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 44 633 93 67; siegenthaler@kof.ethz.ch

September 8, 2026

Abstract

How do firms evaluate workers applying from a different occupation? We use clickstream data from a recruitment platform, where recruiters screen registered job seekers, alongside application diaries that record applications and hiring outcomes. Movers face a 7% lower contact rate than equally qualified incumbents. A new measure of job requirements overlap explains about four-fifths of this gap, which is 3.3 times larger in homogeneous occupations, where incumbency implies a near-complete match with a vacancy's requirements; experience and credentials close the rest. Firms also adapt to scarcity, lowering the bar for movers as markets tighten; this firm response accounts for about a third of the shift toward movers, similarly whether we watch recruiters make contacts or firms make hires.

Keywords: occupations; occupational mobility; job requirements overlap; labor demand; labor supply; labor market tightness
JEL Codes: J24, J62, J64

Acknowledgements: 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 crucial 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 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 we use in this article, about 70% of all job applications come from 'movers'—applicants who last worked in a different occupation than the occupation of the firm's advertised position—yet only 63% of actual hires do. Descriptively, this "mover gap" implies that firms' hiring decisions offset roughly one-third of the cross-occupational moves that workers attempt.

The lack of scientific attention to the role of firms in occupational mobility is likely due to the steep data demands of isolating their role. Since realized occupational transitions reflect both workers' and firms' decisions, researchers must observe both who applies to firms' vacancies and whom firms hire. Causally identifying the mover gap also requires a research design that accounts for the fact that job seekers are unlikely to decide randomly whether to change occupations.

This paper analyzes how firms shape occupational transitions: why firms hesitate to hire movers, when they overcome that hesitation, and how much it shapes realized occupational transitions. We draw on a unique combination of matched Swiss administrative and online data sets: we link unemployment insurance (UI) register data with clickstream data from "Job-Room," the recruitment platform of the Swiss Public Employment Service, where recruiters search a database of registered job seekers, and with monthly application diaries that record applications to vacancies and the associated hir-

[1] For 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.

ing outcomes. These data sets reveal the composition of the candidate pool of vacancies and enable monitoring recruiters' contact decisions when screening profiles of candidates with different last occupations and occupational experience. They allow us to estimate the causal effect of job seekers' occupational experience on firms' selection decisions.

Our data document the importance of firms as gatekeepers of occupational mobility. In the application diaries, an incumbent's application is 35.0% more likely to result in a hire than a mover's. On the recruitment platform, recruiters are 27% more likely to contact an incumbent than a mover. However, we find that these unconditional gaps reflect that job seekers who attempt to switch occupations are negatively selected: movers have lower chances of finding employment in any occupation. To show this, we use recruiter clickstream data and exploit the fact that most job seekers are willing to work in several occupations, so they appear in searches for different occupations. We can therefore compare how the same person is treated when she is an incumbent versus when she is a mover, holding her overall quality constant using candidate fixed effects. In addition, the regressions control for search fixed effects and thus compare decisions only among job seekers who are seen by the same recruiter during the same search. This comparison yields an estimated mover gap of 0.77 percentage points, or 7% relative to the baseline contact rate of 11%. The gap is substantially smaller than the raw difference after adding candidate fixed effects, confirming that movers are, on average, negatively selected.

We analyze firms' role through the lens of employer screening under uncertainty. Recruiters cannot observe a candidate's skills directly, so they assess her in part through the average characteristics of her previous occupation: its typical requirements indicate which skills she is likely to possess. A candidate's expected fit therefore rises with the overlap in requirements between her last occupation and the advertised position. Because a mover's last occupation overlaps less with the vacancy than an incumbent's, recruiters expect movers to meet fewer of the position's requirements and contact them less. Testing this notion requires measuring the expected skill distance between the position and both an average mover and an average incumbent, since even an incumbent's last job might have required skills different from those required for the recruiter's position. Since existing indices do not quantify skill overlap between positions *within* the same occupation, we construct a novel measure using data on job requirements extracted from the near-universe of online job postings in Switzerland. We randomly select vacancy pairs and estimate the proportion of job requirements present in both—a proportion we can compute for ads in different occupations and within the same occupation. We validate this measure by showing that it strongly predicts whether two occupations are classified as "related" according to O*Net. It is also highly predictive of the occupations in which job seekers apply.

Using the new measure, we estimate the contact gap in the comparison where movers and incumbents, on average, meet the same share of the position's requirements. Equalizing overlap in this way reduces the gap by 80%, suggesting that movers' lower requirements overlap with the advertised position is the main reason recruiters contact them less. The framework further predicts that within-occupation homogeneity matters: where jobs in an occupation share many requirements, the incumbent label is a stronger signal of overlap, and movers face a correspondingly larger disadvantage. This prediction is strongly confirmed in the data: the contact gap is 3.3 times larger in occupations with homogeneous than with heterogeneous job requirements. Together, these results suggest that recruiters read a candidate's last occupation as information about her expected productivity in the position. Their behavior is consistent with skill-based hiring: recruiters respond to the skills a candidate likely brings to the position, rather than to her formal credentials alone.

We then estimate contact gaps for different types of movers. We find that the contact gap varies sharply across mover profiles, and job requirements overlap, occupation-specific experience, and professional qualifications each reduce it independently. Consistent with economic intuition and theory, the gap is largest for job seekers from distant occupations with no prior occupational experience in the recruiters' searched occupation—a group that accounts for a quarter of applications from other occupations in the application diary data. Such candidates face a contact gap of 17.5%. Moving to a closer occupation, adding occupational experience, and adding a professional qualification each lowers the gap. Return movers who hold a credential in the vacancy occupation and worked there recently face no contact gap at all. XXX These results validate the focus on skills and credentials in models of occupational mobility and confirm that firms impose costs on most job seekers willing to move to other occupations.

The contact gap also depends on characteristics of the target occupation itself. Beyond the within-occupation homogeneity already discussed, we examine features of the target occupation that the literature has shown to matter for occupational transitions, among them occupational licensing (Kleiner & Xu, in press) and labor market tightness (Modestino et al., 2016, 2020; Lochner et al., 2021; Le Barbanchon, Hensvik, & Rathelot, 2023; Ziegler, 2024; Cullen et al., 2025).The contact gap is roughly 45 percent larger in licensed than in unlicensed occupations, with a predicted gap of 10.3% versus 7.1% of the baseline contact rate, consistent with recruiters treating a formal credential as a hard requirement they can verify rather than a soft signal. We also find that recruiters encountering an increased scarcity of incumbent job seekers are substantially more likely to contact applicants from other occupations. While the mover gap is 8.8% of the baseline contact rate in the slackest occupations, it is only 3.4% in the tightest occupations, consistent with recruiters lowering the bar when incumbents are scarce.

Finally, we quantify the role of firms using variation in occupational labor market tightness. As markets tighten, both the contacts recruiters make and the hires firms complete shift toward movers relative to incumbents. We decompose this shift into a firm response – firms lowering the bar for movers – and a change in the composition of the applicant pool as incumbents become scarce. Firms' own screening accounts for about a third of the shift, and, reassuringly, the two data sources – clickstream contacts and realized hires – deliver almost the same firm share.

Our paper makes three 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, 2025). Our analyses also 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 occupational tightness also reduces firms' hesitation to hire movers, thereby creating opportunities for job seekers to transition into that occupation.

Second, our work also 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, 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 show that movers do best when they target occupations that share many job requirements with their last occupation and, less obviously, when they favor occupations in which jobs within an occupation differ widely. They should also aim for occupations that are currently tight, consider returning to a previously held occupation, or earn an occupational credential.

The paper most closely related 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. They find that a long retraining program substantially reduces movers' callback gap relative to otherwise identical incumbents, an effect that is particularly large in tight labor markets. We complement their results by studying the application behavior of real candidates using an approach with sufficient statistical power to differentiate among various mover profiles. This lets us show why firms are reluctant to hire movers. Because we also observe the candidate pool facing a vacancy, we can quantify the roles of both market sides in actual occupational transitions.

Finally, our work also contributes to the literature on measures of occupational similarity, which are relevant for understanding labor market boundaries and labor reallocation. Previous studies use observed occupational transitions (Schubert et al., 2024; del Rio-Chanona et al., 2021; Schmutte, 2014; Belot et al., 2019, 2025), task overlap between occupations[2], surveys among workers (Gathmann & Schönberg, 2010), or vocational education and training curricula (Eggenberger et al., 2018) to measure occupational distance. In contrast, our measure 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 skill requirements and could be used to estimate the similarity of jobs between regions, industries, or companies. Our results demonstrate substantial variation in job requirement overlap within occupations and show that the mover gap is more than three times as large in homogeneous 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 of firms' role in occupational transitions. Section 3 describes the data and our new measure of job requirements overlap. Section 4 estimates the contact gap, including how it varies with the mover's profile and with characteristics of the target occupation. Section 5 examines how labor market tightness shapes the contact gap and quantifies the role of firms in occupational transitions. Section 6 concludes.

2 Screening Job Seekers from Different Occupations

This section develops a simple framework that formalizes recruiters' decisions when screening applicants from different occupations under uncertainty. The framework follows the logic of employer screening under statistical discrimination (Phelps, 1972): recruiters cannot observe a candidate's skills directly, so they assess an individual candidate in part through the average characteristics of her previous occupation. The framework also motivates our empirical strategy.

The candidate's previous occupation is an informative signal because of the way occupations organize work. A large literature assumes that occupations produce output

[2]Most 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, 2025; Macaluso, 2025; Lyshol, 2022), or the Operational Directory of Trades and Jobs ROME (Goos et al., 2019).

[3]Leping (2009) measures skill overlap between firms based on job requirements in vacancies from an online job-search platform in Estonia. A key advantage of our measure is that our skill and task classification considers thousands of categories, whereas Leping (2009) considers only three broad skill groups. Other related papers are Bloesch et al. (2022), who construct a measure of within-firm, across-position task differentiation based on US job posting data from Burning Glass Technologies (now Lightcast), and Djumalieva et al. (2018), who 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 is that our measure is also well-defined for the overlap between different jobs *within* an occupation.

through occupation-specific bundling of tasks that in turn require specific skills (Robinson, 2018).[4] A candidate's previous occupation therefore carries information about the skills she likely acquired on the job: the more its requirements overlap with those of the advertised position, the more of the position's requirements she is likely to meet. 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. The candidate's (unobserved) skill vector is $a \in \{0,1\}^K$, where $a_k = 1$ if the candidate possesses skill k. The number of required skills the candidate meets is[5]

$$x'a = \sum_{k=1}^{K} x_k a_k. \tag{1}$$

Under full information, the recruiter would contact the candidate whenever the matched skills clear a threshold, $x'a \geq \tau_s$.[6] This threshold is specific to the search s because it depends on the recruiter's selectivity and the candidates she can choose from, since a pool with many strong applicants might let her set a higher bar.

In practice, 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 $p(z)$. The expected number of requirements the candidate meets is

$$m(x, z) = \mathbb{E}[x'a \mid z] = \sum_{k=1}^{K} x_k\, p_k(z) = x'p(z), \tag{2}$$

and the recruiter contacts the candidate if $x'p(z) \geq \tau_s$. The decision now depends on how informative the signals are: the recruiter weighs each requirement by the chance that the candidate meets it. 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] This follows the classic theory of statistical discrimination in hiring, which

[4] Lazear (2009) makes this argument at the firm level: individual skills may be general, but their combination is often specific, so switching is 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; Bohm et al., 2024, among others).

[5] We weight all job requirements equally here, so that $x'a$ simply counts the skills the candidate matches. In a robustness check, we instead attach skill-specific weights and use $\sum_k w_k\, x_k a_k$, letting rare, distinctive requirements matter more than common ones (Section 3.3). All of 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.

[6] The additive form of $x'a$ assumes that a candidate can compensate for a missing requirement by meeting others. In practice, some requirements may be perfectly substitutable, so that recruiters 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 percent larger in licensed than in unlicensed occupations (Section 4.3).

[7] Recruiters could instead hold inaccurate beliefs, for instance exaggerating the skill differences between movers and incumbents. Bohren et al. (2025) show that such inaccurate statistical discrimination

also assumes that employers' beliefs are correct (Phelps, 1972; Aigner & Cain, 1977), and it is consistent with our results below, where measurable signals of candidates' skills explain recruiters' contact decisions well.

Employers can learn from the requirements of an applicant's previous job. Suppose that the likelihood that a candidate possesses a skill is π_1 if the job required the skill and $\pi_0 < \pi_1$ otherwise. Here, π_0 reflects the candidate's general quality, because more skilled job seeker is more likely to possess any required skill. $\pi_1 - \pi_0$ captures the return to applying to a position that requires the same skill as in the previous job, which we assume to be common across candidates. In practice, it is unlikely that employers know the requirement profile on an applicant's prior job, but the requirement profile of the candidate's previous occupation, $\bar{x} \in [0,1]^K$, is both observable and can act as a key signal. Its k-th entry $\bar{x}_k$ is the probability that the candidate's previous job required skill k. Since the recruiter knows only that a share $\bar{x}_k$ of jobs in the previous occupation require skill k, 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 a baseline and an overlap term,

$$x' p(\bar{x}) = \pi_0 |x| + (\pi_1 - \pi_0) x' \bar{x}, \tag{3}$$

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. The overlap term is the additional fit that the previous occupation provides. Contact becomes more likely as the overlap $x'\bar{x}$ rises.

Some candidates the recruiter sees will be incumbents to the vacancy occupation. Equation (3) applies also to those candidates. Let $\bar{x}^I$ denote the requirement profile of an incumbent's previous occupation, which is the occupation the recruiter searches in, and $\bar{x}^M$ that of a mover. Because jobs *within* an occupation also differ in their requirements, the entries of $\bar{x}^I$ lie below one, so even an incumbent's expected overlap falls short of the vacancy's requirements, $x'\bar{x}^I < |x|$. How far it falls short depends on the within-occupation homogeneity of the occupation in question. In occupations where jobs share many requirements, the entries of $\bar{x}^I$ that the vacancy requires are close to one and the advantage over movers is larger because her expected overlap is near-complete.[8]

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 an identical return to overlap

is hard to distinguish from accurate statistical discrimination or taste-based discrimination without data on true productivity differences.

[8]Our empirical measure implements the expected overlap by averaging the Jaccard similarity in requirements of 5,000 randomly drawn pairs of job postings (Section 3.3). Averaging over random posting pairs mirrors the expectation over jobs within an occupation. The Jaccard index additionally normalizes the number of shared requirements by the number of distinct requirements in a pair, which holds differences in the length of postings fixed.

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

$$m^M - m^I = \underbrace{(\pi_0^M - \pi_0^I)\,|x|}_{\text{selection}} + \underbrace{(\pi_1 - \pi_0)\left(x'\bar{x}^M - x'\bar{x}^I\right)}_{\text{occupation}}, \tag{4}$$

the sum of a selection and an occupation term. The occupation term reflects the difference in overlap between movers and incumbents. It is negative on average, since a mover overlaps the vacancy less than an incumbent. The selection term reflects any difference in general quality between the two groups, for instance if lower ability candidates are more likely to switch occupation. In that case, the selection term would be negative and widen the gap.

The recruiter's assessment also has a random component that the signals do not capture, so contact is more likely the higher the expected match: $\Pr(\text{contact} \mid \bar{x}) = F\big(x'p(\bar{x}) - \tau_s\big)$ for an increasing function F. The contact gap between movers and incumbents is then

$$\delta = F\big(m^M - \tau_s\big) - F\big(m^I - \tau_s\big). \tag{5}$$

Since F is increasing, δ has the same sign as $m^M - m^I$ and inherits both the selection and the occupation term of Equation (4).

The contact gap in Equation (5) is the object of interest in our regressions. We regress a contact indicator on a mover indicator. The coefficient recovers the average gap δ from Equation (5). Our main specification includes candidate fixed effects, which hold constant each job seeker's baseline probability of possessing any required skill π_0. Since we compare the same candidate as a mover and as an incumbent, the individual fixed effects absorb most of the selection term in Equation (4).[9] The specification also contains search fixed effects, which absorb everything common to the candidates in a single result list, including the threshold τ_s.

The framework makes predictions about the factors that affect the gap. The gap tracks the difference in overlap, so it should shrink as a mover's overlap with the vacancy rises. It should also be larger when the vacancy occupation is homogeneous, since a typical incumbent then fulfills more job requirements.

A further prediction concerns the threshold τ_s, which also depends on the pool of candidates available to a recruiter. When incumbents are scarce, a recruiter who needs to fill the position may lower τ_s. Lowering the bar raises the probability of contact for movers and incumbents alike, but possibly not by the same amount. If many incumbents already sit above τ_s and would be contacted even at the original bar, lowering the bar

[9] Strictly, the selection term $(\pi_0^M - \pi_0^I)\,|x|$ scales candidate quality by the number of requirements of the vacancy and is therefore a candidate-by-search interaction. Additive fixed effects do not fully absorb such an interaction: candidate fixed effects hold π_0 constant and search fixed effects hold $|x|$ constant, but not their product. The occupation-specific valuation controls introduced in Section 4.1, which interact the vacancy occupation with candidate characteristics, address the remaining variation.

does little to bring in more of them. Movers, by contrast, are more likely to sit just below the original threshold. The lower threshold may thus disproportionately affect movers.[10] In this case, the contact gap is lower in tighter markets.

A final prediction concerns signals of occupation-specific fit beyond the last occupation. In reality, recruiters see more of a candidate's profile than the occupation she comes from: candidates list skills, credentials, and their work history, each of which may speak directly to their fit with the advertised position. A somewhat richer framework incorporates such signals.[11] Each signal of this kind raises the probability that the candidate possesses the vacancy's required skills even when her previous occupation did not require them. The two key signals recruiters observe on Job-Room are a professional degree in the vacancy occupation and prior work experience in it.[12] The richer framework thus predicts that these credentials raise contact rates beyond what overlap alone predicts.

3 Data and Measurement

Our analyses draw on three data sources: clickstream data from a recruitment platform, job seeker application diaries, and online vacancy postings. This section describes each in turn and introduces our measure of job requirements overlap.

3.1 Recruiter Click 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 were collected between March and 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 show that recruiters discount male job seekers who wish to work part-time.

The candidate profiles on the platform belong to job seekers registered with the PES, and the platform covers about 80% of all registered job seekers. Registration is open to everyone, but mandatory for UI recipients. 61% of registered job seekers are UI recipients. Recruiters are identified by a cookie installed in users' browsers when they first visit the

[10]Formally, lowering τ_s narrows the gap whenever $F'(m^M - \tau_s) > F'(m^I - \tau_s)$, that is, whenever more movers than incumbents are close to the margin of being contacted.

