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

Marie-Renée Andreescu (Paris School of Economics)

February 28, 2026

Abstract

Firms are unequal in recruitement opportunities and generally consider that more applicants would allow for more opportunities to find the best fit for a vacancy.However, larger applicant pools may also increase screening costs and generate congestion. This paper studies the causal effect of expanding the set of candidates made available to recruiters through the French Public Employment Service (Pôle emploi) when firms opt into assisted preselection, a service in which caseworkers screen applications and transmit a shortlist to the employer. Using administrative data covering the universe of vacancies and applications on the Pôle emploi job board from 2017 to 2023, I exploit quasi-random assignment of firms to caseworkers to instrument shortlist size with a predetermined measure of caseworker leniency. Instrumental-variable estimates show that increasing the shortlist by one additional candidate raises the probability that a vacancy is filled within three months by 3.1 percentage points. Heterogeneity analyses reveal diminishing marginal returns to additional candidates: effects are largest for small initial shortlists and for lower-wage postings, and attenuate as shortlists grow. No evidence of negative returns to the expansion of the applicant pool is found.

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

JEL codes: J63, J64, J68.

1 Introduction

The use of Internet has decreased the search cost for job seekers who can apply more easily to an increasing number of vacancies (Goldfarb and Tucker, 2019; Birinci et al., 2025). In turn, firms are confronted with an always growing number of applications to sort through: from 2017 to 2025, the average number of applications received by a firm posting a vacancy in France on a single job board increased from 10 to 17. Having a larger pool of candidates has the potential of improving match quality by increasing the chance of finding a suitable worker, but it can also raise screening and coordination costs, potentially creating congestion. Understanding this trade-off is central to employer search: vacancy outcomes depend not only on the distribution of applicants but also on the resources firms devote to evaluating them.

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 might affect both the number of applications processed and the probability of a successful match. This paper overcomes this identification challenge by using a novel judge-design framework applied to the French Public Employment Service (PES), Pôle emploi. The latter 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. The size of the shortlist therefore provides a direct measure of the candidate pool accessible to the recruiter within this channel. What is then the marginal impact of transmitting an additional candidate on vacancy filling and on subsequent firm behavior?

In order to answer to the research question, the analysis leverages a key institutional feature: firms are assigned randomly to caseworkers within the relevant local unemployment agency. Conditional on office-by-time cells, this generates quasi-random variation

in the caseworker responsible for screening a firm's job postings. The paper exploits this attribution mechanism in a judge-design framework. Caseworkers differ in their propensity to transmit candidates, even when operating in the same local labor market and time period. A predetermined measure of caseworker leniencyis constructed from the caseworker's screening decisions in quarters prior to first handling a given firm, and use this measure as an instrument for shortlist size. The empirical strategy estimates two-stage least squares models on vacancies outcomes for firms joining the preselection program between 2017 and 2023. The large administrative dataset allows to track applications, caseworker assignment and job hiring decisions.

The main results show that expanding the shortlist significantly improves matching outcomes. Instrumental-variable estimates imply that one additional transmitted candidate increases the probability that a vacancy is filled within three months by 3.1 percentage points, approximately a 10 percent increase relative to the baseline mean. Larger shortlists also marginally increase establishments' future posting activity and, in particular, the likelihood of returning to assisted preselection in future vacancies. Rather than being a deterrent, it seems that an increased shortlist size encourages firms to take advantage of the public employment services' (PES) labor market policies aimed at firms. Heterogeneity analyzes reveal diminishing marginal returns: effects are largest when initial shortlists are small and decline as the pool grows, consistent with the fact that additional candidates are most valuable when the probability of meeting a hiring threshold is low. However, the value added of a shortlist never becomes negative: there is never a point when adding an applicant is harmful to the firm due to excessive screening costs.

This paper contributes to three strands of litterature. First, the identification is based on a specific active labor market policy (ALMP). A large part of the literature looking at ALMPs focuses on policies that operate on the supply side of the labor market, such as offering training and search assistance for job seekers (Card et al., 2010), rather than on labor demand and firms' recruiting costs. This paper expands the incipient literature on demand-aimed ALMPs (Pallais, 2014; Carranza et al., 2022; Abel et al., 2020). 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

from preselection, vacancy writing assistance and interview support. This experiment took place in France in 2014, predating the current framework and comparing SMEs that were not assisted by caseworkers with ones who were. They emphasize the role of preselection in lowering hiring costs, as firms using this program are more likely to publish new vacancy. This paper builds on their findings by looking at the functioning of the program outside of the pilot phase, where self-selection in the treatment is endogenous. This paper is mostly contributing to the literature on hiring costs, screening frictions and the recruitment process. Search-and-matching models interpret hiring costs as a central driver of labor market frictions on the demand side (**?**). Despite this role in theory, the internal mechanics of how firms recruit and the costs they face remain largely unexplored. The recruitment process seems to be a black box. Recent empirical work has begun to look in more detail at this part of the job-matching problem, by documenting for instance the congestion phenomenon (Horton et al., 2024). Separately, work on information and referrals has shown that the quality of information available to recruiters shapes who gets hired, with referred candidates enjoying significant advantages because they arrive with a credible signal attached (Pallais and Sands, 2016). Other contributions have examined how vacancy design, job requirements, and wage posting affect the composition and volume of applications (Balgova et al., 2025; Belot et al., 2022; Card et al., 2025; Burn et al., 2022). This paper builds on these results to look at the impact of the number of applicants for job-employee matching efficiency and at the existing trade-off with screening costs. 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 agents that are attributed to firms, and not to an individual.

