CES ifo
Working Papers
12744
2026
June 2026

When to Broaden, When to Focus: Job Search and Market Tightness

Jeremias Klaeui

CES ifo

Imprint:

CESifo Working Papers

ISSN 2364-1428 (digital)

Publisher and distributor: Munich Society for the Promotion
of Economic Research - CESifo GmbH

Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740

Email office@cesifo.de
https://www.cesifo.org

Editor: Clemens Fuest

An electronic version of the paper may be downloaded free of charge
· from the CESifo website: www.ifo.de/en/cesifo/publications/cesifo-working-papers
· from the SSRN website: www.ssrn.com/index.cfm/en/cesifo/
· from the RePEc website: https://ideas.repec.org/s/ces/ceswps.html

When to Broaden, When to Focus: Job Search and Market Tightness

Jeremias Klaeui*

June 11, 2026

Abstract

This paper investigates whether directing job search toward tight labor markets improves re-employment. Such targeting is theoretically attractive as it reduces aggregate congestion. I link Swiss unemployment records to clickstream search data to measure which markets jobseekers target. To address selection, I exploit quasi-random assignment to caseworkers differing in their tendency to redirect clients toward tight markets. Moving from the 10th to the 90th percentile raises six-month job-finding by 2.7%. The channels are heterogeneous: jobseekers who become unemployed in slack markets gain from mobility, consistent with prior findings from policies that encourage broader search. The new insight is that jobseekers who become unemployed in tight markets benefit from a narrower focus on this initial market. Effects are orthogonal to other caseworker behaviors and do not come at the expense of job quality. My findings suggest that market tightness offers a simple basis for targeted search advice.

Keywords: job search, labor market tightness, labor supply, matching. JEL Codes: J64, J68, J62

1 Introduction

Preventing long-term unemployment is a major concern for individuals and society. Encouraging jobseekers to broaden their search to related occupations and nearby regions has become a popular tool to this end (Belot et al., 2018). The underlying idea is that jobseekers are employable not just in one occupation but also in similar ones, and that considering a wider set of occupations and regions raises their chances of finding work,

*Klaeui: CREST/ENSAE/Institut Polytechnique de Paris and KOF Swiss Economic Institute, ETH Zurich (email: jeremias.klaeui@ensae.fr). This project benefited from financial support through SNSF grant no. 407740, entitled "What Workers Want". I am very thankful to Rafael Lalive, Roland Rathelot and Michael Siegenthaler for their guidance and continuous feedback on this paper. I thank Ihsaan Bassier, Marion Brouard, Fabrizio Colella, Bruno Crépon, Alexia Delfino, Patrick Kline, John Körtner, Tobias Lehmann, Alan Manning, Andrea Marcucci, Bernhard Nöbauer, Mitia Oberti, Jesse Rothstein, Kristina Schüpbach, Sara Signorelli, Daphné Skandalis, Arne Uhlendorff, Christopher Walters and Josef Zweimüller, as well as the participants of the Copenhagen Workshop on Job Search and Micro Data, CREST internal seminar, UC Berkeley Labor Lunch, PSE Applied Economics Lunch, University of Geneva Institute of Economics and Econometrics Seminar and the CESifo Area Conference on Labor Economics for helpful comments and suggestions.

but jobseekers may not be aware of all the options open to them. Empirical evaluations show that encouraging broader search helps jobseekers in slack markets, where many compete for few jobs and the intervention shifts their attention to alternative markets with better prospects (e.g., Altmann et al., 2026; Belot et al., 2026, 2025). The same advice, however, has no effect when applied uniformly to all jobseekers (Belot et al., 2018; Altmann et al., 2026; Ben Dhia et al., 2022), and can even worsen outcomes when broader search is mandated or subsidized rather than left as a recommendation (van der Klaauw & Vethaak, 2022; Caliendo et al., 2026).[1]

This paper investigates a simple, more targeted alternative: among a jobseeker's initial market and its feasible alternatives, direct search toward the tighter ones. This redirection has an attractive theoretical property: it minimizes congestion. A jobseeker moving from a slack market to a tight one noticeably shortens the long queue in the market they leave, while adding little congestion in the tight market they join. Targeting tight markets therefore raises both individual job-finding and aggregate matching efficiency (Şahin et al., 2014; Marinescu & Rathelot, 2018; Kircher, 2022).

Empirically, redirecting jobseekers toward tighter markets is known to work for those becoming unemployed in slack markets (Belot et al., 2025). For jobseekers in other markets, it is not known whether the same redirection works. Moreover, while jobseekers in slack markets are known to search too narrowly, it is not obvious that this also applies to those in better markets. A prominent information friction targeted by interventions broadening search is that jobseekers do not know which other occupations could fit their skills, likely leading to a too-narrow consideration set. Such a friction is likely less of a problem for jobseekers already in tight markets, who rely less on alternative markets. Thus, their search may already be well targeted, so that an intervention would have no effect or could even have negative impacts if it distorts their search behavior, preventing job matches from forming or producing worse matches. At the same time, a second kind of information friction could matter: jobseekers may not know which markets currently offer the best prospects. Prospects can differ substantially across markets, and jobseekers may struggle to compare them. A jobseeker who underestimates the prospects in their initial market may focus too much on alternatives, leaving room for advice on market prospects to improve outcomes.

Estimating the causal effect of search redirection on job finding faces two main empirical challenges: actually observing the direction of search and endogeneity. The first challenge is measurement: commonly used data sources show only where jobseekers end

[1] Another strand of successful interventions consists of highly personalized recommender systems, for instance Barbanchon et al. (2023); Behaghel et al. (2024); Bächli et al. (2025). However, they come with high data requirements: Barbanchon et al. (2023) requires jobseekers to have clicked on at least three different vacancies; Behaghel et al. (2024) uses predictions of a firm's probability to recruit; Bächli et al. (2025) uses an online assessment of skills and abilities.

up employed, not where they search. Where people end up employed is an equilibrium outcome that reflects both search direction and market tightness, so outcome data alone cannot separate the two. I address this by using clickstream data from job-room.ch, the official job portal of the Swiss public employment service (PES), linked to administrative unemployment-insurance records from the PES. The clicks record which specific job postings jobseekers view, revealing which markets they target before outcomes realize. Markets are defined as cells formed by the interaction of 101 three-digit ISCO occupations and Switzerland's 16 commuting zones. I measure the tightness of each jobseeker's search as the average log vacancy-to-unemployment ratio across the markets in which they click on job postings.[2]

The second challenge is endogeneity: jobseekers self-select into their search strategies. Even conditional on last occupation and location, those targeting slack and tight markets likely differ in unobservables such as skills, preferences, constraints, and private information about their prospects. I address this by exploiting the quasi-random assignment of Swiss jobseekers to caseworkers within local employment offices, as used already by Arni & Schiprowski (2019). This is analogous to "judge-leniency designs" (Goldsmith-Pinkham et al., 2025) which are increasingly applied to caseworkers (Cederlöf et al., 2025; Humlum et al., 2025). In Switzerland, caseworkers regularly meet with their jobseekers and advise on search strategy as part of their role. I measure each caseworker's tendency to direct search toward tight markets using a leave-one-out approach: for each jobseeker, I compute how much the caseworker's *other* clients direct their search toward tight markets, conditioning on those other clients' initial market (the interaction of last occupation, commuting zone, and registration quarter). Conditioning on the initial market accounts for the fact that transition opportunities and their tightness might differ for jobseekers with different backgrounds. The identifying comparison is therefore between two jobseekers in the same office, the same registration quarter, and the same initial occupation-commuting zone market, but assigned to caseworkers who differ in their tendency to redirect jobseekers to tight markets.

Three facts suggest that caseworker redirection toward tighter markets could affect job finding. First, tightness varies substantially even between close-substitute markets: among occupation-region pairs with high worker mobility in the past and wage levels, the tighter market typically has around 150 vacancies per 100 unemployed jobseekers versus 65 in the slacker one, leaving substantial scope for redirection. Second, jobseekers do not appear to perceive these differences: Belot et al. (2025) document that beliefs about job prospects across markets correlate only very weakly with market tightness and application success

[2] To construct market-level tightness, I pair the near-universe of online vacancies with the population of registered unemployed, yielding a vacancy-to-unemployment ratio for each occupation-region cell (similar to Klaeui et al., 2026).

rates.[3] Third, caseworkers are well placed to fill this gap: knowing local labor markets is part of their role, and the PES equips them with tools displaying vacancy and unemployment counts by occupation and region.[4] They also have substantial opportunity to transmit this knowledge through regular meetings, which have been shown to be highly productive (Schiprowski, 2020) and typically include 10–20 minutes spent advising on job search strategies (Egger & Partner, 2013).

My first result is that caseworkers shape clients' search direction: some systematically steer jobseekers toward tighter markets while others do not. A 10% higher search tightness among a caseworker's other clients translates to a 0.8% higher tightness of the jobseeker's own search. To show the transmission in more detail, I combine the exact timing of caseworker meetings with click timestamps. An event study around the first meeting with the caseworker shows that the leave-one-out measure has no effect on click tightness before the meeting but a positive effect after. The zero pre-meeting effect suggests that search tightness is balanced before the caseworker can intervene; the positive effect after the meeting suggests that caseworkers transmit tightness information during counseling. Furthermore, there is no effect on click intensity — the tightness redirection shapes where jobseekers search, not how much.[5]

Turning to the main result, I find that targeting tight markets improves job finding. Moving from a caseworker at the 10th percentile to one at the 90th percentile of the redirection tendency increases job finding within six months by 2.7%. Given the low cost of such advice, the effects are large: they are comparable to raising search effort requirements from the average of 8.3 applications per month to 10.8.[6] For further comparison, a recent experiment using a sophisticated approach to recommend firms that are likely to hire finds an effect of 1% (Behaghel et al., 2024). The effect of tightness targeting persists over time. There is no effect on exit within the first two months of the spell, but thereafter it is significant and positive, the magnitude is 2% at three months, it is highest for five-month job finding (3%) and is still substantial and significant for job finding within a year of the spell start (2.2%).

I next show heterogeneity in the effect by initial market tightness, in the jobseeker's last occupation × commuting zone cell. Unlike the breadth interventions, caseworker redirection improves job finding across the full distribution of initial market tightness, but the channel differs sharply. I observe not only the tightness of the market where jobseekers

[3] Relatedly, Christensen et al. (2026) show that jobseekers also underestimate the dispersion in wages across different job types, suggesting they may similarly underestimate the dispersion in competition.

[4] I further show that caseworkers learn from their recent experiences: an event study documents that when one client finds a job in a given market, the caseworker's other clients subsequently increase their search there, likely due to knowledge transmitted by the caseworker.

[5] A similar pattern appears in Belot et al. (2026): occupational recommendations shifted the composition of search without affecting its overall volume.

[6] Using the estimates of Arni & Schiprowski (2019) and assuming a constant hazard function to convert effects from spell durations to hazard rates.

become unemployed, but also the tightness of the market where they eventually find work. This allows me to ask whether the treatment works by redirecting jobseekers away from slack markets — the channel typically targeted in occupational mobility interventions — or by reinforcing search in markets that are already tight, preventing jobseekers from drifting into segments with worse prospects. For jobseekers starting in slack markets, the treatment works through mobility: it increases the hazard of finding a job in a tighter market by 4%. For jobseekers starting in tight markets, the effect works through reinforced focus: job finding in the initial market increases by 8%. For those in medium-tightness markets, both channels operate: job finding improves in tighter markets by 5% and in the initial market by 5%. Across all three groups, there is no positive effect on finding a job in markets slacker than the initial one. This pattern reconciles the literature's mixed findings on broader search: encouraging mobility benefits slack-market jobseekers, but for tight-market jobseekers the relevant margin is reinforced focus.

As a next result, I show whether the effects predominantly work through occupational or geographical mobility. My results suggest that the main margin is joint mobility: jobseekers redirected toward tighter markets tend to switch both occupation and commuting zone simultaneously. There are also, smaller, effects on occupational mobility but none on purely geographic mobility. While these joint moves involve a switch of commuting zone, the actual distance between the home municipality and the job's municipality is not significantly larger. This suggests that the moves are not far and likely reflect crossing into a neighboring administrative unit rather than expanding search to more distant locations.

Caseworkers influence outcomes through many channels at once. A natural concern is that the tightness tendency may simply proxy for one of these other behaviors. To test this, I compute analogous leave-one-out measures for sanction rate, training referrals, meeting frequency, and meeting duration, and examine their correlations. The tightness tendency is nearly orthogonal to all of them, indicating that caseworkers who redirect toward tighter markets do not systematically differ in how strictly or intensively they manage their clients. To go further, I ask whether any other caseworker behavior produces a similar pattern of effects. Several dimensions are themselves associated with better job-finding outcomes, notably sanction rate and training course referrals. Other, harder-to-observe dimensions such as CV writing advice or motivational coaching could plausibly do the same. However, for all of these "better caseworker" explanations, one would expect job finding to increase uniformly across all market types. This is indeed what I find for the dimensions I can measure: sanction rate, training referrals, stringency in assessment of job search effort all correlate uniformly with job finding in tighter, slacker, and initial markets. Only the tightness tendency produces a different pattern: positive effects in tighter and initial markets, with zero effect in slacker markets.

Finally, I can rule out negative effects on job quality: there is no significant effect on return to unemployment, reemployment earnings, nor the median wage level associated with the found market.

My findings carry direct implications for how employment services should design job search advice. For slack-market jobseekers, encourage broader search toward tighter alternatives; for tight-market jobseekers, reinforce focus on the initial market. Market-level tightness provides a simple way to tailor this advice—it can be constructed from vacancy and unemployment counts that most public employment services already collect. The conceptual framework in Section 2 implies that these gains are unlikely to come at the expense of other jobseekers. Redirection reduces competition where it is fiercest and adds workers where vacancies are abundant enough to absorb them. Jobseekers leaving slack markets relieve pressure on the scarce openings there, while those staying focused in tight markets fill positions that might otherwise go unfilled and simultaneously stop competing for scarce openings in slack alternatives.

My paper contributes to several strands of literature. The first contribution is to the literature on search and mismatch. Kircher (2022) shows theoretically that redirecting search toward tighter markets can be individually beneficial if productivity and search efficiency gaps are not too large, and that the aggregate benefit is increasing in the individual benefit. Şahin et al. (2014) and Marinescu & Rathelot (2018) show that mismatch is reduced when workers are moved to tighter markets. I provide empirical evidence that redirecting toward tight markets benefits jobseekers.

The second contribution is to the literature on broadening job search (Belot et al., 2018; Altmann et al., 2026; Belot et al., 2026, 2025; Ben Dhia et al., 2022; van der Klaauw & Vethaak, 2022; Caliendo et al., 2026; Behaghel et al., 2024; Bächli et al., 2025). They find positive effects of targeting alternative markets in general and tight markets in particular but the positive effects are limited to jobseekers becoming unemployed in slack markets. I add to this literature by showing that targeting tight markets improves job finding across the full distribution of initial market conditions, through mobility for some jobseekers and reinforced focus for others. The paper also relates to a growing literature on recommender systems for job search (Naya et al., 2021; Bied et al., 2026), which combines vacancy-level success probabilities with preferences and emphasizes congestion when many jobseekers receive the same suggestion. I work with a coarser, market-level proxy—tightness—that can be built using widely available data, and I show effects on realized job-finding outcomes complementing the existing studies' simulations and results on clicks and applications.

By using the caseworker tendency as a shifter in search scope, I also contribute to a large body of work on how caseworkers affect jobseekers' outcomes (Behaghel et al., 2014; Huber et al., 2017; Arni & Schiprowski, 2019; van den Berg et al., 2019; Schmieder

& Trenkle, 2020; Dromundo Mokrani & Haramboure, 2022; van den Berg et al., 2023; Humlum et al., 2025; Cederlöf et al., 2025; Rambjer et al., 2026). The detailed jobseeker behavior data allows me to investigate the content of caseworker advice to a degree of detail not achievable with the caseworker actions recorded in standard administrative datasets. Further, my variance-covariance analysis confirms the effects documented in many existing studies within a unified setting and shows that tightness redirection is nearly orthogonal to these other caseworker dimensions.

The paper most closely related to mine is Belot et al. (2025), who conduct a large-scale Dutch RCT providing jobseekers in slack occupations with information about the poor prospects and suggestions of tighter alternatives. They find large effects: employment, hours, and earnings improve by 5–6% after 18 months, and a calibration exercise shows that spillovers are small precisely because tightness is highly imbalanced. Their results confirm that tightness-targeted advice works for slack-market jobseekers. I extend this by studying jobseekers across the full distribution of initial conditions, showing that tightness-targeted advice also improves outcomes for those already in tight markets, but through reinforcing focus rather than encouraging mobility.

A second closely related paper is Altmann et al. (2026), who study online job search advice in a large-scale Danish RCT. They test two treatments: occupational recommendations and vacancy information, which shows the number of local vacancies in the occupations jobseekers are searching for. Both improve employment outcomes, but they document negative spillovers when many jobseekers receive the same advice. Occupational recommendations only help jobseekers in low-tightness occupations; vacancy information also helps those in favorable conditions. They call for research on personalized advice exploiting labor market tightness heterogeneity. I study advice that targets tightness explicitly, advancing on their study in the following ways. Tightness goes beyond vacancy information by capturing not only where demand is but also how many other jobseekers are competing for it. More fundamentally, I show that tightness-directed advice nests the channels through which each of their treatments operates. Like their occupational recommendations, it encourages mobility for jobseekers in slack markets; like their vacancy information, it reinforces search in the initial market for jobseekers in favorable conditions. At the same time, it avoids encouraging broad search for jobseekers who are already in tight markets.

The remainder of this paper is organized as follows. Section 2 presents the conceptual framework. Section 3 describes the institutional setting of Swiss unemployment insurance and the role of caseworkers. Section 4 presents the data sources and documents descriptive facts about search scope and market tightness. Section 5 develops the identification strategy. Section 6 presents balance tests, evidence on search behavior, the main results, heterogeneity analysis, and robustness checks. Section 7 examines interactions with other caseworker dimensions. Section 8 concludes.

2 Conceptual framework

I borrow from the stylized random search model in Kircher (2022)[7] to motivate the empirical analysis. The model features multiple labor market segments in which firms post vacancies and unemployed workers search for jobs. Workers can distribute their search effort across multiple segments with different intensity. In each segment, the number of matches formed depends on the number of vacancies and the number of searchers. The individual job-finding rate therefore depends on market tightness—the ratio of vacancies to unemployed—in the segments a worker targets: as more workers search in the same segment, each faces greater competition and a lower chance of finding a job. When tightness differs across segments, the allocation of search effort can be inefficient: a worker concentrating effort in a slack market forgoes better prospects elsewhere. Redirecting search toward tighter segments can therefore improve individual outcomes.

Beyond improving individual outcomes, search redirection can also raise total matches in the economy. The concavity of the matching function is the key mechanism: moving a worker from a slack market to a tight one generates a large reduction in congestion where she leaves—where many searchers compete for few vacancies—and a small increase in congestion where she arrives, where vacancies are abundant relative to searchers. A key implication is that the individual benefit from redirection and the aggregate benefit are aligned: individual gains arise precisely when tightness is unequal across markets, which is exactly the condition under which aggregate gains are large. Evidence of positive individual effects from search redirection toward tight markets is therefore directly informative about the sign of the aggregate welfare effect. In the extreme case where workers are equally productive across markets, the planner's objective is maximized when tightness is equalized across markets (Kircher, 2022, Proposition 1). In reality workers are likely to be productive only in a few markets (e.g. a data scientist would likely not be a very productive baker). By continuity, it holds that small redirections toward tight markets are, individually and on aggregate, beneficial even if productivity is dispersed.

The alignment between individual and aggregate effects also informs the interpretation of my estimates. My identification exploits that caseworkers differ in how much they redirect search toward tight markets. My results show that this variation exists, so not all caseworkers redirect equally and there is scope for reallocation, given current productivity dispersions.[8] The estimates therefore reflect the marginal effect of redirection in the

[7] Şahin et al. (2014) and Marinescu & Rathelot (2018) reach analogous conclusions through mismatch indices.

[8] Furthermore, Section 4.3 provides descriptive evidence on tightness dispersion even between neigh-

current setting where substantial tightness differences remain. If a policy maker were to make more caseworkers give advice to redirect search toward tight markets, more workers would flow toward tight markets and tightness differences would shrink. As tightness equalizes, the marginal gain from further redirection diminishes—the individual effect I estimate would be smaller under universal adoption. Note that this is inherent to any intervention that targets frictions: the more successfully it reduces mismatch, the less scope remains for further improvement. Crucially, the aggregate effect remains weakly positive throughout this process—gains only vanish once tightness is fully equalized, at which point the planner's objective is maximized.[9] My estimates therefore represent an upper bound on the per-person effect under universal adoption, but the finding of positive individual effects implies that the current allocation is inefficient and that broader adoption of search redirection would improve aggregate welfare.

3 Institutional background

Switzerland provides a particularly suitable institutional setting for studying how caseworkers influence unemployed jobseekers. Caseworkers play a central role in the Swiss unemployment insurance system, with advising as their most time-intensive task. Importantly for my research design, the allocation of jobseekers to caseworkers within an unemployment office is arbitrary. This quasi-random assignment feature generates exogenous variation in exposure to caseworkers with different advising styles, which I exploit in the empirical analysis.

Registration with the Public Employment Service (PES) is mandatory in order to receive unemployment benefits. As in many countries, jobseekers must provide and document search efforts, attend regular caseworker meetings and accept suitable offers or face sanctions. Caseworkers monitor compliance with those rules but also, importantly for this paper, provide advice, which is among their main responsibilities (Lalive et al., 2005; Arni et al., 2013). A detailed time-use study commissioned by the Swiss State Secretariat for Economic Affairs (SECO) documents how PES caseworkers allocate their working hours across activities based on caseworker surveys and personal interviews (Egger & Partner, 2013).

Their responsibilities span administrative processing of unemployment benefits; monitoring compliance and imposing sanctions; advising and developing strategies for jobseekers; and providing placement services, including referrals to active labor market programs

boring markets.

[9] This hinges on caseworkers having up-to-date knowledge of labor market conditions in their domain. With access to the live version of unemployment insurance records and with the internet being an increasingly important channel for vacancies—where most platforms allow users to see vacancy counts—this is not an entirely unreasonable assumption.

(ALMPs).[10] The time-use chart in Appendix Figure C.1 from Egger & Partner (2013), illustrates how caseworkers allocate time across activities, with advising standing out as central. 32.5% of resources are allocated to the advising task, the second-most time-intensive task is case administration with 7.6% of time resources allocated to it. Zooming in more, the study finds that the advising task typically consists of 10–15 minutes per meeting evaluating job search efforts as well as 15–20 minutes developing an individualized strategy to return to the labor market.

Upon registration, each unemployed individual is assigned to a caseworker at their local PES office. Those local unemployment offices (RAV) are responsible for overseeing monitoring and placement activities and also for managing caseworker assignments, with the assignment rules set at the cantonal level. The quasi-random nature of the allocation process in Switzerland was first used as an identification strategy in Arni & Schiprowski (2019). They emphasize that the allocation does not rely on individual employability or performance potential, and the officer in charge of the assignment only observes the same administrative information used in the analysis. As a result, conditional on office and last occupation fixed effects, the assignment of jobseekers to caseworkers can be regarded as good as random.[11]

To go into more detail about the allocation process, Table C.4 in the Appendix reports the results of a self-administered survey of all 26 Swiss cantons on how this assignment is carried out. Of the 21 cantons that responded, ten (covering 44% of the jobseekers in the sample) report assigning new clients to the next free caseworker (in some cases equalizing caseload), six (covering 22%) use equal distribution proportional to contractual hours regardless of current caseload, four (13%) assign cases at the discretion of team or office management, and one (13%) uses software-based random allocation; the five non-responding cantons account for the remaining 8%. Among the assignment methods, equal distribution by administration regardless of caseload and software-based random allocation are the most clearly exogenous, as they do not condition on current caseworker performance or workload. Caseload balancing could in principle create non-random sorting if more effective caseworkers have smaller caseloads—for instance because their clients exit unemployment faster—so that newly registering jobseekers are disproportionately assigned to them. Assignment by management discretion raises similar concerns if managers route certain types of jobseekers to specific caseworkers. As a robustness check, I will

[10] ALMPs consist of basic courses aimed at improving job search effectiveness and self-esteem, language and computer courses, further vocational training, other occupation-specific training programs, and PES-financed wage subsidies for temporary jobs.