[11]Formally, let $e \in \{0, 1\}$ indicate such a signal and suppose it raises the probability that the candidate possesses each skill the vacancy requires: $p_k = \pi_1 \bar{x}_k + \pi_0 (1 - \bar{x}_k) + \lambda e$ with $\lambda > 0$. The expected match becomes $x'p = \pi_0 |x| + (\pi_1 - \pi_0) x'\bar{x} + \lambda e |x|$. Because the signal is specific to a candidate–occupation pair, it varies across searches for the same candidate, so candidate fixed effects do not absorb it.

[12]While we define movers as candidates who *last* worked in a different occupation, some of them might have work experience in the vacancy occupation from earlier employment.

platform after December 2016.[13]

Recruiters usually start by entering the occupation they are hiring for, and most also 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 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." The result list also shows the content of an unstructured text field that contains additional "knowledge, abilities, skills" of job seekers. This information is specific to recruiters' searched occupation. It is only visible if job seekers have filled it out with their caseworkers at the PES. If interested, the recruiter can open the candidate's full profile.

Figure 1 provides a screenshot of the full candidate profile.[14] 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), a potential professional degree or training in that occupation, and whether that degree is Swiss, foreign but recognized in Switzerland, or foreign and not recognized. In addition, recruiters may again see the text field with a job seeker's "knowledge, abilities, skills."

A candidate appears in a recruiter's results whenever the vacancy occupation is one of the occupations listed on her profile. This list of occupations is defined during the initial meeting between the job seeker and her PES caseworker. 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 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.[15]

Our main dependent variable is whether the recruiter clicks the "Show contact details" button on a candidate's profile. In 2017, it was not possible to contact and hire a candidate through Job-Room without clicking this button. Hangartner et al. (2021)

[13] What 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.

[14] Appendix Figure A.1 shows a Screenshot of the original, German version of the profile.

[15] We 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.

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

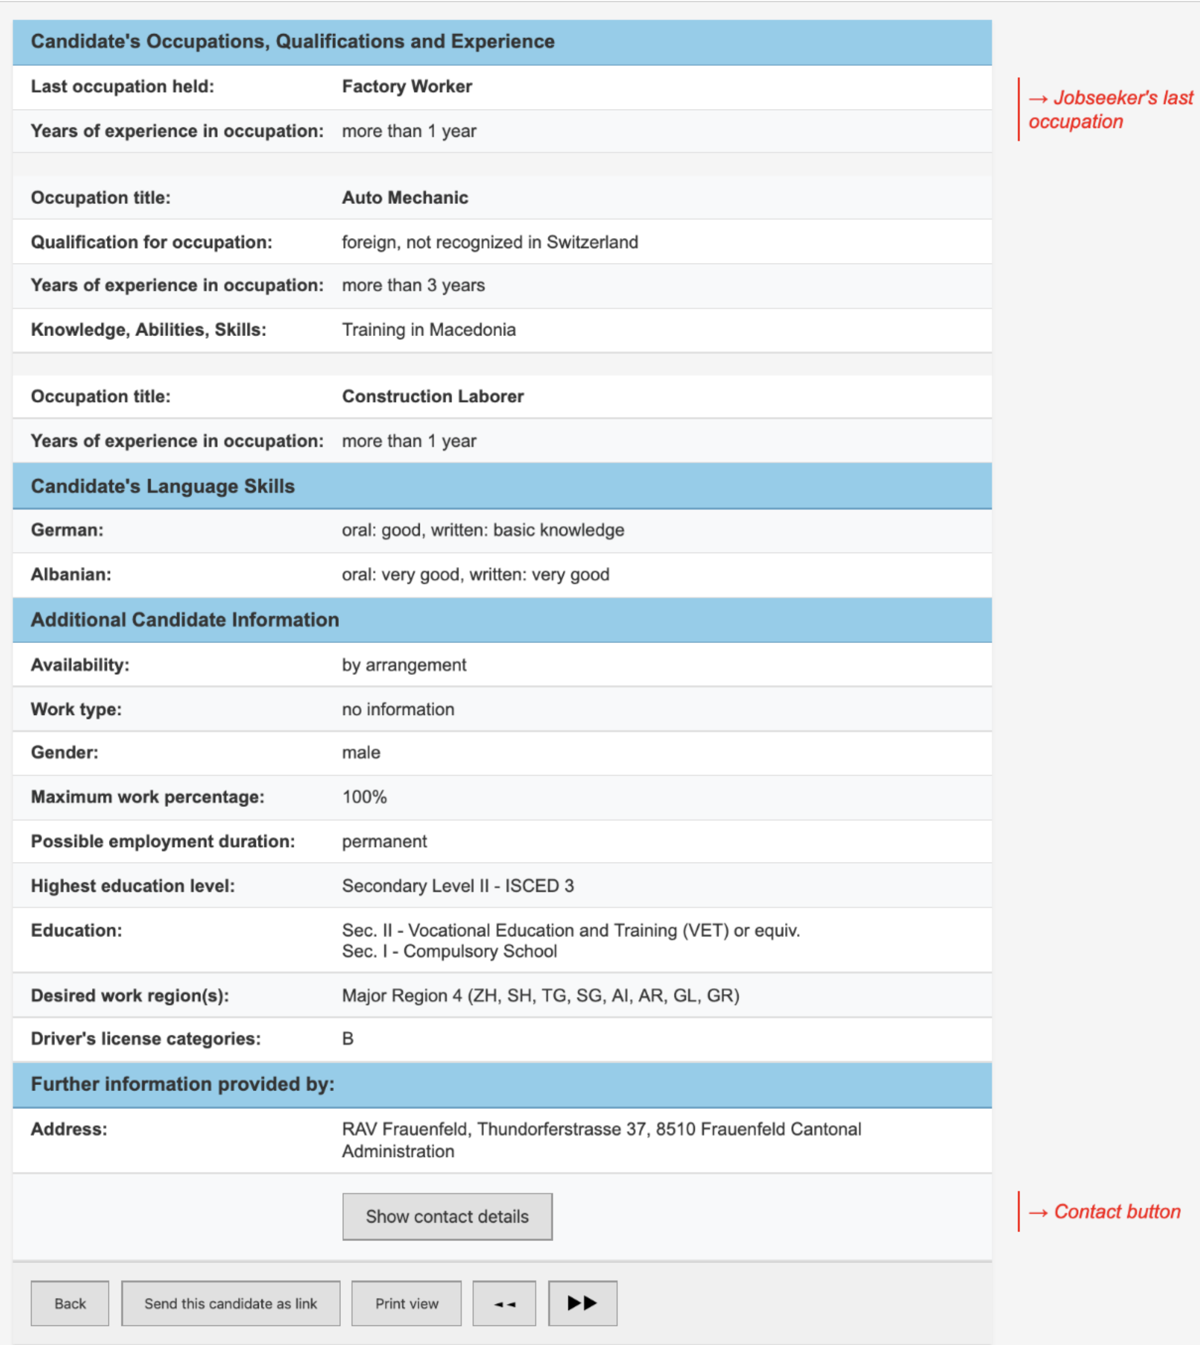

Candidate's Occupations, Qualifications and Experience	
Last occupation held:	Factory Worker
Years of experience in occupation:	more than 1 year
Occupation title:	Auto Mechanic
Qualification for occupation:	foreign, not recognized in Switzerland
Years of experience in occupation:	more than 3 years
Knowledge, 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

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

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 a small number of sample restrictions. The core restriction retains only searches in which the recruiter specified an occupation, since otherwise we cannot tell whether a candidate's last occupation matches the search; this removes 62,777 of 469,779 searches. We also require that candidate profiles report the last occupation, drop searches with missing search-identifier tokens, and exclude the few searches in military occupations. The analysis sample comprises 15,266,233 result-list items from 406,873 recruiter searches.

Table 1 reports descriptive statistics. The top panel shows that the average job seeker lists 2.3 occupations she is willing to work in, with 95% listing their last occupation among them. When comparing incumbents (Column (2)), who last worked in the vacancy occupation, with movers (Column (3)), we note that the two groups are quite similar. Incumbents are slightly younger and more experienced in the vacancy occupation, but many movers also have experience in that occupation, seeking to return to a position they previously held ("return movers"). The middle panel shows that recruiters find, on average, 37.5 candidates per search, of which 19.6 are incumbents. They contact 11% of the candidates in their result lists. Notably, contact rates differ by candidates' occupational background: recruiters contact 12.24% of incumbent job seekers and only 9.64% of movers, which implies an unconditional contact gap of 27%. This gap varies substantially across occupations (Appendix Figure A.2). The bottom panels describe the sample: 37,665 recruiters conduct 406,873 searches and see 325,545 distinct job seekers. Recruiters search disproportionately in craft and trade occupations, but we observe several thousand searches in all major occupational groups (Appendix Table A.1).

3.2 Online Application Diaries

The platform data enables us to analyze the causal effect of job seekers' most recent occupation on recruiters' contact rates. However, it also has two drawbacks. First, the data 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. We therefore supplement the platform data with applications recorded in the application diaries of registered unemployed job seekers, which capture actual applications and their hiring outcomes. 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 directly from the posting. This links each ap-

Table 1. Recruiter clicks: descriptive statistics

	(1) All job seekers	(2) Incumbent job seekers	(3) Job seekers from a different occupation
Job seeker characteristics:			
Number of occupations listed by jobseeker	2.32	—	—
Probability to list last occupation (%)	94.58	—	—
Share female (%)	41.52	42.97	40.59
Average age	39.32	38.71	39.71
Share by highest education (%):			
Tertiary	20.45	20.54	20.38
Upper secondary	51.69	51.50	51.81
Lower secondary	22.06	21.32	22.53
Experience in searched occupation (%):			
More than 3 years	72.16	79.20	64.03
1 to 3 years	16.21	14.31	18.41
Less than 1 year	5.88	4.18	7.84
Experience in last occupation (%):			
More than 3 years	69.73	77.09	61.69
1 to 3 years	19.13	15.33	23.28
Less than 1 year	8.11	4.83	11.68
Recruiter search results:			
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
Count data:			
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 vacancy 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.

plication to the official vacancy identifier. Applications documented through this import function account for 34.7% of all applications documented by the job seekers who use it. We link an application to a hiring outcome by matching the employer named in the UI register against the firms in the diary, a string-matching procedure that achieves 91.4% balanced accuracy on a hand-labeled test set (see Appendix A.1). In our final sample, job seekers have a 1.2% probability of being hired conditional on applying.[16]

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. For these job seekers, we observe 167,024 applications to 96,611 distinct vacancies. After dropping applications with an unmatched employer or an incomplete overlap measure, control or fixed-effect key, our analysis sample is 166,029 applications. The analyses that condition on labor market tightness use the subset of 159,061 applications for which the vacancy's 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 using the Job-Room functionalities are more likely to change occupation upon reemployment compared to the population of job seekers (58% vs. 53%). Apart from that, however, our application diary sample mirrors the characteristics of the registered job seeker population relatively closely.

Descriptively, the application protocols confirm that firms act as gatekeepers of occupational mobility. Movers make up the bulk of the applicant pool: 69.99% of applications come from job seekers whose last job was in an occupation different from the vacancy's. Their applications, however, are markedly less likely to succeed. A mover's application results in a hire 1.12% of the time, against 1.51% for an incumbent applying within her own occupation, so an incumbent's application is 35.0% more likely to end in a hire. In joint terms, 0.78% of all applications are mover hires and 0.45% are incumbent hires; movers account for more hires per application only because they are far more numerous, not because they fare better once they apply. Had movers and incumbents converted applications at the same rate, the pooled rate of 1.23%, movers would have generated 0.49 percentage points more hires per application than incumbents, against the 0.33 points actually observed. Firms' screening therefore offsets a third of the mover advantage that the composition of the applicant pool would otherwise produce.

3.3 Job Requirements Overlap: A New Index

The framework highlights the key role of skill overlap in determining firms' reluctance to hire movers. Testing its predictions requires measuring the overlap of skills between

[16] This number exceeds the 0.9% job finding rate in Zuchuat et al. (2023), who use manually coded data from Swiss application protocols from 2012 and 2013 combined with direct information from the protocols on whether an application was successful. This comparison suggests that our matching procedure successfully identifies most matches.

two jobs in two distinct occupations, and comparing it with the overlap between two jobs within the same occupation. As a consequence, we 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 et al., 2025; Eggenberger & Backes-Gellner, 2023), as these assume that all jobs within an occupation are alike. 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 no viable option either. Building on a few previous studies (Leping, 2009; Djumalieva et al., 2018; Bloesch et al., 2022; Dupuy et al., 2024), we construct a measure of skill overlap based on the overlap in job requirements across job postings, using vacancy data that include requirements extracted from more than 1 million postings. The data cover the near-universe of online job postings in Switzerland between 2016 and 2022 (see Appendix C.1 for details). They assign each job ad to a four-digit ISCO occupation and include a highly detailed list of job requirements extracted from the ad text and the job title.

A concern when using job vacancy data to measure skill 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 various concrete job tasks. Vacancy data are less likely to record basic skills, the foundational skills that facilitate skill acquisition and are likely relevant for hiring. We thus compute overlap across all job requirements an ad lists—that is, all listed tasks, skill requirements, and required educational degrees. Professional certificates and credentials listed on the job ad are more likely to capture the basic skills relevant to the position. We include the listed tasks both because a successful candidate must likely possess the required skills to perform them and because distinguishing skill requirements from job tasks in postings is technically challenging.[17]

The computation of the index is straightforward (see Appendix C.2 for further details and Figure C.1 for a stylized example). For a given pair of occupations, we first draw 5,000 random pairs of postings and compute the Jaccard similarity for each 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. The identical procedure applies to two postings in the same occupation and to two postings in different occupations, which places within- and between-occupation overlap on a common scale. Since not every shared requirement is equally informative about how interchangeable two jobs are, we also compute an alternative version that weights requirements using their inverse document frequency (see Appendix F for details).

[17]Job postings often detail the duties and tasks of the open position, as well as the skill requirements for suitable candidates. For example, the word "programming" could be listed as both a skill requirement and a job task in a job posting.

Figure 2 shows the estimated job requirement overlap for the ten occupations with the most registered job seekers. As expected, overlap is higher within occupations than between them: it averages 5.5% between two distinct occupations and 9.8% within the same occupation.[18] To anchor this difference in overlap, the median ad lists 9 distinct requirements, and the median pair of ads shares about 1.0 of them (a 5.5% Jaccard). The 4 percentage-point difference thus corresponds to roughly 0.7 additional shared requirements. Another key takeaway is that within-occupation overlap varies substantially, from 35.6% for call center agents to 6.1% for manufacturing laborers. This means that jobs in certain occupations differ as much as jobs across occupations, suggesting that job requirements are far from homogeneous in certain occupations.

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

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 over 5,000 randomly drawn pairs of job postings from the two occupations; the diagonal shows the average over 5,000 pairs of postings drawn from the same occupation.

We validate the measure in two ways. First, Appendix Figure C.3 shows that our measure of job requirement overlap lines up with an external classification of related occupations from O*NET. Occupation pairs with less than 4% overlap are almost never classified as related by O*NET, whereas pairs with more than 10% overlap have about a 12% chance of being classified as related. Second, the measure strongly predicts job seekers' applications. Figure 3 relates how often a job seeker applies to an occupation over an unemployment spell to the overlap between that occupation and her most recent

[18] The distribution of overlap across all occupation pairs is in Appendix Figure C.2.

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 almost four 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). Below, we show that the index is highly predictive of recruiters' contact decisions, too.

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

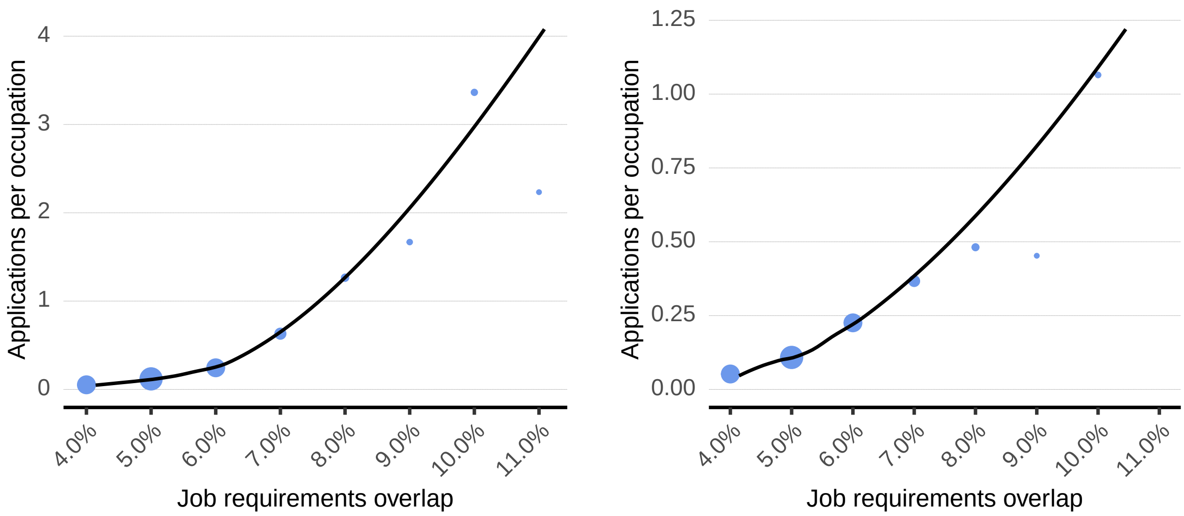

(a) All applications (b) Applications outside the last occupation

Notes: Panels (a) and (b) are binned scatter plots of the number of applications to an occupation over an unemployment spell against the average job requirements overlap between that occupation and the job seeker's last occupation. Each observation is a job seeker and an occupation she applied to. Dot size is proportional to the number of observations and the lines are local linear fits. Panel (a) uses all occupations; Panel (b) restricts to occupations different from the last one. 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 the recruiter clickstream data from Job-Room, the application diary data from the UI register, and the novel measure of job requirement overlap to analyze the size and determinants of the mover gap—the difference in the contact rate between movers and equally qualified incumbents. Section 4.1 introduces the research design. Section 4.2 then establishes the average mover gap and examines whether job seekers who attempt to switch occupations are positively or negatively selected. We then test the core predictions of our conceptual framework: the contact gap should shrink as the overlap between a candidate's last occupation and the vacancy rises, widen with the homogeneity of the vacancy occupation (Section 4.3), and close with occupation-specific experience and qualifications (Section 4.4). The framework's remaining prediction – that the gap should close with labor market tightness – is the subject of Section 5.

