The Labor Hub

Job Market Paper Series No. 13

Advisor: Eric Maurin

DO YOU NEED AN EXTRA APPLICANT? EVIDENCE FROM QUASI-RANDOM CASEWORKER ASSIGNMENT

Marie-Renée ANDREESCU

Do you need an extra applicant? Evidence from quasi-random caseworker assignment

Marie-Renée Andreescu * (Paris School of Economics)

This version: October 20, 2026

Most recent version available here

Abstract

Firms face increasingly large applicant pools: more candidates offer more opportunities to find a suitable worker, but also more applications to assess. This paper studies whether expanding employers' choice sets improves recruitment or makes decisions harder. Using quasi-random assignment of employers to caseworkers who screen applications in the French Employment Service, I estimate the effect of an additional candidate reaching the recruiter. One additional candidate raises the probability of a contract proposal within three months by 10% and of a hire within six months by 7%, with strongly diminishing returns. Longer shortlists have lower average fit but contain more suitable candidates, with no detectable decline in the observable quality of recruited workers. Where candidates are scarce, firms value an extra candidate more but convert fewer proposals into hires.

Keywords: Hiring costs; Choice sets; Public Employment Services; Job Matching; Instrumental Variables.

JEL codes: J63, J64, J68, D83.

*This paper greatly benefited from discussions with Luc Behaghel, Simon Boutin, Thomas Breda, Lucie Briant, Carmen Durrer de la Sota, Leonardo Geretto, Pietro Geuna, Mattis Gilbert, Marc Gurgand, Nina Guyon, Yagan Hazard, Jeremias Klauei, Rafael Lalive, David Margolis, Eric Maurin, Dominique Meurs, Mitia Oberti, Roland Rathelot, Pierre Rousseaux, Joyce Sultan-Parraud, and Todor Tochev, and from helpful comments from participants at the Paris School of Economics Labour Seminar, the Labour Chair Tea Seminar, the IZA-LISER Summer School, the 14th Warwick Economics PhD Conference, The Role of Firms in the Labour Markets Workshop in Lugano, and the Uppsala Workshop in Labour Economics WEI IV. I gratefully acknowledge the support of the EUR grant ANR-17-EURE-0001. All remaining errors are my own.

1. Introduction

Firms searching for prospective employees are confronted with a growing number of applications to sort through. The use of the Internet has decreased search costs for job seekers, who can apply more easily to an increasing number of vacancies (Goldfarb and Tucker, 2019; Birinci et al., 2025). From 2017 to 2025, the average number of applications received by a firm posting a vacancy in France on at least one job board increased from 10 to 17. For a recruiter, each additional candidate is one more chance of finding a suitable worker, but also one more application to read, assess, and possibly interview. A large literature on individual decision making shows that expanding choice sets can degrade decisions: when options multiply, consumers may choose less often and less well, and reducing the choice set can improve outcomes (Huberman and Jiang, 2006; Scheibehenne et al., 2010; Caplin et al., 2019). Whether this logic carries over to firms is an open question. Firms are not individual consumers: they can delegate and spread screening across workers, but they still face monetary screening costs, which search-and-matching theory places at the center of demand-side labor market frictions (Pissarides, 2000; Mortensen and Pissarides, 1994). Yet little is known about how firms adjust their recruitment decisions when the set of candidates they consider expands.

This paper studies the trade-off between expanding the number of candidates made available to recruiters and increasing screening costs. Identifying empirically the causal effect of applicant pool size on hiring outcomes is difficult. The number of candidates a firm considers is endogenous: harder-to-fill vacancies may require larger pools, firm characteristics might make some vacancies more attractive and unobserved elements can affect both the number of applications processed and the probability of a successful match. Firms source candidates through several recruitment channels. This paper focuses on one of these sources, the French Public Employment Service (FPES), Pôle emploi, where the data allow me to observe precisely which candidates reach the recruiter and how the firm responds. Its assisted preselection service provides quasi-experimental variation in the number of candidates made available through this channel, allowing me to overcome the identification challenge by using a judge-design framework. The PES operates a large online job board and offers an employer-directed service called assisted preselection. When a firm opts into preselection for a vacancy, a caseworker screens the applications and transmits a shortlist through the recruiter interface: only shortlisted candidates' applications reach the employer through this channel. The size of the shortlist therefore provides a direct measure of the candidate pool accessible to the recruiter from this source, alongside the candidates the firm may receive through other recruitment channels. Importantly, when a firm first uses preselection, it is assigned a caseworker quasi-randomly within its local agency. Caseworkers differ persistently in how many candidates they transmit. I construct a predetermined measure of each caseworker's leniency from her screening decisions in the quarters before she first handles a given firm, and use it as an instrument for shortlist size

in a judge-design framework (Kling, 2006; Frandsen et al., 2023; Goldsmith-Pinkham et al., 2025). Identification comes from comparing vacancies handled in the same agency and period but assigned to caseworkers with different pre-measured leniency. This method allows me to assess the marginal impact of transmitting an additional candidate on firm decision-making.

When caseworkers transmit candidates to firms, their choice reflects different dimensions: individual preferences for some characteristics, leniency or ability to screen. This is a common characteristic in the literature on multidimensional judges (Frandsen et al., 2023). In this paper, I focus on only one dimension of the caseworker, her *leniency*. Leniency is understood as a lower transmission threshold, i.e., when given a ranked list of candidates, a more lenient caseworker will go further down the list and transmit more people. In order to obtain an accurate identification of the quantity effect, it is of utmost importance for the exclusion restriction to hold. This means establishing that caseworkers who transmit more candidates are not also better at identifying suitable applicants or do not have any specific preference for some candidates. I conduct several tests to make sure that leniency is independent of other dimensions. Mainly, I construct a predetermined measure of how well caseworkers distinguish applicants on observable suitability, i.e., how accurate they are in establishing the match between candidate and vacancy. This measure is not significantly associated with leniency, and controlling for it leaves the estimated effect on contract proposals unchanged. Leniency is also not correlated with preferences for women, foreigners or length of unemployment. These checks support the interpretation of leniency as a lower transmission threshold. Under this interpretation, marginal candidates are applicants who reach the employer because the assigned caseworker applies a less restrictive cutoff and who would otherwise have been screened out.

I organize the analysis around three questions that together describe the firm's decision problem. First, the extensive margin of choice: does an additional transmitted candidate make the firm more likely to select someone from the shortlist instead of using an outside option? Instrumental-variable estimates show that one additional candidate raises the probability that the firm makes a contract proposal from the shortlist within three months by 3.1 percentage points, a 10 percent increase relative to the baseline. The returns are strongly diminishing: the marginal effect is largest when the initial shortlist is small and attenuates steadily as the pool grows, without ever turning systematically negative in the observed range. This is true even after accounting for proportional expansion: a 10 percent larger shortlist raises the probability of a contract proposal by approximately 9.3 percentage points in the bottom quintile of predicted shortlist size, compared with 1.0 percentage point in the top quintile. The gains are concentrated where an extra draw is most likely to matter: low-wage and low-skilled vacancies, where more candidates may be deemed suitable, and tight labor markets, where candidates are scarce.

These gains extend beyond the firm's decision to propose a contract. Linking recruit-

ment decisions to administrative hiring declarations, I find that one additional transmitted candidate raises the probability of hiring from the shortlist within six months by 0.8 percentage points, a 6.5 percent increase relative to the baseline. As with contract proposals, the hiring gains are concentrated among vacancies with small predicted shortlists. Expanding the choice set therefore helps firms recruit, although the smaller hiring response highlights that finding a candidate the employer wishes to hire does not always result in a completed match.

Second, the composition of the choice set: what does the marginal candidate look like? Expansion is not neutral. Candidates added at the margin by more lenient caseworkers are somewhat less experienced and fit the vacancy less well: they are more likely to come from other occupations and less likely to meet the posted educational requirements, so that average fit in the shortlist declines. At the same time, longer shortlists contain more candidates above a basic quality threshold, even as the share of such candidates falls. Expansion therefore adds noise to the pool while simultaneously broadening the set of usable options.

Third, the quality and speed of the firm's selection: does the added noise change final decisions? I find that firms assigned to more lenient caseworkers do not select candidates who are observably different in demographics, labor market quality, fit to the vacancy, or relative rank within the pool. Recruitment timing is also broadly unchanged on average, but this masks a longer process where candidates are scarce. In the tightest labor markets, one additional transmitted candidate increases the time to the last contract proposal by 3.3 days, relative to a baseline of 25 days. This seems to be due to a change in acceptance rates when candidates are scarce, rather than to an inability to process a larger volume of applications. Firms appear able to filter signal from noise, at least when an intermediary has already removed the least relevant applications.

A conceptual framework connects these findings by distinguishing the caseworker's transmission cutoff, which determines who reaches the recruiter, from the firm's inspection frontier, which determines who is worth evaluating. On the one hand, a lower transmission cutoff can reduce average shortlist fit while expanding the number of suitable options. On the other, when firms can rank applications cheaply and inspect selectively, additional candidates provide opportunities they can take or leave. In this case, candidates are inspected only when they might improve on the firm's current best, which allows their presence to weakly raise expected payoff, while never lowering it. However, expansion could also raise the cost of evaluating existing candidates, if the firm chooses to inspect everyone. The benefit of the extra candidate competes then with the increased burden of inspection. Candidate acceptance adds a further tension: in tight labor markets, weaker recruitment alternatives can make additional options more valuable to firms, while candidates' stronger outside opportunities can make matches harder to complete. The framework organizes the evidence on shortlist composition, proposals, hires and

recruitment timing around the trade-off between additional opportunities and the costs of turning them into matches.

This paper contributes to four strands of literature. First, it provides evidence on how firms turn applicant pools into hires. Hiring costs are a central demand-side friction in search-and-matching theory (Mortensen and Pissarides, 1994; Pissarides, 2000), and empirical work has documented their magnitude (Abowd and Kramarz, 2003; Blatter et al., 2012; Muehlemann and Leiser, 2018), but the internal mechanics of recruitment remain largely a black box. Recent contributions have begun opening it, documenting congestion in application flows (Horton et al., 2024), the value of referrals (Pallais and Sands, 2016), and the way vacancy design shapes applicant pools (Belot et al., 2022; Burn et al., 2022; Balgova et al., 2025; Card et al., 2025). Less is known about what actually happens between the moment the firm receives an application and the selection of a worker. Understanding this step is essential as initial screening is increasingly delegated to an intermediary or an algorithm (AI-powered or not): when such systems are used as first-stage screeners, firms no longer make decisions from the raw applicant pool; they rely on an algorithmically curated consideration set. Recent work studies the design and consequences of these screening rules (Blair and Guo, 2026; Cowgill, 2020; Li et al., 2026). My contribution concerns a complementary decision faced by both human and algorithmic screeners: how many candidates should pass the initial screen? I identify the recruitment response to expanding this set and trace the accompanying changes in candidate composition, employer proposals and realized hires.

Second, the paper brings quasi-experimental evidence from firms to the literature on choice sets and costly decision making. Work on rational inattention and consideration sets (Caplin et al., 2019) and on choice overload (Huberman and Jiang, 2006; Scheibehenne et al., 2010) focuses on individual decision makers, while congestion at the firm level is often hypothesized but rarely measured. In an institutional setting, Carvalho et al. (2019) examine the interaction between congestion and match quality in college admissions. I study a related trade-off within recruitment, where an intermediary determines the options available but the firm retains the final choice. Observing both the shortlist and the selected worker allows me to distinguish a decline in average pool fit from a deterioration in the firm's selection. I show that larger, less selective shortlists can increase recruitment without a detectable decline in selected workers' overall observable quality.

Third, the paper adds to the incipient literature on active labor market policies (ALMP) directed at firms. A large part of the literature looking at ALMPs focuses on policies that operate on the supply side of the labor market (Card et al., 2010), rather than on labor demand and firms' recruiting costs (Pallais, 2014; Carranza et al., 2022; Abel et al., 2020). Recent work also shows that the value of search advice depends on labor-market conditions, from the job seekers' side (Klaeui, 2026). This paper examines the complementary firm-side margin, by examining how many candidates recruiters get to consider. The work by Algan

et al. (2020) is particularly important for this paper, as the authors study the randomized introduction of a package of ALMPs aimed at firms, which was composed of preselection, vacancy writing assistance and interview support. This experiment took place in France in 2014, predating the current framework and comparing small and medium-sized enterprises (SMEs) that were not assisted by caseworkers with ones that were. I study the preselection program at scale and in steady state, outside the pilot phase. I identify the effect of relaxing access to an already screened applicant pool and show where additional transmissions generate the largest recruitment gains. I am also able to document increased subsequent use of the service, when the shortlist increases. This provides evidence on how to improve the design of a labor market policy: it is essential to adapt candidate provision to the firm's characteristics and the labour market's constraints.

Finally, this paper takes advantage of the econometric literature and advancements in judge designs (Goldsmith-Pinkham et al., 2025; Frandsen et al., 2023). The use of this identification strategy has moved from only crime studies (Kling, 2006) and has been used for looking at caseworkers' impact on job seekers' enrollment in programs and job search behavior (Humlum et al., 2025; Arni and Schiprowski, 2019). This paper looks at a specific randomization of caseworkers assigned to firms, and not to job seekers. I provide an empirical application of the judge design that makes both the screening decision and the subsequent employer response observable, while explicitly examining the conditions under which caseworker variation identifies the effect of expanding access to candidates.

The remainder of the paper is organized as follows. Section 2 presents the conceptual framework. Section 3 describes the institutional setting. Section 4 presents the data. Section 5 details the empirical strategy. Section 6 reports the effects of shortlist expansion on contract proposals, realized hiring and subsequent recruitment. Section 7 examines changes in shortlist composition, selected candidates and acceptance among proposed candidates. Section 8 studies how recruitment effects and duration vary with labor-market tightness. Section 9 discusses the broader implications of the results. Section 10 concludes.

2. Conceptual framework

This section develops a framework for how firms use candidates transmitted by the public employment service. I differentiate between *screening*, that is, the first filtering step that can be delegated to the PES caseworkers, HR professionals or an algorithm, and *inspection*, which involves bigger costs for the firm, such as interviewing, but allows firms to make a direct assessment of the candidate's potential value added. A longer shortlist provides additional opportunities to find a suitable worker, but can also increase the burden of evaluating applications. Appendix B provides the supporting assumptions and derivations. Appendix C provides numerical illustrations of shortlist expansion, screening precision and candidate acceptance.

2.1. Applicants and caseworker screening

I consider a vacancy j, with a finite pool of applicants. i identifies a candidate and c the caseworker assigned to the vacancy. When recruiting candidate i, the firm can expect a surplus of the form:

$$s_{ij} = m_{ij} + \varepsilon_{ij}, \tag{1}$$

where m_{ij} is the expected fit of the candidate to the vacancy and ε_{ij} is the residual component that can only be revealed through inspection.[1] Fit can reflect relevant experience, location, or wage compatibility.

Before the firm's assessment of the fit of the candidate, having an intermediary assess applicants means that some noise is introduced. Thus the caseworker's signal and transmission decision take the form:

$$\widehat{m}_{ijc} = m_{ij} + \sigma_c \eta_{ijc}, \qquad D_{ijc} = \mathbf{1}\{\widehat{m}_{ijc} \geq \theta_c\}, \qquad \eta_{ijc} \sim \mathcal{N}(0,1). \tag{2}$$

where $\widehat{m}_{ijc}$ is the assessed fit, η_{ijc} is the assessment error, and $\sigma_c > 0$ measures its magnitude. D_{ijc} is equal to one when the application is transmitted.

In order to understand the caseworker's action, two parameters need to be considered that capture the different aspects of screening. On the one hand, *leniency* designates the cutoff θ_c. A more lenient caseworker has a lower θ_c and, for a given applicant pool, transmits weakly more candidates. This is a quantity margin: it changes how many candidates reach the firm. The second dimension is the *screening quality*, that is, the precision with which $\widehat{m}_{ijc}$ tracks true fit m_{ij}, indexed by σ_c. A lower σ_c means the caseworker is better at telling well-matched from poorly matched applicants. This is a quality margin: it changes which candidates reach the firm at a given shortlist size. Better precision then improves average fit among the candidates transmitted. Leniency and screening quality can vary together across caseworkers, so a larger shortlist need not be the result of a lower cutoff alone.

The paper focuses on leniency as an instrument, which is understood as the comparison of different cutoffs, holding noise and assessment errors fixed. A lower cutoff then expands the shortlist without removing previously transmitted candidates. This means that candidates admitted at lower signal values are expected to have a lower fit. Section 5.3 establishes the empirical independence of leniency and screening quality, although it is not assumed at this stage of the model.

2.2. Inspection and the firm's decision rule

Environment The firm receives the shortlist and chooses which candidates to inspect. I assume that firms are able to observe m_{ij} at negligible cost, while learning ε_{ij} requires further evaluation.

[1] The residuals are assumed to have a common distribution function and zero mean, and to be independent across candidates and of fit and caseworker assessment errors.

The cost of inspecting one candidate depends on the workload generated by the full shortlist:

$$\kappa_j(n) = \kappa_{0j} + \lambda \frac{n}{K_j}, \qquad \kappa_{0j} > 0, \quad K_j > 0, \quad \lambda \geq 0. \tag{3}$$

Here, n is the number of shortlisted candidates and κ_{0j} is the baseline cost of evaluating one candidate for vacancy j. Recruitment capacity K_j captures the total staff and organizational resources available to handle applications, so that n/K_j measures workload relative to capacity. Having an extra candidate is not the same for a firm of only one person, who has to stop their activity in order to inspect the extra candidate, versus a firm with a team of HR professionals that can dedicate their time to the process. Then, λ represents how much harder it becomes to assess one candidate when the firm has more applications to handle. When $\lambda = 0$, additional applicants do not make each evaluation more costly. When $\lambda > 0$, a larger shortlist creates a shared burden, such as coordination demands or interruptions, that remains even when some candidates are left uninspected. This is the channel through which shortlist expansion can create congestion.

The firm has an outside option τ_j^0, which is the value of recruiting through another channel or leaving the position unfilled. At any time, the firm has in hand an option, which is its best available surplus, denoted by b_j. Initially, $b_j = \tau_j^0$; it subsequently becomes the maximum of the outside option and the surpluses of inspected candidates. In this setting, candidates accept a proposal and remain available if the firm wishes to return to them. This assumption is relaxed later on. The firm chooses inspections and a final option to maximize expected surplus net of inspection costs.

Decision rule Under these assumptions, the firm can use the optimal sequential-search rule given by Weitzman (1979): the firm evaluates the most promising candidates first and continues only while the expected gain from another evaluation exceeds its cost. This condition is

$$\mathbb{E}\big[\max\{m_{ij} + \varepsilon - b_j, 0\}\big] > \kappa_j(n), \tag{4}$$

The left-hand side measures how much the candidate might improve on that option, on average. A disappointing evaluation does not force the firm to accept a worse match: it can return to its previous best option. The firm moves down the list of candidates ranked by observed m_{ij}. As the firm goes on, the expected gain from further inspection falls. When the next candidate is not worth evaluating anymore, the firm stops.

This decision rule defines an *inspection frontier*: the minimum expected fit that makes a candidate worth evaluating. This threshold is $\underline{m}_j$, defined by

$$\mathbb{E}\Big[\max\{\underline{m}_j + \varepsilon - b_j, 0\}\Big] = \kappa_j(n). \tag{5}$$

At the frontier, the expected gain from evaluating a candidate exactly equals its cost. Higher inspection costs make the firm more selective and raise this threshold. Finding a better

candidate also raises it, because the remaining candidates must now improve on a stronger option. Before any inspection, the firm's benchmark is its outside option. Candidates below this initial frontier are never inspected; those above it may be evaluated if earlier candidates do not make further search unnecessary. The caseworker's transmission cutoff θ_c therefore determines who reaches the firm, while the inspection frontier determines who is worth evaluating.

Additional options and congestion A lower transmission cutoff at unchanged precision expands the firm's options. When $\lambda = 0$, the firm can retain its original search policy and ignore the additional candidates, so its expected net payoff cannot fall. In this case, increasing the shortlist size weakly increases the probability of making a contract proposal. The additional candidate matters when they are worth inspecting, provided that the existing options do not already produce a satisfactory choice. This mechanism can produce smaller gains among vacancies with larger candidate pools. An additional option is less useful when the firm is already likely to find a satisfactory worker.

When $\lambda > 0$, expansion also raises the cost of inspecting candidates who were already available. At a given b_j, the inspection frontier moves up. Some candidates worth inspecting on the shorter list may no longer justify the cost. The effect on recruitment is then ambiguous: the additional opportunities can outweigh this burden, or congestion can discourage enough inspection to reduce the probability of a choice. A longer shortlist consequently need not lead to more inspections.

2.3. Candidate acceptance

The non-acceptance risk If the condition on candidates' acceptance is relaxed, the firm's interest in a candidate need not lead to a completed match. There is a probability $\alpha(m) \in [0, 1]$ that a candidate of fit m would accept the vacancy's terms. Acceptance is assumed to be weakly increasing in fit, such that compatibility in location or wages can make a match attractive to both parties, although employer-valued characteristics need not always increase a worker's willingness to accept.

In this extension, inspection reveals both how suitable a candidate is and whether they would accept the job. A candidate who is willing to accept remains available while the firm evaluates other options. An inspection therefore improves the firm's position only if the candidate is both better than its current best option and willing to accept. The firm continues inspecting when

$$\alpha(m)\,\mathbb{E}\big[\max\{m + \varepsilon - b_j, 0\}\big] > \kappa_j(n), \tag{6}$$

As before, the inspection frontier is the minimum expected fit that makes evaluation

worthwhile. With acceptance risk, the expected gain at this frontier, denoted by $\underline{m}_j^h$, satisfies

$$\alpha(\underline{m}_j^h)\,\mathbb{E}\Big[\max\{\underline{m}_j^h + \varepsilon - b_j, 0\}\Big] = \kappa_j(n). \tag{7}$$

The assumption that acceptance is weakly increasing in fit preserves a single threshold: more promising candidates offer both a larger potential gain and at least as high a probability of accepting.

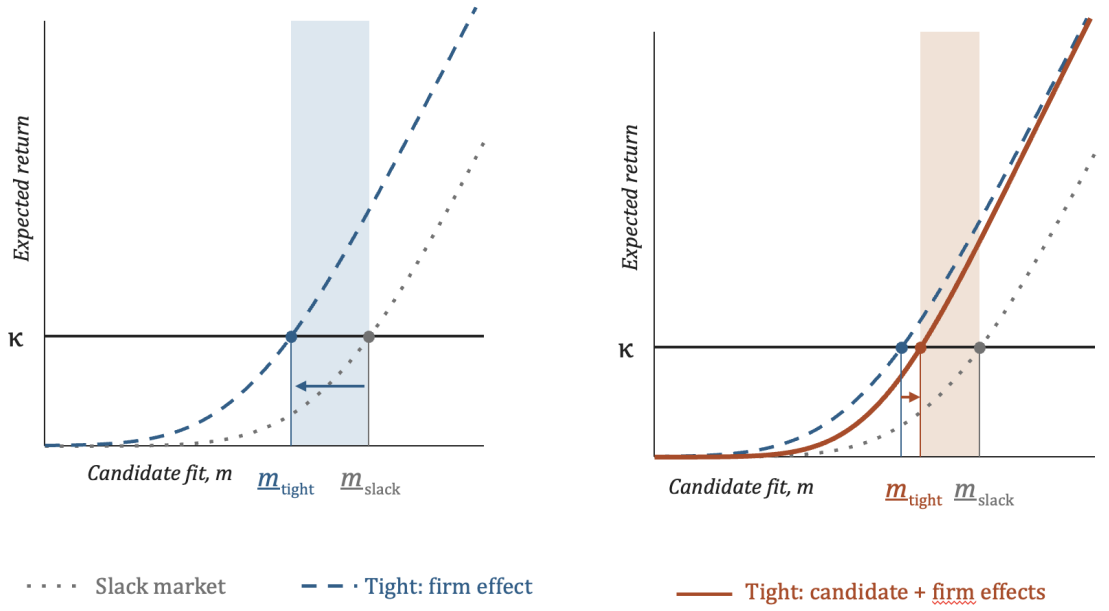
A lower probability of acceptance reduces the expected gain from inspection and raises the frontier, holding costs and the firm's best available option fixed. Further candidates can remain worth evaluating when an earlier candidate would have been suitable but would not accept. A larger shortlist provides alternatives to both unsuitable and unwilling candidates.

The role of labor-market tightness. The value of additional shortlisted candidates also depends on the alternatives available to firms and job seekers. I denote labor-market tightness by ϑ, the number of vacancies relative to job seekers. Figure 1 illustrates two opposing effects, holding the shortlist, vacancy terms and inspection costs fixed. Panel (a) considers the firm's alternatives: when candidates are scarce, recruiting outside the preselection service becomes harder. The firm is therefore more willing to evaluate candidates transmitted by the caseworker, and its inspection frontier shifts to the left. The shaded region contains candidates who become worth inspecting because the firm's alternatives have worsened. Panel (b) adds the candidate's decision. In a tighter market, job seekers may have more competing opportunities and be less likely to accept the vacancy's terms. This reduces the expected return to inspection and shifts the frontier back to the right. In the illustrated case, lower acceptance partly offsets the firm-side effect; more generally, either force can dominate. Candidate scarcity can therefore make access to additional applicants more valuable while making each potential match harder to complete.

2.4. Empirical implications

The framework guides the empirical analysis along three dimensions. First, I examine how caseworker leniency changes the set of candidates available to the firm. A lower transmission cutoff can reduce average shortlist fit while increasing the number of suitable candidates. The analysis therefore considers both the composition of the shortlist and the number of candidates above a fixed fit threshold. Measures of how caseworkers sort applicants on observable characteristics provide complementary evidence on whether leniency also captures differences in screening precision.

Second, I examine whether access to additional candidates changes the probability of a contract proposal or a realized hire. These outcomes capture the firm's willingness to recruit and whether a match is completed. The framework allows additional options to



(A) Firm-side effect

(B) Firm- and candidate-side effects

FIGURE 1. Labor-market tightness and the inspection frontier

Notes: The horizontal axis measures expected candidate fit, m; the vertical axis shows the expected gain from an initial inspection. The horizontal line represents the fixed cost per inspection. Its intersection with expected gain defines the inspection frontier. In panel (a), a weaker firm outside option shifts the frontier left; shaded candidates become worth evaluating. In panel (b), lower candidate acceptance shifts the frontier back to the right.

improve recruitment, but also allows a larger shortlist to discourage inspection if it raises evaluation costs. Comparisons across predicted shortlist sizes and labor-market conditions explore where additional candidates are more valuable.

Third, I examine whom firms select and how long recruitment takes. Selected candidates' characteristics and their position within the shortlist help describe how firms use the options available to them. Proposal dates provide a separate perspective on the recruitment process: additional candidates may require further evaluation or provide alternatives after an unsuccessful proposal. Comparisons of observed hiring among proposed candidates also inform the completion of the match.

Together, these outcomes provide evidence on the mechanisms described by the framework without directly observing in the data inspection decisions or their costs.

3. Institutional setting

This section describes the institutional environment that underlies the empirical design. Three main features are important for the analysis. First, the French Public Employment Service operates a job board that intermediates a large share of national hiring, so that

variation in how it transmits candidates is economically meaningful. Second, its assisted preselection service turns a caseworker into a gatekeeper whose screening style determines how many candidates a firm sees. Third, the assignment of caseworkers to firms is plausibly as good as random within local agencies, which is what permits a judge-design strategy. Sections 3.1–3.2 describe the agency and its job board, Section 3.3 describes assisted preselection, and Section 3.4 details the assignment mechanism.

3.1. The French Public Employment Service's organization

During the sample period (2017–2023), the French Public Employment Service (PES) operated under the name *Pôle emploi*.[2] Created in 2008 through the merger of the public placement agency ANPE and the unemployment insurance agency Assédic, Pôle emploi has a dual mandate: (i) supporting registered job seekers in their search process and (ii) helping firms fill vacancies.

The organization is structured around two main operational branches:

- **Demandeur d'Emploi (Job Seeker Department):** handles job seeker counseling, monitoring, and active labor market programs (ALMPs). Registered job seekers are assigned randomly to a caseworker who provides search assistance and monitors job search activity (Dromundo Mokrani and Haramboure, 2022). Caseworkers typically manage large portfolios (between roughly 70 and 350 job seekers simultaneously).
- **Entreprise (Employer Department):** dedicated to employer relations and recruitment services. This department supports firms in posting vacancies, improving vacancy visibility, and assisting in candidate selection.

The two branches are often working together in agencies and are porous: caseworkers switch frequently between the two of them, according to unobserved criteria. Usually, caseworkers first handle exclusively job seekers before switching to the employer department. As specified in Appendix Table A.1, when the switch occurs, caseworkers still handle some job seekers from their previous portfolio and are assigned new job seekers according to agency congestion, although a smaller number than caseworkers handling exclusively job seekers. Very few vacancies are handled by caseworkers that have their applicants in their portfolio: only 6000 vacancies are encountered where the vacancy and the applicant are followed by the same caseworker. This amounts to 0.7% of cases: excluding them from the sample does not affect final results. Caseworkers can switch between the two departments conditional on agency needs and personal preferences.

[2]Pôle emploi was renamed *France Travail* in January 2024. Since the data for the analysis end in 2023, the former name is used throughout the paper.

3.2. Pôle emploi job board

The French PES has an online job board (pole-emploi.fr) through which job seekers, as well as any person signing up for free, can apply for vacancies. On average, 3 million vacancies are posted yearly on this job board. The vacancies reach the website through two different means: either the firm sends them directly to the PES or the vacancies are collected from other French job boards (for instance "Cadremploi", "emploi collectivités"). Previous research has estimated that the PES job board concentrates offers representing 60 to 75% of all French hirings (Le Barbanchon et al., 2016, 2021).

The volume of applications firms must process has risen markedly. Among vacancies posted on the board, the average number of applications received per vacancy increased from about ten in 2017 to about seventeen in 2025, consistent with the secular decline in application costs for job seekers and the corresponding growth in applications per vacancy (Goldfarb and Tucker, 2019; Birinci et al., 2025). A larger pool offers more chances of a good match but also raises the burden of evaluation, which is precisely the trade-off the paper studies.