[11] My identification strategy uses unemployment × office fixed effects, mimicking the design in Arni & Schiprowski (2019) and additionally last occupation × commuting zone × quarter fixed effects that would remove all sorting that is additively separable in unemployment office and last occupation. In a robustness check, I interact the unemployment office with the last occupation to account for potential non-additive effects.

restrict the sample to cantons using equal distribution or random allocation and show that the results are unchanged on this subsample (35% of spells; Appendix Table C.5).

To redirect jobseekers toward tight markets, caseworkers need information about where demand is high relative to supply. Indeed, as confirmed in personal talks to PES officials, maintaining an up-to-date overview of the local labor market is a core part of caseworkers' professional mandate and several channels provide this information. Caseworkers have access to internal PES software that displays both current job vacancies and registered jobseekers, searchable by occupation and region. While the tools used differ from canton to canton, the same data underlie the public job portal job-room.ch, where vacancy listings are displayed alongside their total count, and anonymized candidate profiles are searchable with the number of matching candidates shown (see Appendix Figure C.2). The combination of vacancy counts and candidate counts effectively reveals market tightness. [12] A further natural channel is the caseworker's personal experience. I provide suggestive evidence for such learning mechanisms in Appendix Figure D.1: after one of a caseworker's clients finds a job in a given occupation, the caseworker's other active clients increase their clicks in that same occupation, suggesting that the caseworker may interpret the unemployment exit as a signal of good prospects and pass that information on to other clients.[13]

4 Data and descriptive statistics

My study combines data from three different sources: i) administrative data about unemployed jobseekers in the Swiss unemployment register, ii) click data from job-room, the official job portal of the Swiss public employment service, iii) data on the near-universe of online vacancies from X28, a web-scraping company.

Unemployment Register. The administrative register, maintained by the State Secretariat for Economic Affairs (SECO), records all individuals registered with the Swiss public employment service (PES) (State Secretariat for Economic Affairs (SECO), 2024a). For each unemployment spell, the data records start and end dates, demographic and human capital characteristics, last occupation before unemployment (coded in both the

[12] As an illustration, Appendix Figure C.2 shows two neighboring sales occupations in Basel: ISCO 332 (sales and purchasing agents and brokers) lists 597 vacancies for 201 candidates, implying a vacancy-to-unemployment ratio of approximately 3, while ISCO 522 (shop salespersons) lists 749 vacancies for 778 candidates, a ratio close to 1. Despite both being sales occupations, tightness differs by a factor of three—information that is visible on the portal and in their internal tools and that a caseworker could use when advising a client in the less favorable market.

[13] The effect already appears one to two months before the client formally deregisters from the PES, consistent with the typical lag between accepting an offer, informing the caseworker, and officially starting the new position.

PES internal classification and ISCO), municipality of residence, and identifiers for the assigned caseworker. The caseworker identifier enables the quasi-random assignment strategy described in Section 5. The data also documents meeting dates and UI benefit sanctions for non-compliance with job search requirements, enrollment in training programs and receipt of wage subsidies for temporary jobs and whether the caseworker judged the monthly search efforts (reported applications) as meeting the effort requirements. These variables allow me to compare the tightness-tendency measure to other caseworker behaviors.

The analysis uses 439,679 unemployment spells with benefit entitlement starting between June 2020 and September 2022. June 2020 is the start of the click data; September 2022 is the latest start date for which the full three-month click window falls within the click sample, which ends in December 2022. I apply several restrictions. First, I require at least one caseworker meeting in the first three months, excluding 2,540 spells (0.6%). Second, I drop 189 spells coming from unemployment after working in military occupations. Third, I exclude 30,377 temporary layoffs (7%) where workers returned to their previous employer, as these cases likely do not involve genuine job search (Liechti et al., 2020). Finally, I exclude spells with missing occupation codes for the last job (7,529 spells) or missing commuting zone of residence (1,746 spells). These restrictions yield 397,298 spells.

I further restrict to spells whose caseworker has at least five clickers in the click data—the variation that identifies the leave-one-out tendency. This restriction excludes only 5,284 spells (1.3%), so the analysis covers essentially the universe of unemployed jobseekers in this period. The regression specification includes occupation $\times$ commuting zone $\times$ quarter fixed effects as well as unemployment office $\times$ year interactions. After dropping the 25,121 singletons—spells or caseworker assignments perfectly predicted by those fixed effects—the final regression sample comprises 366,893 unemployment spells from 2,460 caseworkers across 132 local employment offices.[14]

Click data from job-room.ch. Job-room.ch is the official job portal of the Swiss public employment service, available in German, French, Italian, and English (State Secretariat for Economic Affairs (SECO), 2024b). The portal aggregates job listings from two sources: direct postings by companies (free of charge) and vacancies scraped from other job portals, with the aim of covering as much of the labor market as possible. At any given time, between 70,000 and 110,000 job postings are listed.The platform does not feature paid placements or sponsored listings, so the postings a jobseeker sees reflect purely their search filters and not any platform-side promotion. In addition to job listings,

[14]While the singletons are not a problem in a standard linear regression that uses a within transformation, they are a problem when computing the leave-one-out score detailed in Section 5 as that computation involves inverting the predictor matrix.

the portal allows registered unemployed to log their mandatory application protocols, so that jobseekers can use job-room.ch both to manage their unemployment spell and to search for jobs.

To search for jobs on the portal, jobseekers specify criteria such as occupation, competencies or skills, and work canton or municipality (see Appendix Figure C.3 for a screenshot). Based on these criteria, the portal returns a list of matching vacancies. As jobseekers type, the portal offers auto-completion suggestions drawn from several occupation classifications (ISCO and internal PES lists) and location definitions (municipality and canton) that vary in breadth. A search for a given occupation only yields results that exactly match the entered occupation; if a location filter is specified, the portal implements a default radius of 30km around the location, which can be changed by the jobseeker. To access detailed information about a vacancy—its full description, requirements, and application instructions—the jobseeker must click on the listing (see Appendix Figure C.4 for an example). These clicks are the central measure in this study. Each click reveals that the jobseeker found a particular vacancy worth investigating further, given the occupation and location they searched for. Hence, the clicks also serve as a proxy for the occupation and location criteria the jobseeker entered.

To study how caseworkers shape where jobseekers search, I need a measure of search direction—which occupation–region markets a jobseeker targets. I use the clicks on job postings as clicks capture an early stage of the search process. Opening a vacancy to read more about it signals interest but involves minimal cost, unlike submitting an application, which typically requires a cover letter and entails potential rejection costs. Clicks therefore reveal the set of markets the jobseeker is willing to explore, whereas applications reflect a more selective decision already weighing expected returns and costs, and are closer to an equilibrium outcome.

Two potential concerns about clicks as a measure of search direction deserve discussion. First, job-room.ch captures only a portion of a jobseeker's total search activity—jobseekers may also use other job portals, personal networks, or other informal channels. The measure therefore reflects search direction on this platform rather than the universe of search effort. Second, a click is a relatively low-cost action: a jobseeker may open a vacancy out of curiosity or while casually browsing, without any genuine intent to pursue that occupation or region. If many clicks reflect idle exploration rather than directed search, they could be a noisy measure of the markets a jobseeker actually targets. Despite these concerns, clicks turn out to be a strong proxy for actual search direction. Figure 1a shows that the number of clicks a jobseeker places in an occupation–commuting zone market is strongly predictive of the probability of finding a job in that market after a completed unemployment spell. This suggests that clicks do capture meaningful search intent and that the clicks on the platform are informative about a jobseeker's complete search scope.

I observe all clicks by registered unemployed in the window from 6 June 2020 to 31 December 2022. Each click record contains a timestamp, a person identifier that links the click to the unemployment register, and a job posting identifier. The posting identifier allows me to retrieve the full content of each vacancy from the publicly available API of job-room.ch, including postings that are no longer online at the time of the request. Appendix A.2 describes the API scrape, click deduplication, and the mapping of postings to occupations and commuting zones.

Tightness measure. I measure labor market tightness for cells defined by ISCO 3-digit occupation and commuting zone. Tightness is defined as the stock of vacancies divided by the stock of unemployed jobseekers. The granularity of the cells is chosen to be broad enough to avoid measurement error from classification ambiguity or small denominators, while remaining granular enough to capture job search as a weighted average of different markets. In the results section, I show the robustness of the estimates to different market definitions.

In the computation of the tightness measure I use variables external to job-room.ch. This avoids mechanical dependence between click behavior on the platform and the tightness measure, and yields a measure that can be replicated in other settings using data that is now widely available to PES and researchers. I use the number of jobseekers by occupation and commuting zone from the unemployment records and the number of vacancies from a web-scraped dataset of the near-universe of online vacancies. I also compute a tightness measure using data from the platform, the ratio of job postings to distinct clickers in a market, and show that the results are not sensitive to the choice of tightness measure.

For the denominator, I construct the unemployment count U_{mt} from the administrative register of all jobseekers registered with the Swiss unemployment services. A feature of the PES data is that, as in many other countries, there are multiple searched occupations recorded for every jobseeker. These registered occupations are binding: jobseekers can be sanctioned for refusing a suitable offer in one of them, so they define the markets in which a jobseeker constitutes effective labor supply from the perspective of firms and the PES. I count each jobseeker toward the labor supply in their commuting zone of residence in all occupations listed in the register. To give all the jobseekers an equal weight, I count a jobseeker with k registered occupations with a weight of $1/k$. Finally, I use a 30-day rolling average to smooth out daily volatility and weekday effects; market cells with no jobseekers on a given day are assigned a count of zero.

For the numerator, I use a dataset containing the near-universe of online vacancies in Switzerland, scraped from the web by X28, a human resources company (X28 AG, 2024).

Figure 1: Validation of clicks as a measure for search direction and tightness as a measure of job finding rates

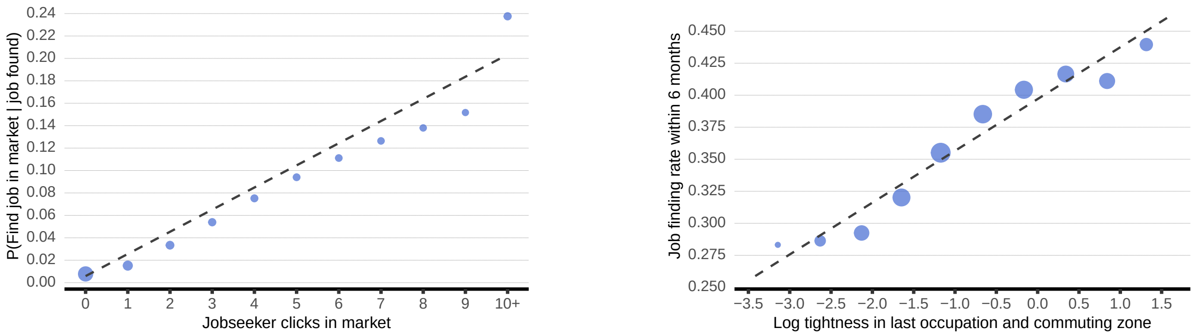

(a) Number of clicks per market and probability that the re-employment job is in the market.

(b) Job finding within 6 months and tightness in last occupation–commuting zone market.

Notes: Both panels use binned scatter plots where each point represents the mean within a bin, with point size proportional to the number of observations. Panel (a) conditions on finding a job and plots the probability of being hired in a market (ISCO 3-digit occupation × commuting zone) against the number of clicks placed there. For each jobseeker, the choice set includes all markets with at least one click plus 50 randomly sampled alternatives from the universe of markets with at least 30 total clicks. Panel (b) relates the log tightness in the jobseeker's last occupation and commuting zone of residence at the time of registration to the probability of finding a job within 6 months. The tightness measure is truncated at the 5th and 95th percentiles.

The data has, among others, been used by Colella (2022) to measure skill-biased technological change and in Klaeui et al. (2026), also to compute tightness measures. The dataset contains the date of posting and the date of removal of each vacancy, allowing me to calculate the stock of vacancies at any point in time (unlike other data providers, such as Lightcast, which measure only the inflow). Each posting is classified by occupation using the same internal system as the unemployment records, which I map to ISCO. The vacancy data covers several million postings between January 2020 and December 2022. I apply a series of cleaning steps and match vacancy postcodes to commuting zones using spatial overlap, both detailed in Appendix A.1. Once I have computed the daily stock at the occupation–commuting zone level, as for the jobseekers, I apply a 30-day rolling average to smooth out daily volatility and weekday effects.

The predictive content of tightness is shown in Figure 1, panel (b): the six-month job-finding rate rises nearly linearly and substantially with log tightness in the jobseeker's last occupation–commuting-zone market. This suggests that tightness is a relevant measure of labor demand that caseworkers could use when advising clients.

4.1 Search tightness index

The search tightness index is the central measure of search direction in this paper. It summarizes where a jobseeker directs their search in a single number that captures the average labor market tightness of targeted markets. As shown in the descriptive statistics, the median jobseeker searches across six occupations and two commuting zones. Two individuals with identical backgrounds — the same last occupation, residence, and registration date — can therefore face different labor demand conditions depending on which mix of markets they target. The intuition of the index is to capture this effective tightness, so that I can then examine how caseworker-induced shifts in it affect search and employment outcomes. Figure 2 illustrates the idea with two stylized search patterns.

For each click, I assign the log tightness of the clicked market (ISCO 3-digit occupation $\times$ commuting zone) at the date of the click, using the smoothed vacancy-to-unemployment ratios described above. I drop clicks in markets with fewer than 30 registered jobseekers (4% of clicks) to ensure reliable tightness measurement. The index is the simple mean of log tightness across all retained clicks. For instance, a jobseeker who clicks twice on a market with log tightness of 1 and three times on a market with log tightness of 2 has an index of $(2 \cdot 1 + 3 \cdot 2)/5 = 1.6$. I restrict attention to the first three months of unemployment and to jobseekers with at least five clicks with non-missing tightness in this window. The three-month restriction captures initial search strategy before jobseekers may endogenously adjust their behavior in response to prolonged unemployment; as spells lengthen, jobseekers might broaden their search toward less selective markets, which would contaminate the index with the consequences of not finding a job rather than the caseworker's initial guidance. I show robustness to the window choice as well as to the 5 clicks cutoff. This index is the variable that enters the leave-one-out caseworker measure constructed in Section 5.

4.2 Outcome variables

The main outcome is whether a jobseeker finds employment within 180 days of registering with the PES. I define job finding from the administrative deregistration records: a spell is coded as ending in employment if the register records a transition to a job. Jobseekers who do not exit to employment within 180 days, whether because they remain registered, exhaust benefits, or leave the register for other reasons, are coded as zero. The baseline analysis focuses on the 180-day horizon; I also report results for shorter and longer windows.

The administrative data allows me to observe not only whether and when a jobseeker finds a job, but also where. I exploit this by decomposing the unconditional job-finding

Figure 2: Schematic illustration of the search tightness index

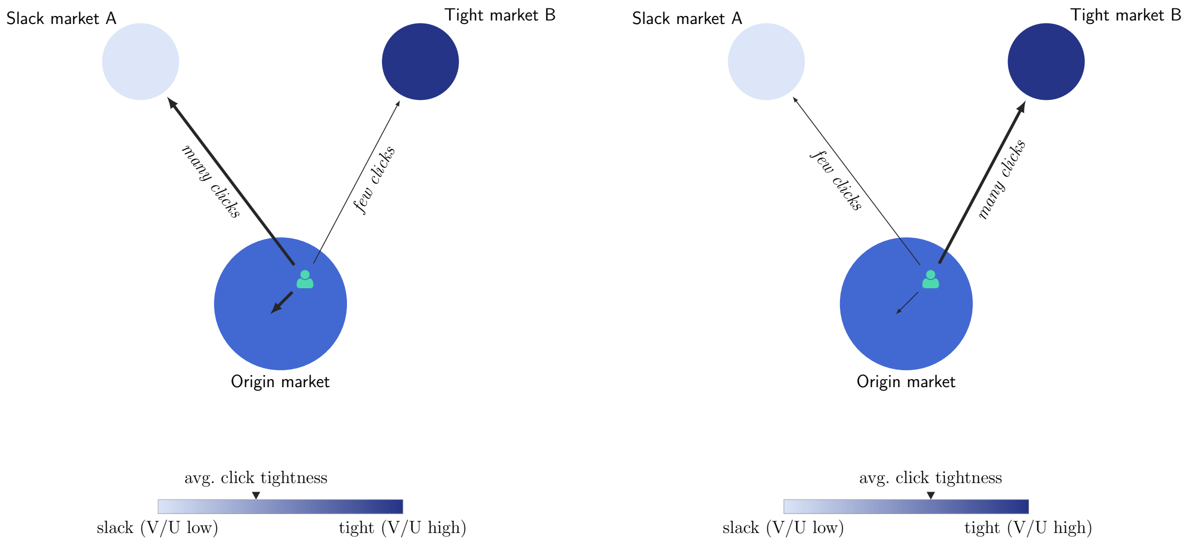

(a) Low search tightness

(b) High search tightness

Notes: Stylized illustration of the search tightness index. The jobseeker (green) sits inside her origin market (last occupation x commuting zone of residence) and clicks on the origin and three other markets, shaded by labor market tightness. Arrow thickness reflects the number of clicks; the bottom bar shows the resulting average click tightness, i.e. the search tightness index. Panel (a): a jobseeker who concentrates clicks on slack markets has a low index. Panel (b): a jobseeker who concentrates clicks on tight markets has a high index. The only difference between the two is the distribution of clicks across the markets.

probability into three mutually exclusive sub-events based on the tightness of the market in which the job is found relative to the jobseeker's initial market. The initial market is the last occupation (ISCO 3-digit) $\times$ commuting zone of residence at registration. The found market is the occupation $\times$ commuting zone of the new job. I measure tightness in both markets at the date of deregistration and classify each exit into three categories: tighter market, initial market, and slacker market. Each category is coded as an indicator equal to one if the jobseeker finds a job of that type within 180 days, and zero otherwise. The three indicators therefore decompose the overall job-finding probability and allow me to test whether caseworker-induced search redirection translates into employment specifically in tighter or the initial market.[15]

[15] Among job finders, 17% have a missing found occupation and 11% have a missing commuting zone; overall, 25% are missing at least one of the two. For these spells I impute the missing found-market cell from the new employer's industry, which is recorded on the register. Within a region, the industry of the new job is informative about the occupation: a retail employer in Geneva most often hires a cashier, not a software engineer. I use the empirical joint distribution of occupations and industries within each region to weight the candidate cells and compute an industry-weighted average vacancy-to-unemployment ratio for the imputed market. Appendix A.5 details the construction. The imputation closes nearly all of the gap, leaving only $\sim$1.5% of job finders without a defined found-market tightness. A robustness specification shows that the effects from regressions that consider the missing values as a separate market (exit to unknown market) instead of imputing the tightness are very similar to the baseline.

Job quality outcomes. To obtain measures of the quality of the job found after a successful unemployment spell, I examine four outcomes conditional on finding employment. *Reemployment earnings* are obtained from the Swiss social security system (AHV/AVS) (Federal Social Insurance Office (FSIO), 2022). The AHV records all employment relationships for the resident population. Earnings are uncapped and include variable pay components such as bonuses and overtime; wage earnings are reported in real time to the social security administration, independently of the income tax system (Martinez et al., 2021). The earnings measure I use is the log of the maximum monthly employment income across the first three full calendar months after deregistration. I further winsorise the resulting distribution at the 5th and 95th percentiles. My cleaning follows Martinez et al. (2021); construction details are provided in Appendix A.3. The AHV data ends in December 2021, so the earnings sample is restricted to deregistrations before October 2021, covering approximately 41% of job finders in the regression sample. The second measure for earnings is the *market wage level*, the log of the median wage in the found occupation $\times$ commuting zone cell. Wages are computed from the pooled 2020 and 2022 Swiss Earnings Structure Survey (ESS) (Swiss Federal Statistical Office (SFSO), 2026), using the standardized full-time-equivalent monthly wage; the median takes the survey weights into account. This measure captures whether redirection steers jobseekers toward higher- or lower-paying markets, independently of individual wage negotiation, and serves as a second proxy for reemployment wages that, unlike the AHV measure, is available for the full sample period. 77.5% of job-finders match to a cell with at least 30 ESS observations; the remaining 22.5% consist almost entirely of spells with missing occupation or region codes for the new job rather than thin ESS cells.

Next, *the return to unemployment* is a binary indicator for re-registering at the PES within 180 or 360 days of deregistration, constructed from subsequent spells in the administrative register. The measure is censored at the last recorded deregistration date (March 2023), so the effective sample is smaller for the 360-day horizon. Finally, *commuting distance* is the log driving time (in hours) between the municipality of residence and the municipality of the found job. I compute driving distances between municipality centroids using OpenRouteService (Heidelberg Institute for Geoinformation Technology (HeiGIT), 2024). For jobseekers who find a job in their municipality of residence, I approximate within-municipality commuting distance by drawing 50 random points within the municipality, computing pairwise driving distances, and taking the average. Construction details are provided in Appendix A.4.

4.3 Descriptive statistics

Table 1 reports characteristics of the 366,893 unemployment spells in the regression sample. Since some jobseekers register more than once, these correspond to 323,224 unique

Table 1: Sample characteristics

	Full sample		Clickers	
	Mean	SD	Mean	SD
Demographics				
Female	0.47	0.5	0.54	0.5
Age 35–49	0.34	0.47	0.35	0.48
Age 50+	0.22	0.41	0.22	0.42
Has children	0.4	0.49	0.4	0.49
Non-Swiss	0.47	0.5	0.41	0.49
Education				
Primary education	0.26	0.44	0.19	0.39
University education	0.17	0.37	0.2	0.4
Employment history				
$>$ 3 years tenure in last job	0.66	0.47	0.67	0.47
Employed in year $t-1$	0.68	0.46	0.63	0.48
Employed in year $t-2$	0.67	0.47	0.61	0.49
Avg. monthly income in year $t-1$ (1,000 CHF)	3.7	3.79	3.42	3.79
Avg. monthly income in year $t-2$ (1,000 CHF)	3.34	3.54	3.06	3.5
Outcomes				
Job finding within 6 months	0.37	0.48	0.4	0.49
Found in tighter market	0.13	0.34	0.15	0.35
Found in origin market	0.11	0.31	0.11	0.31
Found in slacker market	0.12	0.33	0.13	0.34
Found, market unobserved	0.01	0.07	0.01	0.07
N	366,893		101,514	

Notes: The table compares characteristics of the full regression sample and the subset of jobseekers who clicked on at least five job postings in the first three months of unemployment.

jobseekers; below I use "spells" and "jobseekers" interchangeably. The sample is roughly balanced by gender (47% female) and nationality (47% non-Swiss), with 40% having children. About a quarter have only primary education and 17% hold a university degree. Most jobseekers have stable employment histories: 66% had more than three years of tenure in their last job, 68% were employed in the year before registration, and average monthly income in the year before unemployment was 3,700 CHF. Within six months, 38% find employment. Among those finding a job, mobility between markets seems to be the rule rather than the exception: the initial market is the single market with most exits, but only 35% of the six-month exits are to the initial market. The remaining two thirds of the job finders find a job in another market. Slightly more than half of them find a job in tighter markets (10% exit rate) and the other half in slacker markets (9% exit rate).