4.1 Empirical Approach

What is the contact gap for job seekers who come from a different occupation relative to otherwise similarly qualified incumbents? In the framework, this is the contact rate gap δ of Equation (5). The framework shows that the raw contact rate gap combines the occupation term of interest with a selection term. Answering the question thus requires accounting for the possibility that job seekers who attempt to switch occupations may be positively or negatively selected.

A key advantage of the Job-Room clickstream data is that it provides a means to address this selection issue. The idea is to leverage 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 identify the effect of a job seeker's occupational history using 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). By comparing the estimated mover gap with and without candidate fixed effects, we can also test whether those who wish to change occupation are positively or negatively selected.

We estimate the mover gap using variants of the following linear probability model:

$$y_{i,s} = \alpha \, \text{Mover}_{i,s} + \gamma_i + X'_{i,o(s)}\beta + \psi_{\text{rank}(i,s)} + \phi_s + \varepsilon_{i,s}. \tag{6}$$

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 key 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 average contact gap δ in Equation (5)—facing a job seeker from a different occupation relative to an otherwise similar incumbent. Since the specification in this section does not control for occupation-specific skills and credentials, such as experience or a degree in the vacancy occupation, the coefficient measures the contact gap for an average mover. We treat occupation-specific skills and credentials as determinants of the contact gap and analyze them in Sections 4.3 and 4.4.

The key ingredients in terms of causal identification of the mover gap in (6) 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 absorb all unobserved recruiter characteristics, the searched occupation and other search criteria, as well as the prevailing labor market conditions at the time of the search and the candidate pool, including the value of the candidate pool shown to the recruiter and the number of candidates the recruiter contacts. As a result, the search fixed effects control for the search-specific hiring threshold τ_s: we compare only job seekers evaluated against the same decision criteria. The candidate fixed effects, in turn, control for all candidate characteristics that influence contact rates and are constant across recruiter searches. They hold the profile content displayed to recruiters fixed, and with it the candidate's average attractiveness across searches. As a result, they absorb most of the selection term in Equation (4). What varies within a candidate across recruiter searches is her status as an incumbent or a mover.

There are two further elements in Equation (6). The control variables in the vector $X_{i,o(s)}$, added in an extended specification, account for a potential occupation-specific valuation of certain job-seeker attributes and qualifications. To give a concrete example, a recruiter looking for a sales assistant may value a candidate's good German strongly, whereas a recruiter looking for a kitchen helper may weight it much less. To construct controls that capture such occupation-specific valuations, we interact the three-digit vacancy occupation with the candidate's characteristics. 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 there is little difference in the result-list ranking of movers and incumbents.[19]

4.2 Baseline Results

Table 2 presents the baseline mover gap estimates. In this table and in the other candidate-level regressions on the platform data, we cluster standard errors two-way, by recruiter and by candidate. Regressions in the application diary data cluster two-way on the job seeker and the hiring firm.

The regression in 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 incumbent job seekers. The estimated mover gap amounts to -0.77 percentage points, or 7% of the average contact rate of 11%. To benchmark this magnitude, Column (4) reports a specification without candidate fixed effects, in which the effects of candidate characteristics remain identified: the mover gap is equivalent to approximately two-thirds of the contact premium for 'good' or 'very good' German skills, relative to little or no German skills, in the German-speaking part of Switzerland.

[19]Within a search result list, the share of mover candidates rises just slightly with rank, staying essentially flat up to the median list length of about 20 and increasing by only about 1.5 percentage points between rank 20 and 100 (Appendix Figure F.1).

In Column (2), we probe whether the mover gap partly reflects that candidates are incumbents precisely in the occupations that fit them, so recruiters may value job seekers' fixed characteristics more when they are incumbents and less when they are movers. To do so, we let each of the 64 profile characteristics listed in Appendix Table A.2 carry its own effect in every three-digit searched occupation. The estimated gap falls only from -0.77 to -0.71 percentage points: occupation-specific valuation of the characteristics recruiters see accounts for none of the mover gap.

Although the mover gap is quantitatively important, it is substantially smaller than the raw difference in contact rates between movers and incumbents in the Job-Room data (Table 1). This suggests that movers are, on average, negatively selected. We can test the nature of selection directly by dropping the candidate fixed effects from our main specification. As the comparison of Columns (1) and (3) shows, the gap is twice as large without the candidate fixed effects. In the framework's terms, the selection term in Equation (4) is negative: the raw difference between movers and incumbents partly reflects lower candidate quality.

Three tests support the selection-on-observables assumption behind the design. First, we find largely unchanged estimates when replacing the candidate fixed effects, which control for observed and unobserved candidate characteristics, with control variables that account for all candidate information visible to recruiters on the Job-Room profiles (Column 4). This result supports the selection-on-observables assumption because the assumption implies that recruiters' contact decisions should depend solely on observable candidate characteristics. Second, the gap is unchanged when we exclude candidates with entries in the unstructured text field about their knowledge, abilities, and skills (Column 5). This suggests that the hard-to-codify field does not materially drive the estimate. Third, Column 5 also shows that a candidate's age, which we observe in the UI register but recruiters do not see, has no effect conditional on the observables, consistent with recruiters acting only on information from the displayed profile.[20]

The estimate is robust along several further dimensions. First, we find economically and often statistically significant mover gaps in almost all two-digit ISCO occupations (see Appendix Figure F.2), with particularly large barriers to entering jobs as health professionals, information and communication technicians, personal care workers, and handicraft and printing workers. Second, the gap is virtually unchanged when we define a mover at the three-digit rather than four-digit ISCO level (Appendix Table F.1, Column 2). Third, when we restrict the sample to candidate profiles the recruiter either never opened or viewed for at most 20 seconds—decisions in which recruiters are unlikely to have gathered outside information about candidates—the gap is slightly smaller but

[20]This age placebo uses the sample without a filled 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***	-0.0070***	-0.0151***	-0.0085***	-0.0082***
	(0.0003)	(0.0003)	(0.0004)	(0.0003)	(0.0004)
Good German skills				0.0122***	0.0116***
				(0.0009)	(0.0014)
Age = 30-44					-0.0004
					(0.0007)
Age = 45-59					0.0006
					(0.0006)
Age = 60+					0.0005
					(0.0012)
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	0.46506	0.43919	0.44657	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 (6). 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. Candidate fixed effects absorb all fixed characteristics of a job seeker; recruiter search fixed effects are fixed effects for an individual search by a recruiter, so that the gap is identified from candidates appearing in the same result list. Search rank fixed effects are fixed effects for the absolute rank and the relative rank of a job seeker within the list of search results. "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. This regression controls for the impact of good German skills for 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.

still highly significant (Appendix Table F.1, Column 3).[21] Fourth, 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, if anything, rises in relative terms (Appendix Table F.1, Column 4). Finally, Appendix Table F.2 shows that the results are virtually identical when controlling for the characteristics and qualifications of job seekers ranked in the immediate vicinity of the focal candidate. The similarity limits concerns that the effects may be confounded by recruiters primarily contrasting candidates placed adjacent to one another in the result list.[22]

Overall, the findings show that movers face a significant contact gap of 7% relative to otherwise similar incumbents. Accounting for candidate quality halves the gap, indicating that movers are negatively selected. In the next sections, we test the key mediators of the mover gap suggested by our conceptual framework. These analyses shed light on why firms hesitate to hire movers.

4.3 Job Requirements Overlap and the Mover Gap

Our conceptual framework posits that the primary signal a recruiter reads from a candidate's last occupation is how much the job requirements of that occupation overlap with those of the vacancy. Specifically, the framework predicts that the mover gap depends on two factors: the within-occupation homogeneity of the vacancy occupation, which governs how much being an incumbent signals about meeting the vacancy's requirements, and the expected between-occupation overlap $x'\bar{x}$ between the candidate's last occupation and the vacancy.

Figure 4 tests both predictions by enriching the regression model in Equation (6) with the overlap measures developed in Section 3.3. Each panel is based on a separate regression.

Panel (a) of Figure 4 shows the mover gap as a function of the within-occupation overlap in the vacancy occupation. It uses bins that are one percentage point wide and endpoint bins that capture values below 7% and at 12% or more. The figure reveals that the mover gap is markedly larger when the occupation becomes more homogeneous. In occupations in which jobs share many requirements, movers' disadvantage is large. In our framework, this arises because incumbents on average meet many of the position's requirements. In occupations where jobs share few requirements, the mover gap is smaller,

[21] This 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.

[22] In Appendix Table F.2, we control for the leave-one-out mean of each job-seeker characteristic and the mover indicator itself across a centered window of the X neighboring candidates listed just before and just after the focal candidate. These controls test whether the immediate neighborhood composition matters, above and beyond the average competitor quality controlled for by the search fixed effects. We find that the contact gap is essentially unchanged across specifications.

as even incumbents' expected fit is incomplete. Focusing on the endpoints of Panel (a), the mover gap in the most homogeneous occupations is about 3.3 times that in the most heterogeneous occupations. These results indicate that job requirements vary within occupations and that recruiters act on this variation. These findings are in line with other empirical evidence, 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) tests whether the mover gap reflects the shortfall in overlap between the candidate's last occupation and the vacancy. Motivated by the framework,[23] we collapse the two overlap dimensions into one: the overlap gap to incumbents in percentage points, defined as the within-occupation overlap of the vacancy occupation minus the between-occupation overlap of the candidate's last occupation with the vacancy. We observe that the contact gap shrinks steadily as a candidate's overlap approaches the vacancy's within-occupation overlap. In the rightmost bin—movers who, on average, meet the same share of the vacancy's job requirements as incumbents—the gap is only 1.4% of the baseline contact rate. This estimate is 80% smaller than the average mover gap, suggesting that job requirements overlap accounts for almost the entire gap.[24]

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 requirements (Appendix Figure F.4, Appendix Table F.7, and the comparison in Appendix Figure F.5). Under this weighted measure, a lack of overlap accounts for the entire mover gap: movers who meet the same share of the vacancy's job requirements as incumbents face no detectable contact gap.

The results so far establish that overlap matters for the mover gap, but not whether it matters relative to other characteristics of the target occupation. To assess this, we estimate the gap jointly against four characteristics of the vacancy occupation: the tightness in that occupation, whether it is licensed, its wage level, and the typical education of incumbents.[25] Appendix Figure F.3 reports how the mover gap varies with the 5 features and Appendix Table F.3 decomposes each feature's effect into its own effect and the correlation with the other four features, following Gelbach (2016). The overlap gap

[23]In 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 (4)). The framework predicts that the occupation component of the mover gap vanishes as this difference approaches zero.

[24]We arrive at almost the same estimate when computing the reduction due to overlap in a different way: adding flexible controls for the six overlap-gap bins to the baseline specification reduces the mover coefficient from -0.77 to -0.16 percentage points, a reduction of roughly 80% (Table F.8, Columns 1 and 2).

[25]Tightness 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, constructed from data external to Job-Room; Appendix D gives the details. Appendix E describes how licensing, the wage level, and the typical education of incumbents are constructed.

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

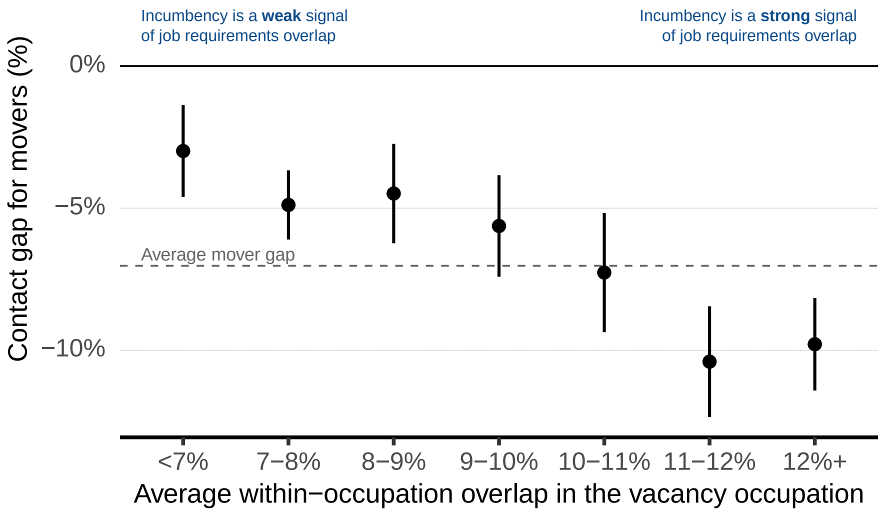

(a) By within-occupation homogeneity of the vacancy occupation

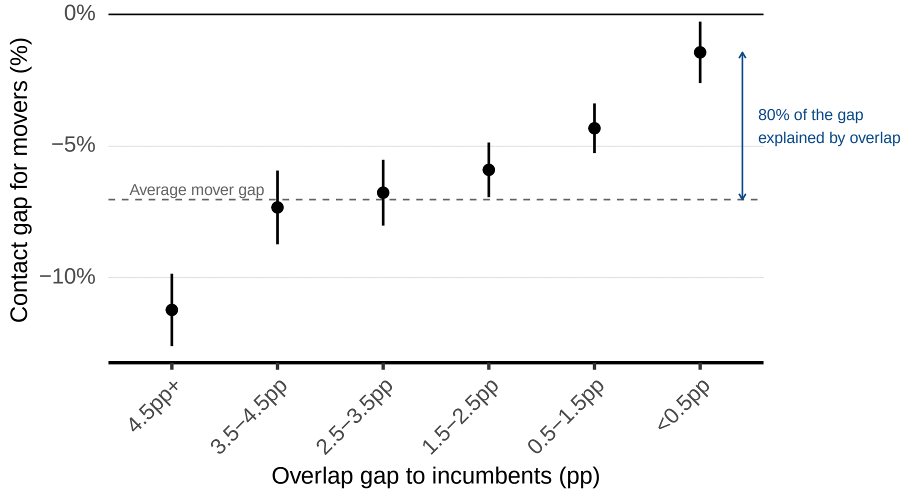

(b) By the overlap gap to the vacancy occupation

Notes: The figure plots the contact gap for movers, as a percent of the baseline contact rate, estimated from Equation (6) enriched with the job requirements overlap measure. Each observation is a candidate profile appearing in a recruiter's result list, and the outcome is whether the recruiter clicks the button revealing the candidate's contact details. Panel (a) interacts the mover indicator with indicators for bins of the within-occupation overlap of the vacancy occupation, in percentage points. Panel (b) adds bins of the overlap gap of the mover to incumbents to the mover indicator, so that each bin's gap combines the mover and the bin coefficient, the gap defined as the within-occupation overlap of the vacancy occupation minus the between-occupation overlap of the candidate's last occupation with the vacancy, in percentage points. The dashed line marks the average mover gap from the baseline regression of Section 4.1. Both regressions include candidate, recruiter-search, absolute-rank, and relative-rank fixed effects. 95% confidence intervals from standard errors clustered two-way on the recruiter and the candidate are shown. Appendix Table F.7, Columns (1) and (3), report the underlying coefficients.

remains first-order when the other features are held fixed.[26] The only feature that has a comparable moderating effect on the mover gap 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 roughly 45 percent larger in licensed occupations. The estimated influence of licensing gets stronger once the others are introduced into the regression, because licensed occupations are disproportionately tight ones, 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 that association is explained by the other features: higher-paying occupations are tighter, and occupations with university-educated incumbents are both tighter and less homogeneous in their job requirements, so the overlap gap to incumbents is smaller there and incumbency carries less of an advantage. Both change the gap little once the other features are accounted for.

There are two potential concerns regarding the external validity of this platform evidence on the importance of overlap. First, contact attempts are an early step in the recruitment process: mover gaps may reappear at the hiring stage even where they are absent at the screening stage. Second, recruiters who search for candidates actively may not be representative of firms more broadly. They may be more open to considering job seekers with a lower fit to the position. Next, we therefore 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 the diary data mirror our empirical design using the Job-Room data as closely as possible. The unit of observation is an application: a registered job seeker applying to a hiring firm for a position in a given occupation. We regress an indicator of whether the application resulted in a hire on the mover indicator and the overlap categories. Instead of search fixed effects, we hold the occupation and canton of the vacancy fixed and additionally control for fixed effects for the hiring firm, so that we compare applications received by the same firm. Instead of job seeker fixed effects, we add fixed effects for the job seeker's last occupation at the two-digit level to absorb differences across occupations of origin, and we control for a rich set of job seeker characteristics.[27]

[26] The other features enter the regression as few dummies, in order to not allow for a more flexible form for the overlap gap we restrict it to be linear instead of re-using the bin-indicators from above. Moreover, the overlap gap describes the match between the mover's last occupation and the vacancy, so one could expect that it is more predictive than the other, purely vacancy-side, characteristic. However, replacing the overlap gap with the within-occupation homogeneity of the vacancy occupation, which does not depend on the job seeker, leaves the picture unchanged: the wage and education gradients still vanish, to a large extent through the same channels (Appendix Table F.4).

[27] The job seeker controls comprise 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 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

The diary analyses depart from the platform design in two respects. First, because a hire is a rare outcome—the mean hire rate is about 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: the average change in an application's predicted hire probability when mover status in a given category is switched on, expressed as a percent of the mean hire rate. Second, the application data lack the platform's within-candidate variation, so the estimates rely on a selection-on-observables assumption. It is less plausible that this assumption holds in the diary than in the platform data: two observationally equivalent applicants can differ along dimensions unobserved to us but observed by recruiters. Since movers are likely negatively selected, the estimated mover gaps may be inflated. We therefore interpret the diary estimates as reflecting both the occupation component and part of the selection term in Equation (4), and abstain from interpreting the estimated coefficients causally.

Figure 5 presents the results, which suggest that within- and between-occupation overlap appear to shape the mover gap at the hiring stage as well. As on the platform, applications from movers are significantly less likely to result in a hire than applications from observationally similar same-occupation candidates. Although the estimates are considerably less precise than on the platform, they display the same overlap gradients. In Panel (a), the hiring gap for movers in occupations with above-median homogeneity is about 1.4 times as large as the gap for movers in more heterogeneous occupations. In Panel (b), applicants from occupations with an overlap gap of half a percentage point or more face a hiring gap of about 49%, roughly twice the gap of about 24% for applicants whose overlap comes within half a point of incumbents'. Both gradients are largely independent of the specification (see Appendix Tables F.5 and F.6): they change little as we add a last-occupation fixed effect and hiring-firm fixed effects, or narrow the comparison to applicants to the very same vacancy. While the two effects in Panel (b) are significantly different from each other, this is not the case in Panel (a). However, if we restrict the gradient to be linear, the slope is negative and statistically significant at least at the 10%-level in all specifications that hold the vacancy canton and occupation fixed (Appendix Table F.5). In general, to the extent that selection on unobservables resembles selection on observables, this stability suggests that omitted candidate characteristics are unlikely to drive the gradients (Oster, 2019).