Matching between candidates and job offers can occur passively like in the case of any job board, as firms post vacancies and job seekers apply directly. However, there are also two active mechanisms put in place to improve vacancy-job seeker matching. On the one hand, caseworkers can send vacancies directly to their managed job seekers to push them to apply. On the other, firms may also request active intermediation services from the Employer Department for the job vacancies that they transmit directly to the Pôle emploi website.

3.3. Preselection as an employer-directed service

The offer of intermediation services became available in 2015: Pôle emploi started targeting services towards firms in order to help the placement of jobseekers and the filling of vacancies (Algan et al., 2020). The services targeted towards firms are either aiming at bringing more candidates to firms by proposing job seeker profiles even outside of existing vacancies (Behaghel et al., 2024) or trying to reduce hiring costs by assisting with vacancy creation and screening applications.

One of the main employer-directed services aimed at reducing hiring costs is *assisted preselection* ("présélection"). When a firm opts for preselection for a given vacancy, Pôle emploi delegates part of the screening process to a caseworker. The process unfolds as follows:

1. The firm posts a vacancy on the Pôle emploi platform.
2. Applicants apply through the platform.
3. If the firm has opted for preselection, a caseworker screens the pool of applications.

4. The caseworker transmits a shortlist of selected candidates to the firm. Only shortlisted candidates' CVs reach the firm if they applied through the Pôle emploi system.

The preselection process is completely free and is proposed for each vacancy: the firm has to tick a box to have caseworkers screen the applications. Entering preselection once does not mean that the firm will automatically be part of it for future vacancies: the choice is made at every vacancy posting, the default option being no preselection.

In the pilot described by Algan et al. (2020), caseworkers have a meeting with the firm when the latter enters an intensive support program. In this meeting, the firm specifies its main recruiting criteria and the shortlist size it is looking for. This is a specific program provided by Pôle emploi through which it offers to firms, sometimes for a fee, a mix of its services including preselection, vacancy writing and a test of applicants' skills. However, outside this specific type of program, firms that enter preselection do not meet face-to-face with the caseworker who screens the applications and criteria for recruitment are sometimes transmitted by email, although not systematically. The preferred shortlist size can be similarly communicated but is in practice rarely filled in by firms. Discussions with caseworkers handling preselection have shown that often they only base themselves on the vacancy and their own understanding of how to best screen the applications.

Importantly, the firm retains the final hiring decision. The role of the caseworker is therefore to reduce hiring costs by filtering the set of candidates considered. Candidates who are preselected are the only ones shown to the firm, which then has full freedom to interview the preselected applicants or not.

3.4. Assignment of caseworkers to firms

A key institutional feature underlying our identification strategy concerns how firms are matched to caseworkers. When a firm opts for preselection for the first time on a vacancy, the vacancy is attributed to a caseworker within the geographically relevant local agency. Assignment operates at the agency level. In practice, each agency has a limited set of caseworkers authorized to handle employer-side preselection (typically three to four at a given point in time), as indicated in the internal Pôle emploi back-office system. Upon the firm's first use of preselection, the internal system attributes the vacancy to one of the eligible caseworkers within the agency who is open to receiving new vacancies in her portfolio. While the allocation is system-based, agency management retains the ability to reassign cases manually. According to qualitative information collected through discussions with caseworkers, such manual reallocations are uncommon and tend to arise in specific operational circumstances (for instance, anticipated departure of a caseworker or extended unavailability). No systematic reassignment patterns are observed in the administrative data.

Once assigned, the same caseworker typically handles subsequent preselected vacancies posted by that firm. Changes in caseworker assignment appear to occur primarily due

to caseworker-side constraints (for instance in the cases of job mobility, extended leave, or internal reorganization), rather than firm requests. Firms do not have a formal mechanism to request a different caseworker. The identity of the individual caseworker in charge of screening is not salient to firms. In the recruiter interface ("Espace recruteur"), firms observe applications transmitted in the name of the agency rather than the individual counselor. Identifying the specific caseworker requires navigating further on the interface and the information published as a point of contact does not correspond to the actual caseworker in charge of screening. Communication from firms is generally directed to the agency, which then internally dispatches messages to the appropriate caseworker.

Taken together, these institutional features imply that the initial attribution of a caseworker to a firm is plausibly as-good-as-random from the firm's perspective. This generates persistent firm–caseworker relationships and allows us to exploit variation in individual caseworker screening styles in a judge-design framework.

4. Data

4.1. Sample construction

Data sources The analysis relies mainly on administrative data from Pôle emploi, the French Public Employment Service, covering the universe of vacancies and applications posted on its online job board since 2016. The dataset includes detailed information on:

- Vacancy characteristics: occupation, contract type, firm size, wage information (when available), and posting dates;
- Applications: applicant identifiers, application dates, and whether the application was shortlisted under preselection;
- Contract offer outcomes: whether the vacancy led to a contract offer and the timing of the proposal;
- Caseworker identifiers and agency affiliation for the vacancies (available from the second half of 2016 onward).
- Applicant characteristics: gender, age, professional experience, unemployment spell length and reservation wage.

The presence of caseworker identifiers allows the linkage of each preselected vacancy to the counselor responsible for screening. The applications dataset comprises a detailed and dated overview of the evolution of the applications that go through the Pôle emploi job board: one is able to track when the application is submitted, screened, sent to the recruiter and when it is ultimately accepted or declined by the firm. This allows me to

track an intermediary outcome between the hiring match and the application process: the existence of candidates that the firm makes an offer to, before the contract is accepted. Having this information allows a rare insight into the firm's final revealed preferences, even if the interviewing decisions are not documented, as they take place only on the firm side. The main limitation of this dataset is that it only comprises applications that are sent through the PES website or by contacting the Pôle emploi agency. There is no possible tracking of the candidates that apply through other means, even if they are job seekers. However, as the focus of the study is on job vacancies that are posted by the recruiter directly on the PES job board and for which preselection is opted for, it is likely that most of the applications will be concentrated in this dataset.

In order to establish which intentions to hire turned into actual contracts, I am matching the application data with social security records, which allows me to determine which of the selected candidates appear in the workforce reported by the employer to the social security authorities. I can thus see the date when the hiring took place. Furthermore, I can observe all hirings that the firms are conducting. By matching with the individual identifiers from the FPES, I can track which hirings took place after applying through the Pôle emploi job board and which are unrelated (for the latter, there is no information on the channel through which they have reached the firm).

Cohort construction The study focuses on establishments that use preselection for the first time between 2017 and 2023. This restriction serves two purposes. First, it allows me to study firms at the moment they enter assisted screening and to track the evolution of their hiring outcomes through time. Second, because caseworkers are assigned at the first use of preselection and typically remain attached to the firm thereafter, this design generates persistent firm–caseworker relationships that are central to our identification strategy. The sample begins in 2017 to ensure that, for the earliest entry cohorts, at least one quarter of prior caseworker activity is observed, allowing the construction of measures of screening style. The sample ends in December 2023, which allows the observation of outcomes for the latest entry cohorts into 2024, as some recruitment processes remain open for several months and may take up to a year to result in a hire. For firms belonging to these entry cohorts, the study observes the full universe of their vacancies and applications posted on the Pôle emploi platform, including both vacancies using preselection and those without assisted screening.

4.2. Outcome Variables

The variables used in the analysis are defined as follows.

Firm choice. The primary outcome is an indicator for whether the firm makes a contract proposal to a shortlisted candidate within three months of posting, the horizon by which

most hiring decisions are made. To account for the robustness of the decision, the same indicator is measured at 14, 30, and 45 days. An intensive-margin measure records the number of candidates to whom the firm proposes a contract, conditional on the firm making any proposal.

Shortlist composition. Several characteristics of the transmitted shortlist are measured at the vacancy level: the demographic composition of shortlisted candidates; broad quality measures such as education and reservation wage; a fit index summarizing how well shortlisted candidates match the vacancy; and the number of shortlisted candidates above a basic quality threshold. The fit index aggregates several dimensions of vacancy–candidate match and includes: an indicator for the correspondence between past experience and vacancy occupation, an indicator of satisfaction of the posted educational requirement, and the distance between reservation wage and posted salary, when the data are available. The index is constructed using the inverse-covariance weighting of Anderson (2008), so that correlated components are not double-counted and the risk of multiple hypothesis testing bias is mitigated. Distinguishing the average fit of the shortlist from the *number* of usable candidates it contains is essential, as the model predicts these move in opposite directions under threshold-driven expansion.

Selected-candidate characteristics. For vacancies that result in a contract proposal, the characteristics of the chosen candidate are measured along the same dimensions as the shortlist: demographics, labor-market quality (reservation wage, education, experience, unemployment duration), fit to the vacancy, and the candidate's relative rank within the transmitted pool. Time to the contract proposal is also recorded. These variables support the test of whether a larger, noisier shortlist degrades the quality or delays the timing of the firm's selection.

Firm engagement. Subsequent firm behavior is measured by whether the establishment posts a new vacancy within three or six months and whether those vacancies again use preselection, capturing how firms update their reliance on the public employment service following a recruitment episode.

Hired-candidate characteristics. For vacancies that have a recorded contract proposal, I merge with the social security records. I track which proposals turned into a contract, the timing of the hire, the characteristics of the actually recruited workers and the duration of the contract. However, the estimates on the stability of the match are very imprecise and will not be the subject of an extended analysis here.

4.3. Descriptive statistics

The administrative data are large-scale. On average, 7,000 caseworkers are active per quarter during the sample period, as some caseworkers rotate or have short tenures, while others follow several vacancies. On average, caseworkers handle 20 vacancies per quarter, which amounts to 300 applications screened. Overall, the dataset contains 15,016 unique caseworkers and covers more than two million vacancies, of which 800,000 involve preselection. The average vacancy receives around 9 applications when posted without preselection and 20 applications when preselection is used. This difference suggests that firms selecting into preselection tend to face thicker applicant pools, consistent with the idea that assisted screening may be particularly valuable when application volumes are high. Between 2017 and 2023, approximately 400,000 firms posted a vacancy with preselection for the first time. Of these, 93% are retained in the estimation sample as they belong to agency–cohort cells with sufficient within-agency variation in caseworker assignment. The remaining firms are excluded because they are attached to agencies with too few active preselecting caseworkers to allow meaningful within-agency comparisons. The data include 67,353 office × entry-quarter cells, providing rich variation in caseworker assignment and screening intensity. Table 1 details the distribution of the sample composition across firm size and vacancy characteristics. While the share of large firms in the cohorts is rather small (under 1%), this replicates the firm size distribution in the French economy (Insee, 2024). 70% of firms that adopt preselection are posting for the first time on the Pôle emploi job board and 45% are hiring for the first time in three years, in social security records. This pattern is consistent with many firms opting into assisted screening immediately upon entering the platform, or joining the platform for this specific service. One possible interpretation is that these firms seek support from the outset, either because they have limited prior experience recruiting or because they turn to the public employment service specifically to obtain assistance with managing applications.

Table 2 describes how observable applicant characteristics evolve along the recruitment funnel. Preselection substantially reduces the number of dossiers that reach the firm: the average vacancy receives 20.15 applications, of which 9.21 are transmitted to the recruiter. On average, shortlisted candidates are very similar to the full applicant pool in terms of age, gender, education, and reservation wages. They are slightly more experienced (5.22 rather than 5.06 years of experience) and have somewhat shorter unemployment spells, but they are marginally less likely to satisfy the posted qualification requirement or to have worked previously in the vacancy's occupation. Thus, at this descriptive stage, caseworker screening appears to reduce the volume of applications substantially while leaving the observable composition of the pool broadly similar.

The final stage of the funnel differs more visibly. Among the 242,093 vacancies for which a contract proposal is observed, selected candidates are younger than both applicants and shortlisted candidates: 13 percent are under age 25, compared with about

TABLE 1. Descriptive Statistics: Firms and Job Offers

	Firms	Offers without preselection	Offers with preselection
Overview			
Number of firms	373,172		
Office × cohort cells	67,353		
First-time posters	0.7		
First-time hirers	0.45		
Have hired more than 3 times	0.27		
Number of job offers	2,075,786	1,266,111	809,675
Average number of job offers per firm	5.6	3.4	2.2
Average number of applications per offer		9	20
Firm size distribution			
1–9 employees (Micro)	0.81	0.78	0.78
10–49 employees (Small)	0.08	0.13	0.17
50–249 employees (Medium)	0.01	0.04	0.03
250+ employees (Large)	0.01	0.01	0.01
Public sector	0.09	0.04	0.01
Contract type distribution			
Fixed-term		0.54	0.40
Permanent		0.46	0.60

Note: This table reports descriptive statistics for firms and job offers that are part of 2017-2023 preselection entry cohorts. "Preselection" refers to postings where firms ask for screening. The characteristics of offers without preselection correspond to the vacancies posted without preselection, although the firm has posted at some time a vacancy with preselection. The first column describes firm-level characteristics, while the second and third columns report statistics at the job-offer level, distinguishing between offers with and without preselection. Firm size and contract type distributions are expressed in shares.

11 percent at the earlier stages. They also have lower average reservation wages, less labor-market experience, and shorter unemployment spells. Selected candidates are less likely to be tertiary educated, to meet the posted qualification requirement, or to have recent experience in the vacancy's occupation. These differences should not be interpreted as causal effects of either preselection or firm choice, since they combine variation across vacancies, occupations, and firms that ultimately do or do not make a proposal. Appendix Table A.2 complements these descriptives by restricting the sample to vacancies that result in a contract proposal through the preselection channel. The table shows that successful candidates are, on average, younger, have shorter unemployment spells, and have less overall and occupation-specific experience than the candidates transmitted to the firm. At the same time, their average quality index is slightly higher. All of these patterns are

purely descriptive.

TABLE 2. Applicant Characteristics Along the Recruitment Funnel

	All applicants	**Shortlisted** applicants	**Selected** applicant
Demographics			
Average age	36.0	36.2	34.0
Share applicants under 25	0.11	0.11	0.13
Share women	0.25	0.25	0.23
Share foreigners	0.07	0.06	0.05
Labor-market quality			
Average reservation wage (EUR/month)	2,214.7	2,187.1	2,061.4
Share tertiary educated (including CAP)	0.13	0.13	0.12
Average experience (years)	5.06	5.22	4.32
Share unemployed over one year	0.23	0.23	0.20
Quality of fit			
Average quality index	0.39	0.39	0.36
Share qualified enough	0.15	0.14	0.12
Average experience within occupation (years)	1.02	0.97	0.66
Share with last position in same occupation as vacancy	0.15	0.14	0.11
Share below posted wage	0.20	0.19	0.18
Number of job offers	809,675	809,675	242,093
Average number of applicants	20.15	9.21	1.51

Notes: The table reports mean applicant characteristics at three stages of the recruitment funnel on the preselection sample: all applicants to a vacancy (column 1), the candidates the caseworker transmits to the firm (column 2), and the candidate the firm ultimately selects (column 3) when a selection takes place. The last row reports the average number of applicants at each stage. All figures are unconditional means and are descriptive rather than causal; the corresponding instrumented estimates of the effect of shortlist size on shortlist composition and on selected-candidate characteristics are reported in Section 7. The quality index follows the inverse-covariance weighting of Anderson (2008). Reading column 1 against column 2 describes how the caseworker reshapes the pool; reading column 2 against column 3 describes the firm's selection.

5. Empirical Strategy

5.1. Empirical object and interpretation of shortlist size

The impact of seeing more candidates from one channel of recruitment is studied. The focus is on vacancies posted on the Pôle emploi job board for which firms opt into assisted preselection. The purpose is not to evaluate the impact of entering preselection versus not entering it, as this decision is highly endogenous and depends on unobservable characteristics of the firm. Instead, the analysis uses the unique institutional features of the

caseworker assignment process to evaluate the trade-off the firm faces between hiring costs and the benefits of having an extra applicant. For each vacancy i, posted by establishment f and handled by caseworker c in office o and period t, the data record the number of candidates transmitted to the recruiter, $X_i \in \mathbb{N}$, and an indicator $Y_i \in \{0, 1\}$ for whether the firm makes a contract proposal to a shortlisted candidate within three months. A vector W_i collects predetermined vacancy and establishment characteristics. Under preselection the recruiter sees only the candidates the caseworker transmits, so X_i measures the candidate pool effectively available to the recruiter through this channel: it is the number of dossiers the firm can review through Pôle emploi. In the notation of the model, X_i is the realized shortlist size n for vacancy i, and Y_i is the firm's selection of a candidate from the shortlist.

The key institutional feature is that, under preselection, the recruiter only sees candidates transmitted by the counselor through the Pôle emploi interface. One can therefore easily interpret the shortlist size X_i as the relevant measure of the Pôle emploi candidate pool effectively available to the recruiter. In particular, X_i proxies the intensity of supply that the public employment service provides to the firm at the screening stage: it is the number of candidate dossiers the recruiter can review through Pôle emploi.

Recruiters may receive candidates through other channels which are not observed in the administrative data. It is impossible to identify for which establishments the public employment service is the dominant recruitment channel. Instead, the estimates should be read as the effect of expanding the *supply of candidates made available to the recruiter through Pôle emploi* for vacancies that are routed through the platform's preselection process. Requiring that the vacancy is posted directly on the Pôle emploi job board and that the firm actively opts into preselection arguably selects vacancies for which the platform is a non-trivial component of the recruitment process. Nonetheless, the magnitude of the effect may vary with the importance of outside recruitment options. The results should be read as effects of altering the supply at the margin for vacancies that rely on the Pôle emploi job board. Since X_i is a count rather than a binary treatment, this object is the average causal response in the sense of Angrist and Imbens (1995), defined and characterized further down.

5.2. Instrumenting strategy: predetermined counselor leniency

Endogeneity concern A naive regression of Y_i on X_i does not recover a causal effect. Caseworkers transmit more candidates when a vacancy appears harder to fill, when the job is more attractive, or when the employer signals stricter criteria; these determinants are imperfectly observed and correlate with the outcome. As a result, ordinary least squares conflates the effect of candidate supply with unobserved vacancy difficulty and employer selectivity, and the bias is plausibly negative, since harder-to-fill vacancies receive larger shortlists.

Instrument definition In order to address the endogeneity concern, this paper exploits the quasi-random variation in counselor screening style. Counselors differ systematically in their propensity to transmit candidates, even when handling similar vacancies in the same local labor market and time period. In order to summarize this propensity, a counselor leniency index is constructed from the caseworker's preselection decisions in the quarters before they first handle a given firm. Let $t_{-1}(f, c)$ denote the last period strictly before caseworker c begins handling preselected vacancies for establishment f. The instrument uses only c's screening behavior observed up to $t_{-1}(f, c)$:

$$\begin{aligned} Z_{fc} &\equiv g\big(\{\text{screening decisions of } c \text{ in periods } \leq t_{-1}(f, c)\}\big), \\ Z_i &= Z_{fc} \quad \text{for all vacancies } i \text{ handled by } (f, c). \end{aligned} \tag{8}$$

Concretely, $g(\cdot)$ is the caseworker's average transmission propensity over the pre-period, residualized on office-by-time cells so that it reflects individual style rather than local conditions, and normalized to mean zero and unit variance.

This predetermination is important for two reasons. First, the current vacancy cannot mechanically influence the instrument used for that vacancy. Second, the instrument is not constructed from contemporaneous outcomes of the same establishment–counselor pair. These properties rule out the own-observation bias that motivates leave-one-out corrections in designs where the instrument is computed from the same realizations of the endogenous variable that enter the first stage (Goldsmith-Pinkham et al., 2025; Frandsen et al., 2023). Restricting attention to caseworkers with prior preselection experience mechanically drops the first firm attributed to a starting caseworker, but this removes fewer than ten percent of observations and yields a more transparent instrument.

Figure 2 summarizes the design. A vacancy's applicant pool is screened by a caseworker assigned quasi-randomly within the local agency. The same pool, handled by a less lenient caseworker, yields a shorter shortlist than when handled by a more lenient one. As the firm chooses only from the transmitted set, the longer shortlist can change which candidate is preferred. The instrument exploits exactly this variation, comparing vacancies whose pools are screened by caseworkers of differing, predetermined leniency.

Econometric specification The two-stage least squares model is

$$X_i = \pi Z_i + \Gamma' W_i + \gamma_{o\times t} + \delta_{\tau(i)} + u_i, \tag{9}$$

$$Y_i = \beta X_i + \Phi' W_i + \gamma_{o\times t} + \delta_{\tau(i)} + \varepsilon_i, \tag{10}$$

where $\gamma_{o\times t}$ are office-by-entry-cohort fixed effects capturing local labor market conditions and agency organization, and $\delta_{\tau(i)}$ are posting-month fixed effects absorbing aggregate seasonality and national shocks. The vector W_i includes predetermined controls, in particular a flexible function of the number of applications received through the platform and timing variables describing when shortlisting occurs. Identification comes from

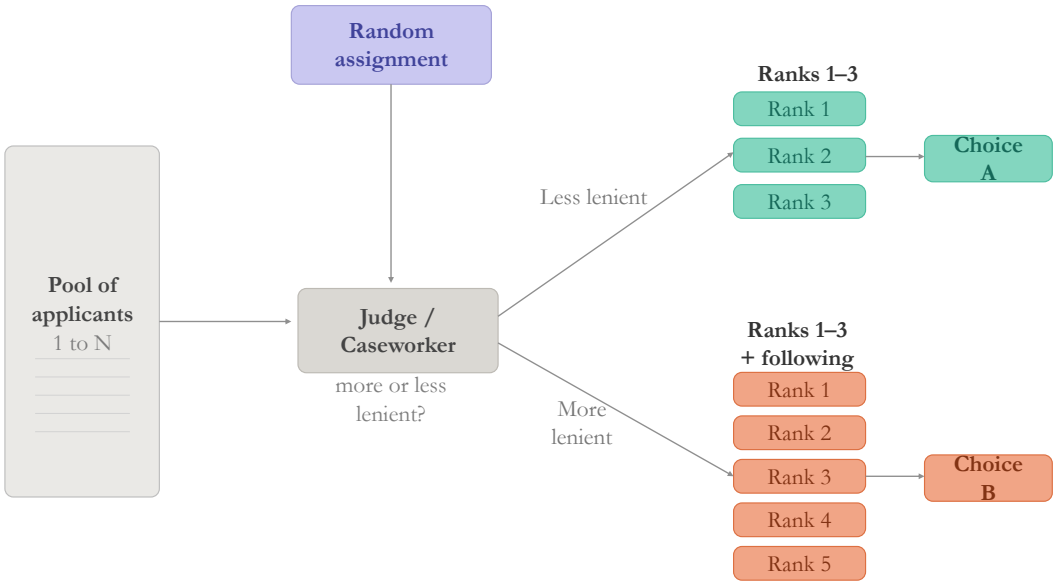


FIGURE 2. The caseworker leniency design

Notes: A vacancy's applicant pool $(1, \ldots, N)$ is screened by a caseworker who is assigned to the firm quasi-randomly within the local agency. A less lenient caseworker transmits a shorter shortlist (ranks 1–3); a more lenient caseworker transmits a longer one (ranks 1–5), making lower-ranked but still credible candidates visible to the firm. Since the firm selects from the transmitted set, the longer shortlist can lead to a different choice (Choice B rather than Choice A). The design isolates the effect of shortlist size by comparing vacancies whose pools are screened by caseworkers of differing, predetermined leniency.

comparing vacancies handled in the same office-by-time cell but assigned to caseworkers with different pre-measured leniency. Standard errors are clustered at the caseworker level, the level at which the instrument varies.[3]

5.3. Identification assumptions

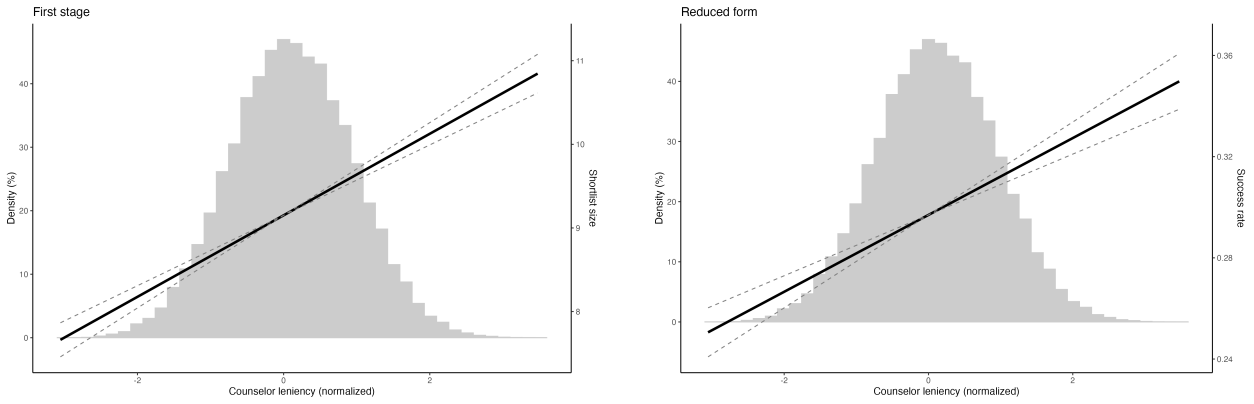
The causal interpretation of β relies on four assumptions.

Assumption A1: Relevance. Caseworker leniency affects the number of candidates transmitted:

$$\pi \neq 0.$$

[3]The instrument Z_i is constructed from the assigned caseworker's past screening decisions and can therefore take different values for the same caseworker when measured in different quarters. Identification relies on differences in predetermined leniency across caseworkers within office-by-time cells. Standard errors are clustered by caseworker to allow for correlation across vacancies handled by the same caseworker, including vacancies observed in different quarters.

This amounts to looking at the first-stage estimates and excluded-instrument F-statistics. Figure 3 (Panel a) provides graphical evidence of this relationship, showing that vacancies assigned to more lenient counselors receive larger shortlists on average. Table 3 reports the corresponding first-stage estimates. The instrument strongly predicts shortlist size in the full sample. The estimated coefficients are close to 0.5, implying that a two-standard-deviation increase in counselor leniency increases shortlist size by one supplementary candidate. The excluded-instrument F-statistic equals 92.4 for the full sample (reported in the first column), well above conventional thresholds for weak instruments. These results indicate that counselor leniency provides a strong source of variation in the number of candidates transmitted to recruiters.



(A) (B)

FIGURE 3. Effect of Counselor Leniency on Shortlist Size (First Stage) and on Vacancy Filling Probability (Reduced Form)

Notes: The figure plots binned averages of outcomes against normalized counselor leniency. Counselor leniency is measured as the pre-determined lagged tendency of counselors to transmit candidates to vacancies and is normalized to have mean zero and unit standard deviation.
Panel (a) shows the first-stage relationship between counselor leniency and shortlist size (number of candidates transmitted to the recruiter). More lenient counselors transmit larger candidate pools.
Panel (b) shows the reduced-form relationship between counselor leniency and contract proposal. Vacancies assigned to more lenient counselors exhibit higher probabilities of the firm extending an offer to a candidate.
Solid lines show linear fits; dashed lines indicate 95% confidence intervals. Gray histograms display the distribution of counselor leniency.

Assumption A2: As-good-as-random assignment within office-by-time cells. Conditional on predetermined characteristics and assignment cells, caseworker leniency is independent of potential shortlist sizes and outcomes:

$$Z_i \perp\!\!\!\perp \left(\{X_i(z)\}_z, \{Y_i(x)\}_x\right) \mid C_i.$$

Here i indexes a vacancy, Z_i is the caseworker's predetermined leniency, and C_i collects the predetermined controls and fixed-effect cell identifiers. The potential shortlist size $X_i(z)$ is the number of candidates transmitted at leniency z, and $Y_i(x)$ is the outcome

TABLE 3. First-Stage Estimates by Subsample

	(1) Whole sample	(2) Only preselection	(3) First offer	(4) First offer × only preselection
Dependent variable: shortlist size				
Caseworker leniency Z_c^{lag}	0.49***	0.51***	0.56***	0.61***
	(0.05)	(0.05)	(0.06)	(0.07)
Observations	809,675	602,564	373,172	283,013
F-statistic (excl. instrument)	92.4	102.2	89.8	86.0
Instrument, 1st percentile	−1.43	−1.40	−1.45	−1.40
Instrument, 99th percentile	1.43	1.40	1.43	1.39
Implied Δ shortlist (p99−p1)	1.39	1.41	1.62	1.71

Notes: Each column reports the first-stage regression of shortlist size on predetermined caseworker leniency, Z_c^{lag}, estimated on a different subsample: the whole sample (1); vacancies receiving only preselection, with no bundled services (2); each firm's first preselected vacancy (3); and the intersection of the two restrictions, first preselected vacancies receiving only preselection (4). All specifications include office-by-cohort and posting-month fixed effects. Standard errors, clustered at the caseworker level, are in parentheses. The F-statistic is the clustered excluded-instrument statistic. The instrument percentiles are computed on the within-fixed-effect residual of Z_c^{lag}, and the implied Δ is $\hat{\pi} \times (p_{99} - p_1)$, expressed in units of shortlist size: a move from the 1st to the 99th percentile of leniency raises the shortlist by between 1.4 and 1.7 candidates. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

with x candidates along the expansion path described below. The observed values satisfy $X_i = X_i(Z_i)$ and $Y_i = Y_i(X_i)$.

This assumption is supported by the institutional process described in Section 3.4: firms do not choose counselors, the counselor identity is not salient in the recruiter interface, and reassignment is rare and driven primarily by counselor-side constraints.

Empirically, I assess this condition using balance tests: predetermined vacancy and establishment characteristics should not systematically vary with Z_i within office-by-time cells. The balance tests are reported in Appendix Tables A.3, A.4, and A.5, where one can find estimates from regressions of the form:

$$F_i = \alpha + \rho Z_i + W_i'\lambda + \gamma_{o\times t} + \delta_{\tau(i)} + u_i,$$

where F_i denotes predetermined establishment and vacancy characteristics measured at the establishment's first observed vacancy posting, or characteristics of the initial applicant pool measured before caseworker screening. Under as-good-as-random assignment, the coefficient ρ should not be significantly different from zero.