The right panel compares the full sample to the 101,514 jobseekers (28%) who clicked on at least five job postings in the first three months. Clickers are more likely to be female, more educated, and less likely to be non-Swiss. Job-finding rates are slightly higher among clickers, but importantly, the decomposition into tighter, origin, and slacker markets is

Figure 3: Tightness dispersion across markets and the direction of jobseeker search

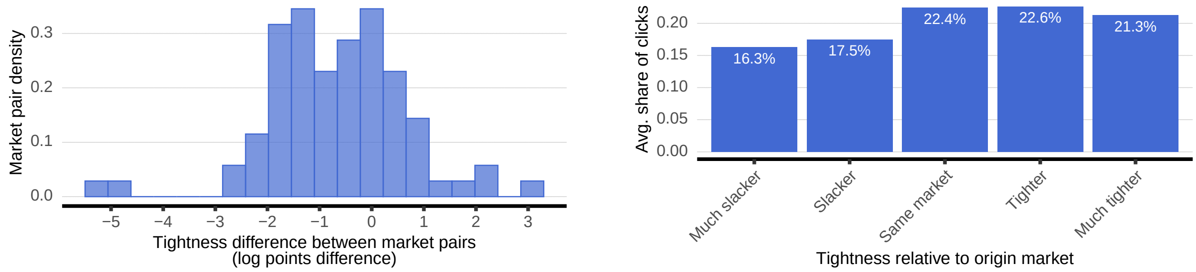

(a) Tightness dispersion across substitutable market pairs

(b) Direction of clicks relative to initial market

Notes: Panel (a) plots the distribution of absolute log tightness differences across market pairs, where a market is an ISCO 3-digit occupation × commuting zone cell. The sample is restricted to pairs with historical transition rates of at least 20% (based on 2004–2019 job-to-job transitions) and whose median market wages differ by less than 5%. The vertical line indicates the mean. Panel (b) shows the distribution of clicks by tightness of the clicked market relative to the jobseeker's initial market (last occupation × residence commuting zone), with tightness measured at the date of the click. Within the tighter and slacker categories, I split at the median tightness difference to distinguish "much tighter" from "tighter" and "much slacker" from "slacker". Tightness values are computed averaging over daily stocks in March 2021 (middle month of the sample). Sample restricted to jobseekers with at least 5 clicks in the first 3 months.

similar across both groups. The comparison matters because the caseworker tightness tendency is measured from clickers, whose clicks reveal which markets a caseworker steers clients toward, but the tendency is then assigned to all jobseekers of that caseworker, including non-clickers, and I estimate effects on outcomes for the full sample. The modest differences between the two groups are therefore reassuring: clickers and non-clickers are similar enough that a tendency estimated from one group's search behavior is likely to reflect advice given to all clients. As a further check, I show that the baseline results hold when restricting the sample to non-clickers only, whose own clicks do not enter the leave-one-out measure at all.

The search behavior on the portal confirms the high mobility patterns seen in the realized transitions. The median clicker places 22 clicks (mean 41) in the first three months of their spell. These clicks spread broadly across markets: the median jobseeker covers six three-digit occupations and two commuting zones, and the initial market absorbs only 22% of clicks. Given that jobseekers routinely search and find jobs across market boundaries, the question of in which direction this mobility should go becomes especially relevant.

The next key question is whether there is meaningful scope for redirecting jobseekers toward tighter markets. This requires both that tightness varies substantially across markets and that jobseekers' current search is not already strictly directed toward tight markets. Panel (a) of Figure 3 shows the distribution of tightness differences across market pairs that are plausibly substitutable from the jobseeker's perspective: pairs with high historical transition rates (at least 20%), median market wages within 5%, and at least 30 unemployed jobseekers in each market. The tightness is measured in the middle of my sample, in March 2021; the results are not sensitive to using other dates. Even among these "close" markets, the median absolute difference in log tightness is 0.84 log points: at two markets each with 100 unemployed jobseekers (close to the sample median), this is the gap between roughly 65 vacancies in the slacker market and 150 in the tighter one.[16]

Panel (b) of Figure 3 examines the direction of clicks relative to the jobseeker's initial market, the last occupation × residence commuting zone. While the initial market is the single market absorbing the largest share of clicks, the majority of search activity is directed elsewhere, and 34% of clicks are directed toward markets that are slacker than the jobseeker's initial market at the time of the click. Combined with the evidence that tightness predicts job finding (Figure 1b), these patterns motivate the analysis of whether caseworker advice can redirect search toward tighter markets and improve employment outcomes.

5 Methodology

The goal of this paper is to estimate the causal effect of redirecting job search toward tight markets on employment outcomes. Naïve comparisons of jobseekers searching in different markets are biased because jobseekers self-select into their search strategies. Even conditional on last occupation, geographical location and further observational characteristics, those targeting slack markets likely differ from those targeting tight markets in unobservable ways, such as skills, preferences, or constraints. To address this selection, I exploit quasi-random assignment of jobseekers to caseworkers who systematically differ in their tendency to redirect clients toward tighter markets. As described in Section 3, caseworkers in Switzerland are assigned within local employment offices based on administrative factors such as caseload balancing and availability at registration but independent of individual characteristics such as employment prospects (Arni & Schiprowski, 2019).

The intuition is the following. Consider two jobseekers with the same last occupation, living in the same commuting zone, registering at the same local office in the same quarter. One is assigned to a caseworker with a tendency of redirecting clients toward tighter markets; the other is not. Since caseworker assignment is as good as random conditional on office, any systematic differences in their subsequent search direction and job finding can be attributed to the caseworker.[17]

[16]Appendix Figure C.5 shows that relaxing the wage restriction raises the median difference to 0.95 log points (a 2.6× ratio), while dropping the transition restriction yields 1.14 log points (a 3.1× ratio).

[17]Since caseworkers act on many dimensions beyond search direction advice (application quality, train-

For the subsample of jobseekers who use the public job portal, hereafter the *clickers*, I observe T_i, the spell-level search tightness index defined in Section 4: the mean log tightness of markets clicked by jobseeker i in the first three months of unemployment. I decompose this index into an additive model with three fixed effects:

$$T_i = \mu_{c(i)} + \alpha_{o(i),\mathrm{yr}(i)} + \phi_{\mathrm{occ}(i)\times\mathrm{cz}(i)\times q(i)} + \varepsilon_i. \tag{1}$$

The object of interest is the caseworker fixed effect $\mu_{c(i)}$, which captures the systematic component of where caseworker $c(i)$'s clients direct their clicks, net of the other terms. I control for two sets of fixed effects $\alpha_{o(i),\mathrm{yr}(i)}$ and $\phi_{\mathrm{occ}(i)\times\mathrm{cz}(i)\times q(i)}$; ε_i is the residual.

The first control, the *unemployment office* $\times$ *year* fixed effect $\alpha_{o(i),\mathrm{yr}(i)}$, captures the conditional random assignment mechanism following Arni & Schiprowski (2019). Within a given office and year, jobseekers are quasi-randomly assigned to one of the available caseworkers. This fixed effect absorbs all office-level variation in caseworker pools, the set of potential caseworkers a jobseeker could have been assigned to.[18] To address the concern that some offices may assign clients not only based on caseworker workload but also on industry or last occupation (Behncke et al., 2010), I show robustness to further interacting this fixed effect with 2-digit last industry and with 2-digit last occupation of the jobseeker.

The second control, the *last occupation* $\times$ *commuting zone* $\times$ *registration quarter* fixed effect $\phi_{\mathrm{occ}(i)\times\mathrm{cz}(i)\times q(i)}$, controls for the feasible choice set of markets each jobseeker can plausibly target. Not all occupational and geographical transitions are equally feasible and desirable. Last occupation and commuting zone are therefore a proxy for the jobseeker's effective choice set: the markets she can plausibly target. I further interact last occupation and commuting zone with the quarter of registration because the choice set and its tightness do not stay fixed over time. Without this fixed effect, $\mu_{c(i)}$ would absorb compositional differences across caseworkers: a caseworker whose clients, due to the finite sample of clients quasi-randomly assigned to caseworkers, disproportionately come from occupation-region cohorts with structurally tighter feasibility sets, or a caseworker with an excess of clients entering unemployment during a boom, would mechanically look more redirection-oriented purely because of who was assigned to her, not because of her advice. I will show robustness to a specification without this fixed effect.

Under quasi-random assignment within office-year and conditional on the market choice-set fixed effect, the caseworker is the only systematic dimension along which observa-

ing referrals, sanction enforcement), the reduced-form effect of the tightness tendency may partly capture correlated caseworker behaviors rather than search redirection per se. Section 7 addresses this concern directly.

[18] Because caseworkers almost always work in a single office, the caseworker fixed effects μ_c are not all separately identified once $\alpha_{o,\mathrm{yr}}$ is included. I therefore omit one caseworker per office; the remaining μ_c are measured relative to that office-specific reference caseworker.

tionally equivalent jobseekers differ. Differences in μ_c across caseworkers therefore reflect caseworker behavior rather than selection of jobseekers.

Estimating $\mu_{c(i)}$ directly and using the estimate for jobseeker i would mechanically correlate the measure with ε_i: own behavior would help define the estimate of her own caseworker. I employ a leave-one-out estimator removing this mechanical correlation by excluding spell i's observation when computing the prediction. As a stronger restriction, Section 6.3.1 reports a cluster leave-out following Frandsen et al. (Forthcoming) that leaves out more than just the focal jobseeker, addressing concerns that outcomes across clients sharing the same last occupation, or registering in the same quarter, might be correlated.

Following Goldsmith-Pinkham et al. (2025), who synthesise the literature on examiner and judge leniency designs, I construct the leave-one-out caseworker tightness tendency using the formula from Kolesár (2013):

$$Z_{-i,c(i)} = \hat{E}_{-i}[T_i \mid \mu, \alpha, \phi] - \hat{E}_{-i}[T_i \mid \alpha, \phi]\,, \tag{2}$$

where $\hat{E}_{-i}[\cdot]$ denotes leave-one-out prediction from a regression excluding spell i. The first term predicts the search tightness for spell i given the caseworker identity and the two sets of control fixed effects. The second term is a normalisation: subtracting the controls-only prediction recentres $Z_{-i,c(i)}$ within each office-year and market choice-set cell, leaving only the caseworker's marginal contribution.[19]

I restrict the sample to caseworkers with at least five clicking clients to ensure reliability of the leave-one-out averages. Section 6.3.1 verifies that the results are stable when the click-count threshold is set to other values.

5.1 Regression equation

The main estimating equation is:

$$Y_i = \beta \cdot Z_{-i,c(i)} + \alpha_{o(i),\mathrm{yr}(i)} + \phi_{\mathrm{occ}(i)\times\mathrm{cz}(i)\times q(i)} + X_i'\gamma + u_i, \tag{3}$$

where Y_i is the outcome for jobseeker spell i, in the baseline a binary indicator of whether the jobseeker found a job within 180 days after registration, the other sections look at

[19]Computing exact leave-one-out predictions requires either re-estimating the regression once per spell, excluding that spell each time, or using a hat-matrix identity to recover leave-one-out predictions from a single full-sample regression. I use the latter approach, which is substantially faster. However, computing individual leverages requires inverting the design matrix, which is infeasible with the high-dimensional fixed effects in my specification. I address this using the Johnson–Lindenstrauss Approximation (JLA) proposed by Kline et al. (2020), which uses random projections to efficiently approximate the leverage values. Technical details are provided in Appendix B.

different horizons and other outcomes. $Z_{-i,c(i)}$ is the leave-one-out caseworker tightness tendency defined in Equation (2). The specification includes the same fixed effects as in Equation (1).

In most specifications, I further control for a range of individual characteristics, X_i, to increase precision. These include gender, age, number of dependent children, a dummy for Swiss nationality, education, tenure in the last occupation, employment and log earnings in the two years before registration, potential benefit duration, detailed last occupation fixed effects (ISCO 4-digit), and municipality of residence fixed effects. I further include employment office $\times$ weekday of first meeting fixed effects to account for potential confounders from caseworkers performing differently across the week combined with endogenous selection into meeting days. Standard errors are clustered at the caseworker level.

The caseworker tendency $Z_{-i,c(i)}$ is computed from clients who click on the job portal, but the measure can be assigned to all clients of that caseworker. This allows me to estimate job-finding effects for the full population of registered unemployed, not just platform users, increasing statistical power. As a robustness check, I show that the results hold when restricting the sample to non-clickers only, ruling out that the effects are specific to platform behavior (Section 6.3.1).

I estimate the reduced-form effect of caseworker tendency on jobseeker outcomes rather than using the caseworker tendency as an instrument for the jobseeker's own search tightness. There are three reasons for this choice. First, the reduced form allows me to estimate effects for all jobseekers assigned to a given caseworker, not only those who click on the job portal.[20] Second, the reduced form avoids the monotonicity assumptions typically required in instrumental variables frameworks (Arni & Schiprowski, 2019; Humlum et al., 2025). Third and most importantly, the reduced-form estimand is itself policy-relevant: it captures how job finding differs when a jobseeker is assigned to a caseworker with a high versus low tendency to redirect search toward tight markets. This corresponds directly to the policy question of what would happen if caseworkers were systematically trained to provide tightness information, or if job platforms displayed market conditions to jobseekers. To ease interpretation, I report magnitudes as the percentage change in the predicted outcome when moving from a 10th- to 90th-percentile caseworker tendency, holding all other variables at their mean effects.

[20] This parallels the leave-one-out logic: for a clicker, the tendency is computed from all *other* clients' clicks, excluding their own; for a non-clicker, the tendency is computed from all clients' clicks, since there is no own behavior to exclude. In both cases, the measure is free of mechanical correlation with the individual's own search choices.

6 Results

This section presents the empirical results. I first check that the caseworker tendency is balanced across predetermined jobseeker characteristics, then show that it redirects clients' own search toward tighter markets, and turn to the central question: whether this redirection translates into faster reemployment. I then examine heterogeneity by initial market tightness, the geographic and occupational margins of mobility, and effects on job quality.

6.1 Balance in covariates

Self-selection into search markets means jobseekers' predetermined characteristics should predict the tightness of the markets they click on. Appendix Table C.7 quantifies this in Panel A and compares it to the relationship between those same characteristics and the caseworker leave-one-out tightness tendency in Panel B. Each row is a separate regression of one predetermined characteristic on the respective tightness measure, conditioning on the office $\times$ year and initial occupation $\times$ commuting zone $\times$ quarter fixed effects as in the main specification from Equation 3.

The analysis documents substantial sorting in own click behavior. Ten of thirteen characteristics show statistically significant relationships after a multiple-testing correction.[21] When looking at the relationship between the characteristics and the caseworker leave-one-out tightness tendency the magnitudes collapse. Eleven of the thirteen characteristics are no longer distinguishable from zero. The five largest Panel A gradients shrink to small fractions of their original size.[22]

Two imbalances remain statistically detectable after the multiple-testing correction. The female gradient is -1.7%, with an adjusted p-value of 0.003, and the university-education gradient is 2.7%, with an adjusted p-value of 0.037. Both magnitudes are economically small. Furthermore, the baseline specification controls for a wide range of personal characteristics and I show that adding those controls does not move the estimates.

A further concern is whether the click-based measure of search direction has the same interpretation across caseworker types. Clicks are a proxy for search effort and direction, but caseworkers could in principle shape how their clients use the platform. If clicks then carry different informational content depending on the assigned caseworker, the tightness

[21] I report Romano-Wolf step-down adjusted p-values across the thirteen covariates, computed via a cluster wild bootstrap on the caseworker level with Rademacher weights and B = 1,999 draws.

[22] For instance, the p10-p90 effect on the likelihood of being a woman drops from -34% to -1.7%. For being a woman with children, it drops from -45% to -1.9%. For holding a university degree, from 179% to 2.7%. For having only primary education, from -49% to -0.2%. For being over fifty, from -15% to -1.2%.

of clicked markets would mix shifts in where jobseekers search with shifts in how they search. Appendix Table C.3 repeats the validation exercise from Figure 1a, regressing an indicator for finding a job in a market on the number of clicks placed there, and interacts the click count with the caseworker's leave-one-out tendency. The interaction is small and insignificant: clicks predict where a jobseeker ultimately finds a job equally well regardless of caseworker type.

6.2 Effects on job search

I first show results from a spell-level regression testing whether a caseworker's leave-one-out tendency predicts the average tightness of a jobseeker's own clicked markets over the first three months of unemployment, mimicking a first-stage regression in a typical instrumental variables setting. The regression uses the same specification as Equation (3), with the jobseeker's own search tightness index as the outcome. Figure 4, Panel (a), presents the results. A 10% higher caseworker leave-one-out score is associated with approximately 0.8% higher tightness of the jobseeker's own clicked markets, significant at all conventional levels (Column (1)). The effect is robust to the full set of individual controls (Column (2)). Column (3) shows no effect on click intensity (clicks per day in the first 90 days), indicating that caseworkers affect the direction of search rather than its volume. This parallels Belot et al. (2026), who find that recommending broader occupational search expanded the occupational scope of job searches but had no effect on their quantity or geographic range.

The spell-level regression establishes that caseworkers affect search targeting at the aggregate level. To shed more light on the mechanism, I combine the exact timing of caseworker meetings with click timestamps. The identifying idea is that if caseworker advice is the channel, effects on click tightness should materialise around the time of meetings, not before. I focus on the first meeting with the assigned caseworker, which usually takes place in the first two weeks after registration.

The main challenge is that because the meeting occurs early in the spell, few clicks are observed beforehand, limiting statistical power in the estimation of what happens before the meeting. Appendix Figure C.6 plots the average number of clicks in the days around the first meeting. Click activity rises gradually in the days leading up to the meeting, spikes on the day of the meeting, and stabilises at a higher level than before, suggesting that meetings have some effect on search intensity, consistent with Schiprowski et al. (2024). I first test whether this increase in activity differs by caseworker type. Appendix Table C.8, Column (1), shows that the overall number of clicks does not differentially increase for jobseekers assigned to high- versus low-redirection caseworkers, confirming once more that caseworkers influence the direction rather than the intensity of search.

I then use this increased activity around the meeting to estimate an event study on a narrow time horizon, using 168-hour (seven-day) bins around the timestamp of the first meeting. I estimate the regression at the click level, so each click contributes one observation. The specification controls for the hour-of-day and weekday of the click in addition to spell, calendar-month, and month-of-spell fixed effects.[23]

Figure 4, Panel (b), shows the results. Pre-meeting coefficients are indistinguishable from zero. After the meeting, the interaction turns positive and becomes significant at the 5%-level from the second week onward. A difference-in-differences estimator pooling all clicks in the pre- and post-meeting periods yields a significant effect of similar magnitude to the spell-level estimate in Figure 4, Panel (a) (Appendix Table C.8, Column (2)).

Taken together, the event study confirms two facts. The flat pre-trend establishes that the endogenous regressor itself, the tightness of the search, was balanced across jobseekers before the caseworker had a chance to intervene. The divergence after the meeting pins down the mechanism: caseworkers redirect search toward tighter markets through information transmitted during counselling. The effect targets the direction of search rather than its intensity, persists for weeks after the initial interaction, and remains significant when aggregating clicks over the first three months of the spell.

6.3 Effects on job finding

Table 2 presents the baseline effects of caseworker redirection on job finding within six months. Column (1) reports the specification without individual controls, including only unemployment office × year and last occupation × commuting zone × quarter fixed effects. This represents the minimal set of controls required for the research design: caseworker assignment is as good as random conditional on office and occupation, and the occupation-region-quarter interaction absorbs differences in choice sets. The coefficient on the caseworker leave-one-out measure is positive and statistically significant: jobseekers assigned to caseworkers with higher redirection tendencies are more likely to find a job within six months. Moving from a caseworker at the 10th percentile to one at the 90th percentile of the leave-one-out measure increases the probability of job finding by 2.7% over the baseline. Column (2) adds the full set of controls detailed in Section 5,

[23]The specification is:

$$Y_{i\tau} = \sum_{k\neq -1} \beta_k \cdot \mathbf{1}[b(\tau) = k] + \sum_{k\neq -1} \gamma_k \cdot \mathbf{1}[b(\tau) = k] \times Z_{-i,c(i)} + \alpha_i + \delta_{b(\tau)} + X'_{i\tau}\theta + \varepsilon_{i\tau} \tag{4}$$

where τ indexes a click placed by jobseeker i, $Y_{i\tau}$ is the log tightness of the clicked market, $b(\tau)$ is the seven-day bin containing click τ, and $Z_{-i,c(i)}$ is the leave-one-out caseworker tightness measure. α_i are jobseeker fixed effects, δ_b are seven-day-bin fixed effects, and $X_{i\tau}$ includes calendar month, month-of-spell, weekday of the click, and hour-of-day of the click. The omitted bin spans days -7 to -1. The sample is restricted to jobseekers with at least three clicks in both the pre- and post-meeting window.

Figure 4: Caseworker effects on job search direction: spell-level and meeting-level evidence

	(1)	(2)	(3)
Dependent Var.:	Own clicks tightness	Own clicks tightness	Clicks per day
Leave-one-out tightness	0.0906***	0.0833***	0.0107
	(0.0217)	(0.0211)	(0.0358)
Unemployment office x Year FE	X	X	X
Last occ. x commuting zone x quarter FE	X	X	X
Jobseeker controls		X	X
Family	OLS	OLS	Poisson
N caseworkers	2,460	2,460	2,460
Observations	101,514	101,514	101,514
R2	0.71233	0.73788	–
Mean of dependent var.	-0.57702	-0.57702	0.50365

(a) Spell-level estimates

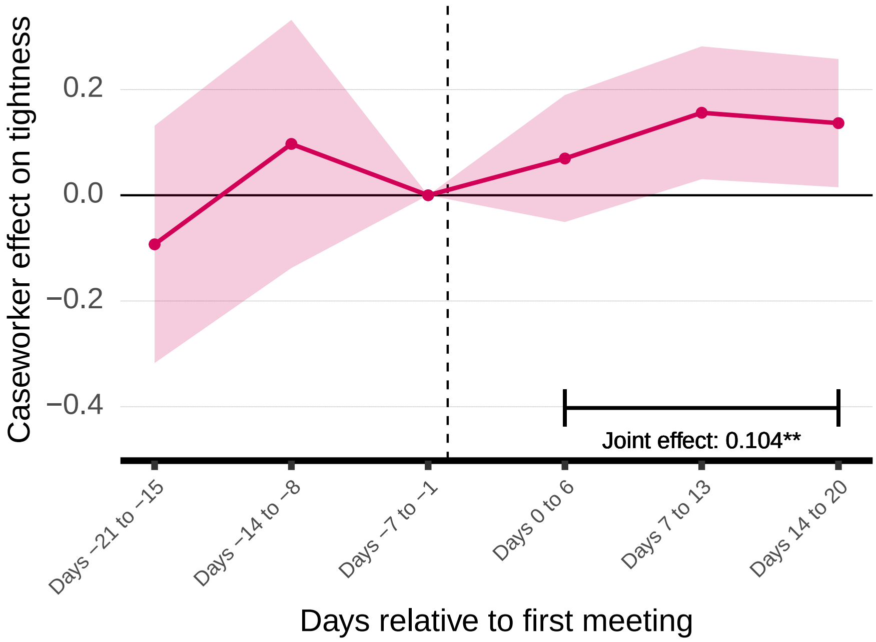

(b) Estimates around first caseworker meeting

Notes: Panel (a) reports estimates from the specification in Equation (3), with the jobseeker's own search tightness index as the outcome in Columns (1) and (2) and the number of clicks per day of unemployment in the first three months in Column (3). Standard errors are clustered at the caseworker level. Panel (b) presents event-study estimates of the interaction between the leave-one-out measure and seven-day bins around the start of the first meeting with the assigned caseworker, using the specification from Equation (4), estimated at the click level. The underbrace indicates the average post-meeting effect from a difference-in-differences specification pooling all post-meeting bins (see Appendix Table C.8, Column (2)). Standard errors are clustered at the spell level. Sample restricted to jobseekers with at least 5 clicks in the first 3 months. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 2: Effect of caseworker-induced search tightness on job finding

	Find job within 6 months				
	(1)	(2)	(3)	(4)	(5)
Dependent Var.:	In any market	In any market	In tighter market	In origin market	In slacker market
Leave-one-out tightness	0.0323***	0.0311***	0.0155***	0.0142**	0.0003
	(0.0096)	(0.0090)	(0.0052)	(0.0063)	(0.0052)
Unemployment office x Year FE	X	X	X	X	X
Last occ. x commuting zone x quarter FE	X	X	X	X	X
Jobseeker controls		X	X	X	X
p10-p90 magnitude	+2.8%	+2.7%	+3.8%	+4.2%	+0.1%
N caseworkers	2,460	2,460	2,460	2,460	2,460
Observations	366,893	366,893	366,893	366,893	366,893
R2	0.08196	0.14161	0.09811	0.09310	0.13211
Mean of dependent var.	0.36554	0.36554	0.13027	0.10688	0.12291

Notes: The table reports linear probability model estimates of the caseworkers tendency to redirect search toward tight markets as described in Equation (3). The dependent variable indicates whether the jobseeker finds employment within 6 months: Columns (1) and (2) show the effect on unemployment exit to any job. Column (3) shows as an outcome only unemployment exits to jobs in markets that are, at deregistration date, tighter than the jobseeker's initial market. Column (4) shows the likelihood of exits to the initial market and Column (5) to slacker markets. Jobseeker controls include gender, age, children, nationality, education, tenure in last occupation, employment and earning in last two years, potential benefit duration, employment-office x weekday of first meeting, municipality FE. The table also reports the effect of going from the 10th to the 90th percentile in the caseworker leave-one-out tendency, expressed as a percentage of the dependent variable. Standard errors are clustered at the caseworker level. Significance levels: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

including spell start month, detailed last occupation, weekday of the first meeting times unemployment office, as well as controls for the employment history and demographics. The estimated effect remains unchanged in magnitude, confirming that the caseworker allocation is orthogonal to any observed personal characteristics that influence job finding. Nevertheless, including controls serves to increase statistical precision, as reflected in the smaller standard errors.