In sum, we find strong evidence in the platform data, and consistent patterns in the application diary data, that the lower likelihood to meet job requirement is a key 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 tracks job requirements overlap

zero for job seekers who were not employed in the respective year, so that their level is absorbed by the corresponding employment indicator.

Figure 5. Application diaries: the hiring gap for movers by overlap

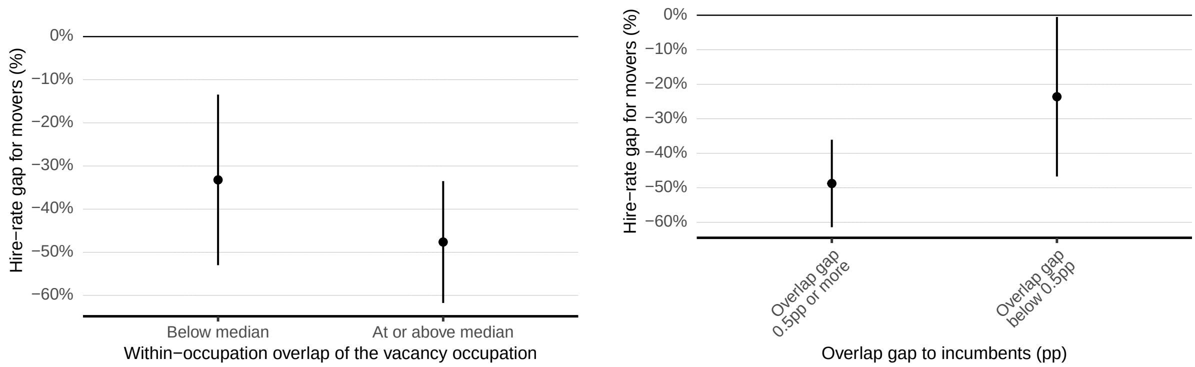

(a) By within-occupation homogeneity of the vacancy occupation

(b) By the overlap gap to the vacancy occupation

Notes: The figure replicates Figure 4 in 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. Each point is the hire-rate gap between movers in the given category and same-occupation applicants, expressed as a percent of the mean hire rate. The gaps are average marginal effects, holding all other characteristics fixed at their mean effects. They come from logistic regressions that include fixed effects for the vacancy occupation by canton, for the job seeker's last occupation at the two-digit level and for the hiring firm, and that control for 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. Panel (a) interacts the mover indicator with the within-occupation overlap of the vacancy occupation, split at its median. Panel (b) adds the overlap gap to the mover indicator, split at half a percentage point. Whiskers are 95% confidence intervals from delta-method standard errors clustered two-way on the job seeker and the hiring firm. The underlying regressions are in Appendix Tables F.5 and F.6, Column (5).

continuously and vanishes for movers whose overlap matches incumbents', suggesting that recruiters use occupation as a signal of underlying skills rather than as a credential in itself.

4.4 Occupational Experience, Qualifications, and the Mover Gap

Job requirements overlap does not exhaust the signals of occupation-specific fit that recruiters observe on Job-Room. As the framework highlights, candidate profiles contain two further such signals: prior work experience in the vacancy occupation and a professional degree in it. An enriched framework predicts that these signals raise contact rates beyond what overlap alone predicts (Section 2).

To analyze this prediction, we now estimate contact gaps separately for nine mover profiles. The profiles combine three groups defined by occupational experience and qualification (no experience in the vacancy occupation, experience without a professional degree, and experience with a degree) with three categories of the overlap gap to incumbents (overlap gap $\geq$2.5pp, between 0.5 and 2.5pp, and $<$0.5pp). The full regression is presented in Column 4 of Appendix Table F.8. Figure 6 reports the contact gap for the nine resulting mover profiles, each relative to the average incumbent.

The figure shows that overlap, experience, and qualifications each close the mover gap independently, and it confirms the dominant role of overlap. Within each experience-qualification group, the mover gap shrinks sharply 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 (mover types (1) vs. (3)). Return movers—those with prior work experience in the vacancy occupation—also face smaller gaps relative to incumbents. For instance, prior occupational experience reduces the gap for movers from distant occupations from 17.5% to 10.6% (mover types (1) vs. (4)). Return movers with a professional qualification in the vacancy occupation (mover types (7)–(9)) have the lowest mover gaps. In fact, for a return mover with a qualification who comes from a very close occupation, the contact gap against the average incumbent is virtually zero: the point estimate is 0.3% and statistically indistinguishable from zero (mover type (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. As discussed in Section 3.1, candidates on Job-Room primarily list occupations in which they have work experience. Job seekers' occupational search scope is thus narrower on Job-Room than in the application diaries. Indeed, the diary shares reveal that many movers attempt relatively distant switches: 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.

More generally, the substantial differences between the two share columns in Figure 6 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 Job-Room data understate the average mover gap? The answer is no, because a second compositional difference works in the opposite direction: the platform overrepresents searches in occupations that are comparatively difficult 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, the bottom row of Figure 6 reports a mover gap from a regression that reweighs the platform observations so that their joint distribution over the nine mover profiles and two-digit vacancy occupations matches the distribution of applications in the application diaries. The reweighted mover gap is similar in magnitude to the baseline estimate (-5.9% versus -7.0%). The two compositional differences thus largely offset each other, and the platform data do not materially misstate the average mover gap in the wider labor market.

Figure 6. Recruiter clicks: the contact gap for nine mover profiles

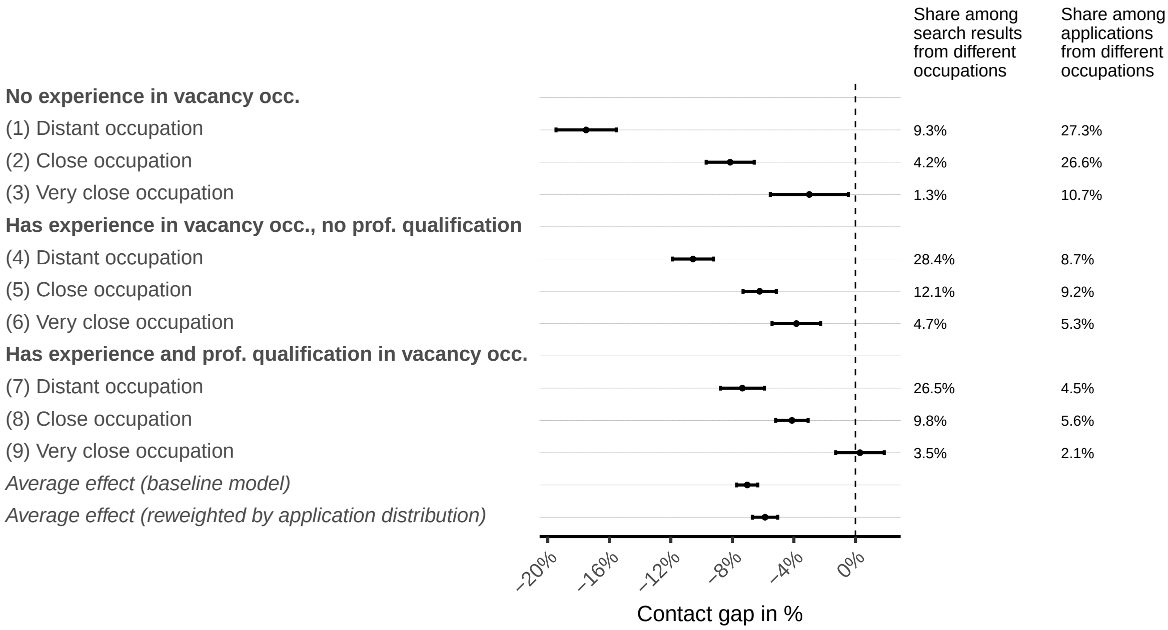

Notes: The figure reports the contact-rate gap, as a percent of the baseline contact rate, for nine mover profiles defined by three experience-and-qualification tiers and three overlap categories (distant, at least 2.5 points below the vacancy's within-occupation overlap; close, 0.5 to 2.5 points; very close, less than 0.5 points). Each observation is a candidate profile appearing in a recruiter's result list, and the outcome is whether the recruiter clicks the button revealing the candidate's contact details. Each gap is relative to the average incumbent, whose experience-and-qualification composition is fixed at the averages observed among incumbents. The last two rows show the estimated average mover gap from the baseline regression and from a reweighted version that aligns the estimation sample with the observed distributions 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 a linear probability model with candidate, recruiter-search, absolute-rank, and relative-rank fixed effects (Equation (6)); standard errors are clustered two-way on the recruiter and the candidate. The underlying regression is Appendix Table F.8, Column (4).

Together, the results in this section suggest that job requirements overlap, occupational experience, and occupation-specific qualifications 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 a key assumption of existing models of occupational mobility. The results also shed light on 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 all 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 screening decision in our framework has a second margin: the threshold τ_s that the match 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 chance of contact 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 mainly affects the movers, who are more likely to be below the initial threshold (Section 2).

In the following, 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 the two results in a decomposition that quantifies the role of firms in occupational transitions.

5.1 The Candidate Pool in Tight Markets

We relate the mover gap to an occupation- and region-specific measure of tightness among incumbents: 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 construction of the measure in detail.

Before turning to recruiters' decisions, we document what tightness means for the candidate pool recruiters face. 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. For the counts of incumbents and movers, we then estimate the following Poisson regression (separately by candidate type $g \in \{I, M\}$):

$$\mathbb{E}\left[n_u^g\right] = \exp\left(\beta_{\text{mid}}^g\, T_u^{\text{mid}} + \beta_{\text{tight}}^g\, T_u^{\text{tight}} + \text{FE}_u\right), \tag{7}$$

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 $tightness_{o(u),\mathcal{C}(u),t(u)}$—the tightness of the occupation and region of search or vacancy u at the time the search takes place (Appendix D)—falls in the middle or the top tercile, with the bottom tercile (slack occupations) as the reference. The Poisson coefficients measure proportional shifts in the number of candidates of type g, in log points. On the platform, the fixed effects FE_u comprise the calendar month, the searched occupation interacted with the set of searched cantons, and a set of recruiter fixed effects, so 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.[28] Standard errors are clustered on the recruiter in the search-level platform regression and on the hiring firm in the vacancy-level diary regression.

Figure 7 presents the results of those regressions. It shows, first, that incumbents thin out 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 measure captures a scarcity of incumbents at the level of the individual search. Second, and not expected by construction, the number of movers is almost unchanged in both panels, suggesting that there is no compensating inflow of movers into tight occupations. This is the opposite of what we would expect if movers were aware of tightness and applied where their chances are better. The stable number of mover applications is suggestive of a limited short-run response to increased occupational tightness.[29] Tighter occupations therefore confront recruiters with a pool that contains fewer incumbents and approximately the same number of movers.

5.2 Screening in Tight Markets

How does this affect recruiter behavior on Job-Room? To answer this question, we return to the candidate-level regressions estimated earlier and augment the baseline model

[28] Note that in this specification we do not add fixed effects for the vacancy occupation and canton because the application diary data has a limited time dimension. The variation in tightness that would be left after controlling for occupation and location would be too weak to identify effects.

[29] On 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, so the composition of movers is much stickier there in the first place.

Figure 7. The candidate pool by labor market tightness

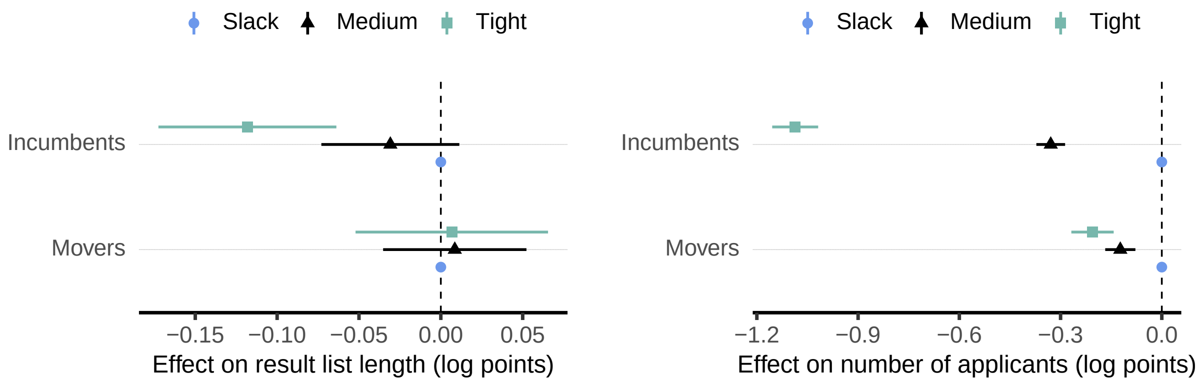

(a) Recruiter result lists (Job-Room platform) (b) Applicant pools (application diaries)

Notes: The figure reports estimates of Equation (7): the shift in the number of candidates of each type across tightness terciles, in log points relative to the slack (bottom) tercile, separately for incumbents (the same last occupation as the vacancy) and movers (a different occupation). 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. The underlying estimates are in Appendix Table F.10. 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 regressions include calendar-month and hiring-firm fixed effects, with standard errors clustered on the hiring firm. The underlying estimates are in Appendix Table F.12. Whiskers are 95% confidence intervals.

(Equation (6)) 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. The level effect of tightness is absorbed by the search fixed effects ϕ_s. The interaction coefficients therefore have a difference-in-differences structure: within each search, the mover gap is the contact-rate difference between movers and incumbents. The interactions, in turn, measure how this within-search gap differs between searches conducted in slack and tight occupations.

Panel (a) of Figure 8 reports the resulting mover gaps in the slackest and the tightest tercile on average and estimated separately for movers whose last occupation overlaps closely with the vacancy occupation and for movers with low overlap. Panel (b) replicates the analysis in the application diary data, using the logit specification of Section 4.3 augmented with the mover-by-tightness interactions. As before, we report average marginal effects.

Panel (a) of Figure 8 shows that tight markets make recruiters far more willing 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. 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, a reduction of roughly 70%.

The diary hiring data point the same way, though the estimates are far less precise (Panel b).[30] The mover hire-rate gap tends to narrow as markets tighten, most clearly for closely related movers, whose tight-market gap is not statistically distinguishable from zero (see also Appendix Table F.11).

Figure 8. The mover gap by labor market tightness and overlap with the vacancy

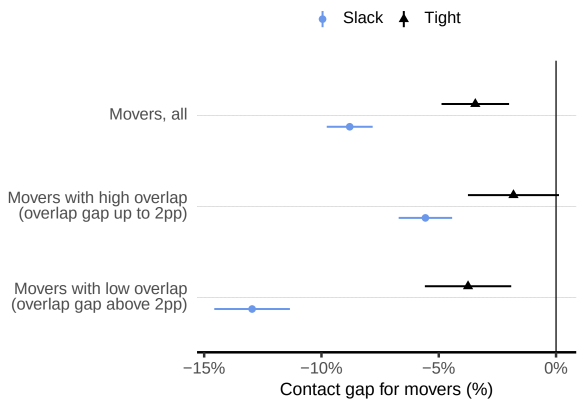

(a) Recruiter contacts (Job-Room)

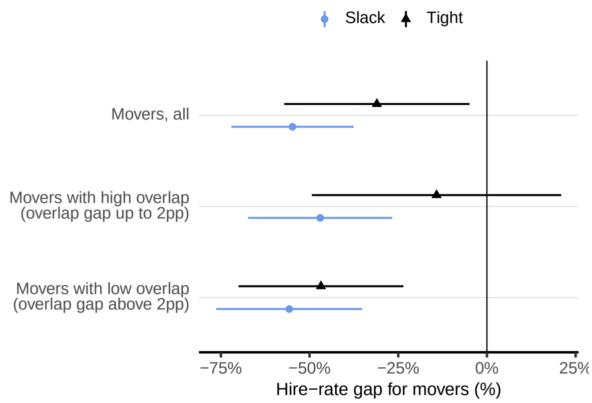

(b) Hires (application diaries)

Notes: The figure shows results from regressions interacting the mover indicator with the tightness in the vacancy's market. Each panel plots the predicted mover gap, as a percent of the baseline rate, in slack (bottom tightness tercile) versus tight (top tightness tercile) occupations, holding all other variables constant. The gaps are shown for all movers and separately for movers with high overlap (the candidate's last occupation within 2 percentage points of the vacancy's within-occupation overlap) and low overlap (more than 2 points below). The overlap dimension is a sample split with one regression per mover group. The middle tightness tercile enters the regressions but is not plotted. In Panel (a), the unit of observation is a candidate profile appearing in a recruiter search on Job-Room. The gap refers to the probability that the recruiter contacts the candidate, measured against same-search incumbents, and is estimated from Equation (6) augmented with interactions between the mover indicator and the tightness terciles, controlling for candidate, recruiter-search, absolute-rank, and relative-rank fixed effects. Standard errors are clustered two-way on the recruiter and the candidate. The underlying regressions are in Appendix Table F.9. In Panel (b), the unit of observation is an application in the application diary data, and the gap refers to the probability that the application leads to the eventual hire, measured against same-occupation applicants, controlling for a wide range of job seeker characteristics and hiring firm fixed effects. Because the hire rate is about 1.2%, we fit logit models with hiring-firm fixed effects instead of a linear probability model and report average marginal effects with delta-method standard errors, clustered two-way on the hiring firm and the job seeker. The underlying regressions are in Appendix Table F.11. Whiskers are 95% confidence intervals.

In sum, the results presented in this section suggest that labor market tightness reduces firms' reluctance to hire movers. These results align with a 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,

[30] As in Section 4.3, we fit the diary gap by logit with a hiring-firm fixed effect and the full job seeker control set, but without a last-occupation fixed effect, which the short diary panel cannot support. The within-overlap regressions provide a test of the fixed effect's importance: there, the regressor varies across the vacancies a single job seeker applies to, so the gradient remains identified when the fixed effect is added—and adding it leaves the results qualitatively unchanged (Appendix Table F.6).

2024; Cullen et al., 2025). Our findings highlight 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.

5.3 Firms' Role in Slack and Tight Markets

The preceding results identify two reasons why tight markets produce more contacts with movers. First, incumbents become scarce, so movers account for a larger share of the candidate pool. Second, firms become more willing to contact a mover relative to an incumbent. We now quantify the contribution of this second, firm-driven channel in two real world settings.

Begin with the recruiter platform. A recruiter draws one candidate at random from a recruiter's result list. The probability of drawing a mover who is contacted depends on both the mover share of the list and movers' contact rate. It is therefore $s_M(T)p_M(T)$, where $s_M(T) = \mathbb{P}(\text{Mover} = 1 \mid T)$ is the mover share and $p_M(T) = \mathbb{P}(\text{contact} = 1 \mid \text{Mover} = 1, T)$ is the contact rate conditional on being a mover. The corresponding probability for incumbents is $[1 - s_M(T)]p_I(T)$.