The remainder of the paper is organized as follows. Section 2 presents a conceptual framework. Section 3 describes the institutional setting. Section 4 presents the data and descriptive statistics. Section 5 details the empirical strategy. Section 6 reports the main results. Section 7 concludes.

2 Conceptual framework

When recruiting workers, firms typically receive a pool of applications and must allocate resources to evaluate candidates before making a hiring decision. A larger applicant pool increases the probability of identifying a suitable match but also raises the costs of screening and interviewing candidates. This trade-off between match quality and evaluation costs is central to the employer search literature, which emphasizes that vacancy outcomes depend both on the distribution of applicant quality and on firms' recruitment behavior (Van Ours and Ridder, 1992). It is also closely related to the broader search-and-matching framework, in which firms incur hiring costs while attempting to fill vacancies (Mortensen and Pissarides, 1994; Pissarides, 2000). This trade-off is illustrated in a simple model where firms receive an exogenous pool of applicants and obtain value from the best available match, while incurring screening costs that increase with the size of the pool. The framework highlights that enlarging the applicant pool generates diminishing returns and may eventually reduce firm surplus when marginal screening costs exceed marginal gains, a mechanism refered to as congestion.

2.1 Environment

Consider a vacancy posted by a firm. Let $N \in \mathbb{N}$ denote the number of candidate applications that reach the firm through a given recruitment channel. The firm chooses how many applications to review and potentially interview. In the institutional setting studied below, N corresponds to the number of candidates transmitted to the recruiter through the public employment service.

Candidate $j \in \{1, \ldots, N\}$ has match quality $q_j \in \mathbb{R}$ for the firm. Qualities are i.i.d. draws from a distribution $F(q; \theta)$, where the parameter θ summarizes vacancy- and firm-specific conditions. The firm receives value $V(q_j)$ from hiring candidate j, with $V'(\cdot) > 0$.

The firm can either (i) hire the best candidate among the N applications it observes, or (ii) not hire through this channel (because it fills the job through outside options or chooses to leave the vacancy unfilled). For conceptual clarity, the hiring decision is modeled as a comparison between the best available candidate and an outside option.

2.2 Hiring decision and the role of pool size

Let

$$q_{(N)} \equiv \max\{q_1, \ldots, q_N\}$$

be the best match quality among the N visible candidates. Let b denote the firm's outside option (the value of hiring through other channels or leaving the position open), expressed in the same units as $V(\cdot)$. The firm hires from the observed pool if and only if

$$V\Big(q_{(N)}\Big) \geq b.$$

Define the probability of a successful hire through the channel as

$$P(N; \theta, b) \equiv \Pr\Big(V(q_{(N)}) \geq b\Big).$$

Because the q_j are i.i.d., the CDF of the maximum is

$$\Pr\big(q_{(N)} \leq x\big) = F(x; \theta)^N.$$

Hence for any x with $F(x; \theta) \in (0, 1)$ and any $N' > N$,

$$\Pr\big(q_{(N')} \leq x\big) = F(x; \theta)^{N'} \leq F(x; \theta)^N = \Pr\big(q_{(N)} \leq x\big),$$

so $q_{(N')}$ first-order stochastically dominates $q_{(N)}$. Since $V(\cdot)$ is increasing, the event $\{V(q_{(N)}) \geq b\}$ is equivalent to $\{q_{(N)} \geq q^\star\}$ with $q^\star = V^{-1}(b)$, implying that $P(N; \theta, b) = \Pr(q_{(N)} \geq q^\star)$ is weakly increasing in N.

A conservative to consider the hiring decision is to look at the quality threshold defines such that:

$$q^\star \equiv V^{-1}(b),$$

so that hiring occurs if $q_{(N)} \geq q^\star$. Since $\Pr(q_{(N)} < q^\star) = F(q^\star; \theta)^N$, it follows that

$$P(N; \theta, b) = 1 - F(q^\star; \theta)^N. \tag{1}$$

Equation (1) illustrates how expanding the visible pool increases hiring success: the firm is more likely to observe at least one candidate above threshold when N is larger.

2.3 Screening costs and the firm's trade-off

Observing and evaluating applications is costly (Mortensen and Pissarides, 1994): it takes recruiting effort, advertisement, screening costs. Here, the main interest is in the screening cost of looking at the applications received and not the cost of posting the vacancy and advertising it. Let $C : \mathbb{N} \to \mathbb{R}_+$ be the screening cost associated with a pool of size N, with

$$C(0) = 0, \qquad C(N+1) - C(N) > 0,$$

and increasing marginal screening costs,

$$\Delta^2 C(N) \equiv \big[C(N+2) - C(N+1)\big] - \big[C(N+1) - C(N)\big] \geq 0.$$

The firm's expected payoff from having access to a pool of size N through this channel is

$$\Pi(N; \theta, b) \equiv \mathbb{E}\Big[\max\{V(q_{(N)}),\, b\}\Big] - C(N). \tag{2}$$

It is useful to rewrite the first term as

$$\mathbb{E}\Big[\max\{V(q_{(N)}),\, b\}\Big] = b + \mathbb{E}\Big[\big(V(q_{(N)}) - b\big)_+\Big],$$

which shows that access to the channel yields a baseline value b plus an *option value* that comes from occasionally drawing a candidate whose value exceeds b.