Columns (3)–(5) distinguish where jobseekers find employment. Outcomes are constructed as indicator variables for different types of reemployment. For each regression, the outcome takes the value one if the jobseeker finds a job within six months in the specified type of market, and zero otherwise. Column (3) equals one only if the job is found in a market that is tighter than the jobseeker's initial market, where the initial market is defined as the combination of their last occupation and commuting zone of residence. Column (4) equals one only if the job is found in the initial market itself, while Column (5) equals one only if the job is found in a market that is slacker than the initial market. The comparisons are based on tightness values measured at the unemployment exit date.

The results show a significant effect on job finding in tighter markets (Column 3). Moving from a p10 to a p90 caseworker increases the probability of finding a job in a tighter market by about 3.8% . This is exactly the pattern predicted by the redirection mechanism: assignment to a caseworker with a high tightness tendency raises the likelihood that clients exit unemployment in tighter markets. Column (4) also shows a significant effect on finding a job in the initial market. Here, the effect size is about 4.2%. This suggests that the treatment effect does not exclusively operate through shifting away from one's original occupation or location. Rather, it also seems to reinforce job finding in the initial market when that market is already favorable. Section 6.4 investigates this channel in more detail. By contrast, Column (5) shows no effect on job finding in slacker markets. While not a formal placebo test, this result provides an important falsification check. If the caseworker effect simply reflected across-the-board changes—such as better CV writing, stronger motivation, greater search intensity, or increased stringency such as higher sanction threats—one would expect to see improved outcomes in slacker markets as well. The absence of an effect there suggests that the treatment operates specifically through redirection toward tight markets. Section 7 adds to this argument by directly testing alternative mechanisms such as search intensity, support behaviors, and sanction rates.

Up to this point, the analysis has focused on the benchmark horizon of six months. Appendix Figure C.7a extends the perspective to other horizons by tracing cumulative job-finding probabilities over the first year of the spell. Panel (a) reports effects on the survival function, i.e. the probability of still being unemployed. The two curves are indistinguishable in the first two months: the effect after 30 days is 1% and not statistically significant. The gap then widens steadily, becoming significant at 120 days and peaking around 150 days at 2.8%. At the six-month benchmark, the effect is 2.7%. It declines somewhat at longer horizons but remains economically meaningful and highly significant at one year (1.9%). The persistence of the effect suggests that caseworker redirection does not merely accelerate job finding by a few weeks; it leads to reemployment that would not have occurred within the first year in the absence of the treatment. Panel (b) provides a period by period hazard rate analysis using period-specific complementary log-log models, which are more robust to functional form assumptions.[24] The results suggest that the effect is not driven by an elevated exit rate in a specific period: rather, throughout the unemployment spell, jobseekers assigned to more redirecting caseworkers are consistently more likely to find employment.

6.3.1 Robustness

The identification strategy relies on quasi-random assignment of jobseekers to caseworkers within employment offices. As discussed in Section 3, assignment methods vary across

[24]Note that this analysis is primarily descriptive as the hazard in period t is estimated conditional on still being unemployed in $t-1$ which itself is an outcome.

cantons: seven cantons use procedures unaffected by caseworker performance, namely distribution strictly proportional to the caseworker's hours worked or software-based random allocation, while others balance caseloads or assign by management decision. The other assignment methods, caseload balancing and assignment by management could introduce endogeneity if more efficient caseworkers handle cases faster and therefore receive more new clients or if the manager—despite only knowing administrative information I control for in the regression— can assign jobseekers in a way that is linked to their employment potential. Appendix Table C.5 restricts the sample to the seven cantons using caseload balance strictly proportional to contractual hours and the software-based random allocation. The coefficients and magnitudes are very comparable to the baseline and, if anything, slightly larger.

A related concern about the assignment is that some employment offices assign caseworkers based on occupation or industry sector (Behncke et al., 2010). If so, the caseworker fixed effect could partly capture occupation- or industry-specific trends. To address this, Columns (1) and (2) of Appendix Table C.11 interact the office $\times$ year fixed effect with the 2-digit NOGA industry of the last job in Column (1) and with the 2-digit ISCO last occupation in Column (2). The magnitudes are 2.4% for both versions respectively, both slightly smaller than the baseline magnitude of 2.7% but very similar overall. Together with the former check, this suggests that the results are not driven by a feature of the caseworker assignment process.

Furthermore, Column (3) of Appendix Table C.11 drops the last occupation $\times$ commuting zone $\times$ quarter FE entirely from the leave-one-out demeaning and the main regression, retaining only the office $\times$ year fixed effect. The magnitude is 3.3%, if anything, slightly larger than the baseline magnitude, suggesting that the results do not crucially rely on the last occupation $\times$ commuting zone $\times$ quarter FE that holds constant the choice set of markets and their tightness.

Columns (4)–(7) of Appendix Table C.11 test robustness to different measures of market tightness. Column (4) uses 4-digit ISCO codes instead of 3-digit for finer occupational granularity, with a magnitude of 1.7%. Column (5) uses 101 labor market regions, the Swiss *Arbeitsmarktregion*, instead of 16 commuting zones, with a magnitude of 1.5%. Column (6) allows the denominator of tightness to reflect actual search behavior using data from the platform, defining tightness as the number of clickers per advertisement, finding a magnitude of 1.6%. The finer market definitions likely introduce more measurement error in tightness, attenuating the estimates. Results remain significant across all alternative definitions.

Column (7) decomposes the caseworker tendency into separate leave-one-out measures for log vacancies and log unemployment. Caseworkers who steer clients toward markets with more vacancies increase job finding by 2.8%, while those who steer toward markets with more competing jobseekers reduce it by -1.9%. Both coefficients are individually significant. There is the caveat though that this is not a conclusive decomposition of the tightness effect: for a given caseworker who has a certain tendency to redirect jobseekers to tight markets, one cannot test whether the effect comes from the vacancy or competition component of the tightness; the identifying variation here comes from differences in how individual caseworkers weight the two components. Nevertheless, the fact that both coefficients are significant and carry the expected signs suggests that both the opportunity side, namely more vacancies, and the competition side, namely fewer competing searchers, contribute to the effect. This complements Altmann et al. (2026), who provide jobseekers with vacancy information but do not account for competition from other searchers; the results here suggest that the competition margin also matters for the effectiveness of search redirection.

A further concern is that the caseworker's effect on clicking could reflect a platform-specific mechanism, for instance caseworkers teaching jobseekers to use job-room.ch more effectively, rather than a genuine change in search direction. In Appendix Table C.12, I restrict the sample to jobseekers with no prior clicks on job-room.ch and estimate the effect of their caseworker's tendency, measured from the caseworker's clicking clients, on job finding. The effect remains positive and statistically significant, with a magnitude of 3.6%, not significantly different from the baseline effect, indicating that caseworker redirection influences employment outcomes through channels beyond platform-specific search behavior.[25]

If many jobseekers arrive to a caseworker at the same time, this could induce a correlation between the tightness of their search and their job-finding rates through common shocks. While the interaction with the quarter of registration in the fixed effects attenuates this problem, I further show robustness by computing a leave-out measure that excludes not just the focal jobseeker's own search but the search of all clients who registered in the same quarter, following the cluster-leave-out idea by Frandsen et al. (Forthcoming). Appendix Table C.13 shows that results remain robust. A related concern applies within occupation: although the occupation $\times$ commuting zone $\times$ quarter fixed effect absorbs market-level tightness shocks, clients of the same caseworker who share the same occupation search in overlapping sets of markets, so the coefficient might be picking up a correlation between

[25] Yap (2024) recently shows that when a constructed leave-out measure such as $Z_{-i,c(i)}$ is used as a regressor, standard cluster-robust standard errors do not fully account for the variability introduced by constructing the measure and can therefore be too small. This concern bites only for observations that enter both the construction of $Z_{-i,c(i)}$ and the outcome regression. In the baseline the overlap is already partial: only clickers, roughly 28% of the regression sample, are used to construct the measure, while the remaining 72% of observations are non-clickers, whose own data does not enter $Z_{-i,c(i)}$. Restricting to non-clickers eliminates the overlap entirely. The fact that the magnitude in Appendix Table C.12 very similar to the baseline and also the null-effect on exit to slacker market persists indicates that the result is not an artifact of within-sample dependence between the constructed regressor and the outcome equation.

occupation-specific search scope and job prospects. I therefore also show that the results are robust to leaving out all clients with the same occupation in Appendix Table C.14.

The estimator requires two sample restrictions: a minimum number of clicks per jobseeker to measure search targeting, and a minimum number of clicking clients per caseworker to estimate the leave-one-out tendency. The baseline sets both thresholds to five. Appendix Table C.15 varies both thresholds across a grid of values: 1, 5, 10, and 30. Results are stable across most combinations, with magnitudes ranging from 1.73% to 3.11%. Only at 30 clickers per caseworker does the effect lose significance, consistent with reduced statistical power and less variation in treatment; this threshold reduces the sample from 366,893 to 291,543 observations and the number of caseworkers from 2,460 to 1,573.

The baseline measures the search tightness index over the first three months of unemployment. Appendix Table C.6 varies this window across 30, 60, 90, 120, and 180 days. Magnitudes are between 2.0% and 2.8% for windows of 60 days and longer; the 30-day window yields a smaller, insignificant estimate as there is much less data to estimate the search tightness measure.

In Appendix Table C.16, I re-estimate the exit hazard using a complementary log-log model with minimal fixed effects. The pattern mirrors the baseline, with a magnitude of 2.7% and effects concentrated in tighter and origin markets and no effect on slacker markets.

A natural concern is that the mid-2020 to end-of-2022 sample period spans parts of the COVID-19 recovery. Appendix Table C.17 re-estimates the baseline on two subsamples: jobseekers registering before 17 February 2022 — the date the Swiss federal government lifted the last broad COVID-19 restrictions — and those registering on or after that date. The baseline magnitudes are 2.6% before and 2.9% after the cutoff, very close to the pooled baseline. Decomposed by destination, the post-cutoff sample shows slightly larger effects on job finding in the initial market and slightly smaller effects on transitions to tighter markets, but these differences between the samples are not statistically significant and are within the noise expected from splitting the sample in two.

Lastly, the tighter, origin, and slacker indicators require both the found and the initial markets to be observed. The $\sim$1.5% of job finders whose found-market cell remains undefined even after the NOGA-2 imputation described in Section 4 enter the baseline regression as zero on all three indicators. Appendix Table C.18 reports a specification that keeps them as a separate *unobserved* outcome rather than imputing them. The coefficients on the three observed market types and on the any-market outcome are very similar to the baseline.

6.4 Heterogeneity by baseline tightness

Figure 5 explores heterogeneity in the effect of caseworker redirection by the tightness of the jobseeker's *initial* market, defined as the commuting zone of residence × last occupation cell. Jobseekers are split into terciles of initial market tightness.[26] Furthermore, the analysis does not only group jobseekers by their initial market but also tracks where reemployment occurs: outcomes distinguish between exits to a job in the same market as the initial one, exits to tighter markets relative to the initial one, or slacker markets. The combination of information on the initial market and on the market of reemployment allows for a detailed examination of the mechanism behind the effect shown in the previous section. If the treatment worked only as a mobility intervention—pushing jobseekers out of slack markets—effects should appear exclusively for those with poor initial prospects. However, there is increasing evidence that search might also be too broad (van der Klaauw & Vethaak, 2022; Caliendo et al., 2026). If this is the case, then redirecting search to tight markets could make jobseekers in favorable initial markets focus more on those markets, narrowing the search scope and increasing job finding rates.

The graph reports the estimated percentage increase in job finding within six months when moving from the 10th to the 90th percentile in the caseworker leave-one-out measure, the underlying regression estimates are shown in Appendix Table C.19, Columns (1)-(4). Before examining the specific channels, Column (1) of Appendix Table C.19 establishes that the overall job-finding effect is positive and similar across all terciles of initial market tightness. If anything, the effect is slightly larger for those in medium or tight initial markets, though the difference is not statistically significant. Figure 5 decomposes this effect by the market in which reemployment occurs.

The graph shows that the effect is not exclusively driven through shifting away from one's original occupation or location. Instead, the channel indeed depends on the conditions in the initial market. For jobseekers starting in slack markets, the treatment operates through an effect on job finding in tighter markets: assignment to a high-redirection caseworker increases the probability of finding a job in a tighter market within six months by about 4%. The probability of reemployment in the initial market or even slacker markets is not affected. This effect is in line with the experiment by Belot et al. (2025), who exclusively target jobseekers in markets with poor prospects, suggesting they broaden their occupational search, and find a 5.2% effect on employment after 18 months. For those becoming unemployed in medium-tightness markets, the effect also works through increased exit rates to tighter markets (about 5%). But in addition, the probability of finding a job in the initial market rises by about 5%. For jobseekers in tight initial markets, the effect is primarily driven by faster exits in the initial market itself, with

[26]The assignment to terciles uses the tightness measured at the unemployment registration date.

Figure 5: Heterogeneous effects of caseworker redirection on job finding by initial market tightness

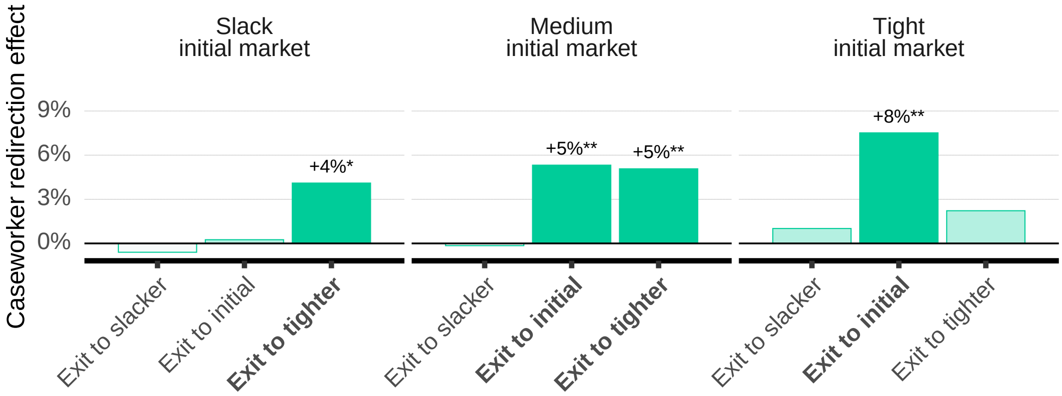

Notes: The figure plots effects of the caseworker leave-one-out tightness tendency on six-month job finding from regressions following Equation (3) with the leave-one-out measure interacted with initial market (ISCO 3-digit × commuting zone) tightness terciles. The graph shows three different outcomes: six-month unemployment exit to a job in the initial market, to a market that is tighter than the initial market and to a market that is slacker than the initial market. The bars show the effect size of going from the 10th to the 90th percentile in the caseworker leave-one-out distribution as a percentage of the baseline probability. The underlying regressions are reported in Appendix Table C.19. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

an 8% effect. Notably, for this group there is a negative effect on job finding in slacker markets, consistent with caseworkers preventing deviation into less favorable markets.

Taken together, this pattern suggests that caseworker advice operates through two complementary channels: redirecting search toward tighter markets and reinforcing search in already favorable ones. A potential implication is that jobseekers are actually not aware of the different job finding prospects in different markets and might underestimate how good they are in their initial market. This finding resonates with the descriptive evidence in Section 4, which showed that a substantial share of jobseeker clicks are directed toward slacker markets. In that context, the heterogeneity analysis suggests that when the initial market is already tight, caseworkers primarily help to prevent jobseekers from deviating into less favorable markets. This distinguishes the mechanism from occupational mobility treatments, which typically emphasize moving workers out of declining sectors. The results here suggest that occupational mobility interventions are not the right approach in every situation. For many jobseekers, especially those already located in favorable markets, the relevant margin is not moving to a new occupation or region but rather sustaining and reinforcing search in their existing tight market.

As a further piece of evidence, I use the click data to examine directly how caseworkers influence the search process itself. Focusing on the subsample of jobseekers who actively use the job-room portal for their job search—defined as those with at least five clicks in the first three months of their spell—allows me to observe in greater detail how caseworker redirection shapes the targeting of different markets, before employment outcomes are realized. Because this subsample covers roughly 28% of the full regression sample, these estimates are noisier and should be interpreted as more suggestive. For each jobseeker, I split the clicks on vacancies in markets that are tighter than, the same as, or slacker than the initial market — the same three categories I used for the reemployment outcomes — and estimate the effect of the caseworker tendency on each count using a pseudo-Poisson maximum likelihood (PPML) regression with the same specification as in Equation (3). Having established that there is no effect on the number of clicks, I add a PPML offset of log(clicks per day) to gain statistical power by holding the jobseeker's total clicks constant.[27] The results, shown in Appendix Figure C.8, confirm the insights from the job-finding analysis. Moving from a p10 to a p90 caseworker, jobseekers starting in slack initial markets click 5% less in the initial market, deterring continued search in the last occupation and commuting zone where prospects are poor. For jobseekers in medium-tightness initial markets, clicks in slacker markets fall by 3% and clicks in tighter markets rise by 2%, both significant, while the focus on the initial market is essentially unchanged. For those in tight initial markets, clicks in tighter markets rise by 3%, while the other two channels are small and insignificant.

6.5 Geographic versus occupational mobility

The previous analysis distinguishes destinations by whether they are tighter, equal, or slacker than the jobseeker's initial market. This section decomposes mobility along its two constituent dimensions—occupation and geography—to understand whether caseworker redirection operates primarily through occupational switches, geographic moves, or a combination of both. I partition job-finding outcomes into four mutually exclusive categories defined by whether the destination shares the jobseeker's last occupation and whether it lies in the same commuting zone (CZ): same occupation & same CZ, same occupation & different CZ, different occupation & same CZ, and different occupation & different CZ. As in the baseline analyses, the outcome is the probability of finding a job in a given category within six months; jobseekers who do not find any job are coded as zero.

Appendix Figure C.9 presents the effect of going from a 10th- to a 90th-percentile caseworker in the leave-one-out measure for each of the four occupation $\times$ commuting zone categories. It shows the average effect as well as heterogeneous effects by initial market tightness tercile, as used in Section 6.4. Three of the categories show significant average

[27] Appendix Table C.1 confirms that the leave-one-out tendency has no effect on the number of clicks within any initial-market tightness tercile.

effects. First, assignment to a high-redirection caseworker increases the probability of finding a job in the same occupation and same commuting zone by 4.1%. Second, it increases the probability of finding a job in a different occupation *and* a different commuting zone by 7.6%, and third, it increases the probability of finding a job in a different occupation in the same commuting zone, albeit with a smaller effect of 3.0% that is only significant at the 10% level. The remaining category —same occupation with a different commuting zone—shows no effect. The relevant margin thus appears to be joint mobility: caseworker redirection does not primarily operate through occupational switches alone or geographic moves alone, but through the combination of both. There are also, smaller, effects on occupational mobility but none on purely geographic mobility.

Looking at the heterogeneity by initial market conditions adds more details to the patterns by initial market documented in Section 6.4. For jobseekers becoming unemployed in slack markets the effect is almost exclusively driven through moves out of the occupation and out of the commuting zone. For jobseekers becoming unemployed in tight markets the effect is predominantly driven by staying both, in the commuting zone and in the occupation, and additionally—to a lesser extent—by staying in the commuting zone but switching the occupation.

6.6 Effects on job quality

A natural concern is whether faster job finding through redirection toward tighter markets comes at the cost of job quality. If jobseekers already target jobs by trading off job finding, which is positively correlated to tightness, and job value, putting more weight on tightness could come at the cost of job value. Alternatively, if there are dispersions in tightness even between similar value markets—as Section 4.3 suggests looking at the wage as a proxy for value—then jobseekers who get informed about these dispersions, can reap the rewards in terms of job finding without trading off job quality.

Table 3 examines three dimensions of job quality, conditional on finding employment during the sample period. First, I consider employment stability, whether redirected jobseekers are more likely to return to unemployment shortly after reemployment. Second, I examine reemployment earnings and the market wage level of the found job, to assess whether faster exits come at the cost of lower pay. Third, I measure commuting distance to the new job, to test whether redirection toward tighter markets pushes jobseekers into geographically more distant employment. Variable definitions follow Section 4. Because social security earnings data ends in December 2021, the earnings sample is restricted to deregistrations before October 2021.