The estimates show no economically or statistically meaningful relationship between counselor leniency and predetermined characteristics. Although a couple of coefficients

are statistically significant due to the large sample size, all point estimates correspond to changes smaller than one percentage point for a one-standard-deviation increase in counselor leniency. These magnitudes are small relative to the baseline averages. The point estimate of β is moreover stable to adding and removing the W_i controls. Vacancies seem to be randomly assigned to more or less lenient caseworkers within office-by-time cells. Caseworkers also do not seem to be specialized in some types of occupations or industries.

Assumption A3: Exclusion restriction. Leniency affects hiring outcomes only through the number of candidates transmitted:

$$Y_i(x, z) = Y_i(x) \quad \forall z.$$

The potential outcome $Y_i(x)$ includes the composition changes that accompany shortlist expansion. The restriction therefore allows average fit to fall as more candidates are transmitted. It requires that, at a given shortlist size, leniency has no further effect through candidate composition, transmission timing, or other services. The screening-precision checks provide evidence on this assumption. This exclusion restriction also requires that counselor leniency does not directly improve hiring conditional on shortlist size (for instance through additional services such as vacancy rewriting, targeted outreach to job seekers, or other forms of employer support that could affect match success independently of the number of dossiers transmitted).

Two features of the institutional setting make this restriction plausible. First, under assisted preselection the counselor's core operational task is narrowly defined as screening applications and transmitting candidates to the recruiter. Second, the firm retains full discretion over subsequent steps (interviews and hiring decisions), so that any counselor influence must operate through the set of candidates made visible to the recruiter.

Nonetheless, because the public employment service can in principle bundle preselection with other forms of employer support, the baseline analysis is complemented with an analysis of the vacancies that only benefit from preselection. The dataset allows me to differentiate between vacancies for which the firm receives *only* the preselection service and those that benefit from a bundle of services, through which the influence of the counselor could also be present. The full IV design on the subsample of "only-preselection" vacancies is re-estimated as a robustness check. Under the assumption that bundled services, when present, are the main threat to exclusion, stability of the estimates in the "only-preselection" subsample strengthens the interpretation that counselor leniency affects hiring outcomes primarily through the number of candidates transmitted. Table 3 shows that the first-stage estimates are stable to the only-preselection sub-sample, with an even stronger F-test than in the whole sample, while the IV estimates provided in Appendix Table A.10 are consistent with the ones from the whole sample. These results support the exclusion restriction assumption.

A second threat to the exclusion restriction is specific to judge designs in which the judge is multidimensional (Frandsen et al., 2023). Caseworker leniency might be correlated with another dimension of caseworker behavior that affects hiring on its own, the most natural candidate being the quality of screening. If more lenient caseworkers were systematically better, or worse, at ordering applicants, then the instrument would shift not only the size of the shortlist but also its composition at a given size, and the 2SLS coefficient would mix the value of an additional candidate with the value of a better- or worse-ranked pool. I address this concern directly by measuring screening precision and showing that it is orthogonal to leniency, that it does not account for caseworkers' effects on hiring beyond the quantity margin, and that the IV estimates are unchanged when it is held fixed.

I measure screening precision by asking whether a caseworker tends to transmit the applicants who appear better suited to the vacancy based on observable characteristics. For each caseworker and quarter, I compute a Mann–Whitney area under the curve (AUC). This statistic is the probability that, among two applicants to the same vacancy, one transmitted and one not transmitted, the transmitted applicant has a higher value of the applicant-quality index defined in Section 4.2.[4] This measure captures how accurately a caseworker selects applicants while being less mechanically tied to the number of applicants transmitted than a regression coefficient of transmission on applicant quality (Fawcett, 2006). To ensure that the measure is predetermined, the value assigned to a caseworker in a given quarter is constructed exclusively from her decisions on vacancies posted before that quarter.

Appendix Table A.6 confirms that leniency and precision capture distinct dimensions of caseworker behavior. To avoid mechanical correlation from common estimation error, I estimate caseworker effects on contract proposals, $\hat{\theta}_Y$, using vacancies posted in even quarters and effects on shortlist size, $\hat{\theta}_X$, using vacancies posted in odd quarters. Leniency is unrelated to precision, and precision does not explain the component of $\hat{\theta}_Y$ left after accounting for leniency. Moreover, the quadratic specification shows no evidence of a nonlinear relationship between $\hat{\theta}_Y$ and $\hat{\theta}_X$.

I conduct a robustness check at the vacancy level, by including the precision measure in the regressions as a supplementary control. The precision measure could only be constructed for a subsample of the vacancies, since some caseworkers did not handle enough candidates in the pre-period to produce an accurate precision measure. Appendix Table A.7 reports the estimates of the robustness check. Within the sample for which precision is observed, controlling for it leaves the 2SLS estimate unchanged at 2.2 percentage points. The lower estimate relative to the baseline of 0.031 appears to reflect sample selection:

[4]The statistic is first computed from comparisons between applicants to the same vacancy and then aggregated across the vacancies handled by the caseworker. An AUC of 0.5 indicates that transmitted and non-transmitted applicants cannot be distinguished on the basis of observable quality, whereas an AUC of 1 indicates that every transmitted applicant is ranked above every non-transmitted applicant.

restricting the sample to caseworkers handling at least ten vacancies reduces the estimate. There is also no statistically significant evidence that the effect of shortlist size varies with precision: the interaction is 0.004 (s.e. 0.010). Interestingly, when looking at the impact of an extra candidate, instrumented by leniency, on the sub-sample with precise and less precise caseworkers, a pattern appears consistent with an interpretation in terms of applicant ordering. An additional candidate raises the probability of a contract proposal by 2.89 percentage points in the subsample where caseworkers are below the median in terms of precision, compared with 1.44 percentage points for the subsample of particularly precise caseworkers. This pattern is consistent with more precise caseworkers identifying suitable applicants earlier in the shortlist. At a given shortlist size, firms may therefore already have access to a suitable candidate, reducing the value of further expansion. With less precise screening, additional candidates may instead compensate for suitable applicants having been missed earlier. Although the confidence intervals overlap, the direction of the point estimates is consistent with this mechanism.

Overall, the evidence supports interpreting leniency as shifting the number of candidates presented to the firm rather than screening quality. Consistently, Section 7.1 shows that more lenient caseworkers transmit shortlists with lower average and minimum applicant fit, as expected when the transmission threshold is relaxed.

Table 4. Testing for Monotonicity with the Baseline Instrument

Dependent variable: Shortlist Size				
	q1	q2	q3	q4
Counselor leniency	0.35***	0.40***	0.52***	0.80***
	(0.06)	(0.08)	(0.10)	(0.09)
Observations	210,027	195,187	202,063	202,398
Dependent var. mean	7.03	8.45	9.83	11.56
Dependent var. SD	9.01	9.65	10.82	12.31
F-statistic (instruments)	35.78	26.94	28.75	75.71

Notes: The sample is partitioned into quartiles based on predicted assignment to treatment (obtained from an OLS regression of treatment on predetermined covariates). Each column reports coefficients from a quartile-specific first-stage regression of treatment on the instrument(s), including fixed effects. Standard errors are clustered at the counselor level. Significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Assumption A4: Monotonicity. Monotonicity requires that more lenient counselors weakly increase the number of transmitted candidates:

$$X_i(z') \geq X_i(z) \quad \text{for } z' > z.$$

This assumption rules out systematic "defiers," that is, vacancies for which assignment to a more lenient counselor would reduce shortlist size. In the present setting, monotonicity is natural if counselor leniency reflects a stable propensity to transmit more candidates rather than strategic or vacancy-specific behavior.

The paper provides empirical support for this assumption in Table 4. Following standard approaches (Humlum et al., 2025; Bhuller et al., 2020), the sample is partitioned into quartiles based on predicted assignment to treatment, obtained from an OLS regression of shortlist size on predetermined characteristics. Within each quartile, the first-stage relationship between counselor leniency and shortlist size is re-estimated.

The first-stage coefficient is positive and statistically significant in all quartiles. Moreover, the magnitude of the effect is increasing across quartiles (from 0.35 in the lowest quartile to 0.80 in the highest quartile), indicating that more lenient counselors consistently transmit more candidates across the entire distribution of predicted shortlist sizes. There is no evidence of a subgroup for which greater leniency reduces transmission.

These results are consistent with the weak monotonicity assumption underlying the IV interpretation (Goldsmith-Pinkham et al., 2025) and suggest that variation in counselor leniency shifts shortlist size in the same direction for all groups. Appendix Table A.8 implements a complementary *reverse-subsample* check, following the logic of Bhuller et al. (2020) and Humlum et al. (2025). The baseline instrument measures a counselor's predetermined transmission propensity using her prior screening decisions. In contrast, for each predicted-shortlist-size quartile q, the reverse instrument is constructed from the counselor's lagged screening decisions for vacancies belonging to the *other three* quartiles, excluding vacancies in quartile q. The resulting measure therefore asks whether a counselor who is relatively more likely to transmit candidates for other types of vacancies is also more likely to do so for vacancies in the target quartile.

This exercise tests a necessary implication of a stable one-dimensional leniency ranking across vacancy types: counselors who are more lenient outside a given subgroup should also transmit more candidates within that subgroup. The estimates are positive and statistically significant in all four quartiles, ranging from 0.62 in the lowest quartile to 1.36 in the highest quartile. Thus, the direction of the first stage is preserved when counselor leniency is measured using screening decisions from other parts of the predicted-shortlist distribution. Together with Table 4, these results provide suggestive evidence against systematic reversals in counselor behavior across vacancy types. Nonetheless, monotonicity remains an identifying assumption that cannot be directly verified (Goldsmith-Pinkham et al., 2025; Frandsen et al., 2023).

Proposition 1 (Average causal response with a count treatment)**.** *Suppose Assumptions A1–A4 hold in the estimation sample, together with consistency and finite required moments. For each observed variable* $V_i \in \{Z_i, X_i, Y_i\}$*, we write* $\widetilde{V}_i = V_i - \mathbb{E}[V_i \mid C_i]$*. If* $\mathbb{E}[\widetilde{Z}_i \widetilde{X}_i] > 0$*, the*

conditional IV estimand can be written as

$$\beta_C^{\mathrm{IV}} = \frac{\mathbb{E}[\widetilde{Z}_i \widetilde{Y}_i]}{\mathbb{E}[\widetilde{Z}_i \widetilde{X}_i]} = \sum_{x \geq 1} \omega_x \mu_x, \qquad \omega_x \geq 0, \quad \sum_{x \geq 1} \omega_x = 1. \tag{11}$$

Here x denotes the increase from $x-1$ to x candidates. The term μ_x is the IV-weighted average causal increment at that margin among vacancies induced to cross it by leniency, and ω_x is the total weight placed on that margin.

This follows from the average-causal-response interpretation of Angrist and Imbens (1995). The coefficient averages the effects of an additional candidate over the shortlist margins changed by the instrument. It includes the accompanying changes in composition and does not generally equal the effect of adding one candidate to an arbitrary vacancy.

The proposition uses conditional-mean adjustment. It also applies to the reported linear 2SLS specification if the included controls and fixed effects correctly represent $\mathbb{E}[Z_i \mid C_i]$. Saturated adjustment is sufficient, but additive fixed effects alone do not guarantee this condition. The connection to the regression is discussed in Appendix B.6.

6. Increasing the candidate pool decreases the need for an outside option

As the instrument seems robust, I look at the impact of shortlist size on the firm's recruitment decisions. Section 6.1 presents the average effects on contract proposals and realized hiring. Section 6.2 documents the diminishing returns to an additional candidate for both outcomes. Section 6.3 turns to subsequent firm behavior.

6.1. Average effects on the firm's choice and hiring

Table 5 reports ordinary least squares and two-stage least squares estimates of the effect of shortlist size on contract proposals and realized hiring. The first and third columns measure whether the firm proposes a contract within three months and whether a hire is observed within six months. The second and fourth columns report the number of candidates receiving a proposal and the number hired, conditional on observing at least one proposal or at least one hire, respectively.

Contract proposals The OLS estimates suggest that larger shortlists are weakly associated with a higher probability of a contract proposal. In contrast, the IV estimates imply substantial positive effects. A one-candidate increase in the shortlist raises the probability that the firm proposes a contract within three months by 3.1 percentage points, corresponding to about a 10 percent increase relative to the baseline mean of 0.299.

The effect is also present at shorter horizons. Appendix Table A.9 reports positive IV estimates at fourteen, thirty, and forty-five days, reversing the negative short-run

associations observed in OLS. This pattern is consistent with harder-to-fill vacancies receiving larger shortlists and taking longer to reach a decision. The results show that expansion raises the probability of a proposal by each horizon.

Among vacancies with at least one contract proposal, an additional transmitted candidate is associated with 0.14 more proposals in the IV specification, about nine percent of the baseline of 1.51. A longer shortlist therefore also corresponds to more candidates being chosen among vacancies where a choice is observed. As firms transmit in a bundle to the FPES the contract proposals they made, this variable does not allow me to distinguish the mechanisms at play, if, for instance, after a first proposal is rejected by the candidate, the firm has moved on to another potential hire.

Appendix Table A.11 reports the difference in effects if the sample is limited to the first encounter. I find stronger results for the first time when the firm uses the preselection program, but, conditional on re-using the program, the effects remain of a similar magnitude. Appendix Table A.10 restricts the sample to vacancies receiving only preselection and yields estimates that are very similar to the main results. This reduces the concern that the result is driven by other bundled recruitment services.

Table 5. Effect of Shortlist Size on Vacancy Outcomes

	Contract proposal		*Hiring*	
	Any proposal	**No. of choices**	**Any hiring**	**No. of hirings**
Panel A: 2SLS				
Shortlist size	0.0312***	0.1406**	0.008***	0.082**
	(0.0033)	(0.0460)	(0.002)	(0.027)
Panel B: OLS				
Shortlist size	0.0006***	0.1078***	0.000***	0.064***
	(0.00006)	(0.0098)	(0.000)	(0.006)
Baseline mean	0.299	1.507	0.123	1.383
Observations	809,675	242,093	809,675	99,244
Office × Cohort FE		Yes		
Posting Month FE		Yes		

Notes: Proposal outcomes are measured within three months of posting and hiring outcomes within six months. Columns (1) and (3) use all vacancies. Column (2) is restricted to vacancies with at least one contract proposal, and column (4) to vacancies with at least one realized hire. Baseline means and observation counts refer to the corresponding estimation sample. Panel A instruments shortlist size with predetermined caseworker leniency; Panel B reports OLS estimates. All specifications include office-by-cohort and posting-month fixed effects. Standard errors are clustered at the caseworker level. Stars are based on Holm-adjusted p-values computed across horizons within each model. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Hiring decisions The third column of Table 5 shows that the proposal gains extend to realized recruitment. One additional transmitted candidate raises the probability of a hire within six months by 0.8 percentage points, relative to a mean of 12.3 percent, or approximately 6.5 percent. Appendix Table A.13 reports positive and statistically significant effects at horizons ranging from fifteen days to six months, indicating that the result is not driven only by the longer matching window.

Among vacancies with at least one realized hire, the fourth column reports an increase of 0.082 in the number hired, about six percent of the conditional mean of 1.383. As with the proposal count, this estimate conditions on an outcome affected by shortlist size.

The hiring response is smaller than the proposal response. At the common horizon of three months, one additional candidate raises the probability of a proposal by 3.1 percentage points and the probability of a hire by 0.7 percentage points. This is consistent with the distinction in the model between a candidate whom the firm wishes to recruit and a completed match.

Elasticities. To express the effects in comparable units, I report the IV responses normalized by the mean shortlist size and outcome in the corresponding estimation sample:

$$\widehat{\epsilon}_{Y,n} = \widehat{\beta}_Y \frac{\bar{n}}{\bar{Y}}. \tag{12}$$

Here, Y denotes the recruitment outcome, n is shortlist size, $\widehat{\beta}_Y$ is the IV effect of one additional transmitted candidate, and $\bar{n}$ and $\bar{Y}$ are their sample means. Table A.12 reports an elasticity of 0.96 for the probability of a contract proposal and 0.6 for the probability of a hire within six months. Locally, a ten percent expansion in shortlist size therefore corresponds to increases of approximately 9.6 and 6.4 percent in these probabilities. The elasticity of the unconditional number of distinct hires is 1.13.

6.2. Decreasing returns to an additional candidate

Decreasing returns The conceptual framework suggests that additional candidates can be especially valuable when the firm has few suitable options. Figure 4 brings together the effects of a 10 percent increase in shortlist size on contract proposals and realized hiring across five quintiles of predicted shortlist size. Expressing expansion in proportional terms allows the comparison to account for differences in the scale of the choice set: one additional candidate represents a larger proportional increase for a short shortlist than for a long one. Both profiles decline sharply: the value of an expansion in the number of candidates is concentrated among vacancies with the smallest pools.

In order to recover the estimates by shortlist size, one cannot rely on the realized outcomes that are endogenous to caseworker leniency. Therefore, I construct a predicted shortlist size only from predetermined characteristics. Specifically, predicted shortlist size

is the fitted value from an ordinary least squares regression of shortlist size on covariates that are fixed before the caseworker acts: firm size, occupation tightness and industry, together with the office-by-cohort and posting-month fixed effects, while excluding the instrument or any post-screening quantity. This should ensure that the bins of predicted shortlist size are independent of caseworker leniency: a vacancy's bin reflects the kind of position being filled, not the screening style of the caseworker who happened to be assigned to it. The quintiles therefore distinguish vacancies by the size of the pool they are expected to receive, rather than by the screening decision of the assigned caseworker. Within each quintile, shortlist size is instrumented with predetermined caseworker leniency. Under the conditions of Proposition 1, the estimates recover local average responses for these groups of vacancies, rather than the effect at each exact shortlist size.

A 10 percent increase in shortlist size raises the probability of a contract proposal by approximately 9.3 percentage points in the bottom quintile, compared with 2.2 percentage points in the second quintile and 1.0 percentage point in the top quintile. The corresponding hiring effect is approximately 3.2 percentage points in the bottom quintile and between 0.3 and 0.4 percentage points in the remaining quintiles. The point estimates remain positive throughout. The clearest pattern is therefore the concentration of gains among vacancies with the smallest predicted pools; the proposal response also declines across the remaining quintiles, while the hiring response is relatively similar across them.

Appendix Figure A.1 reports the effects of one additional transmitted candidate. The concentration of gains among vacancies with small predicted shortlists is visible under both measures of expansion.

Where are gains concentrated? Table 6 describes the vacancies in the groups that benefit most from additional candidates. Vacancies in the lowest predicted shortlist-size quintile more often offer short-term contracts than those in the highest quintile (47.6 versus 36.4 percent) and are more often open to beginners (53.3 versus 42.2 percent). Construction and public works account for 12.7 percent of vacancies in the bottom quintile, compared with 3.2 percent in the top quintile. Sales and retail and business support are instead more represented among vacancies with larger predicted pools. The industry composition follows a similar pattern, with construction more common in the bottom quintile and retail more common in the top quintile. Differences in establishment size and recent hiring activity are smaller: micro establishments dominate all five groups, and the share with no observed hire in the preceding year varies little.[5]

Appendix Figures A.6 and A.7 complement this description with separate IV estimates by occupational field and industry. The contract-proposal point estimate is large for construction and public works occupations, consistent with their greater representation

[5]Appendix Tables A.27, A.28, A.29 and A.30 report complementary heterogeneity by firm size, experience requirements, posted salary and prior recruitment.

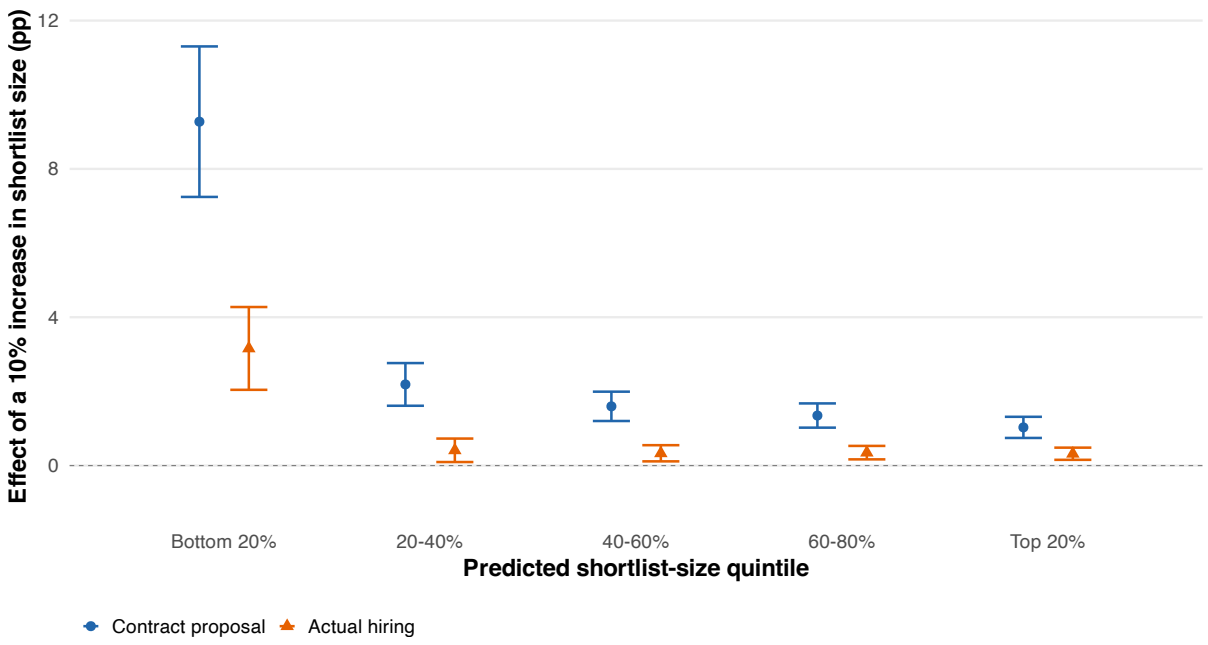


FIGURE 4. Effects of a 10 percent increase in shortlist size on proposals and hiring

Notes: The figure reports IV estimates of the effect of a 10 percent increase in shortlist size on the probability of a contract proposal within three months (blue circles) and a realized hire within six months (orange triangles), by quintile of predicted shortlist size. Effects are expressed in percentage points. Predetermined caseworker leniency provides the instrument for shortlist expansion. Vertical bars show 95% confidence intervals. The quintiles group vacancies using predetermined characteristics and do not correspond to successive expansions of the same vacancy's shortlist.

among vacancies with small predicted pools, but the confidence interval is wide. Positive proposal and hiring responses are more precisely estimated for sales and retail and for hospitality, tourism and leisure occupations. The industry estimates also show positive responses for retail and accommodation and food services. Gains from additional candidates therefore extend to activities that are not especially concentrated in the lowest predicted shortlist-size quintile. These estimates capture returns within broad groups that themselves contain vacancies with different candidate pools, and their uncertainty does not establish a clear ranking across occupations or industries. Together, the results suggest that the size of the expected choice set remains a useful guide to where additional candidates are most valuable, alongside the characteristics of the vacancy.

6.3. Average effects on vacancy publishing probability

Table 7 examines whether shortlist size affects subsequent vacancy posting behavior. The first two columns report whether the establishment publishes a new vacancy within 3 or 6 months, while the last two columns focus on whether the new vacancy uses accompanied preselection.

TABLE 6. Vacancy and applicant characteristics by predicted shortlist-size quintile

	Q1	Q2	Q3	Q4	Q5
Vacancy and employer characteristics (%)					
Short-term contract	47.6	37.3	37.3	38.2	36.4
Beginner position	53.3	40.0	38.1	38.1	42.2
Micro establishment (1–9 employees)	80.1	78.2	77.8	78.1	77.2
No observed hire in prior 12 months	34.0	33.3	33.1	32.3	32.6
Occupational groups (%)					
Blue-collar	30.5	29.6	29.9	31.4	35.3
White-collar	46.5	47.0	49.4	51.3	51.3
Sales and retail	18.2	17.7	19.5	22.1	25.5
Hospitality, tourism and leisure	15.6	18.6	20.0	19.8	18.0
Construction and public works	12.7	12.8	9.4	6.2	3.2
Installation and maintenance	8.9	8.5	6.3	4.3	2.3
Business support	8.7	6.8	9.8	14.6	20.1
Industry sections (%)					
Wholesale and retail trade	17.2	16.1	17.2	19.7	23.6
Construction	12.8	12.5	9.7	7.5	5.3
Accommodation and food services	11.9	14.6	16.1	16.1	15.3
Manufacturing	8.9	9.0	8.6	8.2	7.4
Applicant characteristics					
Women (%)	20.0	19.0	21.3	24.3	26.8
Tertiary educated, excluding CAP (%)	6.7	6.2	6.4	6.3	5.7
Unemployed over one year (%)	21.4	22.8	23.0	23.1	22.8
Mean age (years)	35.9	36.7	36.6	36.2	35.5
Shortlist-size range	1–2	3–5	6–8	9–14	15–100

Notes: Quintiles are formed by ranking vacancies on predicted shortlist size. ROME occupation groups use the first letter of the occupation code; industry NAF sections use the employer's two-digit industry code. Q1 is the quintile with the shortest shortlist; Q5 corresponds to the quintile with the longest shortlist size.

The IV estimates indicate that larger shortlists modestly increase the probability that the establishment posts another vacancy within six months. Shortlist size has a bigger impact on the likelihood that the establishment returns to accompanied preselection: a one-candidate increase raises the probability of a new accompanied-preselection offer within six months by roughly 1.1 percentage points, corresponding to a 2.3 percent increase relative to the baseline mean of 0.486. While the effect is not economically very important, it shows a potential alternative mechanism to the one explored by Algan et al. (2020). The return to the PES services seems to be less the effect of a decrease in hiring cost and more the impact of a successful proposal.

These results suggest that larger candidate pools not only improve immediate hiring outcomes but also increase establishments' propensity to reuse accompanied preselection in subsequent recruitment episodes.

TABLE 7. Effect of Shortlist Size on Subsequent Vacancy Posting

	Any new vacancy		*New preselection (C) vacancy*	
	3 months	6 months	3 months	6 months
Panel A: 2SLS (instrumented by lagged counselor leniency)				
Shortlist size	0.0051	0.0048	0.0102***	0.0111***
	(0.0030)	(0.0029)	(0.0031)	(0.0031)
Panel B: OLS				
Shortlist size	0.0000	0.0008***	0.0006***	0.0014***
	(0.00006)	(0.00006)	(0.00006)	(0.00007)
Baseline mean	0.514	0.605	0.410	0.486
Observations	809,675			
Office × Cohort FE	Yes			
Posting Month FE	Yes			

Notes: Each column reports estimates from a regression where the dependent variable equals one if the establishment posts a new vacancy within 3 or 6 months. Columns (1)–(2) consider any new vacancy; columns (3)–(4) restrict to vacancies using preselection. Panel A reports 2SLS estimates where shortlist size is instrumented using lagged counselor leniency. Panel B reports OLS estimates. All specifications include office-by-cohort and posting-month fixed effects. Standard errors are clustered at the counselor level. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Recruitment across channels To assess whether preselection displaces recruitment through other channels, I use hiring declarations covering all recruitments made by each establishment. Appendix Table A.19 reports estimated increases of 0.20 and 0.63 distinct recruits within six and 24 months per additional candidate on the first shortlist, although

neither estimate is statistically significant. The reduced-form decomposition shows that assignment to a more lenient caseworker increases hires of candidates chosen for the focal first vacancy and subsequent preselected vacancies. Effects on shortlisted-but-not-chosen candidates and recruits who never appeared on a shortlist remain statistically indistinguishable from zero. The estimates therefore document gains in recruitment through preselection, but it does not seem these gains displace recruitment through other channels.

Re-use of preselection The increase in hires associated with subsequent preselected vacancies appears to operate through both greater re-use of the program and larger later shortlists. A one-standard-deviation increase in focal-caseworker leniency is associated with 0.013 additional subsequent preselection vacancies within six months and 0.027 within 24 months, while the corresponding increase in the number of subsequent shortlist positions is 0.39 and 0.85. The accounting exercise in Appendix Table A.20 assigns a similar contribution to the extensive and intensive margins, although together they explain only about 40 percent of the reduced-form increase in hires chosen on subsequent vacancies. Appendix Table A.21 provides no evidence that the causal return to an additional candidate becomes larger on later vacancies: the pooled estimate for vacancies ranked second or higher is smaller and less precise than the estimate for the focal vacancy. These results are consistent with establishments returning to preselection, but do not support a strong learning or dynamic-complementarity interpretation.

7. The change in the choice set and its impact on selection

The previous section established that a larger shortlist raises the probability that the firm chooses a candidate from within the presented choice set. This section turns to the question of how expansion changes the *composition* of the transmitted shortlist and whether the added noise degrades the firm's *selection*.

7.1. Shortlists become noisier

Table 8 reports 2SLS estimates of the effect of one additional transmitted candidate on the characteristics of the shortlist, grouped into the three families of the model: demographics, labor-market quality, and quality of fit. Three patterns emerge.

First, expansion is close to neutral on demographics. The shares of women, foreigners, and young applicants are essentially unchanged by an additional transmitted candidate; the only demographic that moves is average age, which falls slightly. A more lenient caseworker does not transmit a demographically different pool, which is reassuring for the interpretation of the instrument and rules out the concern that leniency proxies for screening along protected characteristics.

Second, broad measures of applicant quality are stable, but longer shortlists have somewhat lower average fit. Shortlist expansion induced by caseworker leniency leaves average education and reservation wages largely unchanged, while reducing average experience and the share of candidates from the vacancy's own occupation. The quality of fit index declines slightly, and the *minimum* quality in the pool falls more sharply. These changes in shortlist composition are consistent with less restrictive screening admitting candidates from further down the fit distribution. This interpretation concerns the candidates added through leniency-induced expansion as a group; the data do not identify which individual candidates would otherwise have been screened out.

