The Labor Hub

Job Market Paper Series No. 4

Advisor: Roland Rathelot

WHEN TO BROADEN, WHEN TO FOCUS: JOB SEARCH AND MARKET TIGHTNESS

Jeremias Klaeui

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

Jeremias Klaeui*

October 23, 2026
Click here for the latest version

Abstract

Navigating the labor market is hard and many public employment services advise jobseekers. Yet little is known about how this advice changes where they search and their chances of finding a job. This paper investigates whether advice directing their search toward tight markets improves re-employment. I link Swiss unemployment records to clickstream data from the public job portal, which show where jobseekers search before they find a job. I exploit quasi-random assignment to caseworkers differing in their tendency to give such advice. The tendency is unrelated to other caseworker behaviors. Moving from a 10th to a 90th percentile caseworker raises six-month job finding by 2.8%, without lowering job quality. Jobseekers from slack markets gain from mobility to tighter markets, while those from tight markets gain from focusing on their market, suggesting that not only advice that broadens search but also advice that focuses it can be beneficial. A back-of-the-envelope calculation indicates that the advice would still direct search toward tighter markets if all caseworkers gave it, suggesting scope for wider use.

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. Studies have investigated a wide range of active labor market policies that aim to help the

*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.

unemployed, many of which have been found ineffective (Lalive et al., 2008; Card et al., 2010; Crépon & van den Berg, 2016; Card et al., 2018). Jobseekers face substantial information frictions and likely do not search optimally (Krueger & Mueller, 2016; Altmann et al., 2018; Belot et al., 2019; Mueller et al., 2021; Jäger et al., 2024; Christensen et al., 2026). A major task of public employment services (PES) is to counsel and inform jobseekers. In Switzerland, for instance, running the public employment offices cost 456 million Swiss francs (478 million euros) in 2024, about 2,500 Swiss francs (2,600 euros) per registered jobseeker and year, and caseworkers spend 33% of their working time on advising jobseekers. In a typical meeting, they spend 15–20 minutes discussing an individualized job search strategy (Egger & Partner, 2013; SECO, 2025a,b). Yet little is known about whether this advice changes where jobseekers search and helps them find jobs.

A common idea in the literature is that jobseekers should search in occupations and regions where they can work and where there are many jobs and little competition (Şahin et al., 2014; Marinescu & Rathelot, 2018; Kircher, 2022). Labor market tightness, the ratio of vacancies to unemployed jobseekers in an occupation-region market, combines the number of jobs and the competition for them. Tightness differs substantially even between two markets that are close substitutes and pay similar wages.[1] For a jobseeker who can work in both markets, the tighter one has fewer competitors per vacancy. Searching there should raise the chance of finding a job. A jobseeker who is uncertain about how tight the two markets are may devote too much search to the slacker one.

This paper investigates whether advice directing job search toward tight markets improves re-employment. I use unique data, linking the administrative unemployment records of the Swiss PES to the clickstream of its official job portal, which records every job posting a jobseeker views. The data show not only whether jobseekers find a job but where they search before: the occupations and regions they are interested in. I can thus observe the search strategies of jobseekers, measure how much caseworkers change them and estimate whether this caseworker influence raises job finding. To measure where jobseekers search, I define a market as the interaction of 113 three-digit ISCO occupations and Switzerland's 16 commuting zones. For each market, I compute tightness from the near-universe of online vacancies and the population of registered unemployed. I measure the tightness of a jobseeker's search as the average log tightness of the markets in which the jobseeker clicks on job postings.

From this search tightness, I extract the caseworker's component: the tendency to redirect search toward tight markets. Caseworkers are trained to give such advice: the caseworker

[1] Among occupation-region pairs with a historical transition rate of at least 20% and median market wages within 5% of each other, the median absolute difference in log tightness is 0.85 log points. In the median pair, the tighter market thus has about 2.3 times as many vacancies per unemployed jobseeker as the slacker one; see Section 3.3.

diploma guidelines describe advising jobseekers to stay in their market if it is good and to consider related markets otherwise (HR Swiss Exams, 2023). At the same time, caseworkers cover many other topics in meetings and, in a survey, report a lot of freedom in choosing them (Arni & Schiprowski, 2015), so the tendency to redirect to tight markets can differ between them.

I measure the caseworker's tendency to redirect search toward tight markets with a leave-one-out estimator (e.g. Kling, 2006; Humlum et al., 2025). For each jobseeker, I use only the search tightness of the caseworker's other clients, so that the measure reflects the caseworker and not the jobseeker. Jobseekers are assigned to a caseworker at their local employment office, and this assignment is as good as random (Arni & Schiprowski, 2019). A survey I conducted among all cantons shows that most assign new jobseekers by fixed rules, such as to the next free caseworker. I compare only jobseekers who register at the same office in the same year, among whom the assignment is as good as random. In addition, I compare only jobseekers with the same registration quarter and the same initial market: the same last occupation before unemployment and the same commuting zone of residence. These jobseekers have the same tighter and slacker markets to choose from, so differences in where they search reflect the caseworker rather than where they start. The identifying comparison is therefore between two jobseekers becoming unemployed in the same quarter, with the same last occupation, living in the same commuting zone and registering at the same office, but one is assigned to a caseworker who tends to make clients search in tight markets and the other not. I show that among such jobseekers, the caseworker's tendency is balanced with respect to the jobseeker's characteristics.

Before turning to the results, I describe how jobseekers use the portal and where they search. A substantial share of jobseekers uses the PES portal for their job search: 28% click on at least five job postings in the first three months of unemployment. This is enough to measure the caseworker's tendency to redirect jobseekers toward tight markets for almost all jobseekers, including those who do not click: 98.7% of jobseekers have a caseworker with at least five such clickers. Clicks are a good measure of search direction: where jobseekers click strongly predicts where they find a job. The register data also show that job finding is closely related to tightness: the six-month job-finding rate rises nearly linearly with the log tightness of the jobseeker's initial market. Jobseekers do search across markets: only a fifth of the postings they view are in their initial market. Yet more than a third of clicks go to markets that are slacker than the initial market. This leaves room for advice that directs search toward tight markets.

My first result is that caseworkers shape where their clients search: some systematically steer search toward tight markets while others do not. Being assigned to a caseworker at the 90th rather than the 10th percentile of this tendency leads a jobseeker to search in markets that are 3.8% tighter on average. To further investigate when jobseekers change

their search, I combine the exact timing of the first meeting with the caseworker with click timestamps. An event study provides suggestive evidence that the change starts with the meeting. Before the meeting, the effect of the tendency on click tightness is zero but imprecisely estimated; after it is positive and clearly significant. The tendency has no effect on the number of clicks: it shapes where jobseekers search, not how much.[2]

My second result is that assignment to a caseworker with a higher tendency improves job finding. Being assigned to a caseworker at the 90th rather than the 10th percentile of this tendency raises a jobseeker's probability of finding a job within six months by about one percentage point, or 2.8% above the baseline job-finding rate. A back-of-the-envelope calculation with averages from application data on the Swiss unemployed (Lalive et al., 2025) shows that this is about the increase in job finding one would expect from the change in search (Section 5.2). For comparison of the effect size, large-scale experiments that advise all jobseekers on their search find employment effects of up to 1%. An online job search assistance platform has no detectable effect on employment (Ben Dhia et al., 2022). Recommending vacancies based on the job ads a jobseeker has viewed raises employment within six months by 0.6% (Barbanchon et al., 2023) and recommending firms likely to hire raises job finding for short-term contracts by 1% (Behaghel et al., 2024). As with the tendency to redirect search, I measure caseworkers' tendencies in other behaviors, such as assigning clients to training courses or sanctioning them with benefit cuts. Comparing caseworkers at the 90th and 10th percentiles of each tendency, the effect of the tendency to redirect search toward tight markets on job finding is about 70% of the effect of the training tendency and about 30% of the effect of the sanction tendency. Compared with training courses and benefit cuts, advice on where to search is a light-touch intervention.

The effect is robust across cantons with different assignment practices and to leave-one-out measures that account for shocks common to clients who register at the same time or come from the same occupation, and it holds for jobseekers who do not use the job portal themselves.[3] Neither vacancy referrals nor the other caseworker behaviors I measure drive the effect: the tendency does not predict finding a job through a referral and its effect is virtually unchanged when I control for meetings, meeting duration, training courses, benefit sanctions and monitoring leniency. The markets where jobseekers find work provide a further test against the concern that the tendency captures some other, unobserved caseworker quality, such as motivation or help with applications. Training courses and sanctions raise job finding also in markets slacker than the initial one, and an unobserved caseworker quality would likely do the same. The tendency, by contrast,

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

[3] But other jobseekers with the same caseworker do use the portal, allowing me to measure the caseworker's tendency.

does not raise job finding in slacker markets, consistent with advice that redirects search toward tight markets.

My third result is that the gains in job finding do not come at the expense of job quality. This is not obvious: some markets may be tight because nobody wants to work there, for example because wages are low or working conditions are hard. Advice that directs search toward these markets would then trade faster job finding for worse jobs. I look at several dimensions of job quality: the stability of employment, earnings, the market-level median wage, commuting time and a revealed-preference measure of the value of an occupation. The advice raises the probability of finding a stable job, a job in a market with higher wages than the initial market and a job in an occupation with a higher value than the last one, but not the probability of finding an unstable job, a job in a market with lower wages or a job in an occupation with a lower value. Training courses and sanctions, by contrast, raise the probability of finding both better and worse jobs (Section 6.4).

My fourth result is that assignment to a caseworker with a higher tendency improves job finding across the full distribution of initial market tightness, but the channel differs sharply. For jobseekers starting in slack markets, the advice raises job finding only in tighter markets: they gain from mobility. For those starting in tight markets, it raises job finding only in their initial market: they gain from focusing on it. For those starting in medium-tightness markets, both channels are at work: the advice moves their search away from slacker markets and raises job finding both in tighter markets and in their initial market. The click data show the same pattern: search shifts toward the markets in which the additional jobs are found. The effect of caseworkers steering their clients' search toward tight markets also differs with their other behaviors: the increases in job finding are most pronounced among caseworkers who tend to meet their clients often and who tend to assign them to training, and less pronounced among caseworkers who tend to sanction often.

The estimates show the effect of assigning one jobseeker to a caseworker with a higher tendency. If the PES instructed every caseworker to act like a caseworker with a high redirection tendency, many jobseekers would change where they search at the same time. In matching models with several labor markets, when many jobseekers move their search from slack to tight markets, the total number of matches rises. The relief in the markets they leave, where many jobseekers compete for few jobs, is substantial, while the additional congestion in the markets they join is minimal. This holds until, in the extreme case, tightness is equalized across markets (Şahin et al., 2014; Marinescu & Rathelot, 2018; Kircher, 2022). This motivates two questions. First, would universal redirection overshoot, moving so many jobseekers that the markets they join become slacker than the markets they leave, which would lower the number of matches again in the model? Second, how much more or less would a jobseeker's search tightness rise under universal

adoption, compared with when only her own caseworker redirects her search? To answer both, I approximate how the tightness of every market would change if every jobseeker had a caseworker at the 90th percentile of the tendency, holding vacancies fixed. The tightness ranking of two markets searched by the same jobseeker reverses in only 0.15% of pairs, so redirection does not overshoot. Search tightness still rises for jobseekers starting in slack, medium and tight markets alike: for a jobseeker who starts in a slack market, by 23% more than when only her own caseworker redirects her search, because others leave the markets she searches in, and for one who starts in a tight market, by about two thirds as much, because her markets receive search from others.

Contributions. This paper makes three contributions. First, it contributes to the literature on broadening job search. Encouraging jobseekers to search in related occupations and nearby regions has been found to help those in slack markets (Belot et al., 2025; Altmann et al., 2026; Belot et al., 2026), but has no effect when applied uniformly to all jobseekers (Belot et al., 2019; Ben Dhia et al., 2022; Altmann et al., 2026) and can even worsen outcomes when broader search is mandated or subsidized (van der Klaauw & Vethaak, 2022; Caliendo et al., 2026).[4] Advice based on tightness distinguishes when to broaden and when to focus. For a jobseeker whose initial market is slack, many related markets are tighter, so the advice is to broaden the search. For a jobseeker whose initial market is tight, few are, so the advice is to focus on that market. Such advice can therefore be given to all jobseekers, not only to those in slack markets. Belot et al. (2025) advise jobseekers in slack occupations to consider tighter alternatives and find large employment gains. I extend this evidence to jobseekers in all markets, using variation in the advice the PES already gives in its daily work.[5]

Second, the paper provides empirical evidence consistent with a mechanism in the literature on search and mismatch, in which reallocating search toward tighter markets raises the number of matches (Şahin et al., 2014; Marinescu & Rathelot, 2018; Kircher, 2022). In these models, the employment gains require that search can move between markets that differ in tightness but not substantially in productivity or utility. The effects I find, without losses in job quality, suggest that such gains exist. Moreover, existing empirical studies show effects of tightness on how jobseekers search and firms hire (Ma & Samaniego de la Parra, 2021; Altmann et al., 2024; Ziegler, 2024; Klaeui et al., 2026). I do the reverse: caseworkers change the allocation of search across markets, and I show the effects

[4] Another strand of interventions uses highly personalized recommender systems (see e.g. Barbanchon et al., 2023; Behaghel et al., 2024; Bächli et al., 2025). These systems use jobseekers' past clicks, predictions of which firms are likely to hire, or online assessments of skills and abilities. Tightness offers another way to personalize advice: it depends only on the jobseeker's market, so it requires little data.

[5] Consistent with this, Altmann et al. (2026) find that information on the number of vacancies in jobseekers' core occupations makes them focus their applications on these occupations and raises their working hours and earnings.

on matches.

Third, the paper contributes to the literature on how caseworkers affect jobseekers. This literature studies the effects of meetings, search requirements, sanctions, vacancy referrals and program assignments, as well as caseworkers' overall value added (Behaghel et al., 2014; Huber et al., 2017; Arni & Schiprowski, 2019; van den Berg et al., 2019; Schiprowski, 2020; 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). Because the clickstream data reveal where clients search, I can show that caseworkers also differ systematically in where they make their clients search. Estimating several caseworker dimensions in the same setting, I find that this is an important additional dimension of caseworker behavior, unrelated to sanctions, training courses and meetings.

Section 2 describes the institutional setting and Section 3 the data. Section 4 presents the identification strategy. Section 5 presents the effects on search and job finding, Section 6 the caseworker mechanisms, including the correlation with other caseworker behaviors, and Section 7 the aggregate effects on search tightness. Section 8 concludes.

2 Institutional background

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).

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 (ALMPs).[6] A time-use study commissioned by the Swiss State Secretariat for Economic Affairs (SECO) documents, based on caseworker surveys and personal interviews, how caseworkers allocate their working hours across these activities (Egger & Partner, 2013); Appendix Figure E.1 reproduces its time-use chart. 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.

[6] 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.

Caseworkers are trained to take labor demand into account when advising on where to search. Every caseworker has to obtain, within five years of being hired, a federal certificate as a human resources specialist in public placement and counseling, or an equivalent qualification (Art. 119b AVIV). The exam guidelines for this certificate, for instance, describe the following typical work situation: "If the situation on the job market for this occupation is good, they set a strategy to find a new job in the previous occupation. If, by contrast, there are few vacancies, they consider together with the jobseeker, based on the jobseeker's competences, in which related occupational fields they could work" (HR Swiss Exams, 2023, p. 31, own translation).[7] The guidelines do not say how to judge whether vacancies are few, and they do not mention the region of search. How to apply this advice is left to the caseworker. In a survey of caseworkers in the canton of Fribourg, caseworkers report the most freedom in choosing the content of their meetings and the least in monitoring and sanctioning jobseekers: on a scale from 1 (very large freedom) to 7 (very small freedom), they rate their freedom at 2.7 for meeting content, 3.0 for referrals to active labor market programs, 3.4 for meeting frequency and 4.5 for monitoring and sanctions (Arni & Schiprowski, 2015, Figure A.4.3).

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.[8]

To go into more detail about the allocation process, Table E.1 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

[7] In July 2021, the ordinance changed the name of the required certificate from "Eidgenössischer Fachausweis Personalberatung" to "HR-Fachmann/HR-Fachfrau mit eidgenössischem Fachausweis, Fachrichtung Öffentliche Personalvermittlung und -beratung". The 2017 edition of the exam guidelines (HR Swiss Exams, 2016), in force during my sample period, contains the same passage.

[8] My identification strategy uses unemployment office × year 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.

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 restrict the sample to cantons using equal distribution or random allocation, which cover 35% of spells, and find an effect that is not significantly different from the baseline (Column (3) of Table 2).

To redirect jobseekers toward tight markets, caseworkers need information about where demand is high relative to supply. Indeed, as confirmed in personal talks with 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 differ in where they get this information. Asked where they mainly inform themselves about current labor market developments, open vacancies and skills in demand, caseworkers in the canton of Fribourg name print media and the internet (68%), colleagues at their office (49%), their own prior work experience (42%), employers (22%) and other sources (19%; multiple answers possible) (Arni & Schiprowski, 2015, Figure A.4.2). 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 E.2). The combination of vacancy counts and candidate counts effectively reveals market tightness.[9] A further channel mentioned in the survey is the caseworker's personal experience. I provide suggestive evidence for such learning mechanisms in Appendix Figure A.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.[10]

[9] As an illustration, Appendix Figure E.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.

[10] 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

3 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, and 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 record start and end dates, demographic, education and occupational history characteristics (coded in both the PES internal classification and ISCO), municipality of residence, and identifiers for the assigned caseworker. The data also document 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; I use these variables to compare the tightness tendency to other caseworker behaviors.

The analysis uses 439,679 unemployment spells with benefit entitlement starting between July 2020 and September 2022. The click data start in June 2020; 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. 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 caseworker tendency is defined only for spells whose occupation $\times$ commuting zone $\times$ quarter cell and unemployment office $\times$ year contain at least two clickers. This drops 25,124 spells, most of them non-clickers in small cells, and the final regression sample comprises 366,890 unemployment spells from 2,460 caseworkers across 132 local employment offices.[11]

the new position.

[11] The leave-one-out tendency in Section 4 is estimated on the clickers with these fixed effects. It is not

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.[12] 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, the portal allows registered unemployed to log their mandatory application protocols.

To search for jobs on the portal, jobseekers specify criteria such as occupation, competencies or skills, and work canton or municipality (see Appendix Figure E.3, Panel (a), 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 30 km 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 E.3, Panel (b), for an example). Each click thus reveals the occupation and location of a vacancy the jobseeker chose to look at, and serves as a proxy for the search criteria they entered.

One concern is that 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 analysis assumes that the direction of search on the platform is representative of the direction of a jobseeker's search overall. A second concern is that 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. 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

defined for a clicker who is alone in a cell, and the fixed effect is not estimated for a cell without clickers.

[12] An empty search on `https://job-room.ch/home/job-seeker`, with no criteria specified, yields the full number of postings available.

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 B.2 describes the API scrape, click deduplication, and the mapping of postings to occupations and commuting zones.

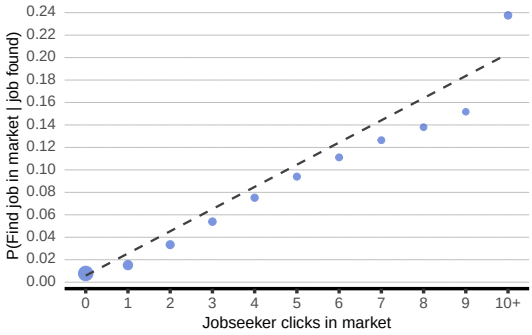
Tightness measure. I measure labor market tightness for cells defined by ISCO 3-digit occupation and commuting zone. Tightness is defined as the stock of vacancies divided by the stock of unemployed jobseekers. The granularity of the cells is chosen to be broad enough to avoid measurement error from 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 are now widely available to PES and researchers. 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.

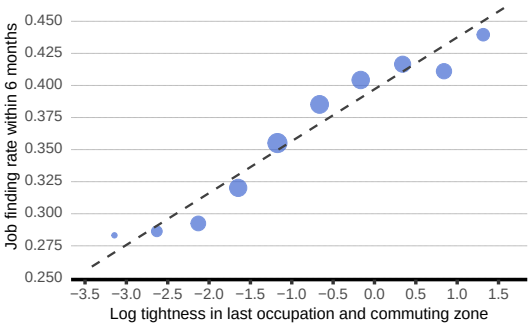
For the denominator, I construct the unemployment count U_{mt} from the unemployment register. A feature of the PES data is that, as in many other countries, there are multiple searched occupations recorded for every jobseeker. These registered occupations are binding: jobseekers can be sanctioned for 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 $1/k$. 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). The data have, 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 cover several million postings between January 2016 and December 2022. I

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) uses job finders with at least five clicks in the first three months and plots the probability of being hired in a market (ISCO 3-digit occupation × commuting zone) against the number of clicks placed there, top-coded at 10. 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, for the 397,298 spells before the caseworker restriction (Section 3) with a tightness measure for this market. The tightness measure is truncated at the 5th and 95th percentiles. The dashed lines are linear fits.

apply a series of cleaning steps and match vacancy postcodes to commuting zones using spatial overlap, both detailed in Appendix B.1. For both jobseekers and vacancies, I apply a 30-day rolling average to the daily stock at the occupation–commuting zone level 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.

3.1 Search tightness index

The search tightness index summarizes where a jobseeker directs their search in a single number that captures the average labor market tightness of targeted markets. Jobseekers search across several occupations and commuting zones (Section 3.3), so two individuals with identical backgrounds — the same last occupation, residence, and registration date — can face different labor demand conditions depending on which mix of markets they target. The index captures this effective tightness.

For each click, I assign the log tightness of the clicked market (ISCO 3-digit occupation ×

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. 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 minimum of five clicks per jobseeker. This index is the variable that enters the leave-one-out caseworker measure constructed in Section 4.

3.2 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. I also report results for shorter and longer windows.

I decompose job finding into three mutually exclusive categories based on 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.[13] If the found market is the initial market, the exit is an exit to the initial market. Otherwise, I measure tightness in both markets at the date of deregistration and classify the exit as an exit to a tighter or to a 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. In the regression sample, 23.3% of job finders within six months find a job whose market is not observed, or whose market or initial market has no tightness measure. I leave these spells out of the analysis by found market, so that the three categories add up to job finding. Section 5.3 shows that the caseworker tendency has no statistically significant effect on finding a job with missing market information, and that the results are very similar when I impute the missing markets from the partial information on them: for most of these job finders, the data record the commuting zone and the industry of the new job but not the occupation, or the occupation and the industry but not the commuting zone (Appendix B.6).

[13]For the 312 place names of new employers that do not exactly match a municipality, I obtain the municipality from ChatGPT (OpenAI, 2024); Appendix B.5 describes the matching.