The results show no effect on job stability or reemployment earnings. Moving from a p10 to a p90 caseworker does not change the probability of returning to unemployment

Table 3: Effect of caseworker redirection on job quality

	Return to unemployment		Earnings		Distance
	(1)	(2)	(3)	(4)	(5)
Dependent Var.:	180 day	360 days	log(Earnings)	log(Market wage)	log(Hours)
Leave-one-out tightness	-0.0062	-0.0072	0.0257	0.0019	0.0098
	(0.0112)	(0.0154)	(0.0313)	(0.0052)	(0.0233)
Leave-one-out tightness x Medium initial market	0.0107	0.0077	-0.0023	0.0102	-0.0442
	(0.0151)	(0.0204)	(0.0442)	(0.0071)	(0.0302)
Leave-one-out tightness x Tight initial market	0.0151	-0.0032	-0.0617	0.0046	-0.0257
	(0.0149)	(0.0207)	(0.0490)	(0.0069)	(0.0301)
Tightness tercile FE	X	X	X	X	X
Unemployment office x Year FE	X	X	X	X	X
Last occ. x commuting zone x quarter FE	X	X	X	X	X
Jobseeker controls	X	X	X	X	X
--	--------	--------	-------------	----------------	----------
Average effect	0.0026	-0.0060	0.0066	0.0069**	-0.0142
	(0.0062)	(0.0088)	(0.0196)	(0.0032)	(0.0124)
Magnitude any market	+0.6%	-0.8%	+0.2%	+0.2%	-0.4%
Magnitude Slack initial market	-1.2%	-0.8%	+0.8%	+0.1%	+0.3%
Magnitude Medium initial market	+1.0%	+0.1%	+0.7%	+0.4%	-1.1%
Magnitude Tight initial market	+2.0%	-1.5%	-1.1%	+0.2%	-0.5%
N caseworkers	2,457	2,430	2,355	2,459	2,460
Observations	194,349	141,099	77,729	187,237	221,373
R2	0.08692	0.12337	0.28544	0.64770	0.22329
Mean of dependent var.	0.14764	0.24329	8.6165	8.7266	-0.89425

Notes: The table reports estimates from the specification in Equation (3), with job quality outcomes as dependent variables, conditional on finding employment. Columns (1)–(2): return to unemployment is an indicator for re-registering within 180 or 360 days of reemployment. Columns (3)–(4): reemployment earnings are log monthly earnings in the new job; market wage is the log median monthly full-time equivalent wage in the found market. Column (5): distance is the log driving time between the municipality of residence and the municipality of the found job. The top panel shows interactions with terciles of initial market tightness; the bottom panel shows the average effect. The table also reports the effect of going from the 10th to the 90th percentile in the caseworker leave-one-out tendency, expressed as a percentage of the dependent variable. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

within 180 or 360 days, and the average effect on log monthly earnings is close to zero and statistically insignificant. For the market wage level of the found job, I find a statistically significant, positive but very small effect: the average coefficient on the leave-one-out tightness is significant at the 1% level, the implied p10-to-p90 caseworker magnitude is around 0.2%. Commuting distance is also unaffected: the average effect is close to zero and statistically insignificant, indicating that while redirection is associated with more commuting zone switches, the effect on the actual distance they need to commute is negligible.

The heterogeneity analysis does not reveal significant differentiation across initial-market terciles for any of the job-quality outcomes. For reemployment earnings, the p10-to-p90 caseworker magnitudes are insignificant across the board—0.8% for jobseekers from slack initial markets, 0.7% for medium, and -1.1% for tight initial markets. Though insignificant, the negative effect on earnings of jobseekers starting in tight markets might be worth noting. However, the corresponding effect on the market wage level, which has a higher sample coverage, is slightly positive, pointing toward measurement error rather than an economically meaningful result. The interaction terms for job stability and commuting distance likewise show no significant heterogeneity.

Taken together, these results suggest that the faster job finding induced by caseworker redirection does not come at the expense of job quality. Redirection seems to steer jobseekers toward genuinely favorable labor markets in terms of job-finding prospects rather than pushing them into low-quality segments where tightness is high because nobody wants to work there.

6.7 Caseworker characteristics

What distinguishes caseworkers who redirect their clients toward tighter markets? Table C.21 in the appendix shows the average value of several caseworker characteristics predicted at the 10th and 90th percentiles of the leave-one-out tightness tendency, and regresses each characteristic on the tendency using the baseline specification. While I do not have any background information on the caseworker in my data, I can infer several dimensions from their interactions with jobseekers. I examine caseload, experience, and the occupational and industry diversity of their client pools.

There is no substantial difference between high- and low-redirection caseworkers in any of the dimensions considered. The only coefficient individually significant at the 5% level is the number of distinct 2-digit industries served: a P90 caseworker serves clients from about one more industry than a P10 caseworker (44.1 versus 43.4), a small difference that could reflect broader industry exposure helping caseworkers accumulate knowledge about opportunities across different segments of the labor market. However, the difference is not significant once I adjust for multiple hypothesis testing (q-value $= 0.131$). The tendency to redirect jobseekers toward tighter markets thus appears largely idiosyncratic to individual caseworkers rather than systematically related to observable characteristics. This is in line with Cederlöf et al. (2025) who find that most observable caseworker characteristics—education, cognitive and non-cognitive ability, local labor market ties, and recruitment background—are unrelated to caseworker value-added across all performance dimensions, with only experience emerging as a significant predictor. As in the teacher value-added literature (Rivkin et al., 2005; Rockoff et al., 2011), what caseworkers do appears to matter far more than who they are. As discussed in Section 3 and Appendix D, caseworkers appear to learn from recent experience: when one client finds a job in a given market, the caseworker's other clients subsequently increase their search there.

7 Interaction with other caseworker dimensions

Search redirection is one of several dimensions along which caseworkers differ. One concern might hence be that the tightness tendency is simply correlated with some caseworker behaviors that actually matter but does not have an effect itself. To test for this, I construct leave-one-out measures across multiple dimensions and examine their variance-covariance structure.

The other dimensions of caseworker behavior measurable in my data are the following. The number of meetings (mean monthly rate in first 90 days), average meeting duration (log hours), and training course participation (pooling Basic courses, language courses, computer courses, and occupation-specific courses, see Gerfin & Lechner (2002) Table A1 for a list). Then, two dimensions related to caseworker stringency: the probability with which the caseworker deems the jobseeker's monthly search effort protocol sufficient[28], and the probability of the caseworker issuing a benefit sanction in the first three months.

For outcomes, I construct value-added measures for: overall job finding, job finding in tighter markets, job finding in slacker markets, job finding in the initial market, all measured at six months after spell start, as in the baseline regressions. I also include the earnings measures conditional on job finding from earlier sections are also included: log reemployment earnings, and reemployment market wage level.

All measures are computed using the leave-one-out methodology described in Section 5, with appropriate fixed effects to absorb office and initial market variation. Correlations are then computed at the unemployment-spell level; because each caseworker measure excludes the focal individual's own outcome, variances and covariances are not inflated by own-observation bias (Kline et al., 2020).

Figure 6 and Appendix Table C.22 present the correlation matrix of leave-one-out caseworker measures. Several patterns emerge. First, many caseworker dimensions are correlated with each other. Meeting rate and meeting duration are negatively correlated ($r = -0.38$), suggesting these are substitutes in how caseworkers allocate their time. Caseworkers who sanction more often also meet more frequently ($r = 0.12$) and more often refer clients to training courses ($r = 0.20$). Sanctions and proof-of-effort acceptance are strongly negatively correlated ($r = -0.64$), as expected given that sanctions are often triggered by insufficient proof of effort.

Second, and crucially for the identification strategy, the tightness redirection tendency is nearly orthogonal to both support ($r \approx 0$ with meetings, duration, and training) and

[28] Arni & Schiprowski (2019) look at the monthly application requirement. This information, as well as the actual number of applications submitted is only available for a subset of cantons. I hence use the probability with which the caseworker deems the jobseeker's monthly search effort protocol sufficient as a measure of stringency in the search effort dimension.

stringency dimensions ($r \approx 0.01$–0.02 with sanctions and proof acceptance). This means that caseworkers who redirect toward tighter markets do not systematically differ in how much they meet with clients, how long those meetings are, whether they refer clients to training, or how strictly they enforce job search requirements. Appendix Table C.23 provides a formal test by regressing each caseworker behavior dimension on the tightness tendency. Tightness is uncorrelated with meeting duration, training referrals, proof-of-effort acceptance, and sanction rate. There is a significant correlation with number of meetings, implying that high-redirection caseworkers meet slightly less frequently; however, the magnitude is small, amounting to a 0.7% reduction in the number of meetings when moving from the 10th to the 90th percentile in the tightness redirection leave-one-out score. Furthermore, since meeting frequency is, if anything, negatively associated with job finding, any resulting bias would work against the main result of a positive effect of redirecting search to tight markets on job finding. For all other dimensions, the coefficients are both statistically insignificant and economically small.

Third, several dimensions are correlated with job-finding value added. Training course participation is also positively associated with job-finding VA ($r = 0.11$), confirming the results from Humlum et al. (2025). More leniency is associated with lower job finding VA, a result similar to Rambjer et al. (2026) who find that stricter monitoring increases job-search intensity, shortens unemployment durations, and raises employment rates. Next, the sanction rate is positively correlated with job-finding VA ($r = 0.27$), consistent with (Lalive et al., 2005; Arni et al., 2013). Between meeting frequency and meeting duration, frequency shows essentially no correlation with overall job finding, while meeting duration shows a negative correlation ($r = -0.10$). As shown in the previous sections, the tightness tendency is positively correlated with job-finding VA ($r = 0.08$).

The key evidence for the tightness mechanism comes from examining how each caseworker dimension correlates with job finding *by destination market type*. The tightness tendency shows a distinctive pattern: it is positively correlated with job finding in tighter markets ($r = 0.09$) and in the initial market ($r = 0.04$), but shows essentially zero correlation with job finding in slacker markets ($r = -0.01$). This is exactly the pattern predicted by the redirection mechanism. In contrast, all other caseworker dimensions that matter for the outcome show *uniform* correlations across market types. The sanction rate is correlated with job finding in tighter markets ($r = 0.16$), slacker markets ($r = 0.15$), and the initial market ($r = 0.10$)—roughly equal in all cases. Training course referrals show a similar pattern: $r = 0.03$ for tighter, $r = 0.07$ for slacker, $r = 0.07$ for initial. If the tightness measure were simply picking up general caseworker quality (motivation, CV help, etc.), I would expect it to show the same uniform pattern as these other dimensions. Instead, only tightness shows the differential pattern consistent with a redirection mechanism.

Figure 6: Correlation matrix of caseworker dimensions

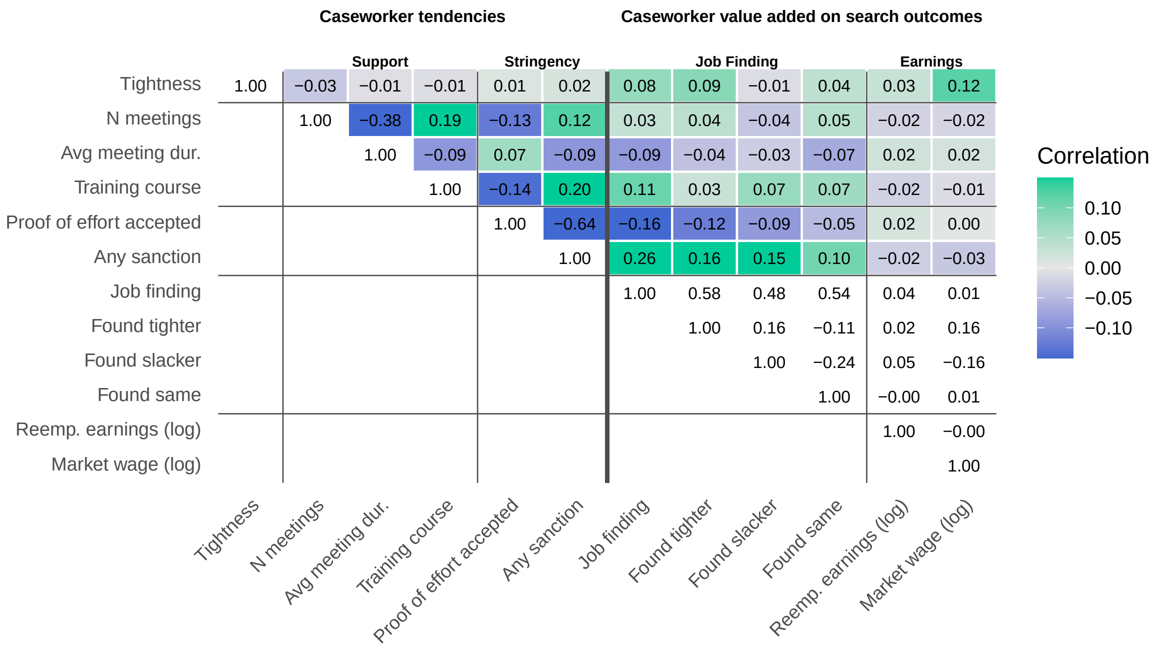

Notes: The figure shows pairwise correlations between leave-one-out caseworker measures. All measures are computed using the methodology described in Section 5. "Tightness" is the main measure of caseworker tendency to redirect toward tight markets. Support measures include number of meetings, average meeting duration, and training course referrals. Stringency measures include proof-of-effort acceptance rate and any sanction. Job-finding VA measures caseworker value added for job finding overall and by destination market type. Wage VA measures value added for reemployment earnings and market wage level conditional on job finding.

8 Conclusion

This paper investigates whether directing jobseekers' search toward tight labor markets improves employment outcomes. Using click-level data from the Swiss public employment service job portal linked to administrative unemployment records, I document that caseworkers differ systematically in how they direct clients' search. Exploiting quasi-random assignment of jobseekers to caseworkers, I find that assignment to a caseworker with a high tendency to redirect toward tight markets increases job finding within six months by 2.7%. Faster job finding does not come at the expense of job quality. A distinctive insight emerges from the heterogeneity analysis: for jobseekers in slack markets, the effect operates through redirection toward tighter markets, but for those starting in tight initial markets, the effect operates through faster exits in the initial market itself. This pattern suggests that occupational mobility interventions are not always the right approach—for jobseekers already in favorable markets, the relevant margin is focusing search rather than expanding it. These results complement recent field experiments targeting jobseekers in poor-prospect occupations with mobility interventions (Belot et al., 2025), while suggesting scope for expanding counseling to those in favorable markets—informing them about their good prospects may sustain search focus and reduce congestion in slack markets.

The findings carry direct policy implications. First, tightness is an effective and simple way to tailor recommendations to the market a jobseeker worked in before unemployment. An attractive feature of tightness is that it uses data that today are widely available to the PES: vacancy counts and jobseeker counts at the market level. Second, my analysis suggests that caseworker counseling can effectively redirect search—a low-cost policy lever already embedded in unemployment insurance systems. However, my method exploits substantial variation in caseworker tendencies. This suggests scope for gains if more caseworkers adopted the behavior of high-redirection ones. Alternatively, tightness information could be provided directly to jobseekers: job portals already have real-time data on vacancies, and the PES also has real-time data on the number of unemployed, and could display market conditions alongside job listings. Together, these recommendations suggest that making labor market tightness information more salient—whether through caseworker training or direct display—could improve the allocation of job search effort across markets.

A natural direction for future research is to test direct provision of tightness information to jobseekers—for instance through a randomized controlled trial on job portals displaying vacancy-to-unemployment indicators, or through targeted caseworker training interventions supported by jobseeker surveys. Such experiments would provide additional causal evidence on the information channel and help disentangle whether the effects documented here operate primarily through information transmission or other aspects such as stressing the impacts of searching in low-prospect markets on job finding.

References

Altmann, S., Glenny, A., Mahlstedt, R., & Sebald, A. (2026). The direct and indirect effects of online job search advice. *Review of Economic Studies*. (Conditionally accepted)

Arni, P., Lalive, R., & Van Ours, J. C. (2013). How effective are unemployment benefit sanctions? looking beyond unemployment exit. *Journal of applied econometrics*, *28*(7), 1153–1178.

Arni, P., & Schiprowski, A. (2019). Job search requirements, effort provision and labor market outcomes. *Journal of Public Economics*, *169*, 65–88.

Bächli, M., Lalive, R., & Pellizzari, M. (2025). Helping jobseekers with recommendations based on skill profiles or past experience: Evidence from a randomized intervention.

Barbanchon, T. L., Hensvik, L., & Rathelot, R. (2023). How can AI improve search and matching? Evidence from 59 million personalized job recommendations.

Behaghel, L., Crépon, B., & Gurgand, M. (2014). Private and public provision of counseling to job seekers: Evidence from a large controlled experiment. *American economic journal: applied economics*, *6*(4), 142–174.

Behaghel, L., Gurgand, M., Dromundo, S., Hazard, Y., & Zuber, T. (2024). The potential of recommender systems for directing job search: a large scale experiment. *Econometrica*. (Conditionally accepted)

Behncke, S., Frölich, M., & Lechner, M. (2010). A caseworker like me–does the similarity between the unemployed and their caseworkers increase job placements? *The Economic Journal*, *120*(549), 1430–1459.

Belot, M., de Koning, B. K., Fouarge, D., Kircher, P., Muller, P., & Phlippen, S. (2025). *Advising job seekers in occupations with poor prospects: A field experiment* (Working paper). IZA.

Belot, M., Kircher, P., & Muller, P. (2018). 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. (2026). Do the long-term unemployed benefit from automated occupational advice during online job search? *The Economic Journal*, *136*(673), 184–206.

Ben Dhia, A., Crépon, B., Mbih, E., Paul-Delvaux, L., Picard, B., & Pons, V. (2022). Can a Website Bring Unemployment Down? Experimental Evidence from France.

Bied, G., Caillou, P., Crépon, B., Gaillac, C., Pérennes, E., & Sebag, M. (2026). A job i like or a job i can get: Designing job recommender systems using field experiments. *arXiv preprint arXiv:2603.21699*.

Caliendo, M., Künn, S., & Mahlstedt, R. (2026). The intended and unintended effects of promoting labor market mobility. *Review of Economics and Statistics*, *108*(2), 421–435.

Cederlöf, J., Söderström, M., & Vikström, J. (2025). The Role of Caseworkers: Job Finding, Job Quality, and Determinants of Value-Added. *Journal of the European Economic Association*. doi: 10.1093/jeea/jvaf055

Christensen, A. N., Harmon, N. A., Settele, S., & Skandalis, D. (2026). *Perceived wage differences and early career job search choices* (mimeo).

Colella, F. (2022). *The Effect of Trade on Skill Requirements: Evidence from Job Postings* (Tech. Rep.). Mimeo.

Dromundo Mokrani, S., & Haramboure, A. (2022). Do job counselors matter? measuring counselors value-added in job search. *Mimeo*.

Egger, D., & Partner, A. (2013). Detailanalyse der unternehmensprozesse, zuständigkeiten, anreiz-und führungssysteme der regionalen arbeitsvermittlungszentren. *Bern: SECO Publikation Arbeitsmarktpolitik*(33), 4–2013.

Federal Office of Topography swisstopo. (2024). *swissBOUNDARIES3D: Administrative Boundaries of Switzerland.* Geospatial dataset. Wabern, Switzerland. Retrieved from https://www.swisstopo.admin.ch/en/landscape-model-swissboundaries3d (Accessed: 2024.)

Federal Social Insurance Office (FSIO). (2022). *AHV/AVS Individual Earnings Records, 2004–2021.* Administrative microdata accessed through the Swiss Federal Statistical Office (SFSO). Bern, Switzerland. (Accessed: 2024.)

Frandsen, B., Leslie, E., & McIntyre, S. (Forthcoming). Cluster jackknife instrumental variables estimation. *Review of Economics and Statistics.*.

Gerfin, M., & Lechner, M. (2002). A microeconometric evaluation of the active labour market policy in switzerland. *The Economic Journal*, *112*(482), 854–893.

Goldsmith-Pinkham, P., Hull, P., & Kolesár, M. (2025). *Leniency designs: An operator's manual* (Tech. Rep.). National Bureau of Economic Research.

Heidelberg Institute for Geoinformation Technology (HeiGIT). (2024). *OpenRouteService Directions API (driving-car profile), based on OpenStreetMap.* Web service. Heidelberg, Germany. Retrieved from https://openrouteservice.org (Accessed: 2024.)

Huber, M., Lechner, M., & Mellace, G. (2017). Why do tougher caseworkers increase employment? the role of program assignment as a causal mechanism. *Review of Economics and Statistics*, *99*(1), 180–183.

Humlum, A., Munch, J., & Rasmussen, M. (2025). What works for the unemployed? evidence from quasi-random caseworker assignments.

Kircher, P. (2022). Job Search in the 21St Century. *Journal of the European Economic Association*, *20*(6), 2317–2352.

Klaeui, J., Kopp, D., Lalive, R., & Siegenthaler, M. (2026). *Adapting to scarcity: The role of firms in occupational transitions* (Tech. Rep.). ROCKWOOL Foundation Berlin (RFBerlin).

Kline, P., Saggio, R., & Sølvsten, M. (2020). Leave-out estimation of variance components. *Econometrica*, *88*(5), 1859–1898.

Kolesár, M. (2013). *Estimation in an instrumental variables model with treatment effect heterogeneity* (Tech. Rep.).

Lalive, R., Van Ours, J. C., & Zweimüller, J. (2005). The effect of benefit sanctions on the duration of unemployment. *Journal of the European economic association*, *3*(6), 1386–1417.

Liechti, D., Suri, M., Möhr, T., Arni, P., & Siegenthaler, M. (2020). Methoden der Stellensuche und Stellensucherfolg. *Grundlagen für die Wirtschaftspolitik Nr. 33. Staatssekretariat für Wirtschaft SECO, Bern, Schweiz.*

Marinescu, I., & Rathelot, R. (2018). Mismatch unemployment and the geography of job search. *American Economic Journal: Macroeconomics*, *10*(3), 42–70.

Martinez, I. Z., Saez, E., & Siegenthaler, M. (2021). Intertemporal labor supply substitution? evidence from the swiss income tax holidays. *American Economic Review*, *111*(2), 506–546.

Naya, V. A., Bied, G., Caillou, P., Crépon, B., Gaillac, C., Pérennes, E., & Sebag, M. (2021). Designing labor market recommender systems: the importance of job seeker preferences and competition. In *4. idsc of iza workshop: Matching workers and jobs online-new developments and opportunities for social science and practice.*

Rambjer, L., Uhlendorff, A., & Vikström, J. (2026). *Monitoring, benefit sanctions and employment outcomes: Evidence from a judge design* (mimeo).

Rivkin, S. G., Hanushek, E. A., & Kain, J. F. (2005). Teachers, schools, and academic achievement. *Econometrica*, *73*(2), 417–458.

Rockoff, J. E., Jacob, B. A., Kane, T. J., & Staiger, D. O. (2011). Can you recognize an effective teacher when you recruit one? *Education Finance and Policy*, *6*(1), 43–74.

Şahin, A., Song, J., Topa, G., & Violante, G. L. (2014). Mismatch unemployment. *American Economic Review*, *104*(11), 3529–3564.

Schiprowski, A. (2020). The role of caseworkers in unemployment insurance: Evidence from unplanned absences. *Journal of Labor Economics*, *38*(4), 1189–1225.

Schiprowski, A., Schmidtke, J., Schmieder, J., & Trenkle, S. (2024). The effects of unemployment insurance caseworkers on job search effort. In *Aea papers and proceedings* (Vol. 114, pp. 567–571).

Schmieder, J. F., & Trenkle, S. (2020). Disincentive effects of unemployment benefits and the role of caseworkers. *Journal of Public Economics*, *182*, 104096.

SFSO. (2018). *Labour market areas 2018* (Explanatory report). Swiss Federal Statistical Office. Retrieved from https://www.bfs.admin.ch/bfs/de/home/statistiken/querschnittsthemen/raeumliche-analysen/raeumliche-gliederungen/analyseregionen.assetdetail.8948838.html (Accessed: 2024-02-18)

State Secretariat for Economic Affairs (SECO). (2024a). *AVAM/ASAL Unemployment Register, 2004–2023.* Administrative microdata accessed through the Swiss Federal Statistical Office (SFSO). Bern, Switzerland. (Accessed: 2024.)

State Secretariat for Economic Affairs (SECO). (2024b). *Job-Room Click Tracking and Public Vacancy API, June 2020–December 2022.* Click-level data accessed through the Swiss Federal Statistical Office (SFSO); vacancy postings retrieved from the job-room.ch public API. Bern, Switzerland. Retrieved from https://job-room.ch (Accessed: 2024.)

Swiss Federal Statistical Office (SFSO). (2026). *Swiss Earnings Structure Survey (LSE), 2020 and 2022 waves.* Survey microdata. Neuchâtel, Switzerland. Retrieved from https://www.bfs.admin.ch/bfs/en/home/statistics/work-income/wages-income-employment-labour-costs/wage-structure-survey.html (Accessed: 2026.)