Third, and most importantly, the *number* of usable candidates rises even as their *share* falls. The share of shortlisted candidates above a basic quality threshold declines slightly, but the count of such candidates increases by about 0.12 per additional transmitted candidate. This is a dilution-in-shares, enrichment-in-levels pattern at the heart of the mechanism: expansion adds noise to the pool while simultaneously broadening the set of credible options the firm can choose from. It also allows us to distinguish the quantity margin from a quality margin. Were lenient caseworkers simply better screeners, average fit would hold steady or improve as the shortlist lengthened; instead, it falls, which is what a lower transmission threshold with more candidates drawn from lower in the fit distribution mechanically produces. The composition results therefore both characterize the mechanism and add support to the exclusion restriction.

7.2. No change in the firm's selection

No change in the characteristics of the contract proposal If the firm could not filter the noise that expansion adds, a longer shortlist should pull the selected candidate down the quality distribution and slow the decision. Table 9 shows that it does neither. The estimates restrict attention to vacancies that result in a hire and report the conditional effect of an additional transmitted candidate on the characteristics of the selected candidate.

The selected candidate's observable quality is unaffected. An additional transmitted candidate leaves the chosen worker's reservation wage, education, experience, and summary fit statistically unchanged; the point estimates are small and, for the quality and fit measures, insignificant after the Holm correction. The demographic composition of hires is likewise stable. In short, the firm does not hire a worse worker when handed a longer, noisier shortlist.Appendix Table A.17 reports complementary estimates using centered outcomes that retain vacancies without a contract proposal, the quality and fit measures remain statistically insignificant after adjustment for multiple testing.

The rank evidence provides some sort of explanation, although one has to be careful as the estimands are not systematically robust. If the firm uses its own information to order the applications, the hired candidate is probably concentrated at the top: an additional transmitted candidate raises the probability that the hire is the highest-quality available

TABLE 8. Effect of Shortlist Size on Shortlist Composition (2SLS)

	Effect of +1 candidate (2SLS)	Baseline mean	Obs.
Demographics			
Average age	−0.235***	36.57	809,675
Share women	0.000	0.218	809,675
Share foreigners	0.000	0.060	809,675
Share under 25	0.001	0.106	809,675
Share older applicants	−0.003**	0.078	809,675
Labor-market quality			
Log average wage	0.002	7.53	809,675
Share tertiary educated (including CAP)	−0.002*	0.135	809,675
Average experience (years)	−0.160***	5.62	809,675
Share unemployed over one year	−0.003*	0.219	809,675
Quality of fit			
Average quality index	−0.011***	0.421	809,675
Minimum quality index in pool	−0.024***	0.074	809,675
Share in same occupation	−0.008***	0.149	809,675
Share qualified enough	−0.009***	0.153	809,675
Average experience within occupation	−0.083***	1.10	809,675
Share above quality threshold	−0.006***	0.166	809,675
Number above quality threshold	**0.117***	1.48	809,675

Notes: Each row reports a separate 2SLS estimate of the effect of one additional transmitted candidate on a characteristic of the shortlisted pool, instrumenting shortlist size with predetermined caseworker leniency. All specifications include office-by-cohort, posting-month, and occupation fixed effects; standard errors are clustered at the caseworker level. Stars are based on Holm-adjusted p-values within outcome family. The quality index follows Anderson (2008). OLS estimates and full diagnostics are in Appendix Table A.14. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

candidate and that she is among the top three by quality, even though, by the caseworker's transmission order, the hire sits well down the list[6]. Firms seem to re-rank the transmitted dossiers on fit and hire near the top of that ranking, regardless of the order in which candidates were sent. This corresponds to the optimal-stopping logic of the model: the firm inspects in descending order of fit, so a longer shortlist raises the chance that a

[6]In Appendix Table A.15, transmission rank refers to the order in which shortlisted applications reach the firm. Caseworkers usually transmit applications in batches, although some arrive at different times. One additional shortlisted candidate raises the selected candidate's transmission rank by approximately 0.8 positions, with no statistically significant change in the probability that the selected candidate is among the first five transmitted. This increase may partly reflect the larger shortlist: transmission order need not correspond to the firm's inspection order, particularly within batches, so these estimates do not establish sequential inspection by firms.

high-quality candidate is present without changing whom the firm ends up choosing among the strongest available.

Finally, expansion does not delay the decision. The effect of an additional candidate on the time to a contract proposal is small and insignificant, consistent with the fact that candidates added at the margin join the bottom of the evaluation queue and are rarely reached, so the number of candidates the firm actually inspects, and hence the time to decision, does not rise in proportion to the shortlist. Taken together, the composition and selection results indicate that shortlist expansion adds noise to the pool but not to the choice. Firms filter signal from noise, choosing workers of unchanged quality without delay, because intermediation has already curated the pool into a set the firm can triage.

TABLE 9. Effect of Shortlist Size on the Selected Candidate (2SLS)

	Effect of +1 candidate (2SLS)	Baseline mean	Obs.
Quality and fit of the hire			
Average quality index	−0.009	0.395	242,093
Has required education	−0.013*	0.143	242,093
Average experience (years)	0.120	4.84	242,093
Reservation wage above 2,000	0.010*	0.107	242,093
Share in same occupation	−0.001	0.142	242,093
Rank of the hire (by quality)			
Hire is top 1 by quality	0.034***	0.231	242,093
Hire is in top 3 by quality	0.028**	0.358	242,093
Hire is in top 5 by quality	0.009	0.408	242,093
Timing			
Time to contract proposal (days)	−1.02	21.0	242,093

Notes: The sample is restricted to vacancies with a successful recruitment. Each row reports a separate 2SLS estimate of the effect of one additional transmitted candidate on a characteristic of the selected candidate, instrumenting shortlist size with predetermined caseworker leniency. All specifications include office-by-cohort, posting-month, and occupation fixed effects; standard errors are clustered at the caseworker level. "Top k by quality" indicates that the hire is among the k highest-quality transmitted candidates. Stars are based on Holm-adjusted p-values within outcome family. Full results, including demographics, the transmission-rank measures, and OLS columns, are in Appendix Table A.15. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

No change in the characteristics of recruited workers The characteristics of candidates receiving a contract proposal describe the firm's selection, but not necessarily the workers who eventually join the establishment. Candidates may decline a proposal, and this second stage could change the composition of realized hires. Appendix Table A.16 therefore

examines the characteristics of workers whose recruitment is observed in the hiring declarations.

There is no evidence that assignment to a more lenient caseworker leads firms to recruit observably different workers. The reduced-form effects on education, experience, reservation wage, and observable match quality are small and statistically indistinguishable from zero. The decline in average fit within the shortlist therefore does not translate into a detectable deterioration in the measured characteristics of recruited workers.

These results are less precise than those for candidates receiving a proposal. The analysis conditions on realized hiring, which is itself affected by shortlist size, and the first stage is weaker in this restricted sample. The estimates should therefore be read as an absence of evidence of deterioration, rather than as precise zero effects or as a comparison of the same vacancies under different shortlist sizes. Together with the proposal results, they are consistent with firms being able to filter the additional noise in longer shortlists, without establishing how they inspect candidates. Appendix Table A.18 reports centered unconditional estimates that retain vacancies without a realized hire. Most characteristics remain statistically insignificant after adjustment for multiple testing, although the centered measure for meeting required education declines.

7.3. Acceptance among proposed candidates

The distinction between contract proposals and realized hires allows me to look at the completion of the match. I link proposed candidates to qualifying hiring declarations and, for this exercise, assume that proposed candidates without a qualifying linked hire refused the proposal. The analysis therefore measures acceptance under this assumption. The descriptive comparisons use 362,600 candidate–vacancy pairs retained by the acceptance model.[7]

Occupational compatibility is associated with a higher probability of acceptance. Among proposed candidates, 26.53 percent are hired when their occupation matches that of the vacancy, compared with 23.78 percent otherwise. The difference of 2.74 percentage points has a caseworker-clustered standard error of 0.47 and remains statistically significant after Holm adjustment across all 26 reported comparisons (Appendix Table A.24). The gap is positive in each tightness quartile, ranging from 2.02 to 3.13 percentage points. All four differences remain significant after adjustment within the occupation-match family. In contrast, there is no statistically detectable difference associated with meeting the vacancy's educational requirement, and reported experience displays no clear monotonic acceptance pattern.

These findings are consistent with the model's distinction between a candidate whom

[7]Appendix Table A.23 and Figure A.3 report in-sample calibration. Figures A.2, A.4 and A.5 compare observed and fitted acceptance by occupational match, wage compatibility and experience. Table A.26 reports the corresponding exploratory wage-gap comparisons.

the firm wishes to recruit and a match that is attractive to both parties. Occupational compatibility provides one observed dimension along which willingness to accept may increase.

The conditional vacancy-level estimates provide a complementary comparison. One additional transmitted candidate is associated with a change of −0.27 percentage points in the share of proposed candidates who are hired, with a standard error of 0.43, relative to a mean of 22.79 percent (Appendix Table A.22). The 2SLS estimate is not statistically significant. A longer shortlist can provide further options when earlier candidates are unsuitable or refuse, without increasing the share of proposals that lead to hiring.

8. Labor-market tightness

The conceptual framework allows the value of an additional candidate to vary with labour-market conditions. A weaker outside option can make more candidates worth considering, while lower acceptance probabilities can limit the hiring response.

Lower acceptance when candidates are scarce Table 10 reports estimates by quartile of local labour-market tightness, measured as the vacancy-to-unemployment ratio averaged across quarters within department and occupation. The proposal effect rises from 1.74 to 4.78 percentage points. This is consistent with the interpretation of an additional candidate being more valuable where candidate pools are thin.

However, the hiring response is less clearly ordered by tightness. The effect on the probability of a hire rises from 0.44 percentage points in the slackest quartile to 1.52 in the third quartile, before falling to 1.01 in the tightest. These estimates indicate a weaker translation of employer interest into completed matches in the tightest markets, even if they are not precise enough to establish a statistically significant difference between the third and fourth quartiles.

The candidate-level acceptance comparisons provide separate descriptive evidence on this stage. In the retained proposal sample, acceptance is 22.97 percent in the tightest quartile and 25.32 percent in the other three quartiles pooled. The difference is −2.35 percentage points (Appendix Table A.25) and remains significant even after adjusting for multiple hypothesis testing. As candidates have stronger outside opportunities in tight markets, they seem to become less likely to accept contract proposals.

Comparisons within broad groups of predicted shortlist size do not establish an additional tightness gradient: none of the pairwise differences survives the reported multiple-testing adjustment, and the weak-instrument-robust confidence sets are often unbounded. The quarterly extension likewise does not identify how the shortlist effect changes when a given market tightens over time, because the first stage for the within-market interaction is weak. The static gradient is therefore interpreted as suggestive

TABLE 10. Shortlist Effects and Elasticities by Labour-Market Tightness

	Q1 (slackest)	Q2	Q3	Q4 (tightest)
Panel A: Contract proposals				
Effect per candidate (pp)	1.74**	2.81***	4.32***	4.78***
	(0.54)	(0.59)	(1.07)	(1.33)
Outcome mean	0.320	0.305	0.299	0.268
Elasticity	0.657	0.905	1.191	1.206
95% CI for elasticity	[0.25, 1.06]	[0.54, 1.27]	[0.61, 1.77]	[0.55, 1.86]
Panel B: Any hire within six months				
Effect per candidate (pp)	0.44	0.92**	1.52**	1.01
	(0.31)	(0.33)	(0.53)	(0.56)
Outcome mean	0.133	0.115	0.114	0.102
Elasticity	0.401	0.783	1.096	0.669
95% CI for elasticity	[-0.16, 0.96]	[0.23, 1.33]	[0.35, 1.85]	[-0.06, 1.39]
Mean shortlist size	12.09	9.80	8.23	6.75
Observations	189,008	188,941	189,014	188,918

Notes: Tightness is the vacancy-to-unemployment ratio averaged across quarters within department and occupation. Quartiles classify persistent market differences. Each estimate is from a separate 2SLS regression with predetermined caseworker leniency as instrument, office-by-cohort and posting-month fixed effects, and caseworker-clustered standard errors. Effects and standard errors are in percentage points (pp). Elasticities and their conventional 95% intervals use within-quartile means; normalizing means are treated as fixed. Stars are based on two-sided tests using the reported standard errors: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. No inference on differences follows from comparing individual confidence intervals.

heterogeneity across markets, rather than as evidence of a causal response to changing tightness.

8.1. Recruitment takes longer in the tightest markets

The previous results concern whether firms propose a contract and whether the proposal leads to a hire. A larger choice set may also affect how long recruitment takes. Even when additional candidates improve the firm's chances of finding a suitable worker, evaluating them or continuing after an unsuccessful proposal can extend the process. The seminal model by Mortensen and Pissarides (1994) theorises that the cost of hiring comes from keeping a vacancy open, no matter the initial number of applicants it receives. Simply lengthening the recruitment process can induce additional costs for the firms. Table 11 reports the effect of one additional shortlisted candidate on the time to the last recorded contract proposal, overall and by quartile of labour-market tightness.

The overall estimate is an increase of 0.48 days, relative to a mean of 22.25 days, and is not statistically significant. The point estimates rise across tightness quartiles, from 0.03 days in the slackest markets to 3.27 days in the tightest. Only the latter is statistically significant at the five-percent level. Relative to a mean duration of 25.27 days in the tightest quartile, this represents an increase of approximately 13 percent.

This pattern suggests that access to additional candidates can come with a longer recruitment process where candidates are scarce. It is consistent with firms continuing to inspect candidates when earlier ones are unsuitable or do not accept, although neither inspection effort nor the reason for the additional time is directly observed. A later last proposal reflects additional proposals rather than a slower first decision.

TABLE 11. Effect of Shortlist Size on Time to the Last Contract Proposal

	Overall	Q1 Slackest	Q2	Q3	Q4 Tightest
Effect of +1 candidate (days)	0.483	0.032	0.276	0.642	3.267*
	(0.366)	(0.394)	(0.653)	(1.130)	(1.606)
Outcome mean	22.25	19.05	22.18	23.93	25.27
Observations	229,744	57,326	54,749	53,892	48,124

Notes: The outcome is time to the last recorded contract proposal, measured in days. Each column reports a separate 2SLS estimate of the effect of one additional shortlisted candidate, instrumenting shortlist size with predetermined caseworker leniency. The sample is restricted to vacancies with a recorded contract proposal and an observed timing outcome. Tightness quartiles run from the slackest (Q1) to the tightest (Q4) markets. Caseworker-clustered standard errors are in parentheses. Stars are based on unadjusted p-values: $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

9. Discussion

Expanding a *curated* shortlist raises the probability of a hire because an extra candidate is an option the firm can take or leave. A compatible framework suggests that the firm evaluates the transmitted dossiers in descending order of fit and stops once further inspection is no longer worthwhile. This allows the choice set to expand without requiring a proportional increase in inspection effort. The evidence shows little change in selected candidates' overall measured quality or in decision time. This section draws out four implications of that interpretation.

Filtering and expansion are complements. The paper isolates a margin distinct from the one studied by Algan et al. (2020), who evaluate a package of recruitment services through which caseworkers assist firms in finding candidates. The margin identified here operates within the curated regime: holding screening precision fixed, it varies the *size* of the shortlist the caseworker transmits. Filtering and expansion can therefore be complementary. A marginal candidate added by a more lenient caseworker has already passed an initial screen and arrives with an informative fit signal, so the firm can rank her cheaply and inspect her only if she is promising. This is why the same expansion that helps here could plausibly harm a firm facing a raw, unranked pool, where additional

applications require attention regardless of their relevance. The value documented in this paper is therefore the value of an additional candidate within a curated pool, and it should be read as conditional on intermediation having already done its filtering work.

This distinction is important in the broader context of algorithmic hiring. Firms that use machine-learning or AI tools to screen CVs at the initial stage often do not face the raw applicant pool at the point of final review: they face a list that has already been curated, ranked, or truncated by a scoring rule. Such tools therefore involve two conceptually distinct choices. The first is the quality and objective of the screening rule itself: which applicant characteristics the algorithm uses, how accurately it predicts fit, and whether it reproduces or corrects patterns in historical hiring decisions (Cowgill, 2020). The second is the quantity margin studied here: conditional on a given screening rule, how many candidates should be advanced to the recruiter. The results in this paper show that predictive ranking alone is not the full design problem. Once a screening technology has generated a credible, rankable set of candidates, the cutoff that determines shortlist size matters in its own right.

Leniency and screening quality. A natural concern in any judge design is that the instrument captures more than the channel of interest: here, more lenient caseworkers might also be better or worse screeners, so that variation in shortlist size is confounded with variation in screening quality. The model distinguishes these two dimensions. Leniency determines how far down a given ranking candidates are transmitted, while screening precision determines which candidates are transmitted at a given shortlist size. These are two dimensions of a potentially multifaceted caseworker, whose decisions may also reflect preferences over applicant characteristics or differences in how they assist employers.

The composition results are consistent with the leniency margin. Shortlist expansion induced by caseworker leniency lowers average fit and minimum quality, while raising the *number* of candidates above a fixed quality threshold. The separate checks using caseworkers' predetermined sorting of applicants on observable characteristics provide additional support for this interpretation. The estimates are therefore interpreted as the effect of expanding the choice set through more lenient screening, including the accompanying changes in candidate composition.

The paper focuses on this particular dimension of caseworker behavior. However, by documenting quasi-random assignment and the screening mechanism through which caseworkers shape firms' choice sets, it also opens the door to research on other dimensions of intermediation. How screening precision, caseworkers' preferences, and employer assistance interact remains an open question and provides an interesting avenue for further research.

The estimates are partial-equilibrium. The judge design identifies the effect of an additional transmitted candidate on the firm to which that candidate is sent, at prevailing market conditions. It is silent on whether the candidates a more lenient caseworker transmits to one firm are thereby diverted from another, or whether firms that hire more readily through preselection crowd out hiring elsewhere. The active-labour-market-policy literature shows that such equilibrium effects can be substantial: placement assistance can displace untreated job seekers (Crépon et al., 2013), while interventions can also generate positive spillovers onto non-targeted workers (Lalive et al., 2015). The distinction between recruitment frictions and job rationing in Michaillat (2012) further explains why improving the matching process need not produce the same employment gains across market conditions.

The establishment-level results extend the analysis beyond the focal vacancy, but do not resolve this market-level question. More lenient caseworker assignment increases subsequent use of preselection and hiring on later preselected vacancies. Effects on total establishment hiring and recruitment outside the observed shortlists remain imprecise. These considerations are especially important when considering tightness heterogeneity. While increasing applicant pools in tight markets might prove costly for firms that take longer to recruit, they could be welfare improving for applicants. The present evidence cannot separate these channels or establish that displacement is greatest in tight markets. In Crépon et al. (2013), displacement from assistance directed at workers is concentrated in weak labour markets. Quantifying the market-level consequences of expanding shortlists would require variation in the intensity of the policy across labour markets and outcomes for firms and workers beyond those directly assisted. The firm-level gains estimated here inform such an assessment but do not themselves deliver it. The aggregate effect also depends on where the additional candidates would otherwise search. Kircher (2022) shows that directing search from slack toward tight markets can improve aggregate outcomes when differences in productivity and search efficiency are sufficiently small. Shortlist expansion could therefore help fill vacancies by making candidates searching in slack markets available to firms in tighter ones, or redistribute candidates already competing within the same market. Distinguishing these mechanisms is left to future work.

Policy. The findings bear directly on how a public employment service should allocate scarce screening capacity. The marginal value of an additional transmitted candidate is large when the predicted shortlist is short and declines as it grows. This pattern carries over from proposals to realized hires. The results therefore suggest that directing additional candidates toward firms with thin candidate pools is more impactful than expanding all shortlists uniformly. A simple counterfactual exercise conducted in Appendix C.4 illustrates the potential recruitment gains from such a reallocation. These calculations hold the number of transmissions fixed, however, rather than their monetary cost: identifying

and assessing an additional candidate may require different amounts of caseworker effort across vacancies.

The same principle applies to digital and algorithmic screening tools. When a public employment service or a firm uses a statistical model, machine-learning system, or AI tool to rank CVs before human review, the relevant policy choice is not only whether to automate first-stage screening, but also where to set the threshold that determines how many candidates are rendered visible. The results suggest that shortlist thresholds should be more permissive when the predicted pool is thin and more selective when the recruiter already has several credible candidates to consider. The tightness heterogeneity is a weaker guide to policy: a larger proposal effect does not establish a larger hiring effect or a larger gain in employment across the market.

An algorithmic system should therefore be evaluated not only on its ability to rank applicants, but also on the composition of the candidates it excludes and advances at alternative thresholds. In particular, changing the shortlist cutoff should be assessed alongside changes in the demographic and skill composition of candidates who reach human review, as well as subsequent hiring outcomes. Within an existing screening process, the cutoff determining shortlist size is itself an important policy lever.

10. Conclusion

This paper studies whether expanding the set of candidates available to a recruiter improves hiring or makes decisions harder. The assisted preselection service of the French Public Employment Service provides a setting in which the candidates reaching the firm are observed. Using quasi-random assignment to caseworkers and their predetermined leniency as an instrument for shortlist size, I estimate the effect of transmitting an additional candidate, including the changes in composition that accompany this expansion.

One additional shortlisted candidate raises the probability of a contract proposal within three months by 3.1 percentage points and of a realized hire within six months by 0.8 percentage points. These represent increases of approximately 10 and 6.5 percent relative to the respective baseline means. The gains are concentrated among vacancies with small predicted shortlists and become smaller across groups with larger expected pools. This heterogeneity is consistent with an additional candidate being most useful when the firm has few suitable options, although it does not trace diminishing returns within the same vacancy.

Expansion also changes the composition of the choice set. Longer shortlists have lower average fit but contain more candidates above a basic quality threshold. Firms therefore receive more noise alongside more usable options. There is no broad deterioration in the measured quality of candidates receiving a proposal, and no statistically detectable change in the characteristics of workers actually recruited. Firms seem to be quite efficient in

filtering the additional noise.

Labour-market tightness brings out the distinction between finding a candidate the firm wishes to recruit and completing the match. The proposal effect is larger in tighter markets, where candidates are scarce, but the hiring response does not follow the same ordering. Due to workers having better outside options, the beneficial effect of having an extra candidate disappears in tight markets. Even worse, in these settings, expanding the choice set means that recruitment takes longer: in the tightest quartile, an additional candidate delays the last recorded proposal by 3.27 days, approximately 13 percent of the mean duration. The extension of the recruitment process is likely to be costly for firms if it is done without considering labour market conditions.

A simple framework of sequential inspection provides one possible interpretation. First-stage screening determines which candidates reach the firm, while the firm chooses which of these candidates to inspect more closely. When dossiers can be ranked cheaply and inspection is optional, an additional candidate can be ignored if evaluating her is not worthwhile. This gives shortlist expansion a nonnegative option value, without requiring inspection effort or decision times to remain unchanged. The lack of data on interviews, actual inspection and hiring costs makes it difficult to empirically establish this mechanism. Candidate acceptance adds a second consideration: weaker recruitment alternatives can make additional candidates worth inspecting, while lower acceptance probabilities reduce the expected return to doing so.

The results suggest directing additional screening capacity towards vacancies with thin expected candidate pools, rather than expanding every shortlist uniformly. They also highlight a distinct margin in the design of human or algorithmic screening: beyond how candidates are ranked, the threshold determining how many reach the recruiter matters.

References

Abel, M., Burger, R., and Piraino, P. (2020). The value of reference letters: Experimental evidence from south africa. *American Economic Journal: Applied Economics*, 12(3):40–71.

Abowd, J. M. and Kramarz, F. (2003). The costs of hiring and separations. *Labour Economics*, 10(5):499–530.

Algan, Y., Crépon, B., and Glover, D. (2020). Are active labor market policies directed at firms effective? evidence from a randomized evaluation with local employment agencies. *J-PAL working paper*.

Anderson, M. L. (2008). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the abecedarian, perry preschool, and early training projects. *Journal of the American statistical Association*, 103(484):1481–1495.

Angrist, J. D. and Imbens, G. W. (1995). Two-stage least squares estimation of average causal effects in models with variable treatment intensity. *Journal of the American statistical Association*, 90(430):431–442.

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

Balgova, M., Tekleselassie, T., Hensel, L., and Witte, M. (2025). Wage information and applicant selection.

Behaghel, L., Dromundo, S., Gurgand, M., Hazard, Y., and Zuber, T. (2024). *The potential of recommender systems for directing job search: A large scale experiment*. JSTOR.

Belot, M., Kircher, P., and Muller, P. (2022). How wage announcements affect job search—a field experiment. *American Economic Journal: Macroeconomics*, 14(4):1–67.

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

Birinci, S., See, K., and Wee, S. L. (2025). Job applications and labour market flows. *Review of Economic Studies*, 92(3):1438–1496.

Blair, P. Q. and Guo, R. (2026). Algorithmic credentialism. Technical report, National Bureau of Economic Research.

Blatter, M., Muehlemann, S., and Schenker, S. (2012). The costs of hiring skilled workers. *European Economic Review*, 56(1):20–35.

Burn, I., Button, P., Corella, L. M., and Neumark, D. (2022). Does ageist language in job ads predict age discrimination in hiring? *Journal of Labor Economics*, 40(3):613–667.

Caplin, A., Dean, M., and Leahy, J. (2019). Rational inattention, optimal consideration sets, and stochastic choice. *The Review of Economic Studies*, 86(3):1061–1094.

Card, D., Colella, F., and Lalive, R. (2025). Gender preferences in job vacancies and workplace gender diversity. *Review of Economic Studies*, 92(4):2437–2471.

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

Carranza, E., Garlick, R., Orkin, K., and Rankin, N. (2022). Job search and hiring with limited information about workseekers' skills. *American Economic Review*, 112(11):3547–3583.

Carvalho, J. R., Magnac, T., and Xiong, Q. (2019). College choice, selection, and allocation mechanisms: A structural empirical analysis. *Quantitative Economics*, 10(3):1233–1277.

Cowgill, B. (2020). Bias and productivity in humans and algorithms: Theory and evidence from resume screening. *Columbia Business School, Columbia University*, 29.

Crépon, B., Duflo, E., Gurgand, M., Rathelot, R., and Zamora, P. (2013). Do labor market policies have displacement effects? evidence from a clustered randomized experiment. *The quarterly journal of economics*, 128(2):531–580.

Dromundo Mokrani, S. and Haramboure, A. (2022). Do job counselors matter? measuring counselors value-added in job search. *Measuring Counselors Value-Added in Job Search*.

Fawcett, T. (2006). An introduction to roc analysis. *Pattern recognition letters*, 27(8):861–874.

Frandsen, B., Lefgren, L., and Leslie, E. (2023). Judging judge fixed effects. *American Economic Review*, 113(1):253–277.

Goldfarb, A. and Tucker, C. (2019). Digital economics. *Journal of economic literature*, 57(1):3–43.

Goldsmith-Pinkham, P., Hull, P., and Kolesár, M. (2025). Leniency designs: An operator's manual. Technical report, National Bureau of Economic Research.

Horton, J., Vasserman, S., and Watt, M. (2024). Reducing congestion in labor markets: A case study in simple market design.

Huberman, G. and Jiang, W. (2006). Offering versus choice in 401 (k) plans: Equity exposure and number of funds. *The Journal of Finance*, 61(2):763–801.

Humlum, A., Munch, J., and Rasmussen, M. (2025). What works for the unemployed? evidence from quasi-random caseworker assignments. Technical report, National Bureau of Economic Research.

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

Klaeui, J. (2026). When to broaden, when to focus: Job search and market tightness.

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

Lalive, R., Landais, C., and Zweimüller, J. (2015). Market externalities of large unemployment insurance extension programs. *The American Economic Review*, pages 3564–3596.

Le Barbanchon, T., Rathelot, R., and Roulet, A. (2016). Applications for job vacancies along the unemployment spell. Working paper.

Le Barbanchon, T., Rathelot, R., and Roulet, A. (2021). Gender differences in job search: Trading off commute against wage. *The Quarterly Journal of Economics*, 136(1):381–426.

Li, D., Raymond, L., and Bergman, P. (2026). Hiring as exploration. *Review of Economic Studies*, 93(2):1200–1240.

Michaillat, P. (2012). Do matching frictions explain unemployment? not in bad times. *American Economic Review*, 102(4):1721–50.

Mortensen, D. T. and Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. *The review of economic studies*, 61(3):397–415.

Muehlemann, S. and Leiser, M. S. (2018). Hiring costs and labor market tightness. *Labour Economics*, 52:122–131.

Pallais, A. (2014). Inefficient hiring in entry-level labor markets. *American Economic Review*, 104(11):3565–3599.

Pallais, A. and Sands, E. G. (2016). Why the referential treatment? evidence from field experiments on referrals. *Journal of Political Economy*, 124(6):1793–1828.

Pissarides, C. A. (2000). *Equilibrium unemployment theory*. MIT press.

Scheibehenne, B., Greifeneder, R., and Todd, P. M. (2010). Can there ever be too many options? a meta-analytic review of choice overload. *Journal of consumer research*, 37(3):409–425.

Weitzman, M. (1979). Optimal search for the best alternative. *Econometrica*, 47(3):641–54.

A. Supplementary tables and figures

A.1. Descriptive statistics

TABLE A.1. Caseworkers' Portfolio Characteristics by Preselection Status

	Preselecting	Non-Preselecting	Diff.
Panel A: Preselection Quarter			
Average portfolio size	43	95	-52
Portfolio composition (%)			
Followed	36	32	4
Guided	28	37	-9
Reinforced	9	13	-4
Personalized	11	5	6
Exhaustive	3	4	-1
Other	13	10	4
Panel B: All Quarters			
Average portfolio size	84	96	-13
Portfolio composition (%)			
Followed	33	31	3
Guided	37	36	0
Reinforced	13	12	0
Personalized	6	5	1
Exhaustive	1	5	-4
Other	10	10	0

Notes: This table compares the characteristics of caseworkers' job-seeker portfolios depending on whether the caseworker is involved in candidate preselection. Panel A restricts the sample to quarters in which a new firm enters the preselection process and is assigned a caseworker. Panel B reports the average number and types of job seekers followed per quarter, across all quarters. Portfolio composition categories are mutually exclusive and expressed as percentages of the total portfolio. The "Diff." column reports the difference between preselecting and non-preselecting caseworkers (Preselecting minus Non-Preselecting). Caseworkers involved in preselection are likely to follow a lower number of job seekers during the quarter when they are assigned a new firm, but on average over their careers they follow almost as many job-seekers as caseworkers who are never involved in preselection.