Job quality outcomes. I measure the quality of the job found in five ways. Section 5.5 uses them as indicators for finding a better or a worse job within six months. *Employment stability* indicates that the jobseeker does not re-register at the PES within 180 or 360 days of deregistration; it is censored at the last recorded deregistration date (March 2023). *Reemployment earnings* are the log of the maximum monthly employment income in the first three full calendar months after deregistration, from the Swiss social security records (AHV/AVS) (Federal Social Insurance Office (FSIO), 2022), which record all employment relationships, include bonuses and overtime, and are uncapped. I winsorize them at the 5th and 95th percentiles and follow the cleaning of Martinez et al. (2021) (Appendix B.3). As the records end in December 2021, earnings cover deregistrations before October 2021, about 39% of job finders. The *market wage* is the median full-time-equivalent monthly wage in the found occupation × commuting zone, from the pooled 2020 and 2022 Swiss Earnings Structure Survey (ESS) (Swiss Federal Statistical Office (SFSO), 2026).[14] The *occupation value* ranks occupations by revealed preference. I apply the firm ranking of Sorkin (2018) to occupations: an occupation has a high value if workers who change jobs move into it more often than they leave it, and more so if they leave other high-value occupations to join it. The Swiss register records only moves through unemployment, which reflect job loss rather than preferences, so I compute the value from the US occupational mobility matrix of Schubert et al. (2024), about 53 million job-to-job switches between 2002 and 2015, crosswalked to 121 ISCO 3-digit occupations. This assumes that occupations rank similarly in the United States and in Switzerland; weighted by Swiss employment, the value has a correlation of 0.77 with the median occupational wage in Switzerland. I standardize it among job finders. *Commuting time* is the log driving time between the municipality of residence and the municipality of the new job, computed with OpenRouteService (Heidelberg Institute for Geoinformation Technology (HeiGIT), 2024) (Appendix B.4).

3.3 Descriptive statistics

Table 1 reports characteristics of the 366,890 unemployment spells in the regression sample, which correspond to 323,218 unique jobseekers; below I use "spells" and "jobseekers" interchangeably. About half of the jobseekers are women and about half are not Swiss; 26% have only primary education, 57% secondary or vocational education and 17% a university degree, and two thirds were employed in the year before registration. Within six months, 37% find employment. Mobility between markets seems to be the rule rather than the exception: 11% of jobseekers find a job in their initial market and 17% in an-

[14] Among job finders, 77.6% match to a cell with at least 30 ESS observations; the remaining 22.4% consist almost entirely of spells with missing occupation or region codes for the new job rather than thin ESS cells.

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	0.09	0.29	0.11	0.31
Found in initial market	0.11	0.31	0.11	0.31
Found in slacker market	0.08	0.27	0.09	0.29
Found, market unknown	0.09	0.28	0.08	0.28
N	366,890		101,511	

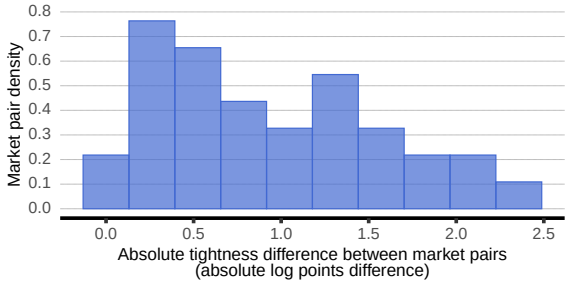
Notes: The table shows characteristics of the full regression sample and the subset of jobseekers who clicked on at least five job postings in the first three months of unemployment.

other market, 9% in a tighter and 8% in a slacker one; for the remaining 9%, the market of the new job is unknown.

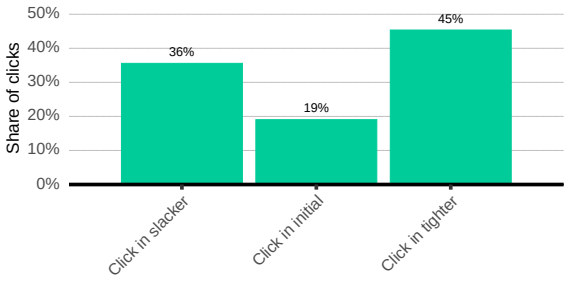
The right panel compares the full sample to the 101,511 jobseekers (28%) who clicked on at least five job postings in the first three months. Overall, clickers are similar to the full sample. They are more often female, better educated and less often non-Swiss, and they were somewhat less often employed in the year before registration. Job-finding rates are slightly higher among clickers and the decomposition into tighter markets, the initial market and slacker markets is similar across both groups. The comparison matters because the caseworker tightness tendency is measured from clickers but assigned to all jobseekers of that caseworker; the modest differences between the two groups suggest that a tendency estimated from clickers reflects 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.

Redirecting jobseekers toward tighter markets requires both that tightness varies substantially across markets and that their current search is not already strictly directed toward

Figure 2: 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 transitions from the last job before unemployment to the job found) and whose median market wages differ by less than 5%. Tightness is the average of daily log tightness over the sample period (June 2020 to December 2022), and pairs are restricted to markets with at least 30 unemployed jobseekers on average. Panel (b) shows the share of clicks in markets that are slacker than the jobseeker's initial market (last occupation × residence commuting zone), in the initial market and in markets that are tighter than the initial market, with the tightness of both markets measured at the date of the click. Sample restricted to jobseekers with at least 5 clicks in the first 3 months of their unemployment spell.

tight markets. Panel (a) of Figure 2 shows the distribution of tightness differences across market pairs that are plausibly substitutable from the jobseeker's perspective, with high historical transition rates and similar median wages (thresholds in the figure notes). Even among these "close" markets, the median absolute difference in log tightness is 0.85 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.[15]

The search behavior on the portal confirms the high mobility seen in the realized transitions, and it shows that search is not already strictly directed toward tight markets. The median clicker places 22 clicks (mean 41) in the first three months of their spell, spread across six three-digit occupations and two commuting zones. Panel (b) of Figure 2 shows where these clicks go relative to the jobseeker's initial market: 19% go to the initial market itself, 45% to markets that are tighter and 36% to markets that are slacker than the initial market at the time of the click.

[15] Appendix Figure E.4 shows that relaxing the wage restriction raises the median difference to 0.97 log points, a 2.6× ratio, while requiring only one past transition yields 1.10 log points, a 3.0× ratio.

4 Methodology

The goal of this paper is to estimate the causal effect of advice that redirects 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 and location. I address this selection by looking at caseworker advice: within local employment offices, jobseekers are assigned to caseworkers as good as randomly and caseworkers likely differ systematically in their tendency to redirect clients toward tighter markets (Section 2 and Arni & Schiprowski, 2019).

To see the intuition, 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 to redirect 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. Caseworkers act on many dimensions beyond search direction advice (application quality, training referrals, sanction enforcement), so the effect of the tightness tendency may partly capture correlated caseworker behaviors rather than search redirection per se. Section 6 addresses this concern directly.

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 3: 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.[16] 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

[16]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.

interacting this fixed effect with 2-digit last industry and with 2-digit last occupation of the jobseeker.

The second control, the *last occupation* × *commuting zone* × *registration quarter* fixed effect $\phi_{\text{occ}(i)\times\text{cz}(i)\times q(i)}$, controls for the feasible choice set of markets each jobseeker can plausibly target. Last occupation and commuting zone proxy for this choice set, and the registration quarter accounts for how it and its tightness change over time. Without this fixed effect, $\mu_{c(i)}$ might absorb differences in the clients' choice sets: a caseworker who, by chance, is assigned more clients from markets with tighter choice sets, or more clients who register during a boom, would look more redirection-oriented without giving different advice.

Estimating $\mu_{c(i)}$ directly and using the estimate for jobseeker i would mechanically correlate the measure with ε_i, because her own search would enter the estimate for her caseworker. Following Goldsmith-Pinkham et al. (2025), I therefore construct a 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 and the two sets of fixed effects. Subtracting the second term, the prediction from the fixed effects alone, recenters $Z_{-i,c(i)}$ within each office-year and choice-set cell, leaving only the caseworker's contribution.[17] As a robustness check, I leave out not only the focal jobseeker but also groups of clients who may be exposed to correlated shocks (Frandsen et al., Forthcoming). I restrict the sample to caseworkers with at least five clicking clients to ensure reliability of the leave-one-out averages.

4.1 Regression equation

The main estimating equation is:

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

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. Other sections use other horizons and 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).

[17] I compute the leave-one-out predictions from a single full-sample regression using the leverages, which I approximate with the Johnson–Lindenstrauss approximation of Kline et al. (2020), because the high-dimensional fixed effects make exact leverages infeasible (Appendix C).

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 × 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.

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 two 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, which increases statistical power.[18] The results also hold for non-clickers alone (Section 5.2.1). Second, 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. It does not directly measure what training caseworkers to provide tightness information, or showing tightness on job platforms, would achieve. To ease interpretation, I report magnitudes as the percentage change in the predicted outcome when moving from a 10th- to 90th-percentile caseworker tendency, relative to the prediction at the 10th percentile and averaging the predictions over the observed values of all other variables.

5 Results

I first show that the caseworker tendency redirects clients' own search toward tighter markets, check that it is balanced across predetermined jobseeker characteristics, and test whether the direction of the effect on own search is likely the same for all jobseekers. I then turn to the central question: whether it also speeds up reemployment. Finally, I examine heterogeneity by initial market tightness, the geographic and occupational margins of mobility, and effects on job quality.

[18] 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.

5.1 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. The regression uses the same specification as Equation (3), with the jobseeker's own search tightness index as the outcome. Figure 3, Panel (a), presents the results. A 10% higher caseworker leave-one-out score is associated with approximately 1.3% higher tightness of the jobseeker's own clicked markets, significant at all conventional levels (Column (1)). The 10th and 90th percentiles of the leave-one-out measure are 0.31 apart in log tightness, so moving from a p10 to a p90 caseworker raises the tightness of the clicked markets by 0.04 in log tightness, or about 4%. The effect is robust to the full set of individual controls (Column (2)), where the same move raises the tightness of the clicked markets by about 3.8%. Column (3) shows no effect on click intensity (clicks per day in unemployment 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. Column (4) shows that the caseworker's tendency also does not change whether a jobseeker clicks at least five times in the first three months, which defines the clicker sample.

To shed more light on the mechanism, I combine the exact timing of caseworker meetings with click timestamps. The idea is that if caseworker advice is the channel, effects on click tightness should materialize 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 E.5 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 stabilizes 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 E.2, 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.[19]

Figure 3, Panel (b), shows the results. Pre-meeting coefficients are close to zero but imprecisely estimated, since few clicks are observed before the first meeting. 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 3, Panel (a). The event study thus provides suggestive evidence that the change in search starts with the first meeting: before the meeting, the tightness of search does not differ detectably with the caseworker's tendency, although the estimates are imprecise. The divergence after the meeting is consistent with caseworkers changing the direction of search in their meetings.

A concern is that caseworkers could shape how their clients use the platform, so that clicks carry different information depending on the caseworker and the tightness of clicked markets mixes shifts in where jobseekers search with shifts in how they search. Appendix Table E.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.

5.1.1 Balance in covariates

Self-selection into search markets means jobseekers' predetermined characteristics should predict the tightness of the markets they click on. Appendix Table E.4 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 × year and initial occupation × commuting zone × quarter fixed effects as in the main specification from Equation (3).

The analysis documents substantial sorting in own click behavior. Nine of thirteen characteristics show statistically significant relationships at the 5% level after a multiple-testing

[19] The specification is:

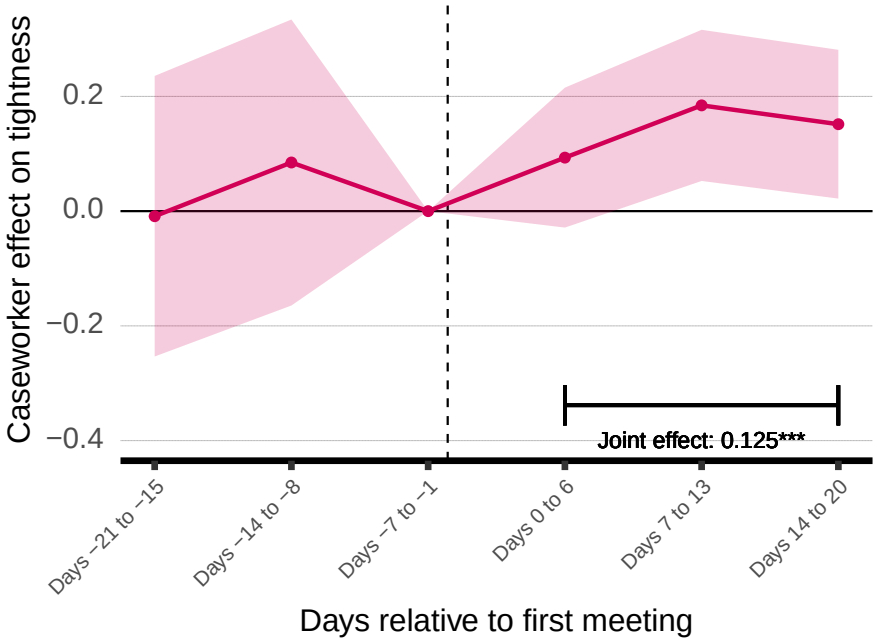
$$Y_{i\tau} = \sum_{k \neq -1} \beta_k \cdot \mathbf{1}[b(\tau) = k] + \sum_{k \neq -1} \gamma_k \cdot \mathbf{1}[b(\tau) = k] \times Z_{-i,c(i)} + \alpha_i + X'_{i\tau}\theta + \varepsilon_{i\tau} \tag{4}$$

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 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 windows.

Figure 3: Caseworker effects on job search direction: spell-level and meeting-level evidence

	Search direction		Search volume	
	(1)	(2)	(3)	(4)
Dependent Var.:	Own clicks tightness	Own clicks tightness	Clicks per day	5+ clicks
Leave-one-out tightness	0.1346***	0.1227***	-0.0125	0.0029
	(0.0213)	(0.0209)	(0.0352)	(0.0086)
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
Family	OLS	OLS	Poisson	OLS
N caseworkers	2,460	2,460	2,460	2,460
Observations	101,511	101,511	101,511	366,890
R2	0.71244	0.73797	–	0.10912
Mean of dependent var.	-0.57700	-0.57700	0.50365	0.27668

(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 of the OLS regressions in Columns (1) and (2), the number of clicks per day of unemployment in the first three months in Column (3), estimated by Poisson pseudo-maximum likelihood, and an indicator for at least five clicks in the first three months in Column (4). Columns (1)–(3) use jobseekers with at least 5 clicks in the first 3 months, Column (4) all jobseekers. 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 pre- and all post-meeting bins (see Appendix Table E.2, Column (2)). Standard errors are clustered at the caseworker level. Sample in Panel (b): jobseekers with at least three clicks in both the pre- and post-meeting windows. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

correction.[20] When looking at the relationship between the characteristics and the caseworker leave-one-out tightness tendency, the magnitudes collapse. Twelve of the thirteen characteristics are no longer distinguishable from zero. The five largest Panel A gradients shrink to small fractions of their original size.[21]

Only the female gradient remains significant after the correction: it is −1.8%, with an adjusted p-value below 0.001; the university-education gradient is 2.3%, with an adjusted p-value of 0.131. 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.

5.1.2 Monotonicity

I test whether caseworkers with a high tendency redirect all their clients toward tighter markets, not only on average, and whether the relationship goes in the same direction for all jobseekers. I use the two tests of Humlum et al. (2025), who adapt those of Bhuller et al. (2020), because their setting is the closest to mine. First, they also study caseworkers. Second, the caseworker does not decide alone what the jobseeker ends up doing. In many judge-leniency designs, the judge takes a decision, such as incarceration, and the data record this decision. In Humlum et al. (2025), the caseworker assigns jobseekers to training, but whether a jobseeker is assigned also depends on how willing the jobseeker is to take part. In my setting, the caseworker gives advice and the jobseeker decides whether to follow it. Random assignment still makes the caseworker tendency unrelated to the jobseeker's characteristics, but jobseekers do not follow the caseworker automatically.

I split clickers into quartiles of their predicted "treatment", in my case the search tightness index, based on the thirteen predetermined characteristics from Appendix Table E.4, and estimate the effect of the caseworker tendency on the search tightness index within each quartile.[22] Monotonicity implies a non-negative effect in every quartile. The first check uses the leave-one-out tendency estimated on the whole sample — the same as used for the main regressions. The second uses a "reverse-sample" tendency, computed only from the caseworker's clients in the other three quartiles.[23]

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

[21] For instance, the p10-p90 effect on the likelihood of being a woman drops from −34% to −1.8%. For being a woman with children, it drops from −45% to −1.7%. For holding a university degree, from 179% to 2.3%. For having only primary education, from −49% to −0.2%. For being over fifty, from −15% to −1.7%.

[22] The prediction regresses the search tightness index on the characteristics, conditioning on the office × year and initial occupation × commuting zone × quarter fixed effects, and uses only the part explained by the characteristics. The coefficients come from a sample split: each is estimated in a separate regression on one quartile, using the specification in Equation (3).

[23] The reverse-sample tendency uses the same formula as the leave-one-out tendency, estimated on the other three quartiles only, so no leave-out correction is needed.

Appendix Figure E.6 shows the results. The effect is largest in the two lowest quartiles, for jobseekers predicted to direct their search least toward tight markets, with 0.13 and 0.16 in both checks compared with 0.08 to 0.10 in the upper quartiles, although the effects are not significantly different from each other. This pattern resembles the finding of Belot et al. (2019) that online search broadening advice has the largest effects for jobseekers who searched narrowly at baseline. The effect is positive and significant at the 10% level in all quartiles and in both checks. The 95% confidence intervals rule out negative effects larger than 0.009 in every quartile, less than a tenth of the average effect.

5.2 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, the minimal set the design requires (Section 4). 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 3.0% over the baseline, or 1.1 percentage points. Column (2) adds the jobseeker controls of Section 4 and is my preferred specification. The estimated effect remains almost unchanged in magnitude, confirming that the caseworker allocation is orthogonal to any observed personal characteristics that influence job finding, and the standard errors are smaller.[24] Moving from a caseworker at the 10th percentile to one at the 90th percentile increases the probability of job finding by 2.8%, or 1.0 percentage point.

A back-of-the-envelope calculation shows that the effect on job finding is of the size one would expect from the change in search. For this calculation, I use data on applications and their success from Lalive et al. (2025). Their data show that jobseekers send 10.5 applications per month and 0.90% of applications lead to an offer. I assume that each vacancy makes exactly one offer and that the caseworker only changes application direction and not application quantity, as the absence of an effect on the number of clicks suggests. Given that, the probability that an application leads to an offer is proportional to the number of vacancies per unemployed in the markets the jobseeker searches in. Going from a 10th to a 90th percentile caseworker increases the tightness of these markets by 3.8% (Section 5.1) and would therefore increase the offer rate per application from 0.883% to

[24]The brackets in Columns (1), (2) and (5) report standard errors from a caseworker block bootstrap that re-estimates the leave-one-out tendency in every draw, so that they also reflect the estimation error in the tendency (Yap, 2024). In Columns (1) and (2), they are about 9% larger than the clustered standard errors, and the effects remain significant at the 5% level.

Table 2: Effect of caseworker redirection on job finding

	Find job within 6 months				
	In any market				In slacker market
	(1) Baseline	(2) Baseline	(3) Strict-assignment cantons	(4) Leave out same quarter	(5) Baseline
Leave-one-out tightness	0.0352***	0.0333***	0.0573***	0.0334***	0.0022
	(0.0099)	(0.0093)	(0.0170)	(0.0077)	(0.0049)
	[0.0108]	[0.0101]			[0.0052]
Unemployment office × year FE	X	X	X	X	X
Last occ. × commuting zone × quarter FE	X	X	X	X	X
Jobseeker controls		X	X	X	X
p10-p90 magnitude	+3.0%	+2.8%	+4.8%	+3.1%	+0.8%
N caseworkers	2,460	2,460	950	2,442	2,460
Observations	366,890	366,890	128,995	366,249	335,701
R2	0.08198	0.14161	0.16126	0.14168	0.11587
Mean of dependent var.	0.36554	0.36554	0.35414	0.36561	0.08815

Notes: The table reports linear probability model estimates of the effect of the caseworker's tendency to redirect search toward tight markets, as described in Equation (3). The outcome variable indicates whether the jobseeker finds employment within 6 months: in Columns (1)–(4) in any market, in Column (5) in a market that is, at the deregistration date, slacker than the jobseeker's initial market; Column (5) leaves out job finders whose market is not observed or has no tightness measure. Column (3) restricts the sample to cantons that assign new jobseekers in proportion to contractual hours or by software-based random allocation. Column (4) replaces the leave-one-out tendency with a leave-out that excludes all of the caseworker's clients registering in the same quarter. Jobseeker controls include gender, age, children, nationality, education, tenure in last occupation, employment and earnings in the last two years, potential benefit duration, detailed last occupation FE, employment office × weekday of first meeting FE, and municipality FE. The table also reports the effect of going from the 10th to the 90th percentile of the caseworker tendency, in percent of the predicted outcome at the 10th percentile. Standard errors clustered at the caseworker level in parentheses. Standard errors from a caseworker-clustered bootstrap (100 draws) that re-estimates the leave-one-out tendency in every draw in brackets (Columns (1), (2) and (5)). Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

0.917%, assuming a symmetric increase around the average rate. With 10.5 applications per month, the probability of at least one offer in a month rises from 8.9% to 9.2%. The probability of at least one offer within six months rises from 42.8% to 44.0%, by 2.8%, exactly the same as the estimated effect on job finding.[25] Given the simple assumptions, the exact match might partly be a coincidence, but the calculation suggests that the estimate is of a plausible size.

Columns (3) and (4) address two main concerns about the design. Column (3) restricts

[25]With an offer rate p per application and independent applications, the probability of at least one offer among n applications is $1-(1-p)^n$, with $n = 10.5$ for one month and $n = 6 \times 10.5 = 63$ for six months. Note that jobseekers who stop applying after receiving an offer would not change this probability, as it is computed as the complement of the probability that all 63 applications fail.

the sample to the seven cantons that assign new clients in proportion to contractual hours or by software-based random allocation. It leaves out the cantons that assign new clients to the next free caseworker or at the discretion of management, where assignment could in principle depend on the caseworker or the jobseeker: caseworkers whose clients leave unemployment faster could receive more new clients, and managers could send certain jobseekers to certain caseworkers (Section 2).[26] The coefficients and magnitudes are larger than in the baseline, with a magnitude of 4.8% against 2.8%, but not significantly different from it.

If many jobseekers arrive at 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). Column (4) of Table 2 shows that results remain robust. A related concern applies within occupation: although the occupation × commuting zone × 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 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, with a magnitude of 2.7% (Column (4) of Appendix Table E.5, Panel A).

Column (5) changes the outcome to finding a job within six months in a market that is slacker than the jobseeker's initial market, the last occupation × commuting zone of residence, with tightness compared at the exit date. The coefficient is 0.0022, a magnitude of 0.8%, and not statistically significant.[27] While not a formal placebo test, this result provides a 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.[28] The absence of an effect there is consistent with advice that redirects search toward tight markets rather than with such across-the-board changes. Section 5.3 shows where the gains occur instead, by the tightness of the initial market, and Section 6 tests alternative mechanisms such as search intensity, support behaviors, and sanction rates.