We define the difference between these two probabilities as the mover gap in contacts:

$$G(T) \;=\; \underbrace{s_M(T)\, p_M(T)}_{\text{a mover is contacted}} \;-\; \underbrace{\big[1 - s_M(T)\big]\, p_I(T)}_{\text{an incumbent is contacted}}, \tag{8}$$

where T denotes tightness in the vacancy occupation.[31] The mover gap in contacts differs from the conditional contact-rate gap, $p_M(T) - p_I(T)$, studied above, because it also incorporates who is available to be contacted. The diary data yield the analogous mover gap in hires: $s_M(T)$ is then the mover share of applications, and $p_M(T)$ and $p_I(T)$ are hire rates conditional on applying.

Figure 9 reports these gaps in the slackest and tightest terciles. To make their magnitudes comparable, each gap is scaled by the pooled probability that a candidate is an incumbent who is contacted (Panel a) or hired (Panel b). Tightness shifts both outcomes toward movers. On the platform, the mover deficit in contacts narrows from -36.1% in slack markets to -20.9% in tight markets. In the diaries, the mover advantage in hires rises from 2.5% to 152.6%.

The opposite signs do not imply that firms favor movers in the diary data. In both data sources, movers advance at a lower conditional rate than incumbents. The signs differ because the pools differ. Movers make up roughly half of platform result lists (46.1% in slack and 49.2% in tight markets), so a randomly drawn profile is less likely

[31] Note that the mover gap in contacts can also be written as $G(T) = s_M(T)[p_M(T) - p_I(T)] - [s_I(T) - s_M(T)]p_I(T)$, where $s_I(T) \equiv 1 - s_M(T)$ is the share of incumbents in the pool. This representation makes clear that both the contact-rate gap, $p_M(T) - p_I(T)$, and the share gap $s_I(T) - s_M(T)$ are driving the gap in actual contacts or hires.

Figure 9. The mover gap in contacts and in hires, and what moves it with tightness

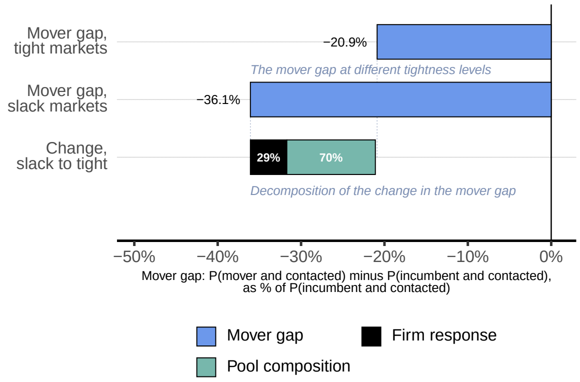

(a) Recruiter contacts

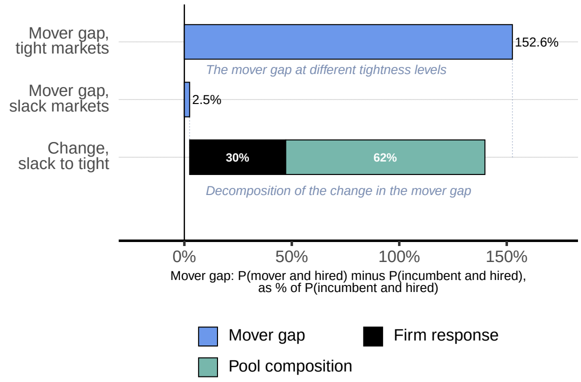

(b) Diary hires

Notes: The mover gap is the probability that a candidate drawn from the pool is a mover and is contacted (Panel a) or hired (Panel b), minus the probability that she is an incumbent and is contacted or hired, as in Equation (8). 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 application diary data. The top two bars give the gap in the slackest and tightest tercile of the tightness of the vacancy occupation. The bottom bar decomposes the change between them by Equation (9) into the firm response and the change in pool composition, each labeled with its share of the change; the bar omits the third term, which only deviates from zero when the mover share in search results or applications is far from half. The underlying coefficients stem from the models behind Figure 7 and Figure 8.

to be a contacted mover than a contacted incumbent. By contrast, movers account for most diary applications (66.3% in slack and 82.6% in tight markets). Their numerical advantage more than offsets their lower hire rate, so a randomly drawn application is more likely to correspond to a hired mover than to a hired incumbent. The level and sign of $G(T)$ thus depend on pool composition. Our object of interest is the change in $G(T)$ as markets tighten.

We decompose the tight-minus-slack change in $G(T)$ in the manner of Oaxaca (1973) and Blinder (1973). [32]

$$\Delta G = \underbrace{\bar{s}_M \, \Delta(p_M - p_I)}_{\text{firm response}} + \underbrace{\left(\bar{p}_M + \bar{p}_I\right) \Delta s_M}_{\text{pool composition}} + \underbrace{\left(2\bar{s}_M - 1\right) \Delta p_I}_{\text{baseline contact rate increase}} , \tag{9}$$

where Δ is the difference between a tight and a slack market.

The decomposition has three parts. First, $\bar{s}_M \Delta(p_M - p_I)$ is the firm's response: firms narrow the conditional rate gap, and the effect is weighted by the average share of movers exposed to that change. Second, $(\bar{p}_M + \bar{p}_I)\Delta s_M$ is the composition effect: incumbents become scarcer and the mover share rises. The two rates enter as a sum

[32] Rewrite the gap as $G = (2s_M - 1)\, p_I + s_M\, \delta$, where $\delta = p_M - p_I$ is the conditional mover–incumbent contact-rate gap of Equation (5). Applying the midpoint product rule, $\Delta(xy) = \bar{x}\Delta y + \bar{y}\Delta x$, to each product yields Equation (9), in which the firm response is $\bar{s}_M \, \Delta\delta$. A bar denotes the midpoint of the slack and tight values, and Δ denotes the tight-minus-slack change.

because shifting one candidate from the incumbent side of the pool to the mover side simultaneously reduces the incumbent contribution and raises the mover contribution. Third, $(2\bar{s}_M - 1)\Delta p_I$ captures a common shift in the baseline contact rate, holding the average pool composition fixed. This term is small when movers and incumbents have similar pool shares (and zero when they have exactly equal shares).

All components come from the preceding analyses. The mover-by-tightness interactions behind Figure 8 identify $\Delta\delta$; the applicant-count models behind Figure 7 identify Δs_M; and the data provide the midpoint shares and rates. The corresponding estimates appear in Column 1 of Appendix Tables F.9 and F.11, and Columns 1 and 2 of Appendix Tables F.10 and F.12. The main figure reports the firm and composition channels; the third term is the distance by which the stacked bar falls short of the tight-market gap.

Interpreting the composition term as a supply-side effect requires that the mover share rises because incumbents become scarce, rather than because movers redirect their applications toward tight occupations. The evidence supports this interpretation: mover applications do not respond to month-on-month changes in tightness, whereas incumbent applications fall sharply (Figure 7 and Appendix Table F.12).

The bottom bar of each panel in Figure 9 applies this decomposition. Most of the shift toward movers is mechanical: incumbents become scarcer in the candidate pool. Firms nevertheless make an economically important contribution. Their screening response accounts for 28.7% of the shift in contacts and 29.8% of the shift in hires. Thus, about one-third of the tightness-driven shift toward movers comes from firms narrowing the conditional rate gap, rather than from pool composition. The close agreement across data sources is notable: it appears both in recruiters' initial contact decisions and in firms' realized hires. The third term accounts for 1.4% of the shift in contacts and 8.5% of the shift in hires.[33]

6 Conclusion

This paper studies how firms use signals when screening candidates from other occupations. Recruiters do not observe a candidate's complete skill set or her match with a particular vacancy. They therefore infer expected fit from observable signals: the requirements associated with her previous occupation, prior experience in the vacancy occupation, and occupational credentials. Linking recruiter clickstream data and application diaries to unemployment-register data allows us to distinguish the information conveyed by these signals from differences in candidates' general employability.

Three findings emerge. First, the average mover faces a contact gap of 7% relative to

[33] The weight on this term, $2\bar{s}_M - 1$, is further from zero the further the mover share of the pool is from one half. On the platform movers are close to half of the result lists, so the term is negligible; in the diary they are around seven in ten applications, so it is not.

an equally qualified incumbent. Job requirements overlap is the dominant explanation: comparing movers and incumbents whose previous occupations imply the same expected overlap with the vacancy reduces the gap by 80%. The value of the occupation signal depends on its precision. The mover gap is 3.3 times larger in homogeneous occupations, where workers from the same occupation are likely to have performed similar jobs and incumbency therefore strongly predicts match quality. In heterogeneous occupations, the occupation label is less informative and movers face a smaller barrier.

Second, recruiters combine several signals rather than relying on the candidate's last occupation alone. Greater requirements overlap, previous experience in the vacancy occupation, and a relevant professional qualification each independently reduce the contact gap. Licensing reinforces the importance of formal signals: the predicted mover gap is 10.3% of the baseline contact rate in licensed occupations, compared with 7.1% in unlicensed occupations.

Third, firms change how they use these signals when incumbent workers become scarce. The mover gap falls from 8.8% of the baseline contact rate in the slackest occupations to 3.4% in the tightest occupations. The response is especially pronounced for movers from distant occupations, whose gap starts largest and falls the most, although movers from nearby occupations gain substantially too. Our decomposition separates this behavioral response from the mechanical change in the candidate pool as incumbents become scarce. Firms' greater willingness to advance movers accounts for about one-third of the tightness-driven shift toward movers; the remainder primarily reflects pool composition. The estimated firm contribution is similar for recruiter contacts and realized hires.

These findings have direct implications for policy. Occupation labels alone are an inadequate guide to feasible transitions. Two occupations with different labels may share most of their requirements, while two jobs within the same occupation may differ considerably. Public employment services could combine vacancy-based overlap measures with information on workers' experience and qualifications to identify target occupations in which a particular mover is likely to pass firms' screening thresholds. The results also suggest directing movers toward heterogeneous occupations, where incumbency is less informative, and toward tight occupations, where firms become more receptive to outside candidates.

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 transferable skills may reduce the mover gap. Job platforms could present verified skills, requirement-level matches, or evidence from prior tasks alongside occupational histories.

Several limitations qualify the conclusions. The platform data measure initial contacts rather than the complete hiring process. The diary data reproduce the main patterns for realized hires, but cover a shorter period and yield less precise estimates. 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. Moreover, although overlap, experience, and credentials explain most of the contact gap, the remaining difference cannot be uniquely attributed to informational frictions. It may reflect unmeasured productivity, inaccurate employer beliefs, other signals correlated with occupational history, or preferences unrelated to expected productivity. The tightness results are also based on observational variation: fixed effects absorb persistent differences across markets, but time-varying demand or composition changes may still confound the estimated response.

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., 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.

Blinder, A. S. (1973). Wage discrimination: Reduced form and structural estimates. Journal of Human Resources, 8(4), 436–455.

Bloesch, 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)

Bohren, J. A., Haggag, K., Imas, A., & Pope, D. G. (2025). Inaccurate statistical discrimination: An identification problem. Review of Economics and Statistics.

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.

Guvenen, 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.

Kircher, P. (2022). Job search in the 21st century. Journal of the European Economic Association, 20(6), 2317–2352.

Klaeui, J. (2025). Guiding jobseekers to where they're wanted most: Evidence from caseworker-driven job search redirection (Draft). CREST and ETH Zurich.

Kleiner, M., & Xu, M. (in press). 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.

Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. Journal of Business & Economic Statistics, 37(2), 187–204.

Phelps, E. S. (1972). The statistical theory of racism and sexism. American Economic Review, 62(4), 659–661.

Poletaev, 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.sbfi.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: Job-Room Data

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

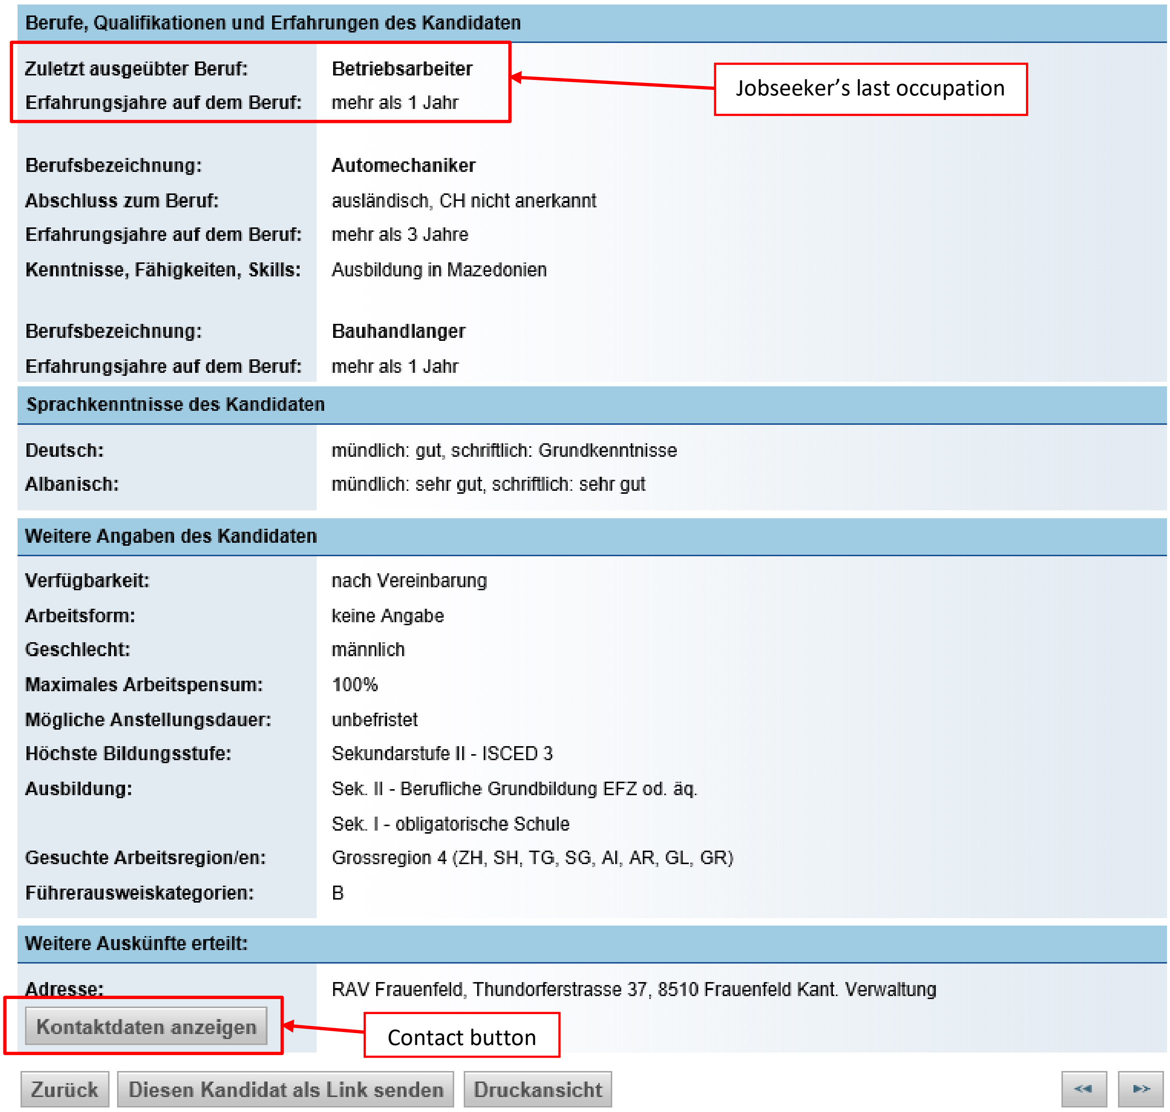

Notes: The figure shows the full candidate profile on Job-Room as seen by a recruiter.

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 vacancy occupation, aggregated at the ISCO-19 one-digit level.

Figure A.2. Recruiter clicks: distribution of the contact rate among candidates from a different occupation

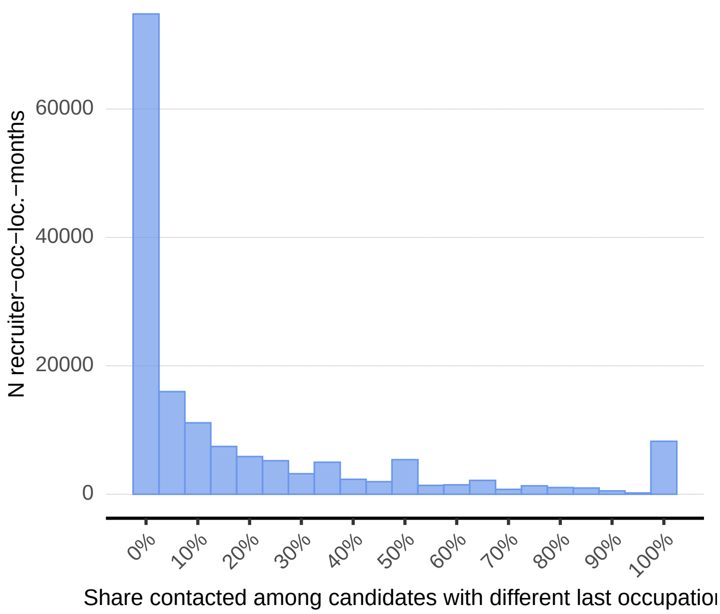

Notes: The figure shows the share of movers contacted by a recruiter, at the level of a recruiter-month-search-criteria combination (vacancy occupation and location), restricted to observations with at least one contacted candidate.

A.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 crucial, 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, and weight each remaining token by its inverse frequency, normalized to sum to one within each name; this 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 (Figure A.3). On a hand-labeled set of 2,000 matched name pairs, the procedure achieves a balanced accuracy of 91.4%.