Swiss Post. (2024). *Swiss Postcode (PLZ) Directory Shapefile.* Geospatial dataset distributed via opendatasoft.com. Bern, Switzerland. Retrieved from https://swisspost.opendatasoft.com (Accessed: 2024.)

van den Berg, G. J., Foerster, H., & Uhlendorff, A. (2023). Structural empirical analysis of vacancy referrals with imperfect monitoring and the strategic use of sickness absence. *IZA Discussion Papers*.

van den Berg, G. J., Hofmann, B., & Uhlendorff, A. (2019). Evaluating vacancy referrals and the roles of sanctions and sickness absence. *The Economic Journal*, *129*(624), 3292–3322.

van der Klaauw, B., & Vethaak, H. (2022). *Empirical evaluation of broader job search requirements for unemployed workers* (Tinbergen Institute Discussion Paper).

X28 AG. (2024). *X28 Online Job Vacancy Database, January 2020–December 2022.* Proprietary web-scraped vacancy data accessed through research agreement. Zurich, Switzerland. Retrieved from https://www.x28.ch (Accessed: 2024.)

Yap, L. (2024). Inference with many weak instruments and heterogeneity. *arXiv preprint arXiv:2408.11193*.

Online Appendix

When to Broaden, When to Focus: Job Search and Market Tightness

A Data appendix

A.1 Vacancy data cleaning

Starting from the full set of online vacancy postings scraped by X28 between January 2016 and December 2022, I apply several cleaning steps. Many vacancies appear on multiple job boards simultaneously. X28 assigns postings to duplicate groups based on textual similarity; I keep only the earliest posting within each group, removing approximately 8% of observations. I also exclude postings classified as apprenticeships, which follow a distinct institutional hiring process and are not part of the regular labor market.

A small share of postings (0.1%) are mapped to more than five occupations due to a crawler error that assigned an abnormally large number of job titles to individual advertisements in certain months. Because there is no reliable prior for which of the on average 27 listed occupations to assign, I drop these postings. I further drop vacancies that cannot be mapped to an ISCO occupation code (5% of remaining postings) and those with missing postcodes that cannot be mapped to a commuting zone (8% of remaining postings).

For the retained vacancies, I match postcodes to commuting zones by intersecting their spatial representations and assigning each postcode the commuting zone with which it overlaps the most. The map shapefiles are obtained from the Swiss Post and the SFSO, respectively (Swiss Post, 2024; SFSO, 2018). For postcodes that cannot be matched this way, I progressively truncate their digits and assign the modal location at each truncated level, ensuring that all retained postcodes are mapped to a commuting zone.

A.2 Click data processing

I scrape the API for every posting listed on the platform during the recording window. Only clicks on the same advertisement that occur on distinct days are retained.[29]

Each job posting's occupation is extracted by the platform provider and reported in the API response. The classification uses the same internal system as the unemployment records, which I map to ISCO. Only 3% of postings match to more than one ISCO occupation at the most granular level, and no posting matches to more than three; in these rare cases, I randomly select one occupation. For location, 89% of postings directly report the municipality of the job. For the remainder, I match the firm's postcode to municipalities using a crosswalk from the Swiss Federal Statistical Office, assigning ambiguous postcodes to the municipality with the highest building overlap. This procedure assigns a municipality to 93% of job postings. I map municipalities to the 16 commuting zones using the official crosswalk.

A.3 Reemployment earnings from AHV/AVS records

This appendix describes the construction of the reemployment earnings measure from Swiss social security (AHV/AVS) records. The raw data and initial cleaning were prepared by Kristina Schüpbach, based on Stata scripts from (Martinez et al., 2021).

[29] Most repeated clicks within a day occur within the same minute, suggesting they reflect technical issues rather than deliberate search behavior.

I first restrict to persons with unemployment spells starting in 2020 or later. Each record contains a person identifier, an employer identifier, a contribution year, a start month, an end month, and the total contributed income over that period. I classify contribution types using the employer identifier prefix: identifiers starting with `99999` correspond to unemployment benefits, `88888` to invalidity insurance, `77777` to maternity and parental leave (EO), `66666` to military service, and `11111` to care allowances. I also take into account cancellation spells: when a positive and a negative entry for the same employer and year cancel out, I delete both records; for partial corrections, I adjust the income accordingly.

I drop records with month codes greater than 12, which likely reflect coding errors. I retain employment, excluding disability, military, and EO and unemployment income. I expand each AHV spell from its start month to its end month, creating a person-month panel. Monthly earnings are computed by dividing the annual contributed income for each spell equally across the months it covers. I sum earnings across all employers within each person-month (approximately 29% of employment months involve more than one employer). I then merge the panel to the unemployment spell data via the person identifier and compute months relative to the deregistration month.

The primary variable is the maximum monthly employment income across the first three full calendar months after deregistration. Monthly income is computed as described above and includes both wage and self-employment income. I restrict the measure to individuals with at least one social security record in any of those three months. Spells with zero employment income despite the unemployment records reporting an exit to employment are treated as missing so that the log outcome is well defined; this concerns 1,620 spells (1.96% of otherwise eligible spells). Because the AHV data end in December 2021, the sample is restricted to deregistrations before October 2021, ensuring that all three post-deregistration months fall within the data coverage. I winsorize the resulting distribution at the 5th and 95th percentiles. The outcome in regressions is the log of the winsorized maximum monthly earnings. I also use the same AHV records to compute average annual income in the year and two years before unemployment, which enter the analysis as descriptive statistics and as control variables.

Among the 231,427 job finders in the regression sample (with at least one caseworker meeting and a non-missing leave-one-out tightness measure), approximately 114,593 deregistered before October 2021 and are therefore eligible for the earnings measure. Among those who deregistered in 2020, 12,848 out of 14,720 spells (87.3%) have a non-missing wage. The remaining missings among eligible spells reflect either no AHV match or zero recorded employment income in the relevant months. Spells with deregistration in October 2021 or later are set to missing by construction, as the required three months of post-deregistration AHV data are not available.

A.4 Commuting times

I obtain the travel distance between municipalities from openrouteservices.org, by the Heidelberg Institute for Geoinformation Technology at Heidelberg University (Heidelberg Institute for Geoinformation Technology (HeiGIT), 2024). Openrouteservices.org is a service that uses OpenStreetMap data to calculate travel distances between two points. I use the 'driving-car' mode of transport, which is the default mode of transport in the service. While many people in Switzerland use trains and other public transport for their commutes, train travel times are much harder to obtain and tend to correlate highly with car driving times. I extract the shortest car travel distance between the centroids of the two municipalities. I obtain the centroids from the swissBOUNDARIES3D shapefile published by swisstopo (Federal Office of Topography swisstopo, 2024). If the municipality of residence and the municipality of a job are the same, I pick 50 random points within the municipality and calculate the average distance between them.

A.5 Imputing missing found-market cells

For job finders with a missing found occupation and/or commuting zone, I impute the cell from the new employer's 2-digit NOGA industry, which is recorded on the unemployment register independently of the occupation and commuting zone. The intuition is that within a given region, the industry of the new job is informative about the occupation: a retail employer in Geneva most often hires a cashier, not a software engineer. I use the empirical joint distribution of occupations, industries, and regions among spells where all three are observed to predict the missing cell.

Formally, let $N(n, r, o)$ denote the count of spells where NOGA-2 industry n, commuting zone r, and ISCO-3 occupation o are all observed. For a spell i with missing o but observed r_i and n_i, I assign an industry-weighted found-market tightness

$$\log \widetilde{\tau}_i \;=\; \log\left(\frac{\sum_o N(n_i, r_i, o)\, V(o, r_i, t_i)}{\sum_o N(n_i, r_i, o)\, U(o, r_i, t_i)}\right), \tag{5}$$

where V and U are the smoothed vacancy and unemployment counts at panel cell (o, r_i, t_i) on the deregistration date t_i, and missing V or U at a candidate cell are treated as zero. Analogous formulas apply when only r is missing (weighting over candidate regions given o and n) and when both are missing (joint weighting over (o, r) pairs given n). When the (NOGA-2, region) cell has no observed candidates, the imputation falls through to the next-coarsest distribution (NOGA-2 marginal, then unconditional).

Imputation shares. Among job finders within 6 months whose found-market cell is incomplete (31,651 spells), 64.7% are imputed using the empirical distribution of occupations given region and industry, 31.9% using the distribution of regions given occupation and industry, and 0.6% using the joint distribution of occupations and regions given industry alone. The remaining 2.9% fall back to coarser distributions (NOGA-2 marginal or unconditional). After imputation, the three-market decomposition is available for 98.5% of job finders; the residual 1.5% enter as a fourth *unobserved* category in the robustness specification (Section 6.3.1, Appendix Table C.18).

B Johnson–Lindenstrauss Approximation

This appendix provides technical details on how leave-one-out predictions are computed efficiently with high-dimensional fixed effects.

B.1 Hat algebra

The key insight, see e.g. Kline et al. (2020), is that leave-one-out predictions can be expressed using leverage (hat) values without re-estimating the model for each observation. For observation i, define the leverage as:

$$h_{ii} = X_i(X'X)^{-1}X_i' \tag{6}$$

where h_{ii} is the i-th diagonal element of the hat matrix $H = X(X'X)^{-1}X'$. The leverage h_{ii} measures how much observation i influences its own fitted value.

Using this, the leave-one-out prediction—the prediction from a model estimated on all observations except i—can be written as:

$$\hat{E}_{-i}[Y|X] = \hat{y}_i - \frac{h_{ii} \cdot \hat{u}_i}{1 - h_{ii}} \tag{7}$$

where $\hat{y}_i = X_i\hat{\beta}$ is the fitted value and $\hat{u}_i = y_i - \hat{y}_i$ is the residual from the full-sample regression. This formula avoids re-estimating the model N times by expressing each leave-one-out prediction as a simple function of full-sample quantities.

B.2 Random projection approximation

Computing exact leverages h_{ii} requires forming the matrix $(X'X)^{-1}$, which is computationally very costly when X includes high-dimensional fixed effects. Following Kline et al. (2020), I approximate leverages using the Johnson–Lindenstrauss Approximation (JLA) with $p = 1{,}000$ Rademacher random projections:

1. Draw random vectors $r^{(s)} \in \{-1/\sqrt{p}, +1/\sqrt{p}\}^n$ for $s = 1, \ldots, p$, where each element is $+1/\sqrt{p}$ or $-1/\sqrt{p}$ with equal probability.

2. For each s, compute $z^{(s)} = X(X'X)^{-1}X'r^{(s)}$. This requires only one linear solve per simulation rather than N inversions.

3. For each simulation, also compute the residual projection $m_i^{(s)} = r_i^{(s)} - z_i^{(s)}$. Approximate the leverage using the ratio estimator:

$$\hat{h}_{ii} = \frac{\sum_{s=1}^{p}\big(z_i^{(s)}\big)^2}{\sum_{s=1}^{p}\big(z_i^{(s)}\big)^2 + \sum_{s=1}^{p}\big(m_i^{(s)}\big)^2} \tag{8}$$

   which is bounded in $[0, 1]$ and has better finite-sample properties than the naive sum $\sum_s (z_i^{(s)})^2$.

Figure C.1: Time allocation of Swiss PES caseworkers across main activities.

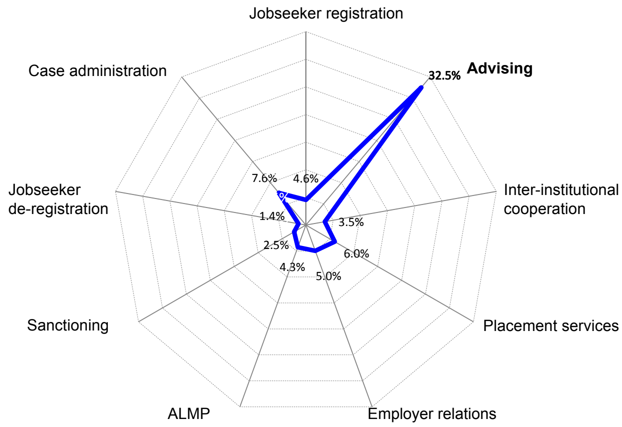

Notes: Adapted from Egger & Partner (2013), own translation.

The intuition is that the random projections provide unbiased estimates of the diagonal elements of the hat matrix through a law of large numbers argument. With $p = 1{,}000$ projections, the approximation error is negligible for inference.

C Further figures and tables

Table C.1: Effects of caseworker redirection on click intensity by initial market tightness

	(1)	(2)
Dependent Var.:	Total clicks	Clicks per day
Leave-one-out tightness	0.0086	-0.0011
	(0.0662)	(0.0656)
Leave-one-out tightness x Tercile 2	-0.0322	0.0162
	(0.0824)	(0.0828)
Leave-one-out tightness x Tercile 3	-0.0037	0.0231
	(0.0835)	(0.0827)
Tightness tercile FE	X	X
Unemployment office x Year FE	X	X
Last occ. x commuting zone x quarter FE	X	X
Jobseeker controls	X	X
N caseworkers	2,460	2,460
Observations	100,528	100,528
Mean of dependent var.	41.991	0.50323

Notes: PPML regressions following Equation (3) with the leave-one-out caseworker tightness tendency interacted with terciles of the jobseeker's initial market tightness (last occupation $\times$ commuting zone at registration). Column (1) uses the total number of clicks over the first three months of the spell; Column (2) uses clicks per day in unemployment. Standard errors clustered at the caseworker level in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C.2: Effects of caseworker redirection on clicks per day in each market type, by initial market tightness

	Average effect (Poisson)			By initial-tightness tercile		
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Var.:	Clicks/day in tighter	Clicks/day in same	Clicks/day in slacker	Clicks/day in tighter	Clicks/day in same	Clicks/day in slacker
Leave-one-out tightness	0.0164	-0.0418	0.0261	-0.0063	-0.1662***	-0.0478
	(0.0488)	(0.0514)	(0.0475)	(0.0217)	(0.0598)	(0.0450)
Leave-one-out tightness x Tercile 2				0.0795**	0.1618**	-0.0599
				(0.0335)	(0.0813)	(0.0554)
Leave-one-out tightness x Tercile 3				0.0953**	0.2524***	0.0913*
				(0.0423)	(0.0869)	(0.0513)
Tightness tercile FE				X	X	X
Unemployment office x Year FE	X	X	X	X	X	X
Last occ. x commuting zone x quarter FE	X	X	X	X	X	X
Jobseeker controls	X	X	X	X	X	X
Magnitude p10-p90 (avg)	+0.5%	-1.3%	+0.8%			
Magnitude p10-p90, slack initial market				-0.2%	-5.1%	-1.5%
Magnitude p10-p90, medium initial market				+2.3%	-0.1%	-3.3%
Magnitude p10-p90, tight initial market				+2.8%	+2.7%	+1.4%
N caseworkers	2,460	2,460	2,460	2,460	2,460	2,460
Observations	100,262	96,930	100,050	100,060	96,805	99,918
Squared Cor.	0.27653	0.24934	0.32764	0.93462	0.64627	0.90530
Mean of dependent var.	0.22250	0.09970	0.17231	0.22236	0.09979	0.17250

Notes: PPML regressions following Equation (3) with an offset of log(clicks per day), so coefficients capture the rate at which clicks are directed to each market type holding total click intensity fixed. The dependent variable in each column is clicks per day in markets that are tighter than, equal in tightness to, or slacker than the jobseeker's initial market (last occupation $\times$ commuting zone). Clicks are restricted to the first three months of the spell. Columns (1)–(3) report the average effect of the leave-one-out caseworker tightness tendency; Columns (4)–(6) interact it with terciles of the jobseeker's initial market tightness. The corresponding p10-to-p90 effects are plotted in Figure C.8. Standard errors clustered at the caseworker level in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C.3: Click-finding relationship by caseworker redirection tendency

	(1)	(2)
Dependent Var.:	Found job in market	Found job in market
Constant	0.0057***	
	(6.21e-5)	
Clicks in market	0.0200***	0.0219***
	(0.0002)	(0.0002)
Clicks $\times$ Leave-one-out tightness	-0.0010	-0.0012
	(0.0011)	(0.0011)
Jobseeker spell FE		X
$\%\Delta\hat{\beta}^{\text{clicks}}$ (P10$\rightarrow$P90 CW)	-1.6%	-1.6%
Observations	3,865,141	3,865,141
R2	0.06308	0.07118

Notes: The dependent variable is an indicator equal to one if the jobseeker finds a job in the market. For each jobseeker who found employment and clicked on at least five postings, the choice set includes all markets (ISCO 3-digit occupation $\times$ commuting zone) in which at least one click was placed, plus 50 randomly sampled alternatives from the universe of markets with at least 30 total clicks. Clicks are top-coded at 10. The interaction term tests whether the per-click job-finding probability varies with the caseworker's leave-one-out tightness tendency. Column (1) includes no fixed effects; Column (2) includes jobseeker fixed effects. Standard errors clustered at the caseworker level in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C.4: Caseworker assignment methods across Swiss cantons

Assignment method	N cantons	Share of sample
Equal distribution by administration (proportional to caseworker's contractual hours, regardless of caseload)	6	22.4%
Caseload balancing (distribute to equalize case burden across caseworkers)	10	44.4%
Random assignment (software-based random allocation)	1	12.6%
Assignment by management (team or office lead decides, potentially based on skills)	4	13.0%
Non-response	5	7.6%

Notes: Results from a survey of all 26 Swiss cantons on how jobseekers are assigned to caseworkers at their local Public Employment Service office. Categories reflect the primary assignment method reported by each canton. The ´Share of sample' column reports the share of unemployment spells in the cantons using the method.

Table C.5: Robustness: Cantons with more strictly quasi-random assignment rules

	Find job within 6 months				
	(1)	(2)	(3)	(4)	(5)
Dependent Var.:	In any market	In any market	In tighter market	In origin market	In slacker market
Leave-one-out tightness	0.0576***	0.0577***	0.0357***	0.0223**	0.0010
	(0.0175)	(0.0165)	(0.0098)	(0.0103)	(0.0085)
Unemployment office x Year FE	X	X	X	X	X
Last occ. x commuting zone x quarter FE	X	X	X	X	X
Jobseeker controls		X	X	X	X
N caseworkers	950	950	950	950	950
Observations	128,996	128,996	128,996	128,996	128,996
R2	0.10040	0.16128	0.12122	0.10798	0.16723
Mean of dependent var.	0.35414	0.35414	0.13646	0.10356	0.10738

Notes: The table reports estimates from the specification in Equation (3), restricting the sample to cantons where a survey of cantonal employment authorities confirms that caseworker assignment follows either equal distribution by administration (proportional to contractual hours, not current caseload) or random software allocation. The dependent variable indicates whether the jobseeker finds employment within 6 months. The table also reports the effect of going from the 10th to the 90th percentile in the caseworker leave-one-out tendency, expressed as a percentage of the dependent variable. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

D Caseworker learning: event-study design

This appendix describes the event-study design used to test whether caseworkers learn from their clients' job-finding outcomes and pass this information on to other clients (Section 7).

The design is a stacked event study. An *event* is defined as a client of caseworker c finding a job in three-digit ISCO occupation o, deregistering at time t. If a caseworker has multiple clients finding jobs in the same occupation, the first occurrence defines the event. Events are restricted to July 2021 through August 2022 to ensure the full estimation window fits the data period. The *treatment group* consists of caseworker c's other active clicking clients in the months around the event, excluding the job finder. The *control group* consists of clients of other caseworkers in the same unemployment office, providing a within-office comparison that absorbs office-level trends in demand for occupation o. Both groups are collapsed to the event–month level: the outcome is the mean number of clicks placed in occupation o, weighted by the number of active clients.

Each event generates a $[-12, +4]$ month window around the deregistration date. I estimate a PPML regression of mean clicks on event-time indicators interacted with a treatment indicator, with event fixed effects, month $\times$ region fixed effects, and month $\times$ occupation fixed effects. Standard errors are clustered at the event-caseworker level. The reference period is $t = -3$.

Table C.6: Robustness: Click window threshold

	(1)	(2)	(3)	(4)	(5)
Leave-one-out tightness	0.0006	0.0333***	0.0312***	0.0279***	0.0258**
	(0.0058)	(0.0080)	(0.0090)	(0.0097)	(0.0107)
Unemployment office x Year FE	X	X	X	X	X
Last occ. x commuting zone x quarter FE	X	X	X	X	X
Jobseeker controls	X	X	X	X	X
--------------------------------------	--------	----------	----------	----------	---------
Click window	30 days	60 days	90 days	120 days	180 days
p10-p90 magnitude	+0.1%	+3.3%	+2.7%	+2.3%	+2.0%
N clickers	38,818	85,723	104,334	114,560	125,361
N caseworkers	2,132	2,414	2,460	2,484	2,501
Observations	295,925	357,894	366,893	370,820	373,879
R2	0.13501	0.14065	0.14162	0.14198	0.14248
Mean of dependent var.	0.36518	0.36588	0.36554	0.36521	0.36494

Notes: The table reports estimates from the specification in Equation (3), varying the window over which clicks are observed for the search tightness index and the caseworker leave-one-out tendency; the baseline uses the first three months of unemployment (90 days, Column (3)). Sample restrictions ($\geq$5 clicks per clicker, $\geq$5 clickers per caseworker) are applied within each window. The dependent variable indicates whether the jobseeker finds employment within 6 months. The table also reports the effect of going from the 10th to the 90th percentile in the caseworker leave-one-out tendency, expressed as a percentage of the dependent variable. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table C.7: Testing Random Assignment: Jobseeker Characteristics and Caseworker Tendencies

	Panel A: own clicks (clickers only)			Panel B: caseworker tendency (full sample)		
Variable	Coef (SE)	RW p-val	p10-p90 magnitude	Coef (SE)	RW p-val	p10-p90 magnitude
Female	-0.07568 (0.00224)	0.000	-33.6%	-0.02605 (0.00688)	0.003	-1.7%
Age 35-49	0.00049 (0.00243)	0.839	+0.4%	-0.00488 (0.00652)	0.907	-0.5%
Age 50+	-0.01276 (0.00211)	0.000	-15.0%	-0.00868 (0.00715)	0.772	-1.2%
Has children	-0.01359 (0.00252)	0.000	-9.3%	0.00288 (0.00684)	0.968	+0.2%
Female x has children	-0.04504 (0.00200)	0.000	-45.1%	-0.01195 (0.00559)	0.269	-1.9%
Non-Swiss	-0.00208 (0.00242)	0.715	-1.4%	0.01043 (0.00791)	0.770	+0.7%
Primary or lower sec. education	-0.04269 (0.00199)	0.000	-49.0%	-0.00139 (0.00811)	0.968	-0.2%
University education	0.06686 (0.00190)	0.000	+178.6%	0.01444 (0.00487)	0.037	+2.7%
3+ years tenure in last occupation	0.00824 (0.00239)	0.004	+3.6%	-0.00828 (0.00856)	0.841	-0.4%
Employed in year t-1	0.00076 (0.00080)	0.715	+0.3%	-0.00302 (0.00245)	0.772	-0.1%
Employed in year t-2	0.00229 (0.00102)	0.096	+1.1%	-0.00640 (0.00320)	0.330	-0.3%
Income in year t-1 (log)	0.03023 (0.00751)	0.001	+1.7%	-0.00544 (0.02326)	0.968	-0.0%
Income in year t-2 (log)	0.03115 (0.00908)	0.004	+1.8%	-0.04096 (0.02875)	0.731	-0.2%
N	101,514			366,893		

Notes: Each row is a separate regression of the predetermined jobseeker characteristic on a tightness measure, with the same fixed effects as the main specification in Equation (3). Panel A uses the jobseeker's own click tightness on the clicker subsample; Panel B uses the caseworker leave-one-out tightness tendency on the full sample. The RW column reports Romano-Wolf step-down multiple-testing-adjusted p-values across the 13 covariates within each panel, computed via cluster wild bootstrap (Rademacher, B=1999) on the caseworker level. The p10–p90 magnitude reports the predicted percent change in the covariate when the tightness measure moves from its 10th to its 90th percentile.