[26] Column (3) also leaves out the cantons that did not answer the survey.

[27] The effect in slacker markets is also significantly smaller than the effect on job finding in any market. On the spells of Column (5), job finding in any market rises by 3.8%, against 0.8% in slacker markets, and the difference is significant at the 1% level in a one-sided test.

[28] This is the case for the two other caseworker actions that raise job finding: caseworkers with a higher tendency to assign training courses or to sanction also raise job finding in slacker markets, by 6.9% and 10.2% (Section 6.4).

Appendix Table E.6 extends the perspective to other horizons by estimating the effect on cumulative job finding over the first year of the spell. The effects after 30 and 60 days are not statistically significant. The gap then widens steadily and becomes significant at 90 days; relative to the baseline probability, the effect peaks around 150 days at 3.1%. At the six-month benchmark, the effect is 2.8%. The relative effect declines somewhat at longer horizons but remains economically meaningful and highly significant at one year, at 2.4% over the baseline probability. The persistence of the effect suggests that caseworker redirection does not merely accelerate job finding by a few weeks; it leads to reemployment that might not have occurred within the first year in the absence of the treatment. Appendix Table E.7 provides a period-by-period hazard rate analysis using period-specific complementary log-log models.[29] 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.

5.2.1 Robustness

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 Column (5) of Appendix Table E.5, Panel A, I restrict the sample to jobseekers with no clicks on job-room.ch in the first three months 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.[30]

Another 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 (2) and (3) of Appendix Table E.5, Panel A, interact the office × year fixed

[29]Note that this analysis is primarily descriptive as the hazard in period t is estimated conditional on not having found a job by $t-1$, which itself is an outcome. The models include the fixed effects of Column (1) of Table 2 but not the jobseeker controls, since the nonlinear model with many fixed effects struggles to converge.

[30]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 do not enter $Z_{-i,c(i)}$. Restricting to non-clickers eliminates the overlap entirely. The fact that the magnitude in Column (5) of Appendix Table E.5, Panel A, is very similar to the baseline indicates that the result is not an artifact of within-sample dependence between the constructed regressor and the outcome equation.

effect with the 2-digit NOGA industry of the last job in Column (2) and with the 2-digit ISCO last occupation in Column (3). The magnitudes are 2.8% and 2.1% respectively, equal to and slightly smaller than the baseline magnitude of 2.8%, very similar overall. Together with Column (3) of Table 2, this suggests that the results are not driven by a feature of the caseworker assignment process.

Columns (1)–(3) of Appendix Table E.5, Panel B, test robustness to different measures of market tightness. Column (1) uses 4-digit ISCO codes instead of 3-digit for finer occupational granularity, showing a magnitude of 2.0%. Column (2) uses 101 more granular labor market regions, the Swiss *Arbeitsmarktregion*, instead of 16 commuting zones, with a magnitude of 2.2%. Column (3) allows the denominator of tightness to reflect actual search behavior using data from the platform, defining tightness as the number of ads per jobseeker with at least one click in the market, finding a magnitude of 1.9%. The finer market definitions likely introduce more measurement error in tightness, attenuating the estimates. The results remain significant across all alternative definitions.

Column (4) of Appendix Table E.5, Panel B, 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 2.6%, while those who steer toward markets with more competing jobseekers reduce it by 2.1%. Both coefficients are individually significant. There is the caveat though that this is not a conclusive decomposition of the effect of the tightness tendency: 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 tendency toward markets with more vacancies and the tendency toward markets with fewer competing jobseekers raise job finding. 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 caseworkers who steer search away from markets with many competing jobseekers also raise job finding.

A natural concern is that the July 2020 to September 2022 sample period spans parts of the COVID-19 recovery. Columns (5) and (6) of Appendix Table E.5, Panel B, re-estimate 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 2.9% before and 2.7% after the cutoff, very close to the pooled baseline.

The estimator requires two sample restrictions: a minimum number of clicks per jobseeker to measure search direction and a minimum number of clicking clients per caseworker

to estimate the leave-one-out tendency. The baseline requires at least five clicks per jobseeker and at least five clicking clients per caseworker. Appendix Table E.8 varies the minimum number of clicks per jobseeker over 1, 5, 10 and 30 and the minimum number of clicking clients per caseworker over 2, 5, 10 and 30. With minimums of up to 10 clicks per jobseeker and up to 10 clicking clients per caseworker, the magnitudes range from 2.6% to 3.0%. With a minimum of 30 clicks per jobseeker, few clicking clients remain per caseworker and the effect on job search is not statistically significant, so these combinations are not informative. With a minimum of 30 clicking clients per caseworker and the baseline minimum of five clicks per jobseeker, however, the effect on job finding is lower and no longer statistically significant, while the effect on job search stays close to the baseline. This suggests heterogeneity in the effect on job finding across caseworkers. Compared with caseworkers with fewer clicking clients, caseworkers with 30 or more have more clients in general and less experience. They also use other instruments differently: they sanction their clients and assign them to training more often and deem fewer monthly search effort protocols sufficient. These differences also hold within employment office (Appendix Table E.9). The effect of the tightness tendency does not vary with the caseworker's caseload or experience (Appendix Table E.10), but it does vary with a caseworker's tendency to use other instruments: Section 6.5 shows that it is larger among caseworkers with a tendency to meet more and are more likely to assign their clients to training and it is much smaller among those with a tendency to sanction more. When I interact the tendency with all five instruments that Section 6.4 looks at, the difference between caseworkers with 30 or more and with fewer clicking clients is no longer statistically significant (Appendix Table E.11, Columns (2) and (4)). The two groups may also differ in respects that I do not observe.

The baseline measures the search tightness index over the first three months of unemployment. Appendix Table E.12 varies this window across 30, 60, 90, 120, and 180 days. Magnitudes are between 2.2% and 3.4% 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.

5.3 Heterogeneity by baseline tightness

Figure 4 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.[31] 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

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

as the initial one, exits to tighter markets relative to the initial one, and exits to slacker markets. 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, if attention is limited, too broad a search can also hurt: jobseekers in tight initial markets may miss opportunities there. For them, advice that redirects search toward tight markets means focusing on the initial market, which narrows their search and raises job finding there.

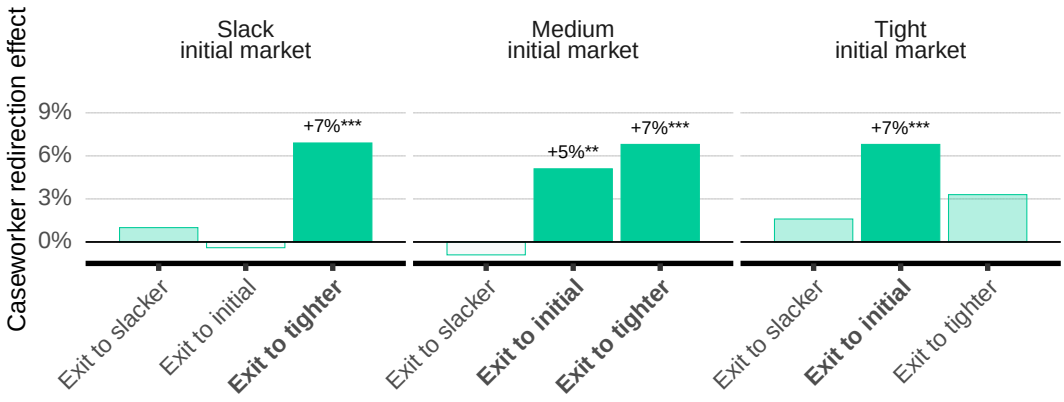
Before examining the specific channels, Column (1) of Appendix Table E.13 establishes that the overall job-finding effect is positive and similar across all terciles of initial market tightness. If anything, the effect is larger for those in medium or tight initial markets, though the difference is not statistically significant. Figure 4 decomposes this effect by the market in which reemployment occurs.

Figure 4 reports the percentage increase in six-month job finding when moving from the 10th to the 90th percentile of the caseworker leave-one-out measure, decomposed by initial market and market of the new job. 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 7%. 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 exit rates to tighter markets that increase by about 7%. 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 a 7% effect.[32]

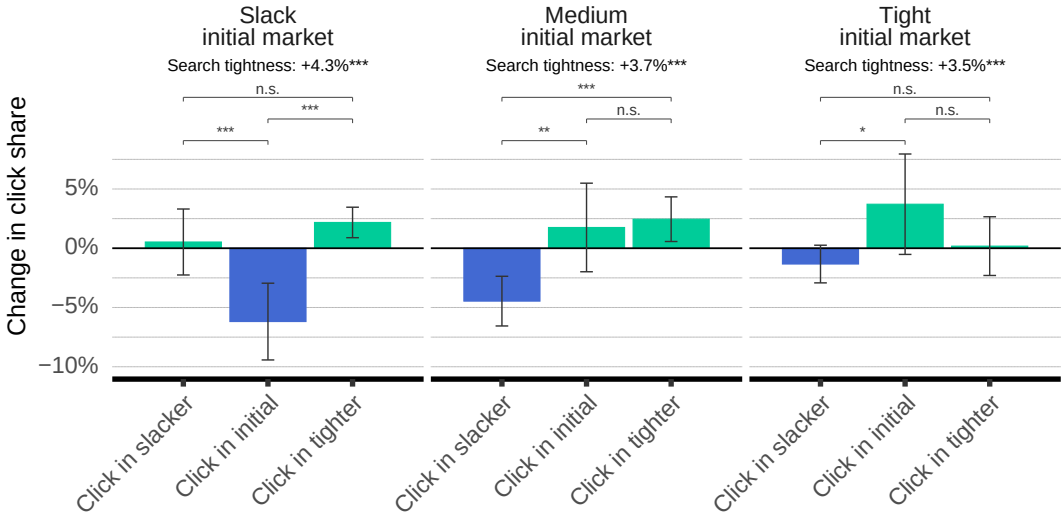
As further evidence, I use the clicks of jobseekers who actively use the job-room portal for their job search, defined as those with at least five clicks in the first three months of their spell, to observe how caseworker redirection shapes the targeting of different markets before employment outcomes are realized. I classify each click by whether it falls in a market that is tighter than, the same as, or slacker than the initial market, the same three categories as for the reemployment outcomes. I estimate a multinomial

[32]These patterns are not mechanical floor or ceiling effects of few slacker markets being available to jobseekers in slack initial markets and few tighter markets to those in tight initial markets. First, substantial shares of job finders from slack initial markets still find a job in a slacker market (14%), and of those from tight initial markets in a tighter market (21%). Second, for jobseekers in medium initial markets nothing is mechanical (25% of job finders find a job in a slacker and 36% in a tighter market), and the effects show the same pattern.

Figure 4: Heterogeneous effects of caseworker redirection on job finding and clicks by initial market tightness



(a) Job finding within 6 months



(b) Click shares by market type

Notes: Panel (a) 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. It 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 percentage effect of going from the 10th to the 90th percentile in the caseworker leave-one-out distribution. The sample leaves out job finders whose market is not observed or has no tightness measure. The underlying regressions are reported in Appendix Table E.13. Panel (b) plots the percent change in the predicted share of clicks in each market type when moving from the 10th to the 90th percentile. The shares come from a multinomial logit over the three market types, estimated separately by tercile on jobseekers with at least five clicks in the first three months. The error bars show 95% confidence intervals of each change; the brackets show pairwise tests on the logit coefficients. The text in each panel reports the effect on average search tightness. The underlying regressions, including the search-tightness effects in Column (5), are reported in Appendix Table E.14. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

logit over the three market types, in which the caseworker leave-one-out tendency shifts the probability of each type.[33] The alternative-specific constants vary by office × year, by last occupation × commuting zone × quarter and by the jobseeker controls, so the comparison is the same as in Equation (3). I estimate the model separately for each tercile of initial market tightness. Because the model conditions on the number of clicks, it describes where clicks go, not how many there are.[34]

Figure 4, Panel (b), shows how the advice changes where clicks go. Each bar is the change in the predicted share of clicks in one market type when moving from a p10 to a p90 caseworker, in percent of the share at the 10th percentile. The advice raises average search tightness — the outcome of Column (2) of the table in Figure 3 — in all three terciles, by 4.3%, 3.7% and 3.5% for jobseekers in slack, medium and tight initial markets, but the clicks move in different directions. Jobseekers in slack initial markets move clicks away from their initial market: its share falls by 6.2%, and the share of tighter markets rises by 2.2%, both significant at the 1% level. Jobseekers in medium initial markets move clicks from slacker to tighter markets: the share of slacker markets falls by 4.5% and the share of tighter markets rises by 2.5%, both significant at the 1% level, while the share of the initial market does not change significantly. Jobseekers in tight initial markets move clicks from slacker markets to their initial market: the share of the initial market rises by 3.7%, significant at the 10% level.[35] The click estimates show the direction in which search moves in each tercile. They are less precise than the job-finding estimates: for jobseekers in tight initial markets, the 95% confidence interval of the change in the share of the initial market ranges from about −0.5% to 8%. This design asks a lot of the data: it estimates a three-by-three matrix of effects, three market types in three terciles, on the clicker sample only. Furthermore, these shares capture only moves between three broad groups of markets. The advice can also move search within a group, for example toward the tightest of the tighter markets.[36]

Taken together, the job-finding results suggest that caseworker advice operates through two complementary channels: redirecting search toward tighter markets and reinforcing

[33] I estimate the multinomial logit model through its Poisson equivalence: a Poisson regression of the click counts by spell and market type with spell fixed effects yields the multinomial logit coefficients (Baker, 1994; Guimaraes, 2004).

[34] Appendix Table E.15 confirms that the leave-one-out tendency has no effect on the number of clicks within any initial-market tightness tercile.

[35] These patterns, too, are likely not mechanical floor or ceiling effects of few slacker markets being available to jobseekers in slack initial markets and few tighter markets to those in tight initial markets. First, substantial shares of clicks still go to slacker markets among jobseekers in slack initial markets (19%) and to tighter markets among those in tight initial markets (27%). Second, for jobseekers in medium initial markets nothing is mechanical (34% of clicks go to slacker and 45% to tighter markets), and the effects show the same pattern.

[36] If only the shares of the three groups changed as in Panel (b), and search within each group stayed as it is, search tightness would rise by about half of the average effect found in Column (2) of the table in Figure 3. The rest comes from moves within the groups. Section 7 describes the computation.

search in already favorable ones. The click results point in the same direction in each tercile, but they are too imprecise to confirm this on their own. The finding distinguishes the mechanism from occupational mobility treatments, which typically emphasize moving workers out of declining sectors, and suggests that such interventions might not be the right approach in every situation.

As a robustness check, Appendix Figure E.7 defines the initial market with the commuting zone of the last job instead of the commuting zone of residence. The pattern across the three terciles is the same.

Lastly, for 23.3% of six-month job finders (8.5% of spells), the found market is not observed or has no tightness measure. The analysis so far leaves these spells out. Appendix Table E.5, Panel A, Column (6), instead treats finding a job in such a market as a separate outcome and shows that the caseworker tendency has no statistically significant effect on it. Furthermore, Appendix Figure E.8 instead imputes the found market of these job finders from the industry of the new employer, which is recorded on the register.[37] The pattern across the three terciles is unchanged.

5.4 Geographic versus occupational mobility

Appendix Figure E.9 splits job finding within six months into four mutually exclusive categories: whether the new job is in the jobseeker's last occupation and whether it is in the commuting zone of residence at registration. Assignment to a high-redirection caseworker increases the probability of finding a job in the same occupation and same commuting zone by 4.2%, in a different occupation and a different commuting zone by 7.9%, and in a different occupation in the same commuting zone by 3.6%. There is no effect on finding a job in the same occupation in a different commuting zone. The relevant margin is thus predominantly joint: the occupation and the commuting zone either both stay the same or both change, and purely geographic mobility is never the margin. When looking at heterogeneity by initial market, the pattern follows Section 5.3: in slack markets the effect comes almost exclusively from moves out of both the occupation and the commuting zone, in tight markets from staying in both.

5.5 Effects on job quality

The faster job finding could come at the cost of job quality. For instance, Dromundo Mokrani & Haramboure (2022) find that the counselors in the French public employment service

[37] I use the empirical joint distribution of occupations and industries within each region to weight the candidate cells, classify each candidate cell as the initial, a tighter or a slacker market, and split the exit across these categories in proportion to the weights. Appendix B.6 details the construction. The imputation leaves 1.0% of job finders with an unknown market.

who foster fast exits from unemployment differ from those who foster stable jobs. For advice on where to search, a particular concern is that some markets may be tight because nobody wants to work there, because wages are low or working conditions are hard.

Table 3 splits the effect on job finding within six months by the type of job found. Each outcome indicates that the jobseeker finds a job within six months and that the job is of the given type; as in Section 5.3, jobseekers who have not found a job within six months are coded as zero. I look at three dimensions of job quality: stability, whether the jobseeker stays out of unemployment for 360 days after finding the job; the market wage, whether the job is in a market with a higher or a lower median wage than the initial market; and the occupation value, whether the job is in an occupation with a higher or a lower value than the last occupation. The occupation value is a revealed preference measure based on job-to-job transitions that captures also non-monetary utility aspects of occupations.[38] For each dimension, I estimate the effect on finding a better and on finding a worse job. If the advice led jobseekers into worse jobs, it would raise the probability of exit in the worse category more than the probability of exit in the better one. Such unconditional probabilities are an attractive outcome because they are defined for all jobseekers. Job quality conditional on finding a job is not: if the advice helps additional jobseekers find a job, the job finders of high- and low-tendency caseworkers differ, and differences in their job quality could reflect who finds a job rather than the jobs themselves.

The advice raises the probability of finding the better job in each dimension, but not the worse one. Moving from a caseworker at the 10th to one at the 90th percentile of the tendency raises the probability of finding a stable job by 2.7%, significant at the 5% level, a job in a market with a higher wage than the initial market by 6.7% and a job in an occupation with a higher value than the last occupation by 7.1%, both significant at the 1% level. The effects on the worse job in each dimension are all not significantly different from zero: jobseekers assigned to high-tightness tendency caseworkers find those jobs at the same rate as those assigned to low-tendency caseworkers. Within each dimension the two effects differ significantly, at the 10% level for stability, the 1% level for wages and the 5% level for occupation values.[39] Jobs in the initial market are the residual category, where the market utility stays constant: the advice raises the probability of finding a job in the initial market by 4.3%, significant at the 1% level (Column (7)). Imputing the market of job finders whose market is not observed from the industry of the new employer shows the same pattern (Appendix Table E.17).

[38] I apply the firm ranking of Sorkin (2018) to occupations: an occupation has a high value if workers who change jobs move into it more often than they leave it. Section 3 describes its construction.

[39] The same outcomes show a different pattern for another caseworker action: the caseworker's tendency to sanction raises the probability of finding both better and worse jobs, including jobs in markets with a lower wage and in occupations with a lower value (Section 6.4).

Table 3: Effect of caseworker redirection on job finding by type of job

	Stable 360 days		Market wage		Occupation value		Same market
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dependent Var.:	Stable	Not stable	Higher	Lower	Higher	Lower	Same market
Leave-one-out tightness	0.0223**	0.0046	0.0210***	0.0006	0.0155***	0.0057	0.0162***
	(0.0089)	(0.0048)	(0.0054)	(0.0052)	(0.0044)	(0.0043)	(0.0061)
Unemployment office x Year FE	X	X	X	X	X	X	X
Last occ. x commuting zone x quarter FE	X	X	X	X	X	X	X
Jobseeker controls	X	X	X	X	X	X	X
p-value equal within pair	0.072	0.072	0.001	0.001	0.050	0.050	
p10-p90 magnitude	+2.7%	+1.7%	+6.7%	+0.2%	+7.1%	+2.5%	+4.3%
N caseworkers	2,404	2,404	2,460	2,460	2,460	2,460	2,460
Observations	222,257	222,257	335,701	335,701	335,701	335,701	335,701
R2	0.15389	0.08243	0.09394	0.10455	0.09165	0.10788	0.10408
Mean of dependent var.	0.25960	0.08475	0.09864	0.09006	0.06906	0.06989	0.11732

Notes: The table reports linear probability model estimates from the specification in Equation (3). Each outcome variable indicates that the jobseeker finds a job within six months and that the job is of the given type; jobseekers without a job within six months, or with a job of another type, are coded as zero. Columns (1) and (2): the jobseeker does not return or returns to unemployment within 360 days of reemployment; the sample is restricted to spells registered by 7 October 2021, so that the 360 days after a job found within six months are observed. Columns (3) and (4): the market wage of the new job, the median monthly full-time equivalent wage in the found occupation × commuting zone, is higher or lower than that of the initial market. Columns (5) and (6): the value of the found occupation is higher or lower than that of the last occupation. Column (7): the job is in the initial market (the residual category for (3) or (4) and (5) or (6)). Columns (3)–(7) leave out job finders whose market is not observed or has no tightness measure (8.5% of the spells), as in Figure 4. Market wages and occupation values are defined in Section 3. The p-value equal within pair tests whether the effects on the two outcomes of a pair are equal. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Panel A of Appendix Table E.18 estimates the effects by tercile of initial market tightness. For none of the six quality outcomes do the effects differ significantly across terciles. For jobseekers who start in slack markets, however, the advice raises job finding in markets with a higher and in markets with a lower wage than the initial market by similar proportions, 5.5% and 6.2%, the former significant at the 5% level and the latter only at the 10% level, and the difference between the two effects is not statistically significant. For these jobseekers, the advice might increase exits to higher- and to lower-wage markets in about equal proportion. In terms of occupation value, the advice still raises finding a job in a higher-value occupation more than in a lower-value one for these jobseekers, by 8.0% against 2.6%, a difference significant at the 10% level.

Having established that the advice raises the probability of finding a better job but not that of finding a worse one, I turn to the jobs found to gauge how large the differences in job quality are. Conditional on finding a job, the differences can be expressed in natural units, such as wages or commuting time. Panel B of Appendix Table E.18 shows that they are small. Moving from a caseworker at the 10th to one at the 90th percentile of the tendency raises the market wage of the new job by 0.3%, significant at the 1% level, and the value of the new occupation by 0.015 standard deviations, significant at the 1% level. Panel B also reports two outcomes that are not observed for the previous job, so that there is no better-or-worse indicator for them: the jobseeker's own monthly earnings in the new job, from social security records, and the commuting time, the driving time from the place of residence to the new job. The effect on monthly earnings is also 0.3% but imprecise, since the records cover only deregistrations before October 2021. The probability of returning to unemployment within 180 or 360 days and the commuting time do not change significantly. This is consistent with the results of Section 5.4: the advice raises job finding in other commuting zones about in proportion to job finding overall, so that the share of jobs in another commuting zone among job finders hardly changes. The advice thus mainly speeds up job finding: it does not lead jobseekers into markets that are tight because nobody wants to work there, and the jobs they find are slightly better.

6 Caseworker mechanisms

This section examines whether caseworkers with a high tendency also tend to use other actions differently and whether they differ in observable characteristics. It also investigates whether the effects of the tightness redirection are heterogeneous across caseworkers who differ in other aspects.