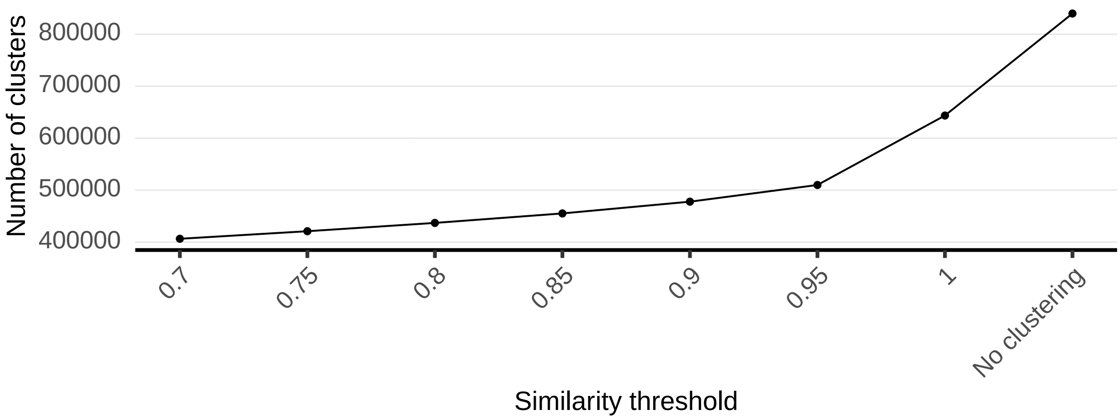

Figure A.3. Application diaries: company name string matching, number of firm clusters as a function of the similarity threshold

Notes: The elbow plot shows the number of clusters as a function of the similarity threshold used in hierarchical clustering of commercial-register entries. The inflection point around 0.95 guides the choice of threshold.

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

Variable	**Description**
Demographics and Availability	
Gender: Male/Female	Gender of the job seeker
Availability	Immediately available or no information
Protected contact information	Contact details hidden from recruiters
Work Preferences	
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
Geographic Mobility	
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)
Language Competencies	
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
Education and Experience	
Education level	Primary, Secondary/vocational, or University
Number of occupations listed	Count of different professions in profile
Skills Field Text Analysis	
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, so job seekers have 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) and import a vacancy's details directly from the posting via a transfer button (Figure B.2). 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, and Klaeui (2025) shows that the occupational and geographic spread of these applications tracks that of the jobs job seekers eventually find.

We focus on job seekers who fill out their protocols through the Job-Room online form and import button, in the second to fourth quarter 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.

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, using the two-stage string-matching and clustering procedure of Appendix A.1, which achieves 91.4% 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 tracks 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, a reassuring comparison given that a higher measured rate may mean that we miss fewer matches.

Figure B.1. Application diaries: online submission form

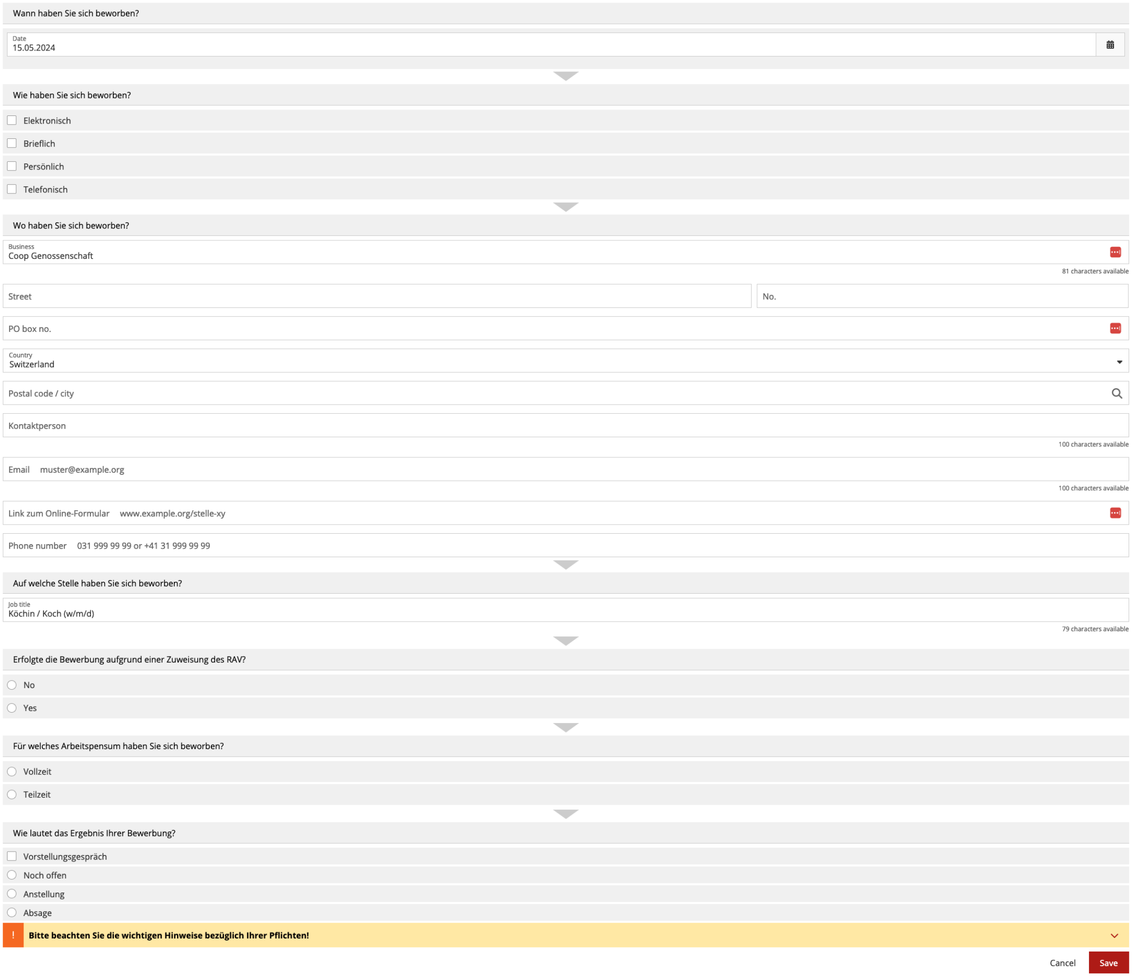
Wann haben Sie sich beworben?

Date
15.05.2024

Wie haben Sie sich beworben?

- Elektronisch
- Brieflich
- Persönlich
- Telefonisch

Wo haben Sie sich beworben?

Business
Coop Genossenschaft
81 characters available

Street | No.

PO box no.

Country
Switzerland

Postal code / city

Kontaktperson
100 characters available

Email muster@example.org
100 characters available

Link zum Online-Formular www.example.org/stelle-xy

Phone number 031 999 99 99 or +41 31 999 99 99

Auf welche Stelle haben Sie sich beworben?

Job title
Köchin / Koch (w/m/d)
79 characters available

Erfolgte die Bewerbung aufgrund einer Zuweisung des RAV?

- No
- Yes

Für welches Arbeitspensum haben Sie sich beworben?

- Vollzeit
- Teilzeit

Wie lautet das Ergebnis Ihrer Bewerbung?

- Vorstellungsgespräch
- Noch offen
- Anstellung
- Absage

! Bitte beachten Sie die wichtigen Hinweise bezüglich Ihrer Pflichten!

Cancel Save

(a) Application entry form

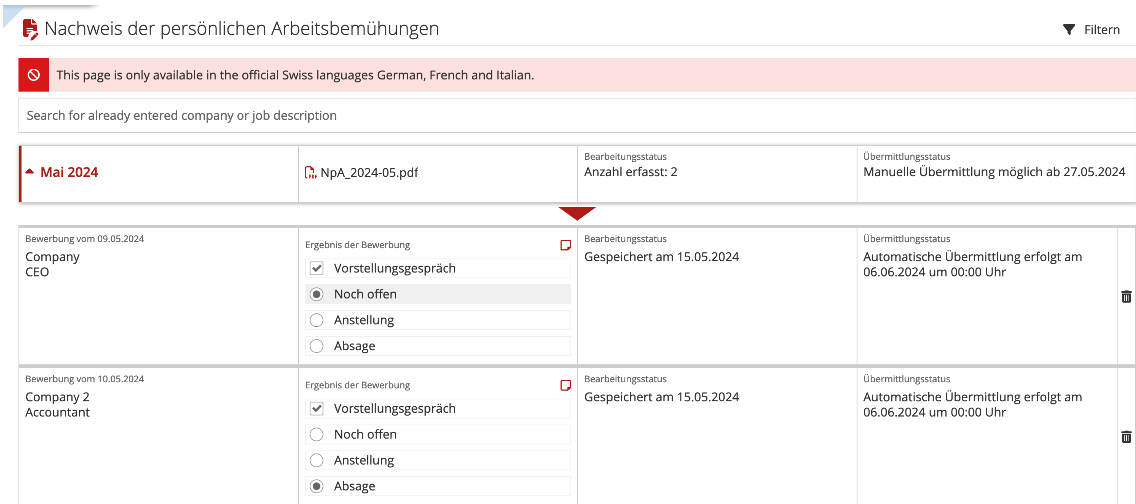
Nachweis der persönlichen Arbeitsbemühungen

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

Notes: The figure shows screenshots of the online forms used to report monthly application efforts to the Swiss PES.

Figure B.2. Application diaries: transfer link on a job ad

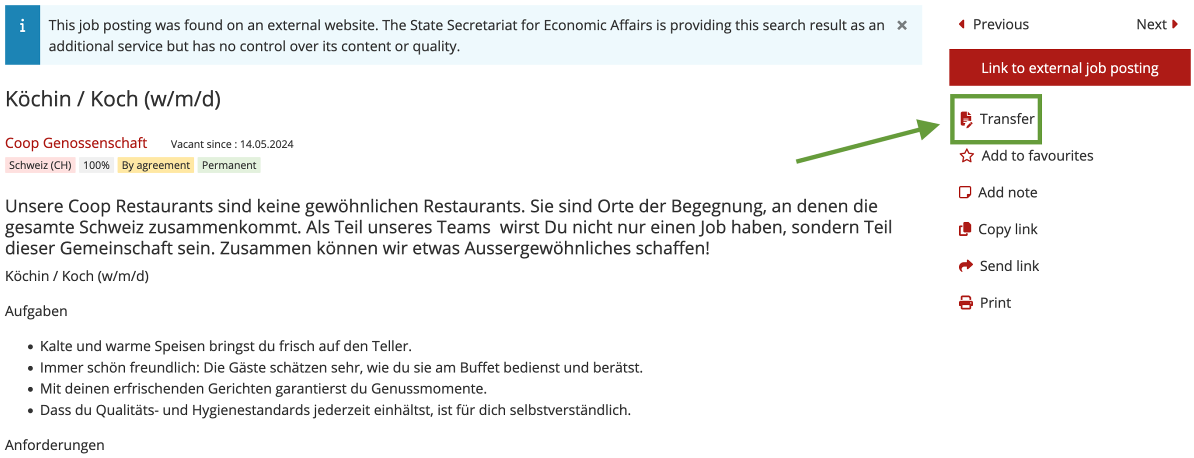

Notes: The screenshot highlights the transfer button on a Job-Room job posting that lets job seekers import the vacancy's details into their monthly application protocol.

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.

C Details on the Index of Overlap in Job Requirements

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, first used by Colella (2022) to analyze the impact of trade on labor demand, 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 key 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 between two occupations 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 postings between 2016 and mid-2022, which covers 80% of all four-digit ISCO occupation pairs.

Figure C.1. Job requirements overlap: an example

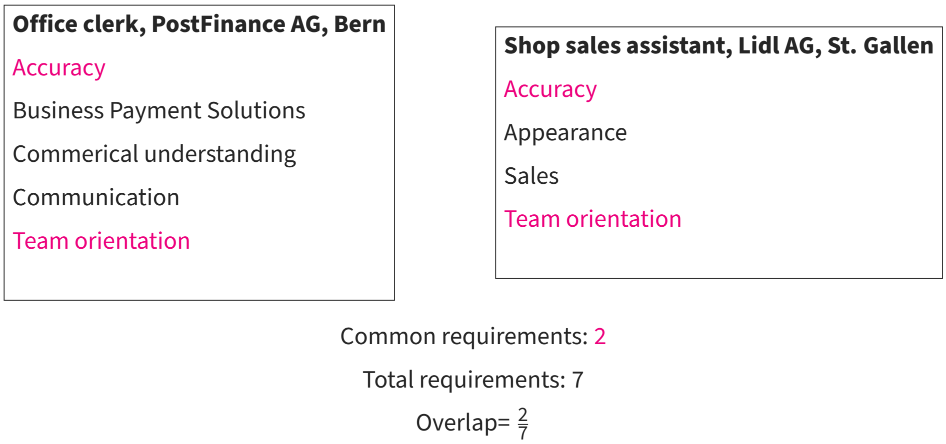

Notes: The figure illustrates job requirements overlap between two example vacancies.

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: first with the average overlap among pairs sharing the same three-digit ISCO occupation, which resolves the majority of cases; then, for the remaining cases, within the two-digit level; and finally 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

	Between (origin, target) pairs		Within-target (vacancy) occupation	
Source	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

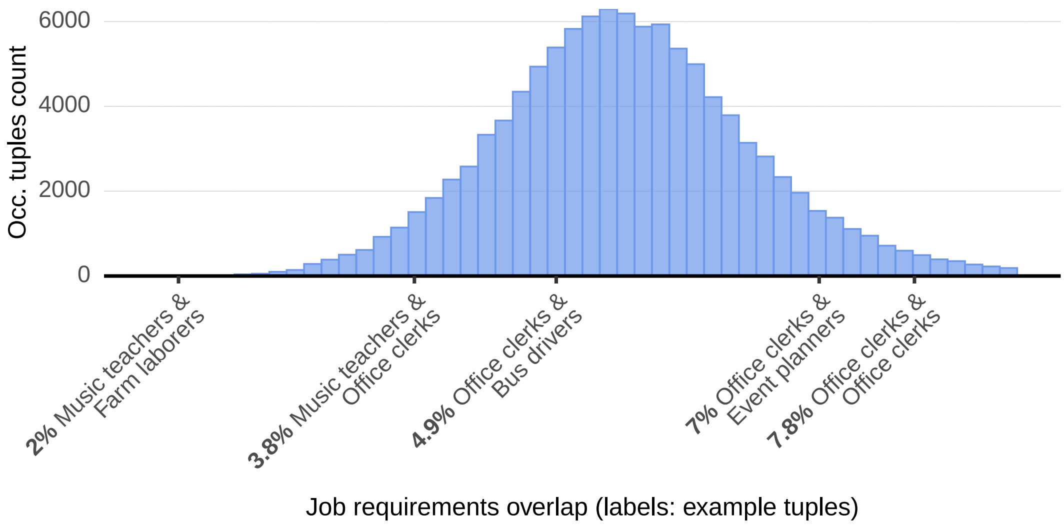

Notes: The histogram shows job requirements overlap across all pairs of four-digit ISCO occupations, including the pairs of an occupation with itself; the axis labels mark example pairs. The graph is truncated at the 99th percentile of the similarity.

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, which identifies a fixed set of related occupations for each occupation based on task, knowledge, and job-title similarity. 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"

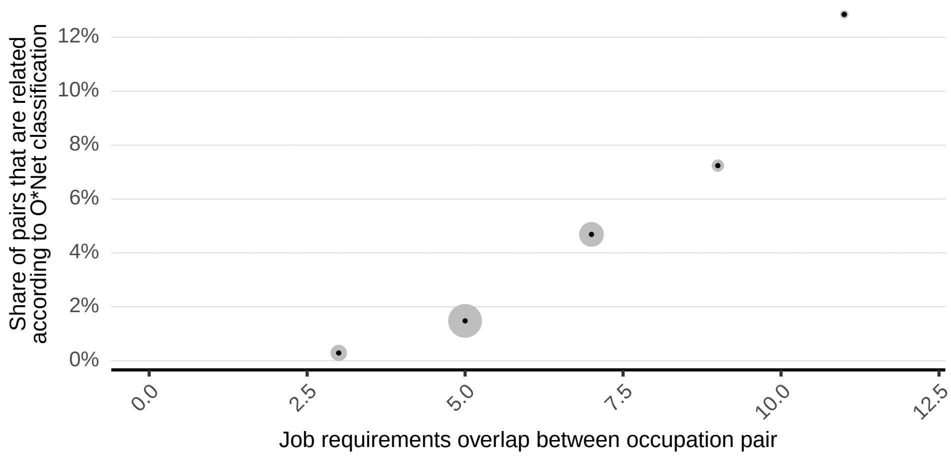

Notes: The figure shows a binned scatter plot of the share of related occupations according to O*NET against our overlap measure. We map O*NET SOC-2019 codes to ISCO-19 via SOC-2010, weight SOC-2019 occupations by their U.S. employment shares (following Belot et al., 2019), and label an ISCO-19 pair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

$$tightness_{o,\mathcal{C},t} = \frac{V_{o,\mathcal{C},t}}{U_{o,\mathcal{C},t}}, \tag{10}$$

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 vacancy 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 shows the distribution of log tightness across recruiters' searched markets, and Figure D.2 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. Recruiter clicks: distribution of the tightness measure across searched markets

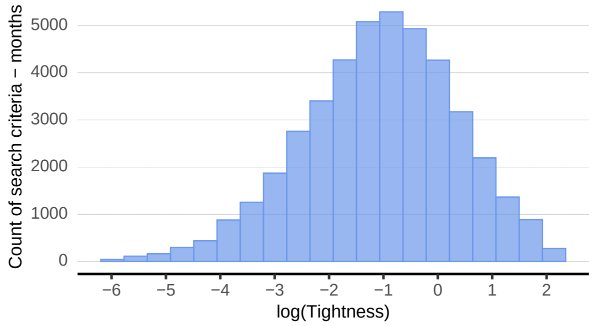

Notes: The histogram shows the log tightness of recruiters' searched markets (occupation × location) across months.

Figure D.2. Temporal variation in labor market tightness across occupation-region cells

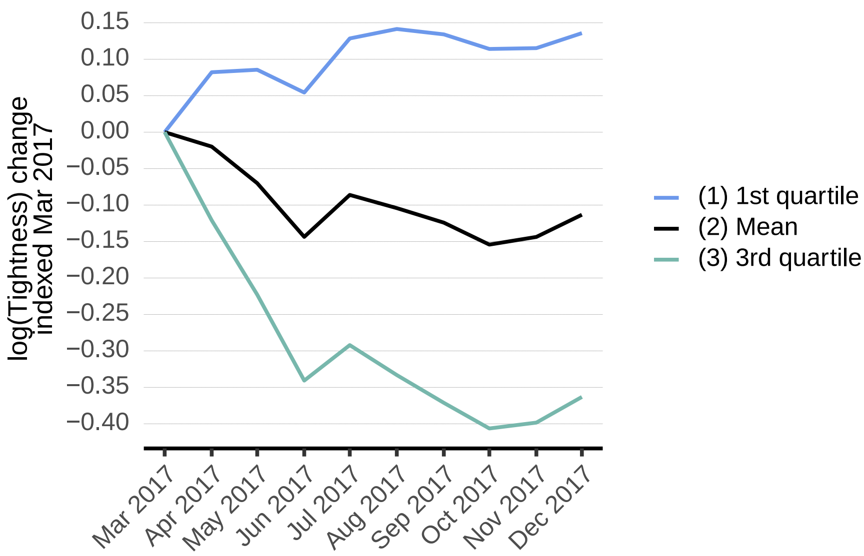

Notes: The figure shows the evolution of tightness at the occupation-region level over time. Each series is indexed by subtracting its log tightness in the initial month (March 2017); the figure plots the 25th percentile, mean, and 75th percentile of these log-differenced values across cells in each month.