TABLE A.2. Applicant Characteristics Along the Recruitment Funnel, Conditional on a Successful Selection

	All applicants	Shortlisted	Successful	Successful – shortlisted
Demographics				
Average applicant age	35.72	35.90	34.01	-1.88
Share under 25	0.114	0.114	0.130	0.016
Share women	0.261	0.260	0.230	-0.030
Share foreigners	0.062	0.059	0.047	-0.012
Labor-market quality				
Average reservation wage (EUR/month)	2,064.89	2,066.82	2,061.38	-5.44
Share with salary above EUR 2,000	0.085	0.085	0.088	0.003
Share tertiary educated (including CAP)	0.130	0.131	0.124	-0.007
Average experience (years)	4.85	5.02	4.32	-0.70
Average unemployment duration (days)	685.08	663.70	588.92	-74.79
Share unemployed over one year	0.234	0.227	0.203	-0.024
Quality of fit				
Average quality index	0.350	0.357	0.360	0.003
Share qualified enough	0.144	0.138	0.121	-0.017
Average experience within occupation (years)	0.90	0.86	0.66	-0.20
Share with last position in same occupation as vacancy	0.143	0.133	0.113	-0.020
Number of job offers	242,093	242,093	242,093	
Average number of persons	19.75	9.79	1.51	–

Notes: This table is restricted to vacancies for which the firm makes a contract proposal to a candidate through the preselection channel. Column 1 reports the mean characteristics of all applicants to these vacancies; Column 2 reports the corresponding mean among candidates transmitted to the firm; and Column 3 reports characteristics of the candidate receiving the contract proposal. Column 4 is the difference between Columns 3 and 2. Candidate-level characteristics in Columns 1–3 are averaged across applicants within the relevant stage of the recruitment funnel. Variables measuring the number or presence of qualified candidates are shortlist-level measures and are therefore not defined for the single successful candidate. Rank measures are defined only for the successful candidate. The table is descriptive and does not condition on observables or use the caseworker-leniency instrument.

A.2. Assumptions

TABLE A.3. Balance on firm characteristics: first vacancy posting

	(1) Counselor effect	(2) Mean	(3) SD
Contract type			
Long-term contract	-0.005	0.546	0.498
	(0.002)		
Short-term contract	0.003	0.405	0.491
	(0.002)		
Qualification			
Blue Collar +	-0.002	0.088	0.284
	(0.001)		
Blue-Collar	-0.000	0.221	0.415
	(0.002)		
Manager	0.001	0.020	0.141
	(0.001)		
Other/Unknown	0.000	0.178	0.382
	(0.001)		
White-Collar	0.001	0.468	0.499
	(0.002)		
White-Collar +	-0.000	0.024	0.154
	(0.001)		
Firm size			
1–9 (Micro)	-0.003	0.808	0.394
	(0.001)		
10–49 (Small)	0.002	0.086	0.281
	(0.001)		
250+ (Large)	0.000	0.008	0.087
	(0.000)		
50–249 (Medium)	0.001	0.017	0.129
	(0.000)		
Unknown/Other	-0.001	0.082	0.274
	(0.001)		
Posted wage quartile			
Bottom 25%	0.004*	0.205	0.403
	(0.001)		
25-50%	-0.001	0.186	0.389
	(0.001)		
50-75%	-0.001	0.217	0.412
	(0.001)		
Top 25%	-0.003	0.170	0.375
	(0.001)		
Unknown	0.001	0.223	0.416
	(0.002)		
Labor-market tightness quartile			
Q1 (low)	0.001	0.250	0.433
	(0.002)		
Q2	0.001	0.250	0.433
	(0.002)		
Q3	0.003	0.250	0.433
	(0.002)		
Q4 (high)	-0.005*	0.250	0.433
	(0.001)		
Obs		373,172	

Notes: Each row reports the coefficient β from the regression $X_i = \alpha + \beta Z_i + \text{FE} + \varepsilon_i$, where Z_i is counselor leniency ($Z_{\text{judge,lag},z}$). All specifications include office-by-cohort and posting-month fixed effects. Standard errors are clustered at the counselor level. Categorical variables are indicator dummies. Significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

TABLE A.4. Balance on industry type: first vacancy posting

	(1) Counselor effect	(2) Mean	(3) SD
Industry			
A. Agriculture, forestry and fishing	0.001	0.028	0.165
	(0.001)		
B. Mining and quarrying	0.000	0.000	0.022
	(0.000)		
C. Manufacturing	0.000	0.071	0.257
	(0.001)		
D. Electricity, gas, steam and air conditioning supply	-0.000	0.002	0.040
	(0.000)		
E. Water supply; sewerage, waste management	0.000	0.005	0.073
	(0.000)		
F. Construction	-0.005	0.110	0.313
	(0.001)		
G. Wholesale and retail trade; repair of motor vehicles	-0.000	0.199	0.399
	(0.002)		
H. Transportation and storage	0.001	0.032	0.177
	(0.001)		
I. Accommodation and food service activities	0.005	0.141	0.348
	(0.002)		
J. Information and communication	0.001	0.018	0.133
	(0.001)		
K. Financial and insurance activities	-0.000	0.024	0.153
	(0.001)		
L. Real estate activities	0.000	0.020	0.139
	(0.001)		
M. Professional, scientific and technical activities	-0.001	0.099	0.298
	(0.001)		
N. Administrative and support service activities	-0.000	0.064	0.245
	(0.001)		
O. Public administration	0.001	0.017	0.129
	(0.000)		
P. Education	0.000	0.021	0.143
	(0.001)		
Q. Human health and social work activities	-0.001	0.061	0.239
	(0.001)		
R. Arts, entertainment and recreation	-0.000	0.017	0.129
	(0.000)		
S. Other service activities	-0.002	0.052	0.222
	(0.001)		
Obs		373,172	

Notes: Each row reports the coefficient β from the regression $X_i = \alpha + \beta Z_i + \text{FE} + \varepsilon_i$, where Z_i is counselor leniency ($Z_{\text{judge,lag,z}}$). All specifications include office-by-cohort and posting-month fixed effects. Standard errors are clustered at the counselor level. Categorical variables are indicator dummies. Significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

TABLE A.5. Balance on initial applicant-pool characteristics

	(1) Counselor effect	(2) Mean	(3) SD
Demographic characteristics			
Share over 50	-0.002	0.077	0.124
	(0.001)		
Share under 25	0.002	0.108	0.161
	(0.002)		
Share women	0.004	0.223	0.234
	(0.003)		
Share foreigners	0.001	0.062	0.113
	(0.001)		
Education and labor-market characteristics			
Average reservation wage	-6.068	2,201.326	2,459.542
	(27.133)		
Share tertiary educated (NV1+NV2)	-0.002	0.066	0.136
	(0.001)		
Share unemployed over one year	-0.003	0.229	0.190
	(0.002)		
Applicant quality			
Average quality index	-0.018*	0.420	0.490
	(0.006)		
Any applicant with quality index above 0.5	0.002	0.719	0.450
	(0.003)		
Obs		373,172	

Notes: Each row reports the coefficient β from the regression $X_i = \alpha + \beta Z_i + \text{FE} + \varepsilon_i$, where Z_i is counselor leniency ($Z_{\text{judge,lag,z}}$). All specifications include office-by-cohort and posting-month fixed effects. Standard errors are clustered at the counselor level. Categorical variables are indicator dummies. Significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The number of observations is 373,172 for average reservation wage and average quality index.

TABLE A.6. Decomposing Caseworker Effects: Leniency versus Screening Precision

	(1)	(2)	(3)	(4)	(5)
Dependent variable:	Hiring effect $\hat{\theta}_Y$			Residual $\hat{\theta}_Y$	Leniency $\hat{\theta}_X$
Leniency $\hat{\theta}_X$	0.0018	0.0021	0.0022*		
	(0.0011)	(0.0011)	(0.0011)		
Leniency squared $\hat{\theta}_X^2$		−0.0003			
		(0.0002)			
Precision (AUC, z)			0.0015	0.0015	0.024
			(0.0012)	(0.0012)	(0.021)
R^2	0.001	0.002	0.003	0.001	0.000
Caseworkers	3,975	3,975	3,617	3,617	3,617
Reliability $\tau^2/(\tau^2 + \bar{s}^2)$: $\hat{\theta}_Y = 0.56$, $\hat{\theta}_X = 0.36$					
Noise-corrected share of Var($\hat{\theta}_Y$) spanned by leniency: 0.002					

Notes: The unit of observation is the caseworker. $\hat{\theta}_Y$ is the caseworker's mean residual effect on contract proposals, estimated from vacancies posted in even quarters, and $\hat{\theta}_X$ is her mean residual effect on shortlist size, estimated from vacancies posted in odd quarters. Residuals are computed within office-by-cohort and posting-month cells. Using separate quarters prevents the two effects from sharing vacancy-level estimation noise. Caseworkers with fewer than thirty vacancies in either split are excluded. Precision is the standardized, office-residualized, empirical-Bayes Mann–Whitney AUC computed from the odd-quarter sample. Columns (1)–(3) regress $\hat{\theta}_Y$ on the listed variables. Column (4) relates precision to the residual from column (1), and column (5) relates leniency to precision. Regressions involving $\hat{\theta}_Y$ are weighted by $1/s_Y^2$. Standard errors are heteroskedasticity robust. Reliability is the estimated signal share of the variance of each caseworker effect. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

TABLE A.7. Exclusion restriction robustness: Caseworker Screening Precision

	Dependent variable: contract proposal within 3 months					
	Robustness			*Heterogeneity by precision*		
	(1) Baseline	(2) Precision sample	(3) + Precision control	(4) Below median	(5) Above median	(6) Inter-action
Shortlist size (2SLS)	0.0312***	0.0221***	0.0217***	0.0289***	0.0144***	0.0202***
	(0.0033)	(0.0040)	(0.0040)	(0.0087)	(0.0042)	(0.0057)
Shortlist size × precision (z)						0.0042
						(0.0102)
Observations	809,675	691,680	691,680	345,842	345,838	691,680
First-stage F	92.4	73.6	73.9	24.1	58.5	
Precision control	No	No	Yes	No	No	Yes
Office × Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Posting Month FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Vacancy-level 2SLS estimates of the effect of shortlist size on the probability of a contract proposal within three months. Shortlist size is instrumented with predetermined caseworker leniency, Z_c^{lag}. All specifications include office-by-cohort and posting-month fixed effects. Precision is the standardized, office-residualized, empirical-Bayes estimate of the caseworker's predetermined Mann–Whitney AUC based on applicant quality. Column (1) reports the baseline estimate. Columns (2) and (3) use the sample with a valid precision measure; column (3) additionally controls for precision. Columns (4) and (5) split this sample at median precision. Column (6) interacts shortlist size with demeaned precision and instruments the two endogenous variables with Z_c^{lag} and its interaction with precision. Standard errors are clustered by caseworker. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

TABLE A.8. Testing for Monotonicity with the Reverse Instrument

Dependent variable: Shortlist Size				
	q1	q2	q3	q4
Reverse instrument	0.62***	0.70***	1.16***	1.36***
	(0.09)	(0.14)	(0.16)	(0.15)
Observations	210,027	195,187	202,063	202,398
Dependent var. mean	7.06	8.37	9.85	11.58
Dependent var. SD	9.08	9.55	10.83	12.32
F-statistic (instrument)	47.25	26.10	52.32	85.69
P-value (F-statistic)	0.000	0.000	0.000	0.000

Notes: This table implements a monotonicity check using a "reverse" instrument constructed from the *lagged* counselor leniency measure. The sample is partitioned into quartiles based on predicted shortlist size, obtained from an OLS regression of *shortlist size* on predetermined characteristics with office × cohort fixed effects. For each quartile, the reverse instrument is defined as the counselor's lagged leniency index computed on observations outside the current estimation window and then aggregated within counselor, so that it varies only at the counselor level and is predetermined relative to the current vacancy. Each column reports a quartile-specific first-stage regression of *shortlist size* on the reverse instrument, including office-by-cohort and posting-month fixed effects. Standard errors are clustered at the counselor level.

A.3. Main results

TABLE A.9. Effect of Shortlist Size on contract offering

	Dependent variable: Contract offered within:			
	14 days	30 days	45 days	3 months
Panel A: 2SLS (instrumented by lagged counselor leniency)				
Shortlist size	0.0126***	0.0188***	0.0226***	0.0312***
	(0.0019)	(0.0025)	(0.0028)	(0.0033)
Panel B: OLS				
Shortlist size	-0.0034***	-0.0031***	-0.0022***	0.0006***
	(0.00010)	(0.00009)	(0.00008)	(0.00006)
Baseline mean	0.086	0.167	0.220	0.299
Observations	809,675			
Office × Cohort FE	Yes			
Posting Month FE	Yes			

Notes: Each column reports estimates from a separate regression where the dependent variable equals one if the vacancy is filled within the indicated time horizon. Panel A reports 2SLS estimates where shortlist size is instrumented using lagged counselor leniency. Panel B reports OLS estimates. All specifications include office-by-cohort and posting-month fixed effects. Standard errors are clustered at the counselor level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

TABLE A.10. Effect of Shortlist Size on Vacancy Outcomes (Only-Preselection Sample)

	Dependent variable: contract proposal within			
	14 days	30 days	45 days	3 months
Panel A: 2SLS (instrumented by lagged caseworker leniency)				
Shortlist size	0.010***	0.018***	0.022***	0.032***
	(0.002)	(0.003)	(0.004)	(0.005)
Panel B: OLS				
Shortlist size	−0.003***	−0.003***	−0.002***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)
Baseline mean	0.081	0.162	0.214	0.290
Observations	602,564			
Office × Cohort FE	Yes			
Posting Month FE	Yes			

Notes: The sample is restricted to vacancies receiving only preselection, with no bundled services. Each column corresponds to a horizon within which the firm proposes a contract to a shortlisted candidate. Panel A reports 2SLS estimates instrumenting shortlist size with predetermined caseworker leniency; Panel B reports OLS. All specifications include office-by-cohort and posting-month fixed effects. Standard errors, clustered at the caseworker level, are in parentheses. The first-stage F is the clustered excluded-instrument statistic. Stars are based on Holm-adjusted p-values computed across horizons within each model; all reported coefficients have raw and Holm-adjusted p-values below 0.001. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

TABLE A.11. OLS and IV Estimates by First-Time and Preselection Subsamples

	(1) First time	(2) Not first time	(3) Only presel.	(4) OP × first	(5) OP × not first
Panel A: 2SLS					
Shortlist size	0.036***	0.026***	0.032***	0.032***	0.033***
	(0.005)	(0.006)	(0.005)	(0.005)	(0.008)
Panel B: OLS					
Shortlist size	0.001***	0.001***	0.001***	0.001***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	373,172	436,503	602,564	283,013	319,551
Dependent var. mean	0.335	0.270	0.290	0.326	0.259
Mean shortlist size	8.96	9.43	9.24	9.03	9.44

Notes: The dependent variable equals one if the firm proposes a contract to a shortlisted candidate within three months. Panel B reports the OLS association between shortlist size and the outcome; Panel A instruments shortlist size with predetermined caseworker leniency, Z_c^{lag}. "OP" abbreviates "only preselection," i.e. vacancies receiving preselection with no bundled services; columns (4) and (5) intersect this restriction with first-time and later vacancies. All specifications include office-by-cohort and posting-month fixed effects, and standard errors are clustered at the caseworker level. The first-stage F is the clustered excluded-instrument statistic. Stars are based on Holm-adjusted p-values, computed separately across the OLS and the 2SLS coefficients; all raw and Holm-adjusted p-values are below 0.001. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

TABLE A.12. IV Responses Expressed as Elasticities at Sample Means

	Contract proposal	*Hiring*	
	Any proposal	**Any hire**	**Number of hires**
IV effect per candidate	0.0312	0.008	0.021
	(0.0033)	(0.002)	(0.004)
Outcome mean	0.299	0.123	0.170
Mean shortlist size	9.21	9.21	9.21
Elasticity	0.961	0.599	1.138
95% confidence interval	[0.762, 1.160]	[0.306, 0.893]	[0.713, 1.562]
Observations	809,675	809,675	809,675

Notes: The proposal outcome retains the original definition in the estimation data. The hiring outcomes are an indicator for any hire and the unconditional number of distinct hires within six months; all three columns include vacancies without success. Effects on indicators are in probability units. Elasticities equal $\widehat{\beta}_Y \bar{n}/\bar{Y}$, using means from the same estimation sample. The 95% intervals rescale the conventional clustered IV intervals, treating the sample means as fixed; they do not incorporate uncertainty in the normalizing means. Shortlist size is instrumented by predetermined caseworker leniency. Specifications include office-by-cohort and posting-month fixed effects; standard errors are clustered by caseworker.

TABLE A.13. Effect of shortlist size on realized hiring

	Horizon after vacancy posting					
	15 days	30 days	45 days	60 days	90 days	Six months
Panel A: Probability of at least one realized hire						
OLS: +1 shortlisted	0.000***	0.000***	0.001***	0.001***	0.001***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
2SLS: +1 shortlisted	0.003**	0.003**	0.004**	0.005**	0.007***	0.008***
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)
Outcome mean	0.044	0.066	0.080	0.089	0.103	0.123
Panel B: Number of distinct realized hires						
OLS: +1 shortlisted	0.003***	0.005***	0.007***	0.009***	0.011***	0.015***
	(0.000)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)
2SLS: +1 shortlisted	0.006**	0.008**	0.011**	0.014***	0.018***	0.021***
	(0.002)	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)
Outcome mean	0.055	0.085	0.106	0.120	0.140	0.170
Observations	809,675	809,675	809,675	809,675	809,675	809,675
First-stage F	92.4	92.4	92.4	92.4	92.4	92.4

Notes: Each coefficient is the effect of one additional shortlisted candidate. Shortlist size is instrumented with the leave-out caseworker-leniency measure. All specifications include office-by-cohort and posting-month fixed effects. Standard errors, clustered by preselection caseworker, are reported in parentheses. Stars use Holm-adjusted p-values within each panel: * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$. The first-stage F-statistic is the squared clustered t-statistic on the excluded instrument.

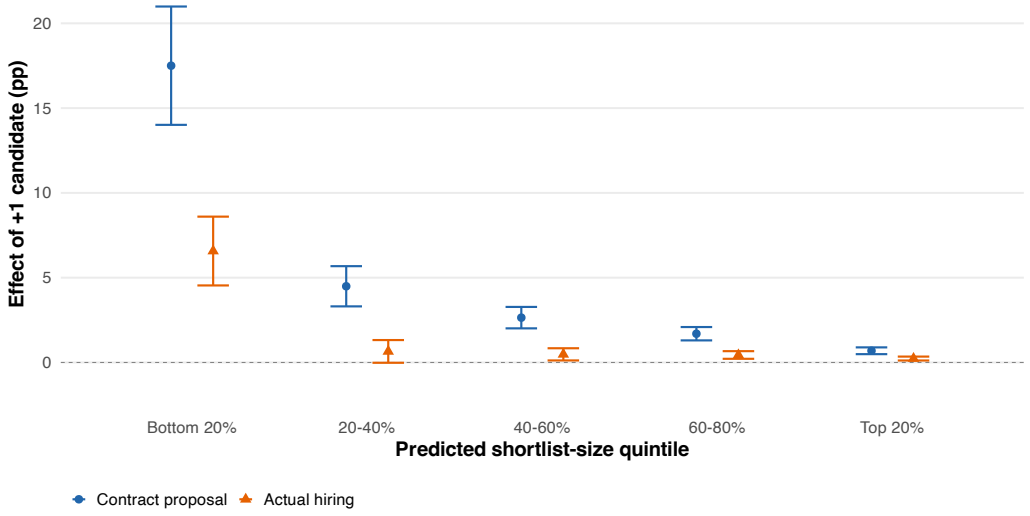


FIGURE A.1. Effects of one additional candidate by predicted shortlist size

Notes: The figure reports 2SLS estimates of the effect of one additional transmitted candidate on the probability of a contract proposal within three months (blue) and a realized hire within six months (orange), by quintile of predicted shortlist size. Effects are expressed in percentage points. Shortlist size is instrumented with predetermined caseworker leniency.

A.4. Candidate characteristics

TABLE A.14. Effect of Shortlist Size on Shortlist Composition: Full Results

	OLS coef.	2SLS coef.	Baseline mean	Obs.
Demographics				
Average age	−0.016***	−0.235***	36.57	809,675
	(0.001)	(0.049)		
Share women	0.0002***	0.000	0.218	809,675
	(2.0e−5)	(0.0011)		
Share foreigners	0.0002***	0.000	0.060	809,675
	(2.0e−5)	(0.0011)		
Share older applicants	−0.0002***	−0.003**	0.078	809,675
	(2.5e−5)	(0.0009)		
Share under 25	–	0.001	0.106	809,675
	–	–		
Labor-market quality				
Log average wage	–	0.002	7.53	809,675
	–	–		
Average reservation wage	−2.959***	−3.362	2,276.9	809,675
	(0.296)	(16.24)		
Average unemployment duration	0.253***	−3.784	649.0	809,675
	(0.059)	(2.915)		
Share with reservation wage above 2,000	−0.0004***	−0.0020	0.107	809,675
	(2.2e−5)	(0.0011)		
Share unemployed over one year	9.3e−5***	−0.003*	0.219	809,675
	(2.1e−5)	(0.0013)		
Share tertiary educated (including CAP)	−0.0003***	−0.002*	0.135	809,675
	(1.5e−5)	(0.0011)		

Continued on next page

TABLE A.14. Effect of Shortlist Size on Shortlist Composition: Full Results (continued)

	OLS coef.	2SLS coef.	Baseline mean	Obs.
Quality of fit				
Average quality index	−0.0011***	−0.011***	0.421	809,675
	(5.6e−5)	(0.0028)		
Minimum quality index in pool	−0.0003***	−0.024***	0.074	809,675
	(2.3e−5)	(0.0027)		
Share in same occupation	−9.2e−5***	−0.008***	0.149	809,675
	(2.7e−5)	(0.0018)		
Share qualified enough	3.8e−5	−0.009***	0.153	809,675
	(2.7e−5)	(0.0013)		
Average experience within occupation	−0.0008***	−0.083***	1.10	809,675
	(0.0002)	(0.0157)		
Average experience (years)	−0.0105***	−0.160***	5.62	809,675
	(0.0006)	(0.0307)		
Share below posted wage	0.0005***	−0.0022	0.176	809,675
	(2.9e−5)	(0.0015)		
Share above quality threshold	−0.0003***	−0.006***	0.166	809,675
	(2.3e−5)	(0.0013)		
Any candidate above quality threshold	0.0156***	0.0027	0.574	809,675
	(0.0005)	(0.0028)		
Number above quality threshold	**0.138****	**0.117****	1.48	809,675
	(0.0033)	(0.0108)		

Notes: Each row reports the coefficient on shortlist size from a separate regression of the listed shortlist characteristic. The OLS column regresses the outcome on shortlist size; the 2SLS column instruments shortlist size with predetermined caseworker leniency. All specifications include office-by-cohort, posting-month, and occupation fixed effects; standard errors, clustered at the caseworker level, are in parentheses. Stars are based on Holm-adjusted p-values computed within outcome family, separately for OLS and 2SLS. The quality index follows Anderson (2008); the threshold is a quality index of 0.5. Significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

TABLE A.15. Effect of Shortlist Size on the Selected Candidate: Full Results

	OLS coef.	2SLS coef.	Baseline mean	Obs.
Intensive margin				
Number of contract proposals	0.1078***	0.1406**	1.507	242,093
	(0.0098)	(0.0460)		
Demographics				
Average age	−0.016***	−0.018	34.32	242,093
	(0.003)	(0.139)		
Share women	0.0021***	0.0030	0.258	242,093
	(0.0001)	(0.0068)		
Share foreigners	0.0033***	0.0063	0.054	242,093
	(0.0002)	(0.0072)		
Share under 25	0.0028***	−0.0011	0.157	242,093
	(9.4e−5)	(0.0057)		
Labor-market quality				
Average reservation wage	−0.526	−22.76	2,191.5	242,093
	(0.779)	(16.24)		
Average unemployment duration	0.280	26.77	552.4	242,093
	(0.164)	(21.24)		
Reservation wage above 2,000	0.0019***	0.010*	0.107	242,093
	(9.1e−5)	(0.0052)		
Share unemployed over one year	0.0038***	0.0053	0.235	242,093
	(0.0002)	(0.0069)		
Share tertiary educated (excluding CAP)	0.0028***	0.0037	0.068	242,093
	(0.0001)	(0.0040)		
Average experience (years)	−0.0071***	0.120	4.84	242,093
	(0.0015)	(0.1935)		

Continued on next page

TABLE A.15. Effect of Shortlist Size on the Selected Candidate: Full Results (continued)

	OLS coef.	2SLS coef.	Baseline mean	Obs.
Quality of fit				
Average quality index	−0.0002	−0.009	0.395	242,093
	(0.0005)	(0.0176)		
Has required education	8.6e−6	−0.013*	0.143	242,093
	(0.0001)	(0.0064)		
Share in same occupation	0.0009***	−0.001	0.142	242,093
	(9.4e−5)	(0.0561)		
Average experience within occupation	−0.0002	−0.0641***	0.726	242,093
	(0.0005)	(0.0057)		
Share below posted wage	0.0033***	0.0092	0.176	242,093
	(0.0001)	(0.0561)		
Rank of the hire (by quality)				
Hire is top 1 by quality	–	0.034***	0.231	242,093
	–	–		
Hire is in top 3 by quality	−0.0050***	0.028**	0.358	242,093
	(0.0003)	(0.0076)		
Hire is in top 5 by quality	−0.0028***	0.009	0.408	242,093
	(0.0002)	(0.0079)		
Transmission rank and timing				
Preselection rank of hire	0.698***	0.7995***	7.266	242,093
	(0.0089)	(0.0089)		
Hire is in transmission top 5	−0.0170***	0.0088	0.585	242,093
	(0.0009)	(0.0079)		
Time to contract proposal (days)	−0.256	−1.02	21.0	242,093
	(0.228)	(0.582)		

Notes: The sample is restricted to vacancies with a successful recruitment. Each row reports the coefficient on shortlist size from a separate regression of the listed characteristic of the selected candidate. The OLS column regresses the outcome on shortlist size; the 2SLS column instruments shortlist size with predetermined caseworker leniency. All specifications include office-by-cohort, posting-month, and occupation fixed effects; standard errors, clustered at the caseworker level, are in parentheses. "Top k by quality" indicates the hire is among the k highest-quality transmitted candidates; "transmission top k" refers to the caseworker's transmission order. Stars are based on Holm-adjusted p-values within outcome family, separately for OLS and 2SLS. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

TABLE A.16. Effect of shortlist size on characteristics of realized hires

Outcome	2SLS estimate	Outcome mean	Observations
Mean education level (0–5)	-0.010	0.967	99,244
	(0.021)		
Mean years of experience	0.088	4.890	99,244
	(0.159)		
Mean monthly reservation wage	-39.046	2,009.977	99,244
	(71.281)		
Years of experience in vacancy occupation	-0.084	0.900	99,244
	(0.064)		
Mean Anderson quality index	-0.017	0.389	99,244
	(0.015)		
Share with exact education match	-0.007	0.079	99,244
	(0.005)		

Notes: Each row reports a separate 2SLS estimate of the effect of one additional shortlisted candidate. Outcomes are defined among vacancies with a realized hire for which the relevant characteristic is observed. Shortlist size is instrumented with the leave-out caseworker-leniency measure. All specifications include office-by-cohort and posting-month fixed effects. Standard errors, clustered by preselection caseworker, are in parentheses. The smaller samples and weaker first stages warrant caution when interpreting these conditional estimates.