6.1 Vacancy referrals

Besides advising jobseekers on where to search, caseworkers can refer them to specific vacancies. The tightness tendency could therefore capture caseworkers who know employers in tight markets and send their clients to them, rather than caseworkers who change where their clients search. The data do not record the exact vacancy a referral concerns, but I observe the number and timing of referrals and whether a referral leads to a job. Among clickers, 35% receive at least one referral by their caseworker, and the average spell has 1.3 referrals.

The most direct test uses whether a referral led to a job. If high-tendency caseworkers raised job finding by sending their clients to employers in tight markets, the tendency would raise the probability of finding a job through a referral. It does not. While the tendency raises job finding overall, its effect on finding a job through a referral is close to zero and statistically insignificant (Column (4) of Appendix Table E.19). The timing of the change in search points the same way. Appendix Figure E.10 repeats the event study of Figure 3, Panel (b), around the first referral instead of the first meeting, controlling for the time since the first meeting. After the first referral, jobseekers click more (Column (1), Appendix Table E.19), but the tightness of the markets they click does not change with the caseworker's tendency (Appendix Figure E.10).[40] The change in the tightness of clicked markets appears after the first meeting, but not after the first referral.

A related explanation is that high-tendency caseworkers help their clients not only with where to search but also with how to search in tight markets, for example through contacts with specific firms or advice on how to apply in these markets. Such knowledge accumulates with experience, so in this case I would expect the effect to increase with caseworker experience. The interaction of the tendency with caseworker experience is small and insignificant, both for the search tightness index and for job finding within six months (Appendix Table E.10; Section 6.5). This is suggestive evidence that the effect does not rely on knowledge of specific employers or application practices that caseworkers build up over time.

6.2 Variation in tightness

Tightness differs across markets because some occupations are tight everywhere and at all times, and because some markets are tight in a particular place or at a particular time. To

[40] To have an exact comparison, Column (3) applies the exact specification of the event study around the first referral to the first meeting. It is more demanding than the specification behind Figure 3, Panel (b): it additionally controls for the time since the first referral, and it measures time in calendar days rather than hours, since referrals are recorded by date only. The estimate is 0.134, close to the pooled estimate of 0.125 for Figure 3, Panel (b).

see along which of these differences the advice moves jobseekers, I split the log tightness of each clicked market into four parts that add up to the total: a national persistent part, the average tightness of the occupation across commuting zones and days; a local persistent part, the average deviation of the occupation in the commuting zone from this level; a national transitory part, the deviation of the occupation's national tightness on the day of the click from its average; and a local transitory part, the remainder. I then estimate Equation (3) with each part, averaged over the clicks of a spell, as the outcome (Appendix Table E.20).

Of the effect of 0.123 on search tightness, 80% comes from the national persistent part, 16% from the local persistent part and 5% from the national transitory part, all significant at the 5% level; the local transitory part contributes nothing. This large share of the national persistent part could simply reflect where tightness differs most. To compare, I compute the differences in search tightness among jobseekers who start in the same market in the same quarter. 77% come from the national persistent part, and 17%, 3% and 2% from the local persistent, national transitory and local transitory parts. The effect of the advice is thus split in almost the same proportions, which suggests that the advice moves jobseekers along the differences in tightness they face rather than toward a particular kind of tightness.

6.3 Who redirects search

What distinguishes caseworkers who direct their clients' search toward tighter markets? I do not observe background information on caseworkers such as personal characteristics or their own occupational experience (Behncke et al., 2010; Huber et al., 2017), but I can infer several characteristics from their interactions with jobseekers: caseload, experience, and the occupational and industry diversity of their client pools. Appendix Table E.21 regresses each characteristic on the tendency using the baseline specification. None differs substantially between high- and low-redirection caseworkers. The only coefficient significant at the 5% level, for the number of industries served, is economically small and not significant after adjusting for multiple hypothesis testing. The tendency thus appears largely idiosyncratic to individual caseworkers. This is in line with Cederlöf et al. (2025), who find that most observable caseworker characteristics are unrelated to caseworker value-added.

6.4 Advice and other caseworker actions

Caseworkers also differ in how they support and monitor their clients. I measure three support dimensions, the number of meetings, meeting duration and training course par-

ticipation, and two stringency dimensions, monitoring leniency and benefit sanctions, with the same leave-one-out method as the tightness tendency (Section 4).[41]

Figure 5 shows the correlations between these caseworker measures, conditional on unemployment office × year and last occupation × commuting zone × quarter. The support and stringency dimensions are correlated with each other, while the tightness tendency is nearly uncorrelated with all of them, with correlations of at most 0.04 in absolute value. Sampling noise in the caseworker measures can bias these correlations. Regressions of the jobseeker's own outcomes, for example whether she was sanctioned or how often she met her caseworker, on the leave-one-out tightness tendency are not affected by this noise, because her own outcome is independent of the noise in the leave-one-out measure, which comes from the other clients (Kline et al., 2020). Appendix Table E.22 reports these regressions. Only the number of meetings differs at the 5% level: jobseekers assigned to a caseworker at the 90th instead of the 10th percentile of the tightness tendency have 0.8% fewer meetings. Since more meetings raise job finding (Schiprowski, 2020), this difference would, if anything, bias the effect of the tightness tendency on job finding downward. Meeting duration, training course participation, monitoring leniency and sanctions differ by at most 1.0%, and none of these differences is statistically significant. This is consistent with the institutional setting (Section 2): advice on where to search is part of the meeting content, over which caseworkers report the most freedom, whereas they report less freedom over program referrals, meeting frequency and sanctions.

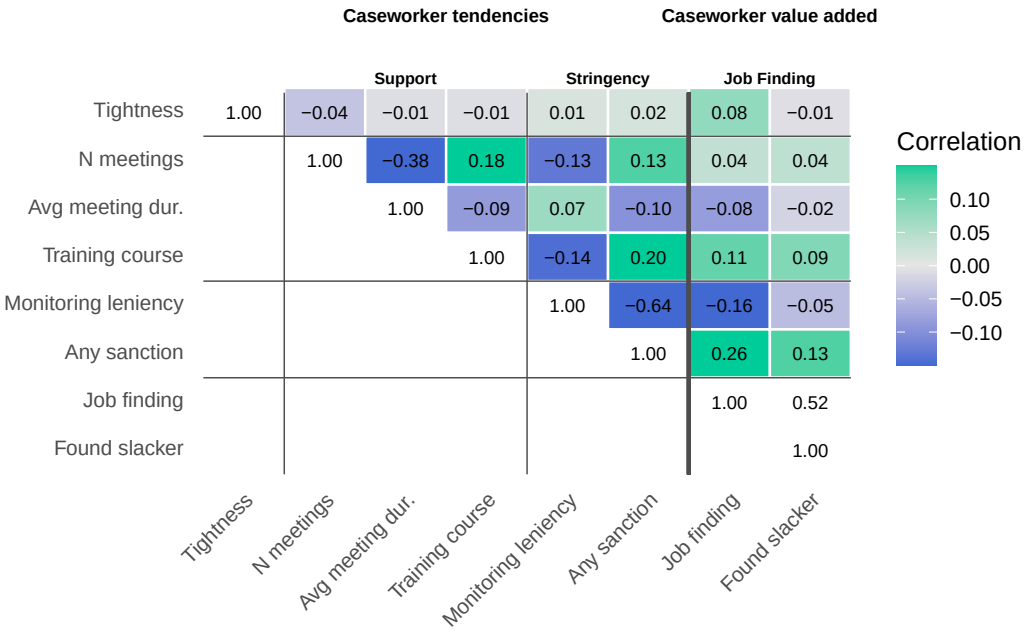
A joint regression tests this directly. I estimate Equation (3) with all six leave-one-out measures as regressors. Figure 6 reports the effect of moving each measure from its 10th to its 90th percentile on job finding within six months, from the joint regression and, for the tightness tendency, from the baseline regression on the same sample. Controlling for all other caseworker dimensions leaves the effect of the tightness tendency on job finding virtually unchanged (Panel (a)). The null effect on exit in slacker markets (Panel (b)) is also unchanged.

The joint regression also shows that among the investigated dimensions only the tightness tendency raises job finding without raising job finding in slacker markets. Training courses and sanctions, the two other dimensions with significant effects on job finding, raise job finding in slacker markets by 6.9% and 10.2% (Panel (b))[42]. If the tightness tendency

[41] The number of meetings is the mean monthly rate in the first 90 days and meeting duration the average log hours per meeting. Training course participation pools basic courses, language courses, computer courses and occupation-specific courses (see Gerfin & Lechner, 2002, Table A1). Monitoring leniency is the probability with which the caseworker deems the jobseeker's monthly search effort protocol sufficient. 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. The rating of the protocols, however, is available throughout. Sanctions are measured by the probability of a benefit sanction in the first three months. For these five measures, the Johnson–Lindenstrauss approximation of the leverages uses 500 instead of 1,000 random projections (Appendix C).

[42] Appendix Figure E.11 shows a similar comparison for one-year job finding.

Figure 5: Correlation matrix of caseworker dimensions

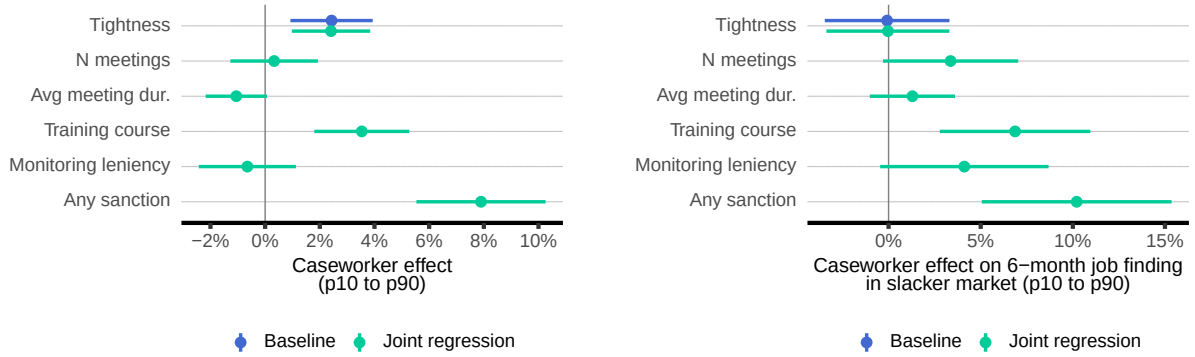


Notes: The figure shows pairwise correlations between caseworker measures. Each measure is the mean outcome of the caseworker's clients, residualized on unemployment office × year and last occupation × commuting zone × quarter fixed effects. "Tightness" uses the log tightness of the clicked markets; "Job finding" and "Found slacker" use job finding within six months in any market and in a slacker market, the latter leaving out job finders whose market is not observed or has no tightness measure. Correlations are computed at the unemployment-spell level.

captured general caseworker quality, such as motivation or help with applications, or caseworker pressure, I would expect it to raise job finding in slacker markets as well, as training and sanctions do.

The dimensions also differ in the type of job they lead to. I estimate the joint regression with indicators for finding a job of a given type within six months as outcomes (Appendix Figure E.12). Sanctions, training courses and more frequent meetings raise the probability of finding both better and worse jobs: jobs in markets with a higher but also with a lower median wage than the initial market, or in occupations with a higher or a lower value than the last occupation. The tightness tendency raises only the probability of finding a job in the better one of these categories. Taken together, the results compare advice with sanctions, the instrument with the largest effect on job finding, on two counts: how much it raises job finding and which jobs it leads to. The effect of the tightness tendency on job finding is about 30% of the effect of the sanction measure. Advice on where to search is a less severe intervention than a sanction, which suspends benefit payments for up to 60 working days (Arni et al., 2013), and unlike a sanction it does not raise job finding in worse jobs.

Figure 6: Effect of the tightness tendency on job finding, controlling for other caseworker tendencies



(a) Job finding in any market

(b) Job finding in a slacker market

Notes: The figure reports the effect of moving each leave-one-out caseworker measure from its 10th to its 90th percentile, in percent of the predicted outcome at the 10th percentile, with 95% confidence intervals. "Joint regression" estimates Equation (3) with all six measures as regressors (Appendix Table E.23, Columns (3) and (6)). "Baseline" includes only the tightness tendency, on the same sample (Columns (1) and (2)). Panel (a): finding a job within six months. Panel (b): finding a job within six months in a market slacker than the initial market; jobseekers without a job, or with a job in another market type, are coded as zero, and job finders whose market is not observed or has no tightness measure are left out. N meetings is the mean monthly number of meetings in the first 90 days, meeting duration the average log meeting duration, training course participation the participation in a training course, monitoring leniency the share of monthly search effort protocols deemed sufficient, and any sanction a benefit sanction in the first three months; each is the leave-one-out caseworker measure of the respective behavior. Sample: spells with all six measures defined (338,885 spells, 2,460 caseworkers). Standard errors are clustered at the caseworker level.

Further, the estimates for training courses and sanctions are in line with other studies. Job finding within six months rises by 9.1 percentage points per training course, confirming the positive effects Humlum et al. (2025) find for classroom training in Denmark. The magnitude is somewhat below the 14 to 16 percentage points they find at a similar horizon.[43] Moving from a caseworker at the 10th to one at the 90th percentile of the sanction tendency raises job finding within six months by 7.9%, in line with the large effects of stricter sanction regimes that Lalive et al. (2005), Arni et al. (2013) and Arni et al. (2022) find for Switzerland. The increase in worse jobs with the sanction tendency is also consistent with those studies. Monitoring leniency has no significant effect on job finding. Rambjer et al. (2026) find that stricter evaluation of search effort shortens unemployment in Sweden, but there sanctions are a main channel of strict evaluation,

[43] I divide the coefficient of the training measure on job finding within six months, 0.075, by its coefficient on the jobseeker's own participation in a training course, 0.820, both from the joint regression (Appendix Table E.23, Column (3), and Appendix Table E.24, Column (5)). Unlike in Humlum et al. (2025), the training measure is correlated with other caseworker dimensions, in particular with sanctions, with a correlation of 0.20. The estimate comes from the joint regression, however, which accounts for these correlations.

and my estimate holds sanctions fixed.[44]

6.5 Heterogeneity by caseworker type

Even though the tightness tendency is nearly uncorrelated with these other dimensions, it is still interesting to explore whether its effect is stronger for some caseworkers than for others. I estimate Equation (3) with the tightness tendency interacted with the five leave-one-out measures of support and stringency, all jointly, and with the five measures as controls. I standardize each measure to facilitate comparisons. Appendix Figure E.13 reports the effect of moving from the 10th to the 90th percentile of the tightness tendency depending on the value of the other dimensions.

The effect of the tightness tendency depends on the caseworker's tendency to sanction. With the sanction measure at its 10th percentile, a sanction risk of about 23%, moving from the 10th to the 90th percentile of the tightness tendency raises job finding within six months by 5.4%. With the sanction measure at its 90th percentile, a sanction risk of about 35%, the effect is -1.3% and not statistically significant. The effect is also larger among caseworkers with a tendency to assign more training courses: 0.5% at a training probability of about 12% and 3.8% at about 27%; the difference is significant at the 5% level. The tightness redirection effect is also larger among caseworkers with a tendency to meet their clients more often: 0.2% at about 1.2 meetings per month and 3.9% at about 1.4 meetings per month, with the difference also significant at the 5% level.

While the caseworker's other dimensions affect job finding, they do not significantly affect the pass-through from the caseworker's tendency to the jobseeker's own search, as shown in Column (5) of Appendix Table E.11. Taken together, these patterns suggest that the other dimensions matter mainly for whether redirected search leads to a job, not for how much the advice redirects search. More frequent meetings and training courses may help jobseekers find a job in the tight markets they search, for example through follow-up on applications or skills for the new market. Under a high sanction risk, jobseekers may take any available job quickly, so where they searched matters less for the job they take. Caseworkers choose how intensively they use each of these dimensions, so these differences describe for which caseworkers advice on where to search pays off; they are not causal effects of these dimensions on the effect of the advice.

Appendix Table E.10 repeats the joint specification with three characteristics of the caseworker instead of the five dimensions: the number of active clients, years of experience and the initial market tightness of the caseworker's clients. None of these interactions

[44] In their setting, central-office caseworkers evaluate the search effort reports, and the jobseeker's local caseworker then decides whether to register a violation. A sanction is thus a consequence of strict evaluation and would be a bad control. In my setting, the same caseworker evaluates the search effort and initiates sanctions, so the data do not allow me to separate the two steps.

is statistically significant, for job search or for job finding. The effect of the advice thus varies with the caseworker's tendencies in the other dimensions, not with the caseworker's workload or experience.

7 Aggregate effects on search tightness

The estimates so far measure the effect of assigning one jobseeker to a caseworker with a higher tendency, at the status quo of caseworker tendencies. A policy that makes the advice systematic, for example by training or instructing caseworkers, changes where many jobseekers search at the same time. This section uses a back-of-the-envelope calculation to investigate how such a policy would spill over onto jobseekers' search.

Another motivation for such a calculation is the following. In a model with several labor markets, Kircher (2022, Proposition 1) shows that moving a small amount of search from a slacker to a tighter market raises total matches, and that total matches are highest when tightness is equal across markets.[45] The concavity of the matching function is the key mechanism: moving a jobseeker from a slack market to a tight one generates a large reduction in congestion in the market the jobseeker leaves, where many searchers compete for few vacancies, and a small increase in congestion in the market the jobseeker joins, where vacancies are abundant relative to searchers. The concern is overshooting: if many jobseekers move search to the same tight markets, these markets could become slacker than the markets they left, and the advice would point to the wrong markets. The calculation checks whether such overshooting occurs.

To do so, I approximate the tightness of every market when every jobseeker gets a caseworker at the 90th percentile of the tendency, a p90 caseworker. The computation has three steps. First, where jobseekers search today. I measure search from the clicks: each clicker has one unit of search, divided equally over the clicks of the first three months, so that the share of search in a market is the share of clicks in that market. Not every jobseeker clicks, so I assume that every jobseeker searches like the clickers with the same initial market and registration quarter.

Second, where jobseekers would search with a p90 caseworker. This requires knowing, for every initial market, how search would be divided over all markets. There are too many pairs of initial and searched markets to estimate the effect of the tendency on each pair. I therefore use two estimates. The multinomial model of Figure 4, Panel (b), gives, for each tercile of initial tightness, how search moves between three groups of markets: markets slacker than the initial market, the initial market, and tighter markets. The

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

effect of the tendency on the search tightness index gives how much the average tightness of the markets searched rises. From these two estimates, I construct search with a p90 caseworker.

The main assumption is which markets jobseekers search with a p90 caseworker. Search could move only between the markets the jobseeker clicks today, or also toward tight markets that nobody clicks today. I assume something in between: a jobseeker can search in markets that the jobseeker does not click today, but markets that no jobseeker with the same initial market and registration quarter clicks today stay unclicked. The two estimates then determine how search moves over these markets. Between slacker markets, the initial market and tighter markets, search moves as the multinomial model implies for the jobseeker's tercile. Within each of these groups, search moves toward the tighter markets. Each market starts from its average share among the clickers with the same initial market and registration quarter, and this share is scaled by the market's tightness, with a strength chosen such that, on average over clickers, the total effect matches the coefficient of the leave-one-out tendency on own search tightness in Column (2) of the table in Figure 3.

Third, the new tightness. Throughout, I assume that vacancies are fixed: firms do not post more or fewer vacancies in response to where jobseekers search. I scale the measured tightness of every market by the ratio of the search of all jobseekers in the market and quarter today to their search in the computed counterfactual:

$$\frac{V_{mt}}{U_{mt}^{\text{p90}}} = \frac{V_{mt}}{U_{mt}} \frac{\sum_i s_{im,q(t)}}{\sum_i s_{im,q(t)}^{\text{p90}}}, \tag{5}$$

where V_{mt} and U_{mt} are the vacancies and registered jobseekers of market m on day t, $q(t)$ is the calendar quarter of day t, and $s_{im,q}$ and $s_{im,q}^{\text{p90}}$ are the search of jobseeker i in market m in quarter q today and with a p90 caseworker. Tightness falls in markets that receive more search and rises in markets that lose search. Appendix D details the computation.

Table 4 shows that there is no overshooting. Column (3) considers pairs of different markets that the same clicker searches in the same quarter. When every jobseeker gets a p90 caseworker, the tighter market of a pair becomes the slacker one in 0.15% of pairs, and in between 0.14% and 0.15% across terciles of initial tightness. The advice therefore still points to the tighter markets when every jobseeker gets it.

The exercise also shows whether the effect of the advice on search tightness, the average log tightness of the markets a jobseeker searches, depends on whether other jobseekers receive the advice too. Columns (1) and (2) compare the increase in search tightness when only the jobseeker's own caseworker changes and when every jobseeker gets a p90 caseworker, with tightness adjusted to the new search of all jobseekers. The answer depends on the initial market. For jobseekers who start in slack markets, the increase

Table 4: Increase in search tightness when every jobseeker gets a p90 caseworker

	Increase in search tightness (log points)		
Initial market	Only own caseworker changes (1)	Every jobseeker gets a p90 caseworker (2)	Tightness ranking reversed (% of pairs) (3)
Slack	2.31	2.84	0.15
Medium	2.18	2.15	0.14
Tight	1.57	1.01	0.14
All jobseekers	2.03	2.02	0.15

Notes: Columns (1) and (2) report the increase in search tightness, in log points (100× log difference), when the jobseeker moves from her current caseworker to a caseworker at the 90th percentile of the leave-one-out tightness tendency. In Column (1), only the jobseeker's own caseworker changes, and market tightness stays as today. In Column (2), every jobseeker gets a p90 caseworker, and market tightness adjusts to the new search, with vacancies fixed. Column (3) reports the percentage of pairs of markets searched by a clicker in a quarter in which the tighter market today becomes the slacker one when every jobseeker gets a p90 caseworker. The rows split jobseekers by the tercile of the tightness of their initial market; the last row covers all jobseekers in the baseline sample. Section 7 and Appendix D describe the computation.

is 23% larger when every jobseeker gets the advice, at 2.84 instead of 2.31 log points, because other jobseekers leave the slack markets they search, relieving congestion. For jobseekers who start in tight markets, it is 36% smaller, at 1.01 instead of 1.57 log points, because the tight markets they search receive search from jobseekers in slacker markets. For jobseekers who start in medium markets, it stays almost the same, at 2.15 instead of 2.18 log points.[46] In every tercile of initial tightness, search tightness thus still rises when every jobseeker gets the advice. For jobseekers who start in tight markets, the increase is about two thirds of the increase when only their own caseworker changes, and in no tercile does it differ by more than 36% from that increase.

8 Conclusion

This paper investigates whether advice 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 raises the average tightness of the markets jobseekers search in by 3.8% and job finding within six months by 2.8%. 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 additional

[46] Note that these changes compare a jobseeker's current caseworker with a p90 caseworker, not a p10 with a p90 caseworker as elsewhere in the paper; they are therefore smaller.

jobs are in tighter markets, but for those starting in tight initial markets, they are in the initial market itself. This pattern suggests that occupational mobility interventions are not always the right approach—for jobseekers already in favorable markets, advice that focuses search, rather than expanding it, raises job finding. These results complement recent field experiments targeting jobseekers in poor-prospect occupations with mobility interventions (Belot et al., 2025).