E Vacancy-Occupation Characteristics

Section 4.3 relates the mover gap to three constructed properties of the vacancy 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, so we 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 vacancy 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 Contact Gap

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

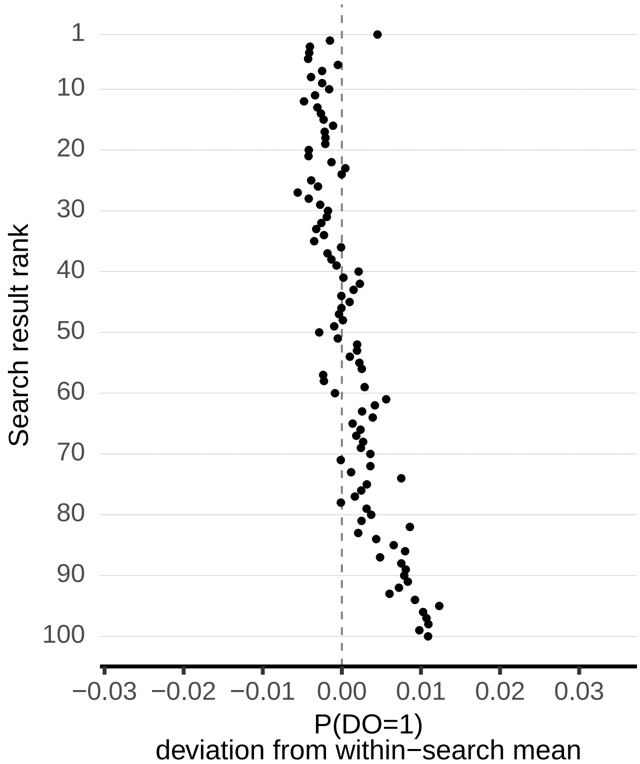

Notes: For each result-list rank, the figure plots the mean of the mover indicator, demeaned within recruiter search: from each candidate's mover indicator we subtract the mean across all candidates in the same search, then average to the rank level. Demeaning removes the fact that searches with more or fewer results can have different baseline mover shares. 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

	Baseline	3-digit ISCO	Not opened or under 20s	Concrete interest
	(1)	(2)	(3)	(4)
Dependent Var.:	Contact click	Contact click	Contact click	Mail or print or fav.
Coefficient	-0.0077***	-0.0076***	-0.0028***	-0.0023***
	(0.0004)	(0.0004)	(0.0003)	(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 (abs. and rel.) 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
N searches	406,873	406,873	387,393	406,873
N recruiters	37,665	37,665	37,043	37,665
Baseline prob.	0.1100	0.1100	0.0618	0.0184
Sample	All	All	Not opened or under 20s	All

Notes: The table reports robustness checks for the baseline in Table 2. Column (1) replicates the baseline (four-digit ISCO). (2) defines a mover at the three-digit ISCO level. (3) restricts to candidates the recruiter either never opened or viewed for at most 20 seconds. (4) uses an outcome combining concrete interest (printing, email click, or favorite) instead of the contact-button click. 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.2. Recruiter clicks: robustness of the recruiter-search fixed effect to leave-out neighborhood controls

	Baseline (X=5 sample)	Leave-out X=1	Leave-out X=3	Leave-out X=5
	(1)	(2)	(3)	(4)
Different occupation	-0.0062***	-0.0062***	-0.0062***	-0.0062***
	(0.0004)	(0.0004)	(0.0004)	(0.0004)
Avg. DO=1 in neighbourhood		-0.0004*	-0.0001	-0.0006
		(0.0002)	(0.0006)	(0.0009)
Avg candidate characteristics in neigbourhood		X	X	X
Candidate FE	X	X	X	X
Recruiter search FE	X	X	X	X
Search rank (abs. and rel.) 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
N searches	262,307	262,307	262,307	262,307
N recruiters	32,997	32,997	32,997	32,997
Baseline prob.	0.0869	0.0869	0.0869	0.0869
Neigbourhood window (each direction)	None	1	3	5

Notes: The table reports the baseline contact-gap regression of Table 2 with leave-out neighborhood controls of job seeker characteristics added on top of the recruiter-search fixed effect. For each candidate at rank r we compute the leave-one-out mean of every job seeker control, and of the mover indicator, across the $2X$ candidates at ranks $r - X, \ldots, r - 1$ and $r + 1, \ldots, r + X$. All columns use the common sample of candidates whose $X = 5$ neighborhood is fully defined (rank ≥ 6 and rank $\leq n_{\text{results}} - 5$). Column (1) is the baseline specification on that sample; (2)–(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 vacancy occupation

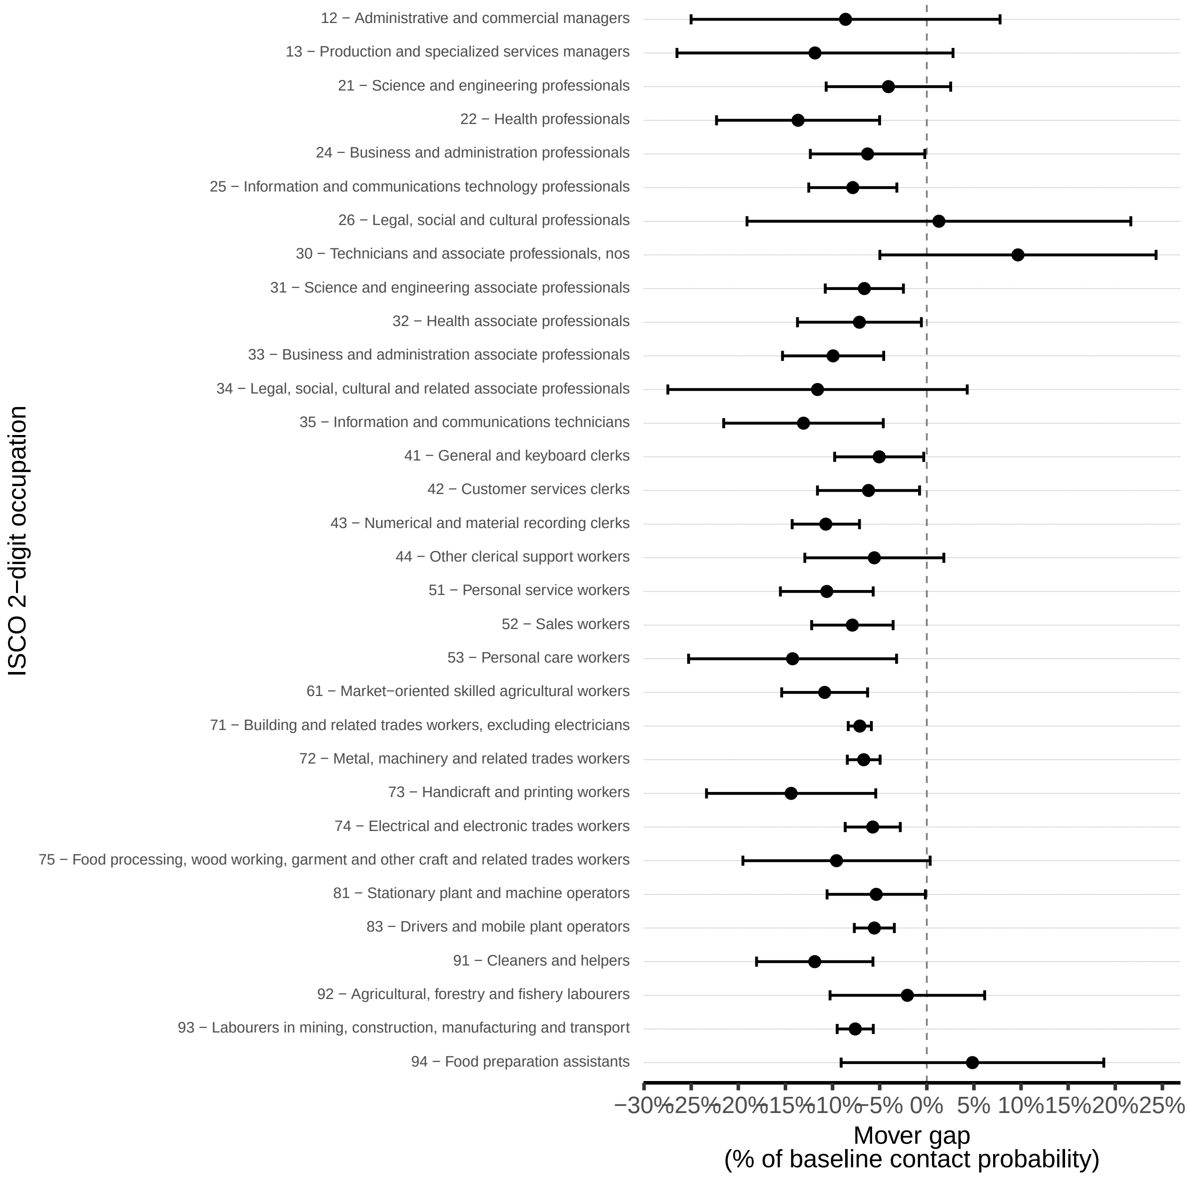

Notes: The figure shows the contact gap interacted with the vacancy occupation, aggregated to the ISCO two-digit level, relative to the occupation-specific baseline contact rate. Each observation is a candidate profile appearing in a recruiter's result list, and the outcome is whether the recruiter clicks the button revealing the candidate's contact details. The rest of the specification follows Table 2, Column (1). Only two-digit occupations with at least 1,000 searches are shown. 95% confidence intervals from standard errors clustered two-way on the recruiter and the candidate are shown.

F.2 Vacancy-Side Heterogeneity

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

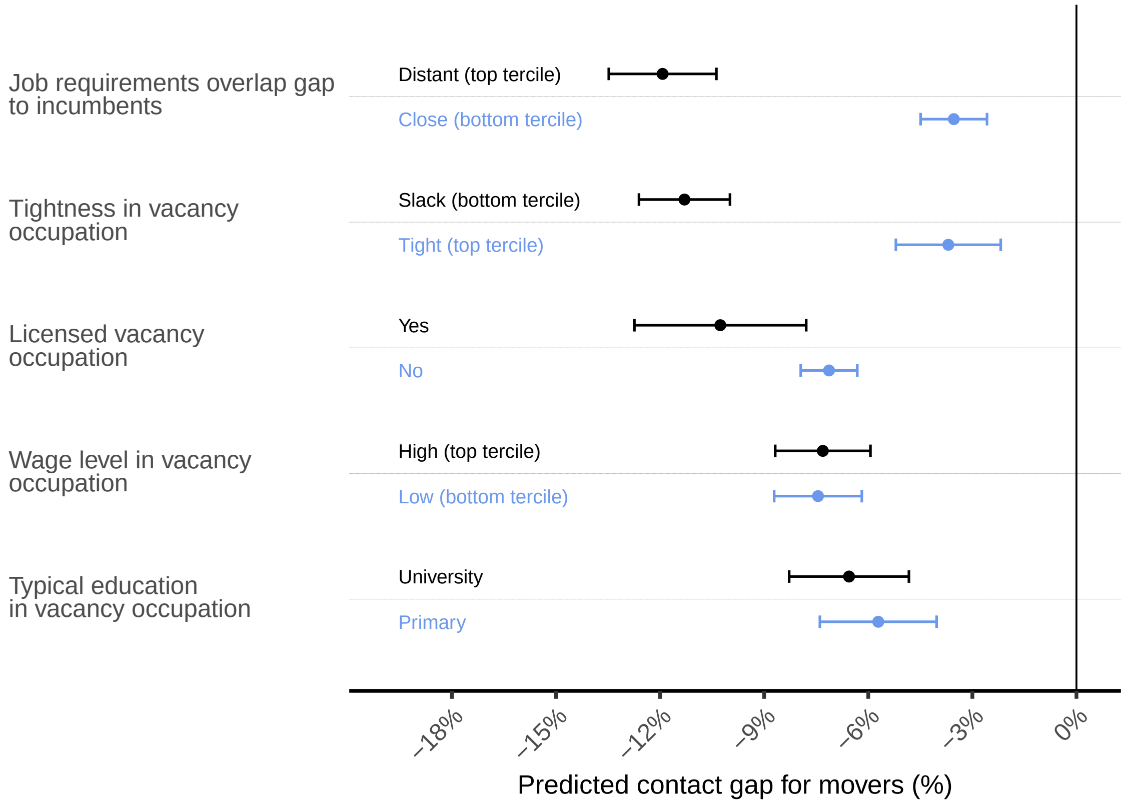

Notes: The figure shows predicted contact gaps for movers, as a percent of the baseline contact rate, along the job requirements overlap gap to incumbents and four characteristics of the vacancy occupation: local, occupational labor market tightness, licensing, the wage level, and the typical education level of incumbents. Each observation is a candidate profile appearing in a recruiter's result list, and the outcome is whether the recruiter clicks the button revealing the candidate's contact details. Each row reports two predictions, at the dimension's low and high values, with the other dimensions held at their sample means: the bottom versus the top tercile of the overlap gap among movers, each evaluated at its tercile mean; slack (bottom tercile) versus tight (top tercile) of log tightness; not licensed versus licensed; bottom versus top tercile of the log wage level; and Primary versus University typical incumbent education (secondary or vocational is the omitted middle category). All dimensions are estimated jointly in one regression of the contact indicator on the mover indicator and its interactions with each dimension, the overlap gap entering additively rather than interacted, since it is zero for incumbents by construction and its coefficient is therefore the mover-specific slope, with candidate, recruiter-search, absolute-rank, and relative-rank fixed effects; the decomposition of these estimates is reported in Appendix Table F.3. Whiskers are 95% confidence intervals from standard errors clustered two-way on the recruiter and the candidate.

Table F.3. Recruiter clicks: Gelbach decomposition of heterogeneity across vacancy occupation characteristics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Change in the mover gap			Contribution from				
Dimension	Univariate reg.	Joint reg.	Difference	Overlap gap	Tightness	Licensing	Wage	Education
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)	–

Notes: The table decomposes how the mover contact gap varies across five dimensions: the job requirements overlap gap to incumbents and four characteristics of the vacancy occupation—local, occupational labor market tightness, licensing, the wage level, and the typical education of incumbents. Each observation is a candidate profile appearing in a recruiter's result list, and the outcome is whether the recruiter clicks the button revealing the candidate's contact details. Each row takes one dimension as the focus. The first two columns report the estimated mover gap in a regression where only the dimension itself is included and in a regression with all five dimensions included. Because the contact gap is negative, a positive value means the gap becomes smaller. The third column reports how much the coefficient changes between the two versions. The remaining columns attribute that change to each of the other dimensions, following Gelbach (2016): each contribution is the product of the coefficient on the contributing dimension in the regression that includes all five, and the coefficient from an auxiliary regression of that contributing dimension on the focal one. The overlap gap enters linearly and is scaled per standard deviation; the other four enter as indicators, and if there are more than two categories we report the difference between the two most extreme ones. The contact gap is expressed as a percentage of the baseline contact rate, so the differences are in percentage points. All regressions include candidate, recruiter-search, absolute-rank, and relative-rank fixed effects and run on the same sample. In some rows the contributions do not add up exactly to the change, because the reported figures are rounded. Standard errors, in parentheses, are clustered two-way on the recruiter and the candidate.

Table F.4. Recruiter clicks: Gelbach decomposition of heterogeneity across vacancy occupation characteristics, with within-occupation overlap instead of the overlap gap

	Change in the mover gap			Contribution from				
Dimension	Univariate reg.	Joint reg.	Difference	Homogeneity	Tightness	Licensing	Wage	Education
Homogeneity (homogenous minus heterogenous)	-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 repeats the decomposition of Table F.3 with the within-occupation homogeneity of the vacancy occupation in place of the overlap gap to incumbents, so that all five dimensions are properties of the vacancy occupation and none depends on the job seeker. The other four dimensions are unchanged: occupational labor market tightness, licensing, the wage level, and the typical education of incumbents. Each observation is a candidate profile appearing in a recruiter's result list, and the outcome is whether the recruiter clicks the button revealing the candidate's contact details. Each row takes one dimension as the focus. The first two columns report the coefficient when estimated in a regression where only the dimension itself is included and in a regression with all five dimensions included. Because the contact gap is negative, a positive value means the gap becomes smaller. The third column reports how much the coefficient changes between the two versions. The remaining columns attribute that change to each of the other dimensions, following Gelbach (2016): each contribution is the product of the coefficient on the contributing dimension in the regression that includes all five, and the coefficient from an auxiliary regression of that contributing dimension on the focal one. All five dimensions enter as indicators, and if there are more than two categories we report the difference between the two most extreme ones. The contact gap is expressed as a percentage of the baseline contact rate, so the differences are in percentage points. All regressions include candidate, recruiter-search, absolute-rank, and relative-rank fixed effects and run on the same sample. In some rows the contributions do not add up exactly to the change, because the reported figures are rounded. Standard errors, in parentheses, are clustered two-way on the recruiter and the candidate.

F.3 Job Requirements Overlap and the Contact Gap

Table F.5. Application diaries: the mover hiring gap by within-occupation overlap, across specifications

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Var.:	Find job at firm	Find job at firm	Find job at firm	Find job at firm	Find job at firm	Find job at firm
Constant	-4.1687***					
	(0.1794)					
Within overlap below median	-0.2105	-0.2613**	-0.2217*	-0.2862***	-0.3581***	-0.3984*
	(0.1313)	(0.1048)	(0.1333)	(0.1000)	(0.1247)	(0.2194)
Within overlap at or above median	-0.3325***	-0.5184***	-0.3379***	-0.4379***	-0.5495***	-0.6792***
	(0.0701)	(0.0831)	(0.0767)	(0.0835)	(0.1032)	(0.2142)
Vacancy occupation x canton		X		X	X	
Jobseeker last occupation, 2-digit			X	X	X	X
Firm FE					X	
Vacancy posting						X
Job seeker characteristics	X	X	X	X	X	X
Observations	166,029	162,831	166,029	162,831	157,180	100,830
Pseudo R2	0.01089	0.14016	0.02076	0.14298	0.38666	0.66337
Base hire rate	0.0124	0.0124	0.0124	0.0124	0.0124	0.0124
p-value, equality of the two categories	0.3168	0.0499	0.2921	0.2101	0.1873	0.3748
Overlap slope, linear specification	-0.0127 (0.0121)	-0.0651** (0.0270)	-0.0137 (0.0116)	-0.0490* (0.0260)	-0.0582* (0.0329)	-0.0546 (0.0496)

Notes: The table estimates the within-occupation overlap diary regression behind Figure 5, Panel (a), across specifications. Each observation is an application by a registered job seeker to a hiring firm, and the outcome is whether the application resulted in the job seeker being hired by that firm. The mover indicator is interacted with whether the within-occupation overlap of the vacancy occupation is below or at or above its median, with same-occupation applicants the omitted reference. Every column controls for 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. Column (2) additionally includes a vacancy-occupation-by-canton fixed effect; (3) a last-occupation fixed effect at the two-digit level; (4) both; (5) adds a hiring-firm fixed effect, the figure specification; (6) replaces the occupation-by-canton and hiring-firm fixed effects with a vacancy fixed effect. The table also reports the p-value of a Wald test that the two category coefficients are equal, and the slope from an otherwise identical specification that enters the within-occupation overlap continuously rather than split at the median, with its standard error in parentheses. All columns come from a logistic regression and report estimated coefficients; the figure reports average marginal effects. 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.