Table C.8: Meeting effects on search targeting: number of clicks and clicked-market tightness

	N clicks	Clicked market tightness
	(1)	(2)
Post meeting	0.8359***	-0.0084
	(0.0138)	(0.0062)
Post meeting * Leave-one-out tightness	0.1055	0.1044**
	(0.0689)	(0.0450)
Sample	Meeting-day panel	Click level
Unemployment spell FE	X	X
Calendar month FE	X	X
Month of spell FE	X	X
Meeting weekday FE	X	
Bin start hour of day FE	X	
Click weekday FE		X
Click hour-of-day FE		X
--------------------------------------	-----------------	-----------
Family	Poisson	OLS
Observations	5,084,932	592,417
Mean of dependent var.	0.39296	-0.61095
p10-p90 magnitude		3.29%
Pseudo R2 / R2	0.2674	0.4978

Notes: This table presents difference-in-differences estimates of the effect of the first caseworker meeting on jobseekers' search activity. Column (1) is a PPML regression of the daily click count on the post-meeting indicator and its interaction with the leave-one-out caseworker tightness measure, estimated on a meeting-day panel covering days -21 to $+20$ around the first meeting. Column (2) is an OLS regression of the log tightness (V/U ratio) of the clicked market on the same regressors, estimated at the click level (one observation per click). The leave-one-out caseworker tightness measure captures how much a caseworker's other clients (excluding the focal jobseeker) search in tight versus slack markets in the first three months of their spell, residualized using unemployment office $\times$ year, last occupation, and temporal fixed effects. Both columns include unemployment-spell, calendar-month, and month-of-spell fixed effects; Column (1) additionally controls for the meeting weekday and bin-start hour-of-day, while Column (2) controls for the click weekday and click hour-of-day. The p10–p90 magnitude reports the change in clicked-market tightness from moving the leave-one-out measure from its 10th to its 90th percentile. Standard errors are clustered at the unemployment-spell level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table C.9: Cumulative Job Finding Effects by Time Horizon

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent Var.:	30 days	60 days	90 days	120 days	150 days	180 days	240 days	365 days
Leave-one-out tightness	0.0017	0.0025	0.0110*	0.0206***	0.0292***	0.0311***	0.0352***	0.0387***
	(0.0017)	(0.0042)	(0.0060)	(0.0072)	(0.0083)	(0.0090)	(0.0100)	(0.0114)
Unemployment office x Year FE	X	X	X	X	X	X	X	X
Detailed last occupation FE	X	X	X	X	X	X	X	X
Last occ. x commuting zone x quarter FE	X	X	X	X	X	X	X	X
Jobseeker controls	X	X	X	X	X	X	X	X
N caseworkers	2,460	2,460	2,460	2,460	2,460	2,460	2,460	2,460
Observations	366,893	366,893	366,893	366,893	366,893	366,893	343,647	299,297
R2	0.04718	0.07979	0.10498	0.12196	0.13504	0.14161	0.14881	0.14096
Mean of dependent var.	0.01827	0.09015	0.17364	0.24880	0.31347	0.36554	0.44699	0.55043

Notes: The table reports estimates from the specification in Equation (3), with the dependent variable indicating whether the jobseeker finds employment within the horizon shown in the column header (30, 60, 90, 120, 150, 180, 240, or 365 days). The sample is restricted to ensure sufficient follow-up time for each horizon, so sample size varies across columns. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table C.10: Period-by-Period Cloglog Analysis

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Dependent Var.:	30 days	60 days	90 days	120 days	150 days	180 days	210 days	240 days	270 days	300 days	330 days	360 days
Leave-one-out tightness	0.1536	0.0082	0.1068**	0.1442***	0.1733***	0.1148*	0.1302*	0.1159	0.2298***	0.1011	0.0204	0.1042
	(0.1004)	(0.0514)	(0.0472)	(0.0503)	(0.0551)	(0.0605)	(0.0670)	(0.0767)	(0.0837)	(0.0940)	(0.1007)	(0.1106)
Unemp. office x Year x Last occupation FE	X	X	X	X	X	X	X	X	X	X	X	X
Observations	152,993	274,200	266,715	238,308	208,595	181,167	154,129	128,856	108,507	93,768	80,507	68,430
Pseudo R2	0.09284	0.07966	0.07915	0.07713	0.07799	0.07954	0.07963	0.08411	0.08780	0.09137	0.09581	0.09463
Mean of dependent var.	0.04341	0.09509	0.11320	0.11375	0.11146	0.10309	0.09935	0.09545	0.09162	0.08774	0.08556	0.08248

Notes: This table presents complementary log-log (cloglog) regression results for discrete-time hazard analysis across monthly periods. Each column represents a 30-day period (Period 1: 0-30 days, Period 2: 31-60 days, etc.). The dependent variable is an indicator for finding a job within the specific period, conditional on not having found a job in previous periods. The key regressor is the leave-one-out caseworker tendency (residualized). All regressions include RAV office-year fixed effects and last occupation (ISCO) fixed effects. Standard errors are clustered at the caseworker level. The sample includes only observations that are "at risk" for each period (have not yet found employment). Significance levels: *** p<0.01, ** p<0.05, * p<0.10.

Table C.11: Robustness: Alternative demeaning and tightness measures

	Alternative demeaning			Alternative tightness measures			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Leave-one-out tightness	0.0250***	0.0261***	0.0141***	0.0206**	0.0162**	0.0318**	
	(0.0082)	(0.0088)	(0.0030)	(0.0085)	(0.0073)	(0.0137)	
Leave-one-out log V							0.0346***
							(0.0101)
Leave-one-out log U							-0.0264**
							(0.0116)
Unemployment office x Year x Last industry FE	X						
Last occ. x commuting zone x quarter FE	X	X		X	X	X	X
Unemployment office x Year x Last 2-dig. occupation FE		X					
Unemployment office x Year FE			X	X	X	X	X
Jobseeker controls	X	X	X	X	X	X	X
Specification	RAV-Yr-Noga	RAV-Yr-Occ2	No choice-set FE	CZ-ISCO4	Loc-Occ	Ads/Click	V/U split
p10-p90 magnitude	+2.4%	+2.4%	+3.3%	+1.9%	+1.6%	+1.8%	V +2.8%, U -1.9%
N caseworkers	2,460	2,460	2,460	2,460	2,460	2,460	2,460
Observations	333,335	349,348	391,785	368,166	368,055	368,731	366,893
R2	0.17736	0.16413	0.10727	0.14194	0.14190	0.14199	0.14162
Mean of dependent var.	0.36716	0.36601	0.36575	0.36569	0.36569	0.36561	0.36554

Notes: The table reports estimates from the specification in Equation (3), with alternative specifications. Columns (1)–(3) use alternative fixed-effect structures: (1) interacts office × year FE with 2-digit industry (NOGA) of the last job; (2) interacts office × year FE with 2-digit last occupation (ISCO); (3) does not include a FE for the last occupation x commuting zone x quarter. Columns (4)–(7) use alternative tightness measures: (4) 4-digit ISCO codes instead of 3-digit; (5) 101 labor market regions instead of 16 commuting zones; (6) number of clickers per advertisement on the platform; (7) decomposes the caseworker tendency into separate leave-one-out measures for log vacancies and log unemployment, entered jointly. The dependent variable indicates whether the jobseeker finds employment within 6 months. The table also reports the effect of going from the 10th to the 90th percentile in the caseworker leave-one-out tendency, expressed as a percentage of the dependent variable. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table C.12: Robustness: Effect on non-clickers sample

	Find job within 6 months (Non-clickers only)				
	(1)	(2)	(3)	(4)	(5)
Dependent Var.:	In any market	In any market	In tighter market	In origin market	In slacker market
Leave-one-out tightness	0.0424***	0.0394***	0.0176***	0.0155**	0.0050
	(0.0112)	(0.0105)	(0.0060)	(0.0073)	(0.0059)
Unemployment office x Year FE	X	X	X	X	X
Last occ. x commuting zone x quarter FE	X	X	X	X	X
Jobseeker controls		X	X	X	X
p10-p90 magnitude	+3.8%	+3.6%	+4.6%	+4.7%	+1.3%
N caseworkers	2,460	2,460	2,460	2,460	2,460
Observations	265,379	265,379	265,379	265,379	265,379
R2	0.09622	0.15879	0.10999	0.10856	0.14204
Mean of dependent var.	0.35404	0.35404	0.12379	0.10618	0.11848

Notes: The table reports estimates from the specification in Equation (3), restricted to jobseekers with no clicks on job-room.ch; the caseworker tendency is measured from other, clicking clients and assigned to non-clickers. The dependent variable indicates whether the jobseeker finds employment within 6 months. The table also reports the effect of going from the 10th to the 90th percentile in the caseworker leave-one-out tendency, expressed as a percentage of the dependent variable. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table C.13: Robustness: Clustered leave-one-out, same registration quarter excluded

	Find job within 6 months				
	(1)	(2)	(3)	(4)	(5)
Dependent Var.:	In any market	In any market	In tighter market	In origin market	In slacker market
Leave-one-out tightness (CW×Quarter cluster)	0.0307***	0.0291***	0.0132***	0.0101**	0.0051
	(0.0078)	(0.0074)	(0.0043)	(0.0051)	(0.0043)
Unemployment office x Year FE	X	X	X	X	X
Last occ. x commuting zone x quarter FE	X	X	X	X	X
Jobseeker controls		X	X	X	X
--	-------------	-------------	----------------	----------------	-----------------
p10-p90 magnitude	+3.1%	+2.9%	+3.7%	+3.4%	+1.5%
N caseworkers	2,460	2,460	2,460	2,460	2,460
Observations	366,938	366,938	366,938	366,938	366,938
R2	0.08200	0.14165	0.09812	0.09309	0.13213
Mean of dependent var.	0.36552	0.36552	0.13026	0.10688	0.12291

Notes: The table reports estimates from the specification in Equation (3), with the leave-one-out measure computed leaving out not just the focal jobseeker but all clients of the same caseworker who registered in the same quarter, following Frandsen et al. (Forthcoming). The dependent variable indicates whether the jobseeker finds employment within 6 months. The table also reports the effect of going from the 10th to the 90th percentile in the caseworker leave-one-out tendency, expressed as a percentage of the dependent variable. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table C.14: Robustness: Clustered leave-one-out, same occupation excluded

	Find job within 6 months				
	(1)	(2)	(3)	(4)	(5)
Dependent Var.:	In any market	In any market	In tighter market	In origin market	In slacker market
Leave-one-out tightness (CW×Occ cluster)	0.0284***	0.0266***	0.0130***	0.0097	0.0030
	(0.0092)	(0.0087)	(0.0050)	(0.0060)	(0.0050)
Unemployment office x Year FE	X	X	X	X	X
Last occ. x commuting zone x quarter FE	X	X	X	X	X
Jobseeker controls		X	X	X	X
--	-------------	-------------	----------------	----------------	-----------------
p10-p90 magnitude	+2.7%	+2.5%	+3.5%	+3.2%	+0.8%
N caseworkers	2,460	2,460	2,460	2,460	2,460
Observations	366,938	366,938	366,938	366,938	366,938
R2	0.08197	0.14162	0.09812	0.09309	0.13212
Mean of dependent var.	0.36552	0.36552	0.13026	0.10688	0.12291

Notes: The table reports estimates from the specification in Equation (3), with the leave-one-out measure computed leaving out not just the focal jobseeker but all clients of the same caseworker with the same last occupation, following Frandsen et al. (Forthcoming). The dependent variable indicates whether the jobseeker finds employment within 6 months. The table also reports the effect of going from the 10th to the 90th percentile in the caseworker leave-one-out tendency, expressed as a percentage of the dependent variable. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table C.15: Robustness: Sensitivity to click and caseworker clicker thresholds

Click Threshold	CW Clicker Threshold	N	Clickers	Caseworkers	Coef (Any)	Magnitude (Any) %	Coef (Tighter)	Magnitude (Tighter) %
1	2	381,394	144,147	2716	0.0344***	2.96	0.017***	4.16
1	5	379,406	143,655	2536	0.0364***	3.11	0.018***	4.35
1	10	375,033	142,561	2373	0.0333***	2.78	0.018***	4.27
1	30	341,686	133,078	1894	0.0224**	1.73	0.0136**	2.96
5	2	369,984	102,044	2653	0.033***	2.92	0.0147***	3.66
5	5	366,893	101,514	2460	0.0311***	2.70	0.0155***	3.81
5	10	360,255	100,296	2283	0.029***	2.46	0.0144***	3.44
5	30	291,543	86,322	1573	0.0041	0.31	0.0061	1.27
10	2	358,002	77,867	2604	0.0298***	2.85	0.0161***	4.31
10	5	353,822	77,324	2405	0.0301***	2.82	0.017***	4.49
10	10	342,117	75,700	2174	0.0278***	2.53	0.014**	3.59
10	30	231,171	57,380	1244	0.0012	0.09	0.008	1.64
30	2	311,642	36,877	2445	0.0168***	2.18	0.008**	2.84
30	5	297,230	36,004	2146	0.016**	1.96	0.0073*	2.47
30	10	253,115	32,734	1670	0.0076	0.84	0.0086	2.61
30	30	48,706	9,570	309	0.0051	0.41	-0.0241	-4.68

Notes: The table reports estimates from the specification in Equation (3), for different combinations of the minimum click threshold (rows) and the minimum number of clicking clients per caseworker (columns); the baseline uses 5 clicks and 5 clickers per caseworker. Each cell reports the coefficient and the implied effect of going from the 10th to the 90th percentile in the caseworker leave-one-out tendency, expressed as a percentage of the dependent variable. The dependent variable indicates whether the jobseeker finds employment within 6 months. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table C.16: Robustness: Complementary log-log specification

	Find job within 6 months (cloglog)			
	(1)	(2)	(3)	(4)
Dependent Var.:	In any market	In tighter market	In origin market	In slacker market
Leave-one-out tightness	0.1171***	0.1467***	0.1387**	0.0069
	(0.0356)	(0.0464)	(0.0697)	(0.0482)
Unemployment office x Year FE	X	X	X	X
Last occ. x commuting zone x quarter FE	X	X	X	X
p10-p90 magnitude	+2.8%	+4.2%	+4.0%	+0.2%
N caseworkers	2,460	2,460	2,460	2,460
Observations	366,893	366,893	366,893	366,893
Squared Cor.	0.08217	0.07082	0.07200	0.11398
Mean of dependent var.	0.36554	0.13027	0.10688	0.12291

Notes: The table reports complementary log-log estimates, replacing the linear probability model in Equation (3). The dependent variable indicates whether the jobseeker finds employment within 6 months. The table also reports the effect of going from the 10th to the 90th percentile in the caseworker leave-one-out tendency, expressed as a percentage of the dependent variable. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table C.17: Robustness: COVID-19 sample split at 17 February 2022

	Registered before 17 Feb 2022				Registered on/after 17 Feb 2022			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent Var.:	In any market	In tighter market	In origin market	In slacker market	In any market	In tighter market	In origin market	In slacker market
Leave-one-out tightness	0.0291***	0.0165***	0.0123*	0.0005	0.0366**	0.0125	0.0220**	-0.0034
	(0.0097)	(0.0056)	(0.0067)	(0.0056)	(0.0147)	(0.0099)	(0.0098)	(0.0099)
Unemployment office x Year FE	X	X	X	X	X	X	X	X
Last occ. x commuting zone x quarter FE	X	X	X	X	X	X	X	X
Jobseeker controls	X	X	X	X	X	X	X	X
p10-p90 magnitude	+2.6%	+4.2%	+3.7%	+0.1%	+2.9%	+2.7%	+6.3%	-0.8%
N caseworkers	2,435	2,435	2,435	2,435	2,102	2,102	2,102	2,102
Observations	282,123	282,123	282,123	282,123	84,770	84,770	84,770	84,770
R2	0.15778	0.10320	0.10036	0.13679	0.14231	0.12574	0.11427	0.16800
Mean of dependent var.	0.35930	0.12703	0.10619	0.12067	0.38630	0.14105	0.10919	0.13036

Notes: The table re-estimates the baseline specification of Equation (3) on two subsamples split by the jobseeker's registration date. Columns (1)–(4) restrict to spells registering before 17 February 2022, when the Swiss federal government lifted the remaining COVID-19 restrictions; Columns (5)–(8) restrict to spells registering on or after that date. Within each subsample, Column (1)/(5) reports the effect on finding a job within 6 months in any market, and Columns (2)–(4)/(6)–(8) decompose this effect by whether the new job is in a tighter, the initial, or a slacker market than the jobseeker's initial market. All specifications include the full controls. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table C.18: Robustness: No imputation of missing found-market cells

	Find job within 6 months					
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Var.:	In any market	In any market	In tighter market	In origin market	In slacker market	Market unknown
Leave-one-out tightness	0.0323***	0.0311***	0.0157***	0.0143**	0.0017	-0.0006
	(0.0096)	(0.0090)	(0.0057)	(0.0063)	(0.0049)	(0.0103)
Unemployment office x Year FE	X	X	X	X	X	X
Last occ. x commuting zone x quarter FE	X	X	X	X	X	X
Jobseeker controls		X	X	X	X	X
p10-p90 magnitude	+2.8%	+2.7%	+5.5%	+4.3%	+0.7%	-0.2%
N caseworkers	2,460	2,460	2,460	2,460	2,460	2,460
Observations	366,893	366,893	366,893	366,893	366,893	366,893
R2	0.08196	0.14161	0.08450	0.09310	0.10439	0.07902
Mean of dependent var.	0.36554	0.36554	0.09253	0.10682	0.08066	0.08554

Notes: The table reports estimates from the specification in Equation (3) without imputing the tightness of reemployment markets with missing occupation or location information. Column (1) reports the effect on finding a job within 6 months in any market under the minimal fixed effects; Column (2) adds the full controls. Columns (3)–(5) decompose the any-market outcome of Column (2) by whether the new job is in a tighter, the initial, or a slacker market than the jobseeker's initial market, using only spells with a directly observed found-market cell (no NOGA-2 imputation). Column (6) reports the effect on the residual outcome of finding a job whose market cell is unobserved. The coefficients in Columns (3)–(6) sum to the any-market coefficient in Column (2) by construction. The table also reports the effect of going from the 10th to the 90th percentile in the caseworker leave-one-out tendency, expressed as a percentage of the dependent variable. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table C.19: Heterogeneity by baseline tightness: job-finding outcomes

	Job finding (within 6 months)			
	(1)	(2)	(3)	(4)
Dependent Var.:	Any market	In tighter	In origin	In slacker
Leave-one-out tightness	0.0179*	0.0180*	0.0009	-0.0024
	(0.0106)	(0.0092)	(0.0077)	(0.0051)
Leave-one-out tightness x Tercile 2	0.0213	0.0028	0.0172*	0.0018
	(0.0145)	(0.0121)	(0.0099)	(0.0093)
Leave-one-out tightness x Tercile 3	0.0181	-0.0094	0.0215**	0.0063
	(0.0155)	(0.0115)	(0.0107)	(0.0129)
Tightness tercile FE	X	X	X	X
Unemployment office x Year FE	X	X	X	X
Last occ. x commuting zone x quarter FE	X	X	X	X
Jobseeker controls	X	X	X	X
Magnitude p10-p90, slack initial market	+1.5%	+4.1%	+0.3%	-0.6%
Magnitude p10-p90, medium initial market	+3.4%	+5.1%	+5.3%	-0.2%
Magnitude p10-p90, tight initial market	+3.4%	+2.2%	+7.5%	+1.0%
N caseworkers	2,460	2,460	2,460	2,460
Observations	362,911	362,911	362,911	362,911
R2	0.14147	0.09732	0.09232	0.13107
Mean of dependent var.	0.36646	0.13140	0.10794	0.12422

Notes: The table reports estimates from the specification in Equation (3), with the leave-one-out measure interacted with terciles of the jobseeker's initial market tightness (tightness in the last occupation $\times$ commuting zone at registration). The dependent variable in Column (1) indicates whether the jobseeker finds employment within 6 months. Column (2) shows as an outcome only unemployment exits to jobs in markets that are, at deregistration date, tighter than the jobseeker's initial market. Column (3) shows the likelihood of exits to the initial market and Column (4) to slacker markets. The table also reports the effect of going from the 10th to the 90th percentile in the caseworker leave-one-out tendency, expressed as a percentage of the dependent variable. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table C.20: Caseworker redirection and job finding by occupation $\times$ geography category

	Find job within 180 days (unconditional)			
	(1)	(2)	(3)	(4)
Dependent Var.:	Same occ, same CZ	Same occ, diff CZ	Diff occ, same CZ	Diff occ, diff CZ
Leave-one-out tightness	0.0012	0.0009	0.0047	0.0152***
	(0.0074)	(0.0046)	(0.0075)	(0.0045)
Leave-one-out tightness x Medium initial market	0.0186*	0.0024	0.0065	-0.0080
	(0.0101)	(0.0062)	(0.0101)	(0.0064)
Leave-one-out tightness x Tight initial market	0.0223**	-0.0068	0.0113	-0.0090
	(0.0108)	(0.0067)	(0.0107)	(0.0072)
Tightness tercile FE	X	X	X	X
Unemployment office x Year FE	X	X	X	X
Last occ. x commuting zone x quarter FE	X	X	X	X
Jobseeker controls	X	X	X	X
--	-----------------	-----------------	-----------------	-----------------
Average effect	0.015**	-0.001	0.011*	0.009***
	(0.006)	(0.003)	(0.006)	(0.003)
N caseworkers	2,460	2,460	2,460	2,460
Mean dep. (slack)	0.10	0.03	0.09	0.03
Mean dep. (medium)	0.13	0.04	0.12	0.04
Mean dep. (tight)	0.13	0.05	0.13	0.05
Observations	333,517	333,517	333,517	333,517
R2	0.10364	0.07915	0.08047	0.07909

Notes: The table reports estimates from the specification in Equation (3), with the leave-one-out measure interacted with terciles of the jobseeker's initial market tightness (last occupation $\times$ commuting zone at registration). The dependent variable in each column is the unconditional probability of finding a job within six months in the given occupation $\times$ commuting zone category; jobseekers who do not find a job within six months are coded as zero. Standard errors are clustered at the caseworker level. Significance levels: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table C.21: Caseworker characteristics by tightness redirection tendency

	P10 caseworker	P90 caseworker	Beta	p-value	q-value
Caseload					
Sample period caseload	194.9	197.0	6.557	0.43	0.516
Clients using job-room	52.4	53.0	2.126	0.407	0.516
Meetings per month	7.0	7.0	0.234	0.43	0.516
Experience					
Years of experience	6.8	6.6	-0.699	0.383	0.516
Client diversity					
Distinct 3-digit occupations	54.8	55.0	0.456	0.736	0.736
Distinct 2-digit industries	43.4	44.1	2.212	0.022**	0.131

Notes: The table compares observable characteristics of caseworkers at the 10th and 90th percentiles of the leave-one-out tightness tendency within their employment office. The first two columns report average predicted values when the tightness tendency is set to the 10th and 90th percentile, respectively, with all other covariates held at their observed values. Beta reports the coefficient from regressions following Equation 3. The q-value is the Benjamini–Hochberg adjusted p-value. The number of distinct occupations refers to the last occupation before unemployment of the assigned jobseekers. Years of experience are calculated as the time between the caseworker's first observed case and the sample start, censored at 16 years.

Table C.22: Correlation matrix of caseworker dimensions

	Tightness	Support			Stringency		Job Finding VA				Reemp. Wage VA	
Variable	Tightness	N meetings	Avg meeting dur.	Training course	Proof of effort accepted	Any sanction	Job finding	Found tighter	Found slacker	Found same	Reemp. earnings (log)	Market wage (log)
Tightness	1.000	-0.033	-0.008	-0.009	0.009	0.015	0.075	0.087	-0.011	0.040	0.030	0.121
N meetings	-0.033	1.000	-0.375	0.189	-0.127	0.124	0.033	0.042	-0.036	0.045	-0.021	-0.019
Avg meeting dur.	-0.008	-0.375	1.000	-0.091	0.066	-0.092	-0.087	-0.041	-0.027	-0.067	0.022	0.018
Training course	-0.009	0.189	-0.091	1.000	-0.144	0.198	0.110	0.028	0.074	0.067	-0.023	-0.012
Proof of effort accepted	0.009	-0.127	0.066	-0.144	1.000	-0.639	-0.157	-0.123	-0.088	-0.048	0.017	0.004
Any sanction	0.015	0.124	-0.092	0.198	-0.639	1.000	0.258	0.163	0.152	0.104	-0.025	-0.034
Job finding	0.075	0.033	-0.087	0.110	-0.157	0.258	1.000	0.577	0.480	0.538	0.042	0.008
Found tighter	0.087	0.042	-0.041	0.028	-0.123	0.163	0.577	1.000	0.164	-0.110	0.024	0.157
Found slacker	-0.011	-0.036	-0.027	0.074	-0.088	0.152	0.480	0.164	1.000	-0.240	0.055	-0.165
Found same	0.040	0.045	-0.067	0.067	-0.048	0.104	0.538	-0.110	-0.240	1.000	-0.004	0.012
Reemp. earnings (log)	0.030	-0.021	0.022	-0.023	0.017	-0.025	0.042	0.024	0.055	-0.004	1.000	-0.003
Market wage (log)	0.121	-0.019	0.018	-0.012	0.004	-0.034	0.008	0.157	-0.165	0.012	-0.003	1.000

Notes: The table shows pairwise correlations between leave-one-out caseworker measures. All measures are computed using the methodology described in Section 5. "Tightness" is the main measure of caseworker tendency to redirect toward tight markets. Support measures include number of meetings, average meeting duration, and training course referrals. Stringency measures include proof-of-effort acceptance rate and any sanction. Job-finding VA measures caseworker value added for job finding overall and by destination market type. Wage VA measures value added for reemployment earnings and market wage level conditional on job finding. Correlations are computed at the unemployment-spell level.

Figure C.2: Vacancy and candidate counts on the job portal for two neighboring sales occupations in Basel

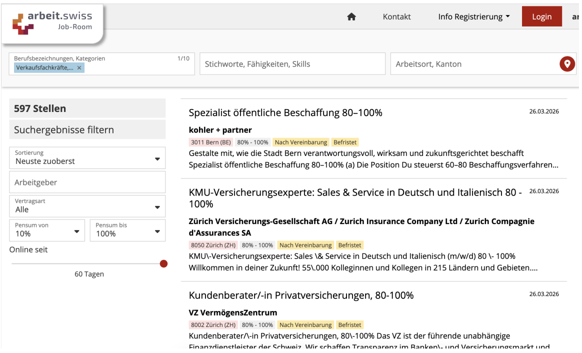

(a) Vacancies: ISCO 332 — Sales and purchasing agents and brokers (597 listings)

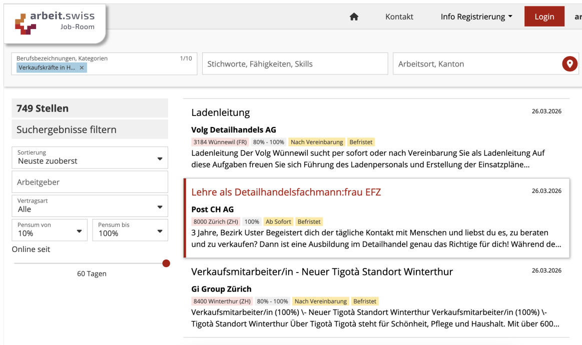

(b) Vacancies: ISCO 522 — Shop salespersons (749 listings)

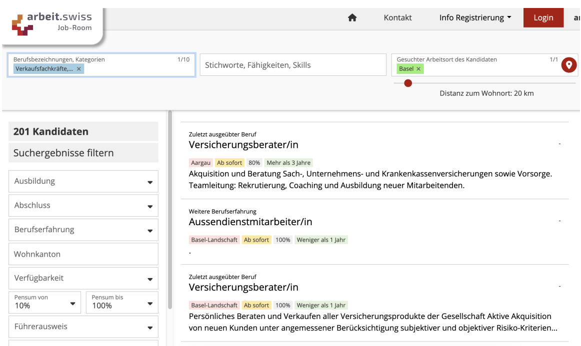

(c) Candidates: ISCO 332 — Sales and purchasing agents and brokers (201 candidates)

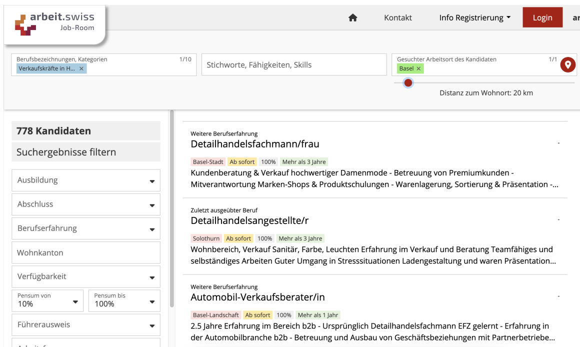

(d) Candidates: ISCO 522 — Shop salespersons (778 candidates)

Notes: Screenshots from the Swiss public employment service job portal (job-room.ch), taken on 26 March 2026, showing vacancy listings and anonymized candidate profiles for two neighboring sales occupations in Basel. ISCO 332 (Sales and purchasing agents and brokers) has a vacancy-to-candidate ratio of approximately $597/201 \approx 3.0$, while ISCO 522 (Shop salespersons) has a ratio of $749/778 \approx 1.0$. Despite both being sales occupations, tightness differs by a factor of three. The vacancy counts and candidate lists are publicly accessible, illustrating the type of labor market information available to caseworkers when advising jobseekers.

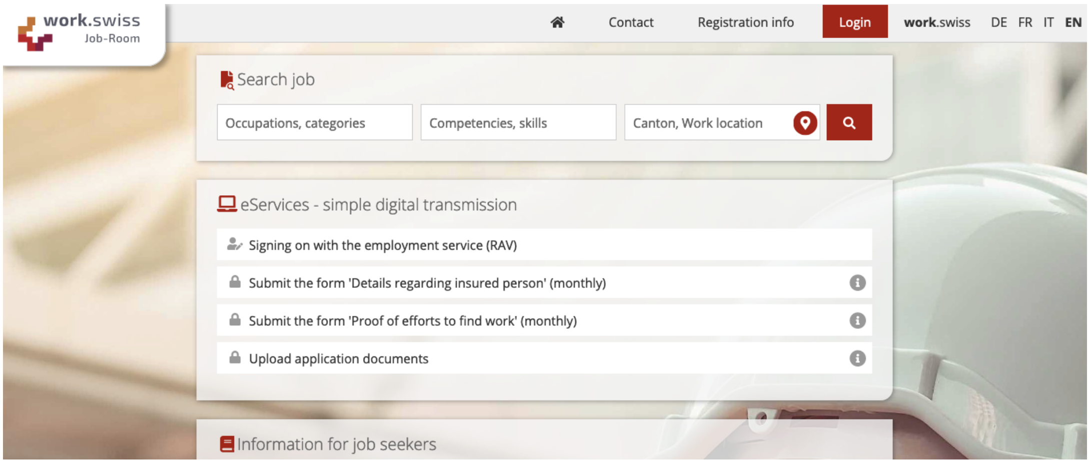

Figure C.3: Screenshot job-room.ch.

Screenshot of the starting page of the job search portal, https://job-room.ch/home/job-seeker. Screenshot taken 10-02-2023

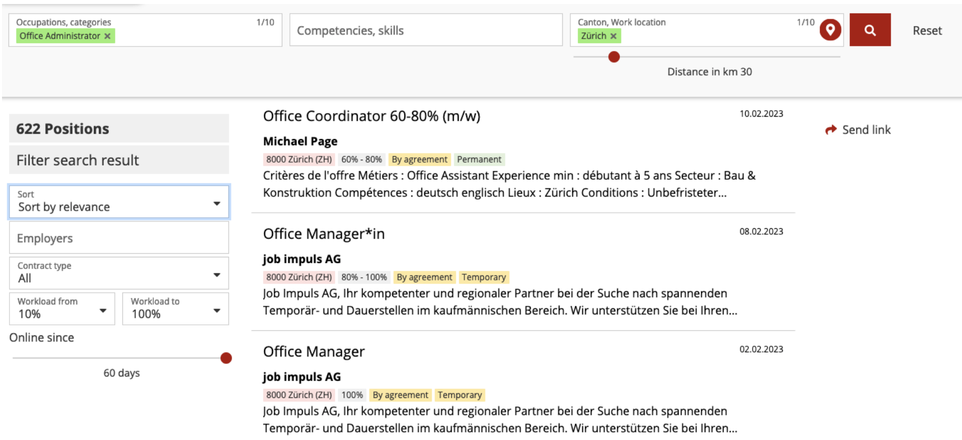

Figure C.4: Screenshot of the result list when searching for "Office managers" in Zurich on the job room (https://job-room.ch/home/job-seeker). Screenshot taken 10-02-2023.

Figure C.5: Tightness dispersion across market pairs: alternative samples

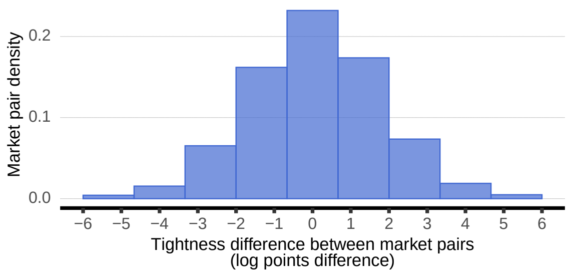

(a) All market pairs

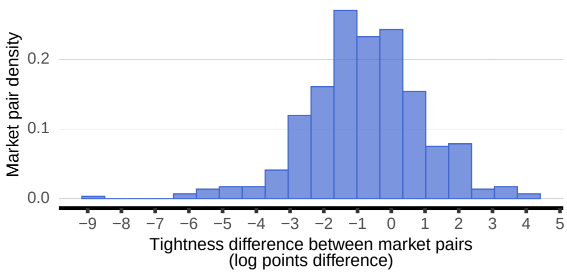

(b) High transition pairs

Notes: Distribution of absolute log tightness differences across market pairs. A market is defined as an ISCO 3-digit occupation × commuting zone cell. Panel (a) includes all market pairs. Panel (b) restricts to pairs with historical transition rates of at least 20% (based on 2004–2019 job-to-job transitions).

Figure C.6: Caseworker effects on job search: Number of clicks around the first meeting

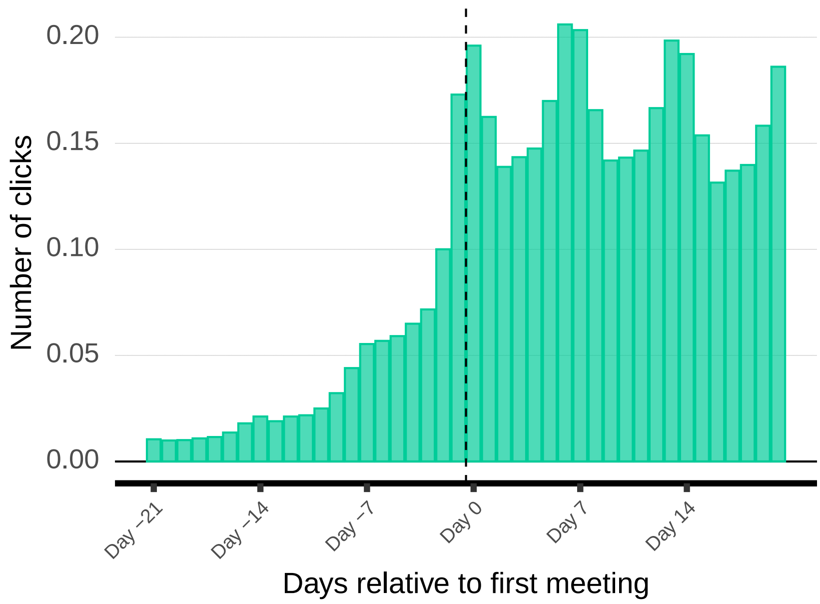

Notes: This figure plots the mean number of clicks per day around the first meeting with the assigned caseworker, over the window from 21 days before to 20 days after the meeting. Daily click counts are winsorized at the 95th percentile before averaging. The dashed vertical line marks the boundary between the day before and the day of the first meeting.

Table C.23: Orthogonality of tightness tendency to other caseworker behaviors

	Support			Stringency	
	(1)	(2)	(3)	(4)	(5)
Dependent Var.:	N meetings	Meeting dur.	Training course	Proof accepted	Any sanction
Tightness tendency	-0.0844**	-0.0087	-0.0061	0.0014	0.0041
	(0.0430)	(0.0180)	(0.0100)	(0.0046)	(0.0106)
Unemployment office x Year FE	X	X	X	X	X
Last occ. x commuting zone x quarter FE	X	X	X	X	X
Jobseeker controls	X	X	X	X	X
Magnitude p10-p90	-0.7%	+0.5%	-1.1%	+0.0%	+0.5%
N caseworkers	2,460	2,460	2,460	2,460	2,460
Observations	366,893	366,893	366,893	339,397	366,893
R2	0.33843	0.40760	0.09096	0.05752	0.07984
Mean of dependent var.	3.6543	-0.59066	0.18139	0.94161	0.27953

Notes: The table reports estimates from the specification in Equation (3), with the dependent variable replaced by caseworker behavior dimensions. Columns (1)–(5) show, respectively, number of meetings, average meeting duration, training course referral, proof-of-effort acceptance rate, and any sanction. The table also reports the effect of going from the 10th to the 90th percentile in the caseworker leave-one-out tendency, expressed as a percentage of the dependent variable. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Figure C.7: Effects over the unemployment spell

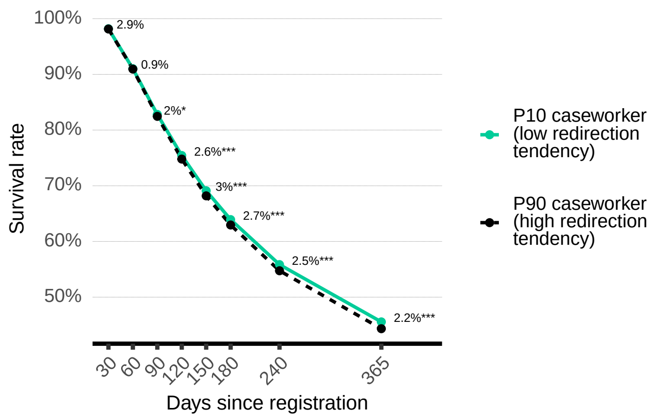

(a) Survival Function

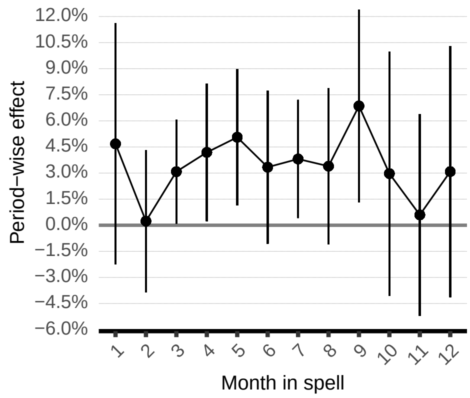

(b) Period-wise Effects

Notes: Panel (a) shows effects on the survival function (1 minus the cumulative job-finding rate). The graphs display predictions obtained by holding all covariates at their sample means and varying only the caseworker's redirection tendency, comparing the 10th and 90th percentiles of the leave-one-out measure. Separate regressions are estimated by period, with samples restricted to ensure complete follow-up at each horizon. Numbers report percentage effects on the exit rate when moving from a p10 to a p90 caseworker. Stars indicate significance levels (*** p<0.01, ** p<0.05, * p<0.10). Panel (b) reports period-specific effects from complementary log-log models. The sample includes only observations that are "at risk" for each period (have not yet found employment). The graphs again show predicted percentage effects of moving from a p10 to a p90 caseworker, with 95% confidence intervals, computed using clustered bootstrap (n=50). The cloglog models include only minimal fixed effects (unemployment office-year and detailed last occupation). Regression results are provided in Appendix Tables C.9 and C.10.

Figure C.8: Heterogeneous effects of caseworker redirection on clicks per day in each market type

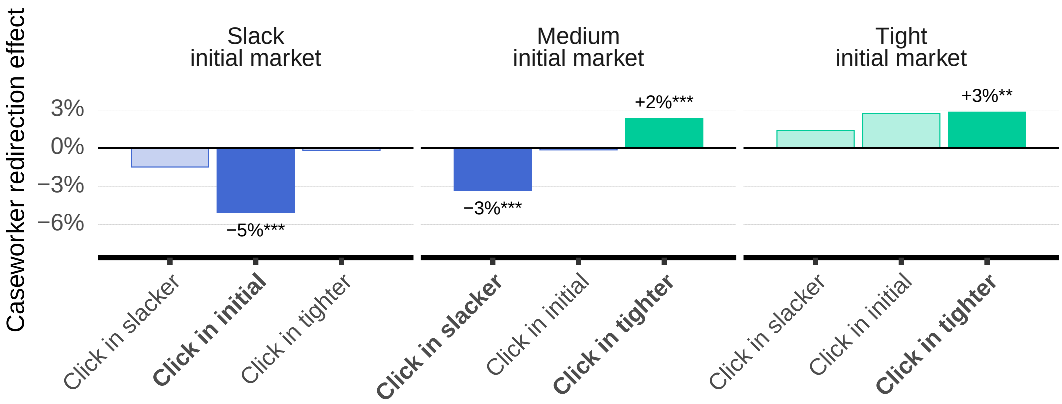

Notes: The figure plots p10-to-p90 effects of the caseworker leave-one-out tightness tendency on clicks per day in each market type, by tercile of the jobseeker's initial market tightness (last occupation × commuting zone at registration). Estimates come from a PPML regression following Equation (3) with the leave-one-out measure interacted with the terciles and an offset of log(clicks per day), so the bars show effects on the rate at which clicks are directed to each market type holding total click intensity fixed. Clicks are restricted to the first three months of the spell. The corresponding effects on job finding are shown in Figure 5 in the main text. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Figure C.9: Caseworker redirection effects on job-finding hazard by occupation × geography category

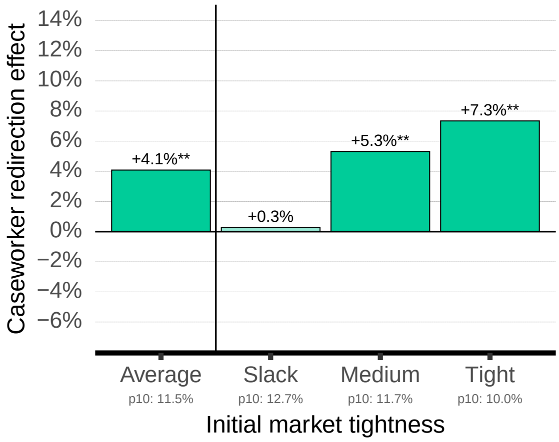

(a) Same occupation, same commuting zone

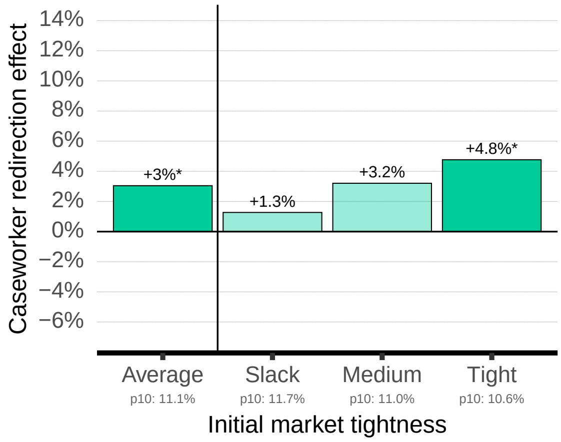

(b) Different occupation, same commuting zone

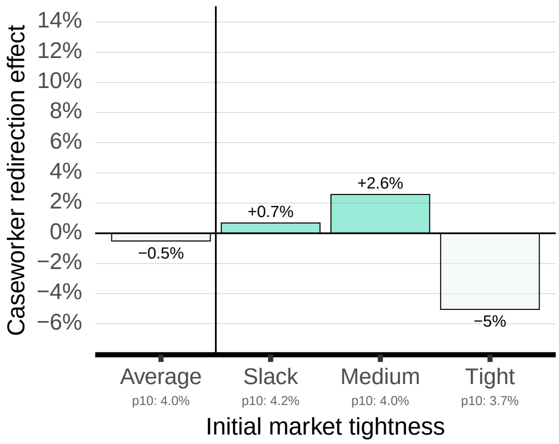

(c) Same occupation, different commuting zone

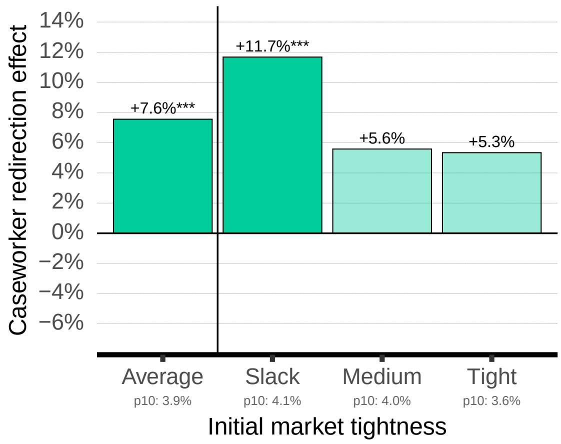

(d) Different occupation, different commuting zone

Notes: The figure plots p10-to-p90 effects of the caseworker leave-one-out tightness tendency on the unconditional probability of finding a job in the given occupation × commuting zone category within six months, by tercile of the jobseeker's initial market tightness (last occupation × commuting zone at registration), from the specification in Equation (3) with the leave-one-out measure interacted with the terciles. The top row conditions on the same commuting zone, the bottom row on a different commuting zone. Jobseekers who do not find any job within six months are coded as zero. Below each x-axis category, the baseline probability at the p10 of the leave-one-out tightness measure is reported. The underlying regressions are reported in Appendix Table C.20. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Figure D.1: Caseworker learning from client job-finding events

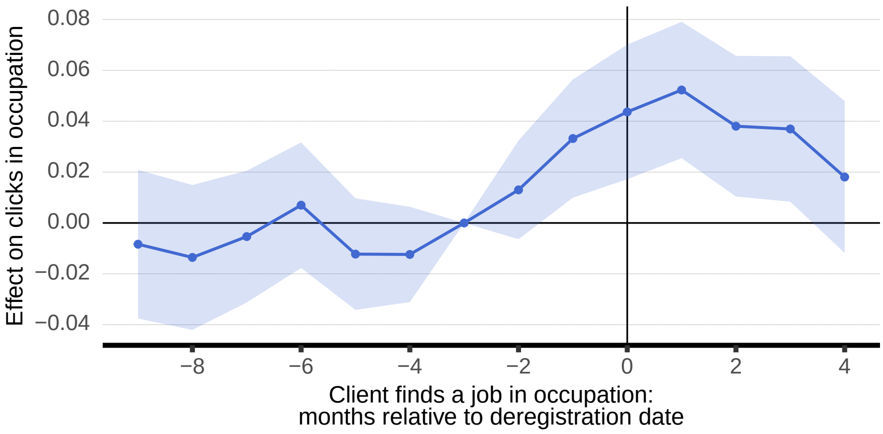

Notes: Event-study estimates of how caseworkers' other clients adjust their search after one client finds a job in a particular occupation. The reference period is $t = -3$. See Appendix D for details on the event-study design. Standard errors are clustered at the event-caseworker level.