The findings speak to three policy objects. The first is caseworker training. The estimates compare caseworkers who already differ in their tendency, so they show what assignment to a high-tendency caseworker achieves. They suggest scope for gains if more caseworkers adopted the behavior of high-tendency ones, and Section 7 shows that the advice would likely still direct search toward tighter markets if every jobseeker received it. The second is showing tightness directly to jobseekers. Job portals already have real-time data on vacancies, and the PES has real-time data on the number of unemployed, so portals could display market conditions alongside job listings. This paper does not test such a display, and the caseworker tendency may capture more than tightness information — for instance convincing jobseekers or slightly more detailed advice than just markets — so whether it achieves similar gains is open and an interesting question for future research. The third is advice to broaden search to other occupations or regions. The results suggest that such advice should depend on the initial market: for jobseekers in slack markets, the additional jobs are in tighter markets, but for those in tight markets, they are in the initial market.

The analysis has limitations. The sample period, July 2020 to September 2022, covers the recovery from the COVID-19 pandemic, a period of unusual reallocation between sectors and many vacancies. The estimates are similar for jobseekers registering before and after the last restrictions were lifted in February 2022 (Section 5.2.1), but both periods belong to the recovery. The analysis also relies on the Swiss institutional context, with mandatory registration, regular caseworker meetings and an official job portal. Results may differ in settings where caseworkers meet jobseekers less often or where tightness information is harder to measure.

More broadly, the paper illustrates what can be learned by combining administrative data with new data sources. As job search moves online, it leaves digital traces of what jobseekers do, such as which postings they view. Linked to administrative records, these traces make it possible to study questions that register data alone cannot answer, such as where jobseekers search before they find a job.

References

Altmann, S., Falk, A., Jäger, S., & Zimmermann, F. (2018). Learning about job search: A field

experiment with job seekers in germany. *Journal of Public Economics*, *164*, 33–49.

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)

Altmann, S., Mahlstedt, R., Rattenborg, M. J., & Sebald, A. (2024). *Which occupations do unemployed workers target? insights from online job search profiles* (Discussion Paper 16696). IZA.

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. (2015). *Die Rolle von Erwartungshaltungen in der Stellensuche und der RAV-Beratung. Teilprojekt 2: Pilotprojekt Jobchancen-Barometer* (SECO Publikation Arbeitsmarktpolitik No. 43). Bern: State Secretariat for Economic Affairs (SECO). (Also published as IZA Research Report No. 70)

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

Arni, P., van den Berg, G. J., & Lalive, R. (2022). Treatment Versus Regime Effects of Carrots and Sticks. *Journal of Business & Economic Statistics*, *40*(1), 111–127. doi: 10.1080/07350015.2020.1784744

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

Baker, S. G. (1994). The multinomial-poisson transformation. *Journal of the Royal Statistical Society: Series D (The Statistician)*, *43*(4), 495–504.

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. (2019). Providing Advice to Jobseekers at Low Cost: An Experimental Study on Online Advice. *The Review of Economic Studies*, *86*(4), 1411–1447.

Belot, M., Kircher, P., & Muller, P. (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.

Bhuller, M., Dahl, G. B., Løken, K. V., & Mogstad, M. (2020). Incarceration, recidivism, and employment. *Journal of Political Economy*, *128*(4), 1269–1324.

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.

Card, D., Kluve, J., & Weber, A. (2010). Active labour market policy evaluations: A meta-analysis. *The economic journal*, *120*(548), F452–F477.

Card, D., Kluve, J., & Weber, A. (2018). What Works? A Meta Analysis of Recent Active Labor Market Program Evaluations. *Journal of the European Economic Association*, *16*(3), 894–931. doi: 10.1093/jeea/jvx028

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.

Crépon, B., & van den Berg, G. J. (2016). Active Labor Market Policies. *Annual Review of Economics*, *8*, 521–546. doi: 10.1146/annurev-economics-080614-115738

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.

Guimaraes, P. (2004). Understanding the multinomial-poisson transformation. *The Stata Journal*, *4*(3), 265–273.

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.)

HR Swiss Exams. (2016). *Wegleitung zur Prüfungsordnung HR-Fachmann / HR-Fachfrau mit eidg. Fachausweis, Fachrichtung B: Öffentliche Personalvermittlung und -beratung. Ausgabe 2017.* Zürich. (Dated 15 July 2016)

HR Swiss Exams. (2023). *Wegleitung zur Prüfungsordnung HR-Fachmann / HR-Fachfrau mit eidg. Fachausweis. Ausgabe 2023.* Zürich. Retrieved from https://www.hrse.ch/download/pictures/a9/fu9ztrm2czjpgkelbh0ju65a2s39w5/wl_hr-fachleute_2023_de.pdf (Dated 1 January 2023)

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.

Jäger, S., Roth, C., Roussille, N., & Schoefer, B. (2024). Worker beliefs about outside options. *The Quarterly Journal of Economics*, *139*(3), 1505–1556.

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.

Kling, J. R. (2006). Incarceration length, employment, and earnings. *American Economic Review*, *96*(3), 863–876.

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

Krueger, A. B., & Mueller, A. I. (2016). A contribution to the empirics of reservation wages. *American Economic Journal: Economic Policy*, *8*(1), 142–179.

Lalive, R., Osikominu, A., Pesaresi, L., Zuchuat, J., & Zweimüller, J. (2025). *Duration dependence in finding a job: Applications, interviews, and job offers* (Tech. Rep.). RF Berlin-CReAM Discussion Paper Series.

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.

Lalive, R., van Ours, J. C., & Zweimüller, J. (2008). The Impact of Active Labour Market Programmes on the Duration of Unemployment in Switzerland. *The Economic Journal*, *118*(525), 235–257. doi: 10.1111/j.1468-0297.2007.02111.x

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*.

Ma, K., & Samaniego de la Parra, B. (2021). *Labor market tightness, recruitment and search behavior* (Tech. Rep.). Working Paper.

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.

Mueller, A. I., Spinnewijn, J., & Topa, G. (2021). Job seekers' perceptions and employment prospects: Heterogeneity, duration dependence, and bias. *American Economic Review*, *111*(1), 324–363.

OpenAI. (2024). *ChatGPT (GPT-4o) [Large language model].* https://chat.openai.com/chat. (Response to queries made in 2024)

Rambjer, L., Uhlendorff, A., & Vikström, J. (2026). *Strictness in the evaluation of job search effort and employment outcomes.* (Working paper, August 2026)

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.

Schubert, G., Stansbury, A., & Taska, B. (2024). Employer concentration and outside options. *Available at SSRN 3599454*.

SECO. (2025a). *Jahresbericht Ausgleichsfonds der Arbeitslosenversicherung 2024* (Jahresrechnung). Bern: State Secretariat for Economic Affairs (SECO). (November 2025)

SECO. (2025b). *Schweizer Arbeitsmarkt 2024: Arbeitskräfteknappheit entspannt sich leicht.* Press release, State Secretariat for Economic Affairs (SECO), Bern. (10 January 2025)

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)

Sorkin, I. (2018). Ranking firms using revealed preference. *The Quarterly Journal of Economics*, *133*(3), 1331–1393. doi: 10.1093/qje/qjy001

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.

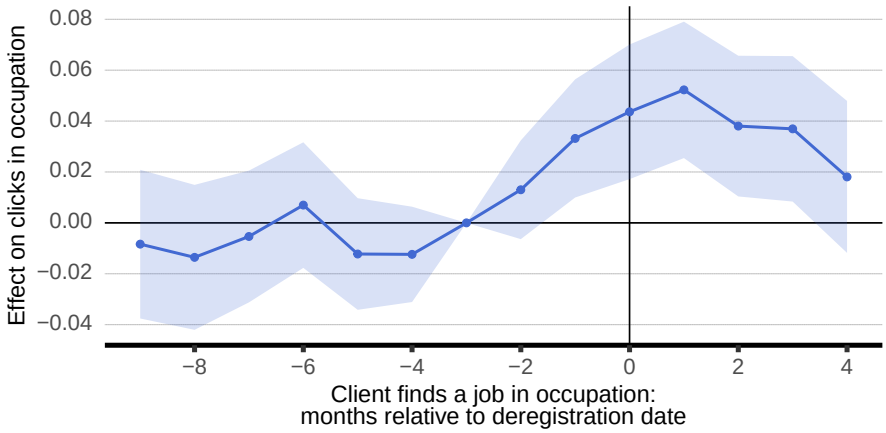
van der Klaauw, B., & Vethaak, H. (2022). *Empirical evaluation of broader job search requirements for unemployed workers* (Tinbergen Institute Discussion Paper).

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*.

Ziegler, L. (2024). *Labor market tightness and hiring outcomes* (Draft). University of Vienna.

Figure A.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 $t = -3$ months. See Appendix A for details on the event-study design. Standard errors are clustered at the event-caseworker level, 95% confidence intervals are shown.

Online Appendix

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

A 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 2).

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 × region fixed effects, and month × occupation fixed effects. Standard errors are clustered at the event-caseworker level. The reference period is $t = -3$.

B Data appendix

B.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 (8.5% 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.

B.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.[47]

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 6% of postings match to more than one ISCO occupation at the most granular level, and no posting matches to more than three; in these cases, I keep the occupation to which most of the posting's internal occupation codes map, breaking ties at random. For location, 97% 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 98% of job postings. I map municipalities to the 16 commuting zones using the official crosswalk.

B.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).

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 income, excluding disability, military, 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

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

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 1,566 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 224,427 job finders in the regression sample with a deregistration date, 88,466 deregistered before October 2021 and are therefore eligible for the earnings measure. Among those who deregistered in 2020, 12,799 out of 14,295 spells (89.5%) have a non-missing wage. The remaining missing values 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.

B.4 Commuting times

I obtain the travel distance between municipalities from openrouteservice.org, by the Heidelberg Institute for Geoinformation Technology at Heidelberg University (Heidelberg Institute for Geoinformation Technology (HeiGIT), 2024). Openrouteservice.org is a service that uses OpenStreetMap data to calculate travel distances between two points. I use the "driving-car" mode of transport. 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.

B.5 Location of the new job

The unemployment register records the location of the new employer as a free-text place name. I match these names to municipalities in three steps. First, I match them to the official municipality names. Second, I match the remaining names to the city names in the job-portal ads that correspond to a single municipality. Third, 312 place names remain, compared with 2,145 municipalities in Switzerland. They do not exactly match a municipality name because of differences in spelling, language or hyphenation. I match them with ChatGPT (GPT-4o) (OpenAI, 2024): I gave it the list of place names and asked it to return the name of the corresponding municipality. The exact prompt was not preserved. ChatGPT found no Swiss municipality for six names, five of which are in Liechtenstein and treated as abroad. I match the sixth and one further name by hand. After all three steps, 0.02% of new jobs with a place name remain unmatched. I map municipalities to commuting zones with the official crosswalk.

B.6 Imputing missing found-market cells

In the baseline, job finders with a missing found occupation and/or commuting zone are left out of the analysis by found market. In the robustness check of Appendix Figure E.8, 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, the candidate found markets are the cells (o, r_i), with weights

$$w_i(o) \;=\; \frac{N(n_i, r_i, o)}{\sum_{o'} N(n_i, r_i, o')}. \tag{6}$$

Each candidate market is classified with the baseline rules: as the initial market if it is the jobseeker's initial market, and otherwise as tighter or slacker by comparing its log tightness $\log\left(V(o, r_i, t_i)/U(o, r_i, t_i)\right)$ with that of the initial market, where V and U are the smoothed vacancy and unemployment counts on the deregistration date t_i. Candidates without a tightness measure count as market unknown. The imputed outcome for each market type is the sum of the weights of the candidates of that type, so the four outcomes add up to one for every imputed job finder. 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 (29,626 spells), 65.1% are imputed using the empirical distribution of occupations given region and industry, 31.5% 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 2.8% fall back to coarser distributions (NOGA-2 marginal or unconditional). In the regression sample, the found market is imputed for 22.1% of job finders within six months. After imputation, 1.6% of job finders remain in the market unknown category, compared with 23.3% in the baseline without imputation (Section 5.3, Appendix Figure E.8).

C Johnson–Lindenstrauss Approximation

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

C.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:

$$h_{ii} = X_i (X'X)^{-1} X_i' \tag{7}$$

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:

$$\hat{E}_{-i}[Y|X] = \hat{y}_i - \frac{h_{ii} \cdot \hat{u}_i}{1 - h_{ii}} \tag{8}$$

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.

C.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:

$$\hat{h}_{ii} = \frac{\sum_{s=1}^{p} \left(z_i^{(s)}\right)^2}{\sum_{s=1}^{p} \left(z_i^{(s)}\right)^2 + \sum_{s=1}^{p} \left(m_i^{(s)}\right)^2} \tag{9}$$

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

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.

D Search with a p90 caseworker

This appendix details how Section 7 computes search and tightness when every jobseeker gets a caseworker at the 90th percentile of the tendency, a p90 caseworker, and how Table 4 compares them with search and tightness today. Every jobseeker below the 90th percentile moves up to it at the same time.[48]

Where a jobseeker searches today follows from the clicks. Each clicker has one unit of search, spread equally over the clicks of the first three months. Since not every jobseeker clicks, every jobseeker, clicker or not, gets the average search of the clickers with the same initial market (3-digit ISCO occupation × commuting zone) who registered in the same quarter, each clicker counting once. Only clicks for which both the clicked and the initial market have a tightness on the day of the click are used, and clicks in markets with fewer than 30 registered jobseekers are dropped, as for search tightness. Jobseekers whose initial market and registration quarter have no clicker are left out.

With a p90 caseworker, the jobseeker keeps the markets of this average and only changes how the one unit of search is split across them. Each market the jobseeker searches is slacker than, as tight as, or tighter than the jobseeker's initial market on the day of the click; the market as tight as the initial market is the initial market itself. These are the three groups of the multinomial model in Figure 4, Panel (b). The first step determines how much search goes to each group. Columns (2)–(4) of Appendix Table E.14 report, for each tercile of initial tightness, how the log odds of searching in the initial market and in tighter markets, relative to slacker markets, change with the tendency. With a p90 caseworker, the share of group k in the search of jobseeker i is

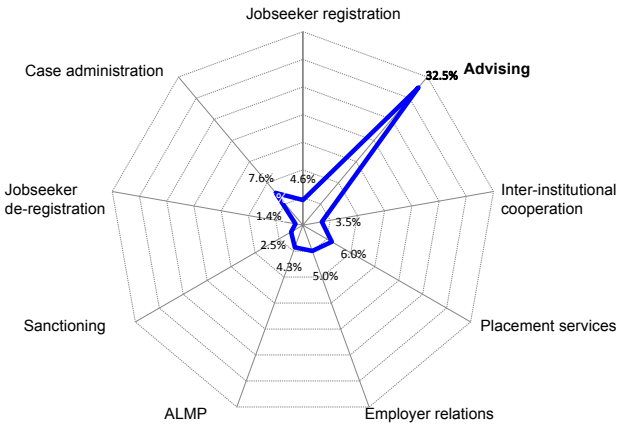
$$s_{ik}^{\mathrm{p90}} = \frac{s_{ik} \exp\left(\delta_{t(i)k} \left(Z^{\mathrm{p90}} - Z_{-i,c(i)}\right)\right)}{\sum_{k'} s_{ik'} \exp\left(\delta_{t(i)k'} \left(Z^{\mathrm{p90}} - Z_{-i,c(i)}\right)\right)}, \tag{10}$$

where k is one of the three groups, s_{ik} is the share of group k in the search of jobseeker i today, $t(i)$ is the tercile of initial tightness of jobseeker i, δ_{tk} is the coefficient of the multinomial model for group k in tercile t, with $\delta_{tk} = 0$ for slacker markets, $Z_{-i,c(i)}$ is the tendency of the caseworker of jobseeker i from Equation (2), and Z^{p90} is its 90th percentile. The odds of the initial market and of tighter markets, relative to slacker markets, are thus multiplied by $\exp\left(\delta_{t(i)k} \left(Z^{\mathrm{p90}} - Z_{-i,c(i)}\right)\right)$, and the denominator rescales the three shares to sum to one. The multinomial model cuts the terciles among clickers; the calculation applies its coefficients to the terciles of the baseline sample. For a jobseeker in a slack initial market, the share of the initial market falls and the share of tighter markets rises, with coefficients of -0.27 on the initial market and 0.07 on tighter markets in Column (2) of Appendix Table E.14. For a jobseeker in a medium initial market, the share of slacker markets falls, with coefficients of 0.27 on the initial market and 0.29 on tighter markets in Column (3) of Appendix Table E.14. For a jobseeker in a tight initial market, the share of the initial market rises, with coefficients of 0.20 on the initial market and 0.06 on tighter markets in Column (4) of Appendix Table E.14. I use the point estimates as they are, whether significant or not.

The second step determines which markets within each group receive the search. Inside each group, search moves toward the tighter markets of the group: each extra log point of tightness raises a market's share in its group by the same proportion,

[48] About 10% of jobseekers have a caseworker above the 90th percentile. For them, the distance to the 90th percentile is negative, and their search tightness falls slightly.

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



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

which grows with the distance to the 90th percentile, and the total share of the group stays at its value from the first step. For a market m in group k, the share in the search of jobseeker i with a p90 caseworker is

$$s_{im}^{\mathrm{p90}} = s_{ik}^{\mathrm{p90}} \frac{s_{im} \exp\big(\lambda \left(Z^{\mathrm{p90}} - Z_{-i,c(i)}\right) \log(V_{mt}/U_{mt})\big)}{\sum_{m' \in k} s_{im'} \exp\big(\lambda \left(Z^{\mathrm{p90}} - Z_{-i,c(i)}\right) \log(V_{m't'}/U_{m't'})\big)}, \tag{11}$$

where s_{im} is the share of market m in the search of jobseeker i today, with a market clicked on several days counted once per day, V_{mt}/U_{mt} is its tightness on the day t of the click, the sum runs over the markets of group k, s_{ik}^{p90} is the share of group k from Equation (10), and λ is the strength of the shift. This step is needed because the multinomial model only says how search moves between the three groups. Within the group of tighter markets, a high-tendency caseworker may also make the jobseeker search in the tightest of them, which the first step alone would miss. The strength of this shift is the same for all terciles and is set so that the search tightness of clickers rises on average by the coefficient of the tendency in Column (2) of the table in Figure 3 times their average distance to the 90th percentile, 1.90 log points. The first step alone gives 0.95 of these 1.90 log points, about half, and the second step gives the rest.

Vacancies stay fixed, so the tightness of a market changes only through the search arriving in it. The search in a market and calendar quarter is the sum of the search of all jobseekers in the sample in that market and quarter. In each market and quarter, new tightness divided by old tightness equals old search divided by new search, as in Equation (5). There, V_{mt} and U_{mt} are the vacancies and registered jobseekers of market m on day t, $s_{im,q}$ is the search of jobseeker i in market m in calendar quarter q, summed over the days of the quarter, $q(t)$ is the quarter of day t, and the sums run over all jobseekers in the sample. The calculation uses only the ratio of search today to search with a p90 caseworker, not the level of search: the measured tightness V_{mt}/U_{mt} is scaled by this ratio, so vacancies stay fixed and the number of jobseekers U_{mt} changes in proportion to the search of the jobseekers in the sample. The calculation rests on several assumptions. Vacancies are fixed: firms do not post more or fewer vacancies in response to where jobseekers search. Search effort is fixed: the caseworker changes where a jobseeker searches, not how much, so each jobseeker has one unit of search before and after. No new markets are added to search: jobseekers shift search only among markets that clickers from the same initial market and registration quarter already search. Clicks measure search: the share of a jobseeker's clicks in a market is the share of their search in that market. Non-clickers search like clickers with the same initial market and registration quarter, and respond to the caseworker in the same way. The jobseekers in the sample are all the searchers in a market: searchers outside the sample, such as registered jobseekers after their first three months and employed or unregistered jobseekers, are not counted. If they search as before, tightness changes less than computed here. There is one round of adjustment per quarter: search responds to the caseworker, but not again to the new tightness. Finally, the outcome is search tightness, not job finding, so no matching function is estimated.

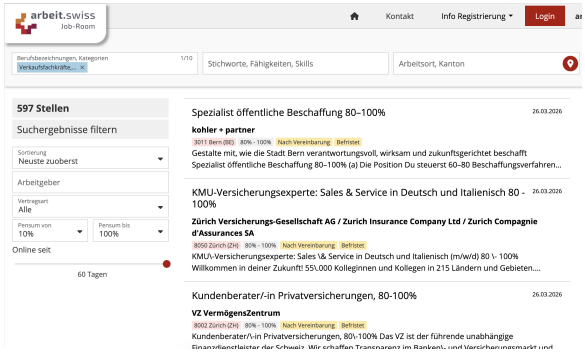
E Further figures and tables

Table E.1: Caseworker assignment methods across Swiss cantons

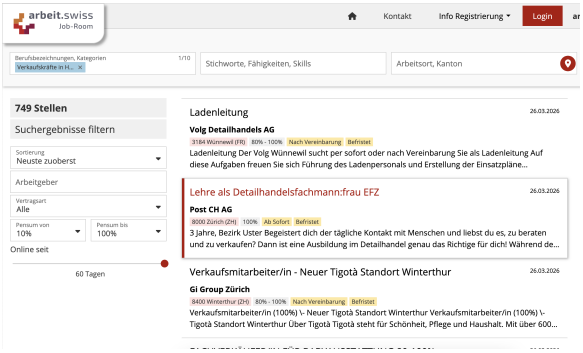
Assignment method	N cantons	Share of sample
Equal distribution by administration (proportional to caseworker's contractual hours, regardless of caseload)	6	22.4%
Caseload balancing (distribute to equalize case burden across caseworkers)	10	44.4%
Random assignment (software-based random allocation)	1	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.