TABLE A.17. Effect of Shortlist Size on Selected Candidates: Centered Unconditional Outcomes

Characteristic X	OLS effect on $S(X-c)$	2SLS effect on $S(X-c)$	Reference mean c	Obs.
Number of contract proposals				
Number of contract proposals	0.0494***	0.0227	1.507	809,675
	(0.0071)	(0.0076)		
Demographics				
Average age	−0.0066***	−0.0051	34.321	809,675
	(0.0007)	(0.0220)		
At least one woman	0.0013***	−0.0005	0.258	809,675
	(4.9×10^{-5})	(0.0011)		
At least one foreign candidate	0.0009***	0.0014	0.054	809,675
	(4.1×10^{-5})	(0.0006)		
At least one candidate aged 25 or younger	0.0014***	0.0015	0.157	809,675
	(0.0001)	(0.0009)		
At least one candidate over 50	0.0006***	0.0005	0.069	809,675
	(4.1×10^{-5})	(0.0006)		
Labor-market quality				
Average monthly reservation wage	−0.7266***	2.8360	2,041.635	809,675
	(0.1821)	(6.5930)		
Log average reservation wage	−0.0001	0.0008	7.440	809,675
	(3.6×10^{-5})	(0.0009)		
Average unemployment duration (days)	0.1090	2.4627	552.437	809,675
	(0.0570)	(1.3577)		
Reservation wage above 2,000	0.0008***	0.0018	0.107	809,675
	(4.0×10^{-5})	(0.0008)		
Reservation wage below 1,500	0.0015***	0.0005	0.213	809,675
	(0.0001)	(0.0011)		
Unemployed over one year	0.0016***	0.0011	0.235	809,675
	(0.0001)	(0.0011)		
Tertiary educated, including CAP	0.0011***	−0.0006	0.148	809,675
	(4.5×10^{-5})	(0.0009)		
Tertiary educated, excluding CAP	0.0008***	−0.0002	0.068	809,675
	(4.0×10^{-5})	(0.0007)		
Average experience (years)	−0.0032***	−0.0009	4.84	809,675
	(0.0004)	(0.0130)		

Continued on next page

TABLE A.17. Effect of Shortlist Size on Selected Candidates: Centered Unconditional Outcomes (continued)

Characteristic X	OLS effect on $S(X-c)$	2SLS effect on $S(X-c)$	Reference mean c	Obs.
Quality of fit				
Average quality index	-0.0003^{***}	-0.0015	0.395	809,675
	(4.1×10^{-5})	(0.0014)		
At least one candidate meeting required education	0.0003^{***}	-0.0026	0.143	809,675
	(3.1×10^{-5})	(0.0010)		
At least one candidate in the same occupation	0.0012^{***}	-0.0005	0.142	809,675
	(4.6×10^{-5})	(0.0009)		
Average experience within occupation (years)	1.4×10^{-5}	-0.0143	0.726	809,675
	(0.0002)	(0.0082)		
Rank by quality				
Selected candidate is in quality top 1	-0.0024^{***}	0.0011	0.231	809,675
	(0.0001)	(0.0012)		
Selected candidate is in quality top 3	-0.0018^{***}	-0.0004	0.358	809,675
	(0.0001)	(0.0012)		
Selected candidate is in quality top 5	-0.0009^{***}	-0.0001	0.408	809,675
	(0.0001)	(0.0012)		
Transmission rank				
Preselection rank of selected candidate	0.2645^{***}	0.1494^{***}	7.495	809,675
	(0.0105)	(0.0264)		
Selected candidate is transmission rank 1	-0.0049^{***}	0.0048	0.299	809,675
	(0.0002)	(0.0019)		
Selected candidate is in transmission top 3	-0.0059^{***}	0.0041	0.457	809,675
	(0.0002)	(0.0019)		
Selected candidate is in transmission top 5	-0.0063^{***}	0.0014	0.575	809,675
	(0.0002)	(0.0016)		
Preselection rank / shortlist size	-0.0013^{***}	0.0009	0.862	809,675
	(4.6×10^{-5})	(0.0009)		

Notes: Each row reports the effect of one additional shortlisted candidate on the centered unconditional outcome $S(X-c)$, where S indicates that at least one candidate receives a contract proposal and X is the listed vacancy-level characteristic of selected candidates. The outcome is zero when $S=0$; missing characteristics when $S=1$ remain missing. The reference mean c is the mean of X among vacancies with selection and observed X in the full eligible sample; it is not the mean of the centered outcome. Stars use the supplied Holm-adjusted p-values : * $p<0.05$, ** $p<0.01$, *** $p<0.001$.

TABLE A.18. Effect of Shortlist Size on Realized Hires: Unconditional Results

Outcome	2SLS estimate	Reference mean c	Obs.
Demographics			
Mean age	0.0063	34.114	809,675
	(0.0120)		
Share women	−0.0001	0.516	809,675
	(0.0005)		
Share under 25	0.0009	0.253	809,675
	(0.0005)		
Share under 35	−0.0006	0.580	809,675
	(0.0005)		
Share over 50	0.0002	0.118	809,675
	(0.0003)		
Share with nationality code other than 01	0.0001	0.034	809,675
	(0.0003)		
Labor-market quality			
Mean education level (0–5)	−0.0039	0.967	809,675
	(0.0018)		
Mean experience (years)	−0.0003	4.890	809,675
	(0.0074)		
Mean monthly reservation wage	−0.2677	2,009.977	809,675
	(3.7079)		
Mean unemployment duration (days)	0.2988	533.146	809,675
	(0.7113)		
Share with reservation wage above 2,000	0.0003	0.138	809,675
	(0.0004)		
Share unemployed over one year	0.0001	0.437	809,675
	(0.0005)		
Share tertiary educated (including CAP)	−0.0002	0.309	809,675
	(0.0005)		
Quality of fit			
Mean experience in vacancy occupation (years)	−0.0136	0.900	809,675
	(0.0051)		
Mean Anderson quality index	−0.0012	0.389	809,675
	(0.0008)		
Minimum Anderson quality index	−0.0015	0.358	809,675
	(0.0008)		

Continued on next page

TABLE A.18. Effect of Shortlist Size on Realized Hires: Unconditional Results (continued)

Outcome	2SLS estimate	Reference mean c	Obs.
Quality of fit (continued)			
Share in the same occupation	−0.0015	0.143	809,675
	(0.0005)		
Share below posted wage	−0.0003	0.188	809,675
	(0.0005)		
Share with exact education match	−0.0006	0.079	809,675
	(0.0004)		
Share overqualified	−0.0011	0.075	809,675
	(0.0004)		
Share meeting or exceeding required education	−0.0017*	0.154	809,675
	(0.0005)		
Share meeting required education with occupation experience	−0.0010*	0.051	809,675
	(0.0003)		
Share with quality index above 0.5	−0.0010	0.357	809,675
	(0.0005)		
Any hire with quality index above 0.5	−0.0007	0.384	809,675
	(0.0005)		
Number of hires with quality index above 0.5	−0.0002	0.423	809,675
	(0.0007)		
Rank			
Share of hires in transmission top 5	−0.0005	0.641	809,675
	(0.0007)		
Timing			
Mean time to recruitment (days)	−0.0016	47.952	809,675
	(0.0824)		

Notes: Each row reports the 2SLS effect of one additional shortlisted candidate on $H(X - c)$, where H indicates a qualifying hire and X is the listed vacancy-level characteristic. Outcomes equal zero without a hire; missing characteristics among hires remain missing. The reference c is the mean of X among vacancies with a hire and observed X in the full eligible sample, not the centered outcome mean. Shortlist size is instrumented with predetermined caseworker leniency, controlling for office-by-cohort and posting-month fixed effects. Parenthesized standard errors are clustered by caseworker. Stars retain the supplied Holm adjustment : * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

TABLE A.19. Effects on establishment-level recruitment and recruitment channels

	6 months	12 months	18 months	24 months
Panel A: Total distinct recruits, 2SLS				
Effect of +1 candidate on the focal shortlist	0.197	0.358	0.476	0.631
	(0.133)	(0.215)	(0.306)	(0.449)
Holm-adjusted p-value	0.387	0.387	0.387	0.387
Outcome mean	6.617	10.714	14.445	17.725
Panel B: Reduced-form decomposition, per SD of caseworker leniency				
Chosen for focal vacancy	0.0131	0.0128	0.0129	0.0129
	(0.0027)	(0.0028)	(0.0028)	(0.0028)
Chosen for subsequent preselected vacancy	0.0068	0.0104	0.0124	0.0145
	(0.0024)	(0.0033)	(0.0040)	(0.0044)
Shortlisted but never chosen	0.0035	0.0034	0.0035	0.0037
	(0.0024)	(0.0025)	(0.0025)	(0.0026)
Never shortlisted	0.0858	0.1714	0.2348	0.3186
	(0.0719)	(0.1164)	(0.1661)	(0.2445)
Total reduced form	0.1092	0.1981	0.2636	0.3497
	(0.0735)	(0.1183)	(0.1678)	(0.2457)
Pre-assignment hiring control	Yes	Yes	Yes	Yes
Observations	373,172	373,172	373,172	373,172

Notes: The unit of observation is the establishment's first preselection vacancy. Panel A reports 2SLS estimates of the effect of first-vacancy shortlist size on the number of distinct persons recruited by the same establishment before each horizon. Panel B reports reduced-form coefficients on standardized caseworker leniency. Its four categories are mutually exclusive and exhaustive, and all five regressions within a column use the same sample; their coefficients therefore add exactly to the total reduced form. All specifications control for distinct persons hired during the 12 months before first preselection and include office-by-cohort and posting-month fixed effects. Standard errors are clustered by preselection caseworker. Panel A p-values are Holm-adjusted across horizons. Repeated contracts for the same identified person are counted once within each horizon.

A.5. Subsequent recruitment

TABLE A.20. Subsequent preselection: extensive- and intensive-margin accounting

	6 months	12 months	18 months	24 months
Inputs				
RF: number of subsequent vacancies	0.0130	0.0179	0.0238	0.0274
RF: total subsequent shortlist positions	0.3912	0.6002	0.7291	0.8506
Mean shortlist per subsequent vacancy	9.1714	9.2725	9.3259	9.3666
Subsequent-vacancy fill rate	0.1198	0.1190	0.1173	0.1147
Later-vacancy IV effect per candidate	0.0045	0.0045	0.0045	0.0045
Accounting decomposition				
Implied extensive-margin hires	0.0016	0.0021	0.0028	0.0031
Implied intensive-margin hires	0.0012	0.0020	0.0023	0.0027
Total accounted hires	0.0028	0.0041	0.0051	0.0058
RF: hires chosen subsequently	0.0068	0.0104	0.0124	0.0145
Reconciliation residual	0.0040	0.0063	0.0073	0.0087

Notes: Reduced-form (RF) estimates are per one-standard-deviation increase in focal-caseworker leniency. Implied extensive-margin hires equal the reduced-form change in subsequent vacancies multiplied by their observed fill rate. The intensive-margin calculation nets from the change in later shortlist positions those positions attributable to additional vacancies and applies the pooled IV estimate for vacancies ranked second or higher. All specifications control for hiring during the 12 months before first preselection. This is an accounting exercise, not an identified mediation decomposition.

TABLE A.21. Effect of shortlist size on realized hiring, by vacancy rank

Vacancy rank	1	2	3	4	5+	2+ pooled
Reduced form	0.0060***	0.0009	0.0030	-0.0042	0.0050	0.0019
	(0.0013)	(0.0016)	(0.0029)	(0.0048)	(0.0037)	(0.0011)
2SLS: +1 shortlisted	0.0109***	0.0028	0.0095	-0.0046	0.0065	0.0045
	(0.0026)	(0.0047)	(0.0093)	(0.0055)	(0.0049)	(0.0027)
Outcome mean	0.150	0.115	0.109	0.103	0.091	0.104
Total vacancies			809,675			
First-stage F	90.6	31.0	6.4	9.7	14.6	37.2

Notes: The outcome is an indicator for a validated DPAE hire assigned to the vacancy. Rank is defined within establishment from its first observed preselection vacancy. Shortlist size is instrumented by the vacancy's caseworker-leniency measure. Specifications control for establishment hiring in the 12 months before first preselection and include office-by-cohort and posting-month fixed effects.

A.6. Acceptance rates

TABLE A.22. Shortlist size and acceptance conditional on a proposal

	Full sample	Labour-market tightness quartile			
	(1)	Q1 (low)	Q2	Q3	Q4 (high)
Panel A: 2SLS					
Shortlist size (pp)	−0.270	0.085	0.719	1.656	−1.155
	(0.427)	(0.574)	(0.875)	(2.416)	(1.444)
Holm-adjusted p-value	1.000	1.000	1.000	1.000	1.000
Panel B: OLS					
Shortlist size (pp)	−0.159***	−0.166***	−0.142***	−0.070*	−0.148***
	(0.012)	(0.019)	(0.020)	(0.027)	(0.025)
Mean acceptance rate (%)	22.79	23.56	22.24	23.25	22.77
Observations	411,565	107,948	97,807	92,229	83,625
First-stage F-statistic	38.73	17.46	13.06	2.32	8.91

Notes: The dependent variable is the number of distinct proposed candidates linked to a qualifying DPAE hire, divided by the number of distinct proposed candidates, among vacancies with at least one proposal that pass the proposal-coverage checks. It measures observed proposal-to-hire conversion. The mean is an unweighted average of vacancy-level rates. Proposal recipients are drawn from the aligned monthly proposal files, without an additional three-month restriction; hires follow the study's matching rules. Coefficients and standard errors are in percentage points per additional shortlisted candidate. Panel A instruments shortlist size with predetermined caseworker leniency; Panel B uses the same retained observations. All specifications include office-by-cohort and posting-month fixed effects. Standard errors, in parentheses, are clustered at the caseworker level. Quartiles refer to occupation-by-department labour-market tightness, from low to high; the full sample also includes vacancies with missing tightness. Stars use Holm-adjusted p-values across the five samples for this outcome, separately for 2SLS and OLS: $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$. Conditioning on a proposal selects on an outcome that shortlist size may affect.

TABLE A.23. In-sample calibration of fitted acceptance probabilities

Prediction decile	Pairs	Fitted (%)	Observed (%)	Difference (pp)
1	36,260	1.61	1.82	+0.21
2	36,260	7.78	8.16	+0.37
3	36,260	13.70	13.08	-0.63
4	36,260	18.35	17.90	-0.46
5	36,260	22.39	22.49	+0.09
6	36,260	26.36	26.22	-0.14
7	36,260	30.37	29.56	-0.81
8	36,260	34.64	35.54	+0.90
9	36,260	40.12	40.57	+0.46
10	36,260	51.54	51.54	0.00

Notes: Deciles use pooled cutoffs of the saved flexible-logit predictions. Quartile-specific summary counts and fitted means are aggregated using their candidate counts; difference is observed minus fitted. All 362,600 retained pairs are included. Across the 40 decile-by-tightness cells, the candidate-weighted mean absolute discrepancy is 0.87 percentage points. This is a grouped in-sample calibration measure, not an individual prediction error or out-of-sample validation. Fitted probabilities include the original controls and fixed effects. No inference for observed-minus-fitted differences is supplied: observed-rate intervals alone do not account for estimation of the predictions.

TABLE A.24. Acceptance and observed candidate–vacancy compatibility

Panel A: Occupational match

	Different (%)	Same (%)	Gap (SE)	p	p_H^F	p_H^A
Pooled	23.78	26.53	2.74 (0.47)	< 0.001	< 0.001	< 0.001
Q1	24.02	27.05	3.03 (0.92)	0.001	0.004	0.026
Q2	24.66	26.69	2.02 (0.92)	0.027	0.029	0.354
Q3	24.84	27.30	2.46 (1.00)	0.014	0.029	0.215
Q4	22.00	25.13	3.13 (0.96)	0.001	0.004	0.026

Panel B: Education requirement

	Below (%)	Meets (%)	Gap (SE)	p	p_H^F	p_H^A
Pooled	24.44	24.72	0.29 (0.63)	0.652	1.000	1.000
Q1	25.79	24.95	-0.84 (1.26)	0.504	1.000	1.000
Q2	25.38	25.35	-0.03 (1.27)	0.984	1.000	1.000
Q3	24.57	25.76	1.19 (1.36)	0.382	1.000	1.000
Q4	22.60	23.04	0.43 (1.09)	0.690	1.000	1.000

Notes: The sample contains 362,600 candidate–vacancy pairs retained by the flexible logit, with 89,519 qualifying hires and 7,869 caseworkers. Acceptance equals a qualifying linked hire; other proposed candidates are treated as refusing under the maintained assumption. These are candidate-weighted raw comparisons, without covariate adjustment. Rates are percentages; gaps and standard errors (SE) are in percentage points. Standard errors are clustered by caseworker, preserving covariance across groups. Two-sided tests use t_{G-1} critical values and the variance correction $G/(G-1)$, where G is the number of caseworkers contributing to the comparison. Gaps are same minus different occupation in Panel A and meets minus below requirement in Panel B. p_H^F adjusts by Holm within the indicated family; p_H^A adjusts across all 26 comparisons (five occupation, five education, twelve tightness and four wage comparisons). Comparisons are exploratory; these corrections do not account for all preceding specification searches. No significance stars are used.

TABLE A.25. Acceptance differences across labour-market tightness groups

Comparison	Q4 (%)	Other (%)	Gap (SE)	p	p_H^F	p_H^A
All proposed candidates						
Q4 versus Q1	22.97	25.06	-2.09 (1.09)	0.054	0.272	0.652
Q4 versus Q2	22.97	25.36	-2.39 (0.86)	0.005	0.053	0.100
Q4 versus Q3	22.97	25.62	-2.65 (0.89)	0.003	0.033	0.063
Q4 versus Q1-Q3 pooled	22.97	25.32	-2.35 (0.77)	0.002	0.027	0.053
Different occupation						
Q4 versus Q1	22.00	24.02	-2.02 (1.32)	0.127	0.272	1.000
Q4 versus Q2	22.00	24.66	-2.66 (1.05)	0.011	0.079	0.180
Q4 versus Q3	22.00	24.84	-2.84 (1.08)	0.009	0.079	0.159
Q4 versus Q1-Q3 pooled	22.00	24.47	-2.47 (0.95)	0.010	0.079	0.162
Same occupation						
Q4 versus Q1	25.13	27.05	-1.92 (1.05)	0.069	0.272	0.685
Q4 versus Q2	25.13	26.69	-1.56 (1.03)	0.129	0.272	1.000
Q4 versus Q3	25.13	27.30	-2.17 (1.15)	0.058	0.272	0.652
Q4 versus Q1-Q3 pooled	25.13	27.00	-1.88 (0.84)	0.025	0.151	0.353

Notes: The sample contains 362,600 candidate–vacancy pairs retained by the flexible logit, with 89,519 qualifying hires and 7,869 caseworkers. Acceptance equals a qualifying linked hire; other proposed candidates are treated as refusing under the maintained assumption. These are candidate-weighted raw comparisons, without covariate adjustment. Rates are percentages; gaps and standard errors (SE) are in percentage points. Standard errors are clustered by caseworker, preserving covariance across groups. Two-sided tests use t_{G-1} critical values and the variance correction $G/(G-1)$, where G is the number of caseworkers contributing to the comparison. Q4 is the tightest quartile. Each gap is Q4 minus the comparison group, using that group's own composition. p_H^F adjusts by Holm within the indicated family; p_H^A adjusts across all 26 comparisons (five occupation, five education, twelve tightness and four wage comparisons). Comparisons are exploratory; these corrections do not account for all preceding specification searches. No significance stars are used.

TABLE A.26. Acceptance across wage-gap bins: exploratory comparisons

Quartile	Low (%)	High (%)	Gap (SE)	p	p_H^F	p_H^A
Q1	24.88	25.01	0.13 (1.36)	0.923	0.923	1.000
Q2	28.40	23.85	-4.55 (1.55)	0.003	0.011	0.066
Q3	25.58	24.17	-1.41 (1.53)	0.355	0.711	1.000
Q4	25.95	21.84	-4.11 (1.37)	0.003	0.011	0.060

Notes: The wage gap is $\log(w_j^{\min}/r_i)$, where $w_j^{\min}$ is the posted minimum wage and r_i is the candidate's reservation wage, expressed on the same cleaned pay scale. Ten equal-width bins divide the common central range $[-0.4463, 0.2877]$, the intersection of quartile-specific 5th–95th percentile ranges. Low combines bins 1–3; high combines bins 8–10. Cell means weight candidate–vacancy pairs equally. Gaps are high minus low, not effects of increasing offered pay. The central-range plotting sample has 46,750 pairs; these tests use only its two tails. Caseworker-clustered SE and t_{G-1} tests use the same convention as Table A.24. p_H^F adjusts by Holm within the indicated family; p_H^A adjusts across all 26 comparisons (five occupation, five education, twelve tightness and four wage comparisons). Comparisons are exploratory; these corrections do not account for all preceding specification searches. No significance stars are used.

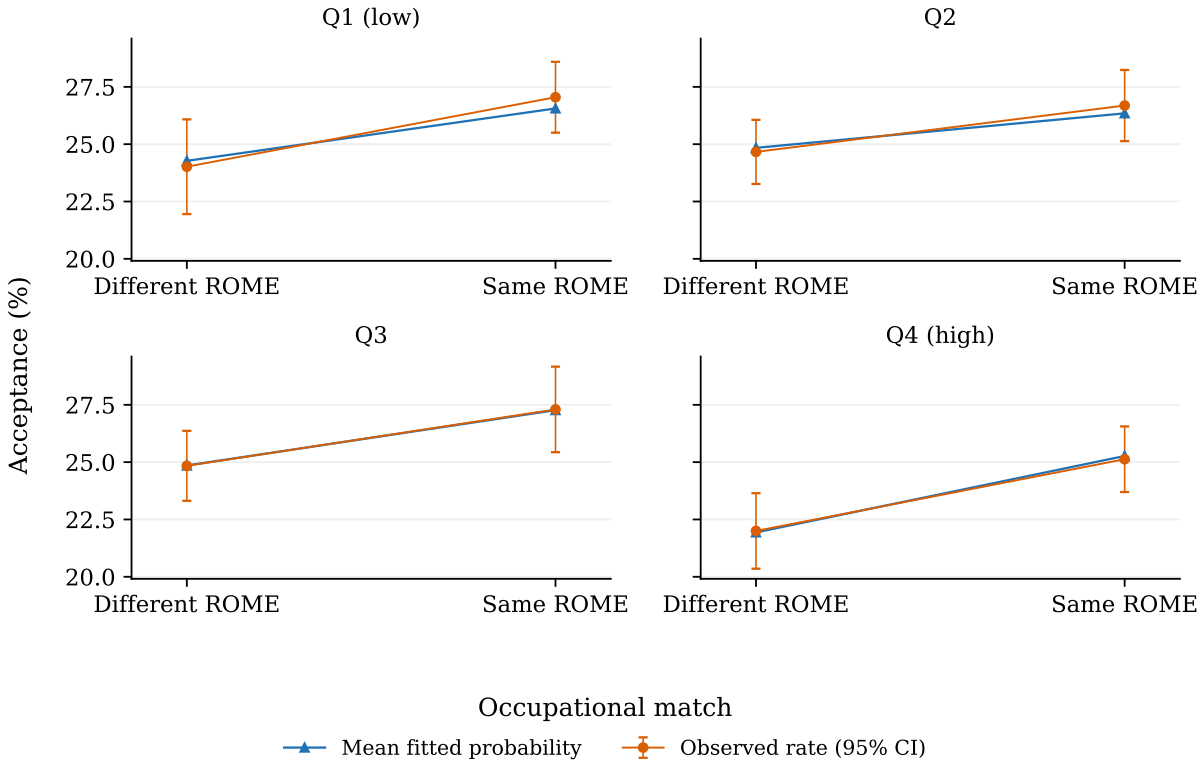


FIGURE A.2. Observed and fitted acceptance by occupational match

Notes: Cells distinguish same versus different candidate and vacancy ROME codes within each tightness quartile. Orange points show observed acceptance; bars are pointwise 95 percent caseworker-clustered intervals. Blue triangles show saved logit predictions averaged within the identical cells. Predictions include the original compatibility proxy, controls and fixed effects. No fitted-probability confidence bands are shown; all comparisons are in sample.

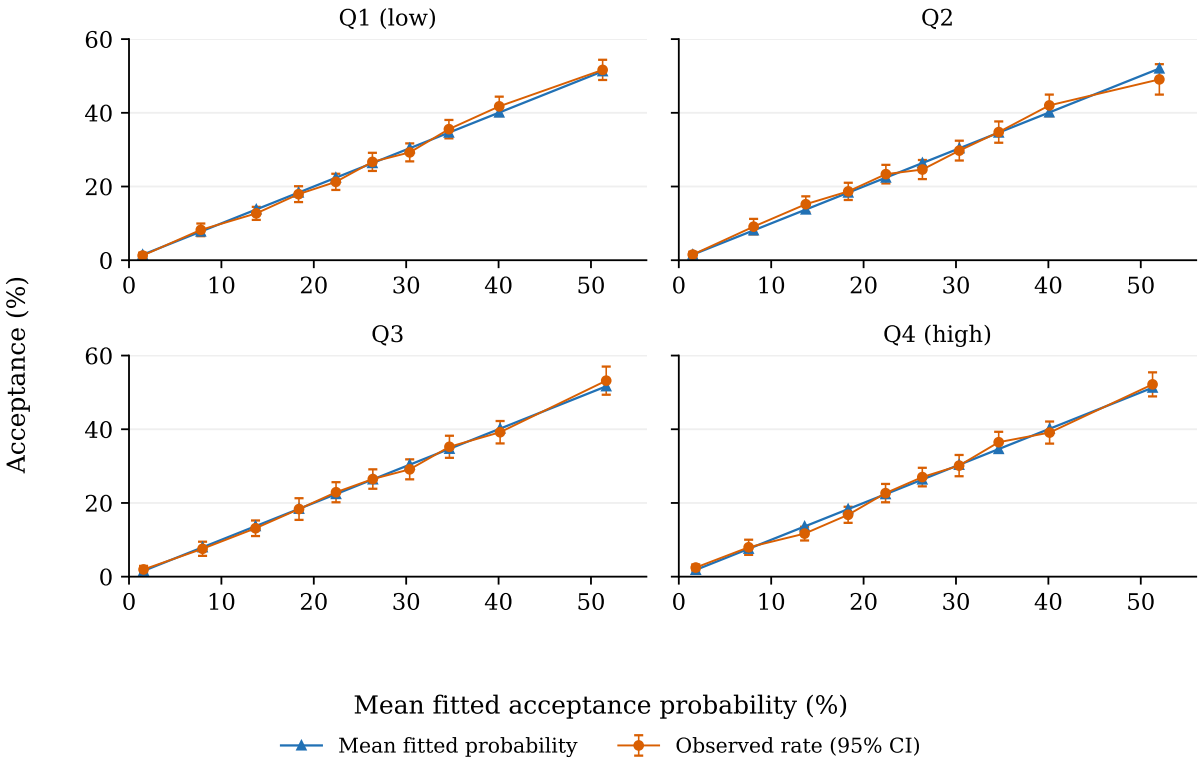


FIGURE A.3. In-sample calibration by market tightness

Notes: Bins use pooled predicted-probability decile cutoffs. The blue line is the equality benchmark because both coordinates are the cell mean fitted probability. Orange points show observed acceptance; bars are pointwise 95 percent caseworker-clustered intervals. Blue triangles show saved logit predictions averaged within the identical cells. Predictions include the original compatibility proxy, controls and fixed effects. No fitted-probability confidence bands are shown; all comparisons are in sample.

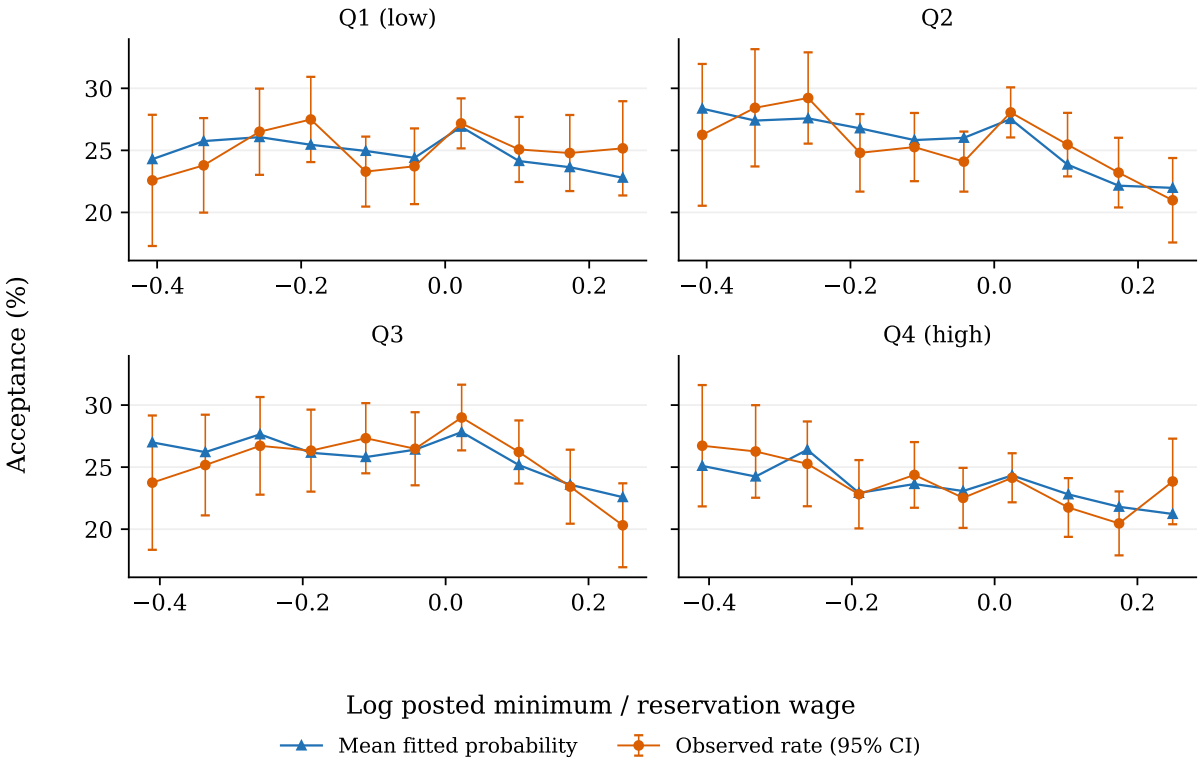


FIGURE A.4. Observed and fitted acceptance across wage-gap bins

Notes: Ten equal-width bins use the intersection of the quartile-specific central ranges. The horizontal coordinate is the mean observed log posted-minimum-to-reservation-wage ratio in each cell. Orange points show observed acceptance; bars are pointwise 95 percent caseworker-clustered intervals. Blue triangles show saved logit predictions averaged within the identical cells. Predictions include the original compatibility proxy, controls and fixed effects. No fitted-probability confidence bands are shown; all comparisons are in sample.