2.4 Congestion and net benefits of additional applicants

Even though N is exogenous in the institutional setting, the payoff function (2) clarifies when an increase in the pool is beneficial versus congesting. Consider the discrete marginal effect of increasing the pool by one candidate:

$$\Delta\Pi(N) \equiv \Pi(N+1; \theta, b) - \Pi(N; \theta, b) = \underbrace{\Delta G(N; \theta, b)}_{\text{marginal option value}} - \underbrace{\Delta C(N)}_{\text{marginal screening cost}}, \tag{3}$$

where $G(N; \theta, b) \equiv \mathbb{E}[\max\{V(q_{(N)}), b\}]$ and $\Delta C(N) \equiv C(N+1) - C(N)$.

A congestion effect arises when marginal screening costs eventually dominate marginal

option value:

$$\Delta\Pi(N) < 0 \quad \text{for sufficiently large } N.$$

Intuitively, as the pool grows, the probability that the *overall best* candidate is substantially better than the current best shrinks (diminishing returns), while screening burdens may rise sharply if attention is scarce or coordination costs increase.

2.5 Existence and uniqueness of an optimal pool size

Although the institutional setting studied below treats N as exogenous, it is instructive to consider the counterfactual problem in which the firm could choose how intensively to process candidates:

$$N^\star(\theta, b) \in \arg\max_{N \in \mathbb{N}} \Pi(N; \theta, b).$$

Existence of a maximizer follows under mild conditions. Let $G(N; \theta, b) \equiv \mathbb{E}\Big[\max\{V(q_{(N)}), b\}\Big]$. If $G(N; \theta, b)$ is bounded above and $C(N) \to \infty$ as $N \to \infty$, then $\Pi(N; \theta, b) = G(N; \theta, b) - C(N) \to -\infty$ as $N \to \infty$, so a finite maximizer exists. Uniqueness, however, is not automatic with discrete N. A sufficient condition is strict discrete concavity of Π, which can be ensured when the marginal option value

$$\Delta G(N; \theta, b) \equiv G(N+1; \theta, b) - G(N; \theta, b)$$

is strictly decreasing in N (diminishing increments) and the marginal screening cost

$$\Delta C(N) \equiv C(N+1) - C(N)$$

is non-decreasing (convex screening costs). Under these conditions, the payoff rises for small pools, reaches a unique peak, and falls thereafter.

The discrete marginal payoff from expanding the pool is

$$\Delta\Pi(N) \equiv \Pi(N+1; \theta, b) - \Pi(N; \theta, b) = \Delta G(N; \theta, b) - \Delta C(N),$$

and congestion arises when $\Delta C(N) > \Delta G(N; \theta, b)$, so that an additional candidate reduces expected payoff. Intuitively, when the pool is already large, the probability that an extra

candidate substantially improves on the current best match shrinks, while screening costs continue to rise. In the empirical setting studied here, firms do not choose N; this discussion instead provides the conceptual foundation for why expansions in pool size may exhibit diminishing returns and can become counterproductive when screening costs are sufficiently steep.

3 Institutional setting

3.1 Pôle emploi and its organizational structure

During the sample period (2017–2023), the French Public Employment Service (PES) operated under the name *Pôle emploi*.[1] 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 are first handling exclusively job seekers before switching to the employer

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

department. As specified in Appendix Table 8, when the switch occurs, caseworkers still handle some job-seekers from their previous portfolio, and get attributed new job seekers according to agency congestion, although a smaller number than caseworkers handling exclusively job seekers. Caseworkers can switch between the two departments indefinitely according to agency needs and personal preferences.

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 freely can apply to 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).

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 equally two active mechanisms put in place to improve vacancy-job seeker matching. On 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 for intermediation services has become available in 2015: Pôle emploi has 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' CV 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 firms specifies their main recruiting criteria and the shortlist size they are 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 the caseworker that 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 by firms. Discussions with caseworkers handling preselection have shown that often they only base themselves on the vacancy and their own understanding on 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 which are preselected are not shown to the firm, which then has full freedom to interview or not the preselected applicants.

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 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;
- Hiring outcomes: whether the vacancy led to a contract offer and the timing of the match;
- Caseworker identifiers and agency affiliation for the vacancies (available from the second semester of 2016 onward).
- Applicant characteristics: gender, age, professional experience, unemployment spell legth, reservation wage and attributed caseworker identifier.

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 application 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. 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.

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 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.

Outcome variables The study focuses on two broad categories of outcomes.

1. **Matching efficiency**: this is measured along two dimensions: whether the vacancy is successfully filled and the timing of the match. The main outcome is an indicator for whether the vacancy results in a contract proposal, within three months of the posting, time by which the hiring decisions are generally made for most postings. Additional measures capture hiring speed (success within 30, 45, or 60 days of posting) and vacancy duration (the number of days the vacancy remains online before closure).

2. **Firm engagement**: subsequent firm behavior is analyzed by looking at the firm's probability of posting a new vacancy within 3 or 6 months and the probability that those vacancies are also using preselection. These outcomes allow to study firms' future reliance on the public employment service.