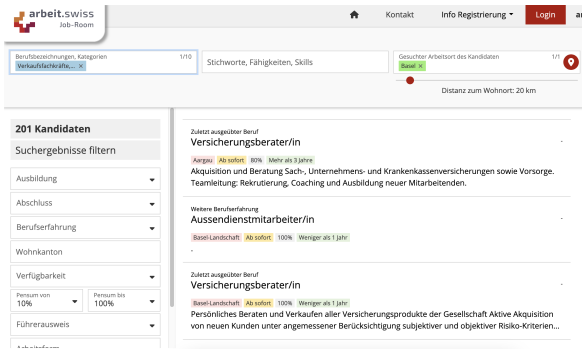
Figure E.2: Vacancy and candidate counts on the job portal for two neighboring sales occupations in Basel



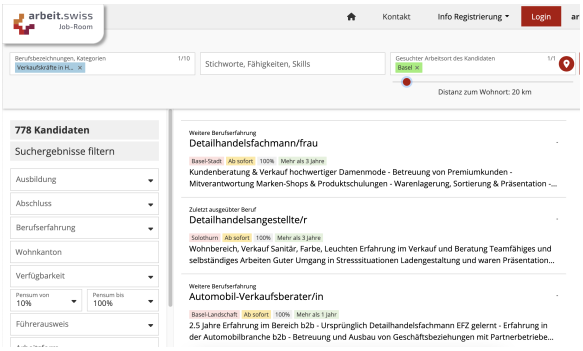
(a) Vacancies: ISCO 332 — Sales and purchasing agents and brokers (597 listings)



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



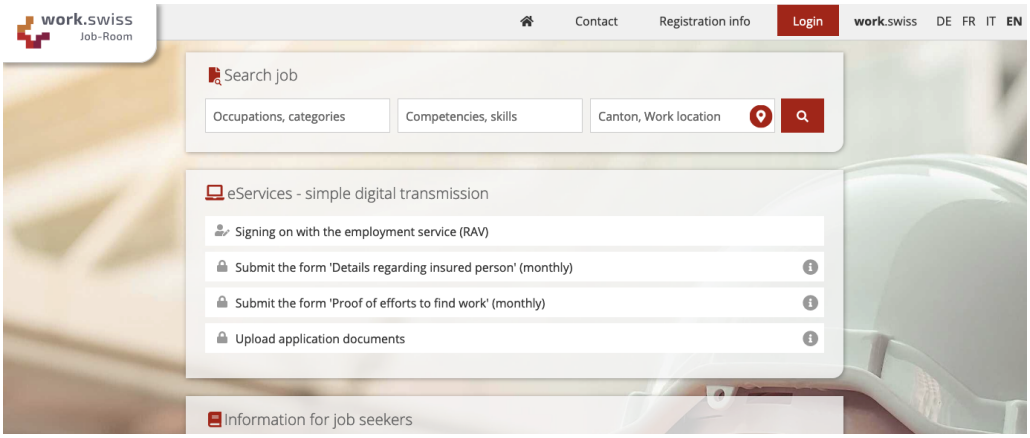
(c) Candidates: ISCO 332 — Sales and purchasing agents and brokers (201 candidates)



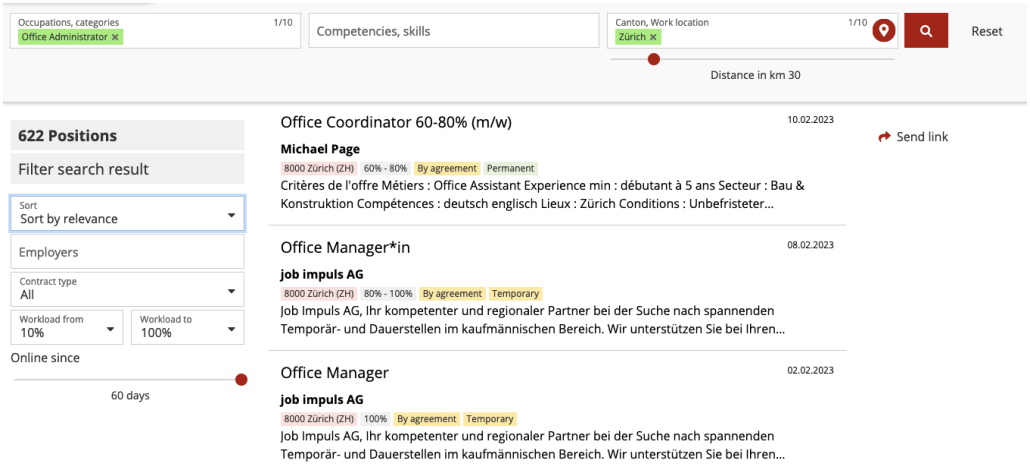
(d) Candidates: ISCO 522 — Shop salespersons (778 candidates)

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 $749/778 \approx 1.0$.

Figure E.3: Screenshots of the job portal job-room.ch



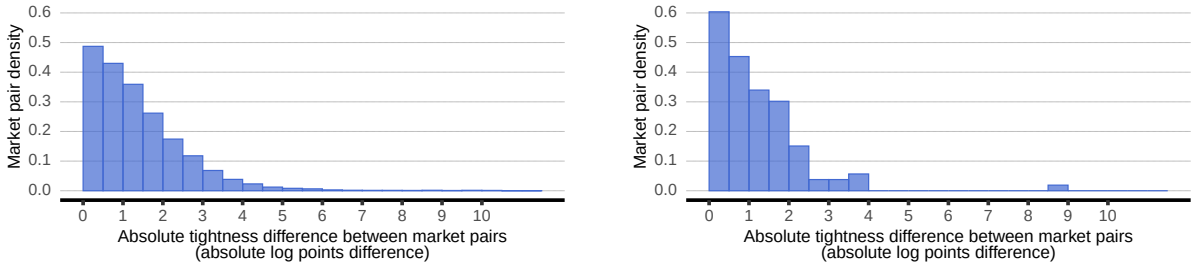
(a) Starting page



(b) Result list for "Office Administrator" in Zurich

Notes: Panel (a) shows the starting page of the job search portal of the Swiss public employment service, `https://job-room.ch/home/job-seeker`. Panel (b) shows the result list when searching for "Office Administrator" in Zurich. Screenshots taken on 10 February 2023.

Figure E.4: Tightness dispersion across market pairs: alternative samples

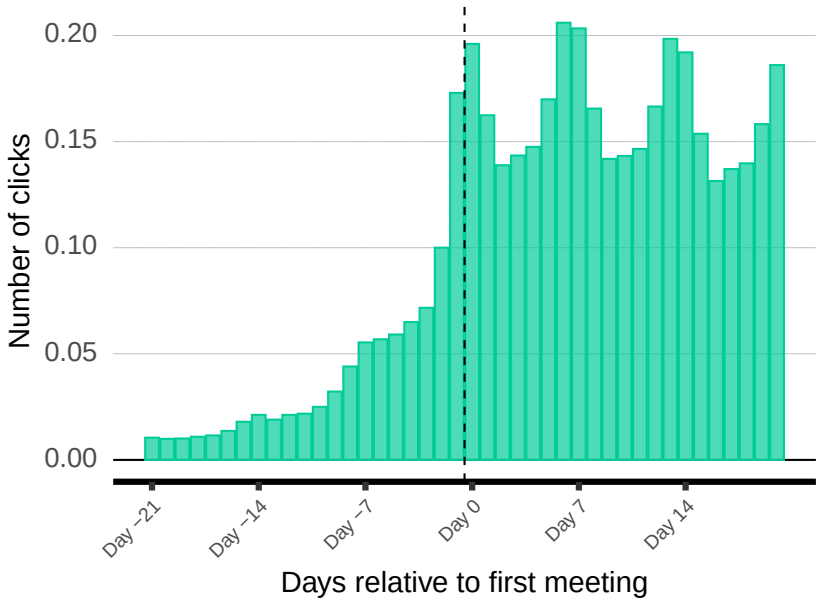


(a) All market pairs

(b) High transition pairs

Notes: Distribution of absolute log tightness differences across pairs of different markets, where a market is an ISCO 3-digit occupation × commuting zone cell. Panel (a) includes all pairs with at least one transition from the first to the second market in 2004–2019, from the last job before unemployment to the job found. Panel (b) restricts to pairs in which at least 20% of these transitions out of the first market go to the second. Tightness is measured as in Figure 2, Panel (a). Both panels use the same bins and axes.

Figure E.5: 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 E.2: Meeting effects on search targeting: number of clicks and clicked-market tightness

	N clicks	Clicked market tightness
	(1)	(2)
Post meeting	0.8357***	-0.0087
	(0.0231)	(0.0062)
Post meeting * Leave-one-out tightness	-0.0602	0.1248***
	(0.1092)	(0.0461)
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	5,084,764	592,417
Mean of dependent var.	0.39296	-0.61095
p10-p90 magnitude		3.65%
Pseudo R2 / R2	0.2674	0.4978

Notes: 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). 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 caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table E.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 × Leave-one-out tightness	-0.0007	-0.0008
	(0.0012)	(0.0012)
Jobseeker spell FE		X
---	---	---
Observations	3,864,951	3,864,951
R2	0.06308	0.07118

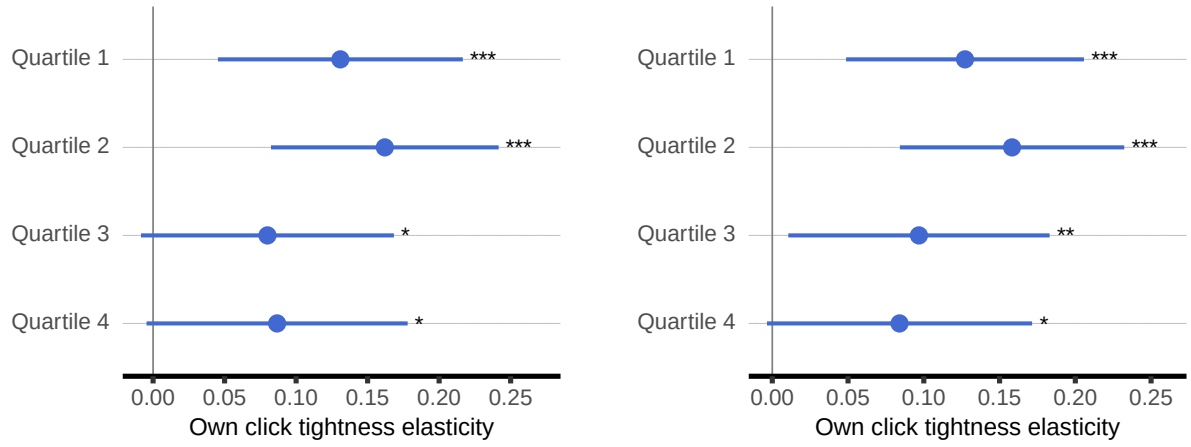
Notes: The outcome variable is an indicator equal to one if the jobseeker finds a job in the market. For each jobseeker who found employment and clicked on at least five postings, the choice set includes all markets (ISCO 3-digit occupation × commuting zone) in which at least one click was placed, plus 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 E.4: 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%	-0.02806 (0.00711)	0.000	-1.8%
Age 35-49	0.00049 (0.00243)	0.839	+0.4%	-0.00337 (0.00664)	0.934	-0.3%
Age 50+	-0.01276 (0.00211)	0.000	-15.0%	-0.01242 (0.00742)	0.567	-1.7%
Has children	-0.01359 (0.00252)	0.000	-9.3%	0.00554 (0.00705)	0.887	+0.4%
Female x has children	-0.04504 (0.00200)	0.000	-45.1%	-0.01105 (0.00578)	0.414	-1.7%
Non-Swiss	-0.00208 (0.00242)	0.715	-1.4%	0.01116 (0.00824)	0.725	+0.7%
Primary or lower sec. education	-0.04269 (0.00199)	0.000	-49.0%	-0.00138 (0.00845)	0.981	-0.2%
University education	0.06686 (0.00190)	0.000	+178.6%	0.01268 (0.00509)	0.131	+2.3%
3+ years tenure in last occupation	0.00824 (0.00239)	0.004	+3.6%	-0.00980 (0.00890)	0.822	-0.5%
Employed in year t-1	0.00076 (0.00080)	0.715	+0.3%	-0.00269 (0.00252)	0.822	-0.1%
Employed in year t-2	0.00229 (0.00102)	0.096	+1.1%	-0.00505 (0.00330)	0.614	-0.2%
Income in year t-1 (log)	0.03023 (0.00751)	0.001	+1.7%	0.00386 (0.02387)	0.981	+0.0%
Income in year t-2 (log)	0.03115 (0.00908)	0.004	+1.8%	-0.02584 (0.02967)	0.887	-0.1%
N	101,511			366,890		

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) at 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.

Figure E.6: Monotonicity: effect on search tightness by quartile of predicted search tightness



(a) Leave-one-out tendency

(b) Reverse-sample tendency

Notes: The figure reports estimates from the specification in Equation (3) with the jobseeker's own search tightness index as the outcome, estimated separately within quartiles of the predicted search tightness index. The prediction uses the thirteen predetermined jobseeker characteristics, net of office × year and initial occupation × commuting zone × quarter fixed effects. Panel (a) shows the coefficient on the caseworker leave-one-out tendency. Panel (b) shows the coefficient on a "reverse-sample" tendency, computed for jobseekers in a given quartile only from the caseworker's clients in the other three quartiles; caseworkers need at least five such clients. Sample: clickers with non-missing characteristics. Bars show 95% confidence intervals. Standard errors are clustered at the caseworker level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table E.5: Robustness of the effect on job finding

Panel A: Assignment and construction of the caseworker tendency

	(1) Baseline	(2) Office × year × industry FE	(3) Office × year × occupation FE	(4) Leave out same occupation	(5) Non-clickers only	(6) Outcome: job in unknown market
Leave-one-out tightness	0.0333***	0.0279***	0.0229***	0.0306***	0.0384***	-0.0018
	(0.0093)	(0.0080)	(0.0089)	(0.0089)	(0.0108)	(0.0109)
Unemployment office × year FE	X			X	X	X
Unemployment office × year × last industry FE		X				
Unemployment office × year × last 2-digit occupation FE			X			
Last occ. × commuting zone × quarter FE	X	X	X	X	X	X
Jobseeker controls	X	X	X	X	X	X
p10-p90 magnitude	+2.8%	+2.8%	+2.1%	+2.7%	+3.6%	-0.6%
N caseworkers	2,460	2,460	2,460	2,460	2,460	2,460
Observations	366,890	332,973	348,521	366,807	225,480	366,890
R2	0.14161	0.17738	0.16389	0.14165	0.16738	0.07692
Mean of dependent var.	0.36554	0.36713	0.36597	0.36554	0.34720	0.08501

Panel B: Tightness measure and sample period

	(1)	(2)	(3)	(4)	(5)	(6)
	Tightness measure				Registered	
	4-digit occupations	Labor market regions	Ads per clicker	V and U separately	Before 17 Feb 2022	From 17 Feb 2022
Leave-one-out tightness	0.0224***	0.0249***	0.0352**		0.0330***	0.0350**
	(0.0086)	(0.0088)	(0.0144)		(0.0099)	(0.0149)
Leave-one-out log vacancies				0.0350***		
				(0.0107)		
Leave-one-out log unemployed				-0.0313**		
				(0.0124)		
Unemployment office × year FE	X	X	X	X	X	X
Last occ. × commuting zone × quarter FE	X	X	X	X	X	X
Jobseeker controls	X	X	X	X	X	X
p10-p90 magnitude	+2.0%	+2.2%	+1.9%	V +2.6%, U −2.1%	+2.9%	+2.7%
N caseworkers	2,460	2,460	2,460	2,460	2,435	2,102
Observations	368,163	368,052	368,728	366,890	282,120	84,770
R2	0.14193	0.14191	0.14198	0.14162	0.15778	0.14230
Mean of dependent var.	0.36568	0.36568	0.36561	0.36554	0.35930	0.38630

Notes: The table reports estimates from the specification in Equation (3) with jobseeker controls. The outcome variable indicates whether the jobseeker finds employment within 6 months in any market, except in Panel A, Column (6). Panel A varies the assignment conditioning and the construction of the caseworker tendency: Column (1) is the baseline (Table 2, Column (2)); Columns (2) and (3) interact the office × year FE with the 2-digit industry (NOGA) and with the 2-digit occupation (ISCO) of the last job; Column (4) computes the leave-out excluding all clients of the same caseworker with the same last occupation; Column (5) restricts the sample to jobseekers without clicks on job-room.ch in the first three months; Column (6) changes the outcome to finding employment within 6 months in a market that is not observed or has no tightness measure (market unknown), without imputation. Panel B varies the tightness measure and the sample period: Column (1) uses 4-digit ISCO occupations instead of 3-digit, Column (2) the 101 labor market regions instead of the 16 commuting zones, Column (3) the number of ads per clicker on the platform; Column (4) replaces the leave-one-out tightness tendency by two leave-one-out tendencies, one for the log number of vacancies and one for the log number of unemployed jobseekers in the clicked markets, entered jointly, and reports the magnitude of each; Columns (5) and (6) split the sample into jobseekers registering before and from 17 February 2022, when the Swiss federal government lifted the remaining COVID-19 restrictions. Standard errors clustered at the caseworker level in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table E.6: 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	0.0028	0.0044	0.0133**	0.0230***	0.0309***	0.0333***	0.0370***	0.0418***
	(0.0017)	(0.0042)	(0.0061)	(0.0074)	(0.0086)	(0.0093)	(0.0104)	(0.0117)
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
p10-p90 magnitude	+4.8%	+1.5%	+2.4%	+2.9%	+3.1%	+2.8%	+2.6%	+2.4%
N caseworkers	2,460	2,460	2,460	2,460	2,460	2,460	2,460	2,460
Observations	366,890	366,890	366,890	366,890	366,890	366,890	343,644	299,294
R2	0.04722	0.07980	0.10497	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	0.44698	0.55043

Notes: The table reports estimates from the specification in Equation (3), with the outcome variable indicating whether the jobseeker finds employment within the horizon shown in the column header. 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 E.7: 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	0.1665	0.0251	0.1220**	0.1638***	0.1890***	0.1014	0.1297*	0.0738	0.2142**	0.1429	0.0611	0.2097**
	(0.1037)	(0.0579)	(0.0530)	(0.0542)	(0.0590)	(0.0632)	(0.0669)	(0.0725)	(0.0840)	(0.0895)	(0.0958)	(0.1033)
Unemployment office x Year FE	X	X	X	X	X	X	X	X	X	X	X	X
Last occ. x commuting zone x quarter FE	X	X	X	X	X	X	X	X	X	X	X	X
Observations	242,992	328,373	311,357	279,758	250,242	223,518	193,737	168,480	149,492	130,423	115,054	102,808
Pseudo R2	0.08382	0.06634	0.06047	0.05887	0.06119	0.05848	0.06002	0.06017	0.06285	0.06370	0.06819	0.07076
Mean of dependent var.	0.02759	0.08031	0.09835	0.09855	0.09476	0.08540	0.08079	0.07489	0.06847	0.06489	0.06190	0.05701

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 outcome variable is an indicator for finding a job within the specific period, conditional on not having found a job in previous periods; spells registered too late to observe the full period are left out. The key regressor is the leave-one-out caseworker tendency of Equation (3). All regressions include the fixed effects of Table 2, Column (1), unemployment office × year and last occupation × commuting zone × quarter, but no jobseeker controls. The model drops fixed-effect groups without any job finding in the period, so the sample size varies across columns. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.10.

Table E.8: Robustness: Sensitivity to click and caseworker clicker thresholds

Minimum clicks per jobseeker	Minimum clicking clients per caseworker	N	Clicking jobseekers	Caseworkers	Coef (Any)	Magnitude (Any) %	Coef (Own clicks tightness)
1	2	381,390	144,143	2716	0.0345*** (0.0082)	2.89	0.1293*** (0.0191)
1	5	379,402	143,651	2536	0.0364*** (0.0088)	3.02	0.1510*** (0.0194)
1	10	375,030	142,558	2373	0.0332*** (0.0092)	2.69	0.1480*** (0.0203)
1	30	341,684	133,076	1894	0.0220** (0.0101)	1.65	0.1549*** (0.0226)
5	2	369,975	102,035	2653	0.0349*** (0.0083)	3.02	0.1204*** (0.0197)
5	5	366,890	101,511	2460	0.0333*** (0.0093)	2.83	0.1227*** (0.0209)
5	10	360,252	100,293	2283	0.0310*** (0.0098)	2.58	0.1347*** (0.0215)
5	30	291,540	86,319	1573	0.0070 (0.0121)	0.50	0.1299*** (0.0260)
10	2	357,997	77,862	2604	0.0310*** (0.0079)	2.91	0.0872*** (0.0194)
10	5	353,819	77,321	2405	0.0314*** (0.0089)	2.89	0.0870*** (0.0206)
10	10	342,114	75,697	2174	0.0291*** (0.0094)	2.62	0.0931*** (0.0213)
10	30	231,167	57,376	1244	0.0042 (0.0131)	0.31	0.0362 (0.0310)
30	2	311,636	36,871	2445	0.0172*** (0.0062)	2.22	0.0181 (0.0197)
30	5	297,225	35,999	2146	0.0170** (0.0073)	2.08	0.0140 (0.0222)
30	10	253,111	32,730	1670	0.0093 (0.0087)	1.03	0.0091 (0.0256)
30	30	48,691	9,555	309	0.0043 (0.0245)	0.35	0.0250 (0.0685)

Notes: The table reports estimates from the specification in Equation (3), for different combinations of the minimum number of clicks per jobseeker and the minimum number of clicking clients per caseworker; the baseline uses at least 5 clicks per jobseeker and at least 5 clicking clients 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 in percent of the predicted outcome at the 10th percentile. The outcome variable indicates whether the jobseeker finds employment within 6 months. The last column reports the effect on job search: the coefficient of the caseworker leave-one-out tendency on the jobseeker's own search tightness index, estimated on clicking jobseekers. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table E.9: Caseworkers with fewer than 30 and with 30 or more clicking clients

	Group mean		30 or more minus fewer than 30 clickers	
Measure	Fewer than 30 clickers	30 or more clickers	Difference	Difference within office
Sanction rate	0.257	0.284	0.026*** (0.004)	0.025*** (0.004)
Training rate	0.172	0.184	0.012*** (0.004)	0.008** (0.004)
Meetings per month	1.38	1.42	0.03*** (0.01)	-0.01 (0.01)
Minutes per meeting	35.0	34.3	-0.7* (0.4)	-0.1 (0.2)
Search effort protocols accepted	0.946	0.941	-0.006*** (0.002)	-0.005*** (0.002)
Active clients per month	54.3	94.7	40.4*** (1.5)	40.5*** (1.3)
Clients, June 2020 to December 2022	79.9	188.2	108.3*** (2.0)	107.5*** (2.0)
Share of clients who click	0.244	0.299	0.055*** (0.004)	0.037*** (0.004)
Years of experience	7.1	6.5	-0.6** (0.3)	-0.9*** (0.3)
N caseworkers	887	1573		

Notes: Means over caseworkers with fewer than 30 and with 30 or more clicking clients (clients with at least five clicks in the first three months) in the baseline sample, and the difference between the two groups from a regression of each measure on an indicator for 30 or more clicking clients, across all caseworkers and with fixed effects for the caseworker's main office; robust standard errors in parentheses. Sanction rate, training rate, meetings per month of unemployment, minutes per meeting and the share of search effort protocols accepted are raw means over the caseworker's clients in the first three months of the spell. Active clients per month: open spells of the caseworker on the 15th of each month, averaged over months with at least one. Experience: years since the caseworker's first closed spell in the data (from 2004) in June 2020. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table E.10: Effects on job search and job finding: interactions with caseworker characteristics