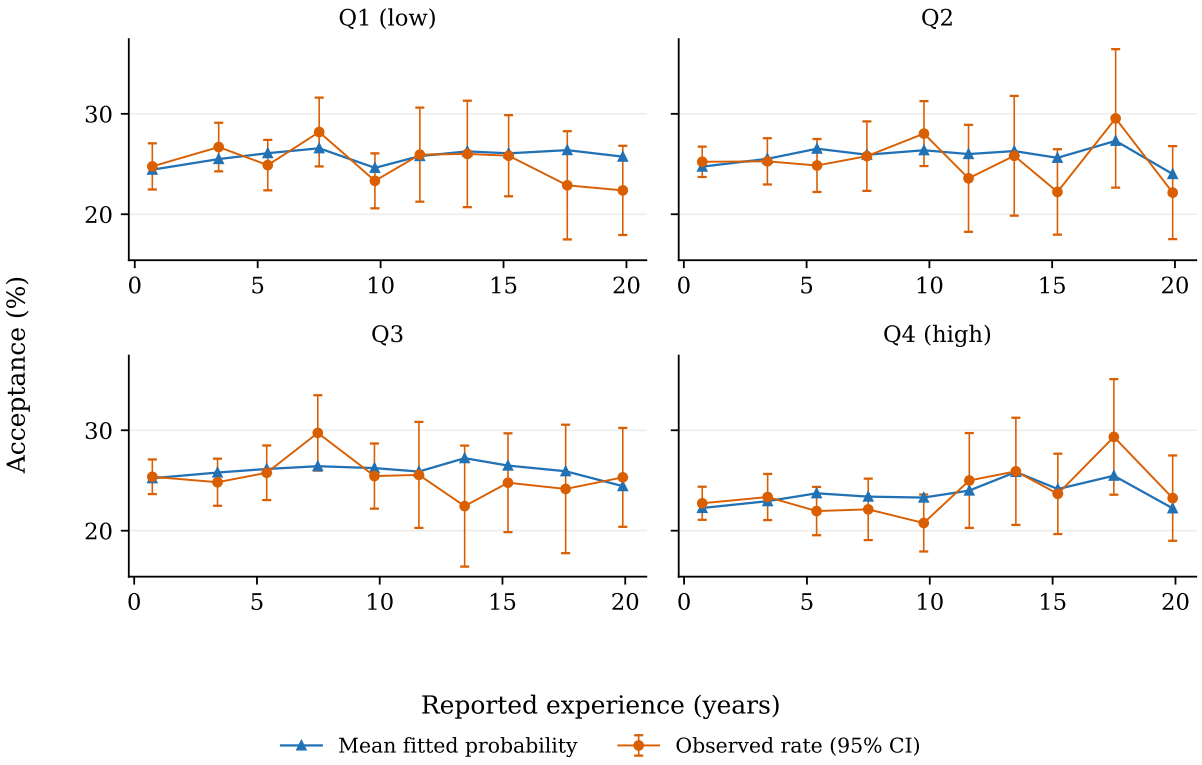


FIGURE A.5. Observed and fitted acceptance across experience bins

Notes: Ten equal-width bins divide the common central range of zero to twenty reported years of experience. Orange points show observed acceptance; bars are pointwise 95 percent caseworker-clustered intervals. Blue triangles show saved logit predictions averaged within the identical cells. Predictions include the original compatibility proxy, controls and fixed effects. No fitted-probability confidence bands are shown; all comparisons are in sample.

A.7. Supplementary heterogeneity dimensions

Firm size. Table A.27 reports the effects of an additional shortlisted candidate by firm size. For firms with one to nine employees, the estimated increases are 2.98 percentage points for contract proposals and 0.76 percentage points for hiring within six months. The larger-firm estimates are much less precise, with first-stage F-statistics of 4.2 or below. Their negative point estimates therefore do not establish adverse effects, and the results do not provide a reliable ranking across firm sizes.

TABLE A.27. Effects of shortlist size on proposals and hiring by firm size

	Full sample	1–9	10–49	50–249	250+	Unknown or other
Panel A: Contract proposal within three months						
OLS effect (pp)	0.060***	0.027**	0.147***	0.158***	0.197**	0.094***
	(0.006)	(0.009)	(0.024)	(0.040)	(0.076)	(0.024)
OLS Holm-adjusted p	–	0.008	< 0.001	< 0.001	0.010	< 0.001
2SLS effect (pp)	3.120***	2.983***	2.114	2.919	−14.677	3.660*
	(0.33)	(0.467)	(1.719)	(7.459)	(110.386)	(1.850)
2SLS Holm-adjusted p	–	< 0.001	0.656	1.000	1.000	0.192
Outcome mean (%)	29.9	30.58	30.31	30.30	33.15	22.39
Equal 2SLS effects, p			0.951			
Panel B: Any hire within six months						
OLS effect (pp)	0.041***	0.017*	0.106***	0.164***	0.163***	0.073***
	(0.006)	(0.007)	(0.019)	(0.037)	(0.047)	(0.017)
OLS Holm-adjusted p	–	0.015	< 0.001	< 0.001	0.001	< 0.001
2SLS effect (pp)	0.800***	0.758**	0.520	−0.980	−13.035	−0.054
	(0.200)	(0.232)	(0.943)	(3.905)	(93.829)	(0.548)
2SLS Holm-adjusted p	–	0.005	1.000	1.000	1.000	1.000
Outcome mean (%)	12.3	13.01	12.16	11.48	9.16	6.18
Equal 2SLS effects, p			0.957			
Observations	809,675	633,328	77,248	19,937	8,505	70,657
First-stage coefficient	0.487	0.466	0.489	0.386	−0.285	0.507
First-stage Wald F	92.4	72.9	4.2	0.4	0.0	8.8

Notes: Outcome effects and standard errors are in percentage points per additional shortlisted candidate; means are in percent. Each column reports separate regressions with office-by-cohort and posting-month fixed effects, without occupation fixed effects. The instrument is predetermined caseworker leniency. Standard errors are clustered by caseworker. Stars use unadjusted p-values: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Holm adjustment is within each dimension, outcome and estimator across subgroups, excluding the pooled column. Equality tests use stacked IV models with group-specific fixed effects and caseworker clustering, excluding unknown groups. First-stage coefficients are in candidates per standard deviation of leniency; F is the squared clustered excluded-instrument t-statistic. Firm-size groups refer to numbers of employees; unknown or other categories are reported separately. The first stage is weak in the larger size groups, limiting the interpretation of their 2SLS estimates and the equality tests.

Experience requirements. Table A.28 compares vacancies open to beginners with those requiring experience. An additional shortlisted candidate raises the probability of a contract proposal by 3.94 percentage points for beginner positions and 2.67 percentage points for experienced positions. The corresponding hiring effects are 1.03 and 0.69 percentage points. These point estimates suggest larger gains for beginner positions, although the tests of equal effects do not reject at the five-percent level for either proposals ($p = 0.073$) or hiring ($p = 0.369$).

TABLE A.28. Effects of shortlist size on proposals and hiring by experience requirement

	Beginner positions		Experienced positions
Panel A: Contract proposal within three months			
OLS effect (pp)	0.021*		0.101***
	(0.010)		(0.012)
OLS Holm-adjusted p	0.062		< 0.001
2SLS effect (pp)	3.939***		2.673***
	(0.733)		(0.455)
2SLS Holm-adjusted p	< 0.001		< 0.001
Outcome mean (%)	33.59		27.13
Equal 2SLS effects, p		0.073	
Panel B: Any hire within six months			
OLS effect (pp)	0.033***		0.048***
	(0.009)		(0.009)
OLS Holm-adjusted p	< 0.001		< 0.001
2SLS effect (pp)	1.026**		0.691**
	(0.353)		(0.219)
2SLS Holm-adjusted p	0.007		0.005
Outcome mean (%)	14.62		10.53
Equal 2SLS effects, p		0.369	
Observations	343,034		466,266
First-stage coefficient	0.459		0.494
First-stage Wald F	49.5		70.3

Notes: Outcome effects and standard errors are in percentage points per additional shortlisted candidate; means are in percent. Each column reports separate regressions with office-by-cohort and posting-month fixed effects, without occupation fixed effects. The instrument is predetermined caseworker leniency. Standard errors are clustered by caseworker. Stars use unadjusted p-values: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Holm adjustment is within each dimension, outcome and estimator across subgroups, excluding the pooled column. Equality tests use stacked IV models with group-specific fixed effects and caseworker clustering, excluding unknown groups. First-stage coefficients are in candidates per standard deviation of leniency; F is the squared clustered excluded-instrument t-statistic. Beginner and experienced positions correspond to the D and E prefixes of the recorded experience-requirement code. The 375 observations with other or unknown codes are omitted. Holm-adjusted p-values retain the original three-group adjustment, including the omitted group; equality tests compare the two displayed groups.

Posted salary. Table A.29 reports larger proposal gains among vacancies offering lower salaries. The estimated effect of one additional candidate declines from 5.60 percentage points in the lowest salary quartile to 1.28 percentage points in the highest. The hiring point estimates follow the same ordering, declining from 1.21 to 0.30 percentage points, but are less precise. The tests of equal effects across the four salary quartiles yield $p = 0.072$ for proposals and $p = 0.647$ for hiring. The proposal estimates therefore suggest greater benefits for lower-salary vacancies, without establishing differences across quartiles at the five-percent level or identifying the mechanism behind this pattern.

TABLE A.29. Effects of shortlist size on proposals and hiring by posted salary

	Full sample	Q1 (lowest)	Q2	Q3	Q4 (highest)	Unknown
Panel A: Contract proposal within three months						
OLS effect (pp)	0.060***	0.009	−0.014	0.023	0.040*	0.118***
	(0.006)	(0.018)	(0.020)	(0.015)	(0.016)	(0.018)
OLS Holm-adjusted p	–	0.961	0.961	0.402	0.042	< 0.001
2SLS effect (pp)	3.120***	5.600***	2.258***	1.707**	1.280*	5.717***
	(0.33)	(1.621)	(0.613)	(0.572)	(0.554)	(1.259)
2SLS Holm-adjusted p	–	0.002	< 0.001	0.006	0.021	< 0.001
Outcome mean (%)	29.9	42.01	34.61	24.31	21.31	27.17
Equal 2SLS effects, p			0.072			
Panel B: Any hire within six months						
OLS effect (pp)	0.041***	0.006	−0.008	0.033**	0.036*	0.059***
	(0.006)	(0.015)	(0.015)	(0.011)	(0.016)	(0.018)
OLS Holm-adjusted p	–	1.000	1.000	0.013	0.083	0.004
2SLS effect (pp)	0.800***	1.210	0.640	0.372	0.301	1.088*
	(0.200)	(0.692)	(0.388)	(0.311)	(0.310)	(0.531)
2SLS Holm-adjusted p	–	0.322	0.322	0.464	0.464	0.202
Outcome mean (%)	12.3	19.86	14.83	9.51	7.25	10.03
Equal 2SLS effects, p			0.647			
Observations	809,675	160,219	151,366	155,145	149,981	192,964
First-stage coefficient	0.487	0.401	0.656	0.631	0.550	0.350
First-stage Wald F	92.4	15.6	38.4	38.3	31.4	38.1

Notes: Outcome effects and standard errors are in percentage points per additional shortlisted candidate; means are in percent. Each column reports separate regressions with office-by-cohort and posting-month fixed effects, without occupation fixed effects. The instrument is predetermined caseworker leniency. Standard errors are clustered by caseworker. Stars use unadjusted p-values: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Holm adjustment is within each dimension, outcome and estimator across subgroups, excluding the pooled column. First-stage coefficients are in candidates per standard deviation of leniency; F is the squared clustered excluded-instrument t-statistic. Q1 and Q4 denote the lowest and highest posted-salary quartiles.

Prior recruitment. Table A.30 compares establishments with no observed hires in the year before their first preselection vacancy with those that hired during that period. The estimated proposal effect is 4.18 percentage points in the former group and 2.87 percentage points in the latter; the corresponding hiring effects are 1.23 and 0.73 percentage points. The point estimates suggest larger gains among establishments without recent recruitment, although the differences are not significant at the five-percent level ($p = 0.062$ for proposals and $p = 0.171$ for hiring). This comparison concerns recent hiring history and does not identify establishments making their first-ever hire.

TABLE A.30. Effects of shortlist size on proposals and hiring by prior recruitment

	Full sample	No hires in prior 12 months	Hired in prior 12 months
Panel A: Contract proposal within three months			
OLS effect (pp)	0.060***	0.074***	0.058***
	(0.006)	(0.014)	(0.009)
OLS Holm-adjusted p	–	< 0.001	< 0.001
2SLS effect (pp)	3.120***	4.181***	2.867***
	(0.33)	(0.674)	(0.506)
2SLS Holm-adjusted p	–	< 0.001	< 0.001
Outcome mean (%)	29.9	31.05	29.27
Equal 2SLS effects, p		0.062	
Panel B: Any hire within six months			
OLS effect (pp)	0.041***	0.037***	0.050***
	(0.006)	(0.010)	(0.007)
OLS Holm-adjusted p	–	< 0.001	< 0.001
2SLS effect (pp)	0.800***	1.230***	0.728**
	(0.200)	(0.322)	(0.244)
2SLS Holm-adjusted p	–	< 0.001	0.003
Outcome mean (%)	12.3	11.18	12.79
Equal 2SLS effects, p		0.171	
Observations	809,675	269,104	540,571
First-stage coefficient	0.487	0.537	0.481
First-stage Wald F	92.4	64.3	61.0

Notes: Outcome effects and standard errors are in percentage points per additional shortlisted candidate; means are in percent. Each column reports separate regressions with office-by-cohort and posting-month fixed effects, without occupation fixed effects. The instrument is predetermined caseworker leniency. Standard errors are clustered by caseworker. Stars use unadjusted p-values: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Holm adjustment is within each dimension, outcome and estimator across subgroups, excluding the pooled column.

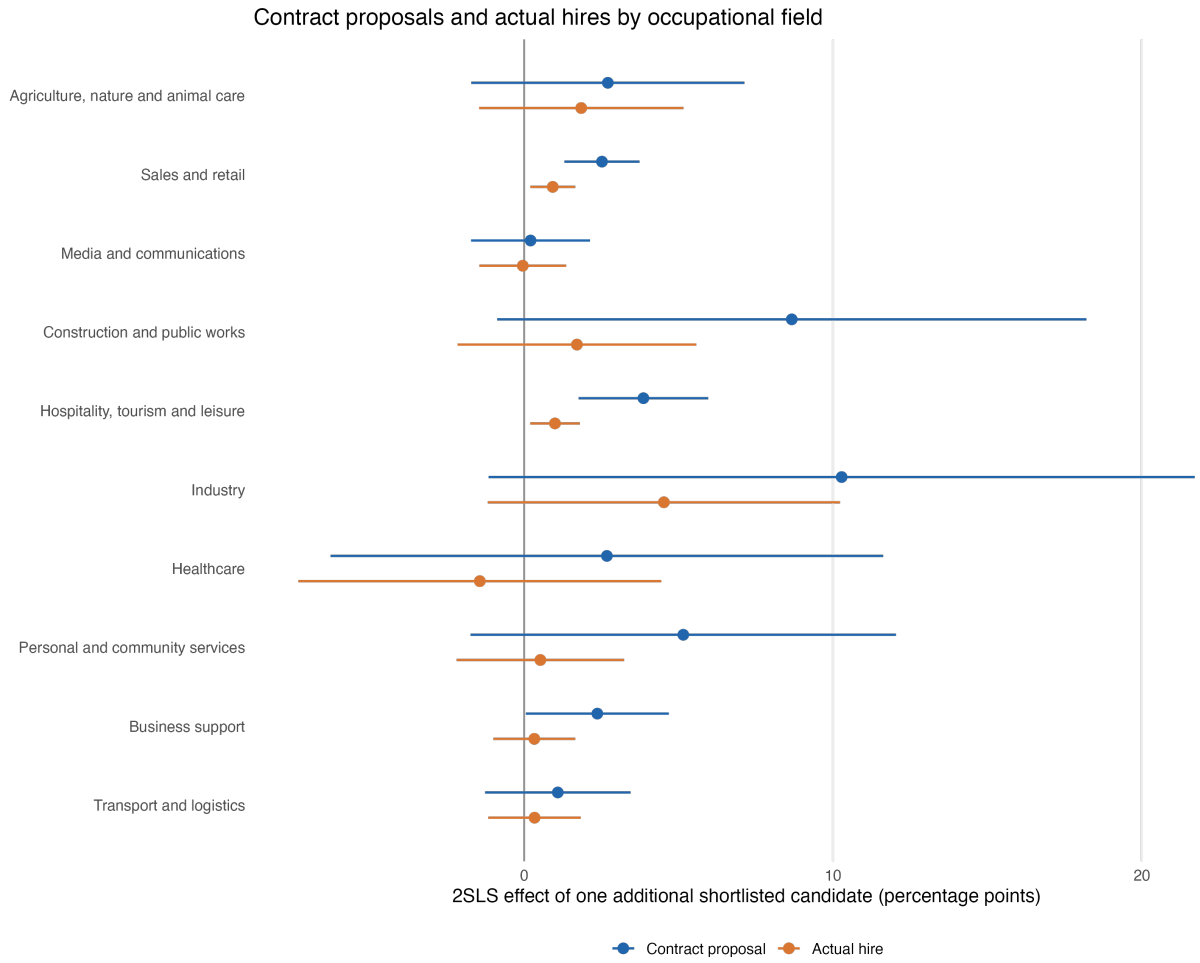


FIGURE A.6. Effects of shortlist size on proposals and hiring by occupational field

Notes: The figure reports separate 2SLS estimates by occupational field, defined using broad ROME groups. Blue points show effects on the probability of a contract proposal within three months; orange points show effects on the probability of a realized hire within six months. Effects refer to one additional shortlisted candidate and are expressed in percentage points. Shortlist size is instrumented with predetermined caseworker leniency. All specifications include office-by-cohort and posting-month fixed effects. Horizontal bars show 95 percent confidence intervals based on standard errors clustered at the caseworker level.

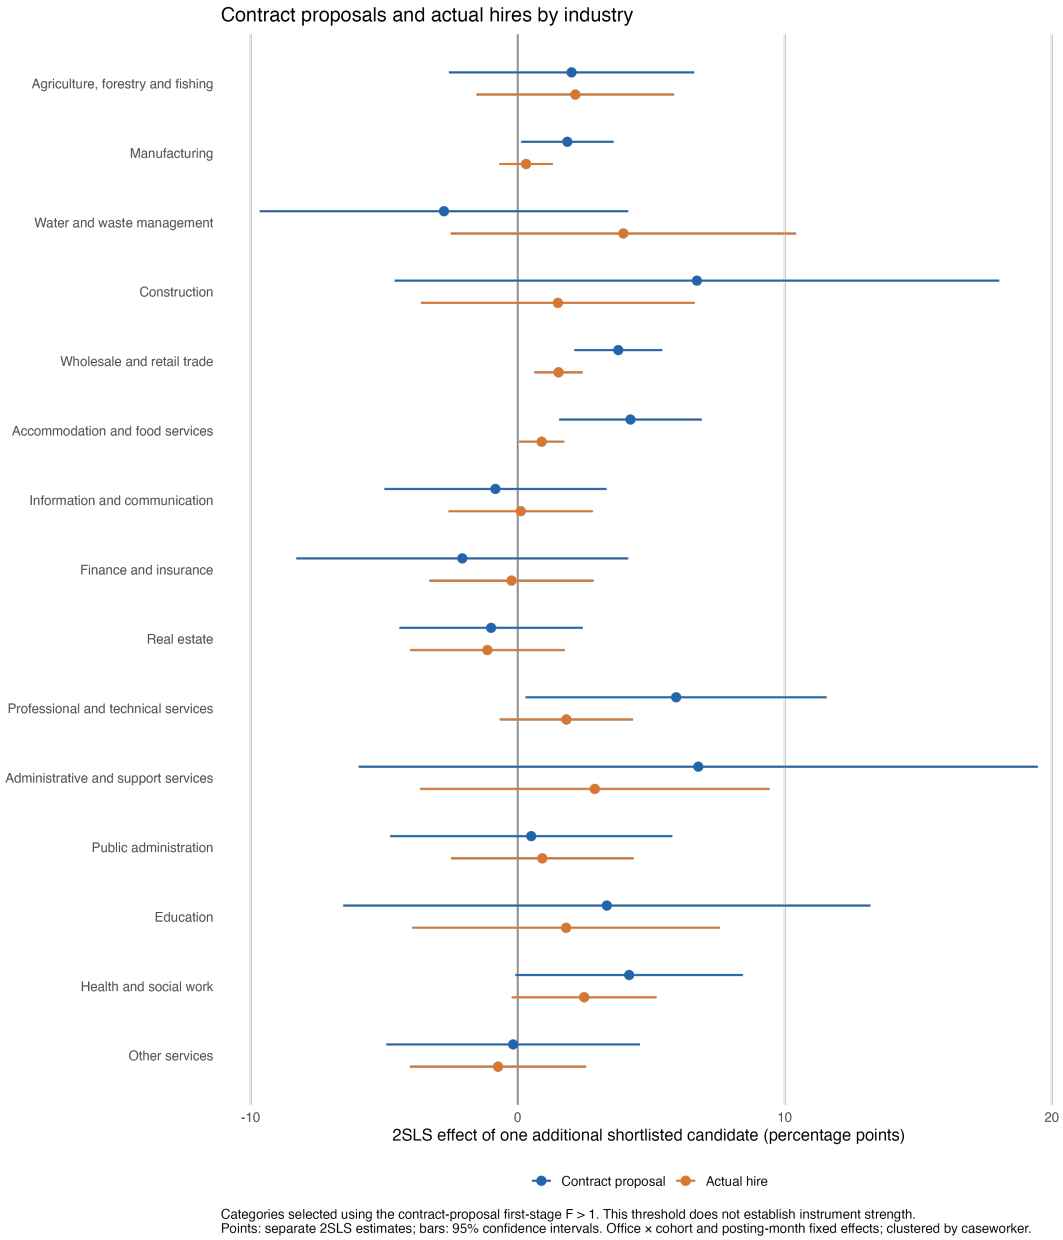


FIGURE A.7. Effects of shortlist size on proposals and hiring by industry

Notes: The figure reports separate 2SLS estimates by industry, defined using NAF sections. Blue points show effects on the probability of a contract proposal within three months; orange points show effects on the probability of a realized hire within six months. Effects refer to one additional shortlisted candidate and are expressed in percentage points. Shortlist size is instrumented with predetermined caseworker leniency. All specifications include office-by-cohort and posting-month fixed effects. Horizontal bars show 95 percent confidence intervals based on standard errors clustered at the caseworker level.

B. Technical details of the conceptual framework and IV interpretation

This appendix sets out the assumptions and arguments supporting the conceptual framework. The search rule follows Weitzman (1979); the IV interpretation follows Angrist and Imbens (1995). The framework describes one vacancy at prevailing market conditions. It does not identify inspection costs from the data or model displacement across employers. Results that require inspection costs to remain unchanged as the shortlist expands are explicitly restricted to $\lambda = 0$.

B.1. Inspection and the stopping rule

For vacancy j, candidate i provides surplus $s_{ij} = m_{ij} + \varepsilon_{ij}$. The firm observes expected fit m_{ij} before inspection and learns the residual ε_{ij} by inspecting. Conditional on the observed fit signals, residuals are independent across candidates, have a common continuous distribution H, and satisfy $\mathbb{E}[\varepsilon_{ij}] = 0$ and $\mathbb{E}[|\varepsilon_{ij}|] < \infty$. They are independent of caseworker assessment errors. The shortlist is finite, previously inspected candidates remain available, and the firm can always retain its outside option τ_j^0. Initially, every candidate is assumed willing to accept.

For a given shortlist of size n, each inspection costs

$$\kappa_j(n) = \kappa_{0j} + \lambda n / K_j, \qquad \kappa_{0j} > 0, \quad K_j > 0, \quad \lambda \geq 0.$$

Here K_j is recruitment capacity and λ measures the effect of shortlist workload on the cost per inspection. This cost remains fixed during search over a given shortlist; it depends on the number transmitted, not the number still uninspected. The firm's objective is expected final surplus less the cost of the inspections undertaken.

The reservation index $z_{ij}(n)$ is defined by

$$\mathbb{E}\big[\max\{m_{ij} + \varepsilon - z_{ij}(n), 0\}\big] = \kappa_j(n), \tag{13}$$

where ε has distribution H. The index is the value of an available alternative that makes inspecting candidate i just worthwhile. Since inspection costs and the residual distribution are common across candidates, higher fit implies a higher index. The optimal rule therefore inspects candidates in decreasing order of fit. Search continues while the highest remaining index exceeds b_j, the best surplus currently available, and stops otherwise. Initially $b_j = \tau_j^0$; after inspection it also incorporates the surpluses of evaluated candidates. The firm stops at equality and chooses the outside option in a tie with it.

For the highest-fit remaining candidate, this rule is equivalent to comparing her expected improvement with the inspection cost:

$$\mathbb{E}\big[\max\{m_{ij} + \varepsilon - b_j, 0\}\big] > \kappa_j(n). \tag{14}$$

The left-hand side increases with fit and decreases with the best available surplus. Its

equality with cost defines the inspection frontier $\underline{m}_j$. The frontier depends on b_j and n, although these arguments are suppressed for readability. It rises when costs increase or a better candidate is found. A candidate below the initial frontier is never inspected, since b_j cannot fall during search. A candidate above it may remain uninspected if an earlier evaluation makes further search unnecessary.

B.2. Transmission cutoffs and shortlist composition

Caseworker c observes $\widehat{m}_{ijc} = m_{ij} + \sigma_c \eta_{ijc}$ and transmits a candidate when this assessment exceeds θ_c. The errors η_{ijc} are independent standard normal draws, independent of fit, and $\sigma_c > 0$ measures assessment noise. The probability of transmission at fit m is

$$\mathbb{P}(D_{ijc} = 1 \mid m_{ij} = m, c) = \Phi\left(\frac{m - \theta_c}{\sigma_c}\right), \tag{15}$$

where Φ is the standard normal distribution function. Lowering the cutoff increases transmission at every fit level.

The cutoff comparison holds the applicant pool, precision and realized assessments fixed. If $S(\theta)$ denotes the transmitted set at cutoff θ, then for $\theta' < \theta$,

$$S(\theta') = S(\theta) \cup \{i : \theta' \leq \widehat{m}_{ijc} < \theta\}. \tag{16}$$

No previously transmitted candidate is removed. Consequently, the number above any fixed fit threshold cannot fall. This threshold is a descriptive measure of suitability and need not equal the firm's inspection frontier.

The implication for average fit concerns expectations. For two fit values $m_2 > m_1$, the ratio of the normal signal densities at assessment v is proportional to $\exp\{v(m_2 - m_1)/\sigma_c^2\}$, which increases with v. Higher assessments therefore imply a higher posterior distribution of fit. For applicants drawn from a common fit distribution with finite absolute mean, this gives

$$\begin{aligned} \mathbb{E}[m \mid \widehat{m} \geq \theta'] &\leq \mathbb{E}[m \mid \widehat{m} \geq \theta], \\ \mathbb{P}(m \geq a \mid \widehat{m} \geq \theta') &\leq \mathbb{P}(m \geq a \mid \widehat{m} \geq \theta), \end{aligned} \tag{17}$$

where a is any fixed suitability threshold and transmission probabilities are positive. For a finite pool of independent draws, these inequalities also describe expected shortlist means and shares conditional on a nonempty shortlist. They need not hold for every realized shortlist average. Expansion can thus lower average fit and the share of suitable candidates while increasing their number.

Precision concerns the ranking itself. For two candidates with $m_{ij} > m_{kj}$, the probability of ranking them correctly is

$$\mathbb{P}(\widehat{m}_{ijc} > \widehat{m}_{kjc} \mid m_{ij}, m_{kj}) = \Phi\left(\frac{m_{ij} - m_{kj}}{\sqrt{2}\sigma_c}\right).$$

This probability rises as assessment noise falls. The nested-set comparison above is therefore a change in the cutoff along a fixed ranking, not a claim that every pair of caseworkers produces nested shortlists. Observable sorting measures can inform the distinction between precision and leniency, but do not establish that all other dimensions of screening are held fixed.

B.3. Expansion with and without a processing burden

When $\lambda = 0$, expanding a shortlist without removing existing candidates leaves the original search strategy feasible at the same cost. The firm can ignore every addition. Its optimized expected net surplus therefore cannot fall.

Under the stopping rule, the probability of selecting someone from the shortlist is also weakly increasing in this benchmark. To see why, consider the candidates whose reservation indices exceed the outside option. The firm selects from the shortlist exactly when at least one of these candidates has realized surplus above the outside option. If none does, search exhausts these candidates and the firm retains its outside option. If one does, either the firm reaches that candidate or it has already found another acceptable option. With unchanged inspection costs, expansion preserves this set and can add to it. This argument holds the existing candidates' fit and residual realizations fixed.

When $\lambda > 0$, an additional transmitted candidate raises the cost of every inspection by λ/K_j. The firm can no longer reproduce its original search strategy at the same cost. Existing candidates' reservation indices fall, and some may cease to be worth inspecting. Additional opportunities and higher costs then work in opposite directions, so neither recruitment nor net surplus has an unconditionally positive response.

Even in the fixed-cost benchmark, the marginal recruitment gain need not decline with every successive addition: caseworker assessments are noisy, so an applicant admitted later can be more promising to the firm than one admitted earlier. Differences in outside options, inspection costs and applicant pools also prevent a common marginal-effect profile across vacancies. The framework does not impose unchanged selected-worker characteristics or inspection effort.

B.4. Acceptance and the inspection frontier

Candidate i is willing to accept with probability $\alpha(m_{ij})$. Conditional on fit, willingness to accept is independent of residual surplus, and the pairs consisting of willingness and residual surplus are independent across candidates. Acceptance is weakly increasing in fit. These are maintained assumptions; employer-valued characteristics need not always make a job more attractive to a worker.

Inspection reveals both surplus and willingness to accept. Willing candidates remain available during search, and b_j now records the best attainable surplus, including the outside option. An unwilling candidate cannot improve this option. Conditional independence

therefore gives the expected gain from inspection as

$$\alpha(m)\,\mathbb{E}\big[\max\{m+\varepsilon-b_j,0\}\big]. \tag{18}$$

The factorization would not generally hold if willingness and residual surplus were dependent conditional on fit.

The reservation-index argument extends by treating an unwilling candidate as providing no improvement over the outside option. For indices above that option, the expected gain in equation (18) equals inspection cost at the reservation index. Since both acceptance and expected improvement are weakly increasing in fit, candidates worth inspecting can still be ordered by decreasing fit. The acceptance-adjusted frontier $\underline{m}_j^h$ satisfies

$$\alpha(\underline{m}_j^h)\,\mathbb{E}\Big[\max\{\underline{m}_j^h+\varepsilon-b_j,0\}\Big]=\kappa_j(n) \tag{19}$$

when a continuous interior crossing exists. More generally, the worthwhile candidates form the upper set for which expected gain strictly exceeds cost; this set can be empty. Candidates with zero acceptance probability are never worth inspecting.

A pointwise reduction in acceptance shrinks the set worth inspecting, holding b_j and costs fixed. At an interior crossing it weakly raises the frontier. An unwilling candidate, however, does not improve b_j, so further candidates may remain worth evaluating. With $\lambda = 0$, nested expansion again weakly increases net surplus and the probability of a completed match, by the same argument as above applied to willing candidates.

This extension is static: it does not specify dates of formal proposals, refusals or responses. It therefore does not deliver a duration equation or identify inspection cost from elapsed time. Nor does the ratio of vacancy-level effects on hires and proposals identify an individual acceptance probability: additional hires can occur at vacancies that would already have made a proposal to someone else, and expansion can change which candidates receive proposals.

B.5. Tightness and the two sides of the match

Labor-market tightness ϑ is the number of vacancies relative to job seekers. The comparison allows the firm's outside option $\tau_j^0(\vartheta)$ and candidate acceptance $\alpha(m;\vartheta)$ to fall as tightness increases. The shortlist, vacancy terms, inspection costs and residual distribution remain fixed. At the initial inspection frontier, the expected gain is

$$\alpha(m;\vartheta)\,\mathbb{E}\Big[\max\{m+\varepsilon-\tau_j^0(\vartheta),0\}\Big].$$

A lower outside option increases this gain, while a lower acceptance probability reduces it. Where the functions are differentiable, the change with tightness at fixed fit is

$$\begin{aligned}\frac{\partial \alpha(m;\vartheta)}{\partial \vartheta}\mathbb{E}\Big[\max\{m+\varepsilon-\tau_j^0(\vartheta),0\}\Big] \\ -\alpha(m;\vartheta)\left[1-H\Big(\tau_j^0(\vartheta)-m\Big)\right]\frac{d\tau_j^0(\vartheta)}{d\vartheta}.\end{aligned} \tag{20}$$

The first term is nonpositive under lower acceptance; the second is nonnegative under a weaker firm outside option. At a regular interior crossing, with expected gain strictly increasing in fit, the frontier falls if the sum is positive and rises if it is negative. These are the two shifts illustrated in Figure 1.

The movement of the frontier does not by itself rank the marginal hiring effect across markets. The value of one more candidate also depends on whether existing candidates already suffice. The framework therefore allows candidate scarcity to increase firms' interest in additional applicants while reducing their willingness to accept, without imposing a larger hiring response in tighter markets. These are partial-equilibrium comparisons for the receiving firm, not statements about net employment creation.

B.6. Interpretation of the IV coefficient

This subsection concerns identification and does not require firms to follow the search rule. Here i indexes a vacancy, X_i is its shortlist size, Z_i is predetermined caseworker leniency, Y_i is a recruitment outcome, and C_i contains predetermined controls and assignment-cell identifiers. Potential shortlist size is $X_i(z)$ and the potential outcome along the relevant expansion path is $Y_i(x)$.

The interpretation requires conditional independence of Z_i from potential treatment and outcome schedules given C_i, exclusion, consistency, a positive first stage and finite required moments. It also requires monotonicity: $X_i(z') \geq X_i(z)$ for $z' > z$ for each vacancy. The exclusion restriction permits the composition changes associated with expansion, but rules out additional effects of caseworker assignment at a given shortlist size. In particular, changing candidate identities, support or transmission timing independently of expansion would violate this restriction. A count treatment does not make different shortlists of the same size equivalent by construction; that equivalence along the instrument-induced path is an identifying assumption.

With conditional-mean adjustment, the population IV estimand is

$$\beta_C^{\mathrm{IV}} = \frac{\mathbb{E}[(Z_i - \mathbb{E}[Z_i \mid C_i])Y_i]}{\mathbb{E}[(Z_i - \mathbb{E}[Z_i \mid C_i])X_i]}. \tag{21}$$

Subtracting the conditional means of Y_i and X_i gives the same ratio. Under the assumptions

above, the average-causal-response interpretation of Angrist and Imbens (1995) yields

$$\beta_C^{\mathrm{IV}} = \sum_{x \geq 1} \omega_x \mu_x, \qquad \omega_x \geq 0, \qquad \sum_{x \geq 1} \omega_x = 1. \tag{22}$$

Here μ_x is an IV-weighted average of $Y_i(x) - Y_i(x-1)$ among vacancies whose shortlist can cross margin x as leniency changes; ω_x is the weight on that margin. This is the interpretation stated in Proposition 1.

The weighting follows by writing each realized outcome as its value at zero candidates plus the increments for the margins crossed:

$$Y_i(X_i) = Y_i(0) + \sum_{x \geq 1} \{Y_i(x) - Y_i(x-1)\} \mathbf{1}\{X_i \geq x\}.$$

Conditional independence removes the baseline term from the IV numerator. For each fixed potential treatment schedule, monotonicity makes $\mathbf{1}\{X_i(z) \geq x\}$ nondecreasing in z, so its conditional covariance with Z_i is nonnegative. These covariances weight the causal increments in the numerator and sum to the first-stage denominator. Normalizing by that positive denominator gives equation (22). The weights depend on the instrument distribution and the margins it shifts, not simply on how frequently each shortlist size occurs.

Connection to the estimated specification. Linear 2SLS removes linear projections on the included controls and fixed effects, rather than unrestricted conditional means. The interpretation above applies when the population linear projection of Z_i on those controls equals $\mathbb{E}[Z_i \mid C_i]$. Fully saturated adjustment is sufficient. Additive office-by-cohort and posting-month fixed effects do not alone guarantee this condition. Without it, the nonnegative-weight interpretation requires further justification.

Scope of the coefficient. The coefficient averages the effects of the additional candidates induced by leniency, including accompanying composition changes. It is not the effect of adding a candidate drawn at random while holding the fit distribution constant. Estimates within predicted-shortlist groups have their own instrument-induced weights and compare different groups of vacancies; they do not trace successive additions to the same vacancy. Likewise, a coefficient conditional on observing a proposal or hire generally combines changes in outcomes with changes in the selected sample. Balance, sorting and monotonicity checks inform the identifying assumptions but do not prove them. Neither the search framework nor these checks establish the exclusion restriction by themselves.

C. Numerical illustrations and recruitment counterfactuals

This appendix illustrates the mechanisms in Section 2 and relates the empirical estimates to changes in the allocation of shortlisted candidates. The first three subsections use illustrative parameters, rather than estimates of the model's primitives. They distinguish the caseworker's transmission cutoff, which determines who reaches the firm, from the firm's inspection frontier, which determines who is worth evaluating. The final subsection presents local recruitment calculations using the estimates reported in the paper.

Common assumptions and notation. For a given vacancy, candidate i provides surplus $s_i = m_i + \varepsilon_i$, where m_i is expected fit, observed by the firm before inspection, and ε_i is revealed through inspection. The vacancy subscript is suppressed throughout the numerical examples. Residuals are independent standard-normal draws, independent of caseworker assessment errors. The functions ϕ and Φ denote the standard-normal density and distribution function. The outside option is denoted by τ^0.

The examples use the fixed-cost benchmark, $\lambda = 0$, so an additional transmitted candidate does not increase the cost of inspecting candidates already available. Each inspection costs $\kappa = \kappa_0 > 0$. Define

$$\psi(u) = \mathbb{E}[(\varepsilon - u)^+] = \phi(u) - u[1 - \Phi(u)], \qquad (x)^+ = \max\{x, 0\}, \qquad \kappa = \psi(1) = 0.08331547. \tag{23}$$

Here u is a residual-surplus cutoff. Because ψ is strictly decreasing, its inverse is well defined on positive expected gains. With certain acceptance, a candidate's reservation index is $z_i = m_i + \rho$, where $\rho = \psi^{-1}(\kappa) = 1$. At the start of search, a candidate is worth inspecting if $z_i > \tau^0$, or equivalently if $m_i > \underline{m} = \tau^0 - \rho$. After an inspection, the outside option in this comparison is replaced by the best surplus available. The firm stops at equality, as in Appendix B.1.

C.1. Adding candidates from further down the ranking

Selection probabilities and inspection effort. Let S_n be the set of n transmitted candidates and let $A_n = \{i \in S_n : m_i + \rho > \tau^0\}$ be the subset initially worth inspecting. Acceptance is certain in this example. Under the sequential inspection rule, the firm selects a candidate from the shortlist exactly when at least one member of A_n has surplus above the outside option. Independence across residuals therefore gives the selection probability

$$P(S_n) = 1 - \prod_{i \in A_n} \Phi(\tau^0 - m_i). \tag{24}$$

An empty product equals one, so $P(S_n) = 0$ when no candidate is worth inspecting. If expansion preserves every existing candidate and adds candidate $n + 1$ with fit $m_{[n+1]}$, the

marginal selection effect is

$$\begin{aligned} g(n) &= P(S_{n+1}) - P(S_n) \\ &= \mathbf{1}\{m_{[n+1]} + \rho > \tau^0\}[1 - \Phi(\tau^0 - m_{[n+1]})] \prod_{i \in A_n} \Phi(\tau^0 - m_i). \end{aligned} \tag{25}$$

The added candidate matters only if worth inspecting and if the existing options would not already yield a selection. This is a probability statement about the final choice, not an assumption that the firm always inspects the added candidate last.

Expected inspection effort follows from the same stopping rule. Order the $L = |A_n|$ initially worthwhile candidates by fit, $m_{(1)} \geq \cdots \geq m_{(L)}$. Parentheses denote the firm's inspection ranking; square brackets denote the order in which candidates enter the shortlist as the transmission cutoff falls. The firm reaches candidate k if every previously inspected surplus is below that candidate's reservation index, $m_{(k)} + \rho$. Consequently,

$$\mathbb{E}[N(S_n)] = \sum_{k=1}^{L} \prod_{\ell<k} \Phi\big(m_{(k)} + \rho - m_{(\ell)}\big) \leq \sum_{k=1}^{L} \Phi(\rho)^{k-1} \leq \frac{1}{1 - \Phi(\rho)}, \tag{26}$$

where $N(S_n)$ is the number inspected and the sum is zero if $L = 0$. The bound uses the decreasing fit ranking and applies to this common-cost, certain-acceptance benchmark. It is not an estimate of inspection effort in the data.

Expansion ordered by fit. The outside option is $\tau^0 = 1$, and candidate k entering the shortlist has fit $m_{[k]} = 0.85 - 0.12(k - 1)$ for $k = 1, \ldots, 12$. Each addition is less promising than the previous one, so the transmission and inspection rankings coincide. The initial inspection frontier is $\underline{m} = 0$, making only candidates 1–8 worth inspecting. Equations (24)–(26) give the quantities shown in Figure C.1. Selection probability rises from 0.4404 with one candidate to 0.9407 with eight. The effect of the next candidate falls from 22.03 percentage points at $n = 1$ to 1.14 percentage points at $n = 7$ and is zero from $n = 8$ onward. Expected inspections reach 3.646, below the uniform bound of 6.303, and remain unchanged when subsequent candidates are below the frontier.

Expansion with an imperfect transmission ranking. The same primitives give a different marginal-effect profile when candidates enter in the order $m_{[1]} = 0.1$, $m_{[2]} = -0.1$ and $m_{[3]} = 1.2$. The successive additions raise selection probability by 18.41, 0.00 and 47.26 percentage points. The second candidate is below the frontier, while the third is more promising than the first and is inspected first once available. Nested expansion still cannot reduce selection probability, but the marginal effect need not decline at every step. This possibility matters when the caseworker ranks noisy assessments rather than the fit observed by the firm.

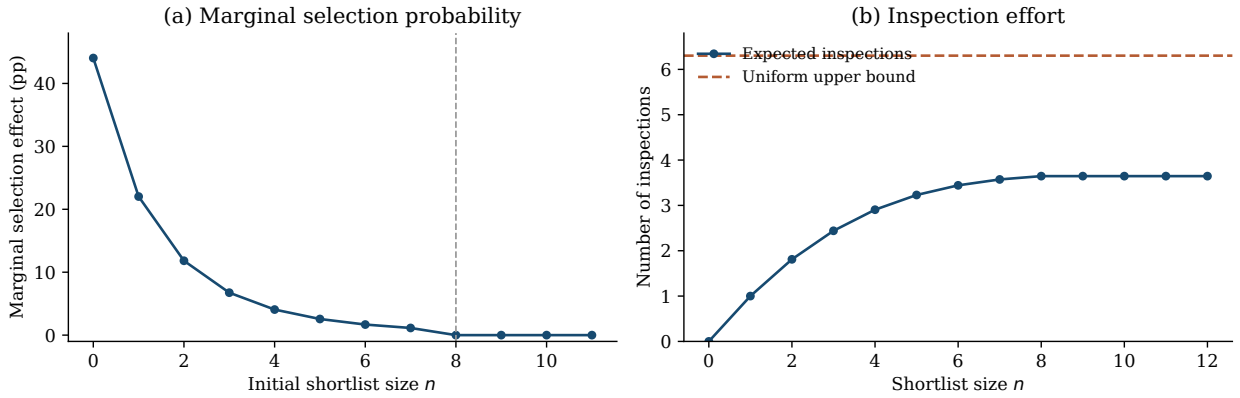


FIGURE C.1. Selection gains and inspection effort when expansion follows fit

Notes: The horizontal axis is shortlist size n. Panel (a) plots $100g(n)$, the increase in selection probability, in percentage points, from adding candidate $n+1$. The vertical line marks $n = 8$, the first size at which the next candidate is below the initial inspection frontier. Panel (b) plots expected inspections $\mathbb{E}[N(S_n)]$ and the bound $1/[1-\Phi(\rho)]$. Residual surplus is standard normal, acceptance is certain, $\lambda = 0$, $\kappa = 0.08331547$, $\rho = 1$, $\tau^0 = 1$, and $m_{[k]} = 0.85 - 0.12(k-1)$.

These examples illustrate a mechanism consistent with larger gains among vacancies with small predicted pools. They do not equate the empirical quintiles with successive additions to one vacancy. Each quintile contains different vacancies and has its own instrument-induced treatment margins, as discussed in Appendix B.6. Nor do these examples establish monotonic gains when $\lambda > 0$: if expansion raises the cost of every inspection, the set initially worth inspecting can shrink.

C.2. Leniency and screening precision

The second illustration separates changes in the transmission cutoff from changes in screening precision. The initial pool contains 100 independent applicants with fit $m \sim \mathcal{N}(0,1)$. A caseworker observes $\widehat{m} = m + \sigma\eta$, where $\eta \sim \mathcal{N}(0,1)$ is independent of fit, and transmits an applicant if $\widehat{m} \geq \theta_r$. The screening-noise standard deviation is either $\sigma = 0.5$ or $\sigma = 1.5$, with the smaller value indicating greater precision. The cutoff is chosen to give transmission probability $r \in \{0.20, 0.50\}$; thus $100r$ is the expected shortlist size, rather than a fixed realized count.

Since $\widehat{m} \sim \mathcal{N}(0, 1+\sigma^2)$, the cutoff and mean transmitted fit are

$$\theta_r = \sqrt{1+\sigma^2}\,\Phi^{-1}(1-r), \qquad \mathbb{E}[m \mid \widehat{m} \geq \theta_r] = \frac{\phi(\Phi^{-1}(1-r))}{r\sqrt{1+\sigma^2}}. \tag{27}$$

Suitability is measured using the fixed descriptive threshold $\bar{m} = 0.5$, which need not coincide with the firm's inspection frontier. The probability that a randomly drawn

applicant is both transmitted and above this threshold is

$$J(\theta_r, \sigma) = \int_{0.5}^{\infty} \phi(m) \left[1 - \Phi\left(\frac{\theta_r - m}{\sigma}\right)\right] dm. \tag{28}$$

The suitability share among transmitted applicants is $J(\theta_r, \sigma)/r$, and the expected number of suitable transmitted applicants is $100J(\theta_r, \sigma)$. Table C.1 reports these quantities for the two precision levels and transmission rates.

TABLE C.1. Composition under alternative cutoffs and screening precision

σ	Rate r	Cutoff θ_r	Mean fit	Share with $m \geq 0.5$	Expected count with $m \geq 0.5$
0.5	0.20	0.941	1.252	0.902	18.04
0.5	0.50	0.000	0.714	0.592	29.58
1.5	0.20	1.517	0.776	0.623	12.46
1.5	0.50	0.000	0.443	0.471	23.55

Notes: Fit and independent assessment errors are standard normal. The parameter σ is the standard deviation of screening noise; r is the probability of transmission; and θ_r is the cutoff yielding that probability. Mean fit and the suitability share refer to transmitted applicants. The last column is the expected number both transmitted and above $m = 0.5$ among 100 initial applicants. Shares are reported as fractions. The suitability threshold is held fixed across rows and is distinct from the transmission cutoff and inspection frontier.

At either precision level, lowering the cutoff reduces average transmitted fit and the suitability share, while increasing the expected number of suitable candidates. At a fixed transmission rate, more precise screening instead improves average fit and the suitability share. These comparisons illustrate two distinct dimensions of caseworker behavior. They do not assume that caseworkers differ only along these dimensions, or establish that leniency is empirically independent of precision. That separation remains part of the identification argument in Section 5.3.

C.3. The two effects of a tighter labor market

The third illustration allows both the firm's outside option and candidates' willingness to accept to change across market conditions. Within each scenario, all candidates accept with the same probability $a \in (0, 1]$, independently across candidates and independently of residual surplus. Thus $\alpha(m; \vartheta) = a$, where ϑ denotes labor-market tightness. Constant acceptance is a special case of the weakly increasing acceptance function in the main framework. Inspection reveals willingness to accept as well as surplus, and willing candidates remain available during search.

At the initial outside option, the expected return to inspecting fit m is $a\psi(\tau^0 - m)$. The

initial hiring inspection frontier is therefore

$$\underline{m}^h = \tau^0 - \psi^{-1}(\kappa/a). \tag{29}$$

For each candidate define

$$p_i^h = \mathbf{1}\{m_i > \underline{m}^h\}\, a[1 - \Phi(\tau^0 - m_i)]. \tag{30}$$

This is the probability that the candidate is willing and has surplus above the outside option, set to zero when inspection is not worthwhile. It is not the probability that the firm ultimately hires that particular candidate. The probability of hiring someone from S_n and the gain from a nested addition are

$$P^h(S_n) = 1 - \prod_{i \in S_n}(1 - p_i^h), \qquad g^h(n) = p_{n+1}^h \prod_{i \in S_n}(1 - p_i^h). \tag{31}$$

These expressions hold with fixed inspection costs and the stated independence assumptions. With identical fit across candidates, $p_i^h = p^h$ and $g^h(n) = p^h(1 - p^h)^n$.

Table C.2. Outside options, acceptance, and marginal hiring effects

Scenario	τ^0	a	$\underline{m}^h$	$m = 0.05$ $g^h(0)$ (pp)	$m = 0.8$ $g^h(0)$ (pp)	$m = 0.8$ $g^h(5)$ (pp)
Slack	1.0	0.80	0.120	0.00	33.66	4.33
Firm-side only	0.6	0.80	−0.280	23.29	46.34	2.06
Both forces	0.6	0.60	−0.117	17.47	34.76	4.11
Larger fall in acceptance	0.6	0.35	0.222	0.00	20.27	6.53

Notes: The parameter τ^0 is the firm's outside option, a is the acceptance probability and $\underline{m}^h$ is the initial hiring inspection frontier. Residual surplus is standard normal; $\lambda = 0$ and $\kappa = 0.08331547$ in every row. The effect $g^h(n)$ is the increase in hiring probability from one additional candidate, reported in percentage points. The $m = 0.05$ column adds a candidate to an empty shortlist. The last two columns use identical fit $m = 0.8$ for all existing and added candidates. Scenarios are illustrative changes in outside options and acceptance, not estimates for empirical tightness quartiles. The model has no calendar-time hiring horizon.

A weaker firm outside option alone lowers the frontier. A moderate decline in acceptance offsets part of this change, while a sufficiently large decline reverses it. A candidate with $m = 0.05$ is consequently below the frontier in the slack scenario, above it in the next two scenarios, and below it again under the largest acceptance decline.

Candidates with $m = 0.8$ are worth inspecting in all four scenarios. The firm-side-only change raises the hiring gain at $n = 0$ but lowers it at $n = 5$: a higher probability of success also makes existing candidates more likely to suffice. Figure C.2 illustrates the frontier shifts and these marginal-effect profiles. The framework therefore permits, but does not guarantee, larger marginal hiring effects in tighter markets.

These calculations concern completed matches in a static search problem. They do not reproduce the empirical sequence of formal proposals and responses, predict elapsed

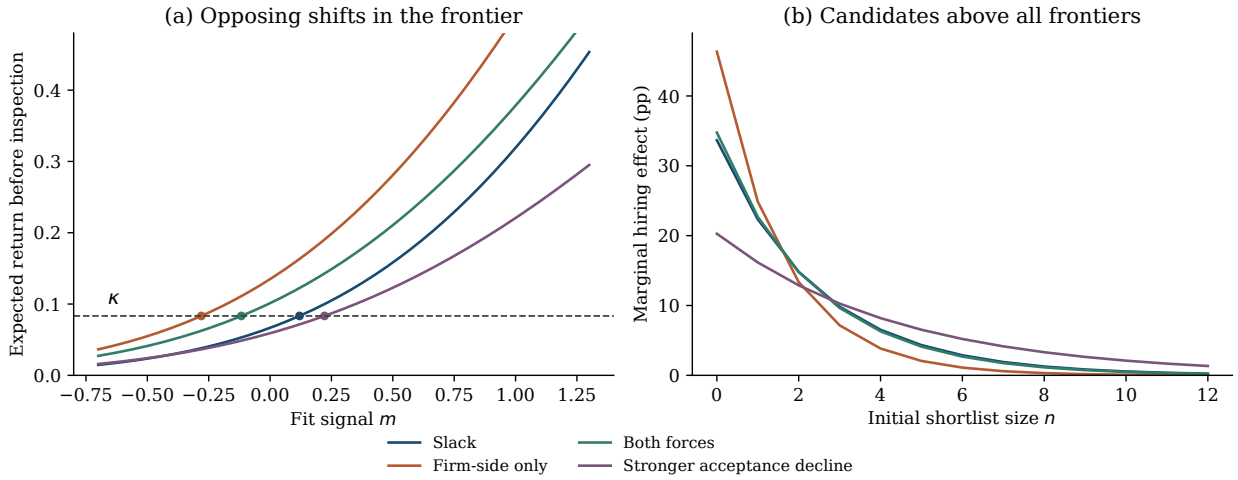


FIGURE C.2. The two tightness channels and the value of an additional candidate

Notes: Panel (a) plots the expected initial inspection gain, $a\psi(\tau^0 - m)$, against candidate fit m. The horizontal line is inspection cost κ, and intersections mark the hiring inspection frontiers. Panel (b) plots $100g^h(n)$, the marginal hiring effect in percentage points, against initial shortlist size n, with all candidates having fit $m = 0.8$. Parameters and scenarios are defined in Table C.2; $\lambda = 0$ throughout.

recruitment time, or identify acceptance from the ratio of vacancy-level hiring and proposal effects. In particular, a longer time to the last proposal can be consistent with further search without providing a direct measure of inspection effort.

C.4. How many additional recruitments can shortlist expansion deliver?

The empirical estimates can be used for local calculations of recruitment gains at prevailing market conditions. The relevant outcomes are the probability of a contract proposal within three months and of a realized hire within six months. These are distinct from the number of people hired, and from net employment creation across the labor market. The calculations below separate an additional candidate on a vacancy from assignment of an establishment to a more lenient caseworker.

One additional candidate per vacancy. Let V be a specified number of vacancies receiving one additional candidate, and let $\widehat{\beta}_P$ and $\widehat{\beta}_H$ denote the IV coefficients for any proposal and any hire, respectively. If the instrument-induced average responses apply to these vacancies and to the candidates added by the policy, the implied changes in expected numbers of vacancies with each outcome are

$$\widehat{\Delta V}_P = V\widehat{\beta}_P, \qquad \widehat{\Delta V}_H = V\widehat{\beta}_H. \tag{32}$$

Table 5 reports $\widehat{\beta}_P = 0.0312$ and $\widehat{\beta}_H = 0.008$. Applied to 1,000 vacancies, these coefficients imply approximately 31 additional vacancies with a proposal within three months and 8 with a hire within six months. The outcomes overlap and must not be added together. The

calculation is a local extrapolation of the IV estimates, rather than an identified effect of universal expansion: a uniform policy need not affect the same vacancies or add the same types of candidates as caseworker leniency.

TABLE C.3. Local recruitment calculations using the reported estimates

Outcome	Coefficient	Scaled change
Panel A: One additional candidate on each of 1,000 vacancies		
Vacancies with a proposal within three months	0.0312	31.2
Vacancies with a hire within six months	0.0080	8.0
Panel B: A one-SD leniency increase for 1,000 entering establishments		
Distinct recruits within 24 months, by recruitment channel		
Chosen for the focal vacancy	0.0129	12.9
Chosen for a subsequent preselected vacancy	0.0145	14.5
Shortlisted but never chosen	0.0037	3.7
Never shortlisted	0.3186	318.6
All channels	0.3497	349.7
Approximate 95% interval, all channels		[−131.9, 831.3]

Notes: Panel A multiplies the unconditional indicator coefficients in Table 5 by 1,000. It counts vacancies with each outcome, not distinct people across vacancies. Panel B multiplies the 24-month reduced-form coefficients in Table A.19 by 1,000. Its four channels are mutually exclusive and exhaustive within an establishment and sum to the total. Distinct recruits are counted within establishment; these counts are not deduplicated across establishments. The interval is the normal approximation $1{,}000(0.3497 \pm 1.96 \times 0.2457)$, using the reported caseworker-clustered standard error. The panels concern different interventions, units and horizons and are not additive. Scaling requires the reported local responses to apply to the proposed intervention and target population, holding market conditions fixed.

Reallocation under a fixed number of transmissions. The empirical heterogeneity suggests examining whether additional candidates could be directed toward vacancies with smaller predicted pools. Let $q \in \{1, \ldots, 5\}$ index the predicted-shortlist quintiles, V_q the number of vacancies in group q, and d_q the average change in candidates transmitted per vacancy in that group. A fixed transmission budget requires

$$\sum_{q=1}^{5} V_q d_q = 0. \tag{33}$$

Let $\widehat{\beta}_{Hq}$ denote the group-specific effect on six-month hiring probability of one additional candidate, expressed in probability units, as in Figure A.1. A local approximation to the

change in vacancies with a hire is

$$\widehat{\Delta V}_H^{\text{reallocation}} = \sum_{q=1}^{5} V_q \widehat{\beta}_{Hq} d_q. \tag{34}$$

For example, moving B transmissions from a donor group q_D to a recipient group q_R gives $B(\widehat{\beta}_{Hq_R} - \widehat{\beta}_{Hq_D})$ additional vacancies with a hire under this approximation. The calculation requires feasible additions and removals, locally stable effects, and the assumption that the expansion estimates also describe the losses from marginal contractions. The same accounting can be applied to proposals with the corresponding per-candidate coefficients.

The effects of a 10 percent expansion in Figure 4 cannot be inserted directly into this per-candidate calculation. A proportional expansion changes the number of candidates by different amounts across vacancies. Moreover, the quintiles compare groups of vacancies; they do not identify the sequence of marginal gains from repeatedly expanding or contracting a single shortlist. A numerical allocation therefore requires the group-specific per-candidate estimates, feasible transmission changes and their joint sampling uncertainty. The reported evidence supports examining targeted expansion, but does not identify a globally optimal allocation or a precise total gain from reallocation.

A fixed number of transmissions is also distinct from a fixed resource budget. Finding and screening an additional candidate may require different amounts of caseworker effort across vacancies. Firm inspection costs may change as well, particularly when $\lambda > 0$. Equation (33) therefore does not imply that a reallocation is costless.

Assignment to more lenient caseworkers. The establishment-level reduced forms provide a separate calculation that is closer to the assignment variation. At 24 months, Table A.19 reports 0.0129 additional distinct recruits chosen for the focal vacancy and 0.0145 chosen for subsequent preselected vacancies per one-standard-deviation increase in predetermined leniency. Their sum is 0.0274, or 27.4 recruits per 1,000 entering establishments. This sum describes those two recruitment channels, rather than an estimate of the total establishment effect. The four-channel total is 0.3497 recruits per establishment, with a standard error of 0.2457, and is imprecisely estimated (Table C.3, Panel B).

These effects concern assignment at entry and subsequent recruitment over 24 months. They cannot be attributed entirely to one extra candidate on the first shortlist. Later hires may reflect changes in program reuse, later shortlist sizes and other consequences of the continuing caseworker relationship. The extensive- and intensive-margin calculations in Table A.20 are an accounting exercise, not an identified mediation decomposition. Adding them to the vacancy-level calculation in Panel A would mix interventions and horizons and could count some recruitment gains twice. Likewise, a one-standard-deviation change in the measured leniency instrument is not a calibrated one-standard-deviation change in

the model's transmission cutoff.

Scope of the calculations. The hiring outcomes already incorporate whether recruitment is observed within the relevant horizon, so the hiring coefficients should not be multiplied by a separate acceptance rate. The acceptance analysis in Section 7.3 does not establish that marginal candidates refuse more often: the conditional effect on the share of proposed candidates hired is not statistically significant. Similarly, the absence of statistically detectable changes in recruited workers' observed characteristics in Table A.16 is not evidence that every dimension of match quality is unchanged.

Finally, these are partial-equilibrium recruitment calculations. As discussed in Section 9, they do not establish whether a hire at one establishment displaces a hire elsewhere, changes the timing of an otherwise completed match, or produces a net employment gain. The framework contains no market-level displacement mechanism and does not rank displacement across tightness conditions. Without further assumptions on spillovers, the firm-level gains are not a formal upper bound on aggregate employment gains. The numerical examples explain why additional options can be useful and why acceptance matters; the empirical calculations quantify local recruitment responses, subject to the intervention and population to which those responses are applied.