4.2 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 small 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 21 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, 80% are retained in the estimation sample because 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 preselection 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 adopt preselection at the time of their first vacancy posted on the Pôle emploi job board. 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 1: DESCRIPTIVE STATISTICS: FIRMS AND JOB OFFERS

	Firms	Offers without Preselection	Offers with Preselection
Overview			
Number of firms	327,769		
Office × cohort cells	67,353		
First-time posters	70%		
Number of job offers	2,075,786	1,266,111	809,675
Average number of job offers per firm		10	3
Average number of applications per offer		9	21
Firm size distribution (%)			
1–9 employees (Micro)	80.9	78.4	78.2
10–49 employees (Small)	8.2	12.6	9.5
50–249 employees (Medium)	1.5	4.0	2.5
250+ employees (Large)	0.7	1.1	1.1
Contract type distribution (%)			
Fixed-term		25.7	39.3
Permanent		45.7	55.3
Other		27.2	5.5

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 offers" without preselection characteristics 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 percentages.

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 institutionnal 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 counselor c in office o and time period t, are observed (i) the number of candidates transmitted to the recruiter through the preselection service, denoted $X_i \in \mathbb{N}$, and (ii) whether the vacancy results in a hire, denoted $Y_i \in \{0, 1\}$. There are also predetermined vacancy and establishment characteristics W_i.

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.

The parameter of interest is the causal effect of increasing the number of supplied candidates on the probability of filling the vacancy:

$$\beta \equiv \frac{\partial}{\partial x} \Pr(Y_i = 1 \mid do(X_i = x)).$$

Economically, β measures the change in hiring success induced by one additional candidate application made visible to the recruiter through the public employment service.

5.2 Instrumenting strategy: predetermined counselor leniency

Endogeneity concern The observed β is highly endogenous, as shortlist size is not randomly chosen. counselors may transmit more candidates when a vacancy appears harder to fill, when the job is more attractive, or when the employer communicates stricter screening criteria. These determinants are only imperfectly observed and may be correlated with the hiring outcome. A naive regression of Y_i on X_i therefore risks confounding the causal effect of candidate supply with unobserved vacancy difficulty and employer selectivity.

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 counselor's past preselection activity, in the quarters before they encounter a new firm. Let's define $t_{-1}(f, c)$ the last period strictly before counselor c starts handling preselected vacancies for establishment f. The instrument is then computed using only counselor c's historical screening behavior observed up to $t_{-1}(f, c)$:

$$Z_{fc} \equiv g(\{\text{screening decisions of } c \text{ in periods } \leq t_{-1}(f, c)\}),$$

$$Z_i = Z_{fc} \text{ for all vacancies } i \text{ handled by } (f, c).$$

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 using the same realizations of the endogenous variable that appear in the first stage (Goldsmith-Pinkham et al., 2025; Frandsen et al., 2023). While this strategy is restricting the sample to counselors that already have some preselection experience, as the first firm attribution to a starting counselor is mechanically dropped, this does not actually affect the sample size a lot (less than 10% of the observations are dropped) and is a way to ensure a more transparent instrument.

Econometric specification The following two-stage least squares model is estimated:

$$X_i = \pi Z_i + \Gamma' W_i + \gamma_{o\times t} + \delta_{\tau(i)} + u_i, \tag{4}$$

$$Y_i = \beta X_i + \Gamma' W_i + \gamma_{o\times t} + \delta_{\tau(i)} + \varepsilon_i, \tag{5}$$

where $\gamma_{o\times t}$ are office-by-time 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 such as a flexible function of the number of applications received through the platform and timing variables describing when shortlisting occurs. Standard errors are clustered at both the establishment and counselor levels.

Identification comes from comparisons of vacancies handled in the same office and time cell but assigned to counselors with different pre-measured leniency.

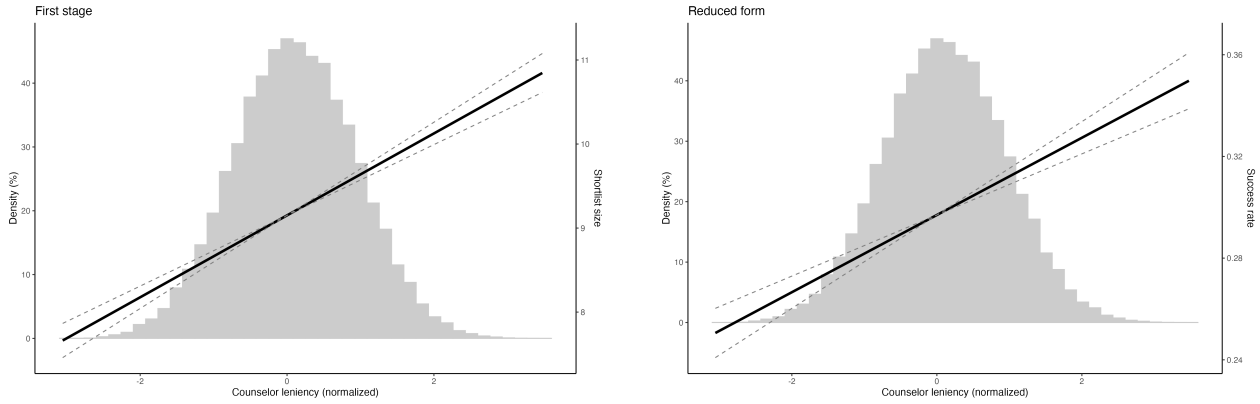
5.3 Identification assumptions

The causal interpretation of β relies on four assumptions.

Assumption 1: Relevance. counselor leniency affects the number of candidates transmitted:

$$\pi \neq 0.$$

This amounts to looking at the first-stage estimates and excluded-instrument F-statistics. Figure 1 (Panel a) provides graphical evidence of this relationship, showing that vacancies assigned to more lenient counselors receive larger shortlists on average. Table 2 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-statistics equals 92.4, 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 1: 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 leave-out mean 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 hiring success. Vacancies assigned to more lenient counselors exhibit higher hiring probabilities.
Solid lines show linear fits; dashed lines indicate 95% confidence intervals. Gray histograms display the distribution of counselor leniency.

Assumption 2: As-good-as-random assignment within office-by-time cells. Conditional on fixed effects and controls, assignment to a more or less lenient counselor is independent of unobserved determinants of hiring success:

$$Z_i \perp \varepsilon_i \mid W_i, \gamma_{o\times t}, \delta_{\tau(i)},$$

TABLE 2: FIRST-STAGE ESTIMATES

	(1) Full sample	(2) Only preselection
Counselor leniency	0.49***	0.51***
	(0.05)	(0.05)
Obs	809,675	602,564
F-stat (excluded instr.)	92.4	102.2
Own-instrument pct 1	-1.43	-1.40
Own-instrument pct 99	1.43	1.40
Implied Δ shortlists (p99–p1)	1.39	1.41

Notes: First stage from regressing shortlist size on counselor leniency with fixed office, entry-cohort and calendar time. Standard errors are clustered by counselor.

where Z_i denotes counselor leniency, W_i vacancy-level controls, $\gamma_{o\times t}$ office-by-cohort fixed effects, and $\delta_{\tau(i)}$ post-by-month fixed effects.

This assumption is supported by the institutional process described in Section 2.4: 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. Table 3 reports estimates from regressions of the form

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

where X_i denotes predetermined characteristics measured at the establishment's first observed vacancy posting. Under as-good-as-random assignment, the coefficient ρ should not be significantly different from zero.

The estimates show no economically meaningful relationship between counselor leniency and predetermined characteristics. Although some 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, as reported in Table 3. Vacancies seem to be randomly assigned to more or less lenient counselors within office-by-time cells.

TABLE 3: BALANCE ON PREDETERMINED CHARACTERISTICS: FIRST VACANCY POSTING

	(1) Counselor effect	(2) Mean	(3) SD
Firm size			
1–9 (Micro)	-0.002	0.802	0.398
	(0.001)		
10–49 (Small)	0.002*	0.084	0.277
	(0.001)		
50–249 (Medium)	0.001**	0.016	0.127
	(0.000)		
250+ (Large)	0.000	0.007	0.086
	(0.000)		
Unknown	-0.001	0.090	0.287
	(0.001)		
Contract type			
Long-term contract	-0.005***	0.540	0.498
	(0.002)		
Short-term contract	0.003	0.412	0.492
	(0.002)		
Qualification			
Blue-Collar	-0.000	0.217	0.412
	(0.002)		
Blue Collar +	-0.002*	0.088	0.283
	(0.001)		
White-Collar	-0.000	0.470	0.499
	(0.002)		
White-Collar +	-0.000	0.025	0.156
	(0.001)		
Manager	0.001***	0.021	0.145
	(0.001)		
Other/Unknown	0.001	0.179	0.383
	(0.001)		
Experience			
Beginner	0.007***	0.429	0.495
	(0.002)		
Experienced	-0.007***	0.571	0.495
	(0.002)		
Published wage quartile			
Bottom 25%	0.004***	0.203	0.402
	(0.001)		
25–50%	-0.001	0.187	0.390
	(0.001)		
50–75%	-0.001	0.216	0.411
	(0.001)		
Top 25%	-0.003***	0.168	0.374
	(0.001)		
Unknown	0.001	0.226	0.419
	(0.002)		
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 post-by-month fixed effects. Standard errors are clustered at the counselor level. Categorical variables are indicator dummies. Significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Assumption 3: Exclusion restriction. Leniency affects hiring outcomes only through the number of candidates transmitted:

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

This exclusion restriction 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 to differentiate between vacancies for which the firm receives *only* the preselection service from those that benefit from a bundle of services, through which the influence of 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 2 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 10 are consistent to the ones from the whole sample. These results support the exclusion restriction assumption.

Assumption 4: 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 partitionned into quartiles based on predicted assignment to treatment, obtained from an OLS regression of shortlist size on predetermined characteristics. Within each quartile, is re-estimated the first-stage relationship between counselor leniency and shortlist size.

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. Table 9 provides an additional check of monotonicity using a reverse instrument based on the lagged counselor leniency measure. The estimates are positive and statistically significant in every quartile. Moreover, the magnitude increases from the lowest to the highest quartile, indicating that counselor leniency shifts shortlist size in the same direction throughout the distribution of predicted shortlist sizes. While monotonicity is not directly testable (Goldsmith-Pinkham et al., 2025; Frandsen et al., 2023), the absence of sign reversals across quartiles provides no empirical indication of systematic defiers and is consistent with interpreting the IV estimates as a LATE for vacancies whose shortlist size is increased by assignment to more lenient counselors.

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
P-value (F-statistic)	0.000	0.000	0.000	0.000

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: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Interpretation Under Assumptions 1–4, β identifies a local average causal effect: the effect of an additional transmitted candidate for vacancies whose shortlist size is shifted by counselor leniency.

6 Effects of increasing the candidate pool

6.1 Average effects on matching probability

Table 5 reports OLS and IV estimates of the effect of shortlist size on vacancy outcomes. The OLS estimates suggest that larger shortlists are weakly associated with higher overall filling rates but lower probabilities of being filled quickly. This pattern likely reflects negative selection: vacancies that are harder to fill may receive larger shortlists and simultaneously exhibit longer durations.

In contrast, the IV estimates imply substantial positive causal effects of shortlist size on hiring success. A one-candidate increase in the shortlist raises the probability that a vacancy is eventually filled by 3.1 percentage points, corresponding to a 10 percent increase relative to the baseline mean of 0.299. Shortlists also significantly increase the probability

of being filled within 14, 30, and 45 days.

The IV coefficients are much larger than the OLS estimates and reverse the negative short-run correlations observed in OLS, consistent with downward bias in naive regressions. These results indicate that expanding the candidate pool causally improves vacancy filling and accelerates hiring. Appendix Table 10 restricts the sample to vacancies with only preselection and yields estimates that are very similar to the main results. The IV coefficients remain positive, statistically significant, and economically meaningful, confirming that the baseline findings are not driven by other mechanisms than the pure increase of the shortlist size.

TABLE 5: EFFECT OF SHORTLIST SIZE ON VACANCY FILLING

	Dependent variable: Vacancy filled 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

6.2 Average effects on vacancy publishing probability

Table 6 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.

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 hiring.

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.

6.3 Decreasing returns to an additional candidate

Figure 2 examines heterogeneity in the IV effect of shortlist size on vacancy filling by bins of initial shortlist size. The figure reports 2SLS estimates of the impact of one additional transmitted candidate on the probability that the vacancy is ultimately filled, estimated separately across intervals of the initial shortlist.

The marginal effect is largest when the initial shortlist is small (0–20 candidates) and declines steadily as the initial pool increases. For vacancies with large initial shortlists (above 60 candidates), the estimated effect becomes small and statistically imprecise. The sample does not provide enough support to estimate effects over 100 candidates, as F-stats fall and the available observations are insufficient.

This pattern provides clear evidence of diminishing marginal returns to enlarging the applicant pool. When few candidates are initially transmitted, an additional candidate substantially raises the probability of drawing a match above the hiring threshold. As the pool grows, the incremental probability that one more candidate improves upon the current best match decreases.

TABLE 6: 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 $\times$ 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Importantly, the marginal effect does not turn systematically negative within the observed range, suggesting that while gains attenuate at higher shortlist sizes, screening congestion does not dominate match-quality improvements at typical candidate volumes. However, this does not mean that firms do not experience congestion. It is likely that once the number of applications exceeds the interior optimum, firms develop alternative screening ways, such as randomly drawing from the given applications or not looking at applications over a certain number, that help mitigate the cost. This might avoid the modeled situation where firms invest equally in screening each application, making the congestion synonymous to exceeding costs over profit. However, the existing data is insufficient for testing this mechanism.

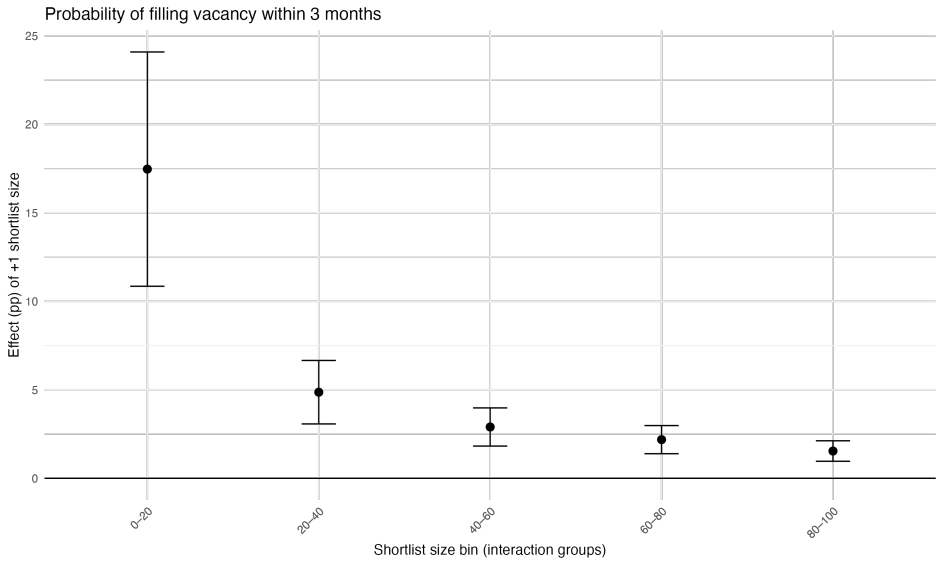


FIGURE 2: HETEROGENEOUS IV EFFECTS OF SHORTLIST SIZE ON VACANCY FILLING

Notes: The figure reports 2SLS estimates of the effect of shortlist size on the probability that a vacancy is filled within three months, across bins of the predicted treatment distribution. Shortlist size is instrumented using lagged counselor leniency. Each point corresponds to a separate regression estimated within percentile bins. Vertical bars denote 95% confidence intervals. All specifications include office-by-cohort and posting-month fixed effects. Standard errors are clustered at the counselor level.

6.4 Heterogeneity in posted wages

Table 7 reports lagged-IV estimates separately by published wage quartile. The first stage is strong in all groups, allowing meaningful comparisons across quartiles. The estimated

effect of an additional shortlisted candidate is positive throughout the wage distribution but declines with the posted wage. In the bottom wage quartile, one additional candidate increases the probability that the vacancy is filled by 5.6 percentage points, whereas the effect is of only 1.3 percentage points in the top quartile. Vacancies with missing wage information display effects similar to the bottom quartile. Overall, these patterns suggest that marginal expansions of the applicant pool are most valuable for lower-wage (and low-information) postings. It is likely that this is due to the differential search for qualitative applicants. Adding any candidate to the shortlist for a low wage and low skilled vacancy can improve the outcome if the firm is looking mainly for a match, no matter the characteristics of the applicant. On the other hand, more demanding vacancies will not benefit from any type of extra candidate - they would take the most value added only if the distance between the vacancy and the extra candidate's characteristics are minimal. This heterogeneity is consistent with an order-statistics mechanism: when the probability mass above the hiring threshold is small, an additional draw increases the chance of reaching the threshold more sharply than when that probability is already high. This result opens the floor to the question of quality of the applicant and the differentiation between having any extra candidate and having the appropriate one.

TABLE 7: HETEROGENEOUS IV EFFECTS BY PUBLISHED WAGE QUARTILE

	Dependent variable: Vacancy filled (within 3 months)				
	Bottom 25%	25–50%	50–75%	Top 25%	Unknown
Effect of +1 shortlisted candidate (2SLS)	0.056***	0.023***	0.017**	0.013*	0.057***
	(0.016)	(0.006)	(0.006)	(0.006)	(0.013)
Observations	160,234	151,367	155,150	149,960	192,964
First-stage F-statistic	15.61	38.36	38.29	31.38	38.14

Notes: Each column reports separate 2SLS estimates within quartiles of the published wage distribution. Shortlist size is instrumented using lagged counselor leniency. All specifications include office-by-cohort and posting-month fixed effects. Standard errors are clustered at the counselor level. The reported F-statistic corresponds to the clustered excluded-instrument first-stage statistic. Significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

7 Conclusion

This paper provides causal evidence that expanding the candidate pool made available to a recruiter raises significantly the probability of filling a vacancy. Using quasi-random variation in caseworker screening leniency as an instrument for shortlist size, the analysis finds that one additional transmitted candidate increases the probability of a successful hire within three months by 3.1 percentage points which amounts to a 10 percent increase at the baseline. These estimates are substantially larger than OLS estimates, confirming that naive regressions are downward biased, probably because vacancies that are harder to fill receive larger shortlists. Two sets of findings extend the baseline result. First, the effects exhibit clear diminishing marginal returns: the gain from one more candidate is largest when the initial pool is small and attenuates steadily as the shortlist grows. The marginal effect does not turn negative within the observed range, suggesting that screening congestion does not dominate match quality gains. Second, larger shortlists increase the probability that firms return to accompanied preselection in subsequent recruitment episodes. This response shows that the gains from an expanded candidate supply are dynamic: firms update their reliance on the public employment service based on prior hiring experiences. The heterogeneity analysis reveals that marginal gains are largest for lower-wage postings, where the probability of drawing a candidate above the hiring threshold is more responsive to an additional observation. These findings have direct implications for the design of labor market intermediation services. Policies that increase the number of candidates transmitted to recruiters through public platforms can raise matching efficiency at the margin. At the same time, the evidence of diminishing returns points to the existence of an interior optimum: the value of increasing the shortlist is highest precisely where it is currently smallest, and the returns to expansion diminish as candidate pools become large. This suggests that caseworker resources may be most productively deployed by ensuring that firms with thin initial shortlists receive more candidates, rather than uniformly expanding applicant numbers across all vacancies. However, these estimates identify a local average treatment effect for vacancies whose shortlist size is shifted by caseworker leniency, and their external validity to vacancies less reliant on the Pôle

emploi platform remains uncertain. The paper also does not directly observe at this point the quality of individual candidates transmitted, so the mechanism is inferred from the aggregate outcome rather than observed at the match level. Future work linking individual candidate characteristics to hiring decisions could shed light on whether leniency differences reflect quantity alone or also composition.

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. 2

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*. 2, 9, 10, 26

Arni, P. and Schiprowski, A. (2019). Job search requirements, effort provision and labor market outcomes. *Journal of Public Economics*, 169:65–88. 3

Balgova, M., Tekleselassie, T., Hensel, L., and Witte, M. (2025). Wage information and applicant selection. 3

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. 9

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. 3

Bhuller, M., Dahl, G. B., Løken, K. V., and Mogstad, M. (2020). Incarceration, recidivism, and employment. *Journal of political economy*, 128(4):1269–1324. 23

Birinci, S., See, K., and Wee, S. L. (2025). Job applications and labour market flows. *Review of Economic Studies*, 92(3):1438–1496. 1

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. 3

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. 3

Card, D., Kluve, J., and Weber, A. (2010). Active labour market policy evaluations: A meta-analysis. *The economic journal*, 120(548):F452–F477. 2

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. 2

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*. 8

Frandsen, B., Lefgren, L., and Leslie, E. (2023). Judging judge fixed effects. *American Economic Review*, 113(1):253–277. 3, 18, 23

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

Goldsmith-Pinkham, P., Hull, P., and Kolesár, M. (2025). Leniency designs: An operator's manual. Technical report, National Bureau of Economic Research. 3, 18, 23

Horton, J., Vasserman, S., and Watt, M. (2024). Reducing congestion in labor markets: A case study in simple market design. 3

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. 3, 23

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

Le Barbanchon, T., Rathelot, R., and Roulet, A. (2016). Applications for job vacancies along the unemployment spell. Working paper. 9

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. 9

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. 4, 6

Pallais, A. (2014). Inefficient hiring in entry-level labor markets. *American Economic Review*, 104(11):3565–3599. 2

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. 3

Pissarides, C. A. (2000). *Equilibrium unemployment theory*. MIT press. 4

Van Ours, J. and Ridder, G. (1992). Vacancies and the recruitment of new employees. *Journal of Labor Economics*, 10(2):138–155. 4

A Supplementary tables

TABLE 8: 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-seekers 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 attributed a caseworker. Panel B reports 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 get attributed a new firm, but on average in their career they follow almost as many job-seekers as caseworkers who are never involved in preselection.

TABLE 9: 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	196,680	206,082	201,153	201,329
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 (without 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 post-by-month fixed effects. Standard errors are clustered at the counselor level.

TABLE 10: EFFECT OF SHORTLIST SIZE ON VACANCY FILLING (PRESELECTION SAMPLE)

	Dependent variable: Vacancy filled within			
	14 days	30 days	45 days	3 months
Panel A: 2SLS (instrumented by lagged counselor leniency)				
Shortlist size	0.0104***	0.0177***	0.0221***	0.0317***
	(0.0020)	(0.0027)	(0.0031)	(0.0037)
Panel B: OLS				
Shortlist size	-0.0030***	-0.0028***	-0.0018***	0.0009***
	(0.00010)	(0.00010)	(0.00009)	(0.00007)
Baseline mean	0.081	0.162	0.214	0.290
Observations	602,564			
Office × Cohort FE	Yes			
Posting Month FE	Yes			

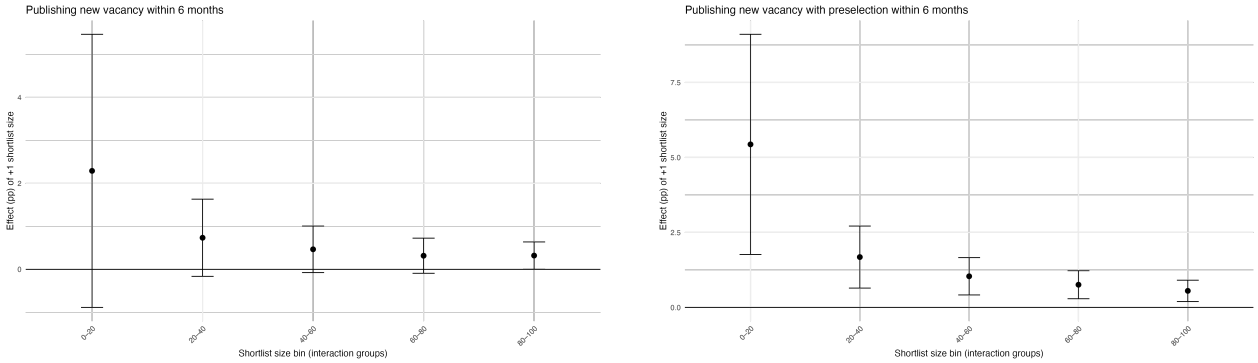
Notes: Sample restricted to vacancies with *only_preselection*=1. Each column reports estimates from a separate regression where the dependent variable equals one if the vacancy is filled within the indicated time horizon. Panel A reports 2SLS estimates where shortlist size is instrumented using lagged 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

TABLE 11: EFFECT OF SHORTLIST SIZE ON SUBSEQUENT VACANCY POSTING (PRESELECTION SAMPLE)

	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	0.0055*	0.0043	0.0108***	0.0103***
	(0.0033)	(0.0032)	(0.0034)	(0.0034)
Panel B: OLS				
Shortlist size	-0.0001*	0.0006***	0.0004***	0.0012***
	(7.21×10^{-5})	(6.99×10^{-5})	(7.36×10^{-5})	(7.91×10^{-5})
Baseline mean	0.513	0.604	0.412	0.487
Observations	602,564			
Office × Cohort FE	Yes			
Posting Month FE	Yes			

Notes: Sample restricted to vacancies with *only_preselection*=1. 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 accompanied preselection (C-vacancies). 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

B Supplementary figures



(a) (b)

FIGURE 3: HETEROGENEOUS IV EFFECTS OF SHORTLIST SIZE ON SUBSEQUENT VACANCY POSTING

Notes: The figure reports 2SLS estimates of the effect of shortlist size on the probability that the establishment posts a subsequent vacancy within six months. Panel (a) considers any new vacancy. Panel (b) restricts attention to vacancies using accompanied preselection. Shortlist size is instrumented using lagged counselor leniency. Estimates are computed within bins of predicted treatment intensity. Vertical bars denote 95% confidence intervals. All specifications include office-by-cohort and posting-month fixed effects. Standard errors are clustered at the counselor level.