	Job search	Job finding
	(1)	(2)
Dependent Var.:	Own clicks tightness	Find job within 6 months
Leave-one-out tightness	0.1148***	0.0332***
	(0.0227)	(0.0102)
× active clients per month (SD)	0.0175	-0.0055
	(0.0214)	(0.0085)
× years of experience (SD)	0.0095	-0.0065
	(0.0206)	(0.0099)
× clients' initial tightness (SD)	0.0168	0.0009
	(0.0194)	(0.0067)
Active clients per month (SD)	0.0027	-0.0021
	(0.0022)	(0.0013)
Years of experience (SD)	-0.0037*	0.0008
	(0.0021)	(0.0013)
Clients' initial tightness (SD)	0.0121***	-0.0010
	(0.0024)	(0.0013)
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	101,511	366,890
R2	0.73807	0.14164
Mean of dependent var.	-0.57700	0.36554

Notes: The table reports estimates from the specification in Equation (3), interacting the caseworker leave-one-out tendency jointly with three caseworker characteristics: active clients per month, years of experience, and the mean initial market tightness of the caseworker's clients. Active clients per month and experience are defined as in Appendix Table E.9; initial market tightness is the log tightness in the client's last occupation × commuting zone at registration. Each characteristic is centered within unemployment office × year and divided by its standard deviation (30.4 active clients, 5.8 years, 0.27 log points), so the coefficient on the leave-one-out tendency is the slope at the office-year mean. The regressions also include the three characteristics themselves. In Column (1), the outcome variable is the jobseeker's own search tightness index, estimated on clicking jobseekers. In Column (2), the outcome 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 E.11: Effect of the tightness tendency, other caseworker dimensions and caseworkers with 30 or more clicking clients

	Find job within 6 months				Own clicks tightness	
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Var.:	Find job within 6 months	Find job within 6 months	Find job within 6 months	Find job within 6 months	Own clicks tightness	Own clicks tightness
Leave-one-out tightness	0.0288***	0.0486***	0.0247***	0.0397***	0.1145***	0.0669*
	(0.0087)	(0.0122)	(0.0086)	(0.0120)	(0.0209)	(0.0399)
× N meetings (SD)			0.0196**	0.0188**	-0.0137	-0.0107
			(0.0088)	(0.0087)	(0.0228)	(0.0230)
× Avg meeting duration (SD)			0.0083	0.0085	0.0064	0.0060
			(0.0093)	(0.0093)	(0.0259)	(0.0259)
× Training course (SD)			0.0161**	0.0161**	0.0124	0.0124
			(0.0079)	(0.0079)	(0.0231)	(0.0232)
× Monitoring leniency (SD)			-0.0137	-0.0121	-0.0127	-0.0150
			(0.0091)	(0.0090)	(0.0254)	(0.0257)
× Any sanction (SD)			-0.0323***	-0.0284***	-0.0361	-0.0425
			(0.0095)	(0.0094)	(0.0250)	(0.0260)
× CW with 30+ clickers		-0.0343**		-0.0248		0.0721
		(0.0170)		(0.0169)		(0.0492)
N meetings (SD)	0.0005	0.0009	0.0006	0.0009	-0.0035	-0.0037*
	(0.0013)	(0.0013)	(0.0013)	(0.0013)	(0.0022)	(0.0022)
Avg meeting duration (SD)	-0.0024*	-0.0024*	-0.0024*	-0.0024*	-0.0014	-0.0014
	(0.0013)	(0.0013)	(0.0014)	(0.0014)	(0.0021)	(0.0021)
Training course (SD)	0.0051***	0.0051***	0.0049***	0.0049***	-0.0004	-0.0004
	(0.0013)	(0.0013)	(0.0013)	(0.0013)	(0.0020)	(0.0020)
Monitoring leniency (SD)	-0.0010	-0.0011	-0.0010	-0.0011	0.0036	0.0036
	(0.0014)	(0.0014)	(0.0014)	(0.0014)	(0.0026)	(0.0026)
Any sanction (SD)	0.0112***	0.0105***	0.0111***	0.0105***	0.0036	0.0041
	(0.0017)	(0.0017)	(0.0017)	(0.0017)	(0.0027)	(0.0027)
CW with 30+ clickers		0.0109***		0.0103***		-0.0098
		(0.0031)		(0.0031)		(0.0066)
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
N caseworkers	2,460	2,460	2,460	2,460	2,460	2,460
Observations	338,885	338,885	338,885	338,885	98,648	98,648
R2	0.14602	0.14611	0.14613	0.14620	0.73984	0.73985
Mean of dependent var.	0.36785	0.36785	0.36785	0.36785	-0.58241	-0.58241

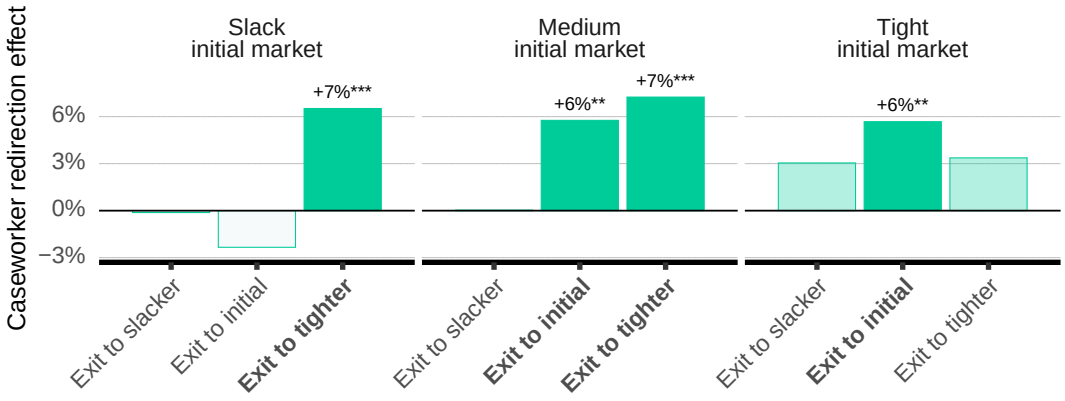
Notes: The table reports estimates from the specification in Equation (3) with the five leave-one-out caseworker measures of support and stringency described in Section 6.4, each standardized across spells. The outcome is finding a job within six months in Columns (1)–(4) and the search tightness index in Columns (5) and (6) (clickers only). Columns (1) and (2) include the tightness tendency and the five measures as regressors, and Column (2) adds an indicator for caseworkers with 30 or more clicking clients and its interaction with the tendency. Columns (3) and (4) interact the tightness tendency with the five measures and Column (4) adds the 30+ indicator's interaction with the tendency. Columns (5) and (6) repeat Columns (3) and (4) for the search tightness index. The sample consists of spells with all five measures defined. Standard errors are clustered at the caseworker level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table E.12: Robustness: Click window threshold

	(1)	(2)	(3)	(4)	(5)
Leave-one-out tightness	0.0009	0.0349***	0.0334***	0.0303***	0.0287***
	(0.0059)	(0.0082)	(0.0093)	(0.0101)	(0.0111)
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%	+3.4%	+2.8%	+2.4%	+2.2%
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	295,919	357,890	366,890	370,817	373,876
R2	0.13500	0.14064	0.14161	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 (at least 5 clicks per jobseeker, at least 5 clicking clients per caseworker) are applied within each window. The outcome 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.

Figure E.7: Heterogeneous effects on job finding by initial market tightness, initial market from the commuting zone of the last job



Notes: The figure repeats Figure 4, Panel (a), with the initial market defined as the last occupation (ISCO 3-digit) × the commuting zone of the last employer instead of the commuting zone of residence. This changes the initial-tightness terciles and the classification of exits into tighter, initial and slacker markets. The workplace of the last job is the place name of the last employer in the unemployment register, matched to municipalities as for the new job; spells without a matched workplace keep the commuting zone of residence. The caseworker tendency, the controls and the fixed effects are unchanged, and job finders whose market is not observed or has no tightness measure are again left out; the occupation × commuting zone × quarter fixed effect stays residence based, since jobseekers are assigned to the employment office of their residence. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table E.13: Heterogeneity by baseline tightness: job-finding outcomes

	Job finding (within 6 months)			
	(1)	(2)	(3)	(4)
Dependent Var.:	Any market	In tighter	In initial	In slacker
Leave-one-out tightness	0.0246*	0.0247***	-0.0012	0.0012
	(0.0136)	(0.0088)	(0.0077)	(0.0047)
Leave-one-out tightness x Tercile 2	0.0181	-0.0001	0.0218**	-0.0036
	(0.0162)	(0.0111)	(0.0106)	(0.0079)
Leave-one-out tightness x Tercile 3	0.0199	-0.0165	0.0294***	0.0071
	(0.0191)	(0.0104)	(0.0112)	(0.0108)
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	+3.1%	+6.9%	-0.4%	+1.0%
Magnitude p10-p90, medium initial market	+4.2%	+6.8%	+5.1%	-0.9%
Magnitude p10-p90, tight initial market	+3.8%	+3.3%	+6.8%	+1.6%
N caseworkers	2,460	2,460	2,460	2,460
Observations	332,531	332,531	332,531	332,531
R2	0.14700	0.09284	0.10400	0.11531
Mean of dependent var.	0.30858	0.10184	0.11778	0.08895

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 × commuting zone at registration). Outcomes: finding a job within 6 months in any market (Column (1)), in a market tighter than the initial market at the deregistration date (Column (2)), in the initial market, i.e. same occupation and commuting zone (Column (3)), in a slacker market (Column (4)). The sample leaves out job finders whose market is not observed or has no tightness measure, so that Columns (2)–(4) add up to Column (1). Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table E.14: Effects of caseworker redirection on the market type of clicks, by initial market tightness

	All	Initial market tightness			Search tightness
	(1)	(2)	(3)	(4)	(5)
Sample (Initial market tightness)	Full sample	Slack	Medium	Tight	
Dependent Var.:	Clicks by market type	Clicks by market type	Clicks by market type	Clicks by market type	Avg. log tightness of clicks
Leave-one-out tightness x Market type = Initial	0.0012	-0.2721***	0.2660**	0.1994*	
	(0.0598)	(0.1042)	(0.1042)	(0.1084)	
Leave-one-out tightness x Market type = Tighter	0.1363***	0.0659	0.2878***	0.0632	
	(0.0417)	(0.0711)	(0.0708)	(0.0749)	
Leave-one-out tightness x Initial market tightness = Slack					0.1377***
					(0.0371)
Leave-one-out tightness x Initial market tightness = Medium					0.1193***
					(0.0291)
Leave-one-out tightness x Initial market tightness = Tight					0.1125***
					(0.0321)
Initial-tightness tercile FE					X
Spell	X	X	X	X	
Market type x Office-year	X	X	X	X	
Market type x Occupation-region-quarter	X	X	X	X	
Unemployment office x Year FE					X
Last occ. x commuting zone x quarter FE					X
Market type x jobseeker controls	X	X	X	X	
Jobseeker controls					X
Varying Slopes:					
Income in year t-1 (log) (Market type)	X	X	X	X	
Income in year t-2 (log) (Market type)	X	X	X	X	
Family	Poisson	Poisson	Poisson	Poisson	OLS
Share change p10-p90, slacker (pp)	-0.5**	+0.1	-1.5***	-0.7	
Share change p10-p90, initial (pp)	-0.3	-1.4***	+0.4	+0.7*	
Share change p10-p90, tighter (pp)	+0.8***	+1.3***	+1.1***	+0.0	
N spells	96,095	31,666	32,262	30,949	100,486
N caseworkers	2,460	2,326	2,414	2,333	2,460
Observations	296,774	98,783	99,657	96,917	100,486
Squared Cor.	0.90433	0.94193	0.88672	0.92822	0.73709

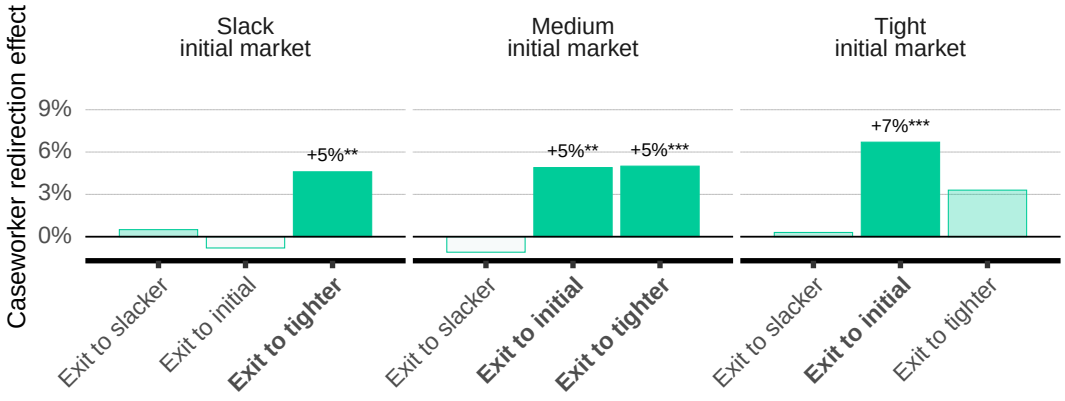
Notes: Multinomial logit over three market types of each click: slacker than, the same as (initial market), or tighter than the jobseeker's initial market (last occupation × commuting zone at registration), with slacker markets as the base. The model is estimated as a Poisson regression of click counts by spell and market type with spell fixed effects. The coefficients are the effects of the leave-one-out caseworker tightness tendency on the log odds of a click in the initial or in a tighter market relative to a slacker market. The alternative-specific constants are interacted with office × year FE, last occupation × commuting zone × quarter FE and the jobseeker controls. Column (1) pools all clickers; Columns (2)–(4) estimate the model separately by tercile of initial market tightness. The bottom rows report the change in the predicted click share, in percentage points, when moving from the 10th to the 90th percentile in the caseworker leave-one-out tendency, with delta-method standard errors; Figure 4, Panel (b), plots the same changes in percent of the predicted share at the 10th percentile. Column (5) reports the effect on the jobseeker's search tightness in the same spells, from an OLS regression with the leave-one-out tendency interacted with the tercile of initial market tightness, the tercile on its own, and the baseline controls and fixed effects. Clicks are restricted to the first three months of the spell; jobseekers with at least five clicks. Standard errors clustered at the caseworker level in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table E.15: Effects of caseworker redirection on click intensity by initial market tightness

	(1)	(2)
Dependent Var.:	Total clicks	Clicks per day
Leave-one-out tightness	-0.0048	-0.0163
	(0.0670)	(0.0672)
Leave-one-out tightness x Tercile 2	-0.0629	-0.0070
	(0.0896)	(0.0903)
Leave-one-out tightness x Tercile 3	-0.0089	0.0215
	(0.0901)	(0.0892)
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	100,525	100,525
Mean of dependent var.	41.992	0.50324

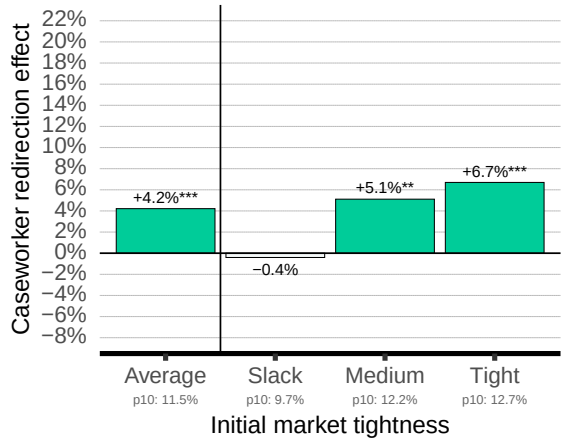
Notes: PPML regressions following Equation (3) with the leave-one-out caseworker tightness tendency interacted with terciles of the jobseeker's initial market tightness (last occupation × commuting zone at registration). Column (1) uses the total number of clicks over the first three months of the spell;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

Figure E.8: Heterogeneous effects on job finding by initial market tightness, with imputed found markets

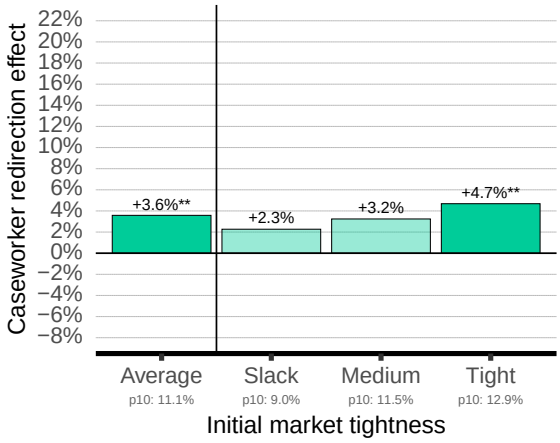


Notes: The figure repeats Figure 4, Panel (a), imputing the market of the new job for job finders whose found occupation or commuting zone is missing (Appendix B.6). After the imputation, 1.0% of job finders remain with an unknown market; this outcome is not plotted. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

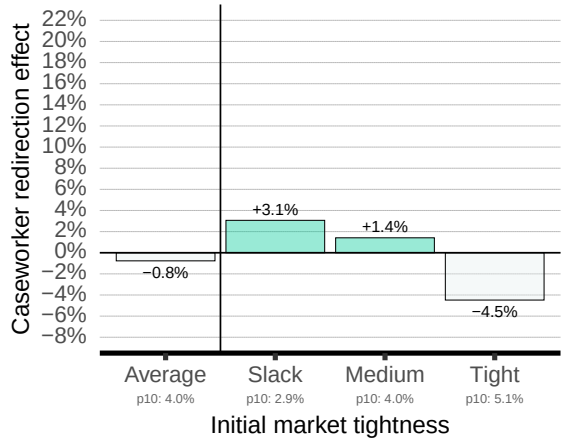
Figure E.9: Caseworker redirection effects on job finding by occupation × geography category



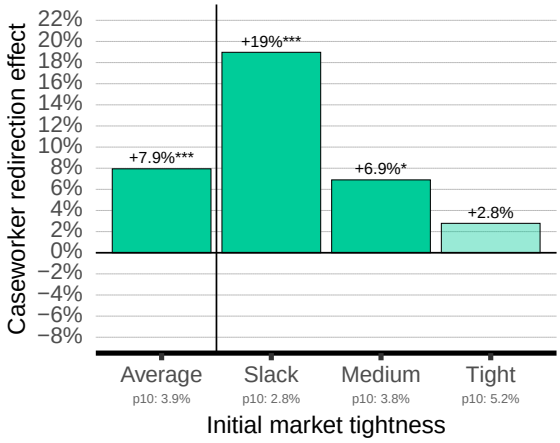
(a) Same occupation, same commuting zone



(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, and job finders whose occupation or commuting zone of the new job is not observed are left out. 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 E.16. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table E.16: Caseworker redirection and job finding by occupation × 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	-0.0013	0.0029	0.0067	0.0173***
	(0.0076)	(0.0048)	(0.0076)	(0.0046)
Leave-one-out tightness x Medium initial market	0.0217**	-0.0011	0.0055	-0.0087
	(0.0106)	(0.0063)	(0.0104)	(0.0065)
Leave-one-out tightness x Tight initial market	0.0291***	-0.0104	0.0130	-0.0126*
	(0.0112)	(0.0069)	(0.0111)	(0.0073)
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	0.016***	-0.001	0.013**	0.010***
	(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	333,514	333,514	333,514	333,514
R2	0.10366	0.07909	0.08048	0.07904

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 × commuting zone at registration). The outcome variable in each column is the unconditional probability of finding a job within six months in the given occupation × commuting zone category; jobseekers who do not find a job within six months are coded as zero. Job finders whose occupation or commuting zone of the new job is not observed are left out. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table E.17: Job quality with imputed found markets

	Market wage		Occupation value		Same market
	(1)	(2)	(3)	(4)	(5)
Dependent Var.:	Higher	Lower	Higher	Lower	Same market
Leave-one-out tightness	0.0225***	-0.0035	0.0147***	-0.0008	0.0148**
	(0.0049)	(0.0050)	(0.0042)	(0.0044)	(0.0060)
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
p-value equal within pair	0.000	0.000	0.003	0.003	
p10-p90 magnitude	+5.5%	-0.9%	+5.1%	-0.3%	+4.1%
N caseworkers	2,460	2,460	2,460	2,460	2,460
Observations	366,890	366,890	366,890	366,890	366,890
R2	0.09647	0.11548	0.10484	0.12087	0.09616
Mean of dependent var.	0.12847	0.12321	0.09063	0.09286	0.11310

Notes: The table repeats Columns (3)–(7) of Table 3, imputing the market of the new job for job finders whose market is not observed (22.1% of the jobs found within six months) from the industry of the new employer, as in Appendix Figure E.8 (Appendix B.6). For these job finders, each outcome is the share of the weight of the candidate markets of the given type; candidates without a market wage or an occupation value count as neither higher nor lower. Job finders whose market is observed are coded as in Table 3, including those whose market has no tightness measure, so no job finders are left out. The p-value equal within pair tests whether the effects on the two outcomes of a pair are equal. The table reports the effect of going from the 10th to the 90th percentile in the caseworker leave-one-out tendency, in percent of the predicted outcome at the 10th percentile. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table E.18: Job quality by initial market tightness and conditional on finding a job

Panel A: Job within six months of the given type, by tercile of initial market tightness

	Stable 360 days		Market wage		Occupation value		Same
	Stable	Not stable	Higher	Lower	Higher	Lower	market
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Leave-one-out tightness × slack initial market	0.0174*	0.0080	0.0169**	0.0098*	0.0180***	0.0037	-0.0013
	(0.0102)	(0.0068)	(0.0084)	(0.0057)	(0.0068)	(0.0049)	(0.0077)
Leave-one-out tightness × medium initial market	0.0230	0.0009	0.0251***	-0.0018	0.0168**	0.0041	0.0205**
	(0.0141)	(0.0083)	(0.0081)	(0.0077)	(0.0074)	(0.0062)	(0.0090)
Leave-one-out tightness × tight initial market	0.0325**	0.0017	0.0216***	-0.0056	0.0121**	0.0097	0.0282***
	(0.0161)	(0.0083)	(0.0079)	(0.0093)	(0.0056)	(0.0081)	(0.0097)
Tightness tercile FE	X	X	X	X	X	X	X
Unemployment office × year FE	X	X	X	X	X	X	X
Last occ. × commuting zone × quarter FE	X	X	X	X	X	X	X
Jobseeker controls	X	X	X	X	X	X	X
p-value equal within pair, slack	0.463	0.463	0.469	0.469	0.066	0.066	
p-value equal within pair, medium	0.181	0.181	0.009	0.009	0.154	0.154	
p-value equal within pair, tight	0.101	0.101	0.017	0.017	0.789	0.789	
p-value equal across terciles	0.686	0.737	0.740	0.258	0.739	0.789	0.020
p10-p90 magnitude, slack	+2.6%	+3.1%	+5.5%	+6.2%	+8.0%	+2.6%	-0.4%
p10-p90 magnitude, medium	+2.5%	+0.3%	+7.9%	-0.6%	+6.6%	+1.8%	+5.1%
p10-p90 magnitude, tight	+3.2%	+0.6%	+6.8%	-1.3%	+6.8%	+3.0%	+6.8%
N caseworkers	2,404	2,404	2,460	2,460	2,460	2,460	2,460
Observations	219,656	219,656	332,531	332,531	332,531	332,531	332,531
R2	0.15389	0.08243	0.09343	0.10406	0.09125	0.10747	0.10400
Mean of dependent var.	0.26046	0.08489	0.09939	0.09083	0.06964	0.07043	0.11778

Panel B: Conditional on finding a job

	Return to unemployment		Earnings		Occupation	Distance
	Within 180 days	Within 360 days	log earnings	log market wage	Value (SD)	log hours
	(1)	(2)	(3)	(4)	(5)	(6)
Leave-one-out tightness	0.0043	-0.0050	0.0086	0.0106***	0.0488***	-0.0085
	(0.0064)	(0.0090)	(0.0200)	(0.0034)	(0.0163)	(0.0135)
Unemployment office × year FE	X	X	X	X	X	X
Last occ. × commuting zone × quarter FE	X	X	X	X	X	X
Jobseeker controls	X	X	X	X	X	X
p10-p90 magnitude	+0.9%	-0.6%	+0.3%	+0.3%	+0.015 SD	-0.3%
N caseworkers	2,457	2,430	2,355	2,457	2,457	2,459
Observations	196,257	142,464	78,474	174,054	184,596	205,767
R2	0.08702	0.12336	0.28604	0.65238	0.51244	0.22649
Mean of dependent var.	0.14784	0.24345	8.6161	8.7267	0.0000	-0.88871

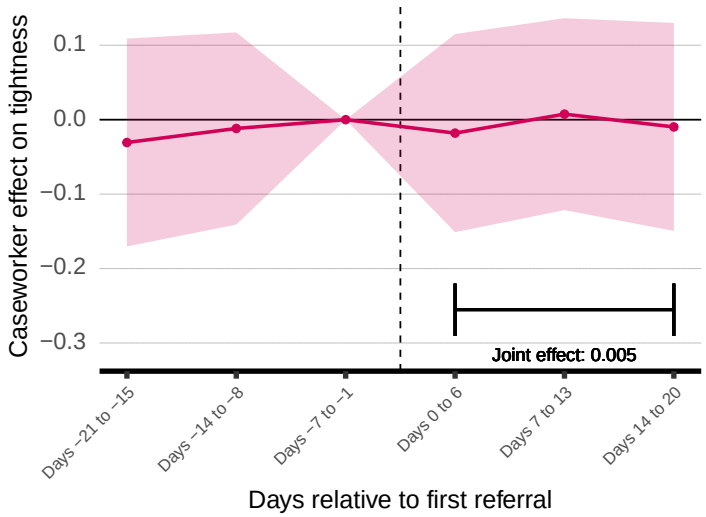
Notes: The table reports estimates from the specification in Equation (3) with jobseeker controls. Panel A repeats Table 3 by tercile of initial market tightness (tightness in the last occupation × commuting zone at registration), with a separate slope of the caseworker leave-one-out tendency in each tercile and tercile fixed effects; the outcomes are defined as in Table 3. The p-values equal within pair test, within each tercile, that the effects on the two outcomes of a pair are equal; the p-value equal across terciles tests that the slope is the same in all three terciles. Panel B conditions on finding a job. Columns (1) and (2): return to unemployment is an indicator for re-registering within 180 or 360 days of reemployment. Column (3): log monthly earnings in the new job, from social security records; since these records end in December 2021, the sample is restricted to deregistrations before October 2021. Column (4): log market wage, the median monthly full-time equivalent wage in the found occupation × commuting zone. Column (5): value of the found occupation, standardized among job finders. Column (6): log driving time between the municipality of residence and the municipality of the new job. Variables are defined in Section 3. Both panels report the effect of going from the 10th to the 90th percentile in the caseworker leave-one-out tendency, in Panel A within each tercile: for the indicators in percent of the predicted outcome at the 10th percentile, for the log outcomes in Panel B, Columns (3), (4) and (6), as the percent change in the level ($e^{\Delta} - 1$, where Δ is the change in the predicted log outcome), and in standard deviations for the occupation value. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table E.19: Search around the first referral and the first meeting, and jobs found through referrals

	Around the first referral		Around the first meeting	Spell level
	(1)	(2)	(3)	(4)
Dependent Var.:	Clicks per day	Clicked market tightness	Clicked market tightness	Job through a referral
Post event	0.0170***	-0.0068	-0.0027	
	(0.0027)	(0.0080)	(0.0071)	
Post event × leave-one-out tightness	-0.0048	0.0054	0.1343***	
	(0.0124)	(0.0478)	(0.0484)	
Leave-one-out tightness				-0.0030
				(0.0057)
Day of week FE	X			
Click day of week FE		X	X	
Unemployment spell FE	X	X	X	
Calendar month FE	X	X	X	
Click hour-of-day FE		X	X	
Distance to the other event FE	X	X	X	
Baseline FE and jobseeker controls				X
Sample	Spell-day panel	Click level	Click level	Spell level
Share of spells with a referral				0.348
Referrals per spell				1.29
N spells	108,299	13,770	13,593	101,511
Observations	4,192,496	603,470	558,700	101,511
Mean of dependent var.	0.19388	-0.94562	-0.60583	0.03714

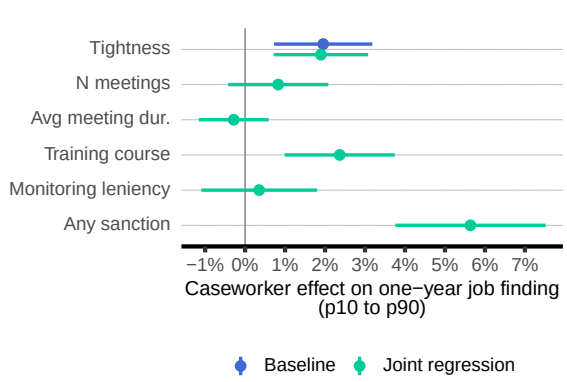
Notes: Columns (1)–(3) report pooled difference-in-differences estimates, in Column (2) for the event study in Appendix Figure E.10: the post-event indicator and its interaction with the leave-one-out tightness tendency. Columns (1) and (2): around the first referral by the jobseeker's own caseworker (application request or candidate proposal), controlling for the time since the first meeting; Column (3): around the first meeting, controlling for the time since the first referral. Window of 21 days before and 20 days after the event, time in calendar days. Column (1): clicks per day in a panel of registered days, winsorized at the 95th percentile of days with at least one click. Columns (2) and (3): log tightness of the clicked market at the click level, jobseekers with at least three clicks with defined tightness before and after the event. Column (4): indicator for finding a job through a referral by the jobseeker's own caseworker, at the spell level, with the specification in Equation (3); sample: clickers in the baseline sample. Standard errors clustered at the caseworker level in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure E.10: Search tightness around the first referral

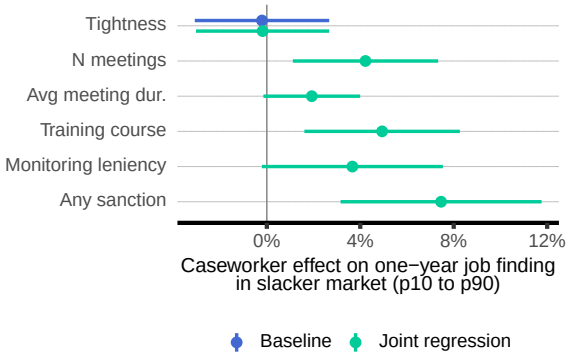


Notes: The figure reports event-study estimates of the interaction between the leave-one-out tightness tendency and seven-day bins around the first referral by the jobseeker's own caseworker, estimated at the click level as in Equation (4). Timing is in calendar days; day 0 is the day of the referral and the omitted bin spans days −7 to −1. The outcome is the log tightness of the clicked market. The specification includes spell, calendar-month, click-weekday and click-hour fixed effects and seven-day bins of the time since the first meeting. The joint effect pools all post-referral bins in a difference-in-differences specification (0.005, standard error 0.048; Appendix Table E.19, Column (2)). Sample: jobseekers with at least three clicks with defined tightness before and after the first referral. Shaded areas are 95% confidence intervals; standard errors are clustered at the caseworker level.

Figure E.11: Longer horizon: effect of the tightness tendency on job finding within one year, controlling for other caseworker tendencies



(a) Job finding within one year in any market



(b) Job finding within one year in a slacker market

Notes: The figure repeats Figure 6 at a longer horizon: the outcome is finding a job within one year (365 days) instead of six months. Panel (a): finding a job within one year. Panel (b): finding a job within one year in a market slacker than the initial market; jobseekers without a job, or with a job in another market type, are coded as zero, and job finders whose market is not observed or has no tightness measure are left out. Spells registered after 31 March 2022, which lack a full year of follow-up, are excluded. Standard errors are clustered at the caseworker level.

Table E.20: Effect of the tightness tendency on the components of search tightness

	Total	Component of own clicks tightness			
	(1)	(2)	(3)	(4)	(5)
Dependent Var.:	Own clicks tightness	National persistent	Local persistent	National transitory	Local transitory
Leave-one-out tightness	0.1227***	0.0976***	0.0193**	0.0065**	-0.0007
	(0.0209)	(0.0174)	(0.0085)	(0.0029)	(0.0022)
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
Share of effect		0.80	0.16	0.05	-0.01
Share of variation across jobseekers within market-quarter		0.77	0.17	0.03	0.02
N caseworkers	2,460	2,460	2,460	2,460	2,460
Observations	101,511	101,511	101,511	101,511	101,511
R2	0.73797	0.57522	0.88912	0.94554	0.51786
Mean of dependent var.	-0.57700	-0.60501	0.01998	0.00381	0.00423

Notes: The table reports estimates from the specification in Equation (3). Column (1): the search tightness index, the mean log tightness of the markets clicked in the first three months. Columns (2)–(5): the four parts of the index, which add up to it. The parts are computed from the tightness of all markets (occupation × commuting zone) on all days of the sample period. The log tightness of each market on each day is split into a national persistent part (average of the occupation over commuting zones and days), a local persistent part (average deviation of the occupation in the commuting zone from the national level), a national transitory part (deviation of the occupation's national tightness on the day from its average) and a local transitory part (remainder). "Share of effect" is each coefficient divided by the coefficient in Column (1). "Share of variation across jobseekers within market-quarter" shows how much of the differences in the search tightness index between jobseekers come from each part. Each share is the covariance of the part with the index divided by the variance of the index, both after removing the controls and fixed effects of the regression, so that only differences between jobseekers with the same initial market and registration quarter remain. The shares add up to one. Comparing them with the shares of the effect shows whether caseworkers shift search along the four parts in the same proportions as the actual search tightness differs between jobseekers. Sample: clickers. Standard errors are clustered at the caseworker level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table E.21: Caseworker characteristics by tightness redirection tendency

	P10 caseworker	P90 caseworker	Beta	p-value	q-value
Caseload					
Sample period caseload	194.9	197.0	6.667	0.449	0.578
Clients using job-room	52.4	53.0	1.970	0.482	0.578
Meetings per month	7.0	7.0	0.238	0.449	0.578
Experience					
Years of experience	6.8	6.6	-0.665	0.434	0.578
Client diversity					
Distinct 3-digit occupations	54.8	55.0	0.609	0.67	0.67
Distinct 2-digit industries	43.3	44.1	2.428	0.017**	0.104

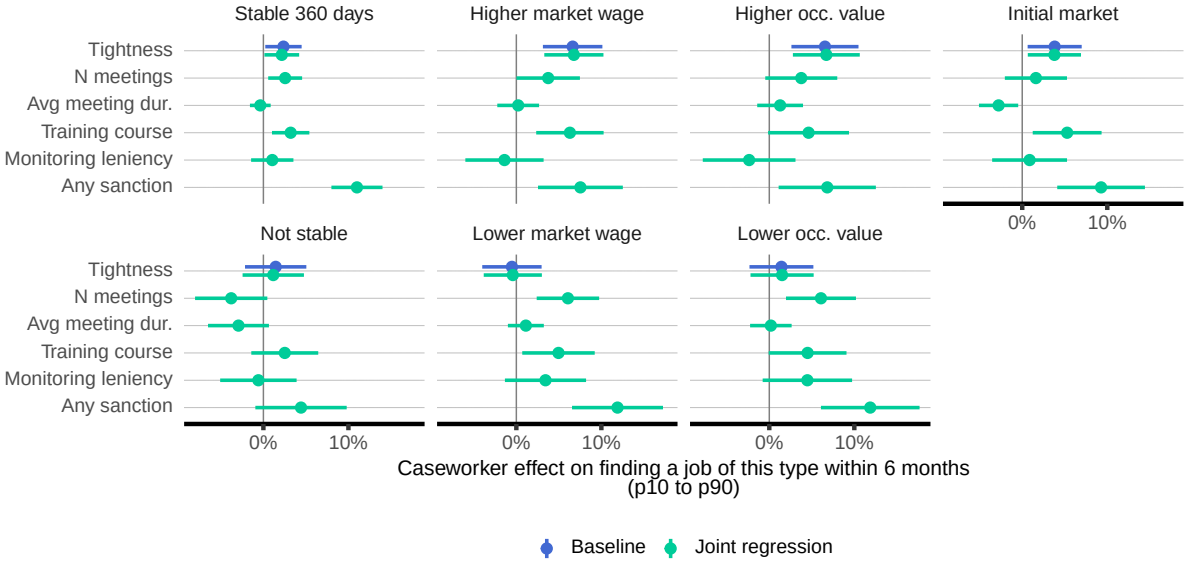
Notes: The table compares observable characteristics of caseworkers at the 10th and 90th percentiles of the leave-one-out tightness tendency within their employment office. The first two columns report average predicted values when the tightness tendency is set to the 10th and 90th percentile, respectively, with all other covariates held at their observed values. Beta reports the coefficient from regressions following Equation (3), with the caseworker characteristic as an outcome. 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. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table E.22: Tightness tendency and the jobseeker's own meetings, training, monitoring and sanctions

	Support			Stringency	
	(1)	(2)	(3)	(4)	(5)
Dependent Var.:	N meetings	Meeting dur.	Training course	Monitoring leniency	Any sanction
Tightness tendency	-0.0926**	-0.0102	-0.0057	0.0026	0.0068
	(0.0454)	(0.0193)	(0.0106)	(0.0048)	(0.0110)
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	-0.8%	+0.5%	-1.0%	+0.1%	+0.7%
N caseworkers	2,460	2,460	2,460	2,460	2,460
Observations	366,890	366,890	366,890	339,394	366,890
R2	0.33846	0.40760	0.09097	0.05753	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 jobseeker controls. The outcome variables are the jobseeker's own outcomes in the first three months: the number of meetings per month, the average log duration of a meeting, participation in a training course, the share of monthly search effort protocols rated sufficient, and a benefit sanction. The regressor is the leave-one-out tightness tendency. The magnitude row reports the change in the outcome when the tendency moves from the 10th to the 90th percentile, in percent of the predicted outcome at the 10th percentile. Standard errors are clustered at the caseworker level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Figure E.12: Caseworker dimensions and the type of job found



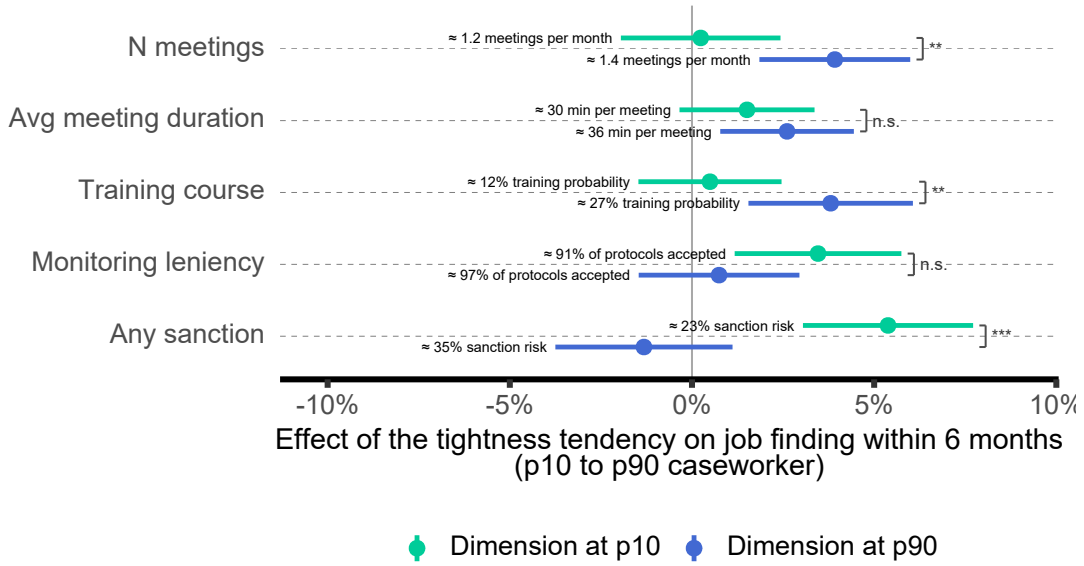
Notes: The figure reports the effect of moving each leave-one-out caseworker measure from its 10th to its 90th percentile on the probability of finding a job of a given type within six months, in percent of the predicted outcome at the 10th percentile, with 95% confidence intervals. Outcomes as in Table 3. Estimates as in Figure 6. Standard errors are clustered at the caseworker level.

Table E.23: Job finding on all caseworker dimensions jointly

	Find job within 6 months					
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Var.:	In any market	In slacker market	In any market	In tighter market	In initial market	In slacker market
Leave-one-out tightness	0.0290***	-0.0002	0.0288***	0.0198***	0.0144**	-0.0001
	(0.0092)	(0.0050)	(0.0087)	(0.0059)	(0.0061)	(0.0050)
Leave-one-out N meetings			0.0017	0.0086***	0.0027	0.0042*
			(0.0043)	(0.0027)	(0.0031)	(0.0023)
Leave-one-out meeting dur.			-0.0200*	0.0015	-0.0171**	0.0059
			(0.0109)	(0.0062)	(0.0072)	(0.0054)
Leave-one-out training course			0.0746***	0.0265**	0.0354**	0.0345***
			(0.0187)	(0.0123)	(0.0138)	(0.0105)
Leave-one-out monitoring leniency			-0.0326	-0.0269	0.0138	0.0486*
			(0.0454)	(0.0352)	(0.0357)	(0.0276)
Leave-one-out any sanction			0.1604***	0.0509***	0.0599***	0.0496***
			(0.0245)	(0.0148)	(0.0170)	(0.0128)
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
p10-p90: tightness	+2.4%	-0.1%	+2.4%	+6.1%	+3.8%	-0.0%
p10-p90: N meetings			+0.3%	+6.1%	+1.6%	+3.4%
p10-p90: meeting dur.			-1.1%	+0.3%	-2.8%	+1.3%
p10-p90: training course			+3.5%	+4.5%	+5.3%	+6.9%
p10-p90: monitoring leniency			-0.6%	-1.9%	+0.9%	+4.1%
p10-p90: any sanction			+7.9%	+9.1%	+9.3%	+10.2%
N caseworkers	2,460	2,460	2,460	2,460	2,460	2,460
Observations	338,885	310,036	338,885	310,036	310,036	310,036
R2	0.14515	0.12051	0.14602	0.09607	0.10690	0.12072
Mean of dependent var.	0.36785	0.08891	0.36785	0.10240	0.11773	0.08891

Notes: The table reports estimates from the specification in Equation (3). Columns (1) and (2) include only the leave-one-out tightness tendency; Columns (3)–(6) include all six leave-one-out caseworker measures. The outcome is finding a job within six months: in any market in Columns (1) and (3), in a slacker market in Columns (2) and (6), in a market tighter than the initial market in Column (4), and in the initial market itself, with the same occupation and commuting zone, in Column (5). For the market-type outcomes, jobseekers without a job, or with a job in another market type, are coded as zero, and job finders whose market is not observed or has no tightness measure are left out. The p10–p90 rows report the effect of moving each measure from its 10th to its 90th percentile, in percent of the predicted outcome at the 10th percentile. Sample: spells with all six measures defined. Standard errors are clustered at the caseworker level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure E.13: Effect of the tightness tendency by the caseworker's other dimensions



Notes: The figure reports the effect of moving from the 10th to the 90th percentile of the leave-one-out tightness tendency on job finding within six months, in percent of mean job finding (0.368), with 95% confidence intervals. Estimates from Equation (3) with the tightness tendency interacted jointly with five leave-one-out caseworker measures of Section 6.4 and the measures as controls; each measure is standardized. Appendix Table E.11, Column (3), reports the regression estimates. For each measure, the effect is evaluated with this measure at its 10th or 90th percentile and the other measures at their means. Labels: the jobseeker's corresponding behavior in the first three months, predicted at the 10th and 90th percentile of the caseworker measure with the other covariates at their means (meetings per month of unemployment in the first 90 days, minutes per meeting, probability of a training course, share of monthly search effort protocols accepted, probability of a sanction). Brackets: significance of the difference between the two effects, i.e. of the interaction (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$, n.s. not significant). Sample: spells with all five measures defined (338,885 spells, 2,460 caseworkers). Standard errors are clustered at the caseworker level.

Table E.24: Caseworker measures and the jobseeker's own search, training and sanctions

	Each measure alone			All six measures jointly		
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Var.:	Own clicks tightness	Own training course	Own sanction	Own clicks tightness	Own training course	Own sanction
Leave-one-out tightness	0.1185***			0.1168***	-0.0045	0.0034
	(0.0209)			(0.0210)	(0.0035)	(0.0045)
Leave-one-out training course		0.8345***		-0.0037	0.8199***	0.0493***
		(0.0095)		(0.0292)	(0.0101)	(0.0085)
Leave-one-out any sanction			0.7055***	0.0511	0.0090	0.5706***
			(0.0104)	(0.0382)	(0.0088)	(0.0143)
Leave-one-out N meetings				-0.0114	0.0100***	-0.0001
				(0.0071)	(0.0017)	(0.0020)
Leave-one-out meeting dur.				-0.0095	0.0028	-0.0100*
				(0.0173)	(0.0041)	(0.0057)
Leave-one-out monitoring leniency				0.1103	-0.0738***	-0.4363***
				(0.0854)	(0.0189)	(0.0283)
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
p10-p90 change in outcome	+0.036	+0.143	+0.123	+0.036	+0.141	+0.099
N caseworkers	2,460	2,460	2,460	2,460	2,460	2,460
Observations	98,648	338,885	338,885	98,648	338,885	338,885
R2	0.73981	0.11561	0.09183	0.73982	0.11573	0.09241
Mean of dependent var.	-0.58241	0.19111	0.28641	-0.58241	0.19111	0.28641

Notes: The table reports estimates from the specification in Equation (3). Outcomes: the search tightness index (clickers only), participation in a training course in the first three months, and a benefit sanction in the first three months. Columns (1)–(3) include only the matching leave-one-out caseworker measure; Columns (4)–(6) include all six measures. The p10–p90 row reports the change in the outcome when the matching caseworker measure moves from its 10th to its 90th percentile (log points for the search tightness index, probability for training and sanctions). Sample: spells with all six measures defined. Standard errors are clustered at the caseworker level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.