Table F.6. Application diaries: the mover hiring gap by the overlap gap to incumbents, across specifications

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Var.:	Find job at firm	Find job at firm	Find job at firm	Find job at firm	Find job at firm	Find job at firm
Constant	-4.1532***					
	(0.1758)					
Different occupation	-0.0568	-0.1352	-0.1249	-0.1540	-0.2404*	-0.2709
	(0.1631)	(0.0974)	(0.1450)	(0.1019)	(0.1313)	(0.2358)
Overlap gap to incumbents 0.5pp or more	-0.2745**	-0.3531***	-0.2056*	-0.2942***	-0.3144**	-0.3826*
	(0.1347)	(0.0890)	(0.1082)	(0.0926)	(0.1257)	(0.2285)
Vacancy occupation x canton		X		X	X	
Jobseeker last occupation, 2-digit			X	X	X	X
Firm FE					X	
Vacancy posting						X
Job seeker characteristics	X	X	X	X	X	X
Observations	166,029	162,831	166,029	162,831	157,180	100,830
Pseudo R2	0.01141	0.14078	0.02095	0.14338	0.38700	0.66364
Base hire rate	0.0124	0.0124	0.0124	0.0124	0.0124	0.0124
p-value, equality of the two categories	0.0415	0.0001	0.0574	0.0015	0.0124	0.0941
Overlap slope, linear specification	-0.0232 (0.0165)	-0.0671*** (0.0225)	-0.0180 (0.0125)	-0.0555*** (0.0215)	-0.0756*** (0.0271)	-0.1056** (0.0436)

Notes: The table estimates the one-axis overlap-gap diary regression behind Figure 5, Panel (b), across specifications. Each observation is an application by a registered job seeker to a hiring firm, and the outcome is whether the application resulted in the job seeker being hired by that firm. Movers are split at a half-point overlap gap to incumbents (close versus distant), with same-occupation applicants the omitted reference, the bins entering additively alongside the mover indicator so that each category's gap combines the two coefficients. Every column controls for 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. Column (2) additionally includes a vacancy-occupation-by-canton fixed effect; (3) a last-occupation fixed effect at the two-digit level; (4) both; (5) adds a hiring-firm fixed effect, the figure specification; (6) replaces the occupation-by-canton and hiring-firm fixed effects with a vacancy fixed effect. The table also reports the p-value of a Wald test that the two category coefficients are equal, and the slope from an otherwise identical specification that enters the overlap gap continuously, with its standard error in parentheses. All columns come from a logistic regression and report estimated coefficients; the figure reports average marginal effects. 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.

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

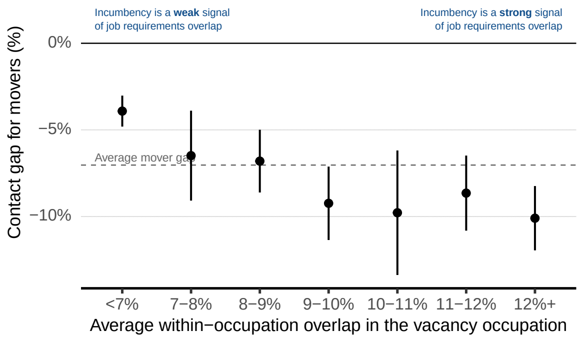

(a) By within-occupation homogeneity of the vacancy occupation

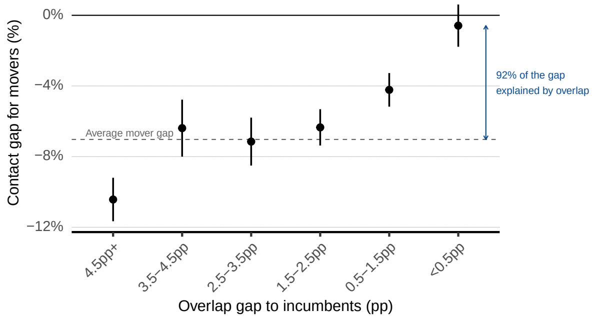

(b) By the overlap gap to the vacancy occupation

Notes: The figure replicates Figure 4 using an inverse-document-frequency-weighted measure of job requirements overlap, which down-weights requirements that appear on many ads: 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, so that 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. Each observation is a candidate profile appearing in a recruiter's result list, and the outcome is whether the recruiter clicks the button revealing the candidate's contact details. Panel (a) interacts the mover indicator with indicators for bins of the within-occupation overlap of the vacancy occupation, in percentage points. Panel (b) adds bins of the overlap gap of the mover to incumbents to the mover indicator, so that each bin's gap combines the mover and the bin coefficient, the gap defined as the within-occupation overlap of the vacancy occupation minus the between-occupation overlap of the candidate's last occupation with the vacancy, in percentage points. The dashed line marks the average mover gap from the baseline regression of Section 4.1. Both regressions include candidate, recruiter-search, absolute-rank, and relative-rank fixed effects. 95% confidence intervals from standard errors clustered two-way on the recruiter and the candidate are shown. Appendix Table F.7, Columns (2) and (4), report the underlying coefficients.

Table F.7. 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***	-0.0043***		
	(0.0009)	(0.0005)		
Within overlap 7-8%	-0.0054***	-0.0071***		
	(0.0007)	(0.0015)		
Within overlap 8-9%	-0.0049***	-0.0075***		
	(0.0010)	(0.0010)		
Within overlap 9-10%	-0.0062***	-0.0102***		
	(0.0010)	(0.0012)		
Within overlap 10-11%	-0.0080***	-0.0108***		
	(0.0012)	(0.0020)		
Within overlap 11-12%	-0.0114***	-0.0095***		
	(0.0011)	(0.0012)		
Within overlap 12% or more	-0.0108***	-0.0111***		
	(0.0009)	(0.0010)		
Mover, overlap closest to within (reference)			-0.0016**	-0.0006
			(0.0007)	(0.0007)
Overlap gap 0.5-1.5pp below within			-0.0032***	-0.0040***
			(0.0007)	(0.0007)
Overlap gap 1.5-2.5pp below within			-0.0049***	-0.0063***
			(0.0008)	(0.0008)
Overlap gap 2.5-3.5pp below within			-0.0059***	-0.0072***
			(0.0010)	(0.0010)
Overlap gap 3.5-4.5pp below within			-0.0065***	-0.0064***
			(0.0011)	(0.0012)
Overlap gap >4.5pp below within			-0.0108***	-0.0108***
			(0.0010)	(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 estimates the two overlap specifications of Section 4.3, each with the unweighted and with an inverse-document-frequency-weighted measure of job requirements overlap. 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. Columns (1) and (2) interact the mover indicator with seven bins of the within-occupation overlap of the vacancy occupation. Columns (3) and (4) add six bins of the overlap gap to the mover indicator, the gap being the within-occupation overlap of the vacancy occupation minus the overlap between the candidate's last occupation and the vacancy occupation, 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) weight each requirement s by $idf_s = \log(N/n_s)$, where n_s is the number of ads listing requirement s and N the number of ads, so that 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. 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.

Figure F.5. Job requirements overlap: inverse-document-frequency-weighted versus unweighted

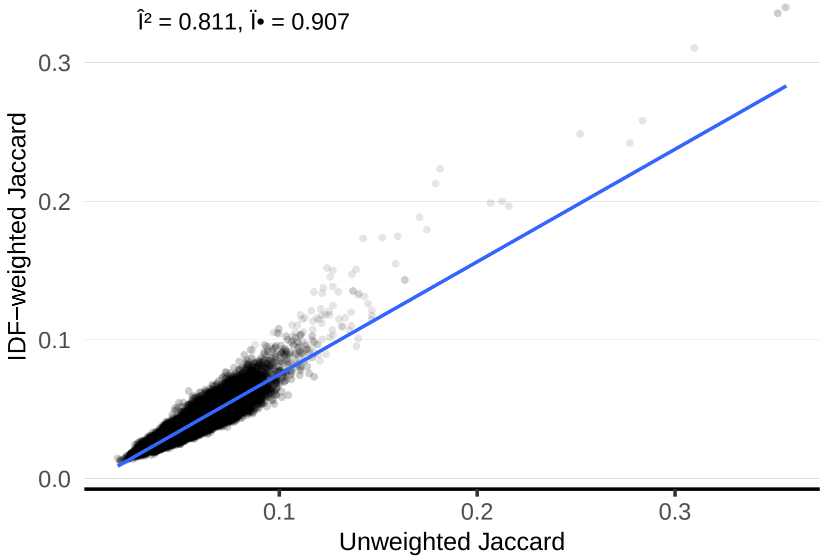

Notes: The figure plots the inverse-document-frequency-weighted against the unweighted overlap for every pair of four-digit ISCO occupations in the estimation sample. The line is the OLS fit; the annotation reports the slope and the Pearson correlation.

F.4 Mover Profiles and the Signal Hierarchy

Table F.8. 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***	-0.0016**	-0.0068***	-0.0010	0.0007
	(0.0004)	(0.0007)	(0.0004)	(0.0009)	(0.0021)
Last-vacancy overlap 0.5-1.5pp below within		-0.0032***		-0.0045***	-0.0086***
		(0.0007)		(0.0010)	(0.0021)
Last-vacancy overlap 1.5-2.5pp below within		-0.0049***		-0.0053***	-0.0069***
		(0.0008)		(0.0012)	(0.0025)
Last-vacancy overlap 2.5-3.5pp below within		-0.0059***		-0.0057***	-0.0122***
		(0.0010)		(0.0014)	(0.0029)
Last-vacancy overlap 3.5-4.5pp below within		-0.0065***		-0.0073***	-0.0103***
		(0.0011)		(0.0016)	(0.0031)
Last-vacancy overlap >4.5pp below within		-0.0108***		-0.0092***	-0.0062**
		(0.0010)		(0.0014)	(0.0031)
No experience in vacancy occ.			-0.0059***	-0.0036**	-0.0091***
			(0.0007)	(0.0016)	(0.0032)
No prof. qualification in vacancy occ.			-0.0041***		
			(0.0005)		
Experience, no prof. qualification in vacancy occ.				-0.0034***	-0.0062***
				(0.0007)	(0.0012)
Last worked in different occupation x Experience, no prof. qualification in vacancy occ.				-0.0011	-0.0018
				(0.0012)	(0.0025)
Experience, no prof. qualification in vacancy occ. x Last-vacancy overlap 0.5-1.5pp below within				0.0028**	0.0047*
				(0.0013)	(0.0027)
No experience in vacancy occ. x Last-vacancy overlap 0.5-1.5pp below within				-0.0003	0.0057
				(0.0018)	(0.0036)
Experience, no prof. qualification in vacancy occ. x Last-vacancy overlap 1.5-2.5pp below within				0.0016	0.0016
				(0.0015)	(0.0031)
No experience in vacancy occ. x Last-vacancy overlap 1.5-2.5pp below within				-0.0013	0.0039
				(0.0020)	(0.0039)
Experience, no prof. qualification in vacancy occ. x Last-vacancy overlap 2.5-3.5pp below within				0.0014	0.0087**
				(0.0017)	(0.0035)
No experience in vacancy occ. x Last-vacancy overlap 2.5-3.5pp below within				-0.0041*	0.0012
				(0.0022)	(0.0044)
Experience, no prof. qualification in vacancy occ. x Last-vacancy overlap 3.5-4.5pp below within				0.0030	0.0057
				(0.0019)	(0.0038)
No experience in vacancy occ. x Last-vacancy overlap 3.5-4.5pp below within				-0.0032	-0.0020
				(0.0023)	(0.0043)
Experience, no prof. qualification in vacancy occ. x Last-vacancy overlap >4.5pp below within				0.0005	-0.0005
				(0.0015)	(0.0031)
No experience in vacancy occ. x Last-vacancy overlap >4.5pp below within				-0.0092***	-0.0089**
				(0.0020)	(0.0040)
Last worked in diff occ. x 35-45y old					0.0003
					(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	14,927,811	15,263,934	4,623,777
R2	0.46442	0.46443	0.46513	0.46444	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 mechanisms regressions behind Figure 6. 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. Column (1) is the baseline mover indicator; (2) adds the six overlap-gap bins; (3) adds indicators for having no work experience and no professional certificate in the vacancy occupation; (4) replaces those indicators with a three-level tier—experience and certificate, experience without certificate, and no experience—and interacts it with the mover indicator and with the overlap bins; (5) repeats (4) on candidates with no entry in the free-text skills field, adding interactions of the mover indicator with age 35 to 49 and with age 50 and above. Figure 6 is built from Column (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.5 Tightness and Recruiter Behavior

Table F.9. Recruiter clicks: the mover contact gap by tightness and overlap with the vacancy

	(1)	(2)	(3)
Dependent Var.:	Contact click	Contact click	Contact click
Mover (different occupation)	-0.0097***	-0.0061***	-0.0143***
	(0.0006)	(0.0006)	(0.0009)
Mover x middle tightness tercile	0.0019**	0.0038***	0.0041***
	(0.0008)	(0.0009)	(0.0011)
Mover x tight (top tightness tercile)	0.0059***	0.0041***	0.0101***
	(0.0010)	(0.0013)	(0.0014)
Candidate FE	X	X	X
Recruiter search FE	X	X	X
Search rank (abs. and rel.) FE	X	X	X
Relative search rank FE	X	X	X
-------------------------------------	-------------	-------------	-------------
Observations	15,085,933	10,056,312	12,938,748
R2	0.46448	0.48676	0.46859
Mover sample	All movers	Close movers	Distant movers
Base contact rate	0.1102	0.1102	0.1102

Notes: The table reports the linear probability models behind Figure 8, Panel (a). 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 mover indicator is interacted with the tightness terciles of the vacancy occupation, with the slack (bottom) tercile the omitted reference, and each column keeps the incumbent candidates as the within-search reference. Column (1) keeps all movers; (2) keeps movers from a closely related occupation (overlap gap to the vacancy's within-occupation overlap up to 2 percentage points); (3) keeps movers from a distant occupation (gap above 2 points). 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.

Table F.10. Recruiter clicks: the result lists by tightness, search-level regressions

	(1)	(2)
Dependent var.	N results DO=1	N results DO=0
Middle tightness tercile	0.0085	-0.0308
	(0.0223)	(0.0215)
Top tightness tercile	0.0067	-0.1182***
	(0.0300)	(0.0277)
Calendar month FE	X	X
Recruiter vacancy x canton FE	X	X
Recruiter FE	X	X
----------------------------	--------------	--------------
Number of searches	401672	401672
Observations	401,672	401,672
Pseudo R2	0.44834	0.43913

Notes: The table reports the search-level regressions behind Figure 7, Panel (a). Each observation is a recruiter search on Job-Room, and the outcomes are the number of different-occupation (Column (1)) and same-occupation ((2)) candidate profiles the search returns. Both columns come from a Poisson regression with calendar-month, vacancy-occupation-by-location, and recruiter fixed effects, on the tightness terciles, with the slack (bottom) tercile the omitted reference. Standard errors, in parentheses, are clustered on the recruiter. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

F.6 Diary Tightness Regressions

Table F.11. Application diaries: the mover hiring gap by tightness and overlap with the vacancy

	(1)	(2)	(3)
Dependent Var.:	Find job at firm	Find job at firm	Find job at firm
Mover (different occupation)	-0.6152***	-0.5406***	-0.5993***
	(0.1300)	(0.1503)	(0.1468)
Mover x middle tightness tercile	0.1801	0.0506	0.2217
	(0.1584)	(0.1757)	(0.2043)
Mover x tight (top tightness tercile)	0.3341*	0.4025	0.1560
	(0.1842)	(0.2569)	(0.1767)
Middle tightness tercile	-0.0343	-0.0388	-0.0700
	(0.1244)	(0.1254)	(0.1271)
Tight market (top tightness tercile)	0.0994	0.0159	0.0810
	(0.1445)	(0.1414)	(0.1678)
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
Mover sample	All movers	Close movers	Distant movers
Base hire rate	0.0125	0.0130	0.0132

Notes: The table reports the mover-by-tightness regression behind Figure 8, Panel (b), in the application diary data, with the same specification: a hiring-firm fixed effect and the job seeker controls detailed below. Each observation is an application by a registered job seeker to a hiring firm, and the outcome is whether the application resulted in the job seeker being hired by that firm. The mover indicator is interacted with the tightness terciles, with the slack (bottom) tercile the omitted reference, and the tightness main effects enter directly. Each column keeps the same-occupation applicants as the within-firm reference. Column (1) keeps all movers; (2) keeps movers from a closely related occupation (overlap gap to incumbents up to 2 percentage points); (3) keeps movers from a distant occupation (gap above 2 points). Every column controls for 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. All columns come from a logistic regression and report estimated coefficients; the figure reports average marginal effects. 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.

Table F.12. Application diaries: the applicant pool by tightness

	(1)	(2)
Dependent var.	N applicants DO=1	N applicants DO=0
Middle tightness tercile	-0.1230***	-0.3292***
	(0.0228)	(0.0217)
Top tightness tercile	-0.2054***	-1.0864***
	(0.0320)	(0.0348)
Calendar month FE	X	X
Firm FE	X	X
------------------------	-----------------	-----------------
Observations	135,926	135,926
Pseudo R2	0.13947	0.16614
Mean of dependent var.	1.59677	0.67592

Notes: The table reports the diary regressions behind Figure 7, Panel (b). Each observation is a vacancy in the application diary data, and the outcomes are the number of mover (different-occupation) applicants (Column (1)) and same-occupation applicants ((2)) per vacancy. Both columns come from a Poisson regression with calendar-month and hiring-firm fixed effects, on the tightness terciles, with the slack (bottom) tercile the omitted reference. Standard errors, in parentheses, are clustered on the hiring firm. *, **, and *** denote significance at the 10%, 5%, and 1% levels.