupation-by-canton fixed effect, Columns (2) and (5) add a hiring-firm fixed effect, and Columns (3) and (6) replace both with a vacancy fixed effect. The last two rows report the predicted hiring gap for movers, in percentage points, at the 10th and the 90th percentile of the overlap measure among mover applications, evaluated at the average hire rate of incumbents. Standard errors, in parentheses, are clustered two-way on the job seeker and the hiring firm. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Overall, we find strong evidence in the platform data, and consistent patterns in the application diary data, that the lower likelihood of meeting job requirements is a main reason why recruiters are less likely to contact—and ultimately hire—movers. These findings support the basic premise of our conceptual framework. More generally, they are consistent with skill-based hiring: the mover gap moves with job requirements overlap and becomes close to zero on the platform for movers whose overlap matches incumbents’. This suggests that recruiters use occupation as a signal of underlying skills rather than

as a credential in itself.

4.4 Occupational Experience, Credentials, and the Mover Gap


Job requirements overlap does not exhaust the signals of occupation-specific fit that recruiters observe. As discussed in Section 2, candidate profiles contain two further signals: prior work experience in the searched occupation and a credential in it, that is, a professional degree or training in that occupation. These signals may raise contact rates beyond what overlap alone predicts.

We therefore estimate contact gaps separately for nine mover profiles. The profiles combine three groups defined by occupational experience and credentials (no experience in the searched occupation, experience without a credential, and experience with a credential) with three categories of the overlap gap to incumbents. The full regression is presented in Column (4) of Appendix Table F.4. Figure 5 reports the contact gap for each mover profile relative to an average incumbent in the same search.

The figure first confirms the role of overlap in closing the mover gap. Within each experience-credential group, the mover gap shrinks strongly from distant to very close occupations. Within the group without experience, for instance, the gap drops from 17.5% at distant occupations to 3.0% at very close ones (profiles 1 and 3). Second, the figure shows that experience and credentials also moderate the mover gap. For instance, prior occupational experience reduces the gap for movers from distant occupations from 17.5% to 10.5% (profiles 1 and 4). This suggests that return movers—those with prior work experience in the searched occupation—face smaller gaps relative to incumbents. Return movers with a credential in the searched occupation (profiles 7 to 9) have the lowest mover gaps. In fact, for a return mover with a credential who comes from a very close occupation, the contact gap is 0.2% and statistically indistinguishable from zero (profile 9).

The right side of the figure reports the share of each mover profile, both in the Job-Room data (measured among mover candidates in recruiters' search results) and in the application diary data (measured among applications from movers). The diary data are likely a better approximation of how movers are distributed across profiles in the broader labor market. The reason is that candidates on Job-Room primarily list occupations in which they have work experience, which means that job seekers' occupational search scope is narrower on Job-Room than in the application diaries (see the discussion in Section 3.1). Indeed, the diary shares reveal that many movers attempt relatively distant switches: slightly more than a quarter of mover applications go to positions in distant occupations in which the applicant has no prior work experience. Return movers with a degree who come from very close occupations—the profile facing a zero contact gap on Job-Room—account for just 2% of mover applications.

Figure 5. Recruiter clicks: the contact gap for nine mover profiles



Notes: The figure reports the contact gap, as a percent of the baseline contact rate, for nine mover profiles. The profiles combine three experience-and-credential groups with three overlap categories: distant (an overlap gap to incumbents of at least 2.5 points), close (0.5 to 2.5 points), and very close (less than 0.5 points). Each gap is relative to an average incumbent. The last two rows show the average mover gap from the baseline regression and from a reweighted version that aligns the estimation sample with the distribution of mover profiles and two-digit vacancy occupations in the application diary data. The right-hand columns report each profile's share among movers in the Job-Room result lists and in the diary applications. All estimates come from Equation (4) with candidate, recruiter-search, absolute-rank, and relative-rank fixed effects. Error bars show 95% confidence intervals based on standard errors clustered two-way on the recruiter and the candidate. The underlying regression is Appendix Table F.4, Column (4).

The differences between the two fractions in Figure 5 also raise the question of whether the mover gaps estimated from the Job-Room data reflect those in the broader labor market. Does the platform’s lower share of movers from distant occupations imply that the platform estimates understate the average mover gap? A second difference works in the opposite direction: the platform overrepresents searches in occupations that are hard for movers to enter, such as craft and trade occupations (compare Appendix Table A.1 and Appendix Figure F.2). To quantify the net effect of these two opposing effects, the bottom rows of Figure 5 compare the baseline mover gap against a reweighted mover gap. The weights align the joint distribution of the platform observations over the nine mover profiles and two-digit vacancy occupations with the distribution of applications in the application diaries. In percentage points, the reweighted mover gap is close to the baseline estimate (-0.65 versus -0.77 percentage points). However, average contact rates are lower when using the diary data weights (0.06 instead of 0.11). Relative to this lower average contact rate, the reweighted mover gap amounts to 10.4%, compared with 7.0% in the baseline. These results suggest that the platform data slightly understate the relative disadvantage that movers face in the wider labor market because they underrepresent distant movers. We thus report the reweighted mover gap as our headline estimate of the mover gap for registered job seekers in Switzerland.

Together, the results in this section suggest that job requirements overlap, occupational experience, and occupation-specific credentials are the primary reasons why recruiters prefer incumbents over otherwise comparable occupational switchers. The mover gap thus primarily reflects a perceived lack of skills and credentials, supporting the assumption in existing models of occupational mobility that human capital is partly occupation-specific. The results are also informative about the relative weight that firms place on different indicators of human capital and occupational suitability. For instance, recruiters value work experience and professional degrees above and beyond their implications for the overlap between a job seeker’s skills and the position’s requirements. Finally, the findings imply that firms impose costs on the large majority of job seekers willing to move to other occupations.

5 Tightness and the Role of Firms

The determinants of the mover gap examined so far operate through the first element of the screening decision: the signals that shape how many of the vacancy’s requirements a candidate is expected to meet. The decision has a second margin in our framework: the search-specific threshold τ_s that a candidate must clear. As discussed in Section 2, a likely force moving τ_s is the tightness of the vacancy occupation’s labor market: when workers from the occupation itself are scarce, recruiters who need to fill the position may relax their screening standards. A lower τ_s raises the chances for movers and incumbents alike,

but many incumbents likely are above the original bar and would be contacted at either threshold. As a result, the lower bar may mainly increase movers’ chances (Section 2).

We first show how the composition of the candidate pool shifts when incumbents become scarce (Section 5.1). Section 5.2 then documents, in both data sources, how firms’ screening of movers changes as occupations tighten. Finally, Section 5.3 combines both results in a decomposition that attempts to quantify the firms’ importance for occupational transitions.

5.1 The Candidate Pool in Tight Markets

This section analyzes how incumbent scarcity affects the mover gap. We do so by relating the mover gap to an occupation- and region-specific measure of tightness among incumbents. The measure is the stock of online vacancies in an occupation and region relative to the contemporaneous number of registered unemployed job seekers who last worked in that occupation and are willing to work in the region. Appendix D describes the measure’s construction in detail.

Before turning to recruiters’ decisions, we analyze what a scarcity of incumbents implies for the number of incumbents and movers in recruiters’ candidate pools in our datasets. To quantify this, we move from the candidate level to the job level. On the platform, we count the number of incumbent and mover profiles each search returns. In the diary data, we count the number of incumbents and movers applying to each vacancy. We then estimate the following Poisson regression for the count n_u^g of candidates of type $g \in \{I, M\}$, separately for incumbents ($g = I$) and movers ($g = M$):

$$\mathbb{E}[n_u^g] = \exp\left(\beta_{\text{mid}}^g T_u^{\text{mid}} + \beta_{\text{tight}}^g T_u^{\text{tight}} + \text{FE}_u\right), \quad (5)$$

where u indexes recruiter searches on the platform and vacancies in the diary data. The indicators T_u^{mid} and T_u^{tight} equal one if $\text{tightness}_{o(u),c(u),t(u)}$ —the tightness of the occupation and region of search or vacancy u at the time the search takes place—is in the middle or the top tercile, respectively, with the bottom tercile (slack occupations) as the reference. The Poisson coefficients measure proportional shifts in the number of candidates of movers and incumbents, in log points. In the Job-Room data, the fixed effects FE_u comprise the calendar month, recruiters’ searched occupation interacted with the set of searched cantons, and a set of recruiter fixed effects. These fixed effects imply that the tercile effects are identified from temporal variation in tightness within occupation-region cells. In the diary data, where the unit of observation is a vacancy, the fixed effects comprise the calendar month and the hiring firm.²³ Standard errors are clustered on the recruiter

²³Note that in this specification we do not add vacancy-occupation-by-canton fixed effects because the application diary data have a limited time dimension. The variation in tightness left after controlling for occupation and location is too small to identify effects.

in the search-level platform regression and on the hiring firm in the vacancy-level diary regression.

Figure 6 presents the results of those regressions using the platform data (Panel a) and the application data (Panel b). The figure first validates the relevance of the external tightness measure. Incumbents become scarce as markets tighten. In the diary data, for instance, the number of incumbent applicants per vacancy falls by about 1.1 log points, or roughly two-thirds, from the slackest to the tightest tercile. Since tightness is the ratio of vacancies to unemployed incumbents, this pattern is expected—it confirms that our tightness measure captures a search- and vacancy-level scarcity of incumbents. Second, and not expected by construction, the number of movers responds much less to incumbent scarcity. On Job-Room, it does not change across terciles (Panel a). In the application diaries, mover applications per vacancy fall by roughly a fifth from the slackest to the tightest tercile (Panel b). Tight occupations thus do not attract a compensating inflow of movers. This is the opposite of what we would expect if movers applied where their chances are better.²⁴

Increased occupational tightness thus confronts recruiters with a pool that contains far fewer incumbents and a similar or moderately smaller number of movers. The mover share rises because more incumbents drop out than movers.

5.2 Screening in Tight Markets


How does a scarcity of incumbents affect recruiter behavior? To answer this question, we return to the candidate-level regressions estimated earlier. We augment the baseline model (Equation (4)) with interactions between the mover indicator, $\text{Mover}_{i,s}$, and indicators for terciles of tightness in the searched occupation and region at the time of the search. Because tightness does not vary within a search, its level effect is absorbed by the search fixed effects ϕ_s . The heterogeneity in the mover gap is identified from temporal variation in tightness within an occupation and region as well as from cross-sectional variation across occupations and regions. Panel (a) of Figure 7 also shows an estimate that adds searched-occupation-by-mover and search-location-by-mover fixed effects, which focuses on the temporal variation.²⁵

Panel (a) of Figure 7 reports the resulting mover gaps in the slackest and the tightest tercile. We report an average effect as well as effects estimated separately for movers whose last occupation overlaps closely with the vacancy occupation (an overlap gap to incumbents of at most 2 percentage points) and for movers with low overlap (a gap above

²⁴On the platform, the occupations in which a job seeker is willing to work—and thus the searches in which her profile appears—are fixed at registration with the unemployment office and change rarely over the spell, which makes the composition of movers much stickier there in the first place.

²⁵The sample window of the application diary data is too short to allow us to do such an interaction design in those data.

Figure 6. The candidate pool by tightness in an occupation



Notes: The figure reports Poisson estimates of Equation (5): the change in the number of incumbents and movers across tightness terciles, in log points relative to the slack (bottom) tercile. In Panel (a), the unit of observation is a recruiter search on Job-Room, and the outcomes are the numbers of incumbent and mover profiles the search returns. The regressions include calendar-month, vacancy-occupation-by-location, and recruiter fixed effects, with standard errors clustered on the recruiter (Appendix Table F.6, Panel A, Columns (1) and (2)). In Panel (b), the unit of observation is a vacancy in the application diary data, and the outcomes are the numbers of incumbent and mover applications the vacancy receives. The counts include the applications of all job seekers in the diary sample, not only of those with a completed spell (see Appendix B). The regressions include calendar-month and hiring-firm fixed effects, with standard errors clustered on the hiring firm (Appendix Table F.6, Panel B). Error bars show 95% confidence intervals.

2 points). Panel (b) replicates the analysis in the application diary data. The specification mirrors the logit regressions from Section 4.3: we regress the hiring indicator on the mover indicator, the tightness terciles, and their interactions, with hiring-firm fixed effects and the job seeker controls. As before, we report average marginal effects in this panel, expressed as a percent of the mean hire rate.

Panel (a) of Figure 7 shows that occupational tightness increases recruiters' willingness to contact candidates from other occupations. Averaged over all movers, the mover gap falls from 8.8% of the baseline contact rate in the slackest tercile of occupations to 3.4% in the tightest. The decline is almost identical when we add searched-occupation-by-mover and search-location-by-mover fixed effects (second row of Panel (a)). Since this specification identifies the change in the gap only from changes in tightness within occupation-region cells over time, the similarity of the two estimates limits the concern that the tightness results reflect fixed characteristics of occupations or regions that are correlated with tightness.

The bottom part of the figure suggests that incumbent scarcity benefits both close and distant movers. For candidates from a closely related occupation, the gap falls from 5.6% to 1.8% and is no longer distinguishable from zero at the 5% level. Movers from distant occupations start from the largest gap and gain the most in absolute terms: their

gap falls from 13.0% to 3.8% of the baseline rate.

The diary data suggest a similar, albeit less pronounced, pattern for the average effect (Panel b), with imprecise but economically meaningful interaction effects. The gap in the tightest tercile is roughly half the gap in the slackest tercile, a difference that is significant at the 10% level. The narrowing is concentrated among movers from closely related occupations, whose gap in tight markets is no longer distinguishable from zero. For movers from distant occupations, the gap narrows only slightly.²⁶

Figure 7. The mover gap by tightness in an occupation and overlap with the vacancy



Notes: The figure plots the mover gap, as a percent of the baseline rate, in slack (bottom tightness tercile) and tight (top tightness tercile) occupations, from regressions that interact the mover indicator with the tightness terciles. The middle tercile enters the regressions but is not plotted. The gaps are shown for all movers and separately for movers with high overlap (an overlap gap to incumbents of at most 2 percentage points) and low overlap (more than 2 points), each group estimated in a separate regression. In Panel (a), the gap is the contact-rate gap relative to incumbents in the same search, estimated from Equation (4) with candidate, recruiter-search, absolute-rank, and relative-rank fixed effects. Standard errors are clustered two-way on the recruiter and the candidate (Appendix Table F.5, Panel A). The second row of Panel (a) repeats the estimate for all movers, adding searched-occupation-by-mover and search-location-by-mover fixed effects, so that the change in the gap is identified only from temporal variation in tightness within occupations and within locations (Appendix Table F.5, Panel A, Column 2). Since this regression only identifies the difference between the slack and the tight searches, but not the gap itself in the slack market, we set the gap in slack markets to the estimate from the first row for this plot. In Panel (b), the gap is the hire-rate gap relative to incumbents in the application diary data. It comes from logit models with hiring-firm fixed effects and the job seeker controls listed in the notes to Table 3, reported as average marginal effects. Standard errors are delta-method standard errors clustered two-way on the hiring firm and the job seeker (Appendix Table F.5, Panel B). Error bars show 95% confidence intervals.

Overall, the results in this section show that labor market tightness is associated with a reduced reluctance of firms to hire movers. These results are in line with a

²⁶The framework provides a possible explanation for why close and distant movers benefit differently in the two data sources. A lower threshold mainly helps candidates who are just below it. Relative to the baseline rate, both groups of movers face much smaller gaps in contacts than in hires (Section 4.3). In the contact decision, even distant movers may thus sit just below the threshold, whereas in the hiring decision mainly close movers do.

nascent literature showing that firms lower their hiring standards in tight labor markets (Modestino et al., 2016, 2020; Lochner et al., 2021; Le Barbanchon, Ronchi, & Sauvagnat, 2023; Ziegler, 2024; Cullen et al., 2025). Our findings point to an additional dimension of how recruiters react to tightness. They also suggest that tightness in an occupation creates opportunities for job seekers to transition into it. Finally, they suggest that the mover gap partly reflects how selective firms can afford to be, not only how well movers fit the position.

5.3 Firms' Contribution to Occupational Transitions

The previous sections suggest that occupational tightness changes both the share of movers among the candidates firms see and movers' chances of being contacted or hired conditional on appearing in a recruiter's result list or applying to a vacancy. Both changes make firms more likely to hire a mover. We now gauge firms' contribution to movers' improved chances relative to incumbents when an occupation tightens, both on JobRoom and in the application data. These results inform the more general question of how important firms are for occupational transitions.

To measure how much firms' behavior affects movers' chances relative to incumbents, we do not resort to the mover gap studied so far, which is conditional on being in the candidate pool. Whether a job goes to a mover depends also on the number of movers in the candidate pool. A firm that treats movers in tight and slack occupations equally will still hire more movers when incumbents become scarce. Conversely, a firm that becomes more open to movers in a tight occupation contributes little to occupational transitions if hardly any movers are in the pool.


Our object of interest in this section is thus the number of contacted movers minus the number of contacted incumbents, both expressed as a share of the candidate pool. Formally, contacted movers make up a share $s_M(T)p_M(T)$ of the candidate pool, where $s_M(T) = \mathbb{P}(\text{Mover} = 1 \mid T)$ is the share of movers in the pool and $p_M(T) = \mathbb{P}(\text{contact} = 1 \mid \text{Mover} = 1, T)$ is the contact rate among movers, both at tightness T . Contacted incumbents make up a share $[1 - s_M(T)]p_I(T)$, where $p_I(T)$ is the contact rate among incumbents. We call the difference between the two shares the mover gap in contacts,

$$G(T) = \underbrace{s_M(T)p_M(T)}_{\text{a mover is contacted}} - \underbrace{[1 - s_M(T)]p_I(T)}_{\text{an incumbent is contacted}}. \quad (6)$$

In the diary data, the same expression defines the mover gap in hires, with $s_M(T)$ the mover share of applications and $p_M(T)$ and $p_I(T)$ the hire rates conditional on applying. $G(T)$ is positive when the share of movers in the pool is sufficiently high to make up for the conditional mover gap in contacts or hires. $G(T)$ changes both when firms change how they treat movers and when the pool changes.

The two top bars in Figure 8 report $G(T)$ in the tightest and the slackest tercile of occupations. To make the two data sources comparable, we divide $G(T)$ by the probability that a given profile is an incumbent and is contacted (Panel a), or that a given application is from an incumbent and leads to a hire (Panel b).²⁷ The figure shows that in both settings, tightness shifts the outcome toward movers. On the platform, the mover gap in contacts narrows from -36.1% in slack to -20.9% in tight markets. In the diaries, the mover gap in hires rises from 2.5% to 152.6%. The opposite signs of $G(T)$ in the two data sources arise because movers make up about half of the profiles a search returns but the majority of applications in the diary data, both in slack and tight markets. Their higher number implies that the mover gap in hires is positive despite movers' lower hire rate.

Figure 8. The mover gap in contacts and in hires, in tight and slack occupations



Notes: The figure presents the results of decomposing the change in the mover gap in contacts (Panel a) and in hires (Panel b) between a tight and a slack occupation into the firm response and the change in pool composition based on Equation (7). The mover gap in contacts and the mover gap in hires are defined in Equation (6). In Panel (a), the pool is the candidate profiles a recruiter search returns on Job-Room. In Panel (b), it is the applications a vacancy receives in the diary data. The top two bars give the gap in the slackest and tightest tercile of the tightness of the vacancy occupation. The gap is expressed as a percent of the probability of drawing a contacted incumbent (Panel a) or a hired incumbent (Panel b). The bottom bar presents the results of the decomposition, highlighting the share of the change attributed to the pool composition and the firm response. The bar omits the third term, which represents 1.4% of the shift in contacts and 8.5% of the shift in hires. The underlying coefficients come from the models behind Figures 6 and 7.

Our goal now is to decompose the change in $G(T)$ between slack and tight occupations into the contributions of firms (labor demand) and job seekers (labor supply). To isolate the firm contribution, we follow Oaxaca (1973) and Blinder (1973). Their decomposition splits a difference in outcomes between two groups into a part due to differences in characteristics and a part due to differences in how these characteristics are rewarded. Here the two groups are slack and tight markets, the characteristic is the composition of

²⁷We use one common denominator per outcome for both tightness terciles and all terms of the decomposition. The terms therefore still add up exactly to the change in the gap.

the pool, and the reward is the rate at which firms contact or hire each type of candidate.²⁸ The decomposition splits the change in $G(T)$ into three terms,

$$\Delta G = \underbrace{\bar{s}_M \Delta(p_M - p_I)}_{\text{firm response}} + \underbrace{(\bar{p}_M + \bar{p}_I) \Delta s_M}_{\text{pool composition}} + \underbrace{(2\bar{s}_M - 1) \Delta p_I}_{\text{common rate shift}}, \quad (7)$$

where Δ is the difference between a tight and a slack market and a bar denotes the midpoint of the two.

The first term in (7) is firms' response—the object of interest. It is the tightness-induced change in the mover gap conditional on being in the candidate pool, $\Delta(p_M - p_I)$. It is multiplied by the average share of movers in the pool, $\bar{s}_M$, because a change in how firms treat movers matters only for the movers who are actually in the pool. The second term is the change in the pool's composition. When incumbents become scarce, the mover share s_M rises according to our estimates in Section 5.1.²⁹ The third term is a shift in the contact rate common to movers and incumbents that matters whenever the pool is not split evenly between movers and incumbents. A rise in the overall contact rate raises $G(T)$ when movers are the majority and lowers it when they are the minority.

We draw on our previous estimates to implement the decomposition in Equation (7) empirically. The change in the conditional mover gap induced by tightness comes from the mover-by-tightness interactions in the regressions behind Figure 7. The change in the mover share comes from the applicant-count regressions behind Figure 6. The midpoint shares and rates are sample averages.³⁰

The bottom bar of each panel in Figure 8 illustrates the results of the decomposition for the mover gap in contacts (Panel a) and in hires (Panel b). In both settings, most of the shift toward movers comes from the higher mover share in the applicant pool in tight occupations. But the firm contribution to the shift toward movers is sizable. Firms' screening response accounts for 28.7% of the rise in the mover gap in contacts and for 29.8% of the rise in the mover gap in hires. The remaining 1.4% of the shift in

²⁸We evaluate each change at the midpoint between the slack and the tight market. The decomposition then does not depend on which market serves as the reference. Rewrite the gap as $G = (2s_M - 1)p_I + s_M \delta$, where $\delta = p_M - p_I$ is the mover gap conditional on being in the candidate pool, as estimated in Section 5.2. Applying the midpoint product rule $\Delta(xy) = \bar{x}\Delta y + \bar{y}\Delta x$ to each product yields Equation (7). A bar denotes the midpoint of the slack and the tight value, and Δ denotes the difference between them.

²⁹The weight on Δs_M is the sum of the two rates because each candidate who shifts from the incumbent side of the pool to the mover side adds to the probability of drawing a contacted mover at the mover rate and subtracts from the probability of drawing a contacted incumbent at the incumbent rate.

³⁰The corresponding estimates appear in Column (1) of Appendix Table F.5, Panels A and B, and in Columns (1) and (2) of Appendix Table F.6, Panels A and B. The change in the incumbent contact rate from the third term cannot be taken directly from our previous estimations because the search fixed effects absorb the baseline tightness effect in the platform regression. We therefore estimate it in a separate regression of incumbents' contact rates on the tightness terciles, with the same fixed effects as the applicant-count model. These estimates are presented in Column (3) of Appendix Table F.6, Panel A. In the diary data we do not have the equivalent of a search fixed effect: the hiring-firm fixed effect does not absorb tightness. Hence, we directly use the estimates from the regression used in Figure 7.

contacts and 8.5% of the shift in hires come from the third term, the common shift in the contact rate.³¹ The results suggest that firms’ lower reluctance to hire movers contributes substantially to movers’ increased chances of finding a job in tight occupations.

Interpreting these changes as firms’ contribution requires that the change in the mover gap conditional on being in the candidate pool does not reflect a change in movers’ quality. This is plausible in our setting, especially in the platform data. There, we identify the change in the gap from the same candidate appearing in searches of different tightness. The candidate fixed effects absorb time-invariant differences in quality between the movers who appear in tight and in slack searches.³² In the diary data, we hold observable applicant characteristics fixed. Moreover, mover applications per vacancy fall rather than rise in tight occupations, by about a fifth from the slackest to the tightest tercile (Figure 6). In line with evidence that job seekers’ targeted occupations do not react to tightness in feasible alternative occupations (Belot, de Koning, et al., 2025), these results suggest that movers in tight occupations are unlikely to be a substantially different type of job seeker drawn in by better prospects.

Taken together, the decomposition suggests that tight occupations raise movers’ hiring chances mainly because fewer incumbents compete for the same jobs. Firms add to this by becoming more open to movers. Without the firm response, the shift toward movers in tight occupations would be roughly 30% smaller in both data sources.

More generally, the decomposition suggests that firms play a substantial role in shaping workers’ ability to change occupations. Incidentally, the share we attribute to firms is close to the descriptive benchmarks of Section 3. There, we document that movers’ applications are about a quarter less likely to result in a hire than incumbents’ in the diary data (Section 3.2), and that movers’ contact rate on the platform is 21% lower than incumbents’ (Table 1). While these estimates are conceptually distinct and each only approximates firms’ quantitative relevance, all point to an important role for firms in occupational mobility.

6 Conclusion

We analyze whether, why, and when firms are reluctant to hire occupational movers, complementing previous research that focused primarily on workers’ choices. We interpret firms’ role through the lens of a framework of screening under uncertainty: firms prefer incumbents, who last worked in the vacancy occupation, over occupational movers

³¹The weight on this term, $2\bar{s}_M - 1$, grows as the mover share of the pool moves away from one half. On the platform, the term is negligible because movers are close to half of the result lists. In the diary, the term is more important because movers represent about seven in ten applications.

³²In addition, whether a profile appears in a search does not depend on the job seeker’s choices. It depends on the recruiter’s query and on the occupations the job seeker registered with the unemployment office, which rarely change over the spell.

because they do not observe a candidate’s complete skill set. They therefore infer expected fit from the requirements of the candidate’s previous occupation, her experience in the vacancy occupation, and her occupational credentials. Our empirical analyses draw on unemployment insurance records linked with application diaries and clickstream data from a recruitment platform, and we test the framework’s predictions using a new measure of occupational similarity based on overlap in job requirements across vacancies. Together, these data allow us to estimate the causal effect of job seekers’ occupational background on recruiters’ contact decisions as well as to examine the relationship between the mover gap and job requirements overlap.

Five key findings emerge. First, movers are contacted 21% and hired 26% less often than occupational incumbents, conditional on being in the candidate pool. Part of this gap reflects that movers are negatively selected. Holding job-seeker quality constant and reweighting the platform data to the mix of movers and vacancy occupations in the application diaries, the contact gap to incumbents shrinks to 10.4%. Second, movers’ lower job requirements overlap appears to be a core determinant of their disadvantage relative to incumbents. The mover gap falls by 80% if we compare movers and incumbents who, in expectation, meet a similar share of the position’s requirements. Third, exploiting the fact that we can measure job requirements overlap within occupations, we show that the mover gap is 3.3 times larger in the most homogeneous occupations than in the most heterogeneous ones. Fourth, firms adapt to a scarcity of incumbents. The estimated mover gap falls from 8.8% of the baseline contact rate in the slackest occupations to 3.4% in the tightest occupations. Finally, an Oaxaca-Blinder decomposition shows that changes in firm behavior account for about 30% of the relative shift of contacts and hires from incumbents to movers as occupations tighten.

Several limitations qualify some of these conclusions. First, the platform data measure initial contacts rather than the complete hiring process, and recruiters who actively search for candidates on a recruitment platform may be more open to occupational movers to begin with. The diary data reproduce the main patterns for realized hires, but cover a shorter period and yield estimates that are less precise and less plausibly causal. Second, our overlap measure captures the requirements of representative jobs in an occupation, not the skills of an individual worker or the requirements of the exact job she previously held. It may also miss some job requirements, because job ads often omit skills that are implied by the position or by the required degree. Third, our data are unique but not fully representative of the labor market for registered job seekers in Switzerland. On the platform, recruiters search disproportionately in craft and trade occupations, and candidates attempt fewer distant switches than applicants in the diaries, which is why our preferred mover gap estimate reweights the platform estimates. The diary sample is restricted to job seekers who report their applications online and have a known re-employment firm. At the same time, the two sources complement each other. They cover

different stages of the hiring process and different hiring channels, recruiter-initiated search and job seekers' own applications. That the main patterns emerge in both lends credibility to our conclusions.

Some key findings have direct policy implications. Our results, for instance, suggest that occupation labels alone are an inadequate guide to feasible transitions. Our index of job requirements overlap shows that two occupations with different labels sometimes share more job requirements than two jobs within a heterogeneous occupation. Public employment services could combine our vacancy-based overlap measure with information on workers' experience and credentials to identify target occupations in which a particular mover is likely to pass firms' screening thresholds. The results also suggest directing job seekers toward heterogeneous occupations, where incumbency matters less, and toward tight occupations, where firms are more open to occupational movers.

A further implication concerns the information supplied to firms. If recruiters rely on occupation labels because individual skills are difficult to observe, improving the visibility and credibility of skills may reduce the mover gap. Job platforms and application materials could present verified skills or evidence from prior tasks alongside occupational and job histories.

References

- Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In Handbook of labor economics (Vol. 4, pp. 1043–1171). Elsevier.
- Aigner, D. J., & Cain, G. G. (1977). Statistical theories of discrimination in labor markets. Industrial and Labor Relations Review, 30(2), 175–187.
- Alabdulkareem, A., Frank, M. R., Sun, L., AlShebli, B., Hidalgo, C., & Rahwan, I. (2018). Unpacking the polarization of workplace skills. Science advances, 4(7), eaao6030.
- Altmann, S., Glenny, A., Mahlstedt, R., & Sebald, A. (2022). The direct and indirect effects of online job search advice.
- Altmann, S., Mahlstedt, R., Rattenborg, M., & Sebald, A. (2023). Which occupations do unemployed workers target? insights from online job search profiles.
- Autor, D. H., & Handel, M. J. (2013). Putting tasks to the test: Human capital, job tasks, and wages. Journal of labor Economics, 31(S1), S59–S96.
- Azar, J. A., Berry, S. T., & Marinescu, I. (2022). Estimating labor market power.
- Azmat, G., Behaghel, L., Hazard, Y., Rathelot, R., & Sultan, J. (2025). Occupational mobility and retraining: Experimental evidence on firms' hiring preferences. Working Paper.
- Banfi, S., Choi, S., & Villena-Roldán, B. (2019). Deconstructing job search behavior. Available at SSRN 3323545.

- Bannert, M., Klaeui, J., Kopp, D., Siegenthaler, M., Thöni, S., & Winkelmann, J. (2022). Swiss job tracker: New weekly indices of online job postings in switzerland: Methodology (Tech. Rep.). ETH Zurich.
- Behaghel, L., Dromundo, S., Gurgand, M., Hazard, Y., & Zuber, T. (2024). The potential of recommender systems for directing job search: A large scale experiment. IZA Discussion Paper No. 16781.
- Belot, M., de Koning, B. K., Fouarge, D., Kircher, P., Muller, P., & Phlippen, S. (2025). Advising job seekers in occupations with poor prospects: A field experiment (Tech. Rep.). National Bureau of Economic Research.
- Belot, M., Kircher, P., & Muller, P. (2019). Providing advice to jobseekers at low cost: An experimental study on online advice. The Review of Economic Studies, 86(4), 1411–1447.
- Belot, M., Kircher, P., & Muller, P. (2025). Do the long-term unemployed benefit from automated occupational advice during online job search? The Economic Journal, ueaf041.
- Blair, P. Q., & Chung, B. W. (2019). How much of barrier to entry is occupational licensing? British Journal of Industrial Relations, 57(4), 919–943.
- Blinder, A. S. (1973). Wage discrimination: Reduced form and structural estimates. Journal of Human Resources, 8(4), 436–455.
- Blöesch, J., Larsen, B., & Taska, B. (2022). Which workers earn more at productive firms? position specific skills and individual worker hold-up power.
- Bocquet, L. (2024). The network origin of slow labor reallocation (Working Paper). University of Cambridge.
- Bohm, M. J., Etheridge, B., & Irastorza-Fadrique, A. (2024). Heterogeneity of occupational supply elasticities and changes in labour demand. (Working Paper, March 2024)
- Brüll, E. (2025). Matchmaker: Indexing-based record linkage for r [Computer software manual]. Retrieved from <https://github.com/edubruell/MatchMaker> (R package, version available on GitHub at <https://github.com/edubruell/MatchMaker>)
- Brüll, E., Rostam-Afschar, D., & Schlenker, O. (2025). Cut off from new competition: Threat of entry and quality of primary care. Labour Economics, 92, 102669.
- Bächli, M., Lalive, R., & Pellizzari, M. (2025). Helping jobseekers with recommendations based on skill profiles or past experience: Evidence from a randomized intervention (IZA Discussion Papers No. 17713). Institute of Labor Economics (IZA). Retrieved from <https://EconPapers.repec.org/RePEc:iza:izadps:dp17713>
- Colella, F. (2022). The effect of trade on skill requirements: Evidence from job postings (Mimeo). University of Lausanne.
- Cortes, G. M., & Gallipoli, G. (2018). The costs of occupational mobility: An aggregate analysis. Journal of the European Economic Association, 16(2), 275–315.
- Cullen, Z., Hoffman, M., & Koenig, F. (2025). Wage surging: An analysis of firms’ search during hiring difficulties. (Mimeo)

- del Rio-Chanona, R. M., Mealy, P., Beguerisse-Díaz, M., Lafond, F., & Farmer, J. D. (2021). Occupational mobility and automation: a data-driven network model. Journal of The Royal Society Interface, 18(174), 20200898.
- Dhia, A. B., Crépon, B., Mbih, E., Paul-Delvaux, L., Picard, B., & Pons, V. (2022). Can a website bring unemployment down? experimental evidence from france (Working paper). National Bureau of Economic Research.
- Djumalieva, J., Lima, A., & Sleeman, C. (2018). Classifying occupations according to their skill requirements in job advertisements. Economic Statistics Centre of Excellence Discussion Paper, 4, 2018.
- Dupuy, A., Raux, M., & Signorelli, S. (2024). Digitalization, change in skill distance between occupations and worker mobility: A gravity model approach. IZA Discussion Paper 17535.
- Eggenberger, C., & Backes-Gellner, U. (2023). It skills, occupation specificity and job separations. Economics of Education Review, 92, 102333.
- Eggenberger, C., Rinawi, M., & Backes-Gellner, U. (2018). Occupational specificity: A new measurement based on training curricula and its effect on labor market outcomes. Labour Economics, 51, 97–107. doi: 10.1016/j.labeco.2017.11.010
- Gathmann, C., & Schönberg, U. (2010). How General Is Human Capital? A Task-Based Approach. Journal of Labor Economics, 28(1), 1–49. doi: 10.1086/649786
- Gelbach, J. B. (2016). When do covariates matter? and which ones, and how much? Journal of Labor Economics, 34(2), 509–543.
- Goos, M., Rademakers, E., Salomons, A., & Willekens, B. (2019). Markets for jobs and their task overlap. Labour Economics, 61, 101750.
- Güvenen, F., Kuruscu, B., Tanaka, S., & Wiczer, D. (2020). Multidimensional skill mismatch. American Economic Journal: Macroeconomics, 12(1), 210–244.
- Hangartner, D., Kopp, D., & Siegenthaler, M. (2021). Monitoring hiring discrimination through online recruitment platforms. Nature, 589(7843), 572–576.
- Harmon, N., Kristensen, A., Skandalis, D., & Settele, S. (2025). Wage misperceptions and the sorting of young workers to jobs. (Mimeo)
- Herz, B., & Van Rens, T. (2020). Accounting for mismatch unemployment. Journal of the European Economic Association, 18(4), 1619–1654.
- Johnson, J. E. (2026). Occupational licensing in the united states. Journal of Economic Perspectives, 40(1), 167–190.
- Kircher, P. (2022). Job search in the 21st century. Journal of the European Economic Association, 20(6), 2317–2352.
- Klaeui, J. (2026). When to broaden, when to focus: Job search and market tightness (Working Paper).
- Kleiner, M. M. (2000). Occupational licensing. Journal of Economic Perspectives, 14(4), 189–202.

- Kleiner, M. M., & Krueger, A. B. (2013). Analyzing the extent and influence of occupational licensing on the labor market. Journal of Labor Economics, 31(S1), S173–S202.
- Kleiner, M. M., & Xu, M. (forthcoming). Occupational licensing and labor market fluidity. Journal of Labor Economics.
- Kopp, D. (2025). Do recruiters penalize men who prefer low hours? evidence from online labor market data. The Journal of Labor Economics, forthcoming.
- Lazear, E. (2009). Firm-Specific Human Capital: A Skill-Weights Approach. Journal of Political Economy, 117(5), 914–940. doi: 10.1086/648671
- Le Barbanchon, T., Hensvik, L., & Rathelot, R. (2023). How can ai improve search and matching? evidence from 59 million personalized job recommendations (Draft). Bocconi, Uppsala University, Institut Polytechnique de Paris.
- Le Barbanchon, T., Ronchi, M., & Sauvagnat, J. (2023). Hiring difficulties and firms’ growth. SSRN Working Paper No. 4105264.
- Leping, K.-O. (2009). Measuring the specificity of human capital: a skill-based approach. Baltic Journal of Economics, 9(1), 39–54.
- Lise, J., & Postel-Vinay, F. (2020). Multidimensional skills, sorting, and human capital accumulation. American Economic Review, 110(8), 2328–2376.
- Lochner, B., Merkl, C., Stüber, H., & Gürtzgen, N. (2021). Recruiting intensity and hiring practices: Cross-sectional and time-series evidence. Labour Economics, 68, 101939.
- Lyshol, A. (2022). Occupational mobility, mismatch & unemployment insurance. Mismatch & Unemployment Insurance.
- Ma, K., & Samaniego de la Parra, B. (2021). Labor market tightness, recruitment and search behavior (Draft). University of California, Santa Cruz.
- Macaluso, C. (2025). Skill remoteness and post-layoff labor market outcomes. American Economic Journal: Macroeconomics, 17(2), 134–176.
- Marinescu, I., & Wolthoff, R. (2020). Opening the black box of the matching function: The power of words. Journal of Labor Economics, 38(2), 535–568.
- Modestino, A. S., Shoag, D., & Ballance, J. (2016). Downskilling: changes in employer skill requirements over the business cycle. Labour Economics, 41, 333–347.
- Modestino, A. S., Shoag, D., & Ballance, J. (2020, 10). Upskilling: Do Employers Demand Greater Skill When Workers Are Plentiful? The Review of Economics and Statistics, 102(4), 793–805.
- Oaxaca, R. (1973). Male-female wage differentials in urban labor markets. International Economic Review, 14(3), 693–709.
- Phelps, E. S. (1972). The statistical theory of racism and sexism. American Economic Review, 62(4), 659–661.
- Poetaev, M., & Robinson, C. (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.

- Robinson, C. (2018). Occupational mobility, occupation distance, and specific human capital. Journal of Human Resources, 53(2), 513–551.
- Şahin, A., Song, J., Topa, G., & Violante, G. L. (2014). Mismatch unemployment. American Economic Review, 104(11), 3529–3564.
- Schmutte, I. M. (2014). Free to move? a network analytic approach for learning the limits to job mobility. Labour Economics, 29, 49–61.
- Schubert, G., Stansbury, A., & Taska, B. (2024). Employer concentration and outside options. Available at SSRN 3599454.
- Staatssekretariat für Bildung, Forschung und Innovation. (2021, September). Reglementierte berufe und tätigkeiten in der schweiz. Swiss Confederation, Federal Department of Economic Affairs, Education and Research (WBF). Retrieved from <https://www.sbf.admin.ch/diploma>
- van der Klaauw, B., & Vethaak, H. (2022). Empirical evaluation of broader job search requirements for unemployed workers (IZA Discussion Papers).
- Woessmann, L. (2025). Skills and earnings: A multidimensional perspective on human capital. Annual Review of Economics, 17(1), 397–425.
- Yamaguchi, S. (2012). Tasks and heterogeneous human capital. Journal of Labor Economics, 30(1), 1–53.
- Ziegler, L. (2024). Labor market tightness and hiring outcomes (Draft). University of Vienna.
- Zuchuat, J., Lalive, R., Osikominu, A., Pesaresi, L., & Zweimüller, J. (2023). Duration dependence in finding a job: Applications, interviews, and job offers. CEPR Discussion Papers(18600).

Online Appendix

The Role of Firms in Occupational Mobility

A Data Details: Recruiter Clickstream Data

Figure A.1. Recruiter clicks: screenshot of the candidate profile on Job-Room (German)

Berufe, Qualifikationen und Erfahrungen des Kandidaten	
Zuletzt ausgeübter Beruf:	Betriebsarbeiter
Erfahrungsjahre auf dem Beruf:	mehr als 1 Jahr
Berufsbezeichnung:	Automechaniker
Abschluss zum Beruf:	ausländisch, CH nicht anerkannt
Erfahrungsjahre auf dem Beruf:	mehr als 3 Jahre
Kenntnisse, Fähigkeiten, Skills:	Ausbildung in Mazedonien
Berufsbezeichnung:	Bauhandlanger
Erfahrungsjahre auf dem Beruf:	mehr als 1 Jahr
Sprachkenntnisse des Kandidaten	
Deutsch:	mündlich: gut, schriftlich: Grundkenntnisse
Albanisch:	mündlich: sehr gut, schriftlich: sehr gut
Weitere Angaben des Kandidaten	
Verfügbarkeit:	nach Vereinbarung
Arbeitsform:	keine Angabe
Geschlecht:	männlich
Maximales Arbeitspensum:	100%
Mögliche Anstellungsdauer:	unbefristet
Höchste Bildungsstufe:	Sekundarstufe II - ISCED 3
Ausbildung:	Sek. II - Berufliche Grundbildung EFZ od. äq. Sek. I - obligatorische Schule
Gesuchte Arbeitsregion/en:	Grossregion 4 (ZH, SH, TG, SG, AI, AR, GL, GR)
Führerausweiskategorien:	B
Weitere Auskünfte erteilt:	
Adresse:	RAV Frauenfeld, Thundorferstrasse 37, 8510 Frauenfeld Kant. Verwaltung
Kontaktdaten anzeigen	
Zurück Diesen Kandidat als Link senden Druckansicht	

Notes: The figure shows the full candidate profile on Job-Room as seen by a recruiter.

Table A.1. Recruiter clicks: distribution of searches over broad occupation groups

Occupation group	N recruiter searches
Managers	3633
Professionals	30754
Technicians and associate professionals	41838
Clerical support workers	27757
Service and sales workers	22088
Skilled agricultural, forestry and fishery workers	11354
Craft and related trades workers	213945
Plant and machine operators and assemblers	26727
Elementary occupations	28777

Notes: The table reports the number of recruiter searches by searched occupation, aggregated at the ISCO-19 one-digit level.

Table A.2. Recruiter clicks: job seeker profile control variables

Variable	Description
<i>Demographics and Availability</i>	
Gender: Male/Female	Gender of the job seeker
Availability	Immediately available or no information
Protected contact information	Contact details hidden from recruiters
<i>Work Preferences</i>	
Willing to work shifts	Accepts shift work
Willing to work nights	Accepts night work
Willing to work Sundays/holidays	Accepts weekend/holiday work
Willing to work from home	Accepts remote work
Willing to do apprenticeship	Open to apprenticeship positions
Desired work volume	<50%, 50%-90%, or full-time
Contract type	Preference for temporary or permanent contract
<i>Geographic Mobility</i>	
Canton-specific willingness	Willing to work in a canton incl. a major city: ZH, BE, BS, VD, GE (5 dummies)
Geographic mobility	Number of cantons (0, 2-4, 5-7, 8-26)
<i>Language Competencies</i>	
German	Basic, Good, or Very good
French	Basic, Good, or Very good
Italian	Basic, Good, or Very good
English	Basic, Good, or Very good
Other languages	Basic, Good, or Very good
<i>Education and Experience</i>	
Education level	Primary, Secondary/vocational, or University
Number of occupations listed	Count of different professions in profile
<i>Skills Field Text Analysis</i>	
Education mentions	General, tertiary, vocational, or continuing education
Skills mentioned	Soft skills, IT (general/advanced), machine operation, leadership, languages
Experience duration	1-5, 6-10, 10-20, or >20 years
Text characteristics	Language used (German/Italian/other), word count (1-10, 11-20, 21-30, >30)

Notes: The table lists the job seeker control variables extracted from candidate profiles on Job-Room, capturing the information visible to recruiters. Variables under “Skills Field Text Analysis” are extracted from the free-text skills field. Canton abbreviations: ZH (Zurich), BE (Bern), BS (Basel-Stadt), VD (Vaud), GE (Geneva).

B Data Details: Online Application Diaries

This appendix expands on the diary data introduced in Section 3.2. Switzerland’s UI system requires benefit claimants to document their job applications each month in application protocols, recording for each application the date, the position title, and the employer. The protocols are a monitoring tool: caseworkers discuss and validate them, occasionally verify them with employers, and benefit cuts can follow from non-compliance, giving job seekers a strong incentive to report accurately (Zuchuat et al., 2023).

Since 2022, job seekers with a Job-Room account can submit their protocols through

an online form (Figure B.1, Panels (a) and (b)) and import a vacancy’s details directly from the posting via a transfer button (Figure B.1, Panel (c)). When a job seeker uses the import function, the system stores the vacancy’s unique identifier, which links the application to the register record of the opening. Imported applications account for 34.7% of all applications documented by the job seekers who use the import function. Because Job-Room covers nearly the universe of Swiss online vacancies, job seekers can import most of the jobs they apply to. Klaeui (2026) shows that the occupations and locations of vacancies a job seeker looks at on the portal are highly predictive of the occupation and location they find a job in.

We focus on job seekers who fill out their protocols through the Job-Room online form and import button, in the second to fourth quarters of 2022. We start from 23,477 job seekers, of whom 12,230 have a completed spell with a known re-employment firm; for these we observe 167,024 applications to 96,611 distinct vacancies. We drop 779 applications (0.47%) for which the string-match algorithm cannot assign an applied-to firm and a further 216 (0.13%) with an incomplete outcome, overlap measure, control, or fixed-effect key. This leaves an analysis sample of 166,029 applications, which is the sample behind the overlap results of Section 4.3 and the descriptive shares of Section 4.4. The analyses that condition on tightness drop a further 6,968 (4.17%) applications whose vacancy has no canton or whose occupation-by-region cell has no vacancies or no unemployed job seekers, so that tightness is undefined, leaving 159,061 applications. Individual regressions in these samples lose further observations to fixed effects that perfectly predict the outcome; the counts reported in each table are the estimation samples.

The applicant-count regressions (Appendix Table F.6, Panel B) use the applications of all 23,477 job seekers, including those without a completed spell. Their estimation sample contains about 309,000 applications to 135,926 vacancies.

We link each application to a hiring outcome by matching the re-employment firm named in the UI register against the employers in the protocols. We use the two-stage string-matching and clustering procedure of Appendix B.1, which achieves 91.9% balanced accuracy on a hand-labeled test set. The UI register contains the re-employment firm for 88% of job seekers who leave unemployment for a job.

The diary sample is not random: job seekers who use the Job-Room functionalities are somewhat more likely to change occupation upon re-employment than the population of registered job seekers (58% versus 53%). On other observable characteristics the sample resembles the population closely (Table B.1). We estimate a hiring probability of 1.2% conditional on applying, which exceeds the 0.9% rate Zuchuat et al. (2023) obtain from manually coded Swiss protocols. Part of the difference reflects that our sample conditions on job seekers who eventually found a job.

Figure B.1. Application diaries: online submission form and transfer link on a job ad

(a) Application entry form

Nachweis der persönlichen Arbeitsbemühungen

Filtern

This page is only available in the official Swiss languages German, French and Italian.

Search for already entered company or job description

Mai 2024	NpA_2024-05.pdf	Bearbeitungsstatus Anzahl erfasst: 2	Übermittlungsstatus Manuelle Übermittlung möglich ab 27.05.2024
Bewerbung vom 09.05.2024 Company CEO	Ergebnis der Bewerbung ☒ Vorstellungsgespräch ☒ Noch offen ☐ Anstellung ☐ Absage	Bearbeitungsstatus Gespeichert am 15.05.2024	Übermittlungsstatus Automatische Übermittlung erfolgt am 06.06.2024 um 00:00 Uhr
Bewerbung vom 10.05.2024 Company 2 Accountant	Ergebnis der Bewerbung ☒ Vorstellungsgespräch ☐ Noch offen ☐ Anstellung ☒ Absage	Bearbeitungsstatus Gespeichert am 15.05.2024	Übermittlungsstatus Automatische Übermittlung erfolgt am 06.06.2024 um 00:00 Uhr

(b) Overview of monthly proof of effort

This job posting was found on an external website. The State Secretariat for Economic Affairs is providing this search result as an additional service but has no control over its content or quality.

Köchin / Koch (w/m/d)

Coop Genossenschaft Vacant since : 14.05.2024
Schweiz (CH) 100% By agreement Permanent

Unsere Coop Restaurants sind keine gewöhnlichen Restaurants. Sie sind Orte der Begegnung, an denen die gesamte Schweiz zusammenkommt. Als Teil unseres Teams wirst Du nicht nur einen Job haben, sondern Teil dieser Gemeinschaft sein. Zusammen können wir etwas Aussergewöhnliches schaffen!

Köchin / Koch (w/m/d)

Aufgaben

- Kalte und warme Speisen bringst du frisch auf den Teller.
- Immer schön freundlich: Die Gäste schätzen sehr, wie du sie am Buffet bedienst und beräts.
- Mit deinen erfrischenden Gerichten garantierst du Genussmomente.
- Dass du Qualitäts- und Hygienestandards jederzeit einhältst, ist für dich selbstverständlich.

Anforderungen

Previous Next

Link to external job posting

Transfer

Add to favourites

Add note

Copy link

Send link

Print

(c) Transfer button on a Job-Room job posting

Notes: Panels (a) and (b) show screenshots of the online forms used to report monthly application efforts to the Swiss PES. Panel (c) highlights the transfer button on a Job-Room job posting that lets job seekers import the vacancy's details into their monthly application protocol.

Table B.1. Application diaries: composition of the sample of completed spells

	All spells (N = 178 436)		Sample (N = 12 230)	
	Mean	SD	Mean	SD
Female	0.45	0.50	0.50	0.50
Age (at registration)	38.96	11.53	38.81	11.14
Primary education	0.29	0.45	0.27	0.44
Secondary or vocational educ.	0.48	0.50	0.55	0.50
University education	0.16	0.37	0.13	0.33
Non-permanent resident	0.25	0.44	0.27	0.44
> 3 years tenure in last job	0.67	0.47	0.66	0.47
Insured earnings (CHF)	5078.60	2478.61	4792.11	2244.54
Reemployment occ. different from last occ.	0.53	0.50	0.58	0.49

Notes: The table compares the diary sample of job seekers with completed spells and a known re-employment firm to the population of unemployment spells active in 2022 and completed with a known re-employment firm.

B.1 Matching Firm Names from the Application to the Re-employment Firm

To assign free-text employer names to consistent firm identifiers, we implement a two-stage procedure: string matching from the survey data to firms in the commercial register, followed by clustering of similar firm entries within the register itself. The second step is important, because official registers list legal entities and often include multiple entries for what is essentially the same firm group (divisions, regional subsidiaries, or holding entities with nearly identical names). While the string matching provides good matches for similar entries, where a firm has more than one similar-sounding legal entity, it lacks the precision to pick the right one. Our clustering creates higher-level entities that group these entries, and we match two firms if they belong to the same group; this also lets us use the groups for firm fixed effects.

For the string matching in both stages we follow Brüll (2025) and Brüll et al. (2025). Firm names are normalized by uppercasing and removing non-alphanumeric characters while preserving diacritics, then tokenized. We drop tokens appearing in more than one in 500 firms, which carry little discriminatory power. We weight each remaining token by its inverse frequency, normalized to sum to one within each name, which gives more weight to distinctive words (“NOVARTIS” over “LTD”). For the clustering of register entities, we use hierarchical clustering with complete linkage and a minimum similarity threshold of 0.95. On a hand-labeled set of 2,000 matched name pairs, the procedure achieves a balanced accuracy of 91.9%, with a sensitivity of 98.97% and a specificity of 84.77%.

C Details on the Job Requirements Overlap Index

This appendix documents the job requirements overlap measure: the vacancy data it is built from, how the index is computed, how missing values are imputed, and how it validates against an external classification of related occupations.

C.1 Vacancy Data


We measure the overlap in job requirements using data covering the near-universe of online job postings in Switzerland between 2016 and 2022. The data were first used by Colella (2022) to analyze the impact of trade on labor demand. They are collected by the private company x28 AG, which continuously crawls postings from all major job boards and company websites in Switzerland, identifies duplicates, and assigns postings to industry and occupation classifications (Bannert et al., 2022). The data cover at least 90% of all online job postings in Switzerland, and their industry and regional composition is similar to official vacancy statistics (Bannert et al., 2022).

A central feature of the data is that each ad comes with a list of job requirements extracted from its text and title. x28 extracts these by matching the text to a large, manually maintained database of task and skill terms. The ontology distinguishes 23,592 skills or tasks; we restrict to the 1,087 that appear at least once per 5,000 ads. It includes soft and hard skills as well as education requirements, with examples such as “accuracy”, “commercial understanding”, “sales skills”, “team orientation”, or specific IT skills.

C.2 Index Computation

We compute the overlap from postings in the German-speaking part of Switzerland, where the extraction of requirements is more reliable, pooling all postings from 2016 to the second quarter of 2022. The computation proceeds in two steps. First, we randomly select a pair of vacancy postings and compute their Jaccard similarity, the number of requirements mentioned in both ads divided by the number of distinct requirements across both. Figure C.1 gives a stylized example for an office clerk and a shop sales assistant: the two ads list seven distinct requirements, of which two appear in both, for a Jaccard similarity of $\frac{2}{7}$. Second, we repeat this for 5,000 randomly drawn pairs, sampling one ad from each occupation with replacement, and average the similarities. The same procedure yields within-occupation overlap by drawing both ads from the same occupation. We compute the index only for occupations with at least 200 such postings, which covers 80% of all four-digit ISCO occupation pairs.

Figure C.1. Job requirements overlap: an example



Notes: The figure illustrates job requirements overlap between two example vacancies.

C.3 Imputation of Missing Overlap

The sample restriction leaves some occupation pairs without an overlap value. Of the 15,266,233 recruiter search results, 410,429 (about 2.7%) lack a value. We impute these hierarchically along the ISCO structure. We first use the average overlap among pairs sharing the same three-digit ISCO occupation, which resolves the majority of cases. The remaining cases are imputed within the two-digit level, and the last few within the one-digit level. The imputation source for each pair is recorded, for both the unweighted and the inverse-document-frequency-weighted matrices, in Table C.1.

Table C.1. Job requirements overlap: imputation source for missing values


Source	Between (origin, target) pairs		Within-target (vacancy) occupation	
	Unweighted	IDF	Unweighted	IDF
raw	25,196	25,196	301	301
3-digit	1,861	1,861	43	43
2-digit	857	857	4	4
1-digit	9	9	0	0

Notes: The table reports, for the unweighted and the inverse-document-frequency-weighted overlap matrices, how many (origin, target) pairs and target occupations have an overlap value computed directly versus imputed at the three-, two-, or one-digit ISCO level.

C.4 Job Requirements Within and Across Occupations

Figure C.2 shows the distribution of overlap across all occupation pairs, with example pairs marked on the axis. Jobs for office clerks share on average 7.8% of their requirements with other office clerk jobs, whereas a farm laborer job and a music teacher job share only 2%.

Figure C.2. Job requirements overlap: distribution across occupation pairs




Notes: The histogram shows job requirements overlap across all pairs of four-digit ISCO occupations, including the pairs of an occupation with itself; the axis labels mark example pairs. The graph is truncated at the 99th percentile of the overlap.

C.5 Validation: Comparison to O*NET

We validate the overlap measure against the O*NET classification of related occupations, maintained by the U.S. Department of Labor. For each occupation, O*NET identifies ten related occupations based on task, knowledge, and job-title similarity and sorts them into three tiers, “primary-long”, “primary-short”, and “supplemental”. We focus on the primary tiers. O*NET is defined on U.S. SOC codes. We map SOC-2019 codes to ISCO-19 via SOC-2010, weighting SOC occupations by their U.S. employment shares (following Belot et al., 2019), and label an ISCO-19 pair as related if at least one corresponding SOC pair is marked related in O*NET. Figure C.3 plots the share of occupation pairs classified as related by O*NET against our overlap measure. The relationship is clearly positive: as overlap rises, the likelihood that two occupations are classified as related rises as well. This shows that our continuous measure captures the same dimensions of similarity as the discrete O*NET classification while providing a finer scale.

Figure C.3. Job requirements overlap: comparison to the O*NET “related occupations”



Notes: The figure shows a binned scatter plot of the share of related occupations according to O*NET against our overlap measure. We map O*NET SOC-2019 codes to ISCO-19 via SOC-2010, weight SOC-2019 occupations by their U.S. employment shares (following Belot et al., 2019), and label an ISCO-19 pair as related if at least one corresponding SOC-2019 pair is marked related in O*NET.

D Measuring Tightness


To study how recruiters’ openness to movers responds to scarcity, we construct a time-varying, occupation- and region-specific measure of labor market tightness from data external to Job-Room, which ensures that it is not mechanically related to platform search behavior. For each occupation-region cell we measure tightness as the ratio of the average stock of online vacancies to the average stock of registered unemployed,

$$tightness_{o,\mathcal{C},t} = \frac{V_{o,\mathcal{C},t}}{U_{o,\mathcal{C},t}}, \quad (8)$$

where $V_{o,\mathcal{C},t}$ is the 30-day rolling average stock of online vacancies in occupation o and region $\mathcal{C}$ around day t from the x28 data, and $U_{o,\mathcal{C},t}$ is the corresponding average stock of registered unemployed who last worked in o and search in $\mathcal{C}$, from the UI register. Because we use the log of tightness, we lose the small number of cell-months with zero vacancies or zero unemployed. For recruiters, o is the searched occupation, t the day of the search, and $\mathcal{C}$ the canton or broader region specified in the search (Switzerland as a whole when no region is specified).

Despite the short sample period, tightness varies substantially across cells and over time. Figure D.1, Panel (a), shows the distribution of log tightness across recruiters’ searched markets, and Panel (b) shows its evolution over time, plotting the 25th percentile, mean, and 75th percentile of log-differenced tightness across cells in each month.

Figure D.1. Labor market tightness: distribution across searched markets and variation over time



(a) Recruiter clicks: distribution across searched markets (b) Variation over time across occupation-region cells

Notes: Panel (a) shows the log tightness of recruiters' searched markets (occupation $\times$ location) across months. Panel (b) shows the evolution of tightness at the occupation-region level over time. Each series is indexed by subtracting its log tightness in the initial month (March 2017); the panel plots the 25th percentile, mean, and 75th percentile of these log-differenced values across cells in each month.

E Vacancy-Occupation Characteristics

Section 4.3 relates the mover gap to three constructed properties of the searched occupation besides tightness: whether the occupation is licensed, its wage level, and the typical education of its incumbents.

Licensing. Switzerland maintains a register of regulated occupations and activities, covering those whose practice requires a recognized diploma, a license, or an equivalent authorization (Staatssekretariat für Bildung, Forschung und Innovation, 2021). The register is organized by domain and by the legal act that governs each entry rather than by occupational code. We therefore assign four-digit ISCO codes to its entries by hand. An occupation counts as licensed if it matches an entry in the register.

Wage level. We measure the wage level of an occupation as the log median wage of its incumbents, computed from the 2020 wave of the Swiss Structural Earnings Survey.


Typical education of incumbents. We take the education level that occurs most often among incumbents in an occupation, that is, among candidates whose last occupation is the searched occupation, excluding those for whom no education is recorded. The measure has three levels: primary, secondary or vocational, and university, with secondary or vocational as the omitted reference.

F Further Results and Tables

This appendix collects the regression tables and supporting figures referenced in the main text, grouped by the analysis they belong to.

F.1 Baseline Mover Gap

Figure F.1. Recruiter clicks: within-search relationship between result-list rank and mover status




Notes: For each result-list rank, the figure plots the mean of the mover indicator, demeaned within recruiter search and averaged across searches. Demeaning removes differences in the mover share across searches of different lengths. The sample is the estimation sample of Table 2.

Table F.1. Recruiter clicks: robustness checks for the contact gap for job seekers from a different occupation

<i>Panel A: Alternative mover definition, sample, and outcome</i>				
	(1) Baseline	(2) Three-digit ISCO mover	(3) Not opened or viewed under 20s	(4) Print, email, or favorite
Dependent variable	Contact click	Contact click	Contact click	Print, email, or favorite
Mover (different occupation)	-0.0077*** (0.0004)	-0.0076*** (0.0004)	-0.0028*** (0.0003)	-0.0023*** (0.0001)
Effect relative to baseline (%)	-7.0299	-6.8939	-4.5510	-12.4592
Candidate FE	X	X	X	X
Recruiter search FE	X	X	X	X
Search rank (absolute and relative) FE	X	X	X	X
Observations	15,266,233	15,266,233	14,221,859	15,266,233
R2	0.46442	0.46442	0.47637	0.31158
Number of searches	406,873	406,873	387,393	406,873
Number of recruiters	37,665	37,665	37,043	37,665
Baseline probability	0.1100	0.1100	0.0618	0.0184
<i>Panel B: Leave-out neighborhood controls</i>				
	(1) Baseline (neighborhood sample)	(2) Neighborhood window 1	(3) Neighborhood window 3	(4) Neighborhood window 5
Dependent variable	Contact click	Contact click	Contact click	Contact click
Mover (different occupation)	-0.0062*** (0.0004)	-0.0062*** (0.0004)	-0.0062*** (0.0004)	-0.0062*** (0.0004)
Mean mover indicator in neighborhood		-0.0004* (0.0002)	-0.0001 (0.0006)	-0.0006 (0.0009)
Mean candidate characteristics in neighborhood		X	X	X
Candidate FE	X	X	X	X
Recruiter search FE	X	X	X	X
Search rank (absolute and relative) FE	X	X	X	X
Observations	11,941,492	11,941,490	11,941,480	11,936,122
R2	0.46501	0.46502	0.46504	0.46507
Number of searches	262,307	262,307	262,307	262,307
Number of recruiters	32,997	32,997	32,997	32,997
Baseline probability	0.0869	0.0869	0.0869	0.0869

Notes: The table reports robustness checks for the baseline in Table 2. In Panel A, Column (1) replicates the baseline. Column (2) defines a mover at the three-digit rather than the four-digit ISCO level. Column (3) restricts the sample to candidates the recruiter either never opened or viewed for at most 20 seconds. Column (4) replaces the contact click with an indicator for printing, emailing, or favoriting the profile. Panel B adds neighborhood controls to the baseline regression. For each candidate at rank r we compute the mean of every job seeker control, and of the mover indicator, across the $2X$ candidates at ranks $r - X, \dots, r - 1$ and $r + 1, \dots, r + X$. All columns of Panel B use the sample of candidates whose $X = 5$ neighborhood is fully defined ($\text{rank} \geq 6$ and $\text{rank} \leq n_{\text{results}} - 5$). Column (1) is the baseline specification on that sample. Columns (2) to (4) add the controls at $X = 1, 3$, and 5. Standard errors, in parentheses, are clustered two-way on the recruiter and the candidate. *, **, and *** denote significance at the 10%, 5%, and 1% levels.


Figure F.2. Recruiter clicks: the contact gap by searched occupation



Notes: The figure shows the mover gap by searched occupation at the ISCO two-digit level, as a percent of the occupation-specific baseline contact rate. The specification follows Table 2, Column (1), with the mover indicator interacted with the searched occupation. Only two-digit occupations with at least 1,000 searches are shown. Error bars show 95% confidence intervals based on standard errors clustered two-way on the recruiter and the candidate.

F.2 Vacancy-Side Heterogeneity

Figure F.3. Recruiter clicks: predicted contact gap for movers, by the overlap gap and characteristics of the searched occupation



Notes: The figure shows predicted contact gaps for movers, as a percent of the baseline contact rate, from a joint regression interacting the mover indicator with the overlap gap to incumbents and four characteristics of the searched occupation: labor market tightness, licensing, the wage level, and the typical education of incumbents. Each row reports two predictions, at the dimension's low and high values, with the other dimensions held at their sample means. The low and high values are the bottom and top terciles of the overlap gap among movers, the slack and tight terciles of log tightness, not licensed and licensed, the bottom and top terciles of the log wage level, and primary and university education of incumbents. The rest of the regression specification is equivalent to Equation (4). Column (2) of Appendix Table F.2, Panel A, reports the difference between the two predictions in each row, with the overlap gap scaled per standard deviation instead of by tercile. Error bars show 95% confidence intervals based on standard errors clustered two-way on the recruiter and the candidate.

Two patterns in Table F.2, Panel A, deserve comment. The estimated influence of licensing gets stronger once the other dimensions enter the regression, because licensed occupations are disproportionately tight, an effect that works in the opposite direction. The mover gap is smaller in higher-paying occupations and in those whose incumbents hold university degrees, but these associations are explained by the other dimensions. Higher-paying occupations are tighter, and occupations with university-educated incumbents are both tighter and less homogeneous in their job requirements, implying that the overlap gap to incumbents is smaller there and incumbency confers less of an advantage. Both change the gap little once the other dimensions are accounted for. Because the

overlap gap is the only dimension that varies with the candidate as well as the vacancy, it may have an advantage over purely vacancy-side characteristics in this comparison. Replacing it with the within-occupation homogeneity of the searched occupation, which does not depend on the job seeker, leaves the results unchanged: the wage and education gradients still vanish, to a large extent through the same channels (Table F.2, Panel B).


Table F.2. Recruiter clicks: Gelbach decomposition of heterogeneity across searched occupation characteristics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Change in the mover gap			Contribution from				
Dimension	Univariate reg.	Joint reg.	Difference	Overlap	Tightness	Licensing	Wage	Education
<i>Panel A: Overlap gap to incumbents</i>								
Overlap gap (per standard deviation)	-3.5 (0.4)	-3.9 (0.5)	-0.4	–	-0.6 (0.1)	0.0 (0.0)	0.1 (0.2)	0.1 (0.1)
Tightness (tight minus slack tercile)	5.4 (0.9)	7.6 (1.1)	2.2	1.2 (0.2)	–	0.4 (0.2)	0.0 (0.3)	0.6 (0.3)
Licensing (licensed minus not licensed)	-1.3 (1.3)	-3.1 (1.3)	-1.8	0.2 (0.0)	-2.8 (0.4)	–	0.1 (0.3)	0.6 (0.3)
Wage (high minus low tercile)	2.9 (0.7)	0.1 (0.9)	-2.8	-0.3 (0.0)	-2.6 (0.4)	0.2 (0.1)	–	-0.1 (0.4)
Education (university minus primary)	5.2 (1.0)	-0.8 (1.3)	-6.0	-3.1 (0.4)	-3.0 (0.4)	0.5 (0.2)	-0.3 (0.8)	–
<i>Panel B: Within-occupation homogeneity</i>								
Homogeneity (homogeneous minus heterogeneous)	-5.9 (1.0)	-6.4 (1.2)	-0.4	–	-1.3 (0.3)	0.1 (0.0)	0.9 (0.8)	-0.1 (0.3)
Tightness (tight minus slack tercile)	5.4 (0.9)	7.3 (1.1)	1.9	0.7 (0.1)	–	0.4 (0.2)	0.1 (0.3)	0.7 (0.3)
Licensing (licensed minus not licensed)	-1.3 (1.3)	-2.5 (1.3)	-1.2	0.4 (0.1)	-2.6 (0.4)	–	0.3 (0.3)	0.7 (0.3)
Wage (high minus low tercile)	2.9 (0.7)	0.0 (0.9)	-2.9	-0.4 (0.1)	-2.5 (0.4)	0.2 (0.1)	–	-0.2 (0.4)
Education (university minus primary)	5.2 (1.0)	-0.7 (1.3)	-5.9	-2.8 (0.6)	-2.9 (0.4)	0.4 (0.2)	-0.6 (0.8)	–

Notes: The table decomposes how the mover gap varies across five dimensions: a measure of job requirements overlap and four characteristics of the searched occupation (labor market tightness, licensing, the wage level, and the typical education of incumbents). The overlap measure is the overlap gap to incumbents in Panel A and the within-occupation homogeneity of the searched occupation in Panel B. In Panel B, all five dimensions are properties of the searched occupation and none depends on the job seeker. Columns (1) and (2) report how the mover gap changes along the dimension in a regression that includes only that dimension and in a regression that includes all five. Because the mover gap is negative, a positive value means the gap becomes smaller. Column (3) is the difference between the two. The remaining columns attribute that difference to each of the other dimensions, following Gelbach (2016), doing one Gelbach decomposition per row. Each contribution is the product of the coefficient on the contributing dimension in the joint regression and the coefficient from an auxiliary regression of that dimension on the focal one. The overlap gap enters linearly, scaled per standard deviation. The other dimensions, including homogeneity, enter as indicators, and for dimensions with more than two categories the table reports the difference between the two extreme ones. All values are in percent of the baseline contact rate. All regressions include candidate, recruiter-search, absolute-rank, and relative-rank fixed effects and use the same sample. Contributions may not add up exactly because of rounding. Standard errors, in parentheses, are clustered two-way on the recruiter and the candidate.

F.3 Job Requirements Overlap and the Mover Gap

Figure F.4. Recruiter clicks: the contact gap for movers by overlap, inverse-document-frequency-weighted



(a) By within-occupation homogeneity of the searched occupation

(b) By the overlap gap to the searched occupation

Notes: The figure replicates Figure 4 with an inverse-document-frequency-weighted measure of job requirements overlap. Each requirement s is weighted by $idf_s = \log(N/n_s)$, where n_s is the number of ads listing requirement s and N the number of ads. The overlap between two ads is the sum of the weights of the requirements they share divided by the sum of the weights of all requirements appearing in either. The specification, the definition of the overlap gap, the dashed line, and the confidence intervals are as in Figure 4. Appendix Table F.3, Columns (2) and (4), report the underlying coefficients.

Table F.3. Recruiter clicks: the contact gap by job requirements overlap, unweighted and inverse-document-frequency-weighted

	Within	Within (IDF)	Overlap gap	Overlap gap (IDF)
Dependent Var.:	Contact click	Contact click	Contact click	Contact click
Within overlap below 7%	-0.0033*** (0.0009)	-0.0043*** (0.0005)		
Within overlap 7-8%	-0.0054*** (0.0007)	-0.0071*** (0.0015)		
Within overlap 8-9%	-0.0049*** (0.0010)	-0.0075*** (0.0010)		
Within overlap 9-10%	-0.0062*** (0.0010)	-0.0102*** (0.0012)		
Within overlap 10-11%	-0.0080*** (0.0012)	-0.0108*** (0.0020)		
Within overlap 11-12%	-0.0114*** (0.0011)	-0.0095*** (0.0012)		
Within overlap 12% or more	-0.0108*** (0.0009)	-0.0111*** (0.0010)		
Mover, overlap closest to within (reference)			-0.0016** (0.0007)	-0.0006 (0.0007)
Overlap gap 0.5-1.5pp below within			-0.0032*** (0.0007)	-0.0040*** (0.0007)
Overlap gap 1.5-2.5pp below within			-0.0049*** (0.0008)	-0.0063*** (0.0008)
Overlap gap 2.5-3.5pp below within			-0.0059*** (0.0010)	-0.0072*** (0.0010)
Overlap gap 3.5-4.5pp below within			-0.0065*** (0.0011)	-0.0064*** (0.0012)
Overlap gap >4.5pp below within			-0.0108*** (0.0010)	-0.0108*** (0.0010)
Candidate FE	X	X	X	X
Recruiter search FE	X	X	X	X
Search rank (abs. and rel.) FE	X	X	X	X
Relative search rank FE	X	X	X	X
Observations	15,266,233	15,266,233	15,266,233	15,266,233
R2	0.46443	0.46443	0.46443	0.46443

Notes: The table reports the regressions behind Figure 4 and Appendix Figure F.4. Columns (1) and (2) interact the mover indicator with seven bins of the within-occupation overlap of the searched occupation. Columns (3) and (4) add six bins of the overlap gap to incumbents to the mover indicator, with movers within half a percentage point of the within-occupation value the omitted reference. Columns (1) and (3) use the unweighted overlap. Columns (2) and (4) use the inverse-document-frequency-weighted overlap defined in the notes to Appendix Figure F.4. All columns include candidate, recruiter-search, absolute-rank, and relative-rank fixed effects. Standard errors, in parentheses, are clustered two-way on the recruiter and the candidate. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

F.4 Mover Profiles and the Signal Hierarchy

Table F.4. Recruiter clicks: determinants of the contact gap for job seekers from a different occupation

	(1)	(2)	(3)	(4)	(5)
Last worked in different occupation	-0.0077*** (0.0004)	-0.0016** (0.0007)	-0.0069*** (0.0004)	-0.0011 (0.0009)	0.0006 (0.0021)
Last-vacancy overlap 0.5-1.5pp below within		-0.0032*** (0.0007)		-0.0044*** (0.0010)	-0.0085*** (0.0022)
Last-vacancy overlap 1.5-2.5pp below within		-0.0049*** (0.0008)		-0.0052*** (0.0012)	-0.0069*** (0.0025)
Last-vacancy overlap 2.5-3.5pp below within		-0.0059*** (0.0010)		-0.0056*** (0.0014)	-0.0121*** (0.0029)
Last-vacancy overlap 3.5-4.5pp below within		-0.0065*** (0.0011)		-0.0072*** (0.0016)	-0.0103*** (0.0031)
Last-vacancy overlap >4.5pp below within		-0.0108*** (0.0010)		-0.0091*** (0.0014)	-0.0062** (0.0032)
No experience in vacancy occ.			-0.0052*** (0.0006)	-0.0035** (0.0016)	-0.0090*** (0.0032)
No prof. qualification in vacancy occ.			-0.0036*** (0.0005)		
Experience, no prof. qualification in vacancy occ.				-0.0033*** (0.0007)	-0.0060*** (0.0012)
Last worked in different occupation x Experience, no prof. qualification in vacancy occ.				-0.0011 (0.0012)	-0.0018 (0.0025)
Experience, no prof. qualification in vacancy occ. x Last-vacancy overlap 0.5-1.5pp below within				0.0027** (0.0013)	0.0047* (0.0027)
No experience in vacancy occ. x Last-vacancy overlap 0.5-1.5pp below within				-0.0004 (0.0018)	0.0056 (0.0036)
Experience, no prof. qualification in vacancy occ. x Last-vacancy overlap 1.5-2.5pp below within				0.0015 (0.0015)	0.0015 (0.0031)
No experience in vacancy occ. x Last-vacancy overlap 1.5-2.5pp below within				-0.0014 (0.0020)	0.0039 (0.0039)
Experience, no prof. qualification in vacancy occ. x Last-vacancy overlap 2.5-3.5pp below within				0.0014 (0.0017)	0.0086** (0.0035)
No experience in vacancy occ. x Last-vacancy overlap 2.5-3.5pp below within				-0.0042* (0.0022)	0.0011 (0.0044)
Experience, no prof. qualification in vacancy occ. x Last-vacancy overlap 3.5-4.5pp below within				0.0029 (0.0019)	0.0056 (0.0038)
No experience in vacancy occ. x Last-vacancy overlap 3.5-4.5pp below within				-0.0033 (0.0023)	-0.0021 (0.0044)
Experience, no prof. qualification in vacancy occ. x Last-vacancy overlap >4.5pp below within				0.0004 (0.0015)	-0.0006 (0.0031)
No experience in vacancy occ. x Last-vacancy overlap >4.5pp below within				-0.0093*** (0.0020)	-0.0089** (0.0040)
Last worked in diff occ. x 35-45y old					0.0004 (0.0013)
Last worked in diff occ. x 50y old					0.0006 (0.0013)
Recruiter search FE	X	X	X	X	X
Search rank (abs. and rel.) FE	X	X	X	X	X
Candidate FE	X	X	X	X	X
Observations	15,266,233	15,266,233	15,266,233	15,266,233	4,624,254
R2	0.46442	0.46443	0.46443	0.46445	0.50768
N searches	406,873	406,873	406,873	406,873	406,873
N recruiters	37,665	37,665	37,665	37,665	37,665
Baseline prob.	0.1100	0.1100	0.1100	0.1100	0.1100

Notes: The table reports the regressions behind Figure 5. Column (1) is the baseline specification. Column (2) adds the six overlap-gap bins. Column (3) adds indicators for having no work experience and no professional certificate in the searched occupation. Column (4), which the figure is built from, replaces those indicators with three groups, experience and certificate, experience without certificate, and no experience, and interacts them with the mover indicator and with the overlap bins. Column (5) repeats Column (4) on candidates with no entry in the free-text skills field and adds interactions of the mover indicator with age 35 to 49 and with age 50 and above. All columns include candidate, recruiter-search, absolute-rank, and relative-rank fixed effects. Standard errors, in parentheses, are clustered two-way on the recruiter and the candidate. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

F.5 Tightness and Recruiter Behavior

Table F.5. The mover gap by tightness and overlap with the vacancy

	(1) All movers	(2) All movers	(3) Close movers	(4) Distant movers
<i>Panel A: Recruiter clicks (linear probability model)</i>				
Dependent variable	Contact click	Contact click	Contact click	Contact click
Mover (different occupation)	-0.0097*** (0.0006)		-0.0061*** (0.0006)	-0.0143*** (0.0009)
Mover x middle tightness tercile	0.0019** (0.0008)	0.0023* (0.0012)	0.0038*** (0.0009)	0.0041*** (0.0011)
Mover x top tightness tercile	0.0059*** (0.0010)	0.0062*** (0.0019)	0.0041*** (0.0013)	0.0101*** (0.0014)
Candidate FE	X	X	X	X
Recruiter search FE	X	X	X	X
Search rank (absolute and relative) FE	X	X	X	X
Searched occupation x mover FE		X		
Search location x mover FE		X		
Observations	15,085,933	15,085,933	10,056,312	12,938,748
R2	0.46448	0.46457	0.48676	0.46859
Base contact rate	0.1102	0.1102	0.1102	0.1102
<i>Panel B: Application diaries (logit)</i>				
Dependent variable	Hired by the firm		Hired by the firm	Hired by the firm
Mover (different occupation)	-0.6152*** (0.1300)		-0.5406*** (0.1503)	-0.5993*** (0.1468)
Mover x middle tightness tercile	0.1801 (0.1584)		0.0506 (0.1757)	0.2217 (0.2043)
Mover x top tightness tercile	0.3341* (0.1842)		0.4025 (0.2569)	0.1560 (0.1767)
Middle tightness tercile	-0.0343 (0.1244)		-0.0388 (0.1254)	-0.0700 (0.1271)
Top tightness tercile	0.0994 (0.1445)		0.0159 (0.1414)	0.0810 (0.1678)
Hiring firm FE	X		X	X
Job seeker characteristics	X		X	X
Observations	153,363		100,300	96,581
Pseudo R2	0.28600		0.32842	0.29868
Base hire rate	0.0125		0.0130	0.0132

Notes: The table reports the regressions behind Figure 7. Panel A reports the recruiter-click regressions behind Panel (a) of the figure, and Panel B the application-diary regressions behind Panel (b). The mover indicator is interacted with the tightness terciles of the vacancy occupation, with the slack (bottom) tercile the omitted reference. In Panel B, the tightness main effects also enter directly. Each column keeps the incumbents as the reference. Columns (1) and (2) keep all movers, (3) movers from a closely related occupation (overlap gap to incumbents of at most 2 percentage points), and (4) movers from a distant occupation (gap above 2 points). Panel A includes candidate, recruiter-search, absolute-rank, and relative-rank fixed effects, and standard errors are clustered two-way on the recruiter and the candidate. Column (2) of Panel A adds searched-occupation-by-mover and search-location-by-mover fixed effects, which absorb the mover main effect. Panel B includes a hiring-firm fixed effect and the job seeker controls listed in the notes to Table 3. Its coefficients are from logit regressions, and standard errors are clustered two-way on the job seeker and the hiring firm. The figure reports average marginal effects. Standard errors are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table F.6. The candidate pool and the incumbent contact rate by tightness

	(1)	(2)	(3)
	Movers	Incumbents	Incumbents
<i>Panel A: Recruiter clicks</i>			
Unit of observation	Search	Search	Incumbent profile
Dependent variable	Number of profiles	Number of profiles	Contact click
Estimator	Poisson	Poisson	Linear probability
Middle tightness tercile	0.0085 (0.0223)	-0.0308 (0.0215)	-0.0012 (0.0048)
Top tightness tercile	0.0067 (0.0300)	-0.1182*** (0.0277)	-0.0028 (0.0065)
Calendar month FE	X	X	X
Searched occupation x location FE	X	X	X
Recruiter FE	X	X	X
Observations	401,672	401,672	7,909,127
Pseudo R2	0.44834	0.43913	
R2			0.15962
<i>Panel B: Application diaries</i>			
Unit of observation	Vacancy	Vacancy	
Dependent variable	Number of applicants	Number of applicants	
Estimator	Poisson	Poisson	
Middle tightness tercile	-0.1230*** (0.0228)	-0.3292*** (0.0217)	
Top tightness tercile	-0.2054*** (0.0320)	-1.0864*** (0.0348)	
Calendar month FE	X	X	
Hiring firm FE	X	X	
Observations	135,926	135,926	
Pseudo R2	0.13947	0.16614	
Mean of dependent variable	1.59677	0.67592	

Notes: In Panel A, Columns (1) and (2) report the Poisson regressions behind Figure 6, Panel (a). Each observation is a recruiter search on Job-Room, and the outcomes are the number of mover (Column (1)) and incumbent (Column (2)) profiles the search returns. Column (3) reports the regression behind the third term of Equation (7): a linear probability model of the contact indicator on the sample of incumbent profiles. All columns of Panel A include calendar-month, vacancy-occupation-by-location, and recruiter fixed effects, and standard errors are clustered on the recruiter. Panel B reports the Poisson regressions behind Figure 6, Panel (b). Each observation is a vacancy in the application diary data, and the outcomes are the number of mover (Column (1)) and incumbent (Column (2)) applicants per vacancy. The counts include the applications of all job seekers in the diary sample, not only of those with a completed spell (Appendix B). Both columns of Panel B include calendar-month and hiring-firm fixed effects, and standard errors are clustered on the hiring firm. All columns regress on the tightness terciles, with the slack (bottom) tercile the omitted reference. Standard errors